

**BigFix Inventory
User's Guide**



Special notice

Before using this information and the product it supports, read the information in [Notices \(on page mcxlix\)](#).

Edition notice

This edition applies to version 9.5 of BigFix and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

- BigFix Documentation Homepage.....
- BigFix 9 Inventory Documentation..... vii
- Overview.....viii
 - What's new in this release.....viii
 - Keeping BigFix Inventory up-to-date..... xlv
 - Features and functions..... xviii
 - Preview features..... xlix
 - Limitations.....xlix
 - Key concepts..... liv
 - Scenarios..... lxxxvi
 - Notices..... xciv
- Installing.....xcix
 - Planning the installation..... xcix
 - Installing on Windows..... cxxix
 - Installing on Linux..... clii
 - Installation-related tasks..... clxxviii
- Configuring..... clxxxiv
 - Creating users and groups..... clxxxiv
 - Setting up scans to discover software and hardware inventory..... cxci
 - Performing optional configuration..... ccxcvii
 - Classifying the BigFix client used by multiple BigFix products..... ccciii
- Upgrading.....cccv
 - Upgrading to BigFix Inventory 9.2.17..... cccv
 - Migrating from License Metric Tool 9.2.17..... cccxv
- Managing the infrastructure.....cccxi
 - Computer statuses..... cccxi
 - Managing VM managers.....cccviii
 - Administering the server.....cccxcix
 - Administering the database..... cdxxiv
 - Removing inactive computers..... cdxxxvii

Managing data sources.....	cdxxxviii
Managing software inventory and metric utilization.....	cdxliv
Software catalogs.....	cdxliv
Enterprise applications inventory.....	dii
Software classification.....	dxxviii
License metrics.....	dlvi
Software license optimization.....	dcxv
Reports.....	dcxvi
Contracts.....	dclvi
Managing security threats.....	dclx
Collecting file checksums.....	dclx
Preview: Checking Common Vulnerabilities and Exposures (CVEs).....	dclxv
Tutorials.....	dclxix
Tutorial: Reporting subcapacity usage per computer group.....	dclxix
Tutorial: Managing software in the service provider environment.....	dclxxv
Tutorial: Migrating software assignments between two BigFix servers.....	dclxxv
Security.....	dcxc
Flow of data.....	dcxc
Security configuration scenarios.....	dcxciii
Configuring secure communication.....	dcxciii
Assuring compliance with federal encryption standards.....	dccii
Authenticating users with LDAP.....	dccv
Configuring and enabling single sign-on.....	dccxii
Configuring passwords and encryption mechanisms.....	dccxxxvi
Configuring user account lockout.....	dccxliv
Configuring VM Manager Tool to accept trusted VM manager certificates.....	dccxlv
Relays.....	dccxlvii
Troubleshooting and support.....	dccxlviii
Troubleshooting a problem.....	dccxlviii
Troubleshooting software discovery.....	dccl
Tuning performance.....	dcccxxxi
Infrastructure.....	dcccxxxi

Setting up scans and import.....	dcccxxv
Additional tasks for improving performance.....	dcccxli
Tuning performance in medium and large environments.....	dcccxliv
REST API considerations.....	dcccli
Integrating.....	dccclvii
REST API.....	dccclvii
Integrating with IBM Control Desk.....	mcxxxv
Glossary.....	mcxl
Index.....	

BigFix 9 Inventory Documentation

Welcome to the BigFix Inventory documentation, where you can find information about how to install, maintain, and use this application. With BigFix Inventory, you can maintain an up-to-date inventory of software assets that are installed in your infrastructure, gather information about your hardware, and ensure license compliance of your enterprise. You always know what software you have, where it is, and how it is used. By monitoring PVU and RVU MAPC consumption under full and subcapacity licensing terms, not only can you understand current license consumption but you can also predict the future one, thus providing the means to optimize your license costs.



Important: This documentation applies to BigFix Inventory application updates 9.2.1 through 9.2.17. Supported BigFix Platform versions: 9.0 through 9.5.

For documentation of newer release 10.0.0 and onward refer to [BigFix 10 Inventory Documentation](#).

Getting started

[What's new in this release \(on page viii\)](#)

[Key concepts \(on page liv\)](#)

[Installing \(on page xcix\)](#)

[Configuring \(on page clxxxiv\)](#)

[Software requirements \(on page ci\)](#)

Common tasks

[Managing the infrastructure \(on page cccxvi\)](#)

[Managing software inventory and metric utilization \(on page cdxliv\)](#)

[Upgrading \(on page cccv\)](#)

[Tutorials \(on page dclxix\)](#)

[Videos](#)

Troubleshooting and support

[HCL Software Support](#)

More information

[BigFix Forum](#)

Overview

Become familiar with key concepts that are necessary to understand how BigFix Inventory works and learn about features and functions that are introduced in every version of the application.

What's new in this release

BigFix Inventory provides new features and enhancements that facilitate your work with the application. Check out what is new.

Application update 9.2.17

- **9.2.17 Support for new platforms and systems**

The following platforms and systems are now supported:

- Hyper-V 2019
- Linux VMs on Nutanix AVH (IBM Power only)
- Red Hat Enterprise Linux 8 endpoints
- RHEV 4.2 and 4.3 using new API available in RHEV-M 4.x

- Important changes:

- **9.2.17 Removal of the legacy IBM Software Classification panel**

The IBM Software Classification was already deprecated and substituted with Software Classification panel in application update 9.2.13. At this point, the IBM Software Classification panel is entirely removed. The new Software Classification panel allows simple software classification and ensures better performance. It is fully customizable and allows you to create custom classification rules.

- **9.2.17 End of support for MS SQL 2008 and 2008 R2 database for BigFix Inventory server**

The support for MS SQL 2008 and 2008 R2 database for BigFix Inventory server is sunset. It is related to Microsoft officially ceasing the support in July 2019. For more information, see: [Prepare now for end of support in 2019 and 2020 \(Microsoft\)](#). To upgrade to BigFix Inventory 9.2.17, first migrate to newer MS SQL version that is supported by BigFix Inventory. For the list, see: [Software requirements \(on page ci\)](#).

Application update 9.2.16

- **9.2.16 Drill down for better understanding of reported values on Software Summary report extended to software usage statistics**

You can additionally use drill down on Software Summary report to see the detailed information about the software instances that contribute to software usage through the related columns, for example, **Used, Unused, Used in the Last Month** and other summary columns. For more information, see: [Available reports \(on page dcxvi\)](#).

- **9.2.16 Improved reporting of FlexPoints**

- BigFix Inventory reports information about the highest license consumption of FlexPoints over a specified reporting period for an entire FlexPoint bundle as well as the contributing products.
- You can view the license metric utilization trends for the entire bundle, as well as its components on the Metric Quantity History graph. The graph presents information about the values calculated by the system and the values that are declared by the user.
- Two new columns are now available on metric reports: Peak Date and Metric Quantity. For more information, see: [Report columns \(on page dcxxx\)](#).
- The All IBM Metrics report shows only license consumption on a bundle level. To see the detailed information about license consumption for the contributing products, drill down by clicking the name of the bundle.
- Audit snapshot is extended with an additional file that outlines the usage of FlexPoints in your environment. The `flexpoints.csv` file contains information about the FlexPoint bundles and the underlying products.

For more information, see: [IBM FlexPoints \(on page dlx\)](#).

- **9.2.16 Setting metric quantity for license metrics that are not calculated by BigFix Inventory**

You can manually declare the metric quantity for selected license types for which BigFix Inventory does not calculate license metric utilization out of the box. You can also use this option to customize the metric quantity if needed. For more information, see: [Setting and removing the declared metric quantity \(on page dlxxxiii\)](#).

- **9.2.16 Extension of the `computer_hardware` API association**

You can retrieve information about additional hardware parameters, such as the serial number, vendor, model, or type of the physical server. For more information, see: [computer_hardware association \(on page mcxxi\)](#).

- **9.2.16 Custom rules for automatically suppressing components that meet specified criteria**

By specifying the relevant criteria, you can create the custom rule which automatically suppresses the matching components in the future. For more information, see: [Creating and managing custom rules \(on page dl\)](#).

- **9.2.16 Extended support**

- Support for new platforms and systems
 - Certified support for the newest BigFix platform updates: 9.2.18 and 9.5.13
 - Debian 9 x86 and x64
 - SUSE Linux Enterprise Server 15 x86-64
 - SMBIOS3 on Linux
 - server on Windows 2019
 - server on MS SQL 2017 version 14.0.3162.1 for fresh installations

- **9.2.16 Security enhancements**

- Rails is upgraded to version 5.1.7 to maintain security.
- WebSphere Liberty Profile is upgraded to version 19.00.4 to maintain security.
- Update of the Xerces library to version 3.2.1 on AIX to keep the scanner secure and reliable.

Application update 9.2.15.2

- **Security enhancements**

- WebSphere Liberty Profile is upgraded to version 19.00.4 to maintain security.

Application update 9.2.15

- **9.2.15 Preview feature: New Software Summary (Preview) report**

The report provides an overview of software inventory in your environment. Each row in the report represents a set of software instances that are grouped by the selected columns. You can drill down the Software Summary report to view the details of the instances in a group. The report also provides software usage statistics that contain information about usage of the specific software based on the listed instances. You can check how many instances of a product within the group are used and how many are not. You can see an overview of software usage statistics from last month, last quarter, and last year. You can also track the number of software instances within the group for which software usage is monitored and for which it is not. For more information, see: [Available reports \(on page dcxvi\)](#).

- **9.2.15 Extended reporting of usage statistics for Java**

BigFix Inventory introduces a new and improved method of reporting usage for Java Virtual Machine and Runtime Environments. Paths to Java processes that run on a computer are matched with paths in which the software is discovered. Based on this information, BigFix Inventory can determine exactly which Java instance was used. It is possible to distinguish between usage of BigFix and Oracle Java installed on a single computer. For more information, see: [Monitoring and reporting software usage \(on page dcxv\)](#).

Software catalog in version 9.2.15 does not contain usage signatures for Java any more. To report usage of Java, upgrade the BigFix Inventory server to version 9.2.15.

- **9.2.15 Extended support**

- Support for new platforms
 - Windows 2019
 - Oracle SPARC-M8 processor
 - Ubuntu 18.04 on x86
- Support for new virtualization technologies
 - vCenter with external Platform Services Controllers (PSC) on different computers
 - Solaris Kernel zones on Oracle Solaris x86 and SPARC with nested virtualization. For more information, see: [Client Installation on Oracle Solaris \(on page ccxii\)](#)

- **9.2.15 New Software Installations report that is based on the new data model**

The new version of the Software Installations report is introduced. The old report is deprecated and moved to the bottom of the Reports menu. It is planned to be entirely removed in the future. The new report view is based on the new data model and consistent with other reports. For more information, see: [Available reports \(on page dcxvi\)](#).

- **9.2.15 Extension of the To Do list**

The To Do list is extended and additionally displays information about the outdated scanning actions that need to be restarted. For more information, see: [9.2.13 To Do list \(on page lxi\)](#).

- **9.2.15 Documentation extended with new topics**

Documentation is extended with new content, including the following help topics.

- [Keeping BigFix Inventory up-to-date \(on page xlv\)](#)
- [Software license optimization \(on page dcxv\)](#)
- [Tuning performance based on scalability guidelines \(on page dcccxxxi\)](#)

- **9.2.15 New look and feel of the user interface**

The look and feel of the user interface is refreshed. The reports are now more vivid with a new blue action bar, updated action icons and selected buttons.

- **9.2.15 Usability improvements on the management panels**

Several improvements are introduced to simplify the navigation and usability of the selected management panels, including Users, Saved Reports, Advanced Server Settings and Contracts. You can use filters based on the available columns, and save the custom view as a predefined view. You can also export these reports to a [CSV](#) or [PDF](#) file.

Limitation: To filter the Contracts report by the All Computers computer group, use its English name regardless of the interface language. The name of this group is not translated.

- **9.2.15 Filtering reports by using `matches exactly` operator**

You can filter reports and retrieve data from specified columns by using the `matches exactly` operator. This operator is used to exactly match two values, including upper and lowercase characters.

- **9.2.15 New naming convention for the exported reports**

The names of report files that are exported, or received in a scheduled email follow a new, more meaningful naming convention. Additionally, the email subject contains the name of the predefined or custom report that is included.

- **9.2.15 Extended discovery of Oracle Databases certified for Oracle Database 18c: 18.3.**

- **9.2.15 Discovery of new components from Citrix, Hewlett-Packard, Oracle and other**

The software catalog is extended to discover a number of new components from Citrix, Hewlett-Packard, Oracle and other. To discover the components, ensure that you upgrade BigFix Inventory to version 9.2.15.

Below is the list of the most important newly added components. To learn about their exact versions, use the Software Catalog widget in BigFix Inventory to browse the catalog content.

- Added discovery capability for Citrix products
 - Citrix XenApp
 - Citrix XenApp Reports
 - Citrix XenDesktop
- Added discovery capability for Hewlett-Packard products
 - HP Business Process Monitor
 - HP Business Service Management
 - HP Operations Agent
 - HP Operations Manager
- Added discovery capability for Oracle products
 - Oracle Clusterware
 - Oracle Linux
 - Oracle SOA Suite
- **9.2.15 Names of BigFix Inventory fixlets in the BigFix console contain information about version**

Names of the BigFix Inventory fixlets are extended with the information about the currently available version. By default, fixlet version is updated with each release of BigFix Inventory server. For more information, see: [Ensuring currency of scanner actions and fixlets \(on page xlvi\)](#).

- **9.2.15 Security enhancements**

Rails is upgraded to version 5.0.7 to maintain security.

Application update 9.2.14

- **9.2.14 Reporting license usage for Adobe Creative Cloud products on Windows**

BigFix Inventory reports utilization of the Registered User metric for Adobe Creative Cloud products. The information is extracted during regular software scans directly from computers where components of some Adobe Creative Cloud products are installed. Based on the scan results, the Registered User metric is calculated and its utilization is displayed on the All Metrics report. Detailed information about users is shown on the Software Users report.

Thanks to this feature, you can track licenses of Adobe CC products in BigFix Inventory to analyze trends, define thresholds and optimize costs by identifying users of unused instances. For more information, see: [Software Users report \(on page dcxxv\)](#) and [Registered User \(on page dlxxxj\)](#).

- **9.2.14 Reporting license usage for BigFix Rational portfolio**

Information about the usage of Floating User metric is now available on the All Metrics report for BigFix Rational products. The calculations are available thanks to integration with BigFix Rational License Key Server. The usage is provided by the `.slmtag` files. For more information, see: [Other IBM metrics \(on page dlxxiv\)](#).

- **9.2.14 Tagging software for more flexible categorization**

You can now tag software components, or instances of the discovered software to classify them. The tags might, for example, emphasize the ownership of a software product, categorize it as a test or free of charge instance, or mark it as blacklisted or approved. For more information, see: [Adding and removing tags \(on page dcli\)](#).

- **9.2.14 Preview feature: The new dashboard is extended with additional widgets and quick search**

The newest additions to the new BigFix Inventory dashboard include the summary of the most important inventory data, recent changes, discovery gaps, and security issues in one place. The dashboard is also extended with quick search that allows you to easily find information related to file hashes, file names, vulnerabilities and software inventory. For more information, see: [Dashboard \(on page liv\)](#).

- **9.2.14 New predefined reports**

New predefined reports are available under the Software Classification, Scanned File Data and Computers reports. For more information, see: [Available reports \(on page dcxvi\)](#).

-

- **9.2.14 Reporting end of support dates for Microsoft products**

The end of support dates are automatically populated for Microsoft software that according to information provided on 20 November 2018 is announced to be out of support between 30 November 2015 and 1 November 2021.

- **9.2.14 Best practices for configuring VM managers**

Documentation provides best practices for configuring VM managers. For more information, see: [Best practices for configuring VM managers \(on page cccxcvi\)](#).

- **9.2.14 Top navigation bar extended with the  icon for providing feedback**

- **9.2.14 Changes in the Report menu**

The Software Catalog report is renamed to Products & Metrics. A link to Software Components report is added under Reports in the main menu.

- **9.2.14 Users panel is expanded with the information about time and date of the last login**

- **9.2.14 Improvements on the VM Managers panel**

Several improvements are introduced to simplify the navigation and usability of the VM Managers panel. You can filter the report by any available column, and save the custom view as a predefined view. You can also export the report to a [CSV](#) or [PDF](#) file.

- **9.2.14 Migrating software classification with REST API**

The [PUT](#) operation on the `/api/sam/v2/software_instances` REST API allows you to define, change, or migrate your current software classification. It can be used to migrate your software assignments when moving

clients between the instances of BigFix server. For more information, see: [Classifying software \(v2\) \(on page mii\)](#) and [Tutorial: Migrating software assignments between two BigFix servers \(on page dclxxxv\)](#).

- **9.2.14 Retrieving information about the license utilization peak and the contributing software instances through REST API in version 2**

If you have custom tools or internal processes that use the `software_instances` REST API, you can now achieve your goals by using the `v2/license_usage` and `v2/software_instances` REST API. It is best to adjust the tools to use REST API in version 2. For more information, see: [Mapping functions of the software_instances REST API \(on page dccclxx\)](#).

- **9.2.14 Discovery of new components from Citrix, Oracle and SAP and other**

The software catalog is extended to discover a number of new components from Citrix, Oracle, SAP and other. To discover the components, ensure that you upgrade BigFix Inventory to version 9.2.14.

Below is the list of the most important newly added components. To learn about their exact versions, use the Software Catalog widget in BigFix Inventory to browse the catalog content.

- Added discovery capability for Citrix products
 - Citrix Diagnostic Facility
 - Citrix PowerShell Snap-In
- Added discovery capability for Oracle products
 - Oracle Discoverer
 - Solaris
- Added discovery capability for SAP products
 - SAP PowerBuilder
- Added discovery capability for Adobe CC 2019
- Improved usage reporting for Java products

- **9.2.14 Certification of Microsoft Edge browser**

Microsoft Edge is certified as a supported browser for BigFix Inventory. For more information, see: [Software requirements \(on page ci\)](#).

- **9.2.14 Support for new virtualization technologies**

- [Oracle VM Server for x86 \(on page cccxliv\)](#) 3.4.x
- Solaris Kernel zones on x86 without nested virtualization (not yet announced as BigFix subcapacity eligible). For more information, see: [Client Installation on Oracle Solaris \(on page ccxii\)](#).
- SLES 11 for PowerVM on Power9 in Power8 compatibility mode.
- SLES 12 for PowerVM on Power9 in native mode.
- SLES 15 on Power9 with disconnected scanner in P9 native mode.

- **9.2.14 Security enhancements**

- To maintain security, Java is upgraded to the newest version: 8.0.5.25.
- WebSphere Application Server Liberty is upgraded to version 18.0.0.3.
- Update of the Xerces library to version 3.2.1 on Solaris Sparc and Solaris x86 to keep the scanner secure and reliable.

- **9.2.14 Information about the session time out after a period of user inactivity**

The session timeout specifies the period of user inactivity after which the session ends, and the user is required to log in again. Before it happens, you are notified that the session is about to end.

Application update 9.2.13

- **9.2.13 Utilization of the Registered User metric for Microsoft Office 365**

BigFix Inventory reports utilization of the Registered User metric for Microsoft Office 365. The information is extracted during regular software scans directly from computers where components of Office 365 are installed. Based on the scan results, the Registered User metric is calculated and its utilization is displayed on the All Metrics report. Detailed information about users is shown on the Software Users report.

Thanks to this feature, you can track licenses for Office 365 in BigFix Inventory to analyze trends, define thresholds and optimize costs by identifying users of unused instances. For more information, see: [Registered User \(on page dlxxxj\)](#).

- **9.2.13 Reporting additional IBM metrics based on the resource utilization data that is delivered by .slmtag files**

Additional IBM metrics that are based on the resource utilization data are displayed on the All IBM Metrics and All Metrics reports. BigFix Inventory scans computers in your infrastructure to find the software components that are reported with .slmtag files and then aggregates their metric quantity. For more information, see: [Reported license metrics \(on page dlvi\)](#).

- **9.2.13 Version currency of Internet browsers and other selected software that changes frequently**

BigFix Inventory automatically detects all versions of a number of software components on Windows. These components include Internet browsers (Chrome, Edge, Firefox, and Opera) as well as components such as Adobe Air, Adobe Flash Player, Slack, and other. The components are detected by generic signatures while their versions are automatically retrieved from files or packages that caused detection. Such an approach allows for avoiding catalog gaps and ensures full discovery currency.

- **9.2.13 Reporting utilization of FlexPoints**

FlexPoint is a license metric unit that can be used to determine the cost of IBM products that are purchased as part of FlexPoint bundles. Each product from a bundle is licensed based on a different license metric but all metrics are converted into FlexPoints. Products that are currently available as part of FlexPoint bundles are products from the IBM Cloud and IBM Analytics brands. For more information, see: [IBM FlexPoints \(on page dlx\)](#), and [Classifying FlexPoint bundles](#).

- **9.2.13 Preview feature: Visibility of Common Vulnerabilities and Exposures (CVEs)**

- Information about CVEs is shown on the Software Classification report on which CVEs are matched with software components through component detailed versions.
- The details of CVEs are extended to include a link to a description of the CVE in the National Vulnerability Database as well as the version of the CVSS, either 2.0 or 3.0.
- Information about CVEs can be filtered and sorted by the CVE name. The filter shows the software that meets the specified criteria by searching through the full list of CVEs of each reported software instance.

Thanks to the improvements in showing information about CVEs you can better identify, monitor, and analyze potentially vulnerable software and at the same time prevent potential threats.

- **9.2.13 New predefined reports available under Software Classification panel**

The Software Classification report is extended with four predefined views:

- Out of Support Software that shows software installations for which support already ended
- Software Approaching End of Support that shows software installations for which support ends within the next three months
- Vulnerable Software that shows software installations for which at least one CVE was matched
- Vulnerable Software - Recent that shows software installations that were discovered within the last two weeks and for which at least one CVE was matched

The views provide you with a quick overview of software that might pose security threats in your environment.

- **9.2.13 Preview feature: Security dashboard and reporting**

The new security dashboard shows a preview of security features in the form of widgets. To view the dashboard, click **Go to New Dashboard & Reporting** on the home page. For more information, see: [Dashboard \(on page liv\)](#).

- **9.2.13 New To Do list on the top navigation bar**

The top navigation bar is extended with a To Do list that, at present, displays the following information under the  icon:

- Failed imports of data
- Availability of a new version of the application
- The need to configure connections to virtualization hosts for some of the computers

The relevant information is indicated with a number that is displayed on the icon. The number is shown until you take action to address the displayed information. For more information, see: [9.2.13 To Do list \(on page lxi\)](#).

- **9.2.13 Information about the shared processor pool**

BigFix Inventory displays information about the identifier of the shared processor pool under the Shared Pool ID column on the Hardware Inventory, Computers, Software Installations, Oracle Databases and Software Classification reports.

- **9.2.13 REST API associations for quick analysis of security risks**

The `/api/sam/v2/software_instances` REST API is extended with two associations:

- `cve` that allows to retrieve information about Common Vulnerabilities and Exposures (CVEs) that were matched with a particular component through its detailed version
- `release_component_lifecycle` that allows to retrieve information about when a component reaches its end of support

The associations provide you with a method for retrieving information about software that might pose security threats by using REST API.

- **9.2.13 Filtering reports and retrieving data through REST API request by using `is empty` and `is not empty` operators**

You can filter reports and retrieve data from specified columns by using the `is empty` and `is not empty` operators.

- **9.2.13 Support for BigFix Inventory server on MS SQL 2016**

BigFix Inventory server is now supported on MS SQL 2016 for fresh installations. If you used MS SQL 2016 with earlier versions of the product, see: [Software requirements \(on page ci\)](#).

- **9.2.13 Discovery of new components from Citrix, Symantec, Oracle and other**

The software catalog is extended to discover a number of new components from Citrix, Symantec, Oracle and other. To discover the components, ensure that you upgrade BigFix Inventory to version 9.2.13.

Below is the list of the newly added components. To learn about their exact versions, use the Software Catalog widget in BigFix Inventory to browse the catalog content.

- Added discovery capability for Citrix products
 - Citrix Edgesight
 - Citrix Presentation Server
 - Citrix XenDesktop
- Added discovery capability for Symantec products
 - Veritas Cluster Enterprise Agents
 - Veritas Cluster Server
 - Veritas Cluster Server Bundled Agents
 - Veritas Cluster Server Disk Reservation Modules and Utilities
- Added discovery capability for Oracle products
 - Oracle Clusterware
 - Oracle GoldenGate
 - Sun Cluster HA for WebSphere MQ

- **9.2.13 List of software components with end of support dates is updated**

The information about the end of support dates is refreshed for your BigFix software components for which the dates were announced since the last update of the list. The data is based on the information contained in BigFix Lifecycle in September.

- **9.2.13 Performance improvements**

- The import of VM Manager data is optimized to decrease the processing time.

- **9.2.13 Security enhancements**
 - Java is upgraded to the newest version to maintain security.
 - TLS 1.2 is enabled in BigFix Java by default.
 - WebSphere Application Server Liberty is upgraded to version 18.0.0.2 to address CVE-2018-1553.
 - Update of the Xerces library to version 3.2.1 on Windows to keep the scanner secure and reliable.
- **Support for new virtualization technologies**
 - vSphere ESXi 6.7 and vSphere vCenter Server 6.7.
- **9.2.12 End of support for migration from License Metric Tool (LMT) and Tivoli® Asset Discovery for Distributed (TAD4D) version 7.x**

Starting from 1st October 2018, migration from License Metric Tool 7.x and Tivoli® Asset Discovery for Distributed 7.x will no longer be supported. The support will not be available regardless of the version of License Metric Tool 9.x or BigFix Inventory 9.x that you are using. For more information, see: [End of support for migrating from LMT and TAD4D version 7.x](#).

Application update 9.2.12

- **9.2.12 Reporting detailed hardware information**

BigFix Inventory introduces a detailed hardware scan that allows you to retrieve detailed hardware information related to memory, operating systems, storage, processors, partitions, network adapters, SMBIOS data, IP addresses and logical processor capacity data. You can use the collected information to report and monitor hardware in your infrastructure. The new scan gives you a clear overview of your physical assets. For more information, see: [Detailed hardware scan \(on page ccxx\)](#).
- **9.2.12 Preventing security threats with information about Common Vulnerabilities and Exposures (CVE) added to the software catalog**

Common Vulnerabilities and Exposures (CVE) is a list of known security threats that are assigned identification numbers. BigFix Inventory uses CVE that is provided by the National Vulnerability Database to help you identify potential threats in your environment. In this release, BigFix Inventory presents initial integration with National Vulnerability Database as a preview feature. Thanks to the visibility of potentially vulnerable software, it is possible to identify and prevent potential threats. For more information, see: [Preview: Checking Common Vulnerabilities and Exposures \(CVEs\) \(on page dclxv\)](#).
- **9.2.12 Retrieving information about the discovered software component as specified in the Common Platform Enumeration (CPE) dictionary**

Common Platform Enumeration is a standardized naming scheme for software. This information is available in REST API for integration scenarios. BigFix Inventory uses the CPE dictionary to match CVEs and indicate known vulnerabilities in software products. You can analyze and compare the CPEs based on the data feeds provided by BigFix Inventory with other sources, such as CVEs that are available in the National Vulnerability Database. For more information, see: [component_cpe association \(on page mcxv\)](#).
- **9.2.12 Enhancements to reporting end of support dates**

The end of support dates for selected BigFix software products and components for which the information has been announced are now available on the Software Classification panel. Additionally, a new Predefined End of Support column is available on the Software Components report. The column indicates whether the end of support date was provided by BigFix or was specified manually by a user.

- **9.2.12 Using Review Lite script for extended discovery of Oracle databases**

Review Lite is a standard script used by the Oracle auditors to control the number of Oracle licenses within a company. BigFix Inventory allows you to automatically run this script on all computers in your environment. The results are interpreted by BigFix Inventory which provides you with an overview of the report that you are required to deliver to the auditors. For more information, see: [Discovering usage of Oracle Database features \(on page div\)](#).

- **9.2.12 Automatic scanning of remote shared disks**

Until now, remote shared disks could only be scanned with a manual procedure. With the newest release of BigFix Inventory, you can optimize this process and set up automatic scans. As a result, a single computer is designated to scan a specific shared disk and discover the installed software. The data is then automatically populated to all computers on which the same shared disk is mounted. For more information, see: [Discovering software on shared disks \(on page ccxxi\)](#).

- **9.2.12 Collecting utilization data of SAP license metrics with License Administration Workbench 2.0**

Support for collecting SAP license metric utilization data is extended to License Administration Workbench (LAW) 2.0. LAW 2.0 allows the creation and storage of multiple consolidations. SAP Metric Data Collector uses these consolidations to display the SAP utilization data in BigFix Inventory. For more information, see: [Collecting utilization of SAP license metrics \(on page dxv\)](#).

- **9.2.12 Flexible license management by specifying the type of relation between a component and a product as charged or not charged based on the specific license agreement**

BigFix Inventory allows you to specify the individual relation between a product and a component as charged or not charged. For example, when the existing relation between the component and the product is charged, but your license agreement states that the component is not charged.

- **9.2.12 Refreshing scan data on a computer through the Support Page**

When the software scan, capacity scan, or VM Manager Tool scan finishes successfully on a particular computer but its results are not uploaded to the BigFix server, you can force upload of the data. You can do it directly from BigFix Inventory without the need of accessing the BigFix console and running a separate fixlet for each type of scan. For more information, see: [Troubleshooting problems with computers \(on page dccliii\)](#).

- **9.2.12 Resuming a connection to the VM manager**

You can resume a connection to a VM manager that is in the *Invalid credentials - suspended* status by clicking the Resume button on the VM managers panel. For more information, see: [Understanding the VM Managers panel \(on page cccxlvii\)](#).

- **9.2.12 Information about full version of already discovered files**

Scanned File Data report displays the full version of discovered files. After upgrade, the data is imported to BigFix Inventory during a number of consecutive data imports to avoid performance issues. Each consecutive import contains data from 10000 computers.

- **9.2.12 Extension of audit snapshot**

Audit snapshot is extended with the following information.

- The PVU audit snapshot package is extended with an `audit_snapshot_summary.csv` report. The report contains a summary of all software products that were installed in your environment during the period for which the snapshot is generated and which are charged. For more information, see: [Content of the audit snapshot \(on page dxcvii\)](#).
- Charged column is added to the audit snapshot report. It indicates whether the relation between the component and the product is charged based on the license terms.

- **9.2.12 A new link to Learning Resources under Help **

The list of quick links under Help  is updated with the link to Learning Resources on the BigFix Security Learning Academy. The site consists of video tutorials for the product.

- **9.2.12 Security enhancements**

- Java upgrade which addresses the following CVEs: CVE-2018-2633, CVE-2018-2618, CVE-2018-2603, CVE-2018-2602, CVE-2018-2579, CVE-2018-1417, CVE-2018-2783, CVE-2018-2794.
- Update of the Xerces library to version 3.2.1 on Linux to keep the scanner secure and reliable.

- **9.2.12 Support for new virtualization technologies**

- Power VM on Power9 on AIX, IBM i and Linux.
- KVM on IBM Power8.
- Management of Citrix Xen through VM manager.

- **9.2.12 Deprecated content**

The following reports and panels are deprecated with this application update:

- Usage Properties
- Package Properties
- UNIX Package Properties
- Current State Values

Application update 9.2.11

- **9.2.11 Improved license planning and prevention of security risk by setting up and monitoring end of support dates for specific software**

BigFix Inventory provides end of support dates for BigFix software components as announced in BigFix Lifecycle. You can also specify the dates manually, both for IBM and non-IBM components. You can use this

information to easily define license demand for the future. You can also use it for security purposes. For example, to determine whether software installed on a computer that is under investigation is still supported.

- **9.2.11 Utilization of the Virtual Processor Core (VPC) metric for BigFix products**

BigFix Inventory measures utilization of the Virtual Processor Core (VPC) metric. The metric is based on the number of virtual cores (vCPUs) that are available to the product. VPC consumption is reported on the physical server level as well as across servers similarly to other license metrics such as PVU. For more information, see: [IBM Virtual processor core \(VPC\) \(on page dlxxxvi\)](#).

- **9.2.11 Introduction of the Package Data Summary report**

The Package Data report lists software packages that are discovered on the computers in your infrastructure. You can find out what is the name of every package, its version, vendor, description, and type. Every occurrence of the package is listed as a separate row on the report.

This application update introduces the Package Data Summary report. The report provides an aggregated view of the discovered packages with a number of times each package appears in the infrastructure. This number is a hyperlink that brings you to the detailed list of computers where the package is deployed.

From the Package Data Summary report, you can generate software signatures with one click. You can also select multiple packages and create signatures in bulk. As a result, the process of creating signatures from package data is simplified. You can easily create software signatures for up to 20 components at once. Additionally, the report provides information about which packages were already used as signatures so you can focus only on the remaining ones. For more information, see: [Available reports \(on page dcxvi\)](#) and [Creating signatures from package data \(on page cdxcii\)](#).

- **9.2.11 Reporting information about the detailed version of Windows applications**

Installed software components that are detected based on the presence of specific files are reported with a detailed version which is the same as the file version. The information is shown in the Component Detailed Version column on the Software Classification report. The information can help you decide whether versions of installed software components are covered by your entitlements and are in line with the company policy. It can also be used for security purposes. For example, to check whether the latest security patches are installed or whether versions with vulnerabilities exist in the infrastructure.

- **9.2.11 Discovery of new components from Windows, Corel, Adobe, Symantec and SAP**

The software catalog is extended to discover a number of new components from Windows, Corel, Adobe, Symantec and SAP. To discover the components, ensure that you upgrade BigFix Inventory to version 9.2.11.

Below is the list of the newly added components. To learn about their exact versions, use the Software Catalog widget in BigFix Inventory to browse the catalog content.

- Added discovery capability for Windows products
 - Windows Defender
- Added discovery capability for Corel products
 - CorelDraw Graphic Suite

- Added discovery capability for Adobe products
 - Adobe Flash Player
 - Adobe Lightroom CC
 - Adobe Lightroom Classic
- Added discovery capability for Symantec products
 - Altiris Deployment Solution Agent
 - Altiris Patch Management Agent
 - Symantec Encryption Desktop
 - Veritas Enterprise Vault
- Added discovery capability for SAP products
 - Sybase Central
- **9.2.11 Simplification of updating the software catalog**

The process of updating the software catalog is simplified. You simply need to run the Software Catalog Update task that downloads the software catalog, and automatically uploads it to the BigFix Inventory server during the next import of data. Manual upload of the software catalog through the BigFix Inventory user interface is no longer needed. For more information, see: [Updating the software catalog to ensure accuracy of software discovery and reporting \(on page cdxliv\)](#).

- **9.2.11 Optimizing data imports thanks to enabling a new schema of the BigFix database**

The process of importing data from BigFix to BigFix Inventory is optimized thanks to using an improved schema of the BigFix database. The improved schema is enabled by default in fresh installations and during upgrades from earlier versions. The only requirement is to have the BigFix server in version 9.5.5 or higher.

- **9.2.11 Preventing accumulation of unprocessed data in case of failing imports**

When the imports of data are failing for a longer period of time, you can enable skipping certain import steps. In such a case, the import continues even if these steps fail. When the import finishes, it is marked as partial because only data from successfully completed steps is processed. Enabling partial imports allows for quicker system recovery after a problematic import and prevents accumulation of unprocessed data. For more information, see the description of the **enable_partial_imports** parameter in: [Advanced server settings \(on page cdii\)](#).

Application update 9.2.10

- **9.2.10 Reporting usage of products such as Microsoft Office at suite level aggregated per computer**

The Usage per Computer report shows product usage aggregated per computer so that you can easily identify where the product is deployed and how frequently it is used. Based on this information, you can decide whether you want to extend the license for this product or not.

For product suites such as Microsoft Office, the usage is reported for the entire suite rather than separately for Word, PowerPoint, Excel, and other components. By default, the report is calculated regularly after every

import, but you can also configure it to be calculated on demand. For more information, see: [Available reports \(on page dcxvi\)](#).

- **9.2.10 Collecting and measuring utilization of SAP license metrics**

This application update provides an end-to-end automation process for scheduling generation and retrieval of SAP license utilization reports by leveraging the provided ABAP plug-in. The process can now be handled entirely by the Administrator of BigFix Inventory without the need to engage the SAP Administrator.

Once the SAP license utilization reports are processed by BigFix Inventory, the related SAP license usage data is available on the Resource Utilization report. For more information, see: [Measuring and collecting utilization of SAP license metrics \(on page dxx\)](#).

- **9.2.10 Adding products to the software catalog**

You can easily add a custom product and the license metric that it uses to the software catalog. You can then assign the relevant components to this product so that your software classification is correct. For more information, see: [Adding a new product to software catalog \(on page cdxcvi\)](#).

- **9.2.10 Information about imported part numbers that influenced the initial classification is visible on the Software Classification panel**

Part numbers that are available on Passport Advantage represent products that you purchased and license metrics that these products use. When you upload part numbers to BigFix Inventory, they affect the default assignment of components. Information about part numbers which caused that a component was assigned to a particular product is displayed on the Software Classification panel. The information gives you a better understanding of rationale behind the default classification that was proposed.

- **9.2.10 REST API for advanced integration and data sharing**

You can retrieve information about utilization of license metrics, query and set thresholds of the utilization as well as query and set values of custom fields by using the `api/sam/v2/license_usage` REST API. For more information, see: [Retrieval of license metric utilization \(v2\) \(on page mix\)](#) and [Setting license metric thresholds and custom field values \(v2\) \(on page mxix\)](#).

You can also retrieve computer properties, including custom computer properties, by using the `computer_details` association for the `api/sam/v2/computers` REST API. For more information, see: [computer_details association \(on page mcxviii\)](#).

- **9.2.10 New computer status - Outdated VM Manager Data**

When hardware inventory data is not retrieved from a VM manager for more than 30 days, default PVU counting is applied and higher utilization might be reported. To allow the Infrastructure Administrator to troubleshoot problems with VM manager connection before they affect the reports, the *Outdated VM Manager Data* status is introduced. The status indicates that there are problems but license metric utilization is not yet affected. For more information, see: [Computer statuses \(on page cccxvi\)](#).

- **9.2.10 Stand-alone scanner for monitoring isolated HP-UX computers**

If some of your HP-UX computers cannot be managed from a central location due to security reasons, you can discover software and hardware inventory by using disconnected scans. Scan results are placed in files that can be later on used to import the collected data to the BigFix Inventory server. For more information, see: [Discovering software and hardware with disconnected scanner on Windows and UNIX \(on page ccxlviii\)](#).

- **9.2.10 Information about the number of partition cores**

The number of partition cores that were assigned to a virtual machine when the peak of the metric value occurred is shown in the Partition Cores column on the All Metrics report and in the audit snapshot. You can use the information to better understand how particular VMs contributed to license metric consumption. For more information, see: [Report columns \(on page dcxxx\)](#).

- **9.2.10 Discovery of new components from HP, SAP, BMC and other**

The software catalog is extended to discover a number of new components from Oracle, SAP, and Red Hat. To discover the components, ensure that you upgrade BigFix Inventory to version 9.2.10.

Below is the list of the newly added components. To learn about their exact versions, use the Software Catalog widget in BigFix Inventory to browse the catalog content.

- Added discovery capability for HP products
 - HP Application Lifecycle Management
 - HP Business Process Monitor
 - HP GlancePlus
 - HP OpenView Internet Services
 - HP OpenView Performance Agent
 - HP-UX Image Viewer Subsystem
 - HP-UX LAN Link Product
 - Mercury Business Availability Center
 - Mercury Interactive Quick Test Professional
- Added discovery capability for SAP products
 - BusinessObjects Enterprise
 - Crystal Reports .NET Server
 - SAP BusinessObjects WebIntelligence
 - SAP Crystal Reports
 - SAP SQL Anywhere Network Database Server
 - Sybase PowerDesigner
 - Sybase Software Developer Kit
- Added discovery capability for BMC products
 - BMC Atrium Orchestrator
 - BMC Batch Impact Manager
 - BMC BladeLogic RSCD Agent
 - BMC Control-M Agent
 - BMC Control-M CM For Advanced File Transfer

- Added discovery capability for CA products
 - CA Erwin Data Modeler
 - CA Application Management .NET Agent
 - CA ARCserve Backup Diagnostic Utilities
 - CA Clarity and Portfolio Management
 - CA ControlMinder
 - CA eHealth
 - CA Erwin Data Modeler License Server
 - CA Erwin Mart Server
 - CA Plex
 - CA Process Automation Orchestrator
 - CA SiteMinder
 - CA SiteMinder Web Agent
 - CA Spectrum
 - CA Unicenter CA-XCOM Data Transport
 - CA Unicenter DSM Agent + Remote Control Plugin
 - CA Unicenter NSM
 - CA Unicenter Software Delivery
 - CA Universal Job Management Agent
- Added discovery capability for Tibco products
 - Tibco ActiveMatrix Adapter for Database
 - Tibco ActiveMatrix BusinessWorks Service Engine
 - TIBCO Designer
 - Tibco Hawk
 - TIBCO iProcess Workspace
 - TIBCO Spotfire
 - TIBCO Third-Party Core Libraries
- Added discovery capability for Veritas products
 - NetBackup Bare Metal Restore
 - NetBackup for Microsoft Exchange Server
 - NetBackup for Microsoft SQL Server
 - NetBackup for SAP
 - NetBackup for Vmware
 - NetBackup Remote Administration Console
 - Symantec System Recovery

- **9.2.10 Optimizing the import of data**

You can optimize the process of importing data from BigFix to BigFix Inventory. The option is available for BigFix 9.5.5 and higher. To enable the feature, change the value of the **schema_next** parameter to true. For more information, see: [Advanced server settings \(on page cdii\)](#).

- **9.2.10 Serviceability improvements**

To avoid a situation in which the import of data fails due to deadlocks of SQL queries, the step during which the deadlock occurs is retried up to 5 times before the import fails.

- **9.2.10 Support for Docker containers under Red Hat OpenShift**

BigFix Inventory supports software discovery on Docker containers under the Red Hat OpenShift container application platform. For more information, see: [Discovery of software in Docker containers](#).

Application update 9.2.9

- **9.2.9 Reporting of SAP license usage with integration of SAP License Administration Workbench (SLAW) and SAP USMM**

You can have reports of SAP license usage consolidated into BigFix Inventory by integration with the SLAW and USMM tools. The measurements from all SAP systems in the SAP landscape are transferred to the central SAP and combined into one SLAW report. You can integrate the report into BigFix Inventory so that utilization of license metrics for SAP products is shown on the Resource Utilization report. The capability is available for all SAP products that can be measured by SAP SLAW and USMM tools. For more information, see: [Collecting utilization of SAP license metrics \(on page dxv\)](#).

- **9.2.9 Discovery of Oracle Database extended to Solaris**

Reporting of Oracle Databases is extended to Solaris 10 and 11 for all discoverable versions. You can retrieve granular information about the edition (Standard or Enterprise), options, Management Packs (including Oracle Real Application Cluster, RAC), and user concurrent sessions of Oracle Databases that are deployed in your environment. For more information, see: [Extended discovery of Oracle Database \(on page dii\)](#).

- **9.2.9 Discovery of new components from Oracle, SAP, and Red Hat**

The software catalog is extended to discover a number of new components from Oracle, SAP, and Red Hat. To discover the components, ensure that you upgrade BigFix Inventory to version 9.2.9.

Below is the list of the newly added components. To learn about their exact versions, use the Software Catalog widget in BigFix Inventory to browse the catalog content.

- Added discovery capability for Oracle products
 - Oracle Business Process Management Suite
 - Oracle Database Gateway for APPC
 - Oracle Database Gateway for Informix
 - Oracle Database Gateway for Microsoft SQL Server
 - Oracle Database Gateway for ODBC
 - Oracle Database Gateway for Sybase
 - Oracle Database Gateway for Teradata
 - Oracle Database Gateway for WebSphere MQ
 - Oracle JRockit
 - Oracle Procedural Gateway for APPC

- Oracle Transparent Gateway for DRDA
- Oracle Transparent Gateway for Microsoft SQL Server
- Oracle Transparent Gateway for Sybase
- Oracle Transparent Gateway for Teradata
- Added discovery capability for SAP products
 - SAP Crystal Server
 - Sybase Adaptive Server Enterprise
 - Sybase Enterprise Connect Data Access
 - Sybase IQ Server
 - Sybase Replication Server
- Added discovery capability for Veritas products
 - Veritas NetBackup Client
 - Veritas NetBackup Media Server
- Added discovery capability for CA Technologies products
 - CA Unified Infrastructure Management
 - CA Unified Infrastructure Management Snap
- Added discovery capability for Red Hat Technologies products
 - Red Hat JBoss Data Grid
 - Red Hat JBoss Enterprise Application Platform
 - Red Hat Satellite Server
- Added discovery capability for AXWAY products
 - Axway Transfer CFT
 - Axway Transfer Interpel
- Added discovery capability for RES products
 - RES ONE Automation Agent
 - RES ONE Automation Console
 - RES ONE Automation Dispatcher
- **9.2.9 Adding custom fields to license metric reports**

You can extend information that is related to software products by adding custom fields on the All Metrics report and its derivatives. The information can be of any type and therefore can help you better identify ownership and purpose of the discovered products as well as license management information such as purchase orders, product owners or other. After the information is added, you can use it to create filtering and grouping criteria such as contract name, placement, allocation, date, or other. For more information, see: [Adding custom information to license metric reports \(on page dxciiii\)](#).

- **9.2.9 Reporting information about the detailed version of software components**

The installed software components that are detected with software identification tags are reported with the detailed version, on condition that the tag collects this information. The information is shown in the Component Detailed Version column. It can help you decide whether versions of the installed software instances are covered by your entitlements and are in line with the company policy. It can also be used for security purposes, for example to check whether the latest patches are installed for security purposes. For the

components that are discovered with other detection methods, the values in Component Detailed Version and Component Version columns are duplicated.

- **9.2.9 Stand-alone scanner for monitoring isolated AIX computers**

If some of your AIX computers cannot be managed from a central location due to security reasons, you can discover software and hardware inventory by using disconnected scans. Scan results are placed in files that can be later on used to import the collected data to the BigFix Inventory server. For more information, see: [Discovering software and hardware with disconnected scanner on Windows and UNIX \(on page ccxlviii\)](#).

- **9.2.9 Filtering the Package Data report to view the packages that can be used to create new software signatures**

Package Data report is extended with information about which of the packages that are discovered in your infrastructure are already used as software signatures. Thus, it is easier to identify which packages can be used to extend software discovery. To narrow down the Package Data report to packages that can be used to create new software signatures, filter the Package Data report by the Recognized and Caused Detection columns.

- **9.2.9 Improved security of storing VM manager passwords**

You can overwrite the default key that is used to encrypt VM manager passwords or change the default password to the VM Manager Tool keystore. For more information, see: [Improving security of storing VM manager passwords \(on page dccxl\)](#).

- **9.2.9 Configuring LDAP server that uses load balancer to preserve server certificates**

If your LDAP server uses load balancer that dynamically changes the list of hosts, and the connection between LDAP and the BigFix Inventory server is secure, set up the BigFix Inventory server to preserve certificates of the LDAP server.

- **9.2.9 User interface improvements**

The following user interface improvements are introduced:

- The look and feel of the application is improved to be more in line with other BigFix applications.
- A panel that is used for configuring the collection of data that is displayed on the Unrecognized Scan Data report is removed. Collection of the data is now enabled by setting a parameter that is available on the Advanced Server Settings panel.

- **9.2.9 Scanner improvements**

The following scanner improvements are introduced:

- The scanner provides support for a wider range of national characters. Character encoding in scan outputs and in scan configuration files is in the UTF-8 format.
- 64-bit version of the scanner is available for AIX. The 32-bit version is removed.

- **9.2.9 Version of the application build visible in PDF reports**

The PDF reports that are generated in BigFix Inventory contain detailed information about the version of the application. The information can be helpful for troubleshooting purposes.

- **9.2.9 Software catalog enhancements**

The Software Catalog column is replaced with a set of columns that define the organization that provides the definition of a component, product, publisher, release, signature and version. For more information, see: [Report columns \(on page dcxxx\)](#).

- **9.2.9 Changes in PVU calculation on PowerKVM**

Rules of counting subcapacity licenses on PowerKVM systems were clarified and multithreading no longer impacts the number of reported cores. BigFix Inventory is adjusted to the change and thus the number of PVUs that it reports is lower. For more information, see: [Subcapacity \(Virtualization\) License Counting Rules](#).

To ensure that the change in counting rules is reflected in BigFix Inventory, restart the Run Capacity Scan on Virtualization Hosts fixlet and run the import of data. The reported value might not change immediately, as the report shows the peak value.

Application update 9.2.8

- **9.2.8 Utilization of the Single and Dual Physical Processor metrics for Microsoft Windows Server**

BigFix Inventory measures utilization of the Single Physical Processor and Dual Physical Processor metrics for Microsoft Windows Server. You can use the information to calculate the license demand for this product. The feature is supported for Windows Server 2008 R2 Datacenter, Windows Server 2012 Datacenter, and Windows Server 2012 R2 Datacenter. For more information, see: [Microsoft Single and Dual Physical Processor \(on page dlxii\)](#).

- **9.2.8 Utilization of the Oracle Processor Core metric for Oracle Databases**

Reporting of Oracle Databases is enhanced by processing the core factor and measuring utilization of the Oracle Processor Core metric. You can use the information to calculate the license demand for Oracle Databases without the need to manually calculate core factors. The feature is supported for all discoverable versions of Oracle Database (Enterprise Edition). For more information, see: [Oracle Processor Core \(on page dlxx\)](#).

- **9.2.8 Reporting software installed on shared disks based on software templates**

To ensure completeness of software discovery, it is necessary to collect data from shared disks. The solution that was available in the previous application updates could be resource-consuming when shared disks were used extensively. The new solution is based on creating software templates which allows for limiting the impact of the shared disks scan on the performance of the infrastructure.

To report software that is installed on a shared disk, one of the computers on which the disk is mounted is scanned to create a software template. Then, the template is replicated within a computer group that contains

other computers on which the same disk is mounted. For more information, see: [Discovery of software on shared disks](#) and [Step 2b: Optimized mode - Manual scanning of remote shared disks \(on page ccxxvii\)](#).

- **9.2.8 Collecting executable files based on application usage**

To improve software discovery on Linux and Solaris systems, you can use files associated with processes that have been run on the computers in your infrastructure, regardless of their extension. These files, with supplementing information such as path and size, are available on the Scanned File Data report and can be used to create custom discovery and usage signatures for software components. For more information, see: [File system scan \(on page ccvi\)](#).

To use this feature, you require the BigFix server and client in version 9.5.5 or higher. Collecting the files is enabled by default for new installations. If you are upgrading from earlier versions, restart the Initiate Software Scan actions, and start them again by using the new version of the fixlet. Ensure that you run the file system scan and application usage statistics scan.

- **9.2.8 Version 2 of REST APIs for software and hardware inventory**

New REST APIs for software inventory (`api/sam/v2/software_instances`) and hardware inventory (`api/sam/v2/computers`) provide the following improvements:

- Historical information about uninstalled software and removed computers
- Complete information about software classification, including product to which the component is assigned and its license metric as well as the status of confirmation, exclusion, and suppression
- Associations that allow for retrieving additional hardware inventory information, computer health or usage data with a single REST API call
- Improved security to ensure that a user who has access to a subset of computers can retrieve data only from these computers

The APIs are planned to substitute the following existing APIs: `software_inventory` and `computer_systems`. For more information, see: [Retrieval of software inventory \(v2\) \(on page cmxcii\)](#) and [Retrieval of hardware inventory \(v2\) \(on page mxxx\)](#).

- **9.2.8 Assigning a new license metric to an existing product**

When your license terms define that a particular product uses a specific license metric but such a relation is not available in the software catalog, you can assign any of the license metrics available in BigFix Inventory to any of the products. It gives you more flexibility to accommodate the assignment of license metrics to products in accordance with your contract terms. For more information, see: [Assigning an additional metric to a product \(on page cdxcvi\)](#).

- **9.2.8 Configuration of user account lockout**

To increase BigFix Inventory security, you can set conditions under which user account is locked. You can set the maximum number of failed login attempts, the period within which these attempts must occur for the account to be locked, and the period for which the account is locked. By default, user account is locked for 5

minutes after the user attempts to log in to BigFix Inventory more than 10 times within 5 minutes. For more information, see: [Configuring user account lockout \(on page dccxliv\)](#).

- **9.2.8 Stand-alone scanner for monitoring isolated Solaris computers**

If some of your Solaris computers cannot be managed from a central location due to security reasons, you can discover software and hardware inventory by using disconnected scans. Scan results are placed in files that can be later on used to import the collected data to the BigFix Inventory server. For more information, see: [Discovering software and hardware with disconnected scanner on Windows and UNIX \(on page ccxlviii\)](#).

- **9.2.8 Automatic decommissioning of computers that are inactive for a specified period of time**

When infrastructure management policy makes a process of deleting decommissioned clients in BigFix difficult or impossible, you can enable automatic decommissioning of inactive computers directly in BigFix Inventory. Inactive computers are then decommissioned after a specified period of time, and are no longer included in metric calculation. For more information, see: [Automatically decommissioning inactive computers \(on page cdxxxvii\)](#).

- **9.2.8 Fixlet for collecting logs to troubleshoot problems with computers**

If you are asked to provide data for troubleshooting problems with computers in your infrastructure, you can run the Collect Logs from Endpoints fixlet, and provide the generated output to BigFix Support. The output contains log files needed for troubleshooting purposes. For more information, see: [Troubleshooting problems with computers \(on page dccliii\)](#).

- **9.2.8 Viewing all signatures that are defined for a component in the software catalog**

When a component that is installed in your infrastructure is not discovered, you can check signatures that are defined for that component in the software catalog. You can then use this information to troubleshoot software discovery. To find the list of signatures, use the Software Catalog widget and search for the component. The list of all signatures is displayed on the component details page. For more information, see: [Troubleshooting software discovery \(on page dccl\)](#).

- **9.2.8 Improvements of the VM Manager Tool**

VM Manager Tool is further enhanced to attempt to automatically fix common problems with the VM manager connection parameters. For example, it tries using a default port or a different URL suffix as well as changing the VM manager type. It also tries to correct more than one mistake at a time.

- **9.2.8 Improved browsing of the software catalog content**

Browsing the Software Catalog report is improved so that it is more intuitive and consistent.

- **9.2.8 Serviceability improvements**

To improve serviceability of the application, the following enhancements are introduced:

- BigFix Inventory can reconnect to its database when the active connection was reset or closed without the need of restarting the application server.
- Logging of the disconnected scanner on IBM i is improved.

- **9.2.8 Information about supported operating systems can be dynamically generated through SPCR**

Information about operating systems and hypervisors that are supported by each component of BigFix Inventory including the server, the BigFix client, and disconnected scans is moved to the [Software Product Compatibility Reports](#) (SPCR). You can dynamically generate reports with information about specific operating system or hypervisor, as well as selected components. To view predefined reports generated per component and per operating system, see: [System requirements \(on page c\)](#).

- **9.2.8 Support for new operating systems and public clouds**

You can install the BigFix Inventory server on Windows Server 2016. You can also identify computers as running on the Google Compute Engine public cloud. For more information, see: [System requirements \(on page c\)](#).

- **9.2.8 Discovery of new components from Oracle, SAP, and Red Hat**

The software catalog is extended to discover a number of new components from Oracle, SAP, and Red Hat. To discover the components, ensure that you upgrade BigFix Inventory to version 9.2.8.

Below is the list of the newly added components. To learn about their exact versions, use the Software Catalog widget in BigFix Inventory to browse the catalog content.

- Added signatures for Oracle products
 - Oracle Linux
 - Oracle Data Integrator
 - Oracle JDeveloper
 - Oracle Service Bus
- Added signatures for SAP products
 - SAP SQL Anywhere Personal Database Server
 - SAP SQL Anywhere Network Database Server
 - SAP SQL Anywhere High Availability Option
- Added signatures for Red Hat products
 - Red Hat JBoss Fuse
 - Red Hat JBoss Developer Studio
 - Red Hat JBoss Enterprise Application Platform
 - Red Hat JBoss Web Server
- Added signatures for TIBCO Rendezvous
- Added signatures for SUSE Linux Enterprise Server for SAP Applications

Application update 9.2.7

- **9.2.7 New Software Classification panel replaces the BigFix Software Classification panel**

The Software Classification panel introduces a new way of looking at software inventory and relations between the discovered components and licensable products. It provides a flat structure with software

installations broken into single components, which makes it easier to view software assets and manage relations between them. The new panel provides functions such as reassigning, suppressing and sharing components, excluding products as well as confirming software classifications. The panel replaces the existing BigFix Software Classification panel. For more information, see: [Software Classification panel \(on page dxlii\)](#).

- **9.2.7 Creation of custom classification rules**

When you are assigning a component, or excluding a product, you can create a custom classification rule that will be applied when the same component or product is discovered in the future. It allows for automating actions on analogical components or products, and facilitates the process of software classification. For more information, see: [Creating and managing custom rules \(on page dl\)](#).

- **9.2.7 Utilization of the Physical Core and Virtual Core metrics for Microsoft SQL Server with Software Assurance**

BigFix Inventory measures utilization of the Physical Core and Virtual Core metrics for Microsoft SQL Server. You can use this information to calculate the demand for Microsoft SQL Server licenses. The feature is supported for Microsoft SQL Server 2012, 2014, and 2016 (Standard and Enterprise edition) with Software Assurance. For more information, see: [Microsoft Physical Core with SA \(on page dlxiv\)](#) and [Microsoft Virtual Core with SA \(on page dlxvii\)](#).

- **9.2.7 Utilization of license metrics for VMware vCenter and vSphere**

You can retrieve information about license metric utilization for VMware vCenter and vSphere. You can use this information to calculate the demand for VMware licenses. The data is collected by the VM Manager Tool and is displayed on the Resource Utilization report. For more information, see: [Measuring license metric utilization of VMware products \(on page dx\)](#).

- **9.2.7 Discovery of SAP components**

You can retrieve information about SAP components that are installed on Linux and AIX computers in your infrastructure. The discovery is supported for all SAP components that are based on SAP NetWeaver 7.x. For more information, see: [Discovering SAP components \(on page dxiv\)](#).

- **9.2.7 Utilization of the Install Instances metric for BigFix products**

The All Metrics report is extended to include information about the utilization of the Install Instances metric for BigFix products. You can see not only the current number of installed product instances, but also how that number changed over time. You can also set a threshold, to easily monitor when the number of installed instances exceeds your expectations. For more information, see: [License metric utilization \(on page dlxxxvii\)](#), and [Reported license metrics \(on page dlvi\)](#).

- **9.2.7 Overview of license metrics that are used by BigFix products in your infrastructure**

The All Metrics report is extended to include not only products with license metrics for which utilization is calculated, but also all other BigFix products in your infrastructure. Even though metric utilization might not be calculated for these products, it gives you an overview of all IBM metrics for which you should have licenses.

For more information, see: [License metric utilization \(on page dlxxxvii\)](#) and [Reported license metrics \(on page dlvj\)](#).

- **9.2.7 Discovery of software and hardware inventory with disconnected scans on Windows and Linux**

If there is no direct connection between some of the Windows or Linux computers in your infrastructure and the BigFix server, you can discover software and hardware inventory by using disconnected scans. Scripts that are provided in the disconnected scanner package initiate software and capacity scans, and prepare scan results that you later upload to BigFix Inventory. For more information, see: [Discovering software and hardware with disconnected scanner on Windows and UNIX \(on page ccxlviii\)](#).

- **9.2.7 REST API for adding and changing VM managers**

If you have a large number of VM managers to define, you can use REST API requests to add and change VM managers in your infrastructure. For more information, see: [REST API for managing VM managers \(on page cmxlix\)](#).

- **9.2.7 Adding custom part numbers to the catalog**

Custom part numbers can be assigned to products with custom deals, or any other software that is not included in the software catalog by default. Custom part numbers that you add are matched against the part numbers that you obtain from Passport Advantage to improve the accuracy of automated bundling. You can check which of your uploaded part numbers were matched against the catalog part numbers or custom part numbers on the Software Catalog report. For more information, see: [Part numbers \(on page dxxxiv\)](#) and [Adding custom part numbers \(on page dxxxvi\)](#).

- **9.2.7 Extension of two bundling rules: bundling tags and part numbers**

The bundling tags and part numbers rules are extended to increase the accuracy of automated bundling. The rules now use information contained in part numbers to assign components not only to proper products but also to proper license metrics. For more information, see: [Bundling tags \(on page dxxx\)](#) and [Part numbers \(on page dxxxiv\)](#).

- **9.2.7 Two-phase matching of file signatures**

When a file signature used for discovering a component on Windows had size and version specified, the component was discovered only when both properties matched. Applications that are frequently patched required separate signatures for every patch. Currently, if a patch modifies only the size of the file signature, the signature will cause software discovery.

- **9.2.7 Configuring security policy for user passwords**

If your company has a security policy for user passwords, or you want to improve the application security, you can configure BigFix Inventory to require that user passwords fulfill specific requirements. By default, user passwords are required to be at least 8 characters long and contain one lowercase character, one uppercase character, and one digit. For more information, see: [Configuring security policy for user passwords \(on page dccxxvi\)](#).

- **9.2.7 Improvements of the VM Manager Tool**

VM Manager Tool is enhanced to attempt to automatically fix the most common problems with the VM manager connection parameters such as adding default URL suffix, correcting the user name format, and testing the WinRM and PowerShell protocols for Hyper-V VM managers. For information about disabling this functionality, see: [VM Manager Tool settings \(on page ccclxxxii\)](#).

- **9.2.7 New look and feel of the user interface**

Look and feel of the user interface is updated for improved experience. The possibility of selecting all columns while configuring the report view is added.

- **9.2.7 Support for new operating systems and virtualization technologies**

You can use Microsoft Hyper-V Server 2016, Xen Hypervisor, Oracle VM, RHEV-M 3.6 as well as VMware vCenter, and ESXi 6.5 as supported virtualization technologies on selected operating systems. You can also install the BigFix client on AIX VIOS. For more information, see: [System requirements \(on page c\)](#).

- **9.2.7 Discovery of new components from Microsoft, SAP, and Oracle**

The software catalog is extended to discover a number of new components from Microsoft, SAP, and Oracle. To discover the components, ensure that you upload the software catalog published in March 2017, and upgrade BigFix Inventory to version 9.2.7.

Below is the list of the newly added components. To learn about their exact versions, use the Software Catalog widget in BigFix Inventory to browse the catalog content.

- Components of Microsoft SQL Server
 - Analysis Services
 - Integration Services
 - Master Data Services
 - Reporting Services
- Components of SAP
 - SAP Crystal Reports
 - SAP Crystal Reports Viewer
- Components of Oracle
 - Oracle Adapter
 - Oracle B2B
 - Oracle BPEL Process Manager
 - Oracle Business Process Management Suite
 - Oracle Business Rules
 - Oracle Coherence
 - Oracle Human Workflow
 - Oracle Mediator
 - Oracle SOA Suite
 - Oracle Virtual Assembly Builder

- Oracle WebLogic Server Enterprise Edition
- Oracle WebLogic Suite
- Oracle WebLogic Server
- Service Architecture Leveraging Tuxedo
- Tuxedo Jolt
- Tuxedo Message Queue
- Tuxedo Server
- Tuxedo System and Application Monitor Plus
- MySQL Connector C
- MySQL Workbench
- TIBCO Enterprise Message Service

Application update 9.2.6

- **9.2.6 Preview Feature: Utilization of the physical core-based metric for Microsoft SQL Server**

BigFix Inventory measures utilization of the physical core-based licenses for Microsoft SQL Server. You can use this information to calculate the demand for Microsoft SQL Server licenses. The feature is supported for versions 2012, 2014, and 2016 (Standard and Enterprise edition). For more information, see: [Microsoft Physical Core with SA \(on page dlxiv\)](#).

- **9.2.6 Preview Feature: Extension of the Software Classification panel**

The Software Classification panel introduces a new way of looking at software inventory and relations between the discovered components and licensable products. It provides a flat structure with software installations broken into single components, which makes it easier to view software assets and manage relations between them. The preview panel is extended with the following functions:

- Selecting the license metric of a product to which you are assigning the component
- Excluding a product or a component from license calculations, for example when it is a trial version
- Viewing information about components that were installed during a specific period and their assignment

- **9.2.6 Enhanced discovery of Microsoft Exchange editions**

You can retrieve information about editions of Microsoft Exchange Servers that are installed in your infrastructure and verify how many licenses are needed for each edition. For more information, see: [Discovering Microsoft Exchange edition \(on page dix\)](#).

- **9.2.6 Reporting high water mark of Oracle Database user sessions**

BigFix Inventory retrieves the number of user sessions that are registered into the Oracle Databases deployed in your environment, and provides their high water mark. You can use this information to calculate the demand for the Oracle Named User Plus license. For more information, see: [Metering the number of Oracle Database concurrent sessions \(on page dviii\)](#).

- **9.2.6 Measuring application usage on Mac OS X**

Information about application usage for software that is installed on Mac OS X is available on the Software Installations report. You can use this information to verify whether extension of license is needed for particular software instances based on their frequency of use.

To enable collecting application usage on Mac OS X, stop actions created by the Initiate Software Scan task that are currently running. Then, activate the Application Usage Statistics analysis, and start the package data and application usage statistics scans by using the updated task. For more information, see: [Activating the analyses \(on page cci\)](#) and [Initiating software scans \(on page cciii\)](#).

- **9.2.6 Utilization of the Install Seats metric for products discovered by SWID tags**

The All Metrics report is extended to include information about the utilization of the Install Seats metric for non-IBM products that are discovered by SWID tags. You can now see the number of installations of software that is discovered by SWID tags and for which the software catalog does not define license details. For more information about SWID tags, see: [Support for ISO/IEC 19770-2 \(on page lxxxiii\)](#). For more information about the Install Seats metric and the All Metrics report, see: [License metric utilization \(on page dlxxxvii\)](#).

- **9.2.6 Upload of PVU table and software catalog during the upgrade**

During the upgrade, new versions of the software catalog and PVU table are uploaded to BigFix Inventory to ensure that the latest software is properly discovered and that new processors are correctly matched. The versions that are uploaded are the newest ones that were available during the release of the application update to which you are upgrading. For more information, see: [Upgrading to BigFix Inventory 9.2.17 \(on page cccv\)](#).

- **9.2.6 Reduced impact of scans due to the optimization of the scanner cache folder on the endpoints**

The scanner cache folder is used to store information about scanned files and directories in your file system. By knowing the hierarchy of files, the scanner can locate them quicker, which results in shorter scans. The amount of disk space that is needed for the cache depends on the number of files that are being scanned.

Fresh installations of BigFix Inventory use the optimized cache which includes only files that are relevant for the types of software scans that you are running. The optimization results in a reduced size of the cache on the monitored computer. Scanner cache settings are not changed during the upgrade. For more information, see: [Optimizing scanner cache configuration \(on page cclxvi\)](#).

- **9.2.6 Searching the software catalog by part number**

The Software Catalog report provides information about product part numbers and associated license metrics. Thus, you can check which products your part numbers represent and what license metrics they use. This information facilitates the process of software classification. You can also verify whether a particular part number is contained in the current software catalog, and thus influences automated bundling.

- **9.2.6 Best practices for database backup**

Documentation provides best practices for creating database backups. For more information, see: [Backing up and restoring the database \(on page cdxxiv\)](#).

- **9.2.6 Fixlet version visible in the BigFix client logs**

The BigFix client log contains information about versions of fixlets that were run on the specific computer. The information can be used for troubleshooting purposes, for example to determine whether the action should be restarted because it was created by an old version of the fixlet. For information about the location of client logs, see: [Component log files](#).

- **9.2.6 Additional options for migrating from License Metric Tool and Tivoli Asset Discovery for Distributed 7.x**

You can decide whether you want to migrate the following items that used to be migrated by default:

- Software bundlings
- Software exclusions
- VM managers

- **9.2.6 Improved performance of generating CSV reports**

Generation of CSV reports is now faster and consumes less resources.

- **9.2.6 Support for new operating systems and virtualization technologies**

You can detect software that is installed on Windows Server 2016, Ubuntu 16 for IBM Power Systems, and SUSE Linux 12 for IBM Power Systems on Little Endian Architecture. You can use XenServer and Citrix XenServer as supported virtualization technologies on selected operating systems. For more information, see: [System requirements \(on page c\)](#).

Application update 9.2.5

- **9.2.5 Utilization of the Install metric for non-IBM products**

Information about non-IBM products that use the Install metric is available on the All Metrics report. You can see not only the current number of computers on which the product is installed but also how that number changed over time. You can also set a threshold, to easily monitor when the number of installed instances exceeds your expectations. For more information, see: [License metric utilization \(on page dlxxxvii\)](#).

- **9.2.5 Preview Feature: Extension of the Software Classification panel**

The Software Classification panel introduces a new way of looking at software inventory and relations between the discovered components and licensable products. It provides a flat structure with software installations broken into single components, which makes it easier to view software assets and manage relations between them. The preview panel is extended with the following functions:

- Selecting and assigning multiple components at once
- Confirming the assignment of components
- Classifying products from vendors other than BigFix

- **9.2.5 Bundling tags for improved software classification**

Bundling tags facilitate the process of software classification allowing you to save time and effort related to analyzing to which product a particular component belongs. They contain information about the part number of a product to which the instance of a component should be assigned. Thus, they increase accuracy of the automated bundling. For more information, see: [Bundling tags \(on page dxxx\)](#).

- **9.2.5 Software discovery on IBM i**

You can discover software that is installed on the IBM i systems. The discovery is based on a semi-automatic procedure that uses scripts to initiate software and capacity scans. Because there is no direct connection between the IBM i systems and the BigFix server, the scans that are initiated on these systems are called *disconnected scans*. Although the scans run automatically, you must perform several manual actions, such as installing the scanner, or uploading the results to BigFix Inventory. For more information, see: [Discovering software and hardware on IBM i \(on page cclxxi\)](#).

- **9.2.5 Software discovery in Docker containers**

Docker is a platform that allows for automating the deployment of applications inside software containers. BigFix Inventory discovers software that is installed inside Docker containers. It also measures license metric utilization of the discovered BigFix products. For more information, see: [Discovery of software in Docker containers](#).

- **9.2.5 Determining the usage of Virtual Processor Cores (VPC)**

You can use BigFix Inventory to track the consumption of Virtual Processor Cores (VPC). VPC is a unit of measure by which a program can be licensed. It can be considered as a physical processor core, if the server is not partitioned for virtual machines, or as a virtual core assigned to a virtual machine. For more information, see: [Determining the usage of Virtual Processor Cores \(VPC\)](#).

- **9.2.5 Information about software scan duration**

Information about the duration of the software scan is available in the Software Scan Status analysis. You can use this information, for example, to review the scanner timeout settings or verify whether optimization is needed to reduce the length of the scan.

- **9.2.5 New VM manager status - Insufficient rights**

The *Insufficient rights* status indicates that no data was collected from the VM manager because the user that is defined for this VM manager does not have sufficient rights.

- **9.2.5 Information about processor brand string**

Information about the processor brand string is available on the Hardware Inventory report. It contains a full specification of the processor including its brand, model, and speed as read from the computer operating system. You can use this information to apply vendor Core Factor values.

- **9.2.5 Forcing reupload of software scan results**

When the import of software scan results fails, or when software is not reported properly on the Software Installations report, you can use the Force Reupload of Software Scan Results task to troubleshoot the problem.

- **9.2.5 Installing the scanner in an alternative path**

You can choose whether you want to install the scanner in the default directory (on the `C:` or `opt` drive) or the directory where the BigFix client is installed. This option is useful, for example, when the security policy in your company does not allow for installing software on the `C:` or `opt` drive. For more information, see: [Installing the scanner \(on page ccii\)](#).

- **9.2.5 Support for new operating systems and virtualization technologies**

You can detect software that is installed on Debian 8 for x86 and Ubuntu 16 for x86. You can use z/KVM as a supported virtualization technology on SUSE Linux 12.

Application update 9.2.4

- **9.2.4 Setting metric thresholds for BigFix products**

You can specify a metric threshold that is used to show how the metric consumption that you would like to maintain relates to the actual consumption. You can then schedule emails to be sent whenever metric consumption for any product exceeds the threshold. For more information, see: [Setting license metric threshold \(on page dxc\)](#).

- **9.2.4 Sharing feedback**

If you encounter problems with panels and reports or think that some of them might be improved, you can quickly share your feedback by simply clicking the **Send Feedback** button that is in the top-right corner of the user interface. Each panel has a separate button that creates an email message with the panel name, version of the application that you are using, and your comment. For more information about other available communication channels, see: [Give us feedback](#).

- **9.2.4 Preview features**

Preview features allow you to work with the newly implemented functions and to provide your feedback to influence their final form. These features are subject to BigFix Support. For more information, see: [Preview features \(on page xlix\)](#).

- **9.2.4 Improvements for VM managers**

The following improvements are introduced in the area of VM managers:

- Until now, the VM Managers panel listed only VM managers that are managed in central mode. Now, it also lists those that are managed in the distributed mode. Options such as editing, deleting, or testing connection are still applicable only for the former. However, the read-only view of the latter gives you a clear overview of all VM managers that are defined in your infrastructure. For more information, see: [Understanding the VM Managers panel \(on page cccxlvii\)](#).

- VMware vSphere Fault Tolerance allows for reducing virtual machine and application downtime to zero. If a server failure occurs, a live shadow instance of a virtual machine is created. The instance is up-to-date with the primary virtual machine. BigFix Inventory can handle situations when shadow instances of virtual machines are created and can properly calculate PVU utilization of software that is installed on such machines.
- When VM managers in your infrastructure have status other than OK, and you are not sure how to troubleshoot the problem, you can simply click the name of the status. You will be redirected to a documentation page that provides the most common reasons for each status and how to proceed.
- When the security policy in your company does not allow for transferring VM manager credentials through the network, you can configure BigFix Inventory so that the credentials are not transferred to the server. They are also not displayed on the VM Managers panel when you edit the VM manager. For more information, see: [Advanced server settings \(on page cdii\)](#).

- **9.2.4 API to retrieve cluster information**

Information about the cluster to which a computer belongs is available on the Hardware Inventory report. Now, the information can also be retrieved through REST API. It can be then used, for example, to calculate license demand for Oracle or Microsoft. For more information, see: [Retrieval of clusters \(on page cmlxxiv\)](#).

- **9.2.4 Support for BigFix database restoration**

You can enable an import that retrieves all software scan data from the restored BigFix database. Such an import is time-consuming and should be run only when problems with incomplete data occur after the database is restored. When you enable the option, the import runs once and returns to the normal operation mode. For more information, see: [Database problems \(on page dcclxviii\)](#).

- **9.2.4 Optimization of the scanner cache configuration**

The scanner cache folder stores information about scanned files and directories in your file system. By knowing the hierarchy of files, the scanner can locate them quicker, which results in shorter scans. The amount of disk space that is needed for the cache depends on the number of files that are being scanned. If the current location of the cache folder cannot ensure sufficient disk space, you can optimize the cache. Thanks to the optimization, the cache includes only files relevant for the types of software scans that you are running. For more information, see: [Optimizing scanner cache configuration \(on page cclxvi\)](#).

- **9.2.4 Discovery of software that is not in the current software catalog**

When BigFix software is discovered with a software ID tag but is not in the current catalog, a warning sign is displayed on the Software Catalog widget. It indicates that a newer version of the catalog needs to be uploaded. For more information, see: [Support for ISO/IEC 19770-2 \(on page lxxxiii\)](#).

- **9.2.4 Support for new operating systems and virtualization technologies**

The BigFix client can be installed on Oracle Linux 7 and on Power Linux on the Little Endian architecture. Support is added for PowerKVM virtualization technology. Additionally, support for Multiple Shared Processor Pools is extended to Linux. For more information, see: [System requirements \(on page c\)](#)

Application update 9.2.3

- **9.2.3 Support for Unicode**

Starting from V9.5, the BigFix Platform can gather data from BigFix clients that are deployed in different geographical locations and thus use different code pages and languages. The data is encoded into UTF-8 format, and reported back to the BigFix server. As a result, regardless of the system and encoding that the clients use, the data is correctly displayed both in BigFix and BigFix Inventory. For more information, see: [Unicode \(on page lxxxiv\)](#).

- **9.2.3 Collecting file MD5 and SHA256 checksums for security purposes**

Checksums are long strings that describe the content of files and act as their fingerprints. You can enable the calculation of checksums for files on your computers to check their integrity and to ensure that they were not altered or tampered with. For more information, see: [Collecting file checksums \(on page dclx\)](#).

- **9.2.3 Conditional report emails based on the number of rows in the report**

You can configure reports to be emailed to specified recipients regularly or when a specific number of rows is exceeded in the report. This option is useful for alerting the recipients about changes to the key data or when the recipients do not work with BigFix Inventory, yet must have access to the reports. For more information, see: [Scheduling report emails \(on page dclv\)](#).

- **9.2.3 Controlling scan schedule directly from the BigFix Inventory user interface**

The **Scan Configurations** panel is extended to allow for basic configuration of software scans per computer groups directly from the BigFix Inventory user interface. Use this method as an alternative to scheduling scans from the BigFix console.

- **9.2.3 Data visibility and license reporting per computer group limited to authorized users**

Improvements that are introduced in the area of data visibility and license reporting per computer group include:

- Limitation of information that is displayed on the Audit Trail report to actions performed only within the computer group to which the user who views the report has access.
- Possibility of enabling and disabling subcapacity calculations per computer group. For more information, see: [Setting up computer groups \(on page cxc\)](#).
- Reporting historical data for deleted computers within the scope of a computer group.
- Introduction of permissions to upload part numbers and view contracts. For more information, see: [Roles \(on page clxxxiv\)](#).

The improvements give more flexibility and better management of data visibility in situations where computers are divided into groups and each group is managed by a different person. For example, in the service provider's environment or when computers are divided into regions or business units.

- **9.2.3 Replacing Software Knowledge Base Toolkit with BigFix Inventory for managing the software catalog content**

You can use a tool that moves the custom catalog content from Software Knowledge Base Toolkit to BigFix Inventory. Then, you can stop using Software Knowledge Base Toolkit and start managing the catalog content entirely in BigFix Inventory. Managing signatures in one application instead of two saves your time, facilitates the work, and reduces maintenance costs. For more information, see: [Migrating Software Knowledge Base Toolkit](#).

- **9.2.3 Flexible date filtering**

You can use time filters to display data starting from or ending at a particular date. For example, you can view computers that were not seen within the last week or contracts for which maintenance ends within a month. When you combine this feature with conditional report emails, you get instant notifications about changes within your environment including reports on which data fulfills specific time conditions.

- **9.2.3 Collecting host serial numbers**

You can collect serial numbers of VM manager hosts and define the format in which the information is displayed on the reports. For more information, see: [Advanced server settings \(on page cdii\)](#).

- **9.2.3 Forcing capacity scan data uploads**

When the import of capacity scan data fails, you can run a single capacity scan and force the upload of its results to ensure that the data on reports is refreshed. For more information, see: [Initiating the capacity scan on all computers \(on page ccx\)](#).

- **9.2.3 Changing advanced server settings from the application user interface**

You can change advanced settings of the BigFix Inventory server, for example settings related to the configuration of the central VM Manager Tool, on the **Advanced Server Settings** panel. Additionally, a number of new parameters is added. For more information, see: [Configuring advanced server settings \(on page cdi\)](#).

- **9.2.3 Manually adjusting the PVU per core value from the BigFix Inventory user interface**

In earlier versions of BigFix Inventory, the PVU per core value was adjusted by using a fixlet. Starting from this version, the fixlet is deprecated. The PVU value can be changed on the **Hardware Inventory** panel instead. For more information, see: [Changing the PVU per core value \(on page dcccxxiii\)](#).

- **9.2.3 Support for new operating systems and virtualization technologies**

The BigFix client can be installed on Cent OS 6 and 7. Red Hat Enterprise Linux 7 and SUSE Linux Enterprise Server 12 are supported on z/VM. Additionally, support for RHEV-M 3.5 is provided. For more information, see: [System requirements \(on page c\)](#).

Application update 9.2.2

- **9.2.2 License reports per BigFix subcapacity region**

BigFix subcapacity licensing rules define three regions. To properly calculate subcapacity values for products that are installed in multiple regions, generate the audit snapshot for each region. Then, for every product, sum up the subcapacity values from all regions. The value that you obtain is the overall subcapacity usage for the particular product. For more information, see: [Virtualization Capacity License Counting Rules, BigFix subcapacity regions \(on page lxxxviii\)](#) and [Tutorial: Reporting subcapacity usage per computer group \(on page dclxix\)](#).

- **9.2.2 License reports per organizational business unit**

You can manage software and view reports per organizational business unit. To do this, divide the computers in your infrastructure into groups that represent these units. Then, create dedicated users, each with access only to a particular computer group. For more information, see: [Organizational business units \(on page xciii\)](#) and [Tutorial: Reporting subcapacity usage per computer group \(on page dclxix\)](#).

- **9.2.2 License reports in the service provider environment**

A service provider is a company that provides other organizations with access to software installed on machines that are owned and maintained by the service provider. **9.2.2** Available from 9.2.2. To properly calculate subcapacity values in the service provider environment, divide the computers in your infrastructure into groups that represent your customers. Then, manage software and generate audit snapshots separately for each customer. For more information, see: [Service providers \(on page lxxxvi\)](#) and [Tutorial: Managing software in the service provider environment \(on page dclxxv\)](#).

- **9.2.2 Creating extended signatures**

Extended software signatures are based on data elements such as file names, registry keys, or specific strings. They allow for discovering software that otherwise might not be discovered with file or package signatures. You can now create extended software signatures directly in BigFix Inventory without the necessity of installing Software Knowledge Base Toolkit. For more information, see: [Creating extended signatures](#).

- **9.2.2 Software Knowledge Base Toolkit is being sunset**

Software Knowledge Base Toolkit was used by BigFix Inventory users to create complex signatures, and was required at the time when BigFix Inventory could not provide equal capabilities. Now, with the possibility of creating complex signatures directly BFI_unused_file in BigFix Inventory, Software Knowledge Base Toolkit becomes deprecated and will eventually be sunset.

- **9.2.2 Consolidated endpoint information**

Hardware information, such as processor details, is available on the Computers, Software Installations, and Oracle Databases reports. Combination of software and hardware information on a single report facilitates the calculation of processor-based and core-based licenses for non-IBM products.

- **9.2.2 Disabling the calculation of extended software aggregates**

Software aggregates are pieces of data that are gathered during software scans. They are processed during the successive data imports and later on displayed on the Inventory Exploration report. You can disable the calculation of extended software aggregates to reduce the length of data imports in very large environments where multiple computer groups are defined. For more information, see: [Disabling the calculation of extended software aggregates \(on page dcccxljii\)](#).

- **9.2.2 Configuring location of the scanner cache folder**

The scanner cache folder is used to store information about scanned files and directories in your file system. By knowing the hierarchy of files, the scanner can locate them quicker, which results in shorter scans. Change the location of the cache folder if the default one cannot ensure sufficient disk space. For more information, see: [Optimizing scanner cache configuration \(on page cclxvi\)](#).

- **9.2.2 Identifying VM manager duplicates**

You can use the `vmman_collecting_hostnames_enabled` parameter to identify duplicate UUIDs on VMWare vSphere. When you set the parameter to true, host names of virtual machines with duplicate UUIDs are written in the log file. For more information, see: [VM Manager Tool settings \(on page ccclxxxii\)](#).

- **9.2.2 Defining time range of reports in relation to the current date**

You can define the time range of reports in relation to the current date. Then, the report provides you with data for a specific period, for instance the last 2 weeks.

- **9.2.2 REST API for retrieving raw scan results**

You can retrieve raw scan results through REST API to improve the monitoring of your environment by quickly discovering software that was installed or modified recently. Shortened data imports cause that data about installed software reaches the server more quickly. They also allow for increasing the number of imports that you run each day. For more information, see: [REST API for retrieving raw scan results \(on page dcccclxxvii\)](#).

- **9.2.2 Support for new platforms**

The BigFix client can be installed on SUSE Linux Enterprise Server 12 (x86). Additionally, support for Red Hat Enterprise Linux 7 on Power 8 is provided. For more information, see: [System requirements \(on page c\)](#).

Related information

[Software Use Analysis versions](#)

Keeping BigFix Inventory up-to-date

BigFix Inventory updates and content packs are released periodically. Ensure that you upgrade the BigFix Inventory server with each release to take advantage of the new features and apply the application patches.

Periodical updates

Release of a BigFix Inventory application update

BigFix Inventory update is released periodically, typically at the end of each calendar quarter. The update contains a variety of new and improved features as well as the application patches. BigFix Inventory version increments with each release, and is independent of the version of the BigFix platform.

Release of a content pack

Apart from application updates, a content pack is released usually at the end of each month. A content pack provides new discovery capability for a number of software components, and might contain modifications to fixlets. Content pack does not upgrade the BigFix Inventory server.

How to turn on automatic notifications about application updates and content packs

To keep current with BigFix Inventory and receive notifications about application updates, subscribe to [Announcements for BES Administrators](#). After you subscribe, you will receive e-mail notifications related to the announcements.

Information about the latest releases and content packs

- For a list of features added in a particular application updates, see: [What's new in this release \(on page viii\)](#).
- For a list and description of every release and content pack as well as the list of patches, see: [Read Me: Updates and Patches](#).

Currency of BigFix Inventory components

Keeping BigFix Inventory up-to-date

It is important to upgrade BigFix Inventory server with each release. You do not need to upgrade the remaining infrastructure components, such as the BigFix server or clients that frequently. BigFix Inventory is compatible with any supported version of BigFix. Upgrading the BigFix Inventory server resolves the majority of known defects and allows you to take advantage of the newly developed and improved features.

Versioning of the BigFix Inventory components

Versions of BigFix Inventory and the BigFixplatform are independent and they differ.

9.2.15 Starting from application update 9.2.15, the version of the application update and versions of the internal tools of BigFix Inventory, such as, the scanner, SAP Metric Data Collector and VM Manager Tool are unified. In this update the version of each component is 9.2.15. In the future, the version of each component will change to the version of the application update only when this component will change.

Updating BigFix Inventory components

To make sure that your BigFix Inventory architecture is up-to-date, analyze the list of actions according to their priority.

1. [Upgrade to the newest version of BigFix Inventory. \(on page cccv\)](#)
2. Stop your active actions, and then restart them using the newest versions of fixlets.
3. Upgrade the internal BigFix Inventory tools, such as, the scanner, SAP Metric Data Collector and VM Manager Tool. For more information, see: [Installing the scanner \(on page ccii\)](#) and [Updating VM Manager Tool \(on page cccxciv\)](#).
4. Optional: Upgrade the BigFix platform components. This step is optional and mainly depends on your business needs and requirements. Upgrading the platform might be related to the security measures as each update might bring a number of security enhancements. From BigFix Inventory perspective, upgrading BigFix is usually not required for BigFix Inventory to work properly. The majority of the features are independent and work with older versions of the platform.

Ensuring currency of scanner actions and fixlets

Make sure that the scanner actions that are running on your endpoints are updated after each application release. The fixlet names contain the information about version to help you easily determine which actions are activated with the outdated fixlets and should be restarted after the application update.

Versioning of fixlets

9.2.15 Starting from application update 9.2.15, names of all BigFix Inventory fixlets include their version in brackets, for example; **Initiate Software Scan (9.2.15.0)**.

The fixlet version is updated with each release of the product, regardless of whether the definition of the fixlet changed. However, the versions of the internal tools, such as the scanner, SAP Metric Data Collector and VM Manager Tool are updated only when these tools change. Therefore, it might occur that the version of the installation or upgrade fixlet and the version of the tool that it deploys differ.

Best practices

Limit modifications of the fixlet names

Do not modify the entire name of the fixlet, but rather add a meaningful prefix or suffix to it. Keep the version in the fixlet name, so that it is reflected in the name of the initiated action. Thanks to this, you can easily distinguish the actions that are initiated with different version of the fixlet.

Ensure currency of your scanner actions

Your actions are not automatically updated upon installing a new application release. Whenever new version of fixlets are available, restart your current actions. To check what actions to restart, compare the names of the active actions with the names of the relevant fixlets. If the versions that are included in these names are different, restart the actions.

To restart the single scanning action, perform the following steps.

1. Go to the BigFix console, and select **Actions** in the navigation tree.
2. Right-click the chosen action, and select **Stop Action**.

3. Go to **Sites > External Sites > IBM BigFix Inventory v9**.
4. Choose **Fixlets and Tasks**, select the fixlet that you want to restart and click **Take Action**. Make sure that the fixlet parameters are the same as those set in the previously stopped action.



Note: BigFix Inventory server is compatible with the actions that were initiated with the previous versions of the server. However, only by ensuring the currency of your scanner action and fixlets you take advantage of new features and improvements that are implemented in the new release.

Plan ahead

Plan ahead and be ready to restart your current actions every time a new release of BigFix Inventory is installed. The plan should be aligned with your business needs and should not prevent you from upgrading the server.

Features and functions

BigFix Inventory provides useful features for managing virtualized environments. It discovers the software that is installed in your infrastructure, helps you to analyze the consumption data, and allows you to generate reports.

Software discovery and identification

BigFix Inventory scans your infrastructure to determine what software is installed on the monitored computers.

Signature discovery and creation

BigFix Inventory includes a catalog of software products along with a set of rules called signatures. The signatures are matched against the results of the scan to build a complete overview of software that is installed in your infrastructure.

Software usage counting

BigFix Inventory provides you with information about how software that is installed in your enterprise is used.

Reports

Reports contain detailed information about the computers in your infrastructure and the software items that are installed on these computers. By viewing the reports regularly, you can check whether actual installations reflect the software inventory information and ensure that the capacity values are assigned to products correctly.

Contracts

Contracts are used to manage information about licenses for the software products that are installed in your infrastructure and to track spending that is associated with those licenses. A contract contains information about the cost of acquiring and maintaining a license, as well as its entitlement start and end dates.

IBM license compliance

BigFix Inventory provides the information that is needed to ensure compliance with processor value units (PVU), and resource value units (RVU MAPC) audit reports licensing terms.

BigFix Inventory calculates the maximum core capacity of the server that is available to the installed IBM software. The application also determines the number of processor value units (PVU) or processor-core entitlements that are required. If you deploy a supported virtualization technology, the tool provides processor core capacity information for the whole virtualized environment.

IBM license usage monitoring

You can generate PVU and RVU MAPC audit reports that show the product and processor value units. The audit reports can be viewed in the application or printed.

Preview features

BigFix Inventory provides a set of preview features that give you a chance to take a sneak peek at what we are currently developing and researching. Whether it is a small function or an extended report, you can check how it works and tell us about your ideas. The previews can change with each release, which means that you can still influence their final form.

Disclaimer

The preview features are subject to change and can be modified and improved based on the customer needs and feedback. Some previews might gradually subside or be entirely removed from the product.

Providing your feedback

Your feedback is crucial for the development team and your comments and suggestions are very much appreciated. Test the new features and let us know what you think and what can be improved. Use the following communication channels to provide your feedback at: talk2sam@hcl.com.

Limitations

Review the list of BigFix Inventory limitations.

Installation and configuration

- **Support for Distributed Server Architecture (DSA) is limited to one backup server**

Distributed Server Architecture (DSA) allows for setting up multiple BigFix servers that can replicate data between each other. If the main server is disabled, a backup server takes over its role ensuring continuous work. BigFix Inventory supports a setup in which one of the BigFix servers that is defined as a datasource is designated as a backup server. If this server is running, BigFix Inventory works. However, if this server

fails, imports of data also stop working until the server is up and running. BigFix Inventory server does not automatically switch to use the secondary BigFix server.

- **Microsoft™ SQL Server used as the BigFix database cannot be case sensitive**

Only case insensitive Microsoft™ SQL Server database is supported.

- **Changing the host name of the BigFix server is not supported**

The host name of the BigFix server cannot be changed because it is recorded in the license certificate during the installation. To change the host name of BigFix, reinstall the application.

- **Creating clones of virtual machines that are running is not supported**

Cloning virtual machines that are running might produce unexpected results in BigFix Inventory.

- **Language of the installer depends on the system locale**

To run the BigFix Inventory installer in a language other than English, system locale must be set to that language.

- **Permissions that are defined for the Administrator role cannot be changed**

The Administrator role is assigned all available permissions. It prevents from creating separate Administrator roles, for example one responsible for data, the other for users.

Using BigFix Inventory

- **The scanner does not scan paths that contain certain whitespace characters and other special characters**

The scanner does not scan file paths that contain the following whitespace characters:

- Carriage return (CR, character code 13)
- Line feed (LF, character code 10)
- Whitespace characters whose decimal codes are in the following ranges: 0- 8, 11-12, 14-31

It also does not scan shared disks whose mount point paths contain an asterisk (*) or a question mark (?).

To avoid a situation in which software is not discovered, ensure that files paths in your infrastructure do not contain the unsupported characters.

- **Software installed in the `/Libraries` directory on Mac OS X is not discovered**

Only software that is installed in the `/Applications` directory is discovered.

- **Information about Mac OS X computers is not presented on the IBM Capacity Data Completeness widget**

Computers running on Mac OS X are not included in the computer count on the IBM Capacity Data Completeness widget. It happens because the types of scans that produce results displayed on this widget are not run on Mac OS X. However, these computers are included in the computer count on the Scan Health widget. It causes that the overall number of computers displayed on each widget is different.

- **Calculating utilization of the Managed PVU metric requires an additional configuration of BigFix Inventory**

Managed PVU metric is based on the number of processor cores available on computers that are managed by an application, for example by BigFix Spectrum Protect for Virtual Environments. Calculating utilization of the Managed PVU metric is not available out of the box. To calculate its utilization, place an appropriate `isotag` file on the computers that are managed by the particular application.

- **Creating contracts and catalog customizations for IBM products is limited**

IBM provides catalog content that allows for discovering practically all IBM products. Due to internal reasons, entries for some products are duplicated in this catalog. Thus, it is not possible to modify the correct entry and creating contracts for IBM products might not work correctly. Moreover, BigFix Inventory supports contracts only for installation-based licenses which are not applicable to the majority of IBM products. Creating contracts for IBM products and customizing BigFix products in the software catalog is discouraged.

- **Signatures with the same GUID and different rules might be overwritten**

When you import a signature that has the same GUID as a signature that already exists in BigFix Inventory and the new signature has a different number of rules than the existing one, the signatures are not merged. Instead, the new signature overwrites the existing one. The problem does not occur when the imported signature has a different or no GUID.

- **Software is detected twice if a custom catalog entry is the same as an entry in the BigFix catalog or an entry generated from an `isotag` file**

If you create a custom catalog entry which has the same publisher name, software name, and version as an entry in the BigFix catalog or an entry generated from an `isotag` file, both the custom and original entry is visible as detected.

- **`slmtag` files with overlapping entries are not supported**

Start time of one entry in an `slmtag` file cannot overlap with the previously reported period. It should start at the end time of the previous entry or later.

- **Adding excluded directories fails on Solaris when the number of characters in the task exceeds 1000**

When you use the Add Excluded Directories task and the specified list of directories is longer than 1000 characters, the task fails on Solaris.

- **Specific start time cannot be set for actions by using REST API**

BigFix REST API requires that the start time for an action is specified as an offset from the time on the BigFix server. The solution is prone to network latency problems. It might cause that the action starts at time different from the intended one.

Reports

- **Reports show the current PVU per core value instead of the value that occurred during the license peak time**

IBM PVU Subcapacity and All Metrics reports as well as the audit snapshot show the current PVU per core value instead of the value that occurred during the license peak time. It does not affect the calculation of

PVU subcapacity and full capacity. To display the value that occurred during the license peak on the reports, change the start date of the report to a more recent date. Try changing the date until you find the date when the PVU per core value changed. It allows you to have consistent information displayed on the reports.

- **Data on the Software Installations report is inconsistent with data on other reports**

Information on the Software Installations report regarding software installed on shared disks or components installed multiple times on the same computer is inconsistent with information on other reports. The inconsistency occurs also when a column on a report links to a filtered view of the Software Installations report such as in case of the Installed Software column on the Computers report. It happens because the Software Installations report is based on the old data model that does not take shared disks and multiple component instances into account.

9.2.15

Starting from application update 9.2.15, the old Software Installations report is deprecated and substituted with a new report view. It is based on a new data model and consistent with other reports. The limitation no longer applies.

- **Scheduling multiple report e-mails requires ensuring an appropriate interval between each report e-mail**

When you schedule multiple reports to be sent to your email, you must ensure that they are not being sent at the same time. The reason is that each report is managed in a different session that starts when the report is being sent, and ends when the action is completed. Each closed session, however, terminates also other sessions, in this case the delivery of other reports. Therefore, the schedule must be different for each report. The sufficient interval is about 5-10 minutes. You can specify it while setting the start time.

- **Audit Trail limitations**

Audit Trail report cannot be filtered by the Details column and does not contain information about uploads and imports of part numbers. What is more, audit trail does not contain information about software classification when either component, product or both are non-IBM.

- **Non-IBM software discovered both on a local system and a shared disk is shown only as installed locally**

When non-IBM software is installed on a local system and on a shared disk, it is shown only as installed locally on the Software Classification panel.

- **Filtering computers based on their IP addresses does not work under certain conditions**

When you filter your computers based on their IP addresses and specify the relation as `ends with`, no computers are displayed even if some computers match the criteria. To work around this issue, you can choose other relations, such as `contains` or `begins with`. This limitation concerns all reports.

- **Filtering by installation path does not work under certain conditions**

When the installation path is specified as report filtering criteria, and contains `-1`, for example: `C:\My_App-1`, the filter returns inaccurate results. The filter returns not only the content of the specified directory, but also all non-IBM products that are installed on the computer.

- **Filtering causes that the content of the window is scrolled up**

When you filter a report by using the `in_set` filter, the content of the entire window is scrolled up to the top. You have to scroll down to add another filter.

- **Incorrect number of rows can be displayed on reports that are saved as PDF files**

The number of rows that is displayed in the upper right corner of a PDF file can differ from the number of rows that is displayed on the user interface. This problem was only encountered with Mounted Shared Disks column in the Computers report because it contains multi-line entries.

- **Time values are provided in different time zones on the user interface and in PDF or CSV reports**

On some reports, time values are expressed in the local time zone when the report is viewed on the user interface but in the UTC time when the report is exported to a PDF or CSV file. What is more, the format of the date and time is hardcoded in PDF reports.

- **Signature definitions are not available on the Catalog Audit report for custom signatures that were deleted**

When you delete a custom signature from the software catalog, definition of this signature is not available on the Catalog Audit report. The definition section is empty.

- **Non-English characters are incorrectly displayed on the Package Data report for IBM i systems**

If the name of a package that is discovered on an IBM i system contains non-English characters, the characters are incorrectly displayed on the Package Data report.

- **Sorting by name does not work on the predefined reports.**

Security

- **SP800-131a compliance cannot be enabled with Microsoft™ SQL Server**

If the SP800-131a cryptographic standard is enabled in strict mode in BigFix Inventory and the enhanced security mode is configured on the BigFix and MS SQL servers, connection between the BigFix Inventory server and the SQL Server cannot be established.

- **The pvk format of the encrypted private key is not supported**

You can use an encrypted private key in the openSSL or pkcs8 formats.

- **Key pair generated for BigFix Inventory can be used for Web Reports only if the private key is not password-protected**

Other

- **The To Do List displays the names of outdated scanning actions always in English.**

- **Catalog search does not work during the import of data**

- **Fixlet for distributing scanner catalogs to the endpoints is not translated**

The Catalog Download (Version: *version*) fixlet is a custom fixlet that is not delivered with the BigFix Inventory fixlet site. Thus, it is not translated.

- **Process IDs are not present in audit snapshots**
- **Eastern Arabic numerals are not supported**

Key concepts

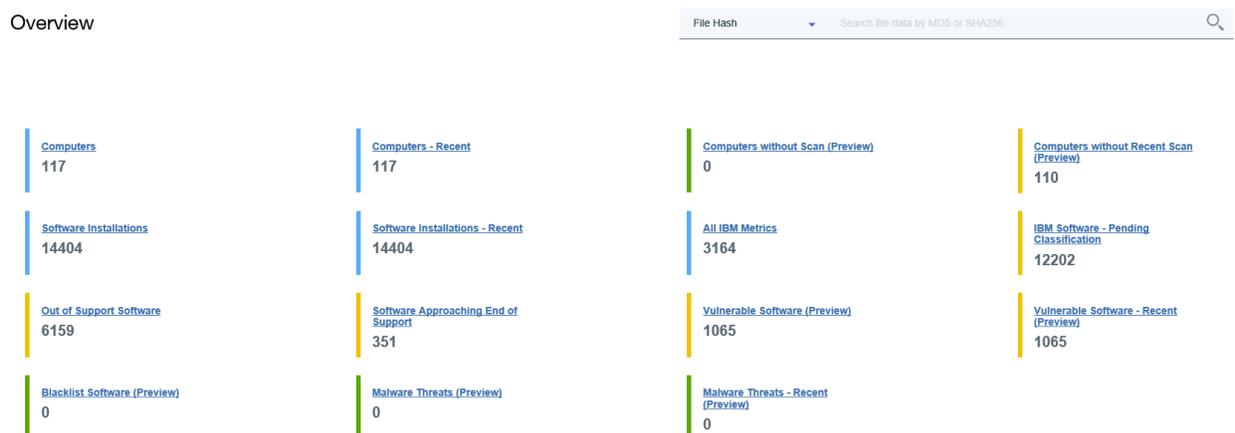
There are several concepts that will help you to understand how BigFix Inventory works and how to use it effectively.

Dashboard

The dashboard provides you with a quick overview of the most important information about scans, computers, and software assets in your infrastructure.

9.2.13 New Dashboard and Reporting

Starting from application update 9.2.13, the BigFix Inventory development team is working on improving the dashboard experience. To view the new dashboard, click **Go to New Dashboard & Reporting** on the home page. To provide feedback about the dashboard, write an email to talk2sam@hcl.com.



The view on the dashboard differs depending on whether you have permissions to view reports on which the dashboard widgets and search are based. Each widget summarizes the number of rows on the respective report or report view. For descriptions and more information about the predefined reports that are a base for widgets and the quick search, see: [Available reports \(on page dcxvi\)](#).

9.2.14 The new dashboard contains a quick search in the upper right corner. It allows you to search your inventory and view the information limited to the defined entity on the appropriate report. You can search your inventory by the following values.

- Search the file data by the file hash.
- Search the file data by the file name.

- Search the vulnerable software by the CVE name.
- Search the software inventory by the component name.

Dashboard

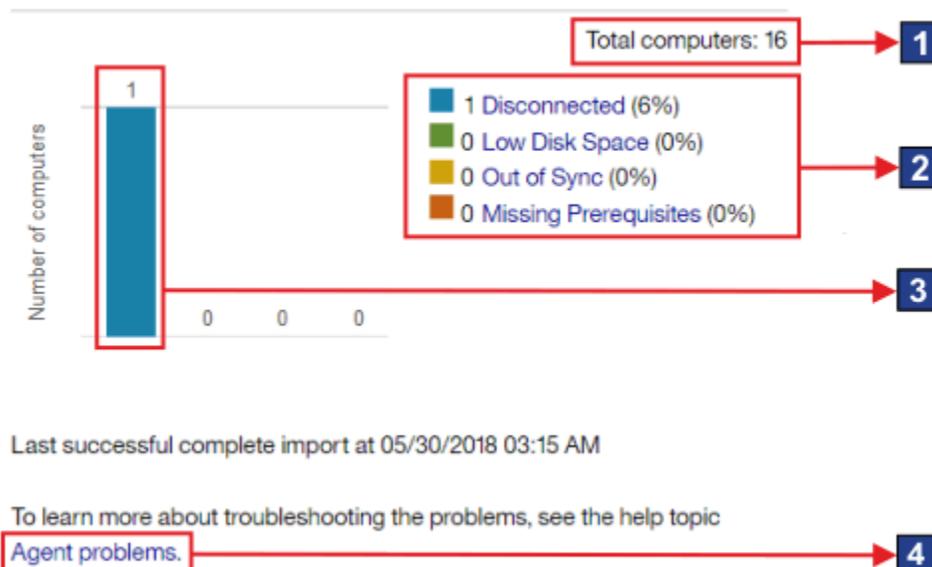
If you use BigFix Inventory 9.2.12 or earlier, your dashboard contains the following widgets.

Deployment Health

The widget shows whether BigFix clients that are installed in your infrastructure are connecting to the BigFix server. It also reports the most common issues that might occur while clients are operating such as problems with disk space or missing scanner prerequisites.

i **Tip:** Read the Deployment Health, Scan Health, and IBM Capacity Data Completeness widgets from left to right and fix issues in this order.

Deployment Health



Elements of the widget

1 The total number of computers to which the user has access. The number is determined by the computer group to which the user is assigned.

Mac OS X The total number of computers includes Mac computers. However, information about deployment health is not collected from these systems. Thus, they are not included in the counts for particular statuses.

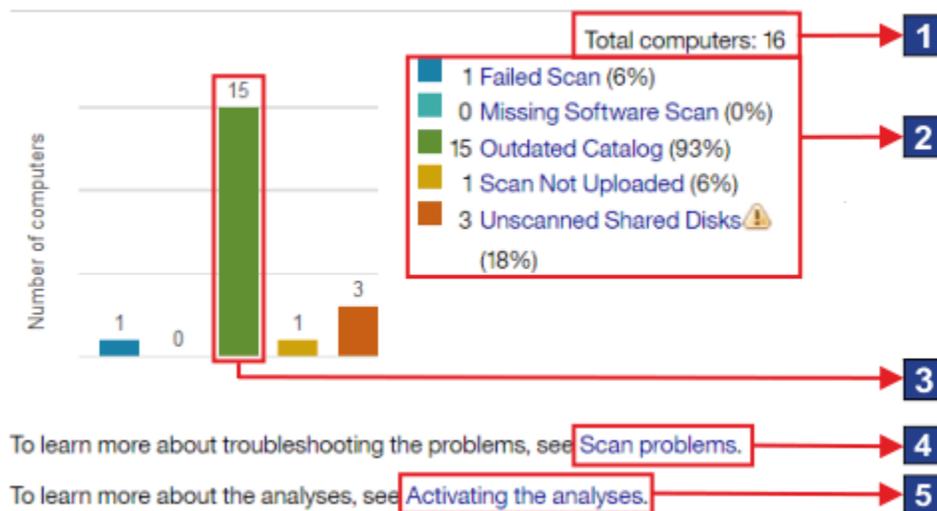
- 2** Links to the report with results narrowed down to computers with the particular status.
- 3** The number of computers with a particular status.
- 4** Link to information about [actions that you can take to resolve the reported problems \(on page dclvii\)](#).

Scan Health

The widget shows the health of scans that are running in your infrastructure. When software scans are not working correctly, the installed software might not be discovered.

- !** **Important:** The Software Scan Status analysis must be activated for the widget to show valid data. If the analysis is not activated, all computers are reported with the Failed Scan, Missing Software Scan, and Outdated Catalog statuses.

Scan Health



Elements of the widget

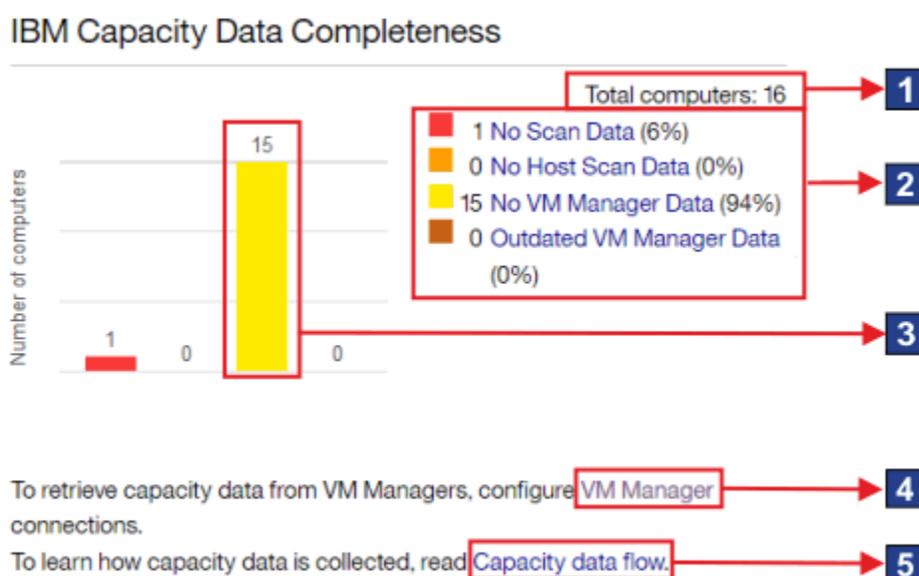
- 1** The total number of computers to which the user has access. The number is determined by the computer group to which the user is assigned.

Mac OS X The total number of computers includes Mac computers. However, only information about the status of the package data scan is collected from these systems. Therefore, Mac computers are included in the count for the Failed Scan status. They are not included in the counts for the remaining statuses.

- 2** Links to the report with results narrowed down to computers with the particular status.
- 3** The number of computers with a particular status.
- 4** Link to information about [scan problems](#) (on page [dcllxxxi](#)).
- 5** Link to information about [activating analyses](#) (on page [cci](#)) required to collect data that is displayed on the widget.

IBM Capacity Data Completeness

The widget shows whether capacity data is correctly gathered from the computers in your infrastructure. The lack of capacity data might impact calculation of PVU consumption.



Elements of the widget

- 1** The total number of computers to which the user has access. The number is determined by the computer group to which the user is assigned.
 - Mac OS X** The total number of computers includes Mac computers. However, information about capacity data is not collected from these systems. Thus, they are not included in the counts for particular statuses.
- 2** Links to the report with results narrowed down to computers with the particular status. For more information about each status, see: [Computer statuses](#) (on page [cccxvi](#)).
- 3** The number of computers with a particular status.

4 A link to the **VM Managers** panel.

5 A link to information about the [flow of capacity data \(on page cccxlv\)](#) between virtual machines, VM Manager Tool, and the BigFix Inventory server.

Inventory Exploration

The widget shows top five publishers with the largest number of defined contracts. The publishers are ordered according to the number of computers on which their software is installed, regardless of the number of contracts.

Inventory Exploration



Elements of the widget

1 The number of products from a particular publisher.

2 A link to the **Inventory Exploration** report.

3 A link to the **Contracts** panel.

Inventory Data

The widget shows a summary of the installed BigFix software as well as computers, and computer groups in your infrastructure.

Inventory Data

11387 Software Installations including titles such as IBM BigFix Platform Agent, Red Hat Enterprise Linux Server, Logical Volume Manager, IBM DB2 Enterprise Server Edition Unlicensed Product Base, IBM BigFix Inventory Server, and IBM Tivoli System Automation for Multiplatforms

0 Computer Groups

16 Computers with operating systems including Linux Red Hat Enterprise Server 7.4 (3.10.0-693.el7.x86_64), Linux Red Hat Enterprise Server 7.0 (3.10.0-123.el7.x86_64), Win2012R2 6.3.9600, Linux SuSE Enterprise Server 12 (3.12.28-4-default), Linux CentOS 7.4.1708 (4.9.87-linuxkit-aufs), and Linux Red Hat Enterprise Server 6.8 (2.6.32-642.el6.x86_64)

Elements of the widget

- 1** A link to the **Software Installations** report.
- 2** A link to the **Computer Groups** report.
- 3** A link to the **Computers** report.

Software Catalog

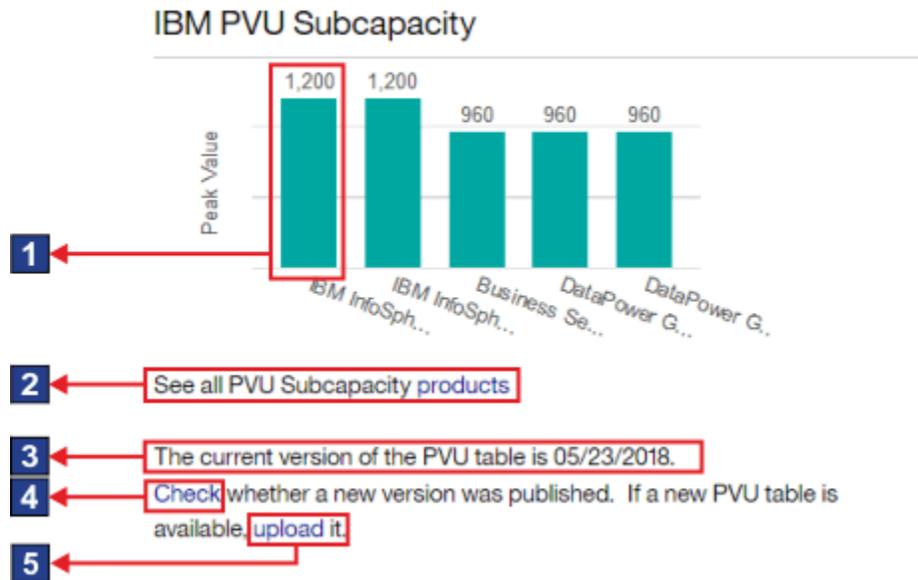
The widget shows information about the content and version of the current software catalog.

IBM PVU Subcapacity

The widget shows products with the highest PVU consumption rate. It shows how many PVUs a product consumes but does not relate this information to your license entitlements. By default, the maximum of five products is displayed.

The accuracy of the displayed data depends on when the scan data was imported, whether the PVU table is up-to-date, and whether software assignment was modified. If any of these factors was changed, an appropriate message is displayed on the widget.

If the widget shows `No data`, the data is not available. It might occur when scan data was not uploaded, the upload of the data has not finished yet, or inventory scans do not work properly. The message is no longer displayed if data from at least one BigFix client is successfully updated.



Elements of the widget

- 1 The PVU consumption rate for a product.
- 2 A link to the entire **IBM PVU Subcapacity** report.
- 3 The current version of the PVU table.
- 4 A link to the website on which you can check whether a new version of the PVU table is available.
- 5 A link to the **Metric Table Upload** panel on which you can upload a new version of the PVU table.

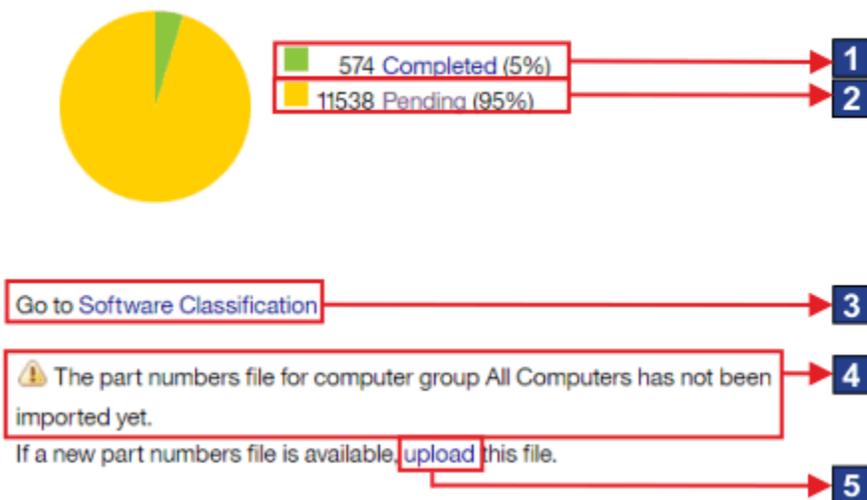
BigFix

The widget shows the number of completed and pending classifications of the software that is installed in your infrastructure.

The accuracy of the displayed data depends on when the scan data was imported and whether the part numbers file is up-to-date. If any of these factors was changed, an appropriate message is displayed on the widget.

If the widget shows `No data`, the data is not available. It might occur when scan data was not uploaded, the upload of the data has not finished yet, or inventory scans do not work properly. The message is no longer displayed if scan data from at least one BigFix client is successfully updated.

IBM Software Classification



Elements of the widget

- 1** A link to the Software Classification panel with results narrowed down to software installations with complete classification.
- 2** A link to the Software Classification panel with results narrowed down to software installations that are still pending classification completion. Go through these items and either confirm their default assignment or reassign them to different products so that all installations in your infrastructure are confirmed.
- 3** A link to the Software Classification panel.
- 4** The date when the last part numbers file was imported.
- 5** A link to the **Part Numbers Upload** panel on which you can upload the part numbers file.

9.2.13 To Do list

9.2.13 Available from 9.2.13. The To Do list provides you with information about items that might require your action including failed imports of data, availability of a new version of BigFix Inventory or the need to configure a connection to VM managers.

To Do items

The number of To Do items is indicated by the number that is displayed on the To Do icon . It is shown until the appropriate action is undertaken, which can be done if you are an Administrator or have the required permissions.

Otherwise, you are advised to contact the BigFix Inventory Administrator to undertake the action for you. To update the To Do List, refresh the BigFix Inventory user interface.

Table 1. Information that is displayed on the To Do list

Type of information	Details
Failed import of data	Provides the date and time of the import that failed as well as the date and time of the last successful import.
Availability of a new version of BigFix Inventory	Provides information about the new version of the BigFix Inventory server that is available as well as the version that you are currently running.
The need to configure connections to virtualization hosts	Provides information about the need of configuring connections to VM managers to avoid overcounting subcapacity licenses. For more information about the impact of not configuring connections to VM managers, see: Managing VM managers (on page cccxviii) .
9.2.15 Outdated scanning actions	<p>Provides information about the scanning actions that are running and which were started using the older versions of the Initiate Software Scan and Upload Software Scan Results fixlets. To address this issue, restart these action in the IBM BigFix Inventory v9 fixlet site. For more information, see: Ensuring currency of scanner actions and fixlets (on page xlvi).</p> <p> Note: The To Do List displays the names of outdated scanning actions always in English.</p>

Changing the scope of the To Do list

If you are an Administrator, you can change which information is shown on the To Do list. Currently, you can disable displaying information about the need of configuring connections to VM managers. To disable the information, go to **Management > Advanced Server Settings**, and change the value of the `todo_display_vm_managers_need` parameter to false.

Extract, Transform, Load (ETL)

The *Extract, Transform, Load* (ETL) is a process in the database usage that combines three database functions that transfer data from one database to another. The first stage, *Extract*, involves reading and extracting data from various source systems. The second stage, *Transform*, converts the data from its original format into the format that meets the requirements of the target database. The last stage, *Load*, saves the new data into the target database, thus finishing the process of transferring the data.

In BigFix Inventory, the *Extract* stage involves extracting data from the BigFix server. The data includes information about the infrastructure, installed agents, and detected software. ETL also checks whether a new software catalog is available, gathers information about the software scan and files that are present on the endpoints, and collects data from VM managers.

The extracted data is then transformed to a single format that can be loaded to the BigFix Inventory database. This stage also involves matching scan data with the software catalog, calculating processor value units (PVUs), processing the capacity scan, and converting information that is contained in the XML files. After the data is extracted and transformed, it is loaded into the database and can be used by BigFix Inventory. The hardest load on the BigFix Inventory server occurs during ETL when the following actions are performed:

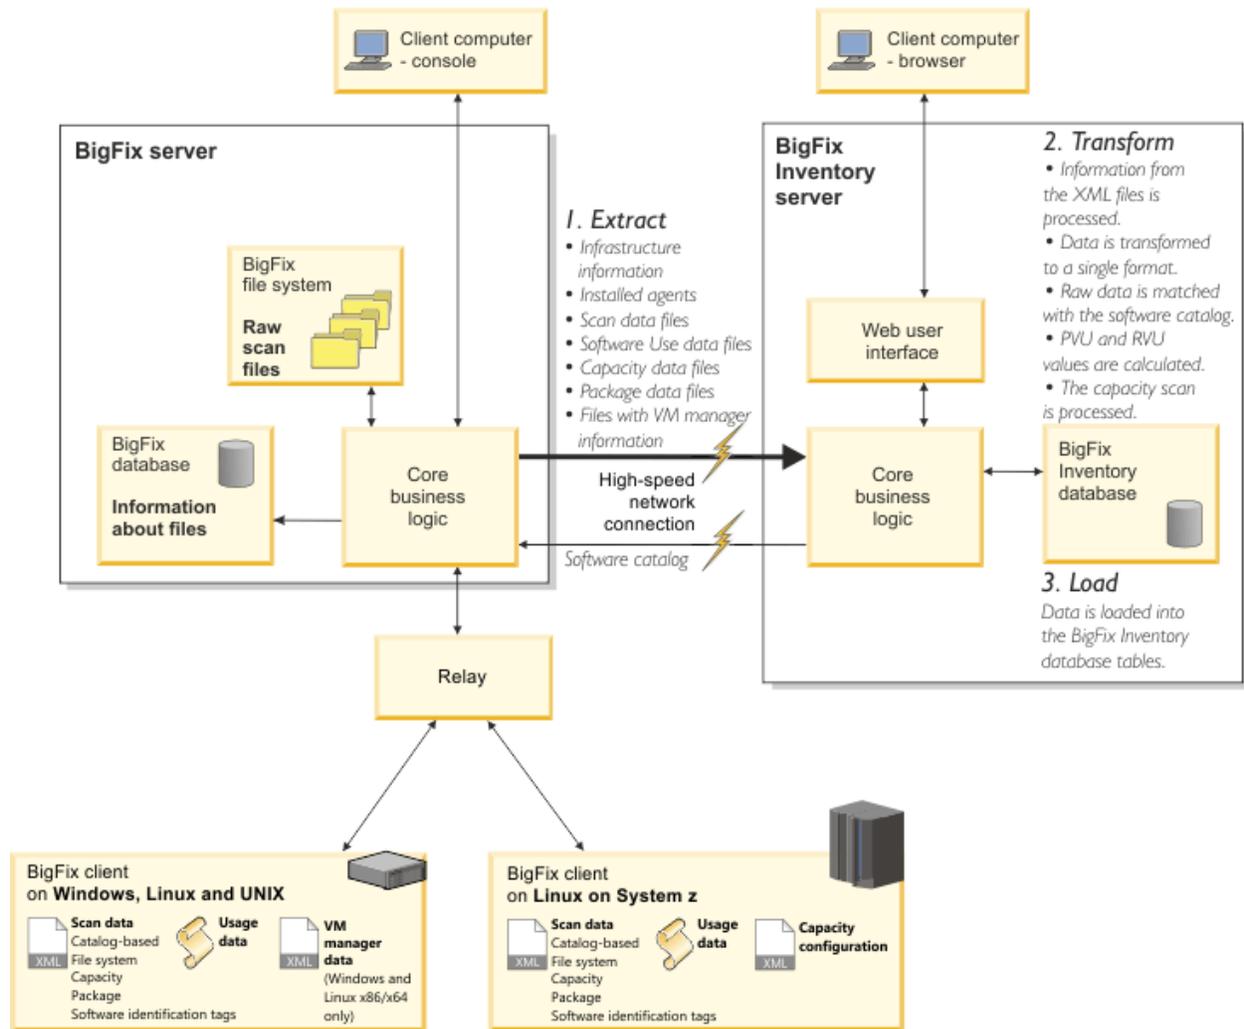
- A large number of small files is retrieved from the BigFix server (Extract).
- Many small and medium files that contain information about installed software packages and process usage data are parsed (Transform).
- The database is populated with the parsed data (Load).

At the same time, BigFix Inventory prunes large volumes of old data that exceeds its data retention period.

Performance of the ETL process depends on the number of scan files, usage analyses, and package analyses that are processed during a single import. The main bottleneck is storage performance because many small files must be read, processed, and written to the BigFix Inventory database in a short time. By properly scheduling scans and distributing them over the computers in your infrastructure, you can reduce the length of the ETL process and improve its performance.

An important factor that influences the duration of the ETL process is the amount of updates on the file system since the last scan. Such operations as security updates or significant system upgrades can cause ETL to run longer, because it has to process information about all modified files. For example, regular updates released by Microsoft on Tuesdays would significantly lengthen the Wednesday import in environments with many Windows platforms.

Extract, Transform, and Load



Software and hardware discovery

BigFix Inventory discovers hardware and software in your IT infrastructure.

With BigFix Inventory you can accomplish the following goals by using software and hardware discovery:

- Manage software inventory
- Monitor license consumption
- Achieve basic license compliance
- Create an inventory of hardware in the enterprise

Infrastructure analysis

The first step, before you start discovering software and hardware, is to obtain the baseline of computers that exists in your environment. Estimate the total number of computers and organize them into computer groups in the BigFix console to avoid overloading the infrastructure.

Discovery of software and hardware

Collect information about your software, license compliance and hardware inventory. It is commonly achieved by activating the capacity and software scans in the BigFix console. For more information, see: [Setting up scans to discover software and hardware inventory \(on page cxciij\)](#).

Discovery	Description
Software inventory	<p>Software inventory is analyzed by the scanner that searches for different types of information to determine whether the software is installed or to measure its usage. You should set up the appropriate software scans during the initial configuration of BigFix Inventory. For more information, see: Types of software scans (on page ccvi).</p> <p>Thanks to BigFix Inventory you can understand and manage your software inventory, especially through the following practices:</p> <ul style="list-style-type: none"> • Discovering IBM and non-IBM software to view and classify your software inventory on the Software Classification panel. • Measuring utilization of license metrics that are used within your organization. You can view the details on the All Metrics report that shows utilization of the assigned license metric by all instances of a particular product. However, it is limited to a subset of license metrics. For more information, see: License metric utilization (on page dlxxxvii). • Estimating resource utilization of products that deliver .slmtags. The Resource Utilization report shows utilization of all license metrics that can be used by a particular product. However, the data is not aggregated and requires further processing. For more information, see: Raw utilization of license metrics (on page dcxi). • Measuring current application usage by monitoring the information about packages and files that exist on the computers in your environment. You can view the usage statistics on the Inventory Exploration and Usage per Computer report. • Discovering software that is installed on Mac OS. To discover software that is installed on Mac OS, run the package data scan on Mac computers. Other types of software scans are not needed. On the reports, the publisher of these products is generic and called Unclassified Mac Software because the information about the publisher is not available directly. In this case, the publisher should be considered as unknown and no manual action is required to set it.

Discovery	Description
Hardware information and capacity data	<p>Thanks to BigFix Inventory you can collect information about the processors and the processor capacity that is assigned to each computer in your infrastructure.</p> <p>Capacity scan</p> <p>Capacity scan collects data about the physical processor, system type (physical or virtual), guest operating system, and logical partitions. It must be activated on every computer, virtualization host and virtual machine. For more information, see: Initiating the capacity scan on all computers (on page ccx).</p> <p>Capacity scan on virtualization hosts</p> <p>To collect capacity data on Xen and KVM hosts, you need to activate an additional fixlet: Run Capacity Scan on Virtualization Hosts. For more information, see: Collecting capacity data from virtualization hosts for Xen and KVM (on page ccxvi).</p> <p>Capacity scan on virtual machines</p> <p>To collect capacity data on VMware, Hyper-V, KVM RHEV-M, and XenServer or Citrix XenServer, you must add relevant VM managers in BigFix Inventory. For more information, see: Adding VM managers for VMware, Hyper-V, KVM with RHEV-M, Xen Server and Nutanix (on page ccxiv).</p> <p>The hardware information and capacity data are displayed on the Hardware Inventory report.</p> <p>9.2.12 Detailed Hardware Scan</p> <p>To collect information about memory, operating systems, storage, processors, partitions, network adapters, SMBIOS data, IP addresses and logical processor capacity data, you need to activate the Collect Detailed Hardware Information fixlet. For more information, see: Detailed hardware scan (on page ccxx).</p> <p>Starting from application update 9.2.12, thanks to BigFix Inventory you can collect detailed information about hardware in your infrastructure. You can retrieve this data with REST API. For more information, see: REST API for retrieving detailed hardware information (v2) (on page mxliv).</p>

Related information

[Installing the scanner \(on page ccii\)](#)

[Initiating software scans \(on page cciii\)](#)

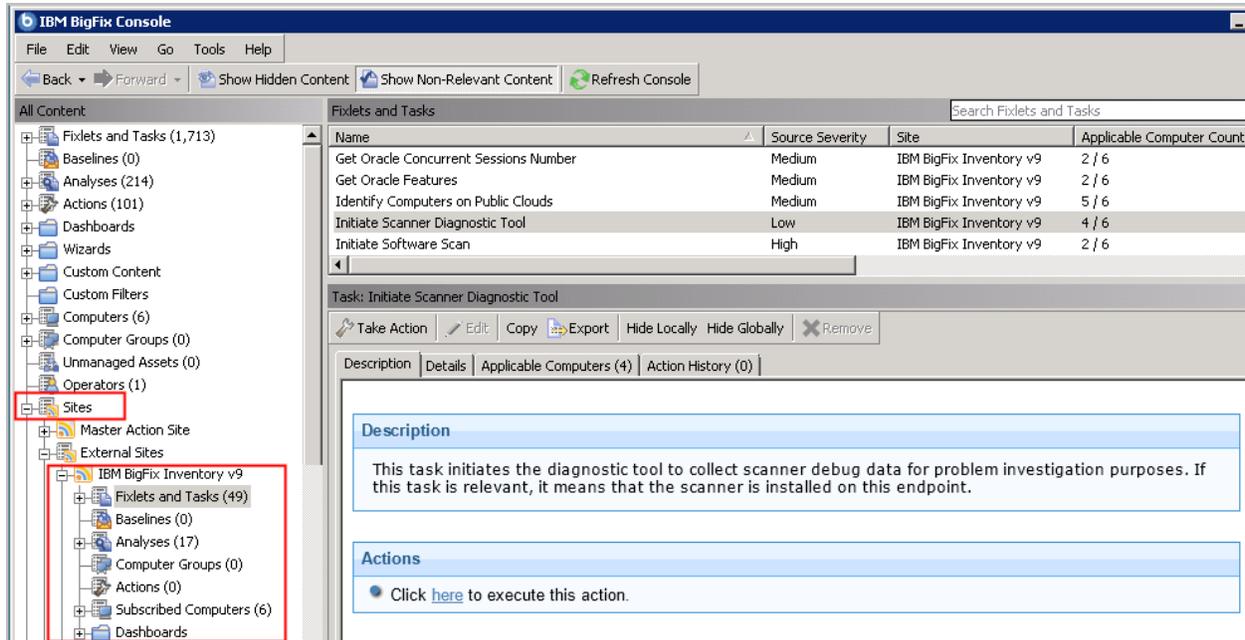
[Uploading software scan results \(on page ccviii\)](#)

Fixlets, tasks, and analyses

Fixlets and tasks are central to BigFix platform. BigFix Inventory, as a part of this platform, uses them to perform required actions on selected computers. Fixlets and tasks define these actions and specify the criteria that make them relevant. Usually, an action is deployed with a simple click.

Overview

Fixlets, tasks, and analyses are available on the dedicated **IBM BigFix Inventory v9** site that can be accessed through the BigFix console.



The screenshot shows the IBM BigFix Console interface. The left sidebar displays a tree view of content, with 'Sites' and 'IBM BigFix Inventory v9' highlighted. The main pane shows a table of Fixlets and Tasks. The selected task, 'Initiate Scanner Diagnostic Tool', is displayed in detail below the table.

Name	Source Severity	Site	Applicable Computer Count
Get Oracle Concurrent Sessions Number	Medium	IBM BigFix Inventory v9	2 / 6
Get Oracle Features	Medium	IBM BigFix Inventory v9	2 / 6
Identify Computers on Public Clouds	Medium	IBM BigFix Inventory v9	5 / 6
Initiate Scanner Diagnostic Tool	Low	IBM BigFix Inventory v9	4 / 6
Initiate Software Scan	High	IBM BigFix Inventory v9	2 / 6

Task: Initiate Scanner Diagnostic Tool

Take Action | Edit | Copy | Export | Hide Locally | Hide Globally | Remove

Description | Details | Applicable Computers (4) | Action History (0)

Description

This task initiates the diagnostic tool to collect scanner debug data for problem investigation purposes. If this task is relevant, it means that the scanner is installed on this endpoint.

Actions

Click [here](#) to execute this action.

For more information, see: [Sites](#) in the BigFix documentation.

Relevance

You can run a fixlet or task only on the computers for which it is relevant. Relevance indicates which computers meet the relevance conditions that are listed for each fixlet or task on the Details tab. The Applicable Computer Count column shows the number of computers that apply. Relevance is determined per computer, and is continuously re-evaluated by BigFix. It helps you understand the sequence of actions. In some cases, it might be required to run one fixlet or task, to enable action defined by another. Relevance of the deployment fixlets indicates that a new component version is available, and you should run it to ensure that the latest code level is used. For more information, see: [Fixlets and tasks](#) in the BigFix documentation.



Note: If you need to run a task on a computer that is not relevant, see: [Checking why a fixlet or task is not relevant \(on page dccccix\)](#).

Analyses

Analyses provide an overview of the most important information about the computers in the environment, the installed scanner version, or status of the last software scan. You can use these details to monitor health and understand if all necessary actions were successfully completed. For more information, see: [Analyses](#) in the BigFix documentation.

Initial configuration

The required configuration fixlets and tasks ensure that BigFix Inventory works properly.

If you selected **Enable default scan schedule for this data source** during initial configuration, the scans are configured automatically. Otherwise, you must configure the scans manually. For more information, see: [Manual scan configuration \(on page cc\)](#).

List of fixlets and tasks

The list and description of fixlets and tasks available in the BigFix console on the **IBM BigFix Inventory v10** site.

Required configuration fixlets and tasks

There are four basic fixlets and tasks that are required to configure BigFix Inventory to collect all necessary data from the computers in your environment.

Table 2. Required configuration fixlets and tasks

Name	Category	Description
Install or Upgrade Scanner	Scanner	BigFix Inventory requires a dedicated scanner to be installed on each endpoint. This task installs the scanner or upgrades it to the latest version on the specified endpoints. For more information, see: Installing the scanner (on page ccii) . Analysis: Scanner Information
Run Capacity Scan and Upload Results	Scanner	After you install the scanner, you need to schedule capacity scans on all endpoints to ensure that the data in BigFix Inventory is complete. By default, the capacity scan is set to run every 30 minutes and usually takes a few seconds. The frequency is required by the IBM pricing for subcapacity reporting. This task runs the capacity scan to retrieve data from the specified endpoints. If the action is completed, the results were gathered and uploaded. For more information, see: Initiating the capacity scan on all computers (on page ccx) .
Initiate Software Scan	Scanner	Use this fixlet to configure software scan that analyzes the file system and registry, and discovers the installed software. The scan can take from few minutes to few hours depending on the number of files

Table 2. Required configuration fixlets and tasks (continued)

Name	Category	Description
		<p>in your file system. You can specify the exact date and time of the scan, not to affect other business applications running on the system. Additionally, you can set a CPU threshold for the scan to minimize its impact, however, that will make the scan last longer.</p> <p>Select all the listed scan types. If the task is not relevant for some endpoints, log in to the BigFix Inventory and run the initial import to ensure that the catalog is propagated to all endpoints. This task is not relevant on endpoints where the scanner is not installed. For more information, see: Initiating software scans (on page cciii).</p> <p>Analysis: Software Scan Status</p>
Upload Software Scan Results	Scanner	<p>This task uploads inventory data that is collected by the software scans from the specified endpoints to the BigFix server. Each endpoint should be targeted only once. For more information, see: Uploading software scan results (on page ccviii).</p>

Troubleshooting fixlets

There are several fixlets and tasks that can help you troubleshoot the most common issues and questions related to the application.

Table 3. Troubleshooting fixlets

Name	Category	Description
Analyze the Relevance of a Fixlet or Task	Troubleshooting	<p>This task creates an analysis that checks which relevance conditions of a particular fixlet or task are not met. Check the results of the analysis to learn why a fixlet or task is not relevant on a particular endpoint. The analysis is created in the Master Action site. For more information, see: Checking why a fixlet or task is not relevant (on page dcccix).</p>
Force Reupload of Software Scan Results	Troubleshooting	<p>This task forces reupload of inventory data that was gathered by the software inventory and the file system scans. The data is reuploaded to BigFix once and no schedule is set. The data is then transferred to BigFix Inventory during the import. The import includes results of the full file system scan, not the delta scan. It might affect the length of the import.</p>

Table 3. Troubleshooting fixlets (continued)

Name	Category	Description
Linux Gather Environment Information	Troubleshooting	This task retrieves information about BigFix Inventory and BigFix environment for troubleshooting purposes. The task is relevant only for the Linux operating system.

Other fixlets

The remaining fixlets and task can help you manage, configure, troubleshoot, deploy and migrate the product and its components, and manage VM Managers.

Table 4. Other fixlets

Name	Category	Description
Add Excluded Directories	Scan Management	This task is used to add entries to the list of directories that are excluded from scanning. After you specify the directories, run the task against the chosen endpoints. The directories are then sent to the scanner that omits them during the software scan. For more information, see: Excluding additional directories (on page ccxliv) .
Add Targeting Exception to Temporarily Disable Scans	Configuration	This task prevents software and capacity scans that are scheduled on the endpoint from running without the need of modifying the entire scan schedule. You can use it to temporarily disable scans on a computer that undergoes maintenance, patching, or some other resource-consuming action that could be disturbed by the software or capacity scan.
Change Scanner Cache Folder	Scanner	This task changes the folder where the scanner stores its data. If this task is relevant, the scanner is installed on all applicable endpoints, but no scans are running. For more information, see: Optimizing scanner cache configuration (on page cclxvi) .
9.2.12 Collect Detailed Hardware Information	Scanner	This task runs a detailed hardware scan on the specified endpoints. It collects information related to memory, operating systems, storage, processors, partitions, network adapters, SMBIOS data, IP addresses and logical partition capacity data. The solution is supported on Windows, Linux x86 and AIX. The task can be run also on other operating systems but the returned results might not be accurate. For more information, see: Detailed hardware scan (on page ccxx) . Analysis: Detailed Hardware Information
Configure File Checksums Collection (MD5/SHA-256)	Scan Management	This task enables the collection of file checksums and is intended for advanced security scenarios only. The MD5 and SHA-256 check-

Table 4. Other fixlets (continued)

Name	Category	Description
		sums are calculated during the file system scan, and are created for each file that is discovered by the software scan. For more information and requirements, see: Collecting file checksums (on page dclx) .
9.2.8 Collect Logs from Endpoints	Troubleshooting	The task collects BigFix client logs, scanner logs, scan configuration logs, VM manager configuration files, and files generated during scans. Run this task when you are requested by BigFix Support to collect log files for troubleshooting purposes.
Configure Scan Cache	Scanner	This task allows for configuring the file system scan cache. By default, the cache is enabled and is cleared after every scan. The cache is shared by different scan types. Disabling the cache will significantly impact scan performance, and it is not recommended, unless absolutely necessary. The task is relevant only on computers where the scanner is installed. For more information, see: Optimizing scanner cache configuration (on page cclxvi) .
Configure Scanner Query Timeout	Scanner	This task allows you to set the maximum time in which the software scan process should finish. As the software scan involves scanning as well as additional evaluation, the overall software scan time might exceed the value provided.
Linux Create Capacity Configuration for Linux on z Systems	Configuration	This task creates a file that contains the manually-entered capacity values. If this task is relevant, it means that no capacity configuration exists on this endpoint yet and the manual configuration is required. If this fixlet is not relevant on a computer, you do not have to perform any manual actions to calculate the capacity values. For more information, see: Creating capacity configuration for Linux on z Systems (on page cclxv) .
Linux Delete Capacity Configuration for Linux on z Systems	Configuration	This task deletes the file that contains the manually-entered capacity values.
Disable Application Usage Statistics	Configuration	This task disables the collection of application usage statistics on endpoints. For more information, see: Disabling the collection of software usage (on page cclxviii) .
9.2.8 Linux Solaris Disable Collecting Executable Files	Configuration	Collecting executable files based on application usage is automatically enabled when you have the BigFix client and server in version 9.5.5 or higher, use Linux or Solaris system, and have the Application Usage Statistics analysis activated. This task disables collection of

Table 4. Other fixlets (continued)

Name	Category	Description
Based on Application Usage		executable files that are based on the list of processes on endpoints. If you run this task, the files associated with these processes are not displayed on the Scanned File Data report.
Discover Remote Shared Disks	Scan Management	<p>This task identifies shared disks that exist in your environment. For more information, see: Discovering software on shared disks (on page ccxxi) and Step 2b: Optimized mode - Manual scanning of remote shared disks (on page ccxxvii).</p> <p>Analysis: Shared Disk Information</p>
Download BigFix Inventory	Deployment	This task downloads the newest BigFix Inventory server installer to the specified endpoint.
Download the Disconnect- ed Scanner Package (version)	Deployment	This task provides a link to the latest disconnected scanner package. The package contains installation and configuration files, and the scripts that are required for software and capacity discovery on Windows and Linux computers. This solution can only be used if the BigFix client cannot be installed due to legal, business, or technical reasons. For more information, see: Discovering software and hardware with disconnected scanner on Windows and UNIX (on page ccxlviii) .
IBM i Download the Disconnected Scanner Package for IBM i (version)	Deployment	This task provides a link to the latest disconnected scanner package for IBM i. The package contains installation and configuration files, and the scripts that are required for software and capacity discovery on IBM i systems. For more information, see: Discovering software and hardware on IBM i (on page cclxxi) .
Linux Edit Capacity Configuration for Linux on z Systems	Configuration	This task edits the file that contains the manually-entered capacity values. If this task is relevant, it means that the capacity configuration already exists on this endpoint.
Edit Scanner Trace Settings	Troubleshooting	This task allows you to set debug traces for the scanner. If this task is relevant, it means that the scanner is installed on an endpoint. For more information, see: Changing scanner trace settings (on page dcccxviii) .
9.2.8 Linux Solaris Enable Collect- ing Executable Files Based on Application Usage	Configuring	This task enables collection of executable files that are based on the list of processes on endpoints. If you run this task, the files associated with these processes will be displayed on the Scanned File Data report. Use this task if you previously disabled the feature using

Table 4. Other fixlets (continued)

Name	Category	Description
		the Disable Collecting Executable Files Based on Application Usage fixlet.
Force VM Manager Tool Scan Results Upload	Troubleshooting	This task forces the upload of capacity data collected by VM Manager Tool or by the Run Capacity Scan on Virtualization Hosts task. The data is uploaded to BigFix and can be transferred to BigFix Inventory by running a data import. The data is uploaded only once and no schedule is set. For more information, see: Forcing the upload of collected data (on page cccxciii) .
Get Microsoft Exchange Edition	Scanner	This task allows you to retrieve information about the editions of Microsoft Exchange servers that are installed on the computers in your infrastructure. For more information, see: Discovering Microsoft Exchange edition (on page dix) .
Get Oracle Concurrent Sessions Number	Scanner	This task allows you to retrieve information about the high water mark value of concurrent sessions for Oracle Databases that are installed on the computers in your infrastructure. For more information, see: Metering the number of Oracle Database concurrent sessions (on page dviii) .
Get Oracle Features	Scanner	This task allows you to retrieve information about the editions and features of Oracle Databases that are installed on the computers in your infrastructure. For more information, see: Extended discovery of Oracle Database (on page dii) .
Get SAP Components	SAP Systems	This task allows you to retrieve information about the SAP components that are installed on the computers in your infrastructure. For more information, see: Discovering SAP components (on page dxiv) .
Identify Computers on Public Clouds	Configuration	This task marks targeted endpoints as computers that are running on a public cloud. It allows the specific license model for the PVU count without hypervisor access. For more information and a list of supported public clouds, see: Identifying computers on public clouds (on page ccxix) .
9.2.9 Initiate Collection of SAP Metric Data	SAP Systems	This task runs the SAP Metric Data Collector. The tool collects license metric utilization data from a single SAP landscape. For more information, see: Collecting utilization of SAP license metrics (on page dxv) .

Table 4. Other fixlets (continued)

Name	Category	Description
9.2.8 Initiate Software Scan on Shared Disks	Scanner	<p>This task runs the catalog-based scan, file system scan, and software identification tags scan to discover software that is installed on shared disks. Scan results are used to create software templates that represent the discovered software. The templates are later on shared between all computers on which the disks are mounted to ensure that the software is reported on these computers. For more information, see: Discovering software on shared disks (on page ccxxi) and Step 2b: Optimized mode - Manual scanning of remote shared disks (on page ccxxvii).</p> <p>Analysis: Status of Shared Disks Software Scan</p>
Install Additional VM Manager Tool <i>version number</i> (OPTIONAL)	VM Managers	<p>This task installs an additional VM Manager Tool in the specified version on any computer. The tool should only be used to manage VM managers that are in separated networks and cannot be reached by the BigFix server. For more information, see: Managing VM managers (on page cccxviii).</p>
9.2.9 Windows Linux Install SAP Metric Data Collector	SAP Systems	<p>This task installs the SAP Metric Data Collector in the specified version on a selected computer. The tool collects license metric utilization data from a single SAP landscape. For more information, see: Collecting utilization of SAP license metrics (on page dxv).</p>
Install VM Manager Tool <i>version number</i>	VM Managers	<p>This task installs VM Manager Tool in the specified version on the BigFix server. Use this task if automatic installation failed. For more information, see: (Optional) Installing VM Manager Tool (on page ccclxv).</p>
Refresh Re-Subscribed Endpoints	Scanner	<p>This fixlet detects endpoints that were re-subscribed from a different fixlet site and is used for the coexistence of License Metric Tool and BigFix Inventory. For such endpoints, it forces the upload of capacity data, and refreshes properties to ensure that they are related to the relevant fixlet site. An endpoint is refreshed after each re-subscription. The initial subscription is set after running the software scan.</p>
Remove Capacity Scan Data from Virtualization Hosts	VM Managers	<p>This task applies to the KVM and Xen hosts only. It removes the data that was generated on virtualization hosts during the capacity scan. For more information, see: Collecting capacity data from virtualization hosts for Xen and KVM (on page ccxvi).</p>
Remove Excluded Directories	Scan Management	<p>This task is used to remove entries from the list of directories that are excluded from scanning. After you specify the directories, run the</p>

Table 4. Other fixlets (continued)

Name	Category	Description
		task against the chosen endpoints. The directories are then sent to the scanner that includes them back in the software scan. For more information, see: Including the excluded directories back in scans (on page ccxlv) .
Remove Targeting Exception to Resume Scans	Configuration	This task resumes software and capacity scans that were temporarily disabled by using the Add Targeting Exception to Temporarily Stop Scans task.
Run Capacity Scan on Virtualization Hosts	VM Managers	This task applies to the KVM and Xen hosts only. It runs a capacity scan on virtualization hosts to retrieve capacity data from all virtual machines that are managed by the specified host. Schedule the scan on every host that manages virtual machines which have the BigFix client installed. For more information, see: Collecting capacity data from virtualization hosts for Xen and KVM (on page ccxvi) .
Schedule VM Manager Tool Scan Results Upload	VM Managers	This task schedules regular uploads of capacity data collected by VM Manager Tool or by the Run Capacity Scan on Virtualization Hosts task. The data is uploaded to BigFix and can be transferred to BigFix Inventory by running a data import. For more information, see: Changing the upload schedule (on page cccxcii) .
Solaris Set DSD Mode	Configuration	This task sets the DSD mode on Solaris endpoints. If the Solaris system is in the DSD domain, set this mode so that the PVU data is correctly calculated. For more information, see: Setting the DSD mode on Solaris (on page cclxvi) .
Software Catalog Update (version)	Deployment	<p>9.2.11 Starting from application update 9.2.11, the catalog update process is automatized. This task automatically downloads the software catalog, and uploads it to the BigFix Inventory server during the next import of data. No further action is required.</p> <p>For versions before 9.2.11, this task downloads the latest software catalog to the computer where the BigFix Inventory server is installed. After the file is downloaded, you must upload it to the BigFix Inventory server. For more information, see: Updating the software catalog to ensure accuracy of software discovery and reporting (on page cdxliv).</p>
9.2.9 Uninstall SAP Metric Data Collector	SAP Systems	This task uninstalls the SAP Metric Data Collector from the selected computer. It also removes folders that the tool uses as well as deletes settings of the BigFix client specific to the tool.

Table 4. Other fixlets (continued)

Name	Category	Description
Uninstall Scanner	Scanner	This task uninstalls the scanner and the corresponding configuration files from the specified endpoints.
Uninstall VM Manager Tool	VM Managers	This task uninstalls the VM Manager Tool from the specified endpoints. It also deletes the client settings and folders used by the VM Manager Tool. For more information, see: Uninstalling VM Manager Tool (on page cccxcv) .
Solaris Unset DSD Mode	Configuration	This task unsets the DSD mode on Solaris endpoints. For more information, see: Setting the DSD mode on Solaris (on page cclxvi) .
9.2.9 Update SAP Metric Data Collector	SAP Systems	This task updates the SAP Metric Data Collector to the latest version.
Update VM Manager Tool to version <i>version number</i>	VM Managers	This task updates VM Manager Tool to the specified version. For more information, see: Updating VM Manager Tool (on page cccxciv) .
Upgrade to the latest version of BigFix Inventory 9.x	Deployment	This fixlet upgrades BigFix Inventory to the latest version. It also uploads the latest software catalog and PVU table to the BigFix Inventory server. For more information, see: Upgrading to BigFix Inventory 9.2.17 (on page cccv) .
WARNING: Endpoints Subscribed to Multiple Sites	Troubleshooting	Some of your endpoints are subscribed to both the License Reporting (LMT) v9 and the IBM BigFix Inventory v9 fixlet site, which will lead to issues with software discovery that are not supported. An endpoint can report to only one application at a given moment. Remove the endpoints from one of the sites.

High-water mark

Utilization of license metrics can fluctuate depending on the location and configuration of virtual machines in the physical environment. High-water mark is the peak in the utilization of a license metric by a product during the reporting period. It represents the number of metric units that are needed to license a product during that period.

The following examples are based on sample utilization of PVUs by DB2 Enterprise Server Edition during a four-week period.

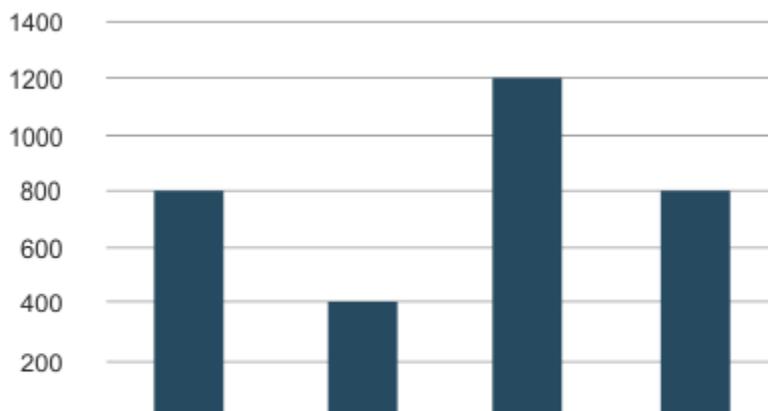
Example 1: High-water mark for a single server with one virtual machine

In this example, DB2 is installed on a single virtual machine. Configuration of this machine changes over time so the number of processor cores that are available to DB2 varies every week. DB2 consumes the following number of PVUs.

- 800 PVUs in the first week
- 400 PVUs in the second week
- 1200 PVUs in the third week
- 800 PVUs in the fourth week

High-water mark for DB2 in this four-week period occurs in the third week and equals 1200 PVUs.

DB2 ESE on server A



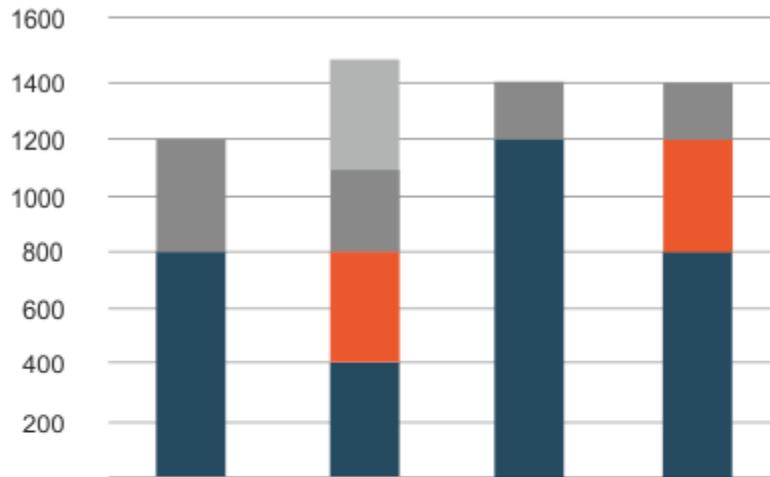
	Week 1	Week 2	Week 3	Week 4
■ Virtual machine1	800	400	1200	800

Example 2: High-water mark for a single server with four virtual machines and no CPU overcommitment

In this example, a server has four virtual machines whose configuration changes over time. The machines do not share any processor cores, and no capping is needed. The server-level utilization of PVU is the sum of PVUs used by DB2 on each virtual machine.

Peak value from a single virtual machine is 1200 PVUs and occurs in the third week. However, DB2 is installed on all virtual machines so values from all virtual machines must be added to calculate the high-water mark. High-water mark for DB2 in this four-week period occurs in the second week and equals 1500 PVUs.

DB2 ESE on server A



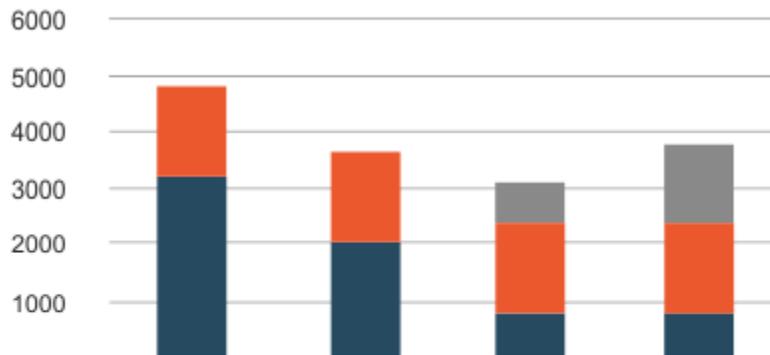
	Week 1	Week 2	Week 3	Week 4
■ Virtual machine 1	800	400	1200	800
■ Virtual machine 2		400		400
■ Virtual machine 3	400	300	200	200
■ Virtual machine 4		400		
Total	1200	1500	1400	1400

Example 3: High-water mark in the entire infrastructure

In this example, there are three servers with many virtual machines. The infrastructure-level PVU utilization is the sum PVUs used by DB2 on each server and utilization on each server is the sum of PVUs used by DB2 on each virtual machine that exists on that server.

High-water mark for DB2 in this four-week period occurs in the first week and equals 4800 PVUs.

DB2 ESE in the entire infrastructure



	Week 1	Week 2	Week 3	Week 4
Server A	3200	2000	800	800
Server B	1600	1600	1600	1600
Server C			700	1500
Total	4800	3600	3100	3900

Related information

[IBM Processor value unit \(PVU\) \(on page dlxxvi\)](#)

[IBM Resource value unit \(RVU MAPC\) \(on page dlxxxiv\)](#)

Main background application tasks

The main goal of BigFix Inventory is to generate PVU, and RVU audit reports based on the collected data. All calculations are done in accordance with PVU, and RVU license pricing rules that are described in the official subcapacity licensing documents.

Aggregation

Aggregation is the main calculation task in BigFix Inventory. The aggregation process is a scheduled background task that is run on a daily basis at a particular hour. By default it is performed when the server time is midnight. The task calculates the PVU, and RVU values based on the data that is collected from the agents during software and capacity scans.

Reaggregation

If the initial software bundles are correct after software discovery and rebundling is unnecessary, aggregation is the only calculation process that is required on the side of the BigFix Inventory server and the data that was once calculated is always correct. However, this is rarely the case. You must always modify some parts of the initial

bundles that are proposed by BigFix Inventory. You must always confirm which bundles are correct for complex products. A complex product is a product that can be bundled with more than one software offering. After rebundling is complete, the PVU and RVU values that were already calculated must be refreshed. Reaggregation tasks were designed to recalculate or refresh PVU and RVU values that were already calculated. Manual actions that might trigger the data reaggregation include:

- Rebundling a software instance from one product to another
- Confirming the default bundle
- Including software instance in PVU, or RVU calculations
- Excluding software instance from PVU, or RVU calculations
- Sharing an instance

These five actions are very basic operations that the application users perform frequently to adjust the bundling data. In addition to manual actions, refresh of the calculated data can also be triggered by automated bundling.

Aggregation versus reaggregation

The aggregation process was designed to calculate the data from many agents for all products over a short time. In contrast, reaggregation process was designed to perform quick recalculation of PVU, and RVU values for a selected subset of products that were already aggregated. Aggregation of all products from all agents is much quicker (even hundreds of times) than reaggregation of the same amount of data. However, when you must recalculate the PVU and RVU values of only one product, reaggregation should be quicker than aggregation, which cannot recalculate the reporting value of only one product, but must reaggregate it for all discovered products simultaneously.

Inventory builder

Inventory builder is another background task that is executed periodically. During this time, the software inventory is built based on the data from the agent software scans. In other words, this task transforms a list of discovered software components to a list of the discovered software products. In most cases, the initial software bundling of detected components performed by the inventory builder has a very low level of confidence.

Automated bundling

Similarly to aggregation, automated bundling is a periodic background task. It is strictly related to the inventory builder task - when the execution of inventory builder ends, automated bundling starts.

When the automated bundling task runs, it determines the best bundle connections to all unconfirmed product instances. If the newly calculated bundles have a higher level of confidence than the current product bundles, automated bundling rebundles those product instances to the new product with the best match. In the vast majority of cases, subsequent automated bundling runs calculate the same level of confidence for most or even for all of the unconfirmed product instances. However, from time to time, especially in large environments with a large percentage of unconfirmed instances where it can actually happen frequently, the newly calculated confidence level can turn out to be higher than the old one. In this case, the software instance is rebundled. The most common reasons why automatic bundling is able to rebundle some old unconfirmed product instances are:

- Import of a new set of part numbers
- Import of a new software catalog
- After manual rebundling or manual confirmation of one product instance, other unconfirmed product instances can be better bundled by using partition or infrastructure collocation rules
- Detection of a new simple software component (a component that can be assigned to only one product) by agents might also change the calculations for other unconfirmed instances due to partition or infrastructure collocation automated bundling rules.

However, automated bundling does not replace the manual work that must be done to confirm or rebundle all unconfirmed product instances. The confidence level that is calculated by automated bundling is supposed to facilitate manual bundling by providing the best potential bundling options for all unconfirmed product instances.

Extract, Transform, Load (ETL)

In general, Extract, Transform, Load (ETL) is the process in database usage that combines three database functions that aim at transferring data from one database and placing it into another. The first stage, Extract, involves reading and extracting data from various source systems. The second one, Transform, converts the data from its original form into the form that meets the requirements of the target database. The last stage, Load, saves the new data into the target database, thus finishing the process of transferring the data.

In BigFix Inventory, the Extract stage involves extracting data from the BigFix server. Such data includes information about the infrastructure, installed agents and detected software. ETL also checks if the new software catalog is available, gathers information about the software scan and files that are present on the endpoints, and collects data from the VM managers.

The extracted data is then transformed to a single format that can be loaded to the BigFix Inventory database. This stage also involves matching scan data with the software catalog, calculating processor value units (PVU), processing the capacity scan, and converting information that is contained in the XML files.

After the data is extracted and transformed, it is loaded to the database and can be used by BigFix Inventory.

Related information

[Bundling Assistant \(on page dxxix\)](#)

Products, components, and bundles

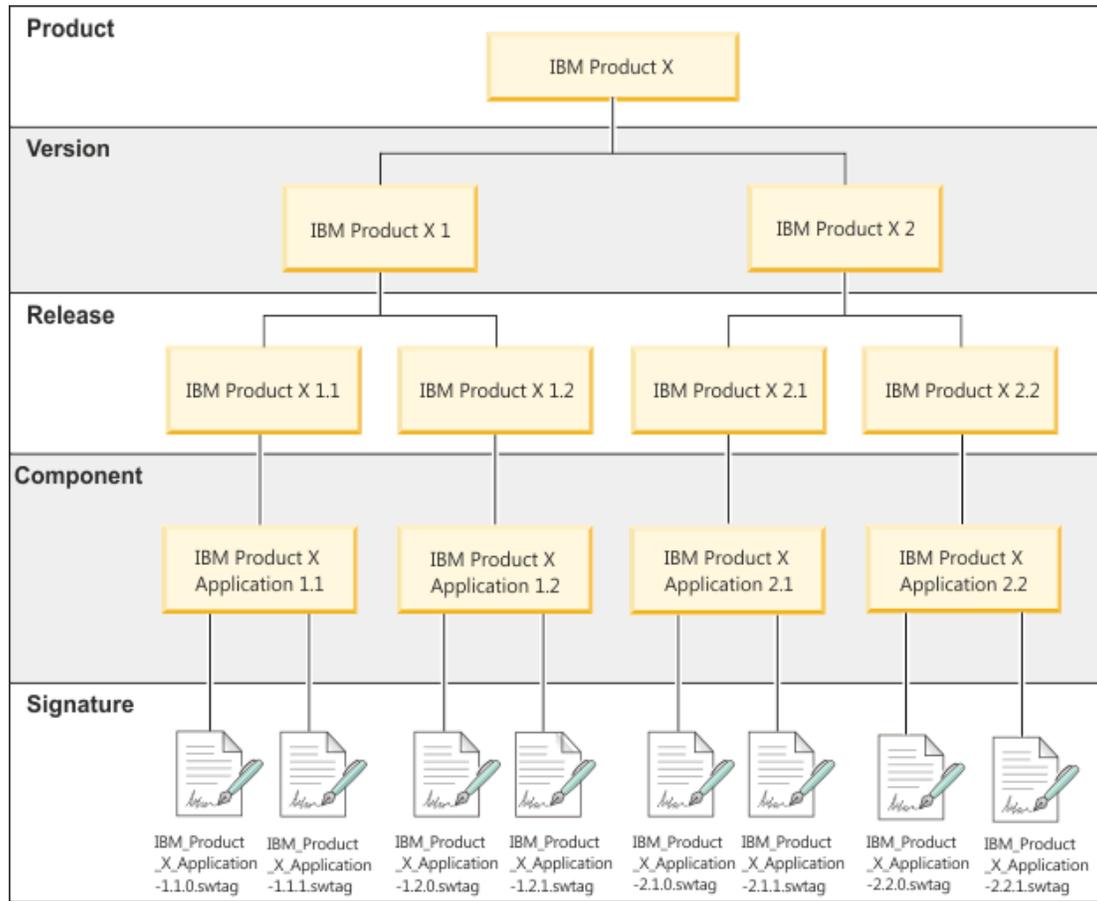
BigFix Inventory discovers software components, and bundles them to products basing on the data contained in the software catalog. Get familiar with these notions to understand the capabilities and performance of BigFix Inventory.

Software catalog

A *software catalog* is a collection of the definitions of software products, their relations and data (signatures) that are used to detect software as installed or used in an IT environment.

An initial software catalog is included in the tool. Be sure to update it regularly, especially if you install new IBM software on your network. IBM delivers and maintains the IBM Software Catalog that includes a set of IBM and non-IBM product definitions and signatures.

The following diagram represents the catalog definition for the product IBM Product X. It shows how the various catalog definitions are related and how they provide a structure that supports identification of the software.



The catalog also contains information indicating what license metric each product uses and how the license quantity is calculated.

The software catalog is a knowledge database that contains information about:

Software product (Offering)

A *software product* is a logical unit of software packaging and sharing that has a managed development and maintenance lifecycle and customer visible attributes. It can be a collection of components, software products whose licensing can be dependent on the licensing of the offering as a whole.

Component

A *component* is a unit of software that can be detected as installed or running on computer systems independently of other software items. It can be a part of a software product, and might be separately identified, but is not individually licensed.

Signature

A *signature* is a unique set of data that is used to discover a given release of the software component. It is used to identify a distributed software release as installed or running on a computer system. There are two types of signatures:

- Inventory recognition signatures detect if a product or component release is installed on the system.
- Usage monitoring signatures detect if a product or component release is running on the system.

A signature can be a file, registry entry, or another identifier, that acts as a fingerprint for identifying software products. The catalog can be augmented with information about custom or proprietary software products in your IT environment.

Hierarchy

A *hierarchy* is a structure that represents software versions, releases, and variations.

Bundle

A *bundle* is set of products that are offered under a single entitlement or license with no dedicated components. In the catalog, a bundle is modeled as a software product with setup relationships to all of the software products that it consists of.

When a file system is scanned, components are discovered by matching the relevant information against signatures that are defined in the catalog. Basing on the match, the most suitable product is selected for each component. You can change the default bundling on the Software Classification panel.

Support for ISO/IEC 19770-2

ISO/IEC 19770-2 is an international standard for creating software identification (SWID) tags. The tags are XML files that are used for discovering and identifying software. They are delivered with software products and contain unique product-related information such as its name, edition, version, whether it is a part of a bundle and more. SWID tags facilitate software discovery and the overall process of software asset management.

Supported schemas

BigFix Inventory supports SWID tags in the 19770-2:2014-CD1 and 19770-2:2015 schemas. For more information about each schema, see the following links:

- <http://standards.iso.org/iso/19770/-2/2015-current/schema.xsd>
- <http://standards.iso.org/iso/19770/-2/2014-CD1/schema.xsd>

BigFix

When BigFix is discovered with an SWID tag that is not contained in the current catalog, a warning sign is displayed on the Software Catalog widget. If you want to see what components were detected, perform the following steps:

1. In the top navigation bar, click **Reports > Software Installations**.
2. Hover over the **Manage Report View** icon , and click **Configure View**. Select the **Component Definition Source** check box to display the column, and add the following filter: `Component Definition Source, equal to, SWID BigFix Tags Without License Info`. Then, click **Submit**.

To be able to bundle such components, upload the latest software catalog.



Tip: **9.2.7** You can create a custom bundling definition on the Software Classification panel.

Limitations

- [Automated bundling \(on page dxxix\)](#) might not be available for BigFix components that are discovered with the SWID tag. To assign such a component to a product, upload the latest software catalog to BigFix Inventory. If the catalog contains the relation, automatic bundling is done. You have to either confirm the suggested relation or, if the relation is not correct, reassign the component to a different product.
- When you create a contract for a product that is discovered by the SWID tag and then upload a catalog that contains information about this product, the product is removed from the contract.

Related information

[Bundling tags \(on page dxxx\)](#)

Unicode

9.2.3 Available from 9.2.3. Starting from V9.5, BigFix Platform has the capability to gather data from BigFix clients deployed with different code pages and languages, encode the data into UTF-8 format, and report it back to the BigFix server. As a result, regardless of the system and encoding that your clients use, the data is correctly displayed both in BigFix and in BigFix Inventory. This capability is useful when your environment has clients with different code pages and data that contains non-ASCII characters.

This capability has the following main assumptions:

- The data is correctly displayed only if the characters belong to the code page of a particular client. Your clients can use various code pages and all of them are supported by BigFix, but from the perspective of a single client, all outgoing data must be consistent with the code page of this client. For example, if you install software whose name contains the `®` character (which is specific to Windows 1252 code page) on a client that uses the Windows 1250 code page, this character will not be correctly displayed in BigFix Inventory, because it does not belong to the code page of the client.

- Encoding into UTF-8 format is supported only for the bottom-up data flow, that is, for the data that flows from the BigFix clients to the BigFix server. The top-down data flow does not support this encoding, which means that all actions sent towards the clients, for example through Fixlets, cannot contain non-ASCII characters. The top-down data flow always uses the code page of the BigFix server.

Prerequisites

- BigFix server 9.5 or higher
- BigFix agent 9.5 or higher
- BigFix Inventory 9.2.3 or higher

Limitations

- Results of the package data scan

All results of the package data scan, that is, packages and data retrieved from the registry, are omitted if they contain characters that do not belong to the code page of the client on which they are installed. Such results are filtered out on the client side and are not uploaded to BigFix. Other types of the software scan are not affected.

- Names of all items created in the BigFix interface

All items created in the BigFix interface, such as custom sites, Fixlets and analyses, computer groups, and so on, will not be correctly displayed if they contain characters that do not belong to the code page of the BigFix server. All these items are a part of the top-down data flow, which does not support encoding into UTF-8 format.

- BigFix Inventory Fixlets that has user-specified paths

Avoid using non-ASCII characters while specifying paths in BigFix Inventory Fixlets, because such paths might not be accepted. The affected Fixlets are: Add Excluded Directories, Remove Excluded Directories, and Change Scanner Cache Folder.

- VM managers credentials

Avoid using non-ASCII characters in user names that are used to log in to VM managers, because the VM Manager Tool that connects to those VM managers does not support encoding into UTF-8 format.

- Subscription e-mails

Avoid using non-ASCII characters in email addresses, which are used as the recipients of scheduled reports.

Related information

[Unicode in BigFix](#)

Scenarios

To understand the business goals that you can achieve with BigFix Inventory, familiarize yourself with the main user scenarios. Each scenario presents a real-life example of using the application and lists personas who are typically involved in achieving a business goal.

9.2.2 Service providers

9.2.2 Available from 9.2.2. To properly calculate subcapacity values in the service provider environment, divide the computers in your infrastructure into groups that represent your customers. Then, manage software and generate audit snapshots separately for each customer.

A service provider is a company that provides other organizations with access to software installed on machines that are owned and maintained by the service provider. There are three typical scenarios that are used in the service provider environment.

Scenario 1: Software that is installed on one virtual machine is used by one customer

In this scenario, all products that are installed on one virtual machine are used by one customer. To properly calculate PVU values for each customer, perform the following steps:

1. Configure VM managers that control virtual machines in your environment.

 **Tip:** If each customer should be allowed to view information about all VM managers in the infrastructure, use [basic VM management \(central\)](#) (on page [cccxlvi](#)). If each customer should be allowed to only view information about VM managers that control the virtual machines that he uses, use [advanced VM management \(distributed\)](#) (on page [ccclxvii](#)).

2. If you grant your customers access to the BigFix console to allow them to run scans and perform other actions on computers, create a site, a computer group, and an operator that is dedicated to each customer. Then, provide the customer with credentials for the dedicated operator. This way, you ensure that each customer has access only to computers on which the software that he uses is installed.
3. Create computer groups that represent your customers in BigFix Inventory. Ensure that each group contains only machines with software that is used by one customer. Avoid creating computer groups that overlap as it might distort the obtained results.
4. Bundle software and generate reports per customer:
 - If you give customers access to BigFix Inventory or when each customer account is managed by a separate person, create dedicated BigFix Inventory users, each with access to one computer group. Then, provide each customer with credentials for the user dedicated to their company. Each customer can manage software and generate reports for the software that they use.

 **Restriction:** To ensure that each user has access to information only about computers that he uses, do not assign the user with the Administrator role. By default, this role has access to

 the All Computers group. Thus, a user who is assigned this role can view information about all computers in the infrastructure.

- If you manage customer software, generate a separate audit snapshot for each computer group. Then, provide each customer with their report.

For detailed instructions, see: [Tutorial: Managing software in the service provider environment \(on page dclxxv\)](#).

Scenario 2: Some of the software that is installed on a virtual machine is used by the service provider

In this scenario, some of the software products that are installed on a virtual machine are used by the service provider, not the customer. To ensure that the customer is not charged for these products, perform the following steps for each computer group:

1. Open the All Metrics report and filter out the product that is not used by the customer.
2. Set the report as the default view for the computer group to which your user is assigned.
3. Generate the audit snapshot. It will not contain the software that is not used by the customer.

 **Restriction:** All instances of the product that you filter out must be used by the service provider.

For detailed instructions, see: [Tutorial: Managing software in the service provider environment \(on page dclxxv\)](#).

Scenario 3: Software that is installed on one virtual machine is used by many customers

In this scenario, multiple customers access software that is installed on one virtual machine. This scenario is not supported in BigFix Inventory.

Data visibility

Some of the items in BigFix Inventory are visible globally to all users. Other items are visible only for the user who has access to a particular computer group. If you grant your customers access to BigFix Inventory, review which items are visible globally and which are visible per computer group.

Table 5. Data visibility

The table consists of three columns and 17 rows.

Item	Visible global-ly to all users	Visible per computer group
Data retention period	Yes	
Email server	Yes	
LDAP server	Yes	
PVU table	Yes	

Table 5. Data visibility

The table consists of three columns and 17 rows.

(continued)

Item	Visible globally to all users	Visible per computer group
Retrieved computer properties	Yes	
Software catalog	Yes	
VM managers (in the basic VM management scenario (on page cccxlvii))	Yes	
Audit trail		Yes
Bundling		Yes
Computer groups		Yes
Contracts		Yes
Part numbers		Yes
Reports		Yes
Software exclusions		Yes
VM managers (in the advanced VM management scenario (on page ccclxvii))		Yes
Unrecognized files		Yes

9.2.2 BigFix subcapacity regions

9.2.2 Available from 9.2.2. BigFix subcapacity licensing rules define three regions. To properly calculate subcapacity values for products that are installed in multiple regions, generate the audit snapshot for each region. Then, for every product, sum up the subcapacity values from all regions. The value that you obtain is the overall subcapacity usage for the particular product.

BigFix subcapacity regions

Regions that are defined by BigFix subcapacity licensing rules include:

- Region 1: North America and South America
- Region 2: Europe and Africa
- Region 3: Asia and Australia

For more information, see: [Virtualization Capacity License Counting Rules](#).

Steps to calculate subcapacity values

To properly calculate subcapacity values in an environment that spans over multiple BigFix subcapacity regions, perform the following steps:

1. Divide the computers in your infrastructure into groups that reflect BigFix subcapacity regions. Ensure that each group contains only computers that are located in one region. Do not create computer groups that overlap as it might distort the obtained results.
2. Optional: If software in each region is managed by a different person, create dedicated BigFix Inventory users, each with access only to computers in one region.
3. Import part numbers for each computer group and bundle software within each region.
4. Generate audit snapshots for each region. Ensure that the snapshots cover the same or approximately the same period.
5. For every product, sum up the subcapacity values from all regions. The value that you obtain is the overall subcapacity usage for the particular product.

! **Important:** If you generate the audit snapshot for the All Computers group, the PVU usage might differ from the value that you obtain by summing up PVU values from snapshots generated for each region. PVU usage for the All Computers group is measured at the time when the sum of PVU values from all regions is the highest. However, this value might not reflect the highest PVU usage in each region. To comply with the subcapacity licensing rules, you must know what is the highest PVU usage in each region, and then sum up these values to obtain the overall PVU usage for a product. To better understand the logic, see the example below.

For detailed instructions, see: [Tutorial: Reporting subcapacity usage per computer group \(on page dclxix\)](#).

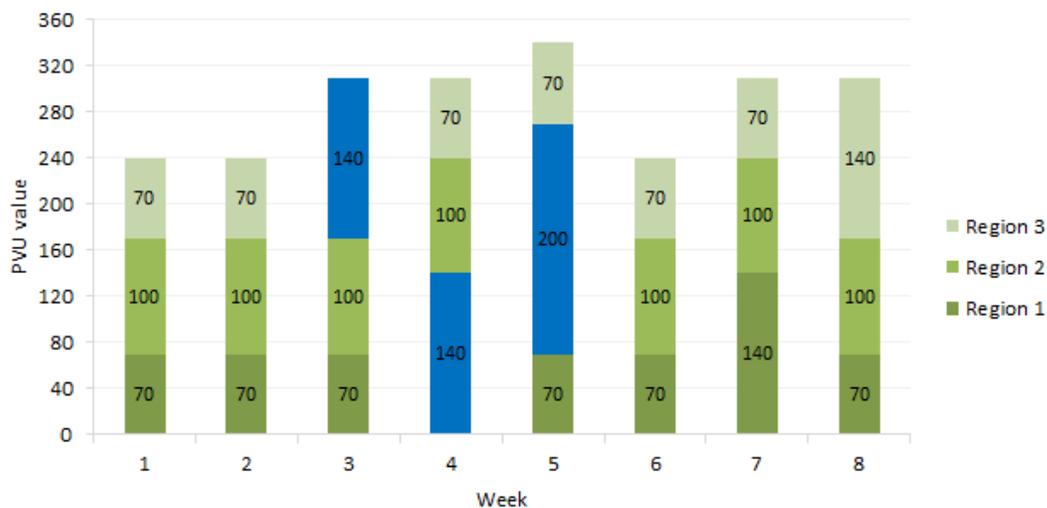
Example

A large international company has offices in Australia, Europe, and North America. To fulfill the subcapacity licensing requirements, computers in this company are divided into three groups. The computer group hierarchy looks as follows.

- All Computers
 - Region 1: North America
 - Region 2: Europe
 - Region 3: Australia

Product A is installed on computers in all regions. To obtain information about the highest PVU usage in each region, the Software Asset Manager generates the audit snapshot for each region. [Figure 1: The highest PVU usage for Product A in each region. \(on page xc\)](#) is a visualization of PVU usage for Product A. The highest PVU usage for each region is marked in blue.

Figure 1. The highest PVU usage for Product A in each region.



According to the generated snapshots, the highest PVU values are:

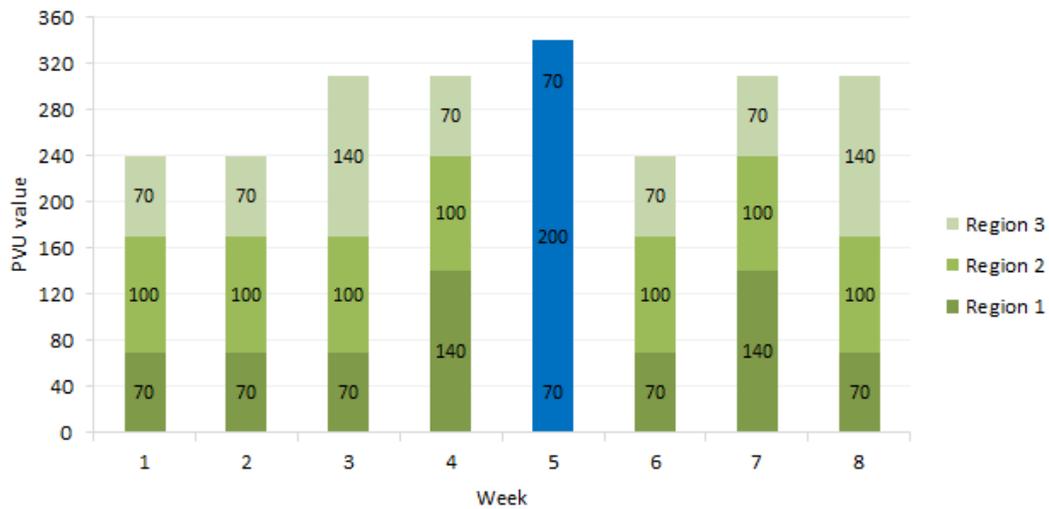
- 140 PVUs in Region 1
- 200 PVUs in Region 2
- 140 PVUs in Region 3

To obtain the overall PVU usage for Product A, the Software Asset Manager sums up the highest PVU values from each region.

$$140 + 200 + 140 = 480 \text{ PVUs}$$

The total PVU usage for Product A is 480 PVUs. If the Software Asset Manager generated a single audit snapshot for the All Computers group, the reported PVU usage would be 340 PVUs. [Figure 2: The highest PVU usage for Product A in the All Computers group. \(on page xci\)](#) is a visualization of PVU usage for Product A calculated on the level of the All Computers group. The highest PVU usage is marked in blue.

Figure 2. The highest PVU usage for Product A in the All Computers group.



The correct PVU value for Product A is the value obtained by summing up the PVU values from three regions, namely 480 PVUs.

Possible setup of the BigFix infrastructure

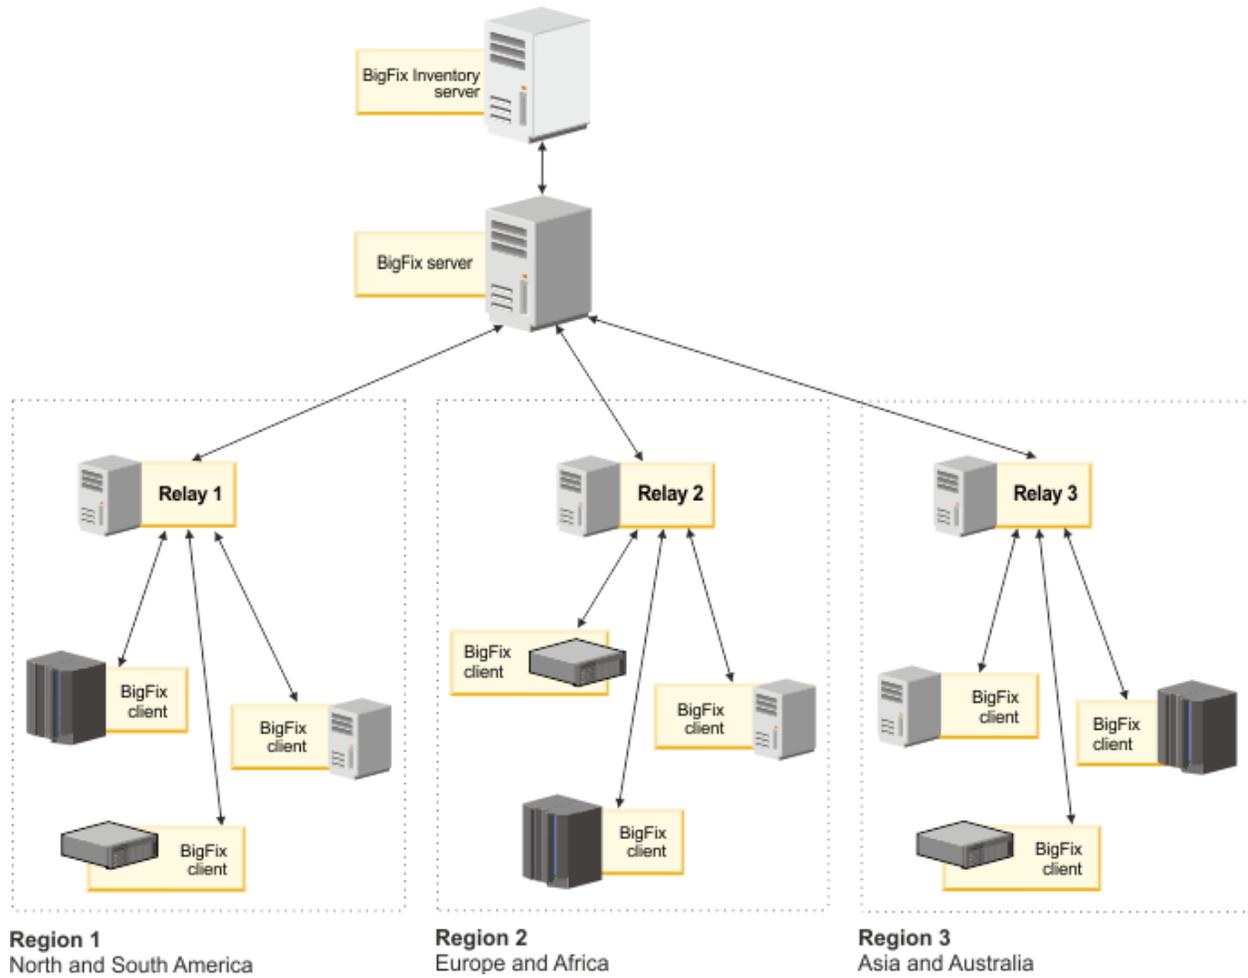
To facilitate infrastructure management and the creation of computer groups, you can set up the BigFix infrastructure in one of the following ways.

! **Important:** These setups are provided as examples only. You can set up the BigFix infrastructure in any way that suits your needs.

One BigFix server with three relays

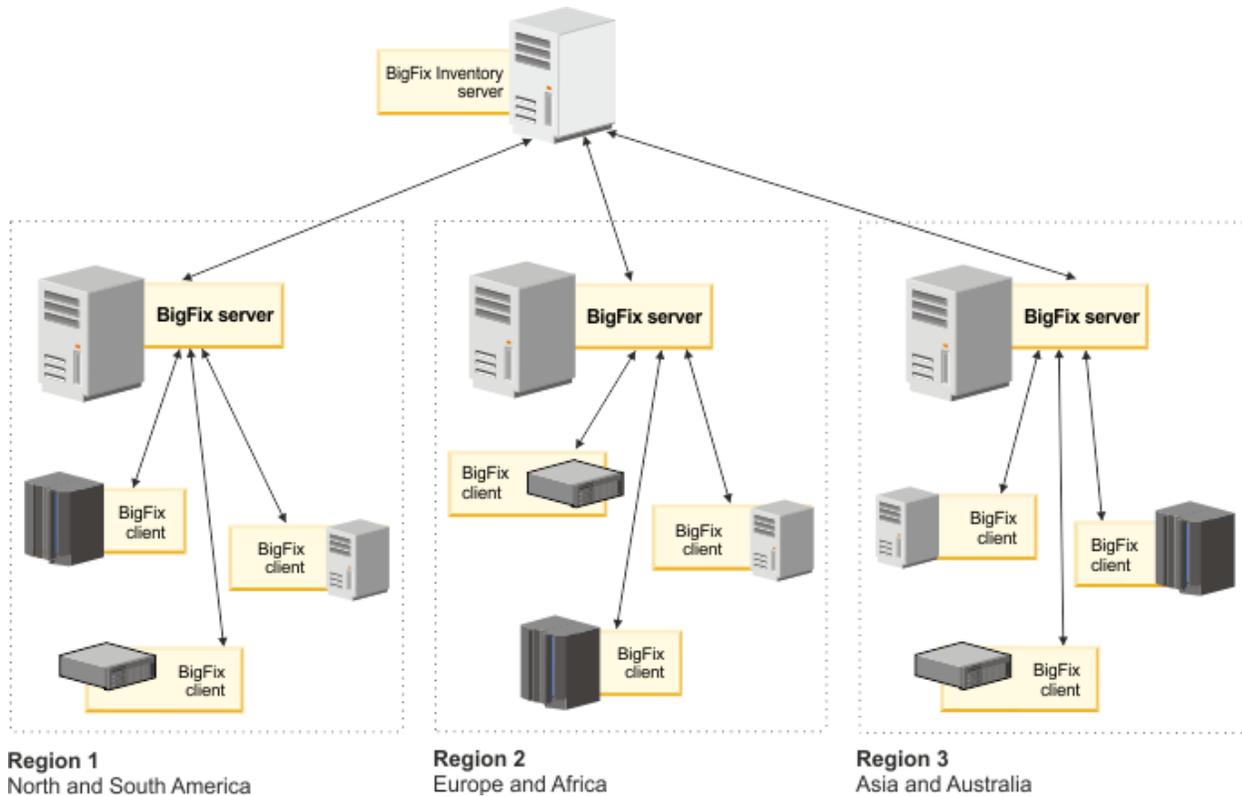
In this case, one BigFix server monitors all computers in your infrastructure. You set up three relays, each dedicated to handling the traffic from computers in one BigFix subcapacity region. Then, in the BigFix console, you create computer groups that reflect the BigFix subcapacity regions. You can later on use these groups to easily create computer groups in BigFix Inventory.

📌 Note: Computer groups created in the BigFix console are used to manage the infrastructure, and run Fixlets against selected computers. They are not automatically copied to BigFix Inventory. However, they can be used as the basis for creating computer groups that define the scope of reports in BigFix Inventory.



Three BigFix servers

In this case, you set up an infrastructure with three BigFix servers. Each server is dedicated to one BigFix subcapacity region and collects data only from the computers in that region. Data from all BigFix servers is imported to one instance of BigFix Inventory in which you create three computer groups. Each computer group reflects one BigFix server.



9.2.2 Organizational business units

9.2.2 Available from 9.2.2. To manage software and view reports per organizational business unit, divide the computers in your infrastructure into groups that represent these units. Then, create dedicated users, each with access to a particular computer group.

To manage software per organization business unit, perform the following steps:

1. Divide the computers in your infrastructure into groups that reflect the departments. Ensure that each group contains only computers from one department. Do not create computer groups that overlap as it might distort the obtained results.
2. If software in each department is managed by a different person, create dedicated BigFix Inventory users, each with access to computers in one department.
3. Import part numbers for each computer group and bundle software within each department.

For detailed instructions, see: [Tutorial: Reporting subcapacity usage per computer group \(on page dclxix\)](#).

Example

A company that is located in the United States has three departments: Finance, HR, and IT. Because employees in each department use software that is typical for their position, the Software Asset Manager wants to manage the software for each department separately. To do this, he asked the IT Architect to create three computer groups in

BigFix Inventory, each containing computers from only one business unit. The computer group hierarchy looks as follows.

- All Computers
 - Computer Group 1: Finance Department
 - Computer Group 2: HR Department
 - Computer Group 3: IT Department

The Software Asset Manager can now import part numbers that are dedicated for each computer group and manage the software within each department separately. He can also generate separate subcapacity reports to get an overview of PVU consumption for BigFix software within each department. However, to create the audit snapshot with the PVU consumption within the entire company he must generate the audit snapshot for the All Computers group.

Notices

This information was developed for products and services offered in the US.

HCL may not offer the products, services, or features discussed in this document in other countries. Consult your local HCL representative for information on the products and services currently available in your area. Any reference to an HCL product, program, or service is not intended to state or imply that only that HCL product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any HCL intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-HCL product, program, or service.

HCL may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

HCL

330 Potrero Ave.

Sunnyvale, CA 94085

USA

Attention: Office of the General Counsel

For license inquiries regarding double-byte character set (DBCS) information, contact the HCL Intellectual Property Department in your country or send inquiries, in writing, to:

HCL

330 Potrero Ave.

Sunnyvale, CA 94085

USA

Attention: Office of the General Counsel

HCL TECHNOLOGIES LTD. PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT,

MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. HCL may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-HCL websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this HCL product and use of those websites is at your own risk.

HCL may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

HCL

330 Potrero Ave.

Sunnyvale, CA 94085

USA

Attention: Office of the General Counsel

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by HCL under terms of the HCL Customer Agreement, HCL International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-HCL products was obtained from the suppliers of those products, their published announcements or other publicly available sources. HCL has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-HCL products. Questions on the capabilities of non-HCL products should be addressed to the suppliers of those products.

Statements regarding HCL's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to HCL, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. HCL, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS," without warranty of any kind. HCL shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from HCL Ltd. Sample Programs.

Trademarks

HCL Technologies Ltd. and HCL Technologies Ltd. logo, and hcl.com are trademarks or registered trademarks of HCL Technologies Ltd., registered in many jurisdictions worldwide. Other product and service names might be trademarks of HCL or other companies.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the HCL website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of HCL.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of HCL.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

HCL reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by HCL, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

HCL MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the HCL website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of HCL.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of HCL.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

HCL reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by HCL, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

HCL MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED

TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Installing

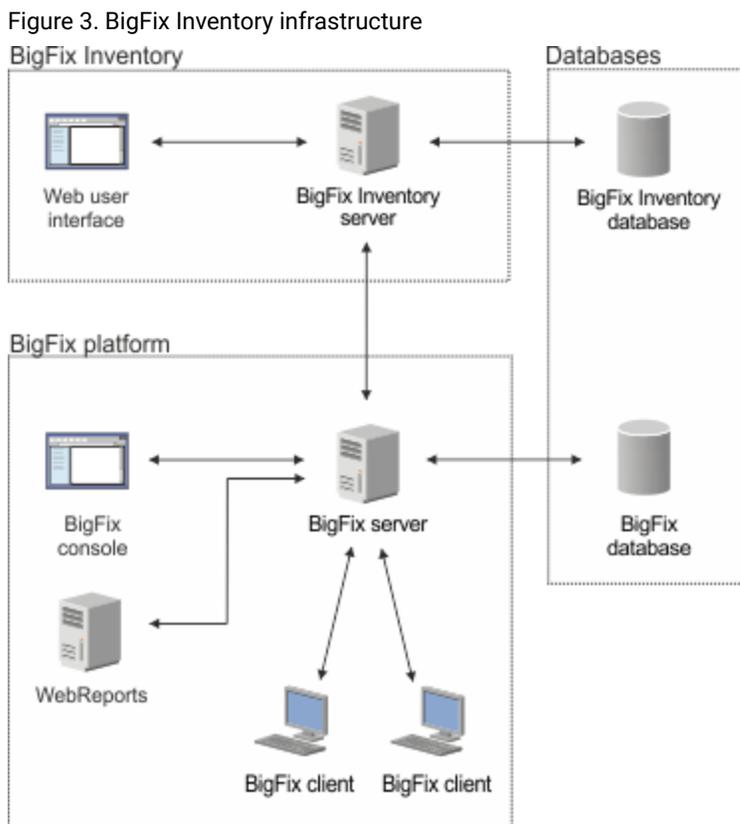
Learn about the requirements and available installation scenarios to ensure that the deployment of BigFix Inventory goes smoothly in your environment.

Planning the installation

Before you start the installation, review information about hardware and software requirements and other considerations to ensure that the installation completes successfully.

BigFix Inventory infrastructure

A complete deployment of BigFix Inventory consists of the BigFix Inventory server and its database, the BigFix server and its database, and a BigFix console. Additionally, a BigFix client must be installed on every computer from which you want to collect software inventory data. Depending on the environment size, you can install all components on a single computer, or distribute them among multiple computers.



BigFix Inventory server

The BigFix Inventory server provides a reporting interface for the inventory data that is collected from the computers managed by BigFix. To process the data in real time and to ensure reporting accuracy, the BigFix Inventory server should run all the time. Users access the application server through a web user interface.

BigFix server

The BigFix server forms the heart of the BigFix infrastructure. It coordinates the flow of data to and from individual computers. The server stores the results of capacity scans for up to five days, and thus should not be stopped for a longer period. It also stores the results of the last successful software scan from each computer. If new software scans are run more than once during server maintenance, not all of them are processed after the restart.

BigFix console

The BigFix console provides a system-wide view of all computers in the infrastructure and their configuration. It is used to target computers with specific actions such as installing the scanner or scheduling software and capacity scans.

Web Reports

Web Reports is a high-level web application that allows for visualizing data collected from the BigFix infrastructure. It is required by BigFix Inventory to connect with the VM Manager Tool that collects data from the VM managers in your infrastructure.

BigFix clients

BigFix client is common for all products that are based on the BigFix platform and should be installed on every computer, including backup and recovery machines, to ensure audit readiness. It provides information about computer properties, and is used to perform tasks such as installing the scanner on the computers. The BigFix client runs on all operating systems supported by BigFix Inventory, except for IBM i that uses disconnected scans instead of the client.

Apart from the BigFix client, a scanner should also be installed on every monitored computer. The scanner is an independent component that is used by BigFix Inventory to collect information about the hardware, as well as the software that is installed on the computers in your infrastructure.

Databases

BigFix server and BigFix Inventory server require their separate database instances. Depending on the size of your environment, these can be two database instances running on a single database engine, or two separate database engines. The data is extracted from the BigFix server database and imported to the BigFix Inventory database during the Extract, Transform, and Load (ETL) import process.

System requirements

Ensure that computers on which you want to install BigFix Inventory components run on supported operating systems and hypervisors and have the prerequisite software installed.

To view the list of supported operating systems for each of BigFix Inventory component refer to following documents:

Inventory 9.2.17: refer to **Application update 9.2.17** section in [What's new in this release \(on page viii\)](#) and the [9.2.16 System Requirements](#) document.

[9.2.16](#)

[9.2.15](#)

[9.2.14](#)

[9.2.13](#)

[9.2.12](#)

[9.2.11](#)

[9.2.10](#)

For information about technologies eligible for PVU full capacity and subcapacity, see also: [Eligible Virtualization Technology and Operating Systems](#) and [Eligible Processor Technology](#).

To learn about BigFix Inventory Platform system requirements, refer to [BigFix Platform 9.5](#).

Software requirements

Ensure that all prerequisite software is installed on the computers in your infrastructure.

Server requirements

For information about operating systems on which BigFix, DB2® and SQL Server are supported, see: [System requirements \(on page c\)](#).

Table 6. Requirements for the BigFix Inventory server

The table consists of four rows and three columns.

	Version	Description
BigFix	<ul style="list-style-type: none"> • 9.5 • 9.2 	<p>For information about the requirements, see:</p> <ul style="list-style-type: none"> • BigFix 9.5 product documentation • BigFix 9.2 product documentation <p> Note: Earlier versions of BigFix are no longer supported and thus cannot be used.</p>
DB2®	<ul style="list-style-type: none"> • DB2® 10.1 with Fix Pack 5 or higher • DB2® 10.5 with Fix Pack 6 or higher (Work- 	<p>To use DB2® as a database, BigFix Inventory must be installed on Linux™.</p> <p>DB2® requirements</p> <p>For information about the database requirements, see:</p>

Table 6. Requirements for the BigFix Inventory server*The table consists of four rows and three columns.***(continued)**

	Version	Description
	group Server Edition, Enterprise Server Edition, or Advanced Enterprise Server Edition)	<ul style="list-style-type: none"> • DB2® 10.1 product documentation • DB2® 10.5 product documentation <p>DB2® edition</p> <p>DB2® Workgroup Server Edition is sufficient to work with BigFix Inventory regardless of the number of endpoints. You can use DB2® Enterprise Server Edition, or Advanced Enterprise Server Edition to have access to advanced features of DB2®.</p> <p>DB2 settings</p> <p>Use a DB2 instance with default settings. Other required database settings are configured by BigFix Inventory installer.</p> <p> Restriction: If you specify non-default DB2 settings, BigFix Inventory might not work properly.</p> <p>Download DB2® software</p> <p>For more information, see: Downloading and installing DB2® (on page clxii).</p>
SQL Server	<ul style="list-style-type: none"> • 2008, 2008 R2 • 2012 • 2014 • 9.2.13 • 2016 	<p>To use SQL Server as a database, BigFix Inventory must be installed on Windows™.</p> <p>SQL Server version</p> <p>9.2.13 Starting from application update 9.2.13, you can use MS SQL 2016 as the database.</p> <p>If you used MS SQL 2016 with earlier versions of BigFix Inventory, your environment was non-compliant from IBM perspective because data that was stored in the database might have been corrupted. Problems with historical data cannot be fixed. However, data collected after the upgrade to BigFix Inventory 9.2.13 is valid and compliant from IBM perspective.</p> <p>SQL Server edition</p>

Table 6. Requirements for the BigFix Inventory server*The table consists of four rows and three columns.***(continued)**

Version	Description
	<p>Restriction:</p> <ul style="list-style-type: none"> • MS SQL 2008 and 2008 R2 does not support the Software Summary (Pre-view) report. The required MS SQL Server version for Software Summary (Pre-view) report is MS SQL 2012 and higher. • The BigFix installer allows for the installation of an evaluation version of MS SQL Server. However, such version expires after 180 days and should not be used in a production environment.
SQL Server settings	<p>Ensure that snapshot isolation is enabled in the SQL Server. For information about the requirements, see the Microsoft™ SQL Server documentation.</p> <p>The support for AlwaysOn solution has not been tested.</p>
	<p>Important: If you use SQL Server 2012, install cumulative update 5 for SQL Server 2012 SP2 to prevent possible corruption of the database.</p>
Permissions and roles	<ul style="list-style-type: none"> • Ensure that the MS SQL Server user has the following permissions. These permissions apply only if the databases were installed with default settings and all customizations and hardening configurations were consulted with BigFix support.

Table 6. Requirements for the BigFix Inventory server
The table consists of four rows and three columns.
 (continued)

	Version	Description
		<ul style="list-style-type: none"> ◦ For the BigFix database (BFEnterprise): <code>CREATE FUNCTION, CREATE SCHEMA, CREATE TABLE, CREATE VIEW, EXECUTE, SELECT</code> ◦ For the Web Reports database (BESReporting): <code>SELECT</code> • Ensure that the MS SQL Server user has the appropriate role to create the BigFix Inventory database. <ul style="list-style-type: none"> ◦ If you create a new database during initial configuration in BigFix Inventory, the user must have the <code>sysadmin</code> role in MS SQL Server. ◦ If you create an MS SQL database manually before the configuration, make sure that the database is empty and use the <code>SQL_Latin1_General_CP1_CS_AS</code> collation. The user for such database must have the <code>db_owner</code> role in MS SQL Server. ◦ If you are using MS SQL Server 2012 and you chose Local System account as the service owner during the installation, give the <code>dbcreator</code> or <code>sysadmin</code> role to the <code>NT AUTHORITY\SYSTEM</code> user in MS SQL Server.
<div style="background-color: #800040; color: white; padding: 2px; display: inline-block;">Linux</div> X server	<ul style="list-style-type: none"> • X11R7.x 	<p>The X server is required if you want to install or uninstall the BigFix Inventory server in interactive mode on Red Hat Enterprise Linux™. It is also required if you want to complete the server configuration by using a browser that is available on the computer where the BigFix Inventory server is installed.</p> <p style="text-align: center;"> Restriction: Xming X Server for Windows™ is not supported.</p>

Other software requirements

Table 7. Other software requirements for BigFix Inventory*The table consists of three rows and two columns.*

Required software	Required versions	Comments
Web Reports		A component of the BigFix platform that is required by BigFix Inventory to configure VM managers. Web Reports can be installed during the installation of BigFix, and must reside on the same server.
Browser	<ul style="list-style-type: none"> • BigFix Inventory is tested in Internet Explorer version 11. Most panels are also displayed correctly in Internet Explorer 10, however you can face some minor issues. It is recommended to use the newest version of Internet Explorer to ensure that all security updates and patches are applied. <p> Note: Internet Explorer must have TLS 1.2 enabled if you want to use TLS 1.2 for secure communication with SSL. You can also use TLS 1.0 for secure communication with SSL.</p> <ul style="list-style-type: none"> • Firefox 17 Extended Support Release (ESR) or higher ESR editions <p> Note: To use TLS 1.2, check whether your version of the browser supports it and whether it is enabled. Otherwise, use TLS 1.0 for secure communication with SSL.</p> <ul style="list-style-type: none"> • Chrome 35 or higher • 9.2.14 Microsoft Edge 	<p>Requirements</p> <ul style="list-style-type: none"> • JavaScript™ must be enabled in the browser. • Minimal supported screen resolution is 1360x768 pixels. • Recommended screen resolution is 1920x1080 pixels. <p>Proxies</p> <p>Proxies are supported and depend on the server configuration. With the default configuration, use SSL proxies. If you disabled SSL on the server, then use HTTP proxies. In big environments, the connections might time out with time-consuming operations (downloading extensive reports). In such a case, increase the timeout limit for communication between the proxy and BigFix Inventory.</p>

Table 7. Other software requirements for BigFix Inventory*The table consists of three rows and two columns.*

(continued)

Required software	Required versions	Comments
	 Restriction: <ul style="list-style-type: none"> ◦ Reports cannot be exported to PDF when multiple columns are selected. Export the data to CSV file instead. ◦ In reports where IP Address column is enabled, the IP address is a link that is activated by an external plug-in, such as Skype. To remove the link, disable the plug-in. 	
Linux KornShell		Original <code>ksh</code> package can be obtained from the Red Hat Enterprise Linux™ installation disc. Open source alternatives, such as <code>pd-ksh</code> and <code>mksh</code> , are not supported.
National language pack for PDF reader		To display PDF reports in a native language, a language pack that displays the native language fonts might be required.

Red Hat packages

The following packages are required if you install the components on Red Hat Enterprise Linux™.

Table 8. Red Hat packages that are required for the installation

Component	Required packages
BigFix	<ul style="list-style-type: none"> • <code>cyrus-sasl-lib.x86_64</code> • <code>krb5-libs.x86_64</code> • <code>libpng12.x86_64</code> (for Red Hat Enterprise Linux 7.x)

Table 8. Red Hat packages that are required for the installation (continued)

Component	Required packages
BigFix Inventory server	<ul style="list-style-type: none"> • <code>libstdc++.i686</code> • <code>libstdc++.x86_64</code> and all dependencies • <code>libXext.x86_64</code> (Web Reports only) • <code>libXrender.x86_64</code> (Web Reports only) • <code>zlib.x86_64</code> (Web Reports only)
DB2®	<ul style="list-style-type: none"> • <code>libstdc++.so.6.0.8</code> • Red Hat binutils

Hardware requirements

During setup, match your optimum deployment size to your hardware specifications. Use the recommendations as a general guidance.

Reference

[BigFix 9.2 requirements](#)

[BigFix 9.5 requirements](#)

Hardware requirements for the server on Windows

Ensure that the computer on which you are installing the BigFix Inventory meets the minimal CPU, and memory requirements for the server and database elements.

Virtualized environment

The ETL (extract, transform, and load) import heavily uses the database resources. BigFix Inventory and SQL Server can be installed on a virtualized environment. However, for large deployments that consist of more than 50,000 computers, it is recommended that dedicated hardware is used.

In a virtual environment for medium size deployments that consist of 10,000 - 50,000 computers, it is recommended that dedicated resources are considered for processor, memory, and virtual disk allocation. The virtual disk that is allocated for the VM should be dedicated RAID storage, with dedicated IO bandwidth for that VM.

Use SQL Server that is dedicated for BigFix Inventory and is not shared with BigFix or other applications. Fine-tuning based on the above mentioned recommendations might be required.

Processor and RAM

The values provided for BigFix were calculated with the assumption that the BigFix server, its database and Web Reports server are installed on a single server and that the BigFix server is dedicated to BigFix Inventory. If you share the BigFix server between other applications, refer to the BigFix documentation for information about hardware requirements: [BigFix requirements](#).

The values provided for BigFix Inventory were calculated for the maximum of five concurrent application users.

Table 9. Processor and RAM requirements for BigFix Inventory

This table consists of four columns, one header row, and four body rows. The second column is further divided into two sub-columns.

Environment size		Topology	Processor	Memory
Small environment	1 server	<ul style="list-style-type: none"> • BigFix, its database, and Web Reports server • BigFix Inventory, and SQL Server 	2-3 GHz, 4 cores	8 GB
	Up to 5 000 endpoints			
Medium environment	2/3 servers**	BigFix, its database, and Web Reports server	2-3 GHz, 4 cores	16 GB
	5 000 - 50 000 endpoints*	BigFix Inventory and SQL Server	2-3 GHz, 4 cores	12 - 24 GB
Large environment	2/3 servers**	BigFix, its database, and Web Reports server	2-3 GHz, 4 - 16 cores	16 - 32 GB
	50 000 - 150 000 endpoints*	BigFix Inventory and SQL Server	2-3 GHz, 8 - 16 cores	32 - 64 GB
Very large environment	2/3 servers**	BigFix, its database, and Web Reports server	2-3 GHz, 16 cores	32 - 64 GB
	More than 150 000 endpoints*	BigFix Inventory and SQL Server	2-3 GHz, 8 - 16 cores	64 - 96 GB***

* For environments with more than 35 000 endpoints, scan groups are required. For more information, see [Tuning performance \(on page dcccxxi\)](#).

** A distributed environment, where BigFix Inventory is separated from the database, is advisable.

*** SQL Server must be throttled to 3/4 RAM capacity.

Disk space

The following tables show disk space requirements for each component.



Important: The BigFix client is also installed on the computers where the components are installed. Ensure that the computers have enough space for the client. For more information, see: [Hardware requirements for the client \(on page cxix\)](#).

BigFix Inventory

Table 10. Disk space requirements for BigFix Inventory installed on Windows

This table consists of four columns, one header row, and three body rows. The second and third body rows are divided in to three rows from the second column.

Default Directory	Space re- quired	Comments
C:\Program Files (x86)\BigFix En- terprise\BES In- stallers\BFI_installer	500 MB	Compressed installer that is downloaded to the selected endpoint from BigFix. It can be deleted after extracting.
C:\Program Files\IBM \BFI	3 GB	BigFix Inventory installation directory. The amount includes space re- quired for future upgrades.
%TEMP%	350 MB	Temporary files used during the installation.
%USERPROFILE%	1 MB	The home directory of the user running the installation.

SQL Server for BigFix Inventory

Table 11. Disk space requirements for SQL Server for BigFix Inventory

This table consists of four columns, one header row, and three body rows. The second and third body rows are divided in to three rows from the second column.

Default Directory	Space re- quired	Comments
C:\Program Files\Mi- crosoft SQL Server	1 GB	Database installation directory.
C:\Program Files\Mi- crosoft SQL Serv- er\MSSQL11.MSSQLSERV- ER\MSSQL\DATA	See Com- ments	<p>Database server instance.</p> <p>The amount of disk space that is required for the database server depends on the number of computers in your environment and the average size of scan files and analyses. It can be calculated according to the following formula*:</p> <ul style="list-style-type: none"> • <i><The number of computers></i> x 1 MB + 6 GB of initial disk space <p>For example:</p> <p>10 000 computers</p> <p>10 000 x 1 MB + 6 GB = 16 GB</p> <p>100 000 computers</p> <p>100 000 x 1 MB + 6 GB = 106 GB</p>
C:\Program Files\Mi- crosoft SQL Serv- er\MSSQL11.MSSQLSERV- ER\MSSQL\DATA	See Com- ments	<p>Database transaction logs.</p> <p>During the data import (ETL process), BigFix Inventory requires some additional free disk space for database server transaction logs. The amount of disk space that is required can be significant because transaction logs store two sets of data:</p> <ul style="list-style-type: none"> • Data that is used for recovery if the ETL fails • Data that is used to create new ETL results <p>The amount of disk space that is necessary for the transaction logs depends on the number of computers in your environment as well as the number of computers in a single scan group from which data is processed during the import.</p>

Table 11. Disk space requirements for SQL Server for BigFix Inventory

This table consists of four columns, one header row, and three body rows. The second and third body rows are divided in to three rows from the second column.

(continued)

Default Directory	Space re- quired	Comments
C:\Program Files\Mi- crosoft SQL Serv- er\MSSQL11.MSSQLSERV- ER\MSSQL\DATA	See Com- ments	<p>To lower the amount of disk space that is necessary for transactions logs, distribute the scans over time so they are processed during several data imports instead of one.</p> <p>The size of transaction logs can be calculated according to the following formula*:</p> <ul style="list-style-type: none"> • <i><The number of computers> x 1.2 MB + <the number of computers in the biggest scan group> x 1.2 MB + 17 GB of initial disk space</i> <p>For example:</p> <p>10 000 computers and 10 000 scan results</p> $10\,000 \times 1.2\text{ MB} + 10\,000 \times 1.2\text{ MB} + 17\text{ GB} = 41\text{ GB}$ <p>100 000 computers and 15 000 scan results</p> $100\,000 \times 1.2\text{ MB} + 15\,000 \times 1.2\text{ MB} + 17\text{ GB} = 155\text{ GB}$ <p>tempdb database that is used to store and manage temporary objects.</p> <p>The size of tempdb can be calculated according to the following formula*:</p> <ul style="list-style-type: none"> • <i><The number of computers> x 0.5 MB</i> <p>For example:</p> <p>15 000 computers</p> $15\,000 \times 0.5\text{ MB} = 7.5\text{ GB}$

*The disk space requirements are calculated based on the average number of files that are discovered in your environment. The values that are provided in the formula assume that the average number of the discovered files is 800 per computer, which is represented by 1 and 1.2 MB. If, according to your estimates, the number of files is higher, increase this value adequately. If you already installed BigFix Inventory, you can use the Scanned File Data and Computers reports for your calculations. Note, that the required initial disk space (15 or 17 GB) is constant. The number of discovered files not only affects the size of the database, but also the server runtime memory

requirements. For more information, see: [Tuning performance in medium and large environments \(on page dcccxliv\)](#). To see the list of extensions that are discovered by BigFix Inventory by default and learn how to decrease the number of detected files, see: [Optimizing the volume of scanned file data \(on page cclxix\)](#).

* The formulas apply to typical environments that are configured to run weekly software scans, daily data imports, and whose endpoints have about 60 software installations each. The results also depend on the amount of data returned by the scans, which means that in some environments the required amount of disk space might be smaller or bigger. In case of irregular data imports or accumulated scans, the required disk space increases.



Note: The retention period in this typical environment is set to the last 7 days. For more information about tuning this parameter, see: [Shortening the retention period gradually to avoid problems with growing database size \(on page dcccxliv\)](#).

BigFix

Table 12. Disk space requirements for BigFix installed on Windows

This table consists of four columns, one header row, and three body rows. The second and third body rows are divided in to three rows from the second column.

Default Directory	Space required	Comments
C:\Program Files\BigFix Enterprise\BES Server	1.8 GB	BigFix installation directory.
C:\Program Files (x86)\BigFix Enterprise\BES Server\www-rootbes\bfmirror\downloads\shal	Around 2 GB	BigFix cache. It stores files that are downloaded by some of the fixlets before they are distributed to the selected endpoints. For example, the BigFix Inventory installer. The required space might increase if you run some of the fixlets multiple times. For example, each time you run the Upgrade to the latest version of BigFix Inventory 9.x , a new installer is downloaded to the cache. Thus, the required space increases. To check whether a fixlet downloads any files and what their size is, log in to the BigFix console, and check the value in the Download Size column for the particular fixlet. Then, ensure that enough disk space is available in the BigFix cache.
C:\Program Files\BigFix Enterprise\BES Console	90 MB	BigFix console.

SQL Server for BigFix

For information about hardware requirements for the SQL Server that is used as the BigFix database, see the [BigFix Version 9.x: Capacity Planning, Performance, and Management Guide](#).

Apart from the disk space that is described in the guide, ensure additional disk space for transaction logs on the computer where the BigFix database is installed. To calculate the required disk space, check how many objects exist in all fixlet sites that you have enabled in the BigFix console. An object is every computer group, analysis, fixlet, and task that exist in the console, including the ones that are not relevant. Every 1000 objects requires 1GB of free disk space. For example, if you have 500 fixlets and tasks, 300 analyses, and 20 computer groups, you have 820 objects in total. Thus, 1 GB of disk space is required.

Reference

[Tuning performance \(on page dcccxxxi\)](#)

Hardware requirements for the server on Linux

Ensure that the computer on which you are installing the BigFix Inventory meets the minimal CPU, and memory requirements for the server and database elements.

Virtualized environment

The ETL (extract, transform, and load) import heavily uses the DB2 database resources. BigFix Inventory and DB2 server can be installed on a virtualized environment. However, for large deployments that consist of 50 000 - 100 000 computers, it is recommended that dedicated hardware is used.

In a virtual environment for medium size deployments that consist of 10 000 - 50 000 computers, it is recommended that dedicated resources are considered for processor, memory, and virtual disk allocation. The virtual disk that is allocated for the VM should be dedicated RAID storage, with dedicated IO bandwidth for that VM.

Use DB2 server that is dedicated for BigFix Inventory and is not shared with BigFix or other applications. Fine-tuning based on the above mentioned recommendations might be required.

Processor and RAM

The values provided for BigFix were calculated with the assumption that the BigFix server, its database and Web Reports server are installed on a single server and that the BigFix server is dedicated to BigFix Inventory. If you share the BigFix server between other applications, refer to the BigFix documentation for information about hardware requirements: [BigFix requirements](#).

The values provided for BigFix Inventory were calculated for the maximum of five concurrent application users.

Table 13. Processor and RAM requirements for BigFix Inventory

The table consists of four columns and four body rows. The second column is further divided into two sub-columns.

Environment size	Topology		Processor	Memory
Small environment	1 server	BigFix, its database, and Web Reports server, BigFix Inventory and DB2	2-3 GHz, 4 cores	8 GB
Up to 5 000 endpoints				
Medium environment	2/3 servers**	BigFix, its database, and Web Reports server BigFix Inventory and DB2	2-3 GHz, 4 cores 2-3 GHz, 4 cores	16 GB 12 - 24 GB
5 000 - 50 000 endpoints*				
Large environment	2/3 servers**	BigFix, its database, and Web Reports server BigFix Inventory and DB2	2-3 GHz, 4 - 16 cores 2-3 GHz, 8 - 16 cores	16 - 32 GB 32 - 64 GB
50 000 - 150 000 endpoints*				
Very large environment	2/3 servers**	BigFix, its database, and Web Reports server BigFix Inventory and DB2	2-3 GHz, 16 cores 2-3 GHz, 8 - 16 cores	32 - 64 GB 64 - 96 GB
More than 150 000 endpoints*				

* For environments with more than 35 000 endpoints, scan groups are required. For more information, see [Tuning performance \(on page dcccxxi\)](#).

** A distributed environment, where BigFix Inventory is separated from the database, is advisable.

Disk space

The following tables show disk space requirements for each component.



Important: The BigFix client is also installed on the computers where the components are installed. Ensure that the computers have enough space for the client. For more information, see: [Hardware requirements for the client \(on page cxix\)](#).

BigFix Inventory

Table 14. Disk space requirements for BigFix Inventory installed on Linux

The table consists of three columns and five body rows. The first body row is further divided into two rows.

Default Directory	Space required	Comments
<code>user_directory/BFI_installer</code>	500 MB	Compressed installer that is downloaded to the selected endpoint from BigFix. The compressed installer can be deleted after it is extracted.
	500 MB	Extracted installation files.
<code>/opt/ibm/BFI</code>	3 GB	BigFix Inventory installation directory. The amount includes space required for future upgrades.
<code>/tmp</code>	350 MB	Temporary files used during the installation.
		 Important: The installer must be able to run executable files in this directory.
<code>\$HOME</code>	1 MB	Home directory of the user running the installation.
<code>/etc</code>	1 MB	Directory that stores scripts that start the server.
<code>/var</code>	10 kB	Directory that stores the installation registry.

DB2 for BigFix Inventory

Table 15. Disk space requirements for DB2 for BigFix Inventory

The table consists of three columns and four body rows.

Default Directory	Space required	Comments
<code>/opt/IBM/db2</code>	1.5 GB	DB2 installation directory.
<code>/var/db2</code>	10 kB	Directory that stores DB2 global registry.
<code>/home/db2inst1</code>	See Comments	Database server instance. The amount of disk space that is required for the database server depends on the number of computers in your environment and the average size of scan files and analysis. It can be calculated according to the following formula*:
		<ul style="list-style-type: none"> • $\langle \text{The number of computers} \rangle \times 1 \text{ MB} + 15 \text{ GB}$ of initial disk space <p>For example:</p>

Table 15. Disk space requirements for DB2 for BigFix Inventory

The table consists of three columns and four body rows.

(continued)

Default Directory	Space required	Comments
/home/db2inst1/ st1/db2inst1/	10 000 computers	10 000 x 1 MB + 15 GB = 25 GB
	100 000 computers	100 000 x 1 MB + 15 GB = 115 GB
	See Comments	Database server transaction logs. During the data import (ETL process), BigFix Inventory requires some additional free disk space for database server transaction logs. The amount of disk space that is required can be significant because transaction logs store two sets of data:
	<ul style="list-style-type: none"> • Data that is used for recovery if the ETL fails • Data that is used to create new ETL results 	
		The amount of disk space that is necessary for the transaction logs depends on the number of computers in your environment as well as the number of computers in a single scan group from which data is processed during the import.
		To lower the amount of disk space that is necessary for transactions logs, distribute the scans over time so they are processed during several data imports instead of one.
		The size of transaction logs can be calculated according to the following formula*:
		<ul style="list-style-type: none"> • <i><The number of computers> x 1.2 MB + <the number of computers in the biggest scan group> x 1.2 MB + 17 GB of initial disk space</i>
		For example:
	10 000 computers and 10 000 scan results	10 000 x 1.2 MB + 10 000 x 1.2 MB + 17 GB = 41 GB
	100 000 computers and 15 000 scan results	100 000 x 1.2 MB + 15 000 x 1.2 MB + 17 GB = 155 GB

Table 15. Disk space requirements for DB2 for BigFix Inventory

The table consists of three columns and four body rows.

(continued)

Default Directory	Space required	Comments
<code>/home/db2fenc1</code>	1 MB	Home directory of the DB2 fenced user.

*The disk space requirements are calculated based on the average number of files that are discovered in your environment. The values that are provided in the formula assume that the average number of the discovered files is 800 per computer, which is represented by 1 and 1.2 MB. If, according to your estimates, the number of files is higher, increase this value adequately. If you already installed BigFix Inventory, you can use the Scanned File Data and Computers reports for your calculations. Note, that the required initial disk space (15 or 17 GB) is constant. The number of discovered files not only affects the size of the database, but also the server runtime memory requirements. For more information, see: [Tuning performance in medium and large environments \(on page dcccxliv\)](#). To see the list of extensions that are discovered by BigFix Inventory by default and learn how to decrease the number of detected files, see: [Optimizing the volume of scanned file data \(on page cclxix\)](#).

* The formulas apply to typical environments that are configured to run weekly software scans, daily data imports, and whose endpoints have about 60 software installations each. The results also depend on the amount of data returned by the scans, which means that in some environments the required amount of disk space might be smaller or bigger. In case of irregular data imports or accumulated scans, the required disk space increases.



Note: The retention period in this typical environment is set to the last 7 days. For more information about tuning this parameter, see: [Shortening the retention period gradually to avoid problems with growing database size \(on page dcccxliv\)](#).

BigFix

Table 16. Disk space requirements for BigFix installed on Linux

The table consists of three columns and five body rows.

Default Directory	Space required	Comments
<code>/var/opt/BESIn- stallers</code>	90 MB	Directory with console and client installers.
<code>/var/opt/BESServer</code>	2 GB	BigFix installation directory.
<code>/var/opt/ BESServer/www-</code>	Around 2 GB	BigFix cache. It stores files that are downloaded by some of the fixlets before they are distributed to the selected endpoints. For example, the BigFix Inventory installer.

Table 16. Disk space requirements for BigFix installed on Linux

The table consists of three columns and five body rows.

(continued)

Default Directory	Space required	Comments
<code>rootbes/bfmirror/downloads/shal/</code>		The required space might increase if you run some of the fixlets multiple times. For example, each time you run the Upgrade to the latest version of BigFix Inventory 9.x , a new installer is downloaded to the cache. Thus, the required space increases. To check whether a fixlet downloads any files and what their size is, log in to the BigFix console, and check the value in the Download Size column for the particular fixlet. Then, ensure that enough disk space is available in the BigFix cache.
<code>/var/opt/BESWebReportsServer</code>	300 MB	Web Reports installation directory.
<code>/opt/BESWebReportsServer</code>	500 MB	Directory with the Web Reports server binaries.
<code>/opt/BESServer</code>	100 MB	Directory with server binaries.
<code>/var/log</code>	1 MB	Directory with log files.
<code>C:\Program Files\BigFix Enterprise\BES Console</code>	90 MB	BigFix console installation directory. The console must be installed on Windows.

For disk speed, see: [Storage performance requirements \(on page dcccxxij\)](#).

DB2 for BigFix

For information about hardware requirements for DB2 that is used as the BigFix database, see the [BigFix Version 9.x: Capacity Planning, Performance, and Management Guide](#).

Apart from the disk space that is described in the guide, ensure additional disk space for transaction logs on the computer where the BigFix database is installed. To calculate the required disk space, check how many objects exist in all fixlet sites that you have enabled in the BigFix console. An object is every computer group, analysis, fixlet, and task that exist in the console, including the ones that are not relevant. Every 1000 objects requires 1GB of free disk space. For example, if you have 500 fixlets and tasks, 300 analyses, and 20 computer groups, you have 820 objects in total. Thus, 1 GB of disk space is required.

Reference

[Tuning performance \(on page dcccxxxi\)](#)

Hardware requirements for the client

Review important information about hardware requirements for the BigFix client and the software and capacity scans that are embedded in it.

Processor and RAM

An BigFix client alone can consume up to 2% of the processing power of one processor core on an endpoint. However, the client is complemented with software and capacity scans that collect necessary software and hardware information from your endpoints. Although the capacity scan reports very low CPU usage, the software scan can consume substantial CPU resources while a scan is in progress. To decrease the impact of a software scan on production system, it can be scheduled to run on the weekends or in the evenings. You can also run the software scan with the CPU threshold option that limits the consumption of your CPU resources.

Table 17. CPU and RAM usage for BigFix clients

This table consists of four columns, one header row, and three body rows.

Component	CPU	RAM	Comments
BigFix client	< 2 %	< 20 MB	For more information, see: https://support.hcltechsw.com/csm?id=bigfix_support .
Software scan	Up to 100 % You can limit the usage by running the scan with the CPU threshold attribute.	< 80 MB	The software scan runs on demand, and can be monitored by checking the following processes: <code>wscansw</code> , <code>wscanfs</code> .
Capacity scan	< 1 %	< 20 MB	The capacity scan runs every 30 minutes, and can be monitored by checking the following process: <code>wscanhw</code> .

Disk space

Ensure that endpoints have enough disk space before you start installing the BigFix clients.



Important: The following disk space requirements do not include the size of files that are downloaded to the endpoint by some of the fixlets. For example, the fixlet **Download BigFix Inventory**, downloads the application installer to the specified endpoint. Thus, more space is required on the targeted endpoint.



To check whether a fixlet downloads any files to the targeted endpoint, log in to the BigFix console, and check the value in the Download Size column for the particular fixlet. Then, ensure that enough disk space is available on the endpoint.

Table 18. Disk space requirements for clients installed on UNIX

This table consists of four columns, one header row, and one body row. The body row is divided in to six rows from the second column.

Operating system	Directory	Space required	Comments
UNIX	<code>/opt/BESClient</code>	80 MB	Client installation directory.
	<code>/var/opt/BESClient</code>	250 MB	Client data directory. This directory contains all scan results.
		200 MB	The additional space is required if the VM Manager Tool is installed on the endpoint.
	<code>/opt/tivoli/cit</code>	50 MB	Scanner installation directory.
	<code>/opt/tivoli/cit/cache_data</code>	100 MB on average	Scanner cache files. The required disk space depends on the number of files, directories, and subdirectories to be scanned. It can be estimated by multiplying the number of files to be scanned by 60 bytes.
	<code>/etc/cit</code>	< 1 MB	Scanner configuration files.
	<code>/tmp</code>	100 MB on average	Scanner temporary files. The required disk space depends on the scanner options that are used (for example, sorting) as well as the number of files, directories, and subdirectories to be scanned. It can be estimated by multiplying the number of files to be scanned by 100 bytes.
	<code>/usr/ibm/tivoli/common/CIT</code>	10 MB	Scanner log files.

Table 19. Disk space requirements for clients installed on Windows

This table consists of four columns, one header row, and one body row. The body row is divided in to five rows from the second column.

Operating system	Directory	Space required	Comments
Windows	C:\Program Files (x86)\BigFix Enterprise\BES Client	150 MB	Client installation and data directory. This directory contains all scan results.
		300 MB	The additional space is required if the VM Manager Tool is installed on the endpoint.
	C:\Program Files\tivoli\cit	20 MB	Scanner installation directory.
	C:\Program Files\tivoli\cit\cache_data	50 MB on average	Scanner cache files. The required disk space depends on the number of files, directories, and subdirectories to be scanned. It can be estimated by multiplying the number of files to be scanned by 60 bytes.
	%WINDIR%	< 1 MB	Scanner configuration files.
	%TMP% or %TEMP%	100 MB on average	Scanner temporary files. The required disk space depends on the scanner options that are used (for example, sorting) as well as the number of files, directories, and subdirectories to be scanned. It can be estimated by multiplying the number of files to be scanned by 100 bytes.
	C:\Program Files\tivoli\ibm\tivoli\common\CIT\logs	10 MB	Scanner log files.

Port requirements

When planning the infrastructure, ensure that port numbers used by BigFix Inventory, BigFix, and the database are free to enable communication between those components.

The following is the list of default ports used by the BigFix Inventory infrastructure. You can change them during the installation of each component:

Table 20. Default ports used by the BigFix Inventory infrastructure

This table consists of three columns, one header rows, and three body rows. The first and third body rows are divided in to two rows from the second column.

Type	Port number	Description
BigFix Inventory	9081	The web browser connects to the server (HTTPS) to display the user interface.
		The BigFix server uses this port to connect to the BigFix Inventory server.
DB2®	50000	The server connects to DB2®.
SQL Server	1433	The server connects to SQL Server.
BigFix	52311	BigFix clients and console connect to the server.
		The BigFix Inventory server uses this port to connect to the BigFix server.

For more information about port numbers and interactions between components of the BigFix Inventory infrastructure, see: [Flow of data \(on page dcxc\)](#).

Make sure that relevant web addresses are accessible from the computer where the BigFix server is installed. For more information, see: [Firewall exceptions \(on page cxxii\)](#).

Firewall exceptions

Some of the fixlets require that the BigFix server connects to the Internet and downloads necessary files and updates. To ensure that they can be downloaded, relevant web addresses must be accessible from the computer where the server is installed. Add those addresses as firewall exceptions and ensure that they are accessible to the proxy server if you are using it.

Ensure that the following web addresses are accessible from the computer where the BigFix server is installed

- esync.bigfix.com
- gatherer.bigfix.com
- software.bigfix.com
- support.bigfix.com
- sync.bigfix.com

The HTTP port 80 must be open for communication.

9.2.12 BigFix Inventory occasionally offers you a customer feedback survey. This quick survey allows you to rate the product and provide valuable feedback that can help us improve our services. To be able to participate in the survey, allow the connection to the following web page from your browser: <https://survey.medallia.eu/>.

Proxy exceptions

If the BigFix server uses a proxy for Internet connection to access fixlet sites or prefetch downloads, the BigFix Inventory server must be added to an exception list to ensure that the catalog propagation works properly.

In an environment with a proxy, propagation of the catalog might fail because the BigFix server connects to the proxy to download the catalog file instead of connecting to the BigFix Inventory server directly. To ensure that the problem does not occur, add the following entries to the proxy exception list in the BigFix server configuration:

- IP address of the BigFix Inventory server
- FQDN host name of the BigFix Inventory server

Additionally if the BigFix server is installed on the same computer as BigFix Inventory server, add the following entries to the proxy exception list:

```
localhost, 127.0.0.1
```

For more information, see: [Setting a proxy connection on the server](#).

Installation accounts

You can install all the infrastructure components as the administrative user. You can also install some components as a non-administrative user but some limitations apply.

Table 21. Installation accounts for the components of the BigFix Inventory infrastructure

Component	User	Limitations
BigFix Inventory server on Linux	root or non-root user	Server that is installed by a non-root user is not registered as a system service. It also cannot be upgraded with a fixlet. It must be upgraded in interactive or silent mode.
BigFix Inventory server on Windows	User with administrative privileges with the log on as a service permission. If you want to use the Windows authentication mode and your database server is remote, ensure that the user that accesses the database is a domain user. This user needs to have the dbcreator or sysadmin rights in MS SQL database.	
DB2®	root or non-root user	For information about limitations that apply when you install DB2® as a non-root user, see the following links.

Table 21. Installation accounts for the components of the BigFix Inventory infrastructure (continued)

Component	User	Limitations
		<ul style="list-style-type: none"> • Non-root installation overview DB2® 10.1 documentation • Non-root installation overview DB2® 11.1 documentation
SQL Server	User with administrative privileges	
BigFix server	User with administrative privileges (Windows) or root (Linux)	
BigFix client	User with administrative privileges (Windows) or root (Linux and UNIX)	The user on Linux and UNIX cannot be a pseudo user.

Coexistence scenarios

This topic describes supported scenarios for the coexistence of BigFix Inventory and License Metric Tool.

Coexistence of and in the same infrastructure

License Metric Tool and BigFix Inventory can coexist within the same infrastructure. They can be connected to the same BigFix server but monitor different sets of computers. You can then use BigFix Inventory to discover non-BigFix software and measure its license metric utilization on selected computers, and use License Metric Tool to monitor computers on which only BigFix is installed. Thanks to connecting both applications to a common BigFix server, you can manage them by using one console. It makes monitoring the environment much easier.

Although License Metric Tool and BigFix Inventory are connected to a common BigFix server, they work as separate applications and must be installed on separate computers. Each application manages a separate set of computers that can report to only one application at a given moment. The applications also use different fixlet sites and dedicated sets of fixlets for installation and configuration. Apart from the common BigFix server, all other aspects, such as software and capacity scans or the software catalog, are specific to each application.

Limitations

- **Separate installation**

BigFix Inventory and License Metric Tool must be installed on separate computers.

- **VM managers**

You can add VM managers through the user interface only for the application that has the BigFix server subscribed to its fixlet site. It happens because VM managers added through the user interface must be managed by the VM Manager Tool that is installed on the BigFix server. For the other application, use advanced VM management. For more information, see: [Advanced VM management \(distributed\) \(on page ccclxvii\)](#).

If you collect VM manager data directly from virtualization hosts by using the Run Capacity Scan on Virtualization Hosts fixlet, the coexistence is possible only if all endpoints from this host report to the same application. For more information about the fixlet, see: [Collecting capacity data from virtualization hosts for Xen and KVM \(on page ccxvi\)](#).

- **Many re-subscriptions**

You might encounter problems with the missing software catalog or scan results if you often re-subscribe endpoints between the two applications. These problems occur only if you re-subscribe endpoints more than once. For example, if you re-subscribe an endpoint from BigFix Inventory to License Metric Tool, and then subscribe it back to BigFix Inventory.

For more information about the coexistence of and in the same infrastructure, see the documentation.

9.2.1 Adding a second BigFix Inventory server

9.2.1 Available from 9.2.1. Multiple instances of BigFix Inventory can coexist within the same infrastructure.

They can be connected to the same BigFix server but monitor different computers that are grouped according to specific criteria. It allows for reporting utilization of license metrics for subsets of computers, managing parts of the environment by separate teams, or distributing the workload by importing data from a smaller number of computers.



Important: You can fulfill the same scenarios by dividing computers into groups in a single instance of BigFix Inventory. It is the recommended approach. For more information, see: [Tutorial: Reporting subcapacity usage per computer group \(on page dclxix\)](#).

To add another instance of BigFix Inventory, install another BigFix Inventory server and connect it to the same BigFix platform. Ensure that both instances of BigFix Inventory use the same fixlet site. From the BigFix perspective, the separation of servers is not visible. The servers use the same fixlet site, computers, and tasks that are used to scan the environment and discover software. The separation occurs during the import of data from BigFix to BigFix Inventory servers.

To import data from different sets of computers, divide the computers into groups by giving them unique computer settings. Then, create an analysis that retrieves the values of these settings. Finally, modify the data import so that each BigFix Inventory server imports data only from computers that have a particular setting, and omits the remaining ones. As a result, each BigFix Inventory server manages a separate group of computers.



Note: Consider the following limitation only if you initially installed BigFix Inventory version 9.2.3 or earlier, regardless of whether you upgraded to a newer version since:

The coexisting servers must use the same database software, either SQL Server or DB2, and have the same software catalog uploaded. It is required because you will later disable the catalog delivery in one of the servers, so that a common catalog is always delivered to all computers. Different signatures in catalogs



for SQL Server and DB2 would cause problems in software discovery. For more information, see [Additional considerations \(on page cxxviii\)](#).

Complete the following procedures to enable the coexistence.

9.2.1 Dividing computers between servers

9.2.1 Available since 9.2.1. Divide your computers between two servers by specifying unique computer settings. Based on these settings, the servers choose specific computers to import the data from. As a result, each of your servers manages a separate group of computers. Plan the division, and then mark your computers as either *BF1* or *BF2*.

This procedure shows only one way of dividing your computers between two servers. You can specify also other settings, or use relevance expressions to divide the computers according to their operating systems, computer groups, or any custom criteria. For more information about creating properties and analyses that are required to retrieve them, see [Creating retrieved properties](#) and [Creating analyses](#).

1. Log in to the BigFix console.
2. Add new computer settings to your computers. Based on these settings, your computers are divided between two instances of BigFix Inventory.
 - a. In the navigation tree, click **Computers**.
 - b. Select a group of computers that you want to add to one of the servers.
 - c. Right-click these computers, and then click **Edit Computer Settings**.
 - d. Select the **Custom Setting** check box, and enter the following values:

Setting Name: *Instance*
Setting Value: *BF1* or *BF2*

The value, *BF1* or *BF2*, depends on the instance that you want to add a specific computer to. Specify *BF1* for a group of computers, and *BF2* for another group.
 - e. Click **OK** to save the setting. Repeat the preceding steps until all your computers have a value *BF1* or *BF2* specified.
3. Create a custom analysis to retrieve the specified settings from your computers.
 - a. To create an analysis, click **Tools > Create New Analysis**.
 - b. Specify the name as *Instance Assignment*.
 - c. Click the **Properties** tab, and then click **Add Property**.
 - d. Specify the name of the property as *Instance*.

e. In the relevance field, enter the following string:

```
value of setting "Instance" of client|"
```

This relevance expression returns the value of the setting *Instance* that you added to your computers. It can return either *BF1* or *BF2*.

f. Click **OK** to save the analysis. The analysis is added to your Master Action Site.

4. Verify that the analysis is returning correct values.

a. Click **Sites > Master Action Site**.

b. Select your custom **Instance Assignment** analysis.

c. Click the **Results** tab. The analysis shows a property called *Instance*, which returns either *BF1* or *BF2*.

If you did not set this property for some computers, the value is empty.

You divided your computers based on specific computer settings. You can now modify the data import settings in both instances of BigFix Inventory to specify whether the servers should import data about either *BF1* or *BF2* computers.

9.2.1 Modifying servers to include specific computers

9.2.1 Available from 9.2.1. Modify the data import settings so that each of your servers imports the data only from a specific group of computers. The servers choose only those computers that meet the specified criteria, for example, that are marked with a value *BF1*, but not *BF2*. Other computers are omitted during the import.

The main element that decides which computers report to which server is the custom analysis that you created to retrieve some information about your computers, in this case specific values like *BF1* or *BF2*. The server checks the analysis results to verify values that are attached to your computers. If a computer has a value that meets the server criteria, then data from this computer is imported. Otherwise, the computer is discarded.

1. Go to `installation_directory/wlp/usr/servers/server1/config`.
2. Edit the `etl_settings.yml` file and specify the following properties and values:

scoped_etl_enabled: 1

Specifies whether a scoped data import should be enabled. If set to *0*, data from all computers is imported.

scoped_etl_site_name: ActionSite

Specifies the name of the Fixlet site that contains your custom analysis. *ActionSite* stands for Master Action Site.

scoped_etl_site_is_custom: false

Specifies whether the Fixlet site is a custom site.

scoped_etl_analysis_name: Instance Assignment

Specifies the name of your custom analysis.

scoped_etl_analysis_property_id: 1

Specifies the ID of the analysis property *Instance*. If you added just one property, it has the ID of 1. However, you can check it by exporting the analysis to a `bes` file and opening it in a text editor.

scoped_etl_based_on: MATCH

Specifies conditions of a scoped data import. In this case, you can enter any value that is different from the default `SITE_URL_FOR_SITE_NAME`. The default value would import data from all computers in a Fixlet site.

scoped_etl_resultstext_match: BF1%

Specifies the value based on which the server selects its computers. Specify *BF1* for one server, and *BF2* for another. A value returned by the analysis might include a space or a new line at the end. Enter a percent sign (%) after the value to include such occurrences instead of discarding a computer.

scoped_etl_overwrite_on_migration: false

Specifies whether to overwrite these settings with default ones after an upgrade or migration.

3. Restart the server.
4. Run a data import.

You modified the servers to import data only from specific computers. You can open the **Computers** report to verify that only chosen computers are shown.

Review additional considerations and apply them to your environment to avoid problems with coexistence.

Additional considerations

To ensure that no problems occur while working with coexisting servers, review the following considerations and apply them to your environment.

Software catalogs

To ensure that software on your computers is discovered based on the same software signatures, [upload the same software catalogs \(on page cdxliv\)](#) to both servers. You can then disable the catalog delivery on one of them to avoid overwriting. The remaining server will deliver the catalog to all computers, regardless of the server that they report to, because the computers are in the same Fixlet site.

To disable the catalog delivery, run the following query on the database:

```
UPDATE DBO.SYSTEM_GLOBALS SET VALUE = '--- false
' WHERE NAME='catalog_download_actions_enabled';
```



Note: Consider the following limitation only if you initially installed BigFix Inventory version 9.2.3 or earlier, regardless of whether you upgraded to a newer version since:



Use the same database software for both servers, either DB2 or SQL Server. Signatures are different between these two types, and one software catalog might cause problems with discovery.

VM managers

You can use the user interface to add VM managers only for one of the servers, namely the one that has the BigFix server assigned to it. This is because VM managers added in the user interface must be managed by the VM Manager Tool that is installed on the BigFix server. For the other server, you can use [advanced VM management \(on page ccclxvii\)](#), in which you install additional tools on any computer.

Changing assignments

If you want to change a computer assignment, for example from *BF1* to *BF2*, ensure that you edit the existing setting instead of creating a new one. To edit the setting, complete the following steps:

1. In the BigFix console, click **Computers**.
2. Right-click a single computer, and then click **Edit Computer Settings**.
3. Select the *Instance* setting, click **Edit**, and specify a new value.

Missing scan results

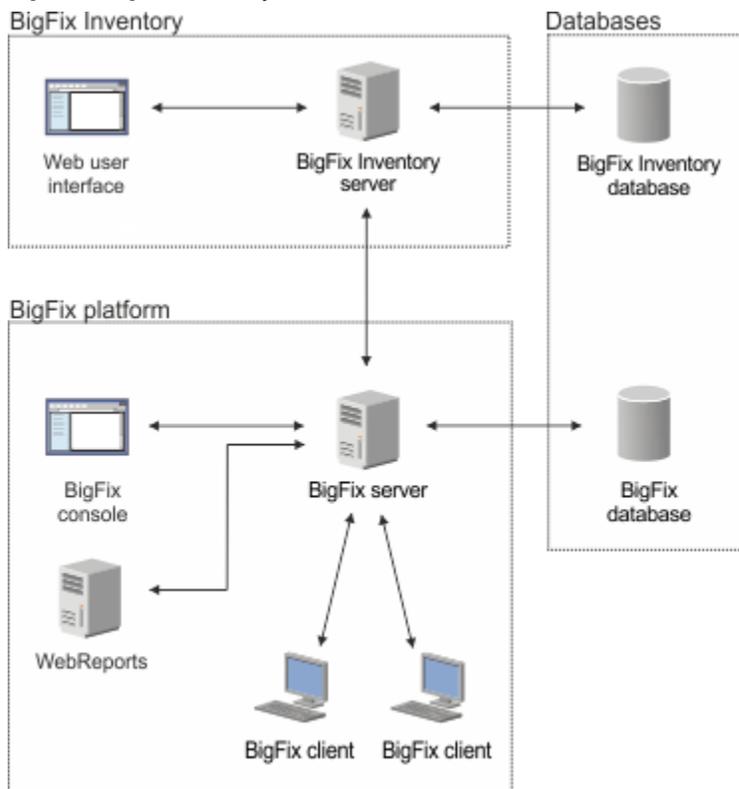
Missing scan results might occur after changing the assignments of your computers. Such a change is not always detected by the server, which as a result does not upload the scan results. To correct this issue, you must reinstall the scanner and rerun the scans. After you do this, the old scan results are deleted, and the new ones are collected and uploaded. To reinstall the scanner, complete the following steps:

1. In the BigFix console, stop all actions that are running on the problematic endpoints.
2. Open the Fixlet site.
3. From **Fixlets and Tasks**, select **Uninstall Scanner**, and click **Take Action**.
4. Select all problematic computers, and click **OK**. Wait until the action is complete.
5. Reinstall the scanner and rerun the scans. For more information, see: [Setting up scans \(on page cxci\)](#).

Installing on Windows

A complete deployment of BigFix Inventory on Windows requires the installation of three components: the BigFix platform, the BigFix Inventory server, and an MS SQL database. Additionally, a BigFix client must be installed on every computer from which you want to collect software inventory data. You can install all components from scratch or add the BigFix Inventory server to the existing BigFix infrastructure. Depending on the environment size, you can install all components on a single computer, or distribute them among multiple computers.

Figure 4. BigFix Inventory infrastructure



Step-by-step installation guide for Windows

Use the step-by-step installation guide to ensure that you complete all steps necessary to successfully install BigFix Inventory on Windows.

Table 22. Step-by-step guide for installing and configuring BigFix Inventory

Stage	Installation step
1.	<p>Plan the installation to ensure that the computer on which you want to install BigFix Inventory fulfills the following requirements:</p> <ul style="list-style-type: none"> • Runs on a supported operating system (on page c) • Fulfills the minimum hardware requirements (on page cvii) • Has sufficient disk space (on page cvii) • Has the required software (on page ci) installed
2.	<p>Install the BigFix platform</p> <ul style="list-style-type: none"> • Create the license authorization file (on page cxxxii) • Download the BigFix installer (on page cxxxiii) • Install the BigFix server, console and the first client (on page cxxxiii)

Table 22. Step-by-step guide for installing and configuring BigFix Inventory (continued)

Stage	Installation step
	<ul style="list-style-type: none"> • Install the BigFix client (on page cxxxvi) on every computer from which you want to collect software inventory data
	<p> Tip: You can install the clients at this point or proceed with the installation of BigFix Inventory and install the clients later on.</p>
3	<p>Install BigFix Inventory</p> <ul style="list-style-type: none"> • Install MS SQL Server • Enable the BigFix Inventory fixlet site (on page cxxxviii) • Install the BigFix Inventory server (on page cxli) • Perform initial configuration (on page cxlvii) • Verify the current application and catalog version (on page clxxxii)
4	<p>Create users and groups to give users access to BigFix Inventory and to define computers for which they can see data.</p> <ul style="list-style-type: none"> • Set up users (on page clxxxix) • Set up computer groups (on page cxc)
5	<p>Set up scans to discover software and hardware inventory in your infrastructure.</p> <p> Important: The steps are not required if you enabled the default scan configuration during the installation.</p> <ul style="list-style-type: none"> • Optional: Distribute scans for improved performance (on page cc) • Activate the required analyses (on page cci) • Install the scanner • Initiate software scans (on page cciii) • Schedule uploads of software scan results (on page ccviii) • Collect capacity data (on page ccx)

Installing the BigFix platform on Windows

The BigFix platform is the core of the BigFix Inventory infrastructure. It consists of a server that coordinates the flow of data to and from the monitored computers, an administrative console, and clients that collect data from the computers in your infrastructure. Before you can install BigFix Inventory, install all components of the BigFix platform.

REFERENCE

[BigFix Inventory infrastructure \(on page xcix\)](#)

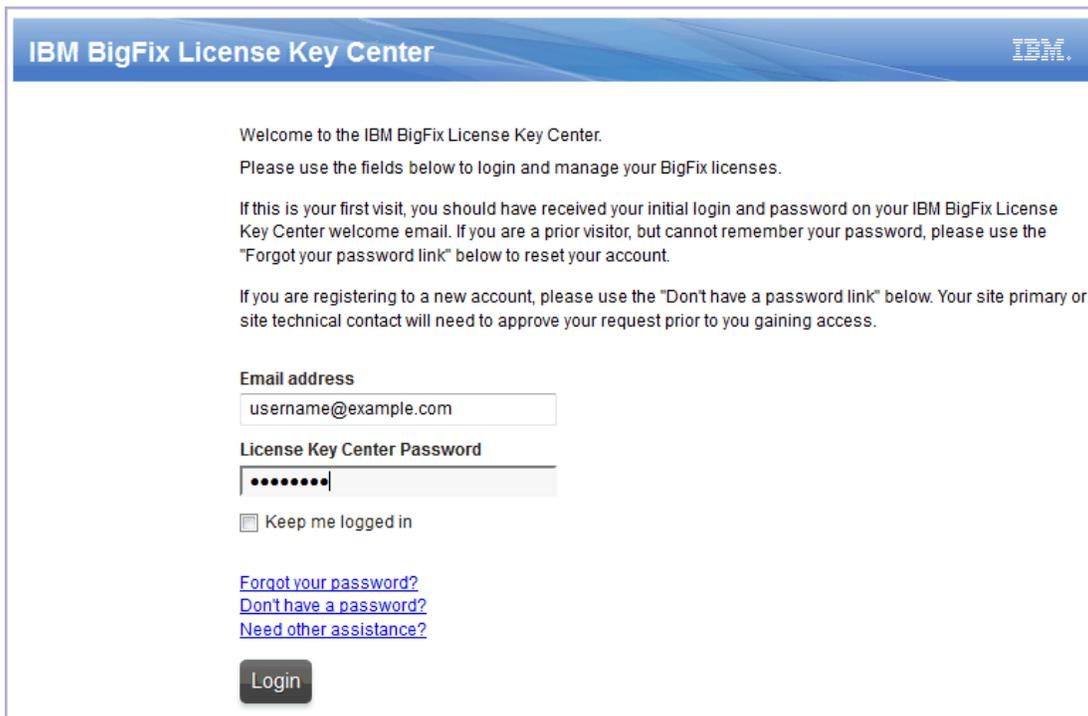
[Software requirements \(on page ci\)](#)

[Hardware requirements for the server on Linux \(on page cxiii\)](#)

Creating the license authorization file

After you order BigFix Inventory from Passport Advantage, you receive a welcome email with information how to access the BigFix License Key Center. Go to the Key Center and create the license authorization file that contains deployment and licensing information. The file is needed during the installation of the BigFix platform.

1. Go to the [License Key Center](#).
2. Enter your email address and the password that you received in the welcome email.



The screenshot shows the IBM BigFix License Key Center login page. The page has a blue header with the text "IBM BigFix License Key Center" and the IBM logo. The main content area is white and contains the following text:

Welcome to the IBM BigFix License Key Center.
Please use the fields below to login and manage your BigFix licenses.

If this is your first visit, you should have received your initial login and password on your IBM BigFix License Key Center welcome email. If you are a prior visitor, but cannot remember your password, please use the "Forgot your password link" below to reset your account.

If you are registering to a new account, please use the "Don't have a password link" below. Your site primary or site technical contact will need to approve your request prior to you gaining access.

Email address
username@example.com

License Key Center Password
••••••••

Keep me logged in

[Forgot your password?](#)
[Don't have a password?](#)
[Need other assistance?](#)

Login

3. For each product that you purchased, specify the allocated client quantity. If you leave 0, you cannot install the particular product.
4. To generate the license authorization file, click **Create Serial Number**.
5. To download the file, click **Download**.

The `LicenseAuthorization.BESLicenseAuthorization` file is downloaded to your computer. The file is needed during the installation of the BigFix platform.

[Download the BigFix installer. \(on page cxxxiii\)](#)

REFERENCE

[Passport Advantage](#)

Downloading the BigFix installer for Windows

To obtain the BigFix installer for Windows, go to the BigFix download center and download the Installation Generator.

1. Go to <http://support.bigfix.com/bes/release/>, and choose the latest version of the BigFix platform.
2. Download the Installation Generator.

Installation Generator

Operating System	Download
Windows	Download

3. To save the installer, click **Save File**.

Move the installer to the computer on which you want to install the BigFix platform and [start the installation \(on page cxxxiii\)](#).

Installing BigFix on Windows

To install the BigFix platform on Windows, run the installation wizard that will guide you through the installation of the BigFix server, console, and the first BigFix client.

The BigFix platform installed on Windows requires MS SQL Server as a database. Ensure that you have MS SQL Server installed and running before you start the installation of BigFix. For information about supported versions of MS SQL Server, see: [Software requirements \(on page ci\)](#).



Restriction:

The BigFix installer allows for the installation of an evaluation version of MS SQL Server. However, such version expires after 180 days and should not be used in a production environment.

The following procedure shows the most common installation scenario for Windows. If it does not fit your environment, you need more information about each step, or other installation tasks, see: [Installing on Windows systems](#) in the BigFix documentation.

1. Request a license certificate.
 - a. To start the installation, run the `BigFix-BES-version.exe` file.
 - b. When prompted, choose the **Production** installation.

- c. After you read and accept the license agreement, select **I want to install with an BigFix license authorization file**.
- d. Specify the location of the [license authorization file \(on page cxxxii\)](#), and click **Next**.
- e. Enter the DNS name or IP address of the computer on which you want to install the BigFix platform, and click **Next**.
- f. Create a public/private key pair that will be used to authorize all users of the BigFix console. Enter the password and choose the key size, then click **Create**.

Request License

Before finishing your license request, you must create a public/private key pair, which will be used to create and authorize all users of the IBM BigFix Console.

Please provide the key size that you would like to use, as well as a password that will be used to encrypt the private key file.

It is important that the key file is stored in a secured location so that it is not compromised or lost.

Password:

Verify password:

Key size:

< Back Create Cancel

- g. Save the created files in a secure folder, and click **OK**.



Attention: If you lose the private key file, a new license certificate needs to be created, which requires a new installation.

- h. Submit the request to BigFix to obtain the license certificate.
 - If the computer can access the Internet, select the first option, and click **Request**. The request is submitted automatically.
 - If the computer cannot access the Internet, select the second option. The `request.BESLicenseRequest` file is generated, and saved to a chosen folder. Copy the file to a computer with Internet access. Go to [BES License Request Handler](#), and submit the request file. In return, the `license.crt` file is saved to your computer. Copy the file back to the computer where you are installing BigFix. Return to the installation, and click **Import** to import the certificate. Then click **Create**.

2. Enter the parameters of the masthead file that contains configuration and license information. Then, click **OK**.

- i** **Tip:** Default settings are suitable for most deployments. For detailed information about the parameters, see: [Requesting a license certificate and creating the masthead](#).

Advanced Masthead Parameters

The default values for these parameters should be suitable for most IBM BigFix deployments. For further information about the implications of these parameters, please contact a IBM BigFix support technician.

Server Port Number:

Gathering Interval:

Initial Action Lock: minutes

Action Lock Controller:

Exempt the following site URL from action locking:

Require use of FIPS 140-2 compliant cryptography.

Allow use of Unicode filenames in archives.

3. Choose the destination folder for the installers of the BigFix components, and click **Next**. After the component installers are deployed, click **Finish**.

The Installation Guide is launched to lead you through the installation of BigFix components: server, console, and client.

4. Install the BigFix server.
 - a. In the Installation Guide, click **Install Server**, and click **Install or Upgrade the Server on this computer**.
 - b. Follow instructions in the installation wizard. When prompted, specify the location of the `license.pvk` file, and enter the password that you specified in step 1.f (on page cxxxiv). Click **OK** to continue.
 - c. Create an account that you want to use to log in to the BigFix console, and click **OK**.
 - d. To complete the installation, click **Finish**.
5. Install the BigFix console.
 - a. In the Installation Guide, click **Install Console**, and click **Install or Upgrade the Console on this computer**.
 - b. Follow instructions in the installation wizard.

- c. To complete the installation, click **Finish**.

 **Tip:** You can log in to the console with the user that you created in step [4.c \(on page cxxxv\)](#).

6. Install the BigFix client on the computer on which you installed the BigFix server.
 - a. In the Installation Guide, click **Install Clients**, and click **Install or Upgrade the Client on this computer**.
 - b. Follow instructions in the installation wizard.
 - c. To complete the installation, click **Finish**.
7. If you want to install the BigFix Inventory server on a different computer, install the BigFix client on that computer.

You can proceed with the [installation of the BigFix client \(on page cxxxvi\)](#) on every computer from which you want to collect software inventory data or with the [installation of the BigFix Inventory server \(on page cxxxvii\)](#). You can perform these tasks in any order.

Installing the BigFix clients

Install the BigFix client on every computer in your network that you want to monitor, including the backup and recovery machines.

Installation methods

Methods for installing the clients vary depending on the operating system. Even if you install the BigFix server on Linux™, you might need to install some of the clients on Windows™ if your network consists of such computers. For more information, see:

- [Installing clients on Windows™ computers](#)
- [Installing clients on Linux™ and UNIX™ computers](#)

If you are not sure which installation method to choose, install the client manually.

Additional installation requirements

- If you are using HP Integrity VM, install the client on the virtual machines and their host operating systems.
- If you are using Solaris Containers/Zones or Logical Domains (LDOM), install the client in the global zone on the control domain and in other global zones. For more information, see: [Client Installation on Oracle Solaris \(on page ccxii\)](#).

Installation packages

The BigFix installation image available on the Passport Advantage® contains client installers. You can also download client installation packages from the [BigFix® support web page](#).

! **Important:** Not all of the operating systems for which installation packages are provided as part of the BigFix installation image are supported by BigFix Inventory. For more information, see: [System requirements \(on page c\)](#).

Client vs scanner

The BigFix client is common for all BigFix products that are based on the BigFix platform, and is used to perform various tasks on the endpoints. It provides information about a set of computer properties like its host name or the IP address. It is also used to install the scanner on the computers, and enables scanner management through fixlets.

The scanner is an independent component that is used by BigFix Inventory. It collects information about the hardware, as well as the software that is installed on the computers in your infrastructure. The data is sent to the BigFix server. You can then import it to the BigFix Inventory server.

Apart from installing the BigFix client, install also the scanner on every computer that you want to monitor to ensure that all necessary data is collected. For more information, see: [Setting up scans to discover software and hardware inventory \(on page cxci\)](#).

Reference

[Mac OS X - client installation instructions](#)

Installing BigFix Inventory on Windows

Before you start the installation of BigFix Inventory on Windows, ensure that an MS SQL Server is installed in your infrastructure. Then, enable a dedicated fixlet site in the BigFix console and download the BigFix Inventory installer. Next, run the installation in interactive or silent mode.

REFERENCE

[BigFix Inventory infrastructure \(on page xcix\)](#)

[Software requirements \(on page ci\)](#)

[Hardware requirements for the server on Linux \(on page cxiii\)](#)

Installing MS SQL Server for BigFix Inventory on Windows

BigFix Inventory that is installed on Windows requires MS SQL Server as a database. However, the software is not delivered as part of BigFix Inventory. Ensure that you install MS SQL Server before you proceed with the installation of the BigFix Inventory server.

The database can be installed on the same computer as the BigFix Inventory server or on a separate one. You can also reuse the MS SQL Server that you installed for the BigFix platform.

For detailed installation instructions, see: [Installing SQL Server](#).

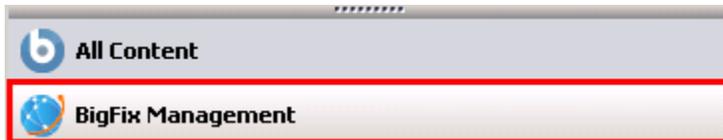
Enabling the fixlet site

A fixlet site is a collection of fixlets, tasks, and analyses that are related to a particular BigFix application. To get access to the content that is specific to BigFix Inventory, enable the application fixlet site. The procedure differs depending on whether the computer where the BigFix server is installed has the Internet access or not.

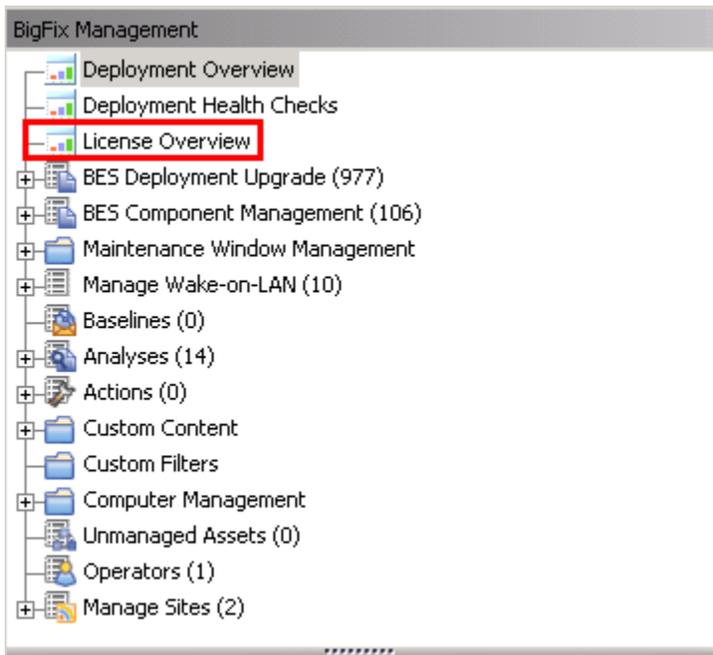
Enabling the fixlet site on Windows with Internet access

If the BigFix server can access the Internet, enable the BigFix Inventory site from the BigFix console. Then, subscribe computers that you want to monitor to this site so that its content becomes applicable on these computers.

1. Log in to the BigFix console.
2. In the bottom-left corner of the console, click **BigFix Management**.



3. In the left navigation panel, click **License Overview**.



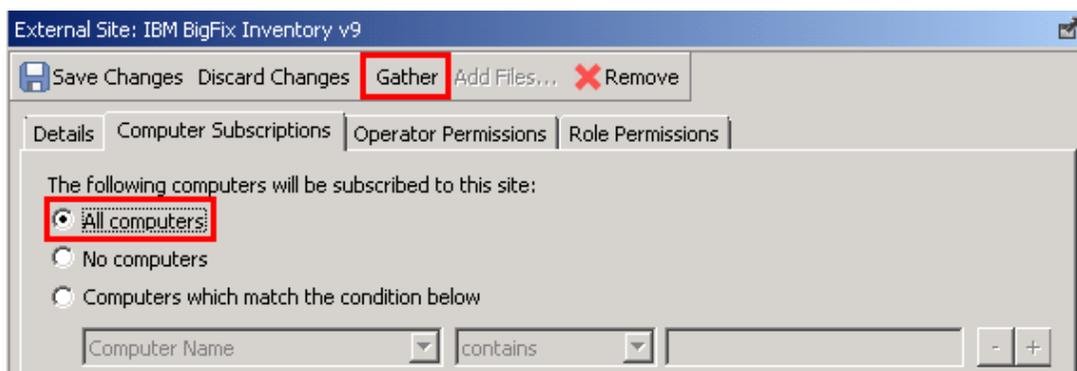
4. In the pane on the right, locate the entry called **Inventory**, and accept the license agreement.
5. From the list of available sites, enable the **IBM BigFix Inventory v9** site.
The content of the site is downloaded to the BigFix server.

 **Tip:** If there is no site called **Inventory**, select the site called **BigFix for Software Use Analysis v9**.

6. Subscribe all computers that you want to monitor to the fixlet site so that its content becomes applicable to these computers.

- a. In the bottom-left corner of the console, click **All Content**.
- b. In the left navigation panel, expand **Sites > External Sites**, and open the **IBM BigFix Inventory v9** site.
- c. In the pane on the right, open the **Computer Subscriptions** tab, and select **All Computers**.

 **Tip:** If the option is not active, click **Gather** to download the content of the fixlet site.



- d. Click **Save Changes**.

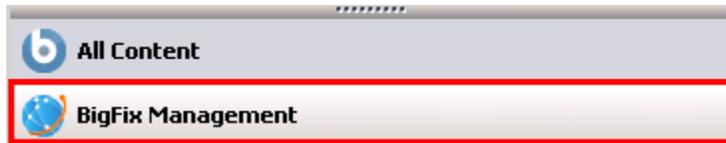
You enabled the BigFix Inventory site and subscribed computers that you want to monitor to this site.

Download the BigFix Inventory installer to a selected computer and [start the installation \(on page cxli\)](#).

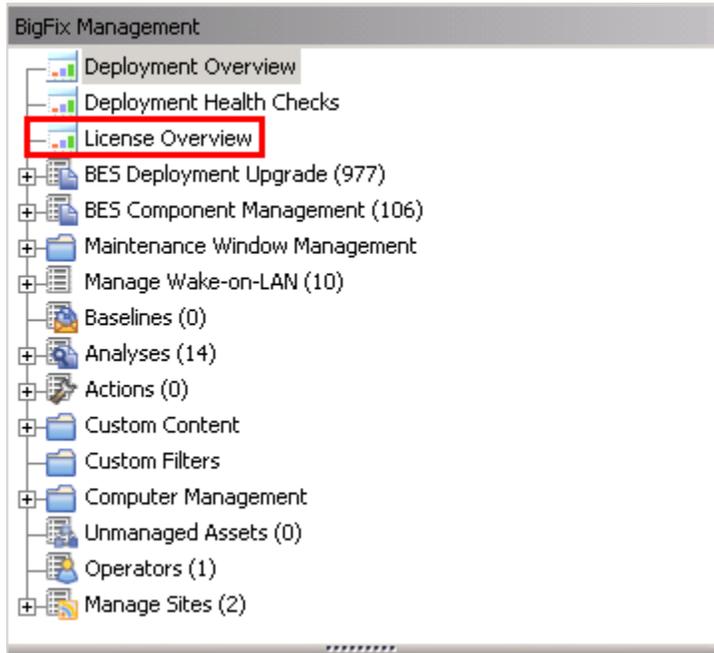
Enabling the fixlet site on Windows without Internet access

If the BigFix server cannot access the Internet, use the Airgap tool to enable the fixlet site. When the site is enabled and the content loaded, use the BES Download Cacher to download and cache the files on the BigFix server.

1. Create a request file to obtain the list of fixlet sites to which you are entitled.
 - a. Go to the BigFix server installation directory, by default `C:\Program Files (x86)\BigFix Enterprise\BES Server`, and run the `BESAirgapTool.exe` file. When prompted, save the files to a dedicated folder, for example `Airgap`.
 - b. Copy the created files to a computer with Internet access, and run the `BESAirgapTool.exe` file on that computer. This action exchanges the request file for a response file.
 - c. Copy the `AirgapResponse` file back to the computer where the BigFix server is installed, and place it in the `Airgap` folder. Run `BESAirgapTool.exe` to load the response to the BigFix server.
2. Enable the BigFix Inventory site.
 - a. Log in to the BigFix console.
 - b. In the bottom-left corner of the console, click **BigFix Management**.



- c. In the left navigation panel, click **License Overview**.



- d. In the pane on the right, locate the entry called **Inventory**, and accept the license agreement.
- e. From the list of available sites, enable the **IBM BigFix Inventory v9** site.

i **Tip:** If there is no site called **Inventory**, select the site called **BigFix for Software Use Analysis v9**.

3. To create a request file to load the content into the site, repeat step 1 (on page cxxxix).
4. Subscribe all computers that you want to monitor to the BigFix Inventory fixlet site so that its content becomes applicable to these computers.
 - a. Log in to the BigFix console.
 - b. In the left navigation panel, click **Sites > External Sites > IBM BigFix Inventory v9**.
 - c. In the pane on the right, open the **Computer Subscriptions** tab, and select **All Computers**. Click **Save Changes**.
The **BigFix Inventory.efxm** file is created on the server.
5. Cache the content of the fixlet site on the BigFix server.

- a. On the computer where the BigFix server is installed, go to: `install_dir\BES Server\wwwrootbes\bfsites`. Copy the `BigFix Inventory.efxm` file to a Windows computer with Internet access, and place it in the `C:\BigFix` directory.
- b. In the `C:\BigFix` directory create a folder called `downloads`.
- c. Run the BES Download Cacher with the following command:

```
BESDownloadCacher.exe -m "C:\BigFix\BigFix Inventory.efxm"
-x C:\BigFix\downloads
```

The BES Download Cacher downloads approximately 1 GB of required files.

- d. **Optional:** The default cache size is enough if you use only the **IBM BigFix Inventory v9** fixlet site. However, if you plan to run fixlets from other sites, such as **BES Support**, increase the cache size so that the BigFix server does not try to delete any files.
 - i. Log in to the BigFix console.
 - ii. In the left navigation bar, click **Computers** and right-click the computer on which the BigFix server is installed. Then, click **Edit Computer Settings**.
 - iii. Increase the value of the `_BESGather_Download_CacheLimitMB` setting. If the setting is not on the list, add it and specify the value in MB.



Tip: The size depends on each fixlet site, however you might need to increase it to at least a couple of gigabytes.

- e. Copy the contents of the `downloads` folder from the computer with Internet access to the following directory on the computer where the BigFix server is installed:

```
install_dir\BES Server\wwwrootbes\bfmirror\downloads\shal
```

You enabled the BigFix Inventory fixlet site and loaded its content.

Download the BigFix Inventory installer to a selected computer and [start the installation \(on page cxli\)](#).

Installing the server on Windows

After you enable the fixlet site, download the BigFix Inventory installer, and proceed with the installation. You can install the server in interactive or silent mode.

Installing the server on Windows in interactive mode

To install the BigFix Inventory server on Windows, download the application installer to a selected computer and run the `setup-server-windows-x86_64.bat` script. Then, follow instructions in the installation wizard.

- Ensure that the BigFix client is installed on the computer on which you want to install the BigFix Inventory server.
- Log in as user with administrative privileges and run the installer with the **Run as Administrator** option.
- Disable or reduce the User Account Control in Windows to avoid warnings during the installation.

1. Download the BigFix Inventory installer.

- a. Log in to the BigFix console.
- b. In the navigation panel, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
- c. In the upper right pane, select **Download BigFix Inventory**, and click **Take Action**.
- d. Select the computer to which you want to download the installer, and click **OK**.

Wait until the status of the action is Completed. The installer is downloaded to the **BES Installers \BFI_installer** directory on the selected computer. By default, the directory is in the following location: **C:\Program Files (x86)\BigFix Enterprise\BES Installers \BFI_installer**.

 **Tip:** If the computer is not available, ensure that it has the BigFix client installed and is subscribed to the BigFix Inventory fixlet site.

2. Log in to the computer where you downloaded the installer as a user with administrative privileges.
3. Go to the directory where the installer was downloaded, and extract the installation files.
4. To start the installation, run the **setup-server-windows-x86_64.bat** file.

Run the installer with the **Run as Administrator** option.

 **Tip:** If you encounter a slow installation, check the following items:

- Check the performance and speed of the hard disk.
- Check whether the antivirus is scanning each file separately as it might slow down packaging. If this is the case, turn off the antivirus. If the antivirus cannot be turned off, expect the installation to be slower.

5. Select the language of the installation, and click **OK**. The installation wizard starts, and the welcome panel opens. Click **Next**.
If you cannot choose your language in the installation wizard, set the system locale to a chosen language. For more information, see [Troubleshooting \(on page dcclxxiii\)](#).
6. Read and accept the license agreement, and click **Next**.
7. Follow instructions in the installation wizard and provide the required parameters.
8. On the last panel, specify the user account that you want to be used for running the application service.

- If you choose **Local System account**, the service runs under the `NT AUTHORITY\SYSTEM` user, which has all required rights.
- If you choose **Other account**, specify a user with the **log on as a service** right and administrative privileges. The User Account Control in Windows must be reduced or disabled, because it might block the service from starting. The user must be entered as `domain\username`, `machine\username`, or `.\username` if it is a local account.

Specify a service account

Local System account
 Other account:

User name:

Password:

Help

Choose the user account for running the application service. The account should be chosen depending on the authentication mode you are using to access the databases:

SQL Server auth. - choose any account.

! **Important:** To use Windows authentication to access the database, the service owner chosen on this panel must be available both to BigFix Inventory and the relevant database server for which you use this authentication. For local databases, you can use the Local System account. For remote databases, specify a domain user that is shared between the two servers.

9. When the installation is complete, click **Done** to exit the wizard.

i **Tip:** If you encountered problems during the installation, analyze the log file that is in the following directory: `%USERPROFILE%/BFI9.2.16`.

The BigFix Inventory server is installed and the initial configuration opens in the browser. If you do not have a browser installed or want to complete the configuration from a different computer, go to: `https://host_name:port`, where `host_name` and `port` are the values that you specified during the installation.

[Complete the initial configuration \(on page cxlvii\).](#)

TASK

[Resuming a stopped or failed interactive installation \(on page clxxviii\)](#)

REFERENCE

[Installation and upgrade problems \(on page dcclxxi\)](#)

[Server installation and upgrade logs \(on page dccxcviii\)](#)

[Server installation and uninstallation return codes \(on page dccc\)](#)

Installing the server on Windows in silent mode

To install the BigFix Inventory server on Windows, download the application installer to a selected computer, edit parameters in the `install_response.txt` file, and run the installation command.

- Ensure that the BigFix client is installed on the computer on which you want to install the BigFix Inventory server.
- Log in as user with administrative privileges and run the installer with the **Run as Administrator** option.
- Disable or reduce the User Account Control in Windows to avoid warnings during the installation.

1. Download the BigFix Inventory installer.

- Log in to the BigFix console.
- In the navigation panel, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
- In the upper right pane, select **Download BigFix Inventory**, and click **Take Action**.
- Select the computer to which you want to download the installer, and click **OK**.

Wait until the status of the action is Completed. The installer is downloaded to the `BES Installers \BFI_installer` directory on the selected computer. By default, the directory is in the following location: `C:\Program Files (x86)\BigFix Enterprise\BES Installers \BFI_installer`.



Tip: If the computer is not available, ensure that it has the BigFix client installed and is subscribed to the BigFix Inventory fixlet site.

- Log in to the computer where you downloaded the installer as a user with administrative privileges.
- Go to the directory where the installer was downloaded, and extract the installation files.
- Read the license agreement in the `BES Installers\BFI_installer\licenses\LA_language.txt` file.
- Edit the `install_response.txt` response file and adjust it to your installation. Ensure that the **RSP_LICENSE_ACCEPTED** parameter is set to true. If you do not accept the license, the installation fails. For more information about installation parameters, see: [Server installation response file \(on page cxlv\)](#).
- Start the command line and change to the directory with installation files. To start the installation, run the following command.

```
setup-server-windows-x86_64.bat -f response_file_path -i silent
```

Where *response_file_path* is the absolute path to the response file that you are using. For example:

```
setup-server-windows-x86_64.bat -f "C:\Program Files (x86)\BigFix Enterprise\BES
Installers\BFI_installer\install_response.txt" -i silent
```

i **Tip:** Use the `-h` option to view help information about using the script, for example: `setup-server-windows-x86_64.bat -h`.

i **Tip:** If you encounter a slow installation, check the following items:

- Check the performance and speed of the hard disk.
- Check whether the antivirus is scanning each file separately as it might slow down packaging. If this is the case, turn off the antivirus. If the antivirus cannot be turned off, expect the installation to be slower.

Access the BigFix Inventory user interface to [complete the initial configuration \(on page cxlvii\)](#). To access the user interface, go to: `https://host_name:port`, where *host_name* and *port* are the values that you specified during the installation.

TASK

[Resuming a stopped or failed interactive installation \(on page clxxviii\)](#)

REFERENCE

[Installation and upgrade problems \(on page dcclxxi\)](#)

[Server installation and upgrade logs \(on page dccxcviii\)](#)

[Server installation and uninstallation return codes \(on page dccc\)](#)

Server installation response file

The `install_response` file specifies input parameters that are used during the installation of the BigFix Inventory server in silent mode.

Table 23. Response file parameters

Parameter key name	Description	Default
RSP_LICENSE_ACCEPTED	Accepts the license agreement. The installation fails if you do not change the value of the parameter to <code>true</code> .	<code>false</code>
RSP_DISABLE_PREREQ_WARNINGS	Disables prerequisite checking warnings.	<code>false</code> (warnings are enabled)

Table 23. Response file parameters (continued)

Parameter key name	Description	Default
	If the server does not have enough memory or processor cores, the silent installation fails. You can change this behavior by setting the value of this parameter to <code>true</code> . The installation can complete even if the requirements are not fulfilled, however insufficient resources might impact performance.	
RSP_TLM_ROOT	Specifies the installation location. Specify an empty directory in which you want to install the server. If the directory does not exist, it is created.	Windows <code>C:\Program Files\ibm\BFI</code> Linux <code>/opt/ibm/BFI</code>
RSP_TLM_HTTPS_PORT	Specifies the port that is used by the server. If you do not specify the port number, a default value is used. If the selected port is already used by a different application, the installation fails.	9081
RSP_DISABLE_COMMUNICATION_WARNINGS	Disables communication warnings. If any of the ports that you specified in the RSP_TLM_HTTPS_PORT is locked by another application, silent installation fails. To specify a port that is temporarily used but will be available later, set the RSP_DISABLE_COMMUNICATION_WARNINGS parameter to <code>true</code> .	false
Windows RSP_USER_ACCOUNT	Specifies the user account for running the application service. The User Account Control in Windows must be reduced or disabled. Otherwise, it might block the service from starting. If you leave the value <code>current</code> , the service runs under the <code>NT AUTHORITY\SYSTEM</code> user, which has all required rights. If you specify a different user, ensure that the following requirements are met:	current

Table 23. Response file parameters (continued)

Parameter key name	Description	Default
	<ul style="list-style-type: none"> The user has the log on as a service right and administrative privileges. The user is entered as <code>domain\username</code>, <code>machine\username</code>, or <code>.\username</code> if it is a local account <p> Important: To use Windows authentication to access the database, the service owner chosen here must be available both to BigFix Inventory and the relevant database server for which you use this authentication. For local databases, you can use <code>current</code>, but for remote ones it must be a domain user that is shared between the two servers.</p>	
Windows RSP_USER_ACCOUNT_PWD	Specifies the password of the user account for running the application service.	

Performing initial configuration on Windows

During the initial configuration, you create a BigFix Inventory database and the application administrator. You also set up a connection to the BigFix server and database. Optionally, you can configure a connection to the Web Reports database to give the Web Reports users access to BigFix Inventory.

Permissions and roles

- Ensure that the MS SQL Server user has the following permissions. These permissions apply only if the databases were installed with default settings and all customizations and hardening configurations were consulted with BigFix support.
 - For the BigFix database (BFEnterprise): `CREATE FUNCTION, CREATE SCHEMA, CREATE TABLE, CREATE VIEW, EXECUTE, SELECT`
 - For the Web Reports database (BESReporting): `SELECT`
- Ensure that the MS SQL Server user has the appropriate role to create the BigFix Inventory database.

- If you create a new database during initial configuration in BigFix Inventory, the user must have the `sysadmin` role in MS SQL Server.
- If you create an MS SQL database manually before the configuration, make sure that the database is empty and use the `SQL_Latin1_General_CP1_CS_AS` collation. The user for such database must have the `db_owner` role in MS SQL Server.
- If you are using MS SQL Server 2012 and you chose Local System account as the service owner during the installation, give the `dbcreator` or `sysadmin` role to the `NT AUTHORITY\SYSTEM` user in MS SQL Server.

Other considerations

- If you want to use Windows authentication to connect to the databases, the following requirements must be fulfilled:
 - BigFix and BigFix Inventory must be installed on Windows.
 - To use a local user for authentication, BigFix and BigFix Inventory must be installed on the same instance of Windows. The applications can run on different copies if the user is a domain user.
 - The owner of the BigFix Inventory service must also be able to access this database.

1. Create the BigFix Inventory database.

a. Enter the host name of the database server.

If you want to configure a named database instance or specify a non-default port, provide the host name in the following format:

- `hostname\instance_name`, for example `localhost\MyInstance`
- `hostname:port_number`, for example `localhost:1444`

b. Enter the name of the application database.

c. Select the authentication mode.

- Select **Windows Authentication**, to authenticate with a Windows user that you specified as the service owner during the installation of BigFix Inventory.



Restriction: If the MS SQL Server is installed on the same computer as BigFix Inventory, enter the database host name without its domain name (FQDN) or use `localhost` instead. The host name can be specified as `NC1985110` or `localhost`, but not as `NC1985110.domain.com` or `198.50.100`.

- Select **SQL Server Authentication** to authenticate with an MS SQL Server user. This authentication mode must be enabled in MS SQL Server. For more information, see: [Enabling the SQL Server Authentication mode](#).

d. To create the database instance, click **Create**.

2. Create the administrator of BigFix Inventory.

- i** **Tip:** Avoid using admin, administrator, root or a similar name for the administrative account. Such an account might be prone to hacker attacks and locked out if an attacker exceeds the specified number of failed login attempts. For more information, about the account lockout, see: [Configuring user account lockout \(on page dccxliv\)](#).



User Name*
john.doe

Password*
●●●●●●

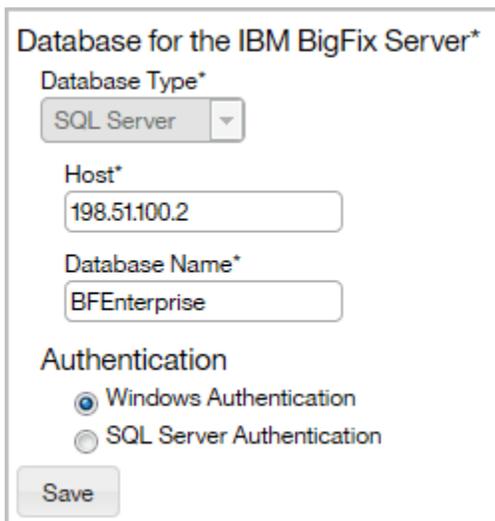
Password Confirmation*
●●●●●●

Create

3. **Optional:** To automatically enable scans that collect data from the computers in your infrastructure, select **Enable default scan schedule for this data source**.

If you enable the default scan schedule, actions that are needed to collect data from the computers in your infrastructure are automatically started on the BigFix server. This option is advised for environments with up to a few thousand computers. For larger environments, finish the installation, divide the computers into groups, and then manually set up scan schedule for each group to avoid performance issues. For more information about the default and manual scan schedule, see: [Setting up scans to discover software and hardware inventory \(on page cxciij\)](#).

4. Configure the connection to the BigFix database. The database stores information about the computers, and data that was discovered on these computers. Specify the host, port, database name, and credentials of the user that can access the BigFix database.



Database for the IBM BigFix Server*

Database Type*
SQL Server

Host*
198.51.100.2

Database Name*
BFEnterprise

Authentication

Windows Authentication

SQL Server Authentication

Save

5. Configure the connection to the BigFix server. The host name or IP address, and the API port number are automatically retrieved from the database. Specify only the administrative user that you created during the installation of BigFix.

IBM BigFix Server*

Authentication (Console Operator)

User Name*

Password*

Disable automatic address lookup



Note: If you do not want to provide the Master Operator, you can create a dedicated BigFix user that fulfills the following requirements:

- Is assigned the BigFix Inventory v9 site
- Is assigned computers that you are going to monitor, and the computer where the BigFix server is installed
- Has the following permissions: *Can use REST API, Can use Console, Custom Content, Can Create Actions*

The option is supported starting from BigFix 9.5.

6. **Optional:** If the BigFix and BigFix Inventory servers are in separated networks, the automatic address lookup might return incorrect address. To disable the lookup, select **Disable automatic address lookup**, and specify the address manually. Then, configure additional environment variables on the BigFix Inventory server. For more information, see: [Configuring servers in separate networks \(on page clxxx\)](#).
7. **Optional:** Configure the connection to the Web Reports database. Specify the database type, host name, database name, and credentials of the Web Reports database user.

98/

Web Reports Database

Database Type*

Host*

Database Name*

Authentication

Windows Authentication

SQL Server Authentication

8. To create connections to the databases, click **Create**.

When the connections are created and configured, a new page opens and a message about the data import is displayed.

9. **Optional:** If your environment consists of more than 50 000 endpoints, complete steps to [enhance the application performance \(on page dcccxliv\)](#) before you run the import.

10. To run the initial import, click **Import Now**.

The import might take a few hours, depending on your hardware capacity.

If you enabled the [default scan schedule \(on page cxcvi\)](#), the collected data might not be displayed in BigFix Inventory after the initial import. Some time is required to finish scans that were initiated during the installation, and to upload scan results to the server. If the reports in BigFix Inventory do not contain any data, wait about an hour until the scans are completed. Then, start another import.

If you did not enable the default scan schedule, [manually configure scans \(on page cxci\)](#) to collect data that is later on displayed on the reports.

REFERENCE

[Database creation logs and return codes \(on page dccxcvii\)](#)

Uninstalling on Windows

To remove BigFix Inventory from your infrastructure, stop the application-specific actions and analyses that are running on the computers and uninstall the scanner. Then, remove the VM Manager Tool. Finally, uninstall the BigFix Inventory server. You can also remove the related database.

Uninstalling the server on Windows in interactive mode

To uninstall the BigFix Inventory server on Windows, run the `uninstall.bat` file. Then, follow instructions in the installation wizard. The wizard does not uninstall the MS SQL Server nor the BigFix server. These components need to be removed separately.

1. Log in to the computer where the BigFix Inventory server is installed as the same user who performed the installation.
2. Go to the `C:\Program Files\IBM\BFI\Uninstall` directory and run the `uninstall.bat` file. Start the uninstallation with the **Run as administrator** option.
3. Follow instructions in the uninstallation wizard. When the uninstallation finishes, click **Done**.

The BigFix Inventory server is uninstalled but the database, user logins, and passwords are preserved. To remove them, [delete the BigFix Inventory database](#). You can also remove the [BigFix server](#).

REFERENCE

[Removing the server manually \(on page dcccxi\)](#)

[Server installation and uninstallation return codes \(on page dccc\)](#)

Uninstalling the server on Windows in silent mode

To uninstall the BigFix Inventory server on Windows, edit parameters in the `uninstall_response.txt` file, and run the uninstallation command. The command does not uninstall the MS SQL Server nor the BigFix server. These components need to be removed separately.

1. Log in to the computer where the BigFix Inventory server is installed as the same user who performed the installation.
2. Go to the directory with the uninstallation response file, open the `uninstall_response.txt` file, and edit the uninstallation parameters. By default, the directory is `C:\Program Files\IBM\BFI\Uninstall`.
3. Start the command line and run the following command.

```
uninstall.bat -f C:\Program Files\IBM\BFI\Uninstall\uninstall_response.txt -i silent
```

The BigFix Inventory server is uninstalled but the database, user logins, and passwords are preserved. To remove them, [delete the BigFix Inventory database](#). You can also remove the [BigFix server](#).

REFERENCE

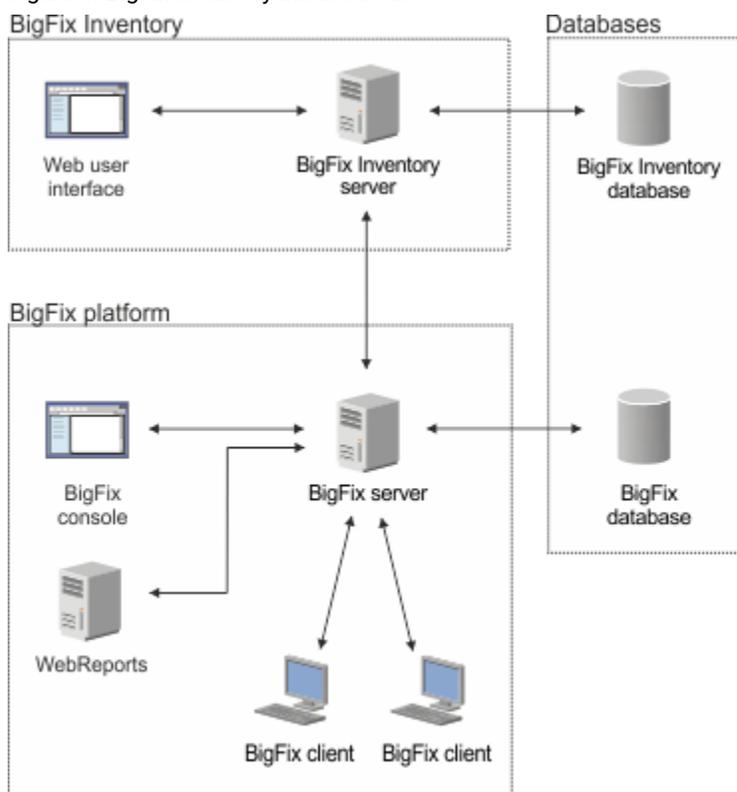
[Removing the server manually \(on page dcccxi\)](#)

[Server installation and uninstallation return codes \(on page dccc\)](#)

Installing on Linux

A complete deployment of BigFix Inventory on Linux requires the installation of three components: the BigFix platform, the BigFix Inventory server, and DB2 database. Additionally, a BigFix client must be installed on every computer from which you want to collect software inventory data. Depending on the environment size, you can install all components on a single computer, or distribute them among multiple computers.

Figure 5. BigFix Inventory infrastructure



Step-by-step installation guide for Linux

Use the step-by-step installation guide to ensure that you complete all steps necessary to successfully install BigFix Inventory on Linux.

Table 24. Step-by-step guide for installing and configuring BigFix Inventory

Stage	Installation step
1.	<p>Plan the installation to ensure that the computer on which you want to install BigFix Inventory fulfills the following requirements:</p> <ul style="list-style-type: none"> • Runs on a supported operating system (on page c) • Fulfills the minimum hardware requirements (on page cvii) • Has sufficient disk space (on page cxiv) • Has the required software (on page ci) installed
2.	<p>Install the BigFix platform</p> <ul style="list-style-type: none"> • Create the license authorization file (on page clv) • Download the BigFix installer (on page clvi) • Install the BigFix server, console and the first client (on page clvi)

Table 24. Step-by-step guide for installing and configuring BigFix Inventory (continued)

Stage	Installation step
	<ul style="list-style-type: none"> • Install the BigFix client (on page clxi) on every computer from which you want to collect software inventory data
	<p> Tip: You can install the clients at this point or proceed with the installation of BigFix Inventory and install the clients later on.</p>
3	Install BigFix Inventory
	<ul style="list-style-type: none"> • Install DB2 (on page clxii) • Enable the BigFix Inventory fixlet site (on page clxiii) • Install the BigFix Inventory server (on page clxviii) • Perform initial configuration (on page clxxiii) • Verify the current application and catalog version (on page clxxxii)
4	Create users and groups to give users access to BigFix Inventory and to define computers for which they can see data.
	<ul style="list-style-type: none"> • Set up users (on page clxxxix) • Set up computer groups (on page cxc)
5	Set up scans to discover software and hardware inventory in your infrastructure.
	<p> Important: The steps are not required if you enabled the default scan configuration during the installation.</p>
	<ul style="list-style-type: none"> • Optional: Distribute scans for improved performance (on page cc) • Activate the required analyses (on page cci) • Install the scanner • Initiate software scans (on page cciii) • Schedule uploads of software scan results (on page ccviii) • Collect capacity data (on page ccx)

Installing the BigFix platform on Linux

The BigFix platform is the core of the BigFix Inventory infrastructure. It consists of a server that coordinates the flow of data to and from the monitored computers, an administrative console, and clients that collect data from the computers in your infrastructure. Before you can install BigFix Inventory, install all components of the BigFix platform.

 **Important:** Even if you set up entire environment on Linux, you need a Windows computer on which you can install the BigFix console.

REFERENCE

[BigFix Inventory infrastructure \(on page xcix\)](#)

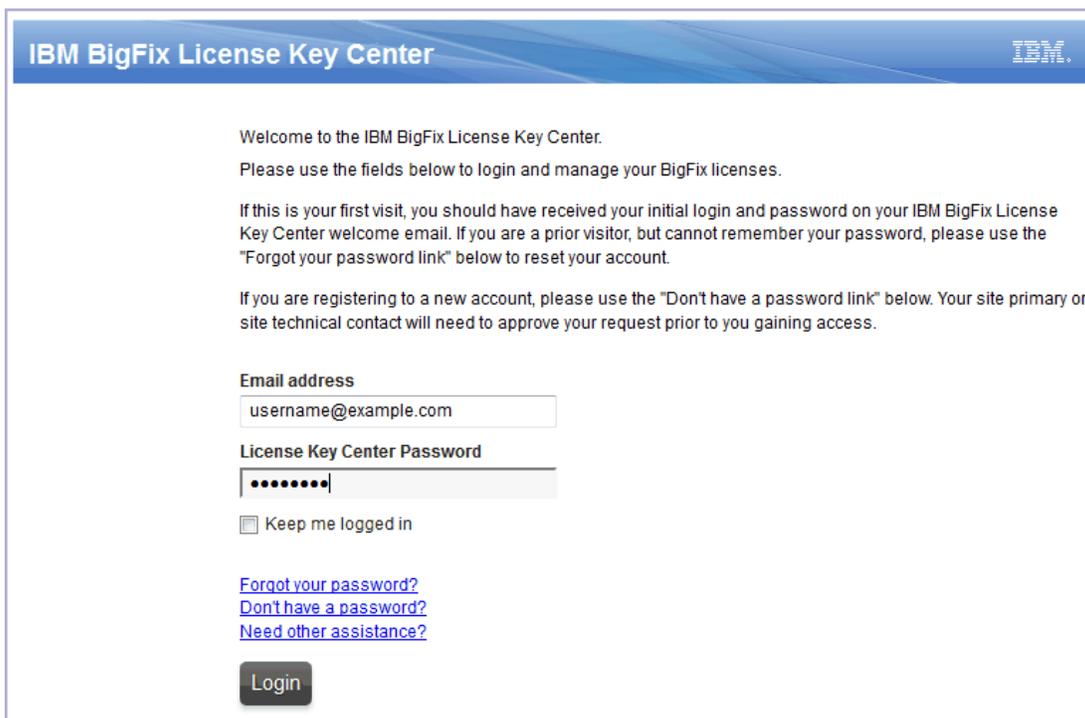
[Software requirements \(on page ci\)](#)

[Hardware requirements for the server on Linux \(on page cxiii\)](#)

Creating the license authorization file

After you order BigFix Inventory from Passport Advantage, you receive a welcome email with information how to access the BigFix License Key Center. Go to the Key Center and create the license authorization file that contains deployment and licensing information. The file is needed during the installation of the BigFix platform.

1. Go to the [License Key Center](#).
2. Enter your email address and the password that you received in the welcome email.



The screenshot shows the IBM BigFix License Key Center login page. The page has a blue header with the text "IBM BigFix License Key Center" and the IBM logo. The main content area is white and contains the following text:

Welcome to the IBM BigFix License Key Center.
Please use the fields below to login and manage your BigFix licenses.

If this is your first visit, you should have received your initial login and password on your IBM BigFix License Key Center welcome email. If you are a prior visitor, but cannot remember your password, please use the "Forgot your password link" below to reset your account.

If you are registering to a new account, please use the "Don't have a password link" below. Your site primary or site technical contact will need to approve your request prior to you gaining access.

Email address
username@example.com

License Key Center Password
••••••••

Keep me logged in

[Forgot your password?](#)
[Don't have a password?](#)
[Need other assistance?](#)

Login

3. For each product that you purchased, specify the allocated client quantity. If you leave 0, you cannot install the particular product.
4. To generate the license authorization file, click **Create Serial Number**.
5. To download the file, click **Download**.

The `LicenseAuthorization.BESLicenseAuthorization` file is downloaded to your computer. The file is needed during the installation of the BigFix platform.

Download the BigFix installer. For more information, see: [Downloading the BigFix installer for Linux \(on page clvi\)](#).

TASK

[Passport Advantage](#)

Downloading the BigFix installer for Linux

To obtain the BigFix installer for Linux, log in to Passport Advantage and download the required packages. If you are planning to install the BigFix platform and DB2 on a single computer, download the package that contains installers for both applications. If you are planning to install the components on separate computers, download the installers separately.

1. Log in to [Passport Advantage](#).
2. On the **Navigation help** tab, click **Software download & media access**.
3. In the lower part of the page, click **BigFix Inventory**.
4. In the window that opens, select your preferred operating system and language for the installation packages.
5. In the lower part of the page, select **Required**, and click **Download**.
6. In the window that opens, select the download location, the preferred download method, select **I agree to the terms and conditions**, and click **Download Now**.
7. On the page that opens, click the displayed links to download the BigFix packages.

Image	Description
BigFix_Pltfm_Install_V95_Lnx_DB2.tgz	BigFix Platform Installer V9.5 for Linux and DB2 Multilingual
DB2_10.5.0.5_limited_Lnx_x86-64.tar.gz	IBM DB2 Server V10.5.0.5 for Linux on AMD64 and Intel EM64T systems (x64) Multilingual

For information about part numbers of the installation packages, see: [BigFix 9.5.0 - Download Information](#).

Move the installer to the computer on which you want to install the BigFix platform and [start the installation \(on page clvi\)](#).

Installing BigFix on Linux

To install the BigFix platform on Linux, run the `install.sh` script and go through the installation of the BigFix server, DB2 database, and the first BigFix client. Then, install the BigFix console on a Windows computer.

The BigFix platform installed on Linux requires DB2 as a database. You can use an existing instance of DB2 or install it together with BigFix.

The following procedure shows the most common installation scenario for Linux. If it does not fit your environment, you need more information about each step, or other installation tasks, see: [Installing on Linux systems](#) in the BigFix documentation.

1. Go to the directory with the installation package, and extract the installation files.

```
tar xvf installation_package
```

2. From the extracted directory, go to `ServerInstaller_n.n.mmn.n-rhe6.x86_64`, and run the installation script.

```
./install.sh
```

3. To choose the production installation, enter 2.

```
Select the type of installation

[1] Evaluation: Request a free evaluation license from BigFix.
This license allows you to install a fully functional copy of the
BigFix on up to 1000 clients, for a period of 30 days.

[2] Production: Install using a production license or an authorization
lf for a production license.

Choose one of the options above or press Enter to accept the default value: [1]
```

4. To accept the license agreement, enter 1.

5. To install all BigFix components, enter 1.

```
Select the BigFix features you want to install:

[1] All Components (server, client, and WebReports)

[2] Server and client Only

[3] WebReports Only

Choose one of the options above or press <Enter> to accept the default value: [1]
```

6. To create a single or master database, enter 1.

```
Select the database replication:

[1] Single or master database

[2] Replicated database

Choose one of the options above or press <Enter> to accept the default value: [1]
```

7. To use a local instance of DB2, enter 1.

```
Select the database:

[1] Use a local database

[2] Use a remote database

Choose one of the options above or press <Enter> to accept the default value: [1]
```

8. Specify the installation location for BigFix.

```
Choose the server's root folder:

Specify the location for the server's root folder or
press <Enter> to accept the default value: /var/opt/BESServer
```

9. Specify the installation location for Web Reports.

```
Choose the WebReports server's root folder:
Specify the location for the WebReports server's root folder or
press <Enter> to accept the default value: /var/opt/BESWebReportsServer
```

10. Enter the port number for Web Reports.

```
Since version 9.5.2, freshly installed Web Reports
is automatically configured to use HTTPS.
Choose the WebReports server's port number:
Specify the port number or press <Enter> to accept the default value: 8083
```

11. The installer checks if DB2 is installed on the server. To install it automatically from the files that are bundled with BigFix, enter 1.

```
DB2 installation check
The installer does not detect if DB2 is installed on the system.
Specify which option corresponds to your installation:
[1] DB2 is not installed, install it.
[2] DB2 is installed, use the installed instance.
[3] Exit from the installation.
Choose one of the options above or press <Enter> to accept the default value: [1]
```

12. Default settings containing DB2 users and installation path are listed. To accept them and proceed with the installation, enter 1.

```
DB2 installation
DB2 will be installed using the following settings:
DB2 instance owner: db2inst1
DB2 fenced user: db2fenc1
DB2 administration server user: dasusr1
DB2 communication port: 50000
DB2 installation directory: /opt/ibm/db2/V10.5
If you want to use different value for these settings specify them
in the installation response file. Refer to the product documentation
for further details.
[1] Proceed to install DB2.
[2] Exit from the installation.
Choose one of the options above or press <Enter> to accept the default value: [1]
```

13. Specify a password for the DB2 administrative users.

```
DB2 administrative user password:
Specify the password for the DB2 administrative users:
```

14. Enter the user name and password for the initial administrative user of BigFix or press Enter to use the default `IBMAdmin`. This user is required to log in to the BigFix console.

15. To run the installation using a license authorization file, enter 1.

```
Choose the type of setup that best suits your needs:
[1] I want to install with a BES license authorization file.
[2] I want to install with a production license that I already have.
[3] I want to install with an existing masthead.
Choose one of the options above or press <Enter> to accept the default value: [1]
```

16. If you are not using a proxy, enter 2.

```
Proxy usage
[1] Use the proxy to access the internet
[2] Do not use the proxy
Choose one of the options above or press <Enter> to accept the default value: [2]
```

17. Specify the location of your license authorization file.

```
License authorization location
Enter the location of the license authorization file that
you received from BigFix or press <Enter> to accept the default value:
./license/LicenseAuthorization.BESLicenseAuthorization
```

18. Specify the DNS name or IP address of the server on which you are performing the installation, or press Enter to accept the default value.

19. Specify the site admin private key password.

20. Specify the key size to encrypt the credentials.

```
Key size level
Provide the key size that you want to use:
[1] 'Min' level (2048 bits)
[2] 'Max' level (4096 bits)
Choose one of the options above or press <Enter> to accept the default value: [2]
```

21. Specify the **License** folder in which the license file is to be generated and saved.

```
Choose the license folder:
Specify a folder for your private key (license.pvk), the license certificate (license.crt),
and the site masthead (masthead.afxm) or press <Enter> to accept the default value: ./license
```

22. Submit the request to BigFix to obtain the license certificate:

- If your server can access the internet, enter 1. The request will be submitted automatically.
- If your server cannot access the Internet, enter 2, and submit the request manually:
 - a. The `request.BESLicenseRequest` is generated and saved to a folder with your license files. Copy this request to a computer with internet access.
 - b. On the computer with internet access, go to [BES License Request Handler](#) and submit the request file.

- c. The `license.crt` file is saved to your computer. Copy it back to the BigFix server.
- d. Return to the installation and enter `1` to import the certificate and continue with the installation.

23. Specify the value of the deployment encoding (FXF encoding) that will be used for the content.

24. To accept the default masthead values, enter `1`.

```
Advanced masthead parameters
The masthead will be created using the following defaults:
Server port number: 52311
Use of FIPS 140-2 compliant cryptography: Disabled
Gather interval: One Day
Initial action lock: Unlocked
Action lock controller: Console
Action lock exemptions: Disabled
Unicode filenames in archives: Enabled
The above default values are suitable for most of BigFix deployments.
[1] Use default values
[2] Use custom values
Choose one of the options above or press <Enter> to accept the default value: [1]
```

25. To run Web Reports as a root, enter `1`.

```
Use root user for WebReports
If you specify true, WebReports service will run with root privileges.
[1] True
[2] False
Choose one of the options above or press <Enter> to accept the default value: [2]
```

26. When the installation is complete, install the BigFix console.

 **Important:** The BigFix console is supported only on Windows.

- a. Go to `/var/opt/BESInstallers`.
- b. Copy the `Console` folder to a Windows computer.
- c. On the Windows computer, run the `setup.exe` script, and follow instructions in the installation wizard.

27. If you want to install the BigFix Inventory server on a different computer, install the BigFix client on that computer.

You can proceed with the [installation of the BigFix client \(on page clxi\)](#) on every computer from which you want to collect software inventory data or with the [installation of the BigFix Inventory server \(on page clxviii\)](#). You can perform these tasks in any order.

Installing the BigFix clients

Install the BigFix client on every computer in your network that you want to monitor, including the backup and recovery machines.

Installation methods

Methods for installing the clients vary depending on the operating system. Even if you install the BigFix server on Linux™, you might need to install some of the clients on Windows™ if your network consists of such computers. For more information, see:

- [Installing clients on Windows™ computers](#)
- [Installing clients on Linux™ and UNIX™ computers](#)

If you are not sure which installation method to choose, install the client manually.

Additional installation requirements

- If you are using HP Integrity VM, install the client on the virtual machines and their host operating systems.
- If you are using Solaris Containers/Zones or Logical Domains (LDOM), install the client in the global zone on the control domain and in other global zones. For more information, see: [Client Installation on Oracle Solaris \(on page cxxii\)](#).

Installation packages

The BigFix installation image available on the Passport Advantage® contains client installers. You can also download client installation packages from the [BigFix® support web page](#).



Important: Not all of the operating systems for which installation packages are provided as part of the BigFix installation image are supported by BigFix Inventory. For more information, see: [System requirements \(on page c\)](#).

Client vs scanner

The BigFix client is common for all BigFix products that are based on the BigFix platform, and is used to perform various tasks on the endpoints. It provides information about a set of computer properties like its host name or the IP address. It is also used to install the scanner on the computers, and enables scanner management through fixlets.

The scanner is an independent component that is used by BigFix Inventory. It collects information about the hardware, as well as the software that is installed on the computers in your infrastructure. The data is sent to the BigFix server. You can then import it to the BigFix Inventory server.

Apart from installing the BigFix client, install also the scanner on every computer that you want to monitor to ensure that all necessary data is collected. For more information, see: [Setting up scans to discover software and hardware inventory \(on page cxci\)](#).

Installing BigFix Inventory on Linux

Before you start the installation of BigFix Inventory on Linux, ensure that DB2 is installed in your infrastructure. Then, enable a dedicated fixlet site in the BigFix console and download the BigFix Inventory installer. Next, run the installation in interactive or silent mode.

REFERENCE

[BigFix Inventory infrastructure \(on page xcix\)](#)

[Software requirements \(on page ci\)](#)

[Hardware requirements for the server on Linux \(on page cxiii\)](#)

Installing DB2 for BigFix Inventory on Linux

BigFix Inventory that is installed on Linux requires DB2 database. You can install the database on the same computer as the BigFix Inventory server or on a separate one. You can also reuse the instance of DB2 that you installed for the BigFix platform. DB2 installation package is delivered together with BigFix Inventory.

The following scenario shows how to install DB2 10.5 that is delivered with BigFix Inventory. If it does not fit your environment, you need more information about each step or other installation options, see the documentation for [DB2 10.1](#) or [DB2 10.5](#).

1. Go to the directory with the installation package, and extract the installation files.

```
tar xvf installation_package
```

2. From the extracted directory, run the installation script.

```
./db2setup
```

3. In the navigation bar of the **DB2 Setup Launchpad**, click **Install a Product**.
4. Scroll down, and click **Install New** under DB2 Workgroup Server Edition.
5. Read and accept the license agreement. Click **Next**.
6. Select the installation type, and click **Next**.
7. Select **Install DB2 Server Edition on this computer and save my settings in a response file**, and click **Next**.
8. Specify the installation directory, and click **Next**.
9. Create a user for the DB2 Administration Server, and click **Next**.
10. Select **Create a DB2 instance**, and click **Next**.
11. Select **Single partition instance**, and click **Next**.
12. Create the DB2 instance owner. This user will be needed when you specify connection details between BigFix and BigFix Inventory. Then, click **Next**.
13. Create the fenced user, and click **Next**.
14. Set up the DB2 server to send notifications or select the second check box to omit this step. Then, click **Next**.
15. Review the settings, and click **Finish**. The installation is started.
16. When the installation is complete, add a license file to extend your license.

- a. From the directory with extracted installation files, go to `/db2/license`.
- b. Copy the `db2wse_o.lic` file to the home directory of the DB2 instance owner, by default `/home/db2inst1`.
- c. Open the terminal and switch the user to the DB2 instance owner.

```
su - db2inst1
```

- d. Add the license file to your installation.

```
db2licm -a db2wse_o.lic
```

The installation of DB2 is complete. The default location is `/opt/ibm/db2`.

Enabling the fixlet site

A fixlet site is a collection of fixlets, tasks, and analyses that are related to a particular BigFix application. To get access to the content that is specific to BigFix Inventory, enable the application fixlet site. The procedure differs depending on whether the computer where the BigFix server is installed has the Internet access or not.

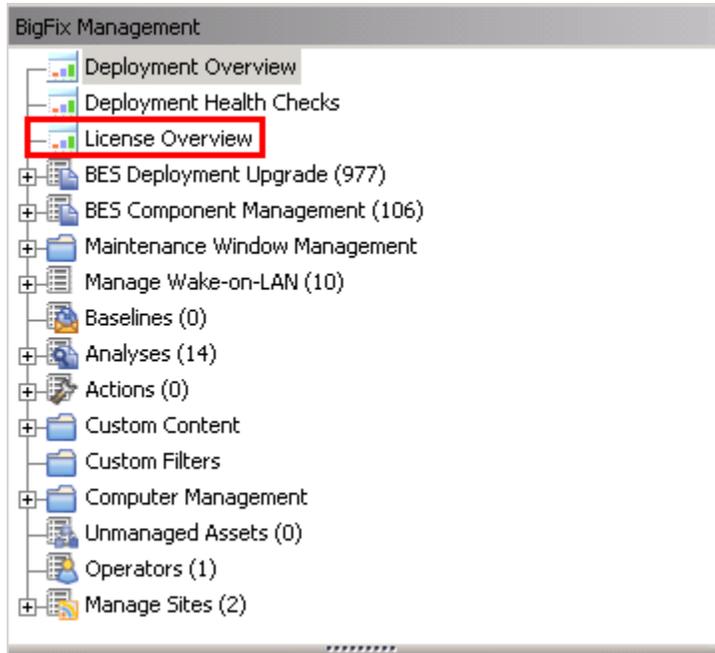
Enabling the fixlet site with Internet access

If the BigFix server can access the Internet, enable the BigFix Inventory site from the BigFix console. Then, subscribe computers that you want to monitor to this site so that its content becomes applicable on these computers.

1. Log in to the BigFix console.
2. In the bottom-left corner of the console, click **BigFix Management**.



3. In the left navigation panel, click **License Overview**.



4. In the pane on the right, locate the entry called **Inventory**, and accept the license agreement.
5. From the list of available sites, enable the **IBM BigFix Inventory v9** site.
The content of the site is downloaded to the BigFix server.

 **Tip:** If there is no site called **Inventory**, select the site called **BigFix for Software Use Analysis v9**.

6. Subscribe all computers that you want to monitor to the fixlet site so that its content becomes applicable to these computers.
 - a. In the bottom-left corner of the console, click **All Content**.
 - b. In the left navigation panel, expand **Sites > External Sites**, and open the **IBM BigFix Inventory v9** site.
 - c. In the pane on the right, open the **Computer Subscriptions** tab, and select **All Computers**.

 **Tip:** If the option is not active, click **Gather** to download the content of the fixlet site.



d. Click **Save Changes**.

You enabled the BigFix Inventory site and subscribed computers that you want to monitor to this site.

Download the BigFix Inventory installer to a selected computer and [start the installation \(on page clxviii\)](#).

Enabling the fixlet site without Internet access

If the BigFix server cannot access the Internet, use the Airgap tool to enable the fixlet site. When the site is enabled and the content loaded, use the BES Download Cacher to download and cache the files on the BigFix server.

To complete the following procedure, you need two computers:

- Linux computer on which the BigFix server is installed
- Windows computer with Internet access

1. Log in to a Windows computer with Internet access and download the Airgap tool.

! **Important:** Ensure that the version of the Airgap tool that you download corresponds with the version of your BigFix server.

2. Create a request file to obtain the list of fixlet sites to which you are entitled.

- Log in to the Linux computer where the BigFix server is installed.
- Go to the BigFix server installation directory, by default `/opt/BESServer/bin`, and run the following command to start the Airgap tool.

```
./Airgap.sh -remotedir /directory
```

Where *directory* is the directory in which the `airgap.tar` file is created.

c. Extract the `airgap.tar` file.

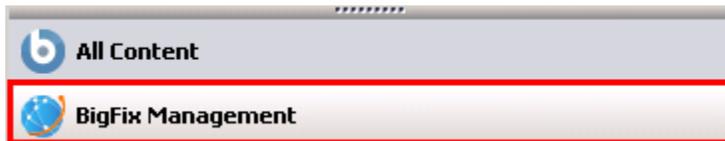
```
tar xvf airgap.tar
```

- d. Copy the extracted `AirgapRequest.xml` file to the Windows computer, and place it in the folder that contains the downloaded `BESAirgapTool.exe` file.
- e. On the Windows computer, run the `BESAirgapTool.exe` file. This action exchanges the request file for a response file.
- f. Copy the `AirgapResponse` file to the Linux computer, and place it in the `/opt/BESServer/bin` directory.
- g. Run the Airgap tool again on the Linux computer.

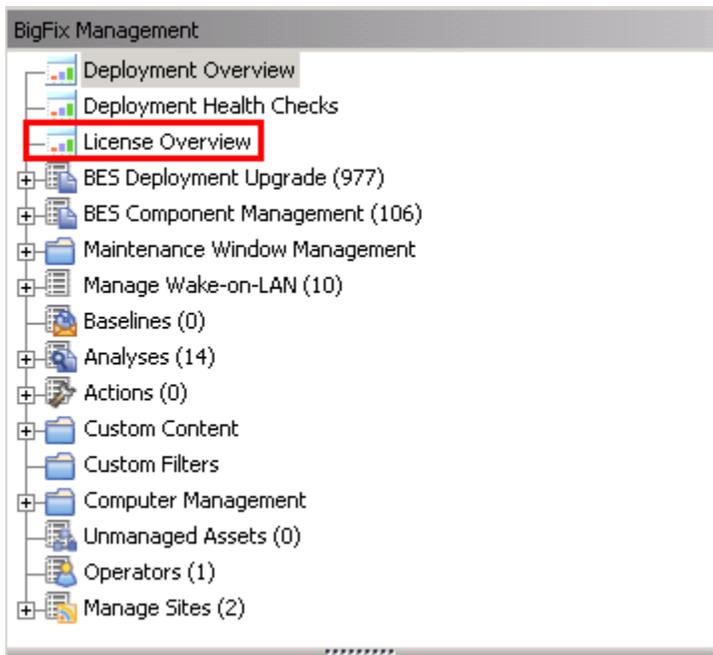
```
./Airgap.sh -run
```

3. Enable the BigFix Inventory site.

- a. Log in to the BigFix console.
- b. In the bottom-left corner of the console, click **BigFix Management**.



- c. In the left navigation panel, click **License Overview**.



- d. In the pane on the right, locate the entry called **Inventory**, and accept the license agreement.

- e. From the list of available sites, enable the **IBM BigFix Inventory v9** site.

 **Tip:** If there is no site called **Inventory**, select the site called **BigFix for Software Use Analysis v9**.

4. To create a request file to load the content into the site, repeat step 2 (on page clxv).
5. Subscribe all computers that you want to monitor to the BigFix Inventory fixlet site so that its content becomes applicable to these computers.

- a. Log in to the BigFix console.
- b. In the left navigation panel, click **Sites > External Sites > IBM BigFix Inventory v9**.
- c. In the pane on the right, open the **Computer Subscriptions** tab, and select **All Computers**. Click **Save Changes**.

The `BigFix Inventory.efxm` file is created on the server.

6. Cache the content of the fixlet site on the BigFix server.

- a. On the computer where the BigFix server is installed, go to the `install_dir/BESServer/wwwroot/bfsites` directory. Copy the `BigFix Inventory.efxm` file to a Windows computer with Internet access, and place it in the `C:\BigFix` directory.
- b. In the `C:\BigFix` directory create a folder called `downloads`.
- c. Run the BES Download Cacher with the following command:

```
BESDownloadCacher.exe -m "C:\BigFix\BigFix Inventory.efxm"
-x C:\BigFix\downloads
```

The BES Download Cacher downloads approximately 1 GB of required files.

- d. **Optional:** The default cache size is enough if you use only the **IBM BigFix Inventory v9** fixlet site. However, if you plan to run fixlets from other sites, such as **BES Support**, increase the cache size so that the BigFix server does not try to delete any files.
 - i. Log in to the BigFix console.
 - ii. In the left navigation panel, click **Computers** and right-click the computer on which the BigFix server is installed. Then, click **Edit Computer Settings**.
 - iii. Increase the value of the `_BESGather_Download_CacheLimitMB` setting. If the setting is not on the list, add it and specify the value in MB.

 **Tip:** The size depends on each fixlet site, however you might need to increase it to at least a couple of gigabytes.

- e. Copy the contents of the `downloads` folder from the computer with Internet access to the following directory on the computer where the BigFix server is installed:

```
install_dir/BESServer/wwwrootbes/bfmirror/downloads/shal
```

You enabled the BigFix Inventory fixlet site and loaded its content.

Download the BigFix Inventory installer to a selected computer and [start the installation \(on page clxviii\)](#).

Installing the server on Linux

After you enable the fixlet site, download the BigFix Inventory installer, and proceed with the installation. You can install the server in interactive or silent mode.

Installing the server on Linux in interactive mode

To install the BigFix Inventory server on Linux, download the application installer to a selected computer and run the `setup-server-linux-x86_64.sh` script. Then, follow instructions in the installation wizard.

- Ensure that the BigFix client is installed on the computer on which you want to install the BigFix Inventory server.
- Ensure that the computer meets the following requirements:
 - A graphical user interface is available
 - X server is configured
 - `DISPLAY` variable is set properly
- Start the installation as root. Otherwise, the server is not registered as a system service.
- The use of `sudo` is not supported.

1. Download the BigFix Inventory installer.

- a. Log in to the BigFix console.
- b. In the navigation bar, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
- c. In the upper right pane, select **Download BigFix Inventory**, and click **Take Action**.
- d. Select a computer to which you want to download the installer, and click **OK**.
The installer is downloaded to the `/user/BFI_installer` directory on the selected computer. By default: `/root/BFI_installer`.

2. Log in as root to the computer where you downloaded the installer.

3. Go to the directory where the installer was downloaded, and extract the installation files.

```
tar xvf BFI-server-version-timestamp-linux-x86_64.tar.gz
```

4. To start the installation, run the `setup-server-linux-x86_64.sh` script.



Tip: If the installation is slow, check the following items:

- Check the performance and speed of the hard disk.
- Check whether the antivirus is scanning each file separately as it might slow down packaging. If this is the case, turn off the antivirus. If the antivirus cannot be turned off, expect the installation to be slower.

5. Select the language of the installation, and click **OK**. The installation wizard starts, and the welcome panel opens. Click **Next**.

If you cannot choose your language in the installation wizard, set the system locale to a chosen language. For more information, see: [Troubleshooting \(on page dcclxxiii\)](#).

6. Read and accept the license agreement, and click **Next**.

7. Follow instructions in the wizard and provide the required installation parameters.

8. When the installation is complete, click **Done** to exit the wizard.



Tip: If you encounter problems during the installation, analyze the log file that is in the `$HOME \BF19.2.16` directory.

The BigFix Inventory server is installed and the initial configuration opens in the browser. If you do not have a browser installed or want to complete the configuration from a different computer, go to: `https://host_name:port`, where `host_name` and `port` are the values that you specified during the installation.

[Complete the initial configuration \(on page clxxiii\)](#).

TASK

[Resuming a stopped or failed interactive installation \(on page clxxviii\)](#)

REFERENCE

[Installation and upgrade problems \(on page dcclxxi\)](#)

[Server installation and upgrade logs \(on page dccxcviii\)](#)

[Server installation and uninstallation return codes \(on page dccc\)](#)

Installing the server on Linux in silent mode

To install the BigFix Inventory server on Linux, download the application installer to a selected computer, edit parameters in the `install_response.txt` file, and run the installation command.

- Ensure that the BigFix client is installed on the computer on which you want to install the BigFix Inventory server.
- Start the installation as root to register the server as a system service and to be able to upgrade it later with a fixlet. Otherwise you can start the installation as non-root user.
- The use of `sudo` is not supported.

1. Download the BigFix Inventory installer.

- Log in to the BigFix console.
- In the navigation bar, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
- In the upper right pane, select **Download BigFix Inventory**, and click **Take Action**.
- Select a computer to which you want to download the installer, and click **OK**.
The installer is downloaded to the `/user/BFI_installer` directory on the selected computer. By default: `/root/BFI_installer`.

2. Log in as root to the computer where you downloaded the installer.

3. Go to the directory where the installer was downloaded, and extract the installation files.

```
tar xvf installation_package
```

4. Read the license agreement in the `/user/BFI_installer/licenses/LA_language.txt` file.

5. Edit the `install_response.txt` response file and adjust it to your installation. Ensure that the **RSP_LICENSE_ACCEPTED** parameter is set to true. If you do not accept the license, the installation fails. For more information about installation parameters, see: [Server installation response file \(on page clxxi\)](#).

6. Change to the directory with installation files. To start the installation, run the following command.

```
setup-server-linux-x86_64.sh -f response_file_path -i silent
```

Where `response_file_path` is the absolute path to the response file that you are using. For example:

```
setup-server-linux-x86_64.sh -f /root/BFI_installer/install_response.txt -i silent
```

 **Tip:** Use the `-h` option to view help information about using the script, for example: `setup-server-linux-x86_64.sh -h`.

 **Tip:** If the installation is slow, check the following items:

- Check the performance and speed of the hard disk.
- Check whether the antivirus is scanning each file separately as it might slow down packaging. If this is the case, turn off the antivirus. If the antivirus cannot be turned off, expect the installation to be slower.

Access the BigFix Inventory user interface to [complete the initial configuration \(on page clxxiii\)](#). To access the user interface, go to: `https://host_name:port`, where `host_name` and `port` are the values that you specified during the installation.

REFERENCE

[Installation and upgrade problems \(on page dclxxi\)](#)

[Server installation and upgrade logs \(on page dccxcviii\)](#)

[Server installation and uninstallation return codes \(on page dccc\)](#)

Server installation response file

The `install_response` file specifies input parameters that are used during the installation of the BigFix Inventory server in silent mode.

Table 25. Response file parameters

Parameter key name	Description	Default
RSP_LICENSE_ACCEPTED	Accepts the license agreement. The installation fails if you do not change the value of the parameter to <code>true</code> .	<code>false</code>
RSP_DISABLE_PREREQ_WARNINGS	Disables prerequisite checking warnings. If the server does not have enough memory or processor cores, the silent installation fails. You can change this behavior by setting the value of this parameter to <code>true</code> . The installation can complete even if the requirements are not fulfilled, however insufficient resources might impact performance.	<code>false</code> (warnings are enabled)
RSP_TLM_ROOT	Specifies the installation location. Specify an empty directory in which you want to install the server. If the directory does not exist, it is created.	Windows <code>C:\Program Files\ibm\BFI</code> Linux <code>/opt/ibm/BFI</code>
RSP_TLM_HTTPS_PORT	Specifies the port that is used by the server. If you do not specify the port number, a default value is used. If the selected port is already used by a different application, the installation fails.	9081

Table 25. Response file parameters (continued)

Parameter key name	Description	Default
RSP_DISABLE_COMMUNICATION_WARNINGS	<p>Disables communication warnings.</p> <p>If any of the ports that you specified in the RSP_TLM_HTTPS_PORT is locked by another application, silent installation fails. To specify a port that is temporarily used but will be available later, set the RSP_DISABLE_COMMUNICATION_WARNINGS parameter to <code>true</code>.</p>	<code>false</code>
Windows RSP_USER_ACCOUNT	<p>Specifies the user account for running the application service.</p> <p>The User Account Control in Windows must be reduced or disabled. Otherwise, it might block the service from starting.</p> <p>If you leave the value <code>current</code>, the service runs under the <code>NT AUTHORITY\SYSTEM</code> user, which has all required rights. If you specify a different user, ensure that the following requirements are met:</p> <ul style="list-style-type: none"> • The user has the log on as a service right and administrative privileges. • The user is entered as <code>domain\username</code>, <code>machine\username</code>, or <code>.\username</code> if it is a local account <p> Important: To use Windows authentication to access the database, the service owner chosen here must be available both to BigFix Inventory and the relevant database server for which you use this authentication. For local databases, you can use <code>current</code>, but for remote ones it must be a domain user that is shared between the two servers.</p>	<code>current</code>
Windows RSP_USER_ACCOUNT_PWD	Specifies the password of the user account for running the application service.	

Performing initial configuration on Linux

During the initial configuration, you create a BigFix Inventory database and the application administrator. You also set up a connection to the BigFix server and database. Optionally, you can configure a connection to the Web Reports database to give the Web Reports users access to BigFix Inventory.

- Ensure that the DB2 user has the following permissions. These permissions apply only if the databases were installed with default settings, and all customizations and hardening configurations were consulted with BigFix support.
 - For the BigFix database (BFENT): **DBAUTH**
 - For the Web Reports database (BESREPOR): **DATAACCESS**
- To reuse DB2 that you installed for BigFix with BigFix Inventory, increase the number of active databases to at least 3. To increase the number, log in as the DB2 instance owner, and run the following command. Then, restart the database server.

```
db2 update dbm cfg using NUMDB number_of_active_databases
```

1. Create the BigFix Inventory database.

- If the DB2 server is installed on the same computer as the BigFix Inventory server, perform the following steps.
 - a. Select **The application and the database are on the same computer**. The host name, port number, and the default name of the database are filled in automatically. If you want to specify a different name of the database, you can choose any name that meets the [DB2 naming requirements](#).
 - b. Provide credentials of the operating systems user that will be used for connecting to the database. The user can be an instance owner.

 **Tip:** Create a dedicated user for connecting to DB2.

- If the DB2 server is installed on a different computer than the BigFix Inventory server, perform the following steps.
 - c. To create the database, click **Create**.
- If the DB2 server is installed on a different computer than the BigFix Inventory server, perform the following steps.
 - a. Clear **The application and the database are on the same computer** check box.
 - b. Provide the host name and port number of the computer where the DB2 server is installed, and specify the name for the application database. You can use the default name or choose any name that meets the [DB2 naming requirements](#).
 - c. Provide credentials of the operating systems user that will be used for connecting to the database. The user can be an instance owner.

 **Tip:** Create a dedicated user for connecting to DB2.

- d. To download a script that is used for creating the database, click **Download Script**.

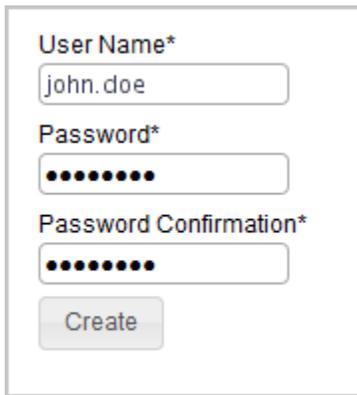
- e. Move the script to the computer where the DB2 server is installed and run it. The script creates a database that can be accessed by the user that you specified in substep c.

! **Important:** The user who runs the script must have the `SYSADM` authority. You can use the DB2 instance owner.

- f. When the database is created, return to the computer on which you are configuring BigFix Inventory, and click **Create**.

2. Create the administrator of BigFix Inventory.

i **Tip:** Avoid using `admin`, `administrator`, `root` or a similar name for the administrative account. Such an account might be prone to hacker attacks and locked out if an attacker exceeds the specified number of failed login attempts. For more information, about the account lockout, see: [Configuring user account lockout \(on page dccxliv\)](#).



The image shows a user creation form with the following fields and controls:

- User Name***: A text input field containing the value "john.doe".
- Password***: A password input field with masked characters (dots).
- Password Confirmation***: A second password input field with masked characters (dots).
- Create**: A button located below the password confirmation field.

3. **Optional:** To automatically enable scans that collect data from the computers in your infrastructure, select **Enable default scan schedule for this data source**.

If you enable the default scan schedule, actions that are needed to collect data from the computers in your infrastructure are automatically started on the BigFix server. This option is advised for environments with up to a few thousand computers. For larger environments, finish the installation, divide the computers into groups, and then manually set up scan schedule for each group to avoid performance issues. For more information about the default and manual scan schedule, see: [Setting up scans to discover software and hardware inventory \(on page cxciij\)](#).

4. Configure the connection to the BigFix database. The database stores information about the computers and data that was discovered on these computers. Specify the host, port, database name, and credentials of the user that can access the BigFix database.

Database for the IBM BigFix Server*

Database Type*
 DB2

Host*
 198.51.100.2

Port*
 50000

Database Name*
 BFENT

Authentication

User Name*
 db2inst1

Password*
 ●●●●●●●●

Create

5. Configure the connection to the BigFix server. The host name or IP address, and the API port number are automatically retrieved from the database. Specify only the administrative user that you created during the installation of BigFix.

IBM BigFix Server*

Authentication (Console Operator)

User Name*
 IEMAdmin

Password*
 ●●●●●●●●

Disable automatic address lookup

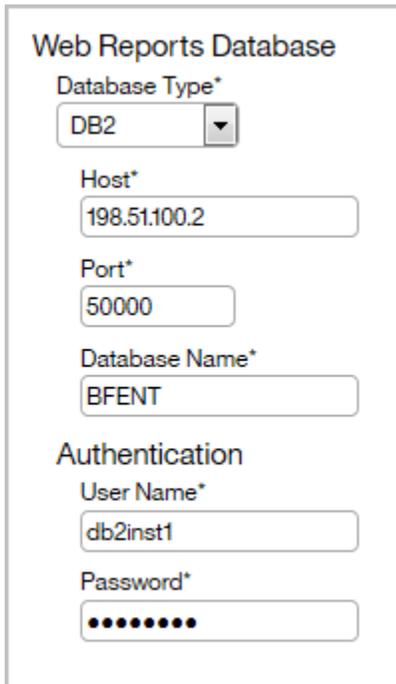


Note: If you do not want to provide the Master Operator, you can create a dedicated BigFix user that fulfills the following requirements:

- Is assigned the BigFix Inventory v9 site
- Is assigned computers that you are going to monitor, and the computer where the BigFix server is installed
- Has the following permissions: *Can use REST API, Can use Console, Custom Content, Can Create Actions*

The option is supported starting from BigFix 9.5.

- Optional:** If the BigFix and BigFix Inventory servers are in separated networks, the automatic address lookup might return incorrect address. To disable the lookup, select **Disable automatic address lookup**, and specify the address manually. Then, configure additional environment variables on the BigFix Inventory server. For more information, see: [Configuring servers in separate networks \(on page clxxx\)](#).
- Optional:** Configure the connection to the Web Reports database. Specify the database type, host name, database name, and credentials of the Web Reports database user.



The image shows a configuration form titled "Web Reports Database". It contains the following fields:

- Database Type***: A dropdown menu with "DB2" selected.
- Host***: A text input field containing "198.51.100.2".
- Port***: A text input field containing "50000".
- Database Name***: A text input field containing "BFENT".
- Authentication**: A section header.
- User Name***: A text input field containing "db2inst1".
- Password***: A text input field with 10 dots representing a masked password.

- To create connections to the databases, click **Create**.

When the connections are created and configured, a new page opens and a message about the data import is displayed.

- Optional:** If your environment consists of more than 50 000 endpoints, complete steps to [enhance the application performance \(on page dcccxliv\)](#) before you run the import.
- To run the initial import, click **Import Now**.

The import might take a few hours, depending on your hardware capacity.

If you enabled the [default scan schedule \(on page cxcvi\)](#), the collected data might not be displayed in BigFix Inventory after the initial import. Some time is required to finish scans that were initiated during the installation, and to upload scan results to the server. If the reports in BigFix Inventory do not contain any data, wait about an hour until the scans are completed. Then, start another import.

If you did not enable the default scan schedule, [manually configure scans \(on page cxciii\)](#) to collect data that is later on displayed on the reports.

REFERENCE

[Database creation logs and return codes \(on page dcccxcvii\)](#)

Uninstalling on Linux

To remove BigFix Inventory from your infrastructure, stop the application-specific actions and analyses that are running on the computers and uninstall the scanner. Then, remove the VM Manager Tool. Finally, uninstall the BigFix Inventory server. You can also remove the related database.

Uninstalling the server on Linux in interactive mode

To uninstall the BigFix Inventory server on Linux, run the `uninstall.sh` script file. Then, follow instructions in the installation wizard. The wizard does not uninstall DB2 nor the BigFix server. These components need to be removed separately.

1. Log in to the computer where the BigFix Inventory server is installed as root.
2. Go to the `/opt/ibm/BFI/Uninstall` directory, and run the `uninstall.sh` script.
3. Follow instructions in the uninstallation wizard. When the uninstallation finishes, click **Done**.

The BigFix Inventory server is uninstalled but the database, user logins, and passwords are preserved. To remove them, log in as the DB2 instance owner on the computer where the database is installed, and run the following command.

```
db2 deactivate db database_name
db2 drop db database_name
```

Where *database_name* is the name of the BigFix Inventory database, by default `TEMADB`. You can also remove the [BigFix server](#).

REFERENCE

[Removing the server manually \(on page dcccxxi\)](#)

[Server installation and uninstallation return codes \(on page dccc\)](#)

Uninstalling the server on Linux in silent mode

To uninstall the BigFix Inventory server on Linux, edit parameters in the `uninstall_response.txt` file, and run the uninstallation command. The command does not uninstall DB2 nor the BigFix server. These components need to be removed separately.

1. Log in to the computer where the BigFix Inventory server is installed as root.
2. Go to the `/opt/ibm/BFI/Uninstall` directory, open the `uninstall_response.txt` file, and edit the uninstallation parameters.

3. Start the command line and run the following command.

```
uninstall.sh -f /opt/ibm/BFI/Uninstall/uninstall_response.txt -i silent
```

The BigFix Inventory server is uninstalled but the database, user logins, and passwords are preserved. To remove them, log in as the DB2 instance owner on the computer where the database is installed, and run the following command.

```
db2 deactivate db database_name  
db2 drop db database_name
```

Where *database_name* is the name of the BigFix Inventory database, by default `TEMADB`. You can also remove the [BigFix server](#).

REFERENCE

[Removing the server manually \(on page dcccxxi\)](#)

[Server installation and uninstallation return codes \(on page dccc\)](#)

Installation-related tasks

If you encounter problems during the installation, you can check details related to the problem, fix it, and then resume the installation. If the BigFix and BigFix Inventory servers are in separated networks, you can disable the automatic address lookup to avoid problems with initial configuration. Learn how to check the current application and catalog version to make sure that your installation and initial configuration were successful.

Resuming a stopped or failed interactive installation

You can rerun the interactive installation during the pre-installation and installation stages. You can also use a built-in function to diagnose problems during the installation of the BigFix Inventory server.

There are different ways to rerun the installation and the solution depends on the phase during which the problem occurred. If an error occurs during the pre-installation phase, you can restart the installation and no additional actions are required. If you encounter an error during the installation phase, you have two options:

- You can remove the installation directory before you restart the installation
- You can use a built-in function to diagnose and fix the problem

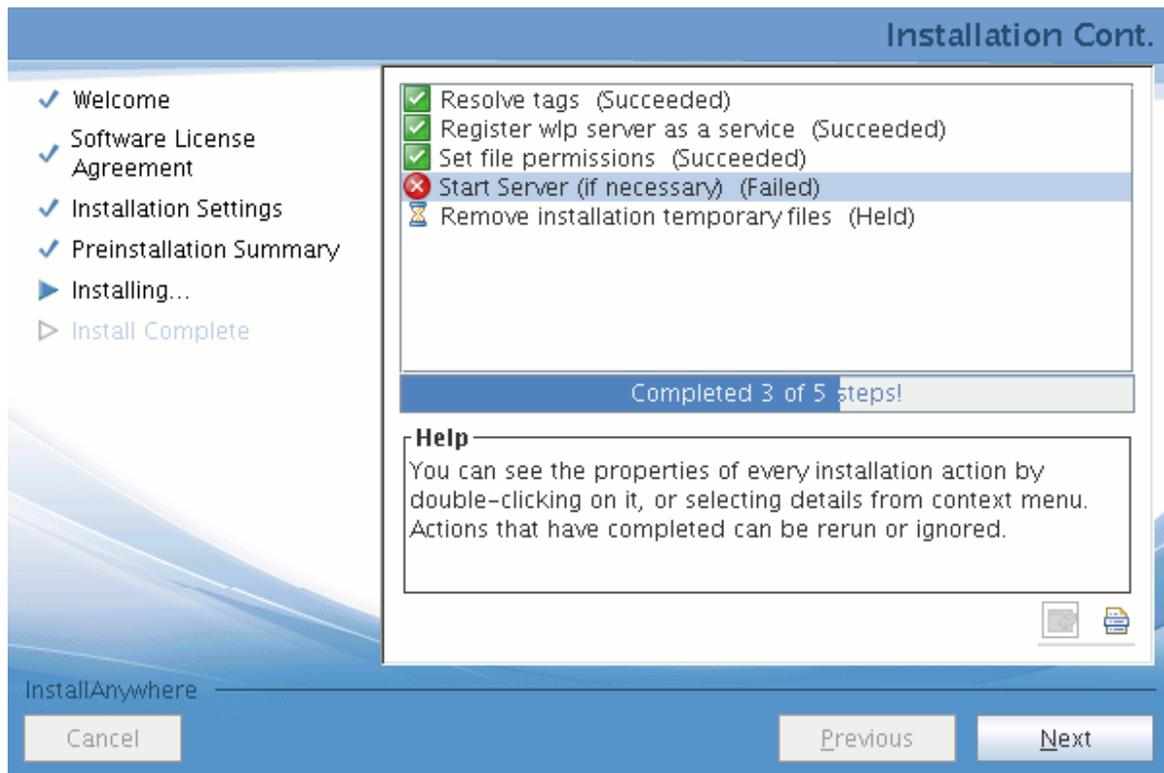
Some configuration steps in the installation depend on other steps. If one of them fails, the execution of the dependent step is also held. If an error occurs, the installation wizard continues running steps that do not depend on the failed one. You can see the list of prerequisites for any step in the step properties dialog. To open the dialog, double-click the step, or right-click it and select **Details**.

You can continue the installation and fix the problem at the end of the installation. You can also end the installation and resolve the problem later, at a convenient time. It is not necessary to specify any special options. It is enough

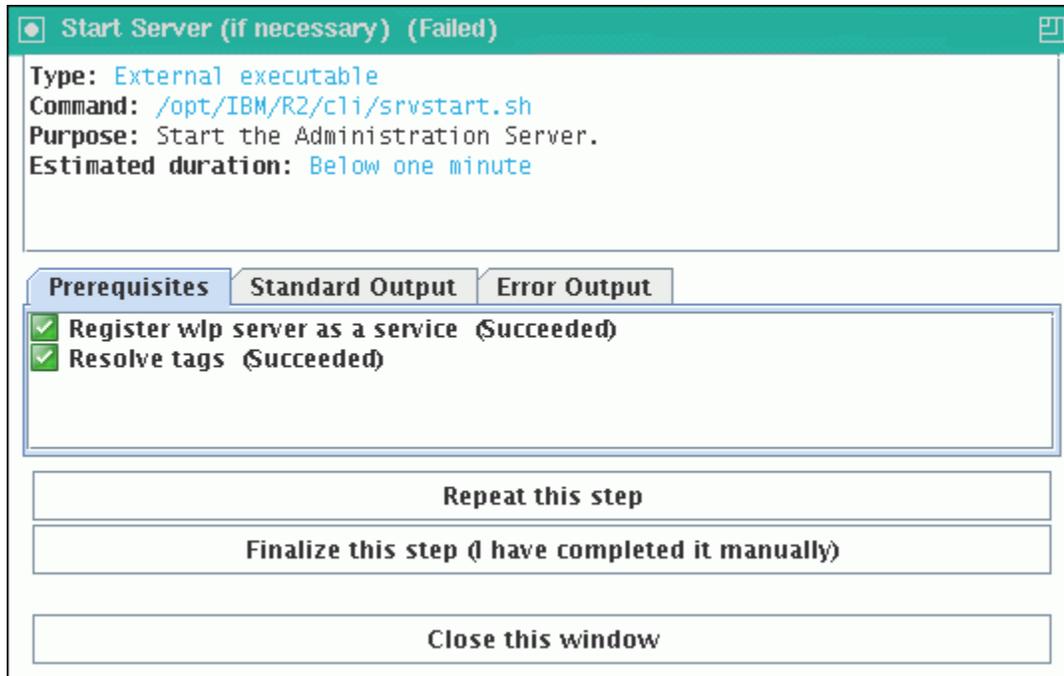
to run the installation wizard again. It detects that the previous configuration attempt failed or was interrupted and starts automatically in resume mode.

If you exited the installation wizard, run it again. It automatically starts the configuration.

1. If you encounter a problem, double-click the line that contains the step name, or right-click it and select **Details**. The line is indicated by a red icon.



2. Review the most important information that is displayed in the top. If a dedicate log file exists, its location is shown.



- Review the information that is shown on the **Standard Output** tab or, if applicable, on the **Dedicated Log** tab to determine the cause of the problem.

! **Important:** To reduce the performance load on the computer, the function that captures the dedicated log file runs with the lowest possible priority. Thus, the **Dedicated Log** tab does not always present the most recent and detailed information. What is more, the end of the log file might not be shown. If a failure occurs, check either the dedicated log whose location can be found in the step description, or the `msg_server.log` file.

- Fix the problem.
- On the installation panel, right-click the line that shows the problem, then click **Set > Ready (rerun the step)**. The installer completes the step and the remaining dependent steps. If you run the failed step outside the installation wizard, mark the step as completed successfully.

🔧 Note: If you cannot diagnose the problem and rerun the step manually, uninstall the product and try to install it again.

- Click **Next**. The **Postinstallation Summary** opens with information about installed components.

Configuring servers in separate networks

When the BigFix and BigFix Inventory servers are installed in separate networks, automatic lookup of their IP addresses might not work correctly. It might happen, for example in networks that use network address translation (NAT). To ensure that IP addresses of the servers are correct and the communication between them works, change the default settings, and manually enter the IP addresses or domain names of the servers.

By default, the address of the BigFix server is automatically retrieved from the database. In case of separate networks, the address might contain a local address of the server that cannot be recognized outside of its network. After you specify the IP address of the BigFix server, you also need to specify the correct address of the BigFix Inventory server to allow the BigFix server to access the software catalog.

1. During the initial configuration of BigFix Inventory, while specifying the connection details of the BigFix server, select **Disable automatic address lookup**, and enter the host name of the BigFix server. The address is always created based on the provided host name and port number.



Note: To access this panel after the installation, log in to BigFix Inventory, and click **Management > Data Sources**.

IBM BigFix Server*

Authentication (Console Operator)

User Name*

Password*

Disable automatic address lookup

Select this option to enter the host name manually, for example when the BigFix Server is in a separated network that uses network address translation (NAT).

Host*

Server API Port*



Tip: To check whether the BigFix host is correct, access the following URL from the BigFix Inventory network: `http://hostname:port/UploadReplication`. Expected message: `Error: no query parameters specified`.

2. On the computer where the BigFix InventoryLicense Metric Tool server is installed, add an environment variable with new addresses that will be used by the BigFix server to access the software catalog.
 - a. Go to one of the following directories, and create the `server.env` file.
 - **Linux** `install_dir/wlp/usr/servers/server1`
 - **Windows** `install_dir\wlp\usr\servers\server1`
 - b. Add the following variable to the file:

SERVER_URL_CATALOG=https://IP_address:port

IP address of the BigFix Inventory server that the BigFix server uses to obtain the software catalog and deliver it to the clients. This catalog is then used during the catalog-based software scan.

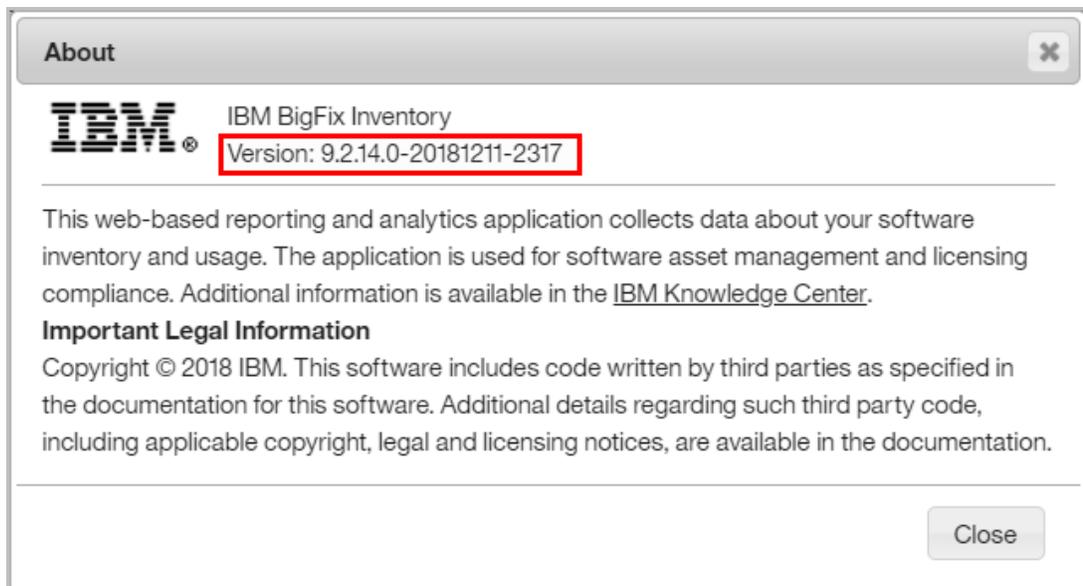
3. Restart the BigFix Inventory server.

Verifying the current version of BigFix Inventory and the software catalog

Check the current version of BigFix Inventory and the software catalog to make sure that the installation was successful.

1. Log in to BigFix Inventory.
2. Check the current version of the BigFix Inventory server.

- a. Hover over the **Help** icon , and click **About**.
- b. Check the version number in the About window.



3. Check the version of the software catalog.

- a. Go to **Management > Catalog Upload**.
- b. Check the version of the catalog under Catalog Version Information and make sure it is the latest available.

Software Catalog: Catalog Upload

To correctly identify software, ensure that the latest version of the catalog is uploaded. You can also upload the charge units file. The catalog, charge units file, or both will be imported during the next automatic [import](#).

Catalog Version Information

Current Catalog Version: 9.2.14.0
Endpoint Scanner Catalog: 1

Upload a Catalog File

Catalog file* No file chosen

Comments

Every time a newly uploaded IBM catalog is imported, the scanner catalogs are automatically distributed to the endpoints. If the automatic distribution fails, you must manually update scanner catalogs. [?](#)

▶ Details

Configuring

After you install BigFix Inventory, configure the application. Create accounts for the users who need access to the application and set up scans to collect software and hardware inventory data from your environment.

Creating users and groups

Each user who has access to BigFix Inventory must be assigned a role and a computer group. The role defines which reports and panels the user can view. The computer group narrows down the scope of these reports and panels to computers that meet certain criteria.

Setting up roles

A role is a collection of permissions that correlates to a list of privileges. Each user is assigned a role according to the privileges the user needs to efficiently operate the application. BigFix Inventory comes with a set of predefined roles. However, you can also create custom roles according to your needs.



You must be an Administrator to perform this task.

The Administrator role is set by default and cannot be modified.

1. In the top navigation bar, click **Management > Roles**.
2. To add a role, click **New**.
3. Specify the name and permissions that you want the new role to have, and click **Create**.

You created a custom user role. You can now create users and assign them suitable roles.

Task

[Setting up users \(on page clxxxix\)](#)

Roles

You can create a set of roles in BigFix Inventory and assign each of them with different permissions. Then, you can assign those roles to particular users, thus making them responsible for different actions. By setting up correct roles, you can easily establish which user has access to particular functions in BigFix Inventory. The roles can be assigned by the Administrator.

Pre-configured roles

BigFix Inventory has six preconfigured roles: Administrators, Auditors, Catalog Managers, Contract Managers, Software Asset Managers, and Infrastructure Administrators. Each of these roles has a different set of permissions that allow the users to perform various actions. The Administrator can edit or delete those roles, or create new ones and assign them permissions of your choice. The following table lists the preconfigured roles and their permissions.

Table 26. Roles preconfigured in BigFix Inventory and their permissions*The table consists of seven columns and 24 rows.*

Permission	Administrator	Auditor	Catalog Manager	Contract Manager	Software	
					Asset Manager	Infrastructure Administrator
Display Preview Panels	Yes					
Edit Server Configuration	Yes					
Manage Catalogs	Yes		Yes			
Manage Computer Groups	Yes					Yes
Manage Computer Properties	Yes					
Manage Contracts	Yes			Yes		
Manage Data Sources	Yes					
Manage Directory Servers	Yes					
Manage Hardware Inventory	Yes					
Manage Imports	Yes				Yes	Yes
Manage Licenses	Yes		Yes		Yes	
Manage Package Properties	Yes					Yes
Manage Part Numbers Uploads	Yes		Yes		Yes	Yes
Manage Roles	Yes					
Manage Scan Configurations	Yes					Yes
Manage Software Classification	Yes				Yes	
Manage Support Data	Yes					
Manage Uploads	Yes		Yes		Yes	Yes
Manage Usage Properties	Yes					Yes
Manage User Provisioning	Yes					
Manage Users	Yes					

Table 26. Roles preconfigured in BigFix Inventory and their permissions

The table consists of seven columns and 24 rows.

(continued)

Permission	Administrator	Auditor	Catalog Manager	Contract Manager	Software	
					Asset Manager	Infrastructure Administrator
Manage VM Managers and Servers	Yes					Yes
View Audit Trail	Yes	Yes			Yes	
View Catalog Audit	Yes	Yes	Yes		Yes	
View Contracts	Yes	Yes		Yes	Yes	
View Endpoints	Yes	Yes	Yes	Yes	Yes	Yes
View Hardware Inventory	Yes	Yes			Yes	Yes
View License Metrics	Yes	Yes		Yes	Yes	
View Raw Data	Yes	Yes		Yes	Yes	Yes
View Shared Disks	Yes	Yes		Yes	Yes	
View Software Catalog and Signatures	Yes	Yes	Yes	Yes	Yes	
View User Details	Yes	Yes			Yes	

Permissions

The following list describes permissions that are available in BigFix Inventory. Find out about each of them so that you can tailor the roles to your needs.



Important: Some of the permissions are specific only for the Administrator and cannot be assigned to other users. These permissions include:

- Display Preview Panels
- Edit Server Configuration
- Manage Computer Properties
- Manage Data Sources
- Manage Directory Servers
- Manage Roles
- Manage Support Data



- Manage User Provisioning
- Manage Users

Table 27. Permissions

Permission	Comments
Display Preview Panels	The user can enable preview features.
Edit Server Configuration	The user can change settings of the BigFix Inventory server.
9.2.14 Manage Catalogs	<p>The user can edit catalog servers, update the catalog as well as create, edit, and delete the content of the custom software catalog.</p> <p>9.2.11 The user can edit the end of support date for software components.</p> <p>9.2.14 The user can add and remove tags on the selected reports.</p>
Manage Computer Groups	The user can create, edit, and delete computer groups.
Manage Computer Properties	The user can specify additional computer properties that are to be gathered from the computers in the infrastructure.
Manage Contracts	<p>The user can create, edit, and delete contracts.</p> <p>9.2.4 The user can set or edit the license threshold to be notified about excessive license consumption.</p> <p>9.2.9 The user can add, edit, and delete custom fields.</p>
Manage Data Sources	The user can add, edit, and delete BigFix servers that are used as data sources by BigFix Inventory.
Manage Directory Servers	The user can add, edit, and delete directory servers that are used for authentication of BigFix Inventory users.
Manage Hardware Inventory (previously: Manage PVU per Core)	The user can change the Oracle core factor or the PVU per core value that is assigned to a processor.
Manage Imports	The user can schedule imports of data, and run them manually.
9.2.10 Manage Licenses	The user can add custom part numbers and new products to the software catalog, and assign additional license metrics to products.
Manage Package Properties	The user can create, edit, and delete application properties that are used to recognize software in your infrastructure.

Table 27. Permissions (continued)

Permission	Comments
9.2.3 Manage Part Numbers Uploads	The user can manage uploads of the part numbers file.
Manage Roles	The user can add, edit, and delete roles that are available in BigFix Inventory.
Manage Scan Configurations	The user can schedule software scans.
Manage Software Classification	The user can reassign software instances between different products, include or exclude them in pricing calculations, and share them between more than one product.
9.2.15 Manage Support Data	The user can access the Computer Support Data panel.
Manage Uploads	The user can manage uploads of the software catalog, and metric tables.
Manage Usage Properties	The user can create, edit, and delete application properties that gather information about the use of software in your infrastructure.
Manage User Provisioning	The user can integrate users from a directory server with BigFix Inventory.
Manage Users	The user can add, edit, and delete users that exist in BigFix Inventory.
Manage VM Managers and Servers	The user can create, edit, and delete VM managers.
View Audit Trail	The user can view the Audit Trail report that contains the history of all actions performed by users.
View Catalog Audit	The user can view information about changes to the custom software catalog.
9.2.3 View Contracts	The user can view contract usage data.
View Endpoints	The user can view information about the installed software as well as scan, registry, and raw data from endpoints.
View Hardware Inventory	The user can view the details of the processors that are used by the software.
View License Metrics	<p>The user can view the list of all software products that are contained in the PVU License Usage reports, the license type and usage for each product, the history of license consumption over the specified time period, and the top license consuming products.</p> <p>9.2.9 The user can view values that are specified in custom fields.</p>
View Raw Data	The user can view:

Table 27. Permissions (continued)

Permission	Comments
	<ul style="list-style-type: none"> • Metering Data report that contains information about the use of the software items • Unrecognized Files report that shows a ranking of the most frequently encountered files • Scanned File Data report that provides information about all files that are detected by the software inventory tool scanner • Package Data report that contains information about all installed packages
9.2.12 View Shared Disks	The user can view information about shared disks that are used in the infrastructure.
View Software Catalog and Signatures	The user can view the software catalog and signatures.
9.2.13 View User Details	The user can view information about current software users and their details for the software that is installed on the computers in the infrastructure.

Setting up users

Create users to give access to BigFix Inventory. Assign each user with a role to determine permissions that the user has, and with a computer group to determine the computers to which the user has access.

 You must be an Administrator to perform this task.

If your company has a security policy for passwords and account management, the password authorization mechanism that is available in BigFix Inventory should be used only for the initial setup.

9.2.7 Starting from application update 9.2.7, you can configure BigFix Inventory to ensure that user passwords fulfill a number of requirements. For more information, see: [Configuring security policy for user passwords \(on page dccxxvi\)](#). If you are using an earlier version of the application, or your password policy cannot be met by configuring the available password requirements, use LDAP or single sign-on. For more information, see: [Authenticating users with LDAP \(on page dccv\)](#) or [Configuring and enabling single sign-on \(on page dccxii\)](#).

1. In the top navigation bar, click **Management > Users**.
2. To add a user, click **New**.
3. Specify the name of the user and the role that you want to assign to that user.
4. Select the computer group to which the user is to have access and the authentication method. Then, click **Create**.

Task

[Setting up roles \(on page clxxxiv\)](#)

Setting up computer groups

Computer groups can be used to narrow down the scope of reports to computers that meet certain criteria or to share software discovered on one computer among all computers in the group, for example when a shared disk is mounted on multiple computers.

 You must have the Manage Computer Groups permission to create a reporting computer group. You must be an Administrator to create a software template group.

Computer groups can be created in the BigFix console or in BigFix Inventory. The groups are not automatically copied between the applications. However, groups created in the BigFix console can be recreated in BigFix Inventory as described in step 3 ([on page cxc](#)) below.

Groups created in the BigFix console are used to manage the infrastructure, assign computers to fixlet sites, and run fixlets against selected computers. Groups created in BigFix Inventory have two types:

- **Reporting.** This type is used to narrow down the scope of reports to computers that meet certain criteria. Such a group can have subgroups and can be assigned to users. It can be also used to assign contracts as well as upload part numbers to a subset of computers.
- **Software template.** This type is used only to share software discovered on one computer among all computers in the group.

The following procedure describes how to create a computer group in BigFix Inventory. For information about creating computer groups in the BigFix console, see: [Computer groups](#).

1. To create a computer group, click **Management > Computer Groups** and then, click **New**.
2. Provide the name and description of the computer group.
3. Specify filters according to which computers are assigned to the group, and click **Create**.

For example, to create a group that is based on a computer group in the BigFix console, choose Data Source Groups, in set, and select the computer group.

Definition

Specify the report filter which matches all of the following conditions:

Data Source Groups

in set

Europe

 **Restriction:** The size of data returned by the specified filter cannot exceed 1 MB.

4. 9.2.8 Choose the type of the computer group.

- Choose **Reporting** to create a group used for limiting the scope of computers visible on reports.

- 9.2.5** Select license metrics for which you want to calculate utilization in this computer group. Disabling the calculations for metrics in which you are not interested can improve import performance, especially in larger environments.

Select license metrics for which you want to calculate utilization within this computer group. Disabling the calculations for metrics in which you are not interested can improve import performance, especially in larger environments.

PVU, RVU MAPC and VPC [?](#)

Install Seats, Install Instances, Microsoft Physical Core with SA, Microsoft Virtual Core with SA, Microsoft Single Processor, Microsoft Dual Processor, Oracle Processor Core and other metrics [?](#)

- 9.2.2** Specify the number of days for which the data is to be calculated. This number defines the default period for which the data that is show on the All Metrics report is calculated and up-to-date.

Aggregate last [▲▼](#) days

For example, when you set the value to 30 days, and then display a report for the last 90 days, the report shows the following data:

- If you display the report for a newly created computer group, the report shows values for the last 30 days and is empty for the rest of the period.
- If you display the report for an existing group, the values might be outdated.

In both cases, recalculate the data.

- Optional: **9.2.3** To configure scans for this computer group, select **Create scan configuration**. Choose the scan frequency, requested start date, and types of software scans that you want to run on the computers in this group.

The option is available only if the definition of the group uses the Data Source Groups filter as the only condition based on which computers are assigned to this group. For more

information about configuring scans from the BigFix Inventory user interface, see: [Changing scan configuration from the BigFix Inventory user interface \(on page cxcvii\)](#).

Create scan configuration ?

Schedule

Frequency

Every day ▼

Requested Start Date

03/04/2016 11:47AM

Scan Types

Catalog-based Scan

File System Scan

Package Data Scan

Software Identification Tags Scan

Application Usage Statistics Scan

Resource Utilization Scan

- Choose **Software Template** to create a group used for sharing software discovered on one computer among all computers in the group. Such a group can be used for example to ensure that software discovered on a shared disk is reported on all computers on which the disk is mounted. Then, select software templates that you want to share between computers in this group.

Software templates

Shared Disk 1

5. To save the computer group, click **Create**.
6. To make the computer group available in BigFix Inventory, wait for the scheduled import or run it manually.

After you create a computer group, and run the import, you can view reports narrowed down to the computers that belong to this group. To do this, go to **Reports > Computer Groups**, and select a computer group. Then, from the upper left corner, select the report that you want to view for this computer group.

Overview ▼

Computer Group: Europe

Send Feedback

Inventory Data
1808 Software Installations including titles such as IBM BigFix Platform Agent, Internet Explorer, Symantec Endpoint Protection, Microsoft .NET Framework, Git, and Java Platform SE

Setting up scans to discover software and hardware inventory

The scanner is an independent, well-defined component that is used by BigFix Inventory. It is installed and managed through the BigFix client, and it enables the software and capacity scans. Software and capacity scans collect data that is later on displayed on the BigFix Inventory reports. For environments with up to a few thousand computers, you can enable the default scan configuration. In this case, the analyses are activated and software and capacity scans as well as uploads of their results are scheduled automatically on the computers that are subscribed to the BigFix Inventory site. For larger environments, it is advisable to divide the computers into groups and manually configure a separate scan schedule for each group to avoid performance issues.

Frequency of scans and uploads of data

By default, software and capacity scans are scheduled with a frequency that meets auditing requirements. If you want to change the default frequency, ensure that the new setup fulfills the minimal requirements and all considerations.

Frequency of scans

Table 28. Frequency of scans

Type of the scan	Default frequency	Minimal frequency	Comments
Capacity scan	Every 30 minutes	Every 30 minutes	<p>In case of subcapacity reporting, it is necessary to capture configuration of vCPU to count the maximum capacity of processor cores for each program concurrently. To ensure that dynamic changes of the vCPU configuration are captured, the capacity scan must run every 30 minutes. You can change the frequency of the scan only if you use BigFix Inventory for purposes other than subcapacity reporting.</p> <p>For information about scheduling the scan, see: Initiating the capacity scan on all computers (on page ccx).</p>
Software scan	Once per week	Once per month	<p>You can adjust the frequency of the software scan to the dynamics with which software inventory changes in your environment. For information about scheduling the scan, see: Initiating software scans (on page ccii).</p>
Scan of VM managers	Every 30 minutes	Every 30 minutes	<p>In case of subcapacity reporting, it is necessary to capture mobility of VMs. To ensure that it is captured, data from VM managers must be collected every 30 minutes. You can change the frequency of collecting data only if you use BigFix Inventory for purposes other than subcapacity reporting.</p>

Table 28. Frequency of scans (continued)

Type of the scan	Default frequency	Minimal frequency	Comments
			For information about scheduling the scan, see: Adding VM managers for VMware, Hyper-V, KVM with RHEV-M, Xen Server and Nutanix (on page cxciv) .
Shared disks scan	Once per week	Once per month	For information about scheduling the scan, see: Discovering software on shared disks (on page ccxxi) .
Disconnected capacity scan	Every 30 minutes	Every 30 minutes	<p>In case of subcapacity reporting, it is necessary to capture configuration of vCPU to count the maximum capacity of processor cores for each program concurrently. To ensure that dynamic changes of the vCPU configuration are captured, the capacity scan must run every 30 minutes. You can change the frequency of the scan only if you use BigFix Inventory for purposes other than subcapacity reporting.</p> <p>The scan is triggered and scheduled automatically after you install the disconnected scanner. For more information, see:</p> <ul style="list-style-type: none"> • IBM i: Installing the scanner and gathering initial data on IBM i systems (on page cclxxvi) • Other platforms: Installing the scanner and gathering initial data (on page cclii)
Disconnected software scan	On demand	Once per month	<p>You can adjust the frequency of the software scan to the dynamics with which software inventory changes in your environment.</p> <p>For information about scheduling the scan, see:</p> <ul style="list-style-type: none"> • IBM i: Running software scans and gathering scan results on IBM i systems (on page cclxxvii) • Other platforms: Running software scans and gathering scan results (on page ccliv)
Detailed hardware scan	On demand		<p>This type of the scan is not required for auditing purposes. Running this scan once per week or less frequently is sufficient to collect detailed hardware inventory.</p> <p>For information about scheduling the scan, see: Detailed hardware scan (on page ccxx).</p>

Frequency of uploading scan results

Upload of data from the scanned computer to the BigFix server occurs according to a frequency that is specific to the type of the scan. The data is then uploaded from the BigFix sever to the BigFix Inventory server during the import. In case of the disconnected scans, the data is uploaded directly to the BigFix Inventory server. It is not stored on the BigFix server.

Table 29. Frequency of uploading scan results

Summary for complex table

Type of up-loaded results	Default frequency	Comments
Capacity scan	When results change in comparison to the results of the previous scan.	Upload of capacity scan results is scheduled automatically when you schedule the capacity scan.
Software scan	When results change in comparison to the results of the previous scan.	For information about scheduling the upload, see: Uploading software scan results (on page ccviii) .
VM managers data	Every 12 hours	Upload of results from VM managers is scheduled automatically when you configure connection to VM managers.
Disconnected capacity scan	On demand	Results of disconnected capacity and software scans are imported directly to the BigFix Inventory server. They should be imported immediately after new results of the software scan are available. For more information, see: <ul style="list-style-type: none"> • IBM i: Importing scan results from IBM i systems to BigFix Inventory (on page cclxxviii) • Other platforms: Importing scan results to BigFix Inventory (on page ccliv)
Disconnected software scan	On demand	
Detailed hardware scan	Once per day	Results of the deep hardware scan are gathered daily by the Detailed Hardware Information analysis. For more information, see: Detailed hardware scan (on page ccxx) .

Default scan configuration

Default scan configuration is advised for environments with up to a few thousand computers. It ensures that the required analyses are activated and software and capacity scans as well as uploads of their results are scheduled automatically. These actions are necessary to collect data from the computers in your infrastructure and to display it on the BigFix Inventory reports. You can enable the default scan configuration during the configuration of BigFix Inventory connections to the databases or when adding another data source.

What happens automatically

When you enable default scan configuration, the following actions are automatically activated and scheduled on each computer that is subscribed to the BigFix Inventory site. The same configuration is also automatically applied to all computers that are subscribed to the BigFix Inventory site at a later time.

- The following analyses are activated:
 - Application Usage Statistics
 - Capacity Configuration for Linux on z Systems
 - Installed UNIX Packages
 - Installed Windows Applications
 - Scanner Information
 - Software Scan Status
 - VM Manager Information
- The scanner is installed.
- All [types of software scans \(on page ccvi\)](#) and uploads of their results are scheduled.
- Capacity scans and uploads of their results are scheduled.

Additionally, when you enable the default scan configuration for the second and successive data sources that you add, the software catalog is distributed to the computers on condition that it was successfully uploaded and imported for the first data source.

What to do next

- After you enable the configuration, run the import of data. The data that is collected from the computers in your infrastructure should be displayed on the BigFix Inventory reports. However, the newly set-up environment needs time to finish the scans and to upload their results to the server. If the reports do not contain any data after the first import, wait about an hour until the scans are completed. Then, run another import.

To change the default scan configuration, use the [Scan Configurations \(on page cxcvii\)](#) panel. Alternatively, you can change the scan configuration through the BigFix console by using the [Initiate Software Scan \(on page cciii\)](#) fixlet.

- [Add VM managers \(on page ccxiv\)](#) that are an important and mandatory addition to the capacity scan. These are the hosts that manage your virtual machines. BigFix Inventory can already evaluate the virtual capacity of your environment, but needs extra data about the physical one to accurately calculate your subcapacity consumption. This data can be collected only from VM managers.

Changing scan configuration from the BigFix Inventory user interface

9.2.3 Available from 9.2.3. The **Scan Configurations** panel allows for basic scan management directly from the BigFix Inventory user interface. Use the panel if you enabled the default scan configuration or as an alternative to scheduling scans from the BigFix console.

The panel is used for configuring the time, frequency, and types of software scans that are run on the endpoints. If you enabled the default scan configuration, all prerequisites are fulfilled and you can manage scans without any further setup. In other cases, ensure that the required analyses are activated, the scanner is installed on the endpoints, uploads of software scan results are scheduled, and capacity scans and uploads of their results are scheduled. You must complete these actions manually from the BigFix console before you can schedule scans from the BigFix Inventory user interface. For more information, see: [Manual scan configuration \(on page cc\)](#).

After the prerequisites mentioned above are fulfilled, divide the computers in your infrastructure into groups in the BigFix console. Then, based on these groups, create computer groups for which you want to schedule software scans in BigFix Inventory. A computer group that you create in BigFix Inventory must fulfill the following conditions to allow for scheduling scans through the **Scan Configurations** panel:

- It has the All Computers group as its parent. It is not a subgroup of another computer group.
- Definition of this group uses a single Data Source Groups filter as the only condition based on which computers are assigned to this group. Data source groups represent computer groups that were earlier created in the BigFix console.

Definition

Specify the report filter which matches all of the following conditions:

Data Source Groups

in set

Europe

✖

+

 You must be an Administrator or an Infrastructure Administrator to perform this task.

When you create scan configuration for a computer group in the BigFix Inventory user interface, the Initiate® Software Scan action is created and sent to the BigFix server. Each time you change scan configuration on the **Scan Configurations** panel, the related action is updated on the BigFix server. When you delete the scan configuration, the related action is by default removed from the BigFix server.

If you configure scans on the **Scan Configurations** panel and then change scan configuration in the BigFix console, information about the change is not reflected on the BigFix Inventory user interface. To avoid duplicated scans or unexpected scan behavior, avoid configuring scans in both places and choose your preferred method instead.

1. [Create computer groups and enable scans for these groups \(on page cxc\)](#).

- To edit scan configuration, click **Management > Scan Configurations**, and select the computer group for which you want to change the scan settings.



Note: Each row represents the schedule and status of the software scan for a particular computer group. The row All computers from *Data source name* represents all computers that report to one BigFix server. The row is created automatically. You can edit or disable the scan schedule for this group.

Status	Name	Enabled	Start Date	Frequency	Catalog-...	File Syst...	Package ...	Software...	Applicati...	Resourc...
<input checked="" type="checkbox"/> Active	All compu...	Yes	02/25/201...	Every 7 days	Yes	Yes	Yes	Yes	Yes	Yes
<input checked="" type="checkbox"/> Active	Europe	Yes	02/25/201...	Every day	Yes	Yes	Yes	Yes	No	No

- Choose the scan frequency and requested start date.

Due to any network latency, the time when the scan actually starts might be several minutes different from the requested start time. The requested start date represents the time and date when the action will be started in accordance to the time zone of the BigFix server.

- Choose the types of software scan to run.

For more information about each type of software scan, see: [Types of software scans \(on page ccvi\)](#).

If the computer group for which you are configuring scans consists entirely of computers that run on Mac OS, run only the package data scan on these computers as other types of software scans are not applicable to Mac OS. If the group consists of computers that run on Mac OS and computers that run on other operating systems, select all types of scans that you want to run. Only the package data scan is run on the Mac OS computers. The remaining types of scans are skipped as not applicable.

Edit Scan Configuration

Name Europe

Data Source Groups Datasource/Europe

Enabled

Schedule

Frequency Every day ▾

Requested Start Date 02/09/2016 12:35PM

Scan Types

- Catalog-based Scan
- File System Scan
- Package Data Scan
- Software Identification Tags Scan
- Application Usage Statistics Scan
- Resource Utilization Scan

5. To save the changes, click **Save and Start Scans**. Then, click **OK** to confirm that you want to restart the scan action that is running on the BigFix server.

! **Important:** When you change a single scan configuration, all scans that were configured from the BigFix Inventory user interface are restarted on the BigFix server.

Scan configuration for a computer group can have one of the following statuses:

- Active - The Initiate® Software Scan action was started on the BigFix server. The software scan runs according to the schedule defined on the **Scan Configurations** panel.
- Stopped - The Initiate® Software Scan action was stopped on the BigFix server. The software scan is not running. You can re-enable the scan on the **Scan Configurations** panel or configure the scan from the BigFix console.
- Server unavailable - The scan schedule cannot be displayed and set because the BigFix server or the Web Reports server is not responding. Follow the [troubleshooting steps \(on page dcclxxxvii\)](#) to resolve the issue.

If you want to disable software scans for a particular computer group, open the scan configuration for this group and clear the **Enabled** check box.

Manual scan configuration

Manual scan configuration is advised for environments with more than a few thousand computers. It requires that analyses are activated and software and capacity scans as well as uploads of their results are scheduled manually after the installation. These actions are necessary to collect data from the computers in your infrastructure and display in on the BigFix Inventory reports. Flexibility of the manual scan configuration allows for avoiding performance issues that might occur when you scan too many computers at the same time.

To manually configure scans, perform the following tasks:

1. Optional: Divide the computers into groups in the BigFix console.

In large environments, it is advisable to divide the computers into groups in the BigFix console and configure a separate scan schedule for each group. For a summary of good practices related to scan distribution, see: [Distribution of scans for improved performance \(on page cc\)](#). For more detailed information, see the [Tuning performance](#).

2. Set up scans to discover software and hardware discovery.
 - a. [Activate the analyses. \(on page cci\)](#)
 - b. [Install the scanner. \(on page ccii\)](#)
 - c. [Initiate software scans. \(on page cciii\)](#)
 - d. [Upload software scan results. \(on page ccviii\)](#)
 - e. [Collect capacity data. \(on page ccx\)](#)

After you successfully schedule software and capacity scans, run the import of data. When the data is displayed on the BigFix Inventory reports, you can further customize the scans. For example, you can exclude directories from software scans or scan remote shared file systems. For more information, see: [Advanced scan configuration \(on page ccxix\)](#).

Distribution of scans for improved performance

Performance of importing data from the BigFix server to BigFix Inventory depends on the number of scan files, usage analyses, and package analyses that are processed during a single import. By properly scheduling scans and distributing them over the computers in your infrastructure, you can reduce the length of the data import.

Use the following guidelines to improve the performance of the data import:

Divide computers into groups

- After you install BigFix Inventory, do not run the scans on all computers in your infrastructure. Divide the computers into groups in the BigFix console and start by gathering the default properties from a single computer group. For more information about creating computer groups in the BigFix console, see: [Computer groups](#).
- Consider creating computer groups that are based on stability. In a stable environment, scans can be run less frequently than once a week.

Schedule scans to run at different time

- Schedule scans to run on different days for different computer groups.
- Avoid a situation in which multiple groups are scanned at the same day or the following day. It might cause that scans and data imports interfere.
- Reduce the frequency of scans. In most cases, it is sufficient to scan the infrastructure once a week, which is the default frequency. In large environments, you can disable the option to automatically run scans and initiate them only when necessary. The minimum scan frequency is once per month.

Limit the amount of gathered data

- In the initial deployment phase, or if you do not need the information, disable the collection of software usage. For more information, see: [Disabling the collection of software usage \(on page cclxviii\)](#).
- Try to limit the number of computer properties that are gathered during scans.

For more information about distributing scans and other actions that you can undertake to improve the application performance, see: [Tuning performance](#).

Activating the analyses

An analysis is a collection of property expressions that is used to view and summarize properties of computers in your infrastructure. Some of the analyses are required to discover installed software and must be activated all the time. Other analyses are used only in specific scenarios or for troubleshooting purposes and can be activated according to your needs.

If you enabled the default scan configuration, the required analyses are activated automatically and this configuration is not required.

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Analyses**.
3. To activate an analysis, right-click it, and click **Activate**.

To ensure that the installed software is properly detected, activate the following analyses on all computers in your infrastructure:

- Installed UNIX Packages
- Installed Windows Applications
- Software Scan Status

To measure software usage, activate also the Application Usage Statistics analysis.

[Install the scanner \(on page ccii\)](#) on the computers on which you want to detect the software.

Installing the scanner

The scanner collects hardware information as well as information about files and packages that exist on the computers in your infrastructure. Install the scanner on every computer on which you want to detect software. If you enabled the default scan configuration, the scanner is installed automatically and this configuration is not required.

- Ensure that the BigFix client is installed and running on the computers on which you want to detect software and that the computers are subscribed to the **IBM BigFix Inventory v9** site.
- To install the scanner on a WPAR, you must install it on the LPAR first.
- **9.2.17** **AIX** Ensure that on AIX the `xlC.rte` library version 13.1.0.0 or higher is installed on the target computer.
- **Linux** Ensure that on Red Hat Enterprise Linux 8.0 the `libns1.so.1` library is installed on the target computer. The last version of the scanner for 32-bit Linux x86 with the `libstdc++.so.5` library is 2.8.0.3000.
- **UNIX** Ensure that the `libstdc++.so.6` library is installed on the target computer.
- **Mac OS X** Software discovery on Mac OS X does not require installing the scanner nor scheduling regular uploads of scan results. To enable the discovery, it is enough to [run the package data scan \(on page cciii\)](#).

The scanner can be used by multiple IBM applications. By default, the scanner that is delivered with BigFix Inventory is installed in the single user mode. It means that if it is used by other applications, it must be run by root (UNIX) or SYSTEM (Windows). If any of the applications requires that the scanner is run by a different user, install it in the regular mode. For more information, see: [Using CIT deployed by LMT/BFI 9 by other IBM products](#).

1. Log in to the BigFix console.
2. In the navigation bar, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. In the upper right pane, select **Install or Upgrade Scanner**.
4. **9.2.5** Select the installation directory. You can choose the default directory, or an alternative one. In the latter case, the scanner is installed in the same path as the BigFix client.



Note: The setting applies only to new installations, you cannot change the path if you are upgrading the existing scanner.

Use default installation directory:

Windows: `C:\Program Files\tivoli\cit`
UNIX: `/opt/tivoli/cit`

Use an alternative installation directory:

Windows: `<BES Client>\CITBin`
UNIX: `<BES Client>/CITBin`

With this setting, the scanner log files will be placed in the `logs` directory under the scanner installation directory.

5. Click **Take Action**.
6. Select the computers on which you want to install the scanner, and click **OK**.



Note: To install the scanner on a WPAR, install it on the LPAR first.

Computer Na...	OS	CPU	Last Report Ti...	Loc
NC91281112...	Linux Red Hat ...	2400 MHz Xeon	2014-06-12 14:...	No
NC91431260...	Linux Red Hat ...	2400 MHz Xeon	2014-06-12 13:...	No
NC91431261...	Linux Red Hat ...	2400 MHz Xeon	2014-06-12 14:...	No
NC91431261...	Linux Red Hat ...	2400 MHz Xeon	2014-06-12 14:...	No



Tip: To ensure that the action is applied on all computers that are added in the future, select **Dynamically target by property**.

You installed the scanner on the computers in your infrastructure. To view information about the installed scanner, for example its version, activate the Scanner Information analysis, and open the **Results** tab.

After you install the scanner, [schedule software scans \(on page cciii\)](#).

Initiating software scans

During the software scans, the scanner collects information about files with particular extensions, package data, and software identification tags to evaluate whether particular software is installed on the computers in your infrastructure. It also gathers information about the running processes to measure software usage. By default, software scans are scheduled to run regularly but you can specify the exact days and times of the scans, or modify their start and end dates.

Backups of software directories that are stored on the endpoints might be reported as separate software instances. It might result in false discoveries and incorrect license consumption. To avoid this problem, either [exclude the backups from scanning \(on page ccxiv\)](#) or compress them with a data compressor.

The scanner does not scan the following elements:

- Archived files such as ZIP or RAR files
- Paths that contain certain whitespace characters and other special characters. To avoid a situation in which software is not discovered, ensure that file paths in your infrastructure do not contain the unsupported characters. For more information, see: [Limitations \(on page I\)](#).

If you enabled the default scan configuration, software scans are initiated by default and this configuration is not required.

Software scan is resource-consuming, and its duration depends on the number of files in the file system. To minimize the impact of the scan, schedule the software scan for the period that is least critical for your environment, for example, at night or during the weekend. Software scan needs to be run with frequency that is adjusted to the dynamics of the environment, but typically few times per each reporting period. For information about default and minimal scan frequency, see: [Frequency of scans and uploads of data \(on page cxciij\)](#).

Additionally, to minimize the impact of the scan, you can set up a CPU threshold for the scan, and optimize the cache configuration. For more information, see: [Optimizing scanner cache configuration \(on page cclxvi\)](#).

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. Select **Initiate Software Scan**.
If the task is not applicable, see: [Scan problems \(on page dcclxxxi\)](#).
4. To ensure complete software discovery, select all types of scans under the **Software Discovery** section.

Software Discovery

To guarantee complete software discovery, select all check boxes below. If you leave any of the check boxes deselected, some of the installed software might not be discovered.

- Catalog-based scan**
Discovers software based on scanner catalogs that are sent to endpoints and matched with the scanner findings.
- File system scan.**
Discovers software with additional row data based on file extensions and processes that were run on the computers. File system scan can retrieve information about checksums. To modify a set of collected file checksums, use [Configure File Checksums Collection \(MD5/SHA-256\)](#).
- Package data scan**
Discovers software based on information about Windows and UNIX packages that exist in the registry.
- Software identification tags scan**
Discovers software based on the existence of software identification tags that exist on the endpoints. This type of scan discovers the majority of IBM software.

Different types of scans search for different types of information to determine whether the software is installed or to measure its usage. Generally, all types of scans should be run regularly. However, you can choose to run different types of scans at different times or distribute the scan schedule over the computers in your environment to improve the performance of the import. For more information, see: [Types of software scans \(on page ccvi\)](#).

Mac OS X

To discover software that is installed on Mac OS X, run the package data scan on Mac computers. To measure application usage, activate the Application Usage Statistics analysis and run the application usage statistics scan. Other types of software scan are not necessary and are ignored on Mac OS X.

5. To see license metric utilization on BigFix Inventory reports, select the relevant options under the **Additional Scan Capabilities** section.

Additional Scan Capabilities

The following options collect additional information about software products.

- Application usage statistics**
Collects information about processes that are running on the endpoints and translates them into usage statistics.
- Resource utilization scan**
Searches for software license metric tags that contain information about types of licenses that can be used by a product and their usage. NOTE: It is recommended to schedule the resource utilization scan to run daily. For more information, see the [product documentation](#).
- User information**
Collects information about users for specific software products to calculate usage of selected user-based licenses.
 - Store and show user names on reports**
By default, the user names are obfuscated with SHA-1. Use this option if you want to show this data.



Note: The resource utilization scan collects information contained in `.slmtag` files. The license metric quantity for these products is based on the `.slmtag` data and is reported on All Metrics and All IBM Metrics panels. To ensure that the reported data is accurate and current, it is recommended to schedule the resource utilization scan to run daily. If you prefer to run the remaining scans at different times, run the Initiate Software Scan fixlet for resource utilization scan separately.

6. **Optional:** To enable scanning of shared disks in basic mode, select **Scan remote shared disks**. The basic mode is advised for environments where a single shared disk is mounted on a single computer or a few computers. However, this mode is only one of the available solutions. Before selecting this option, see: [Discovering software on shared disks \(on page ccxxi\)](#).
7. **Optional:** To limit the amount of processor resources that the scanner consumes, select **Initiate the software scan with CPU threshold**. Specify the consumption limit that is in the range 5 - 100. The higher value you specify, the higher is the consumption limit. For example, if you specify 75, scanner processes use the average of 75% of the processing power of the target computer.



Important: Setting the threshold does not guarantee that the CPU consumption is always below the specified value. It fluctuates around that value, sometimes exceeding it and sometimes dropping below it. Temporary peaks are expected.

Mac OS X

The setting does not apply to Mac OS X because the scanner is not installed on these computers.

CPU threshold

- Initiate the software scan with CPU threshold

CPU threshold value: (range 5-100)

8. To start the scan, click **Take Action**.

9. Select the computers on which you want to initiate the scans.

Only local drives are scanned on the selected computers. For information about scanning shared disks, see: [Discovering software on shared disks \(on page ccxxi\)](#).



Tip: To ensure that the action is applied on all computers that are added in the future, select **Dynamically target by property**.

10. **Optional:** By default, the scans are scheduled to run weekly. If you want to specify the dates and frequency of the scans, open the **Execution** tab. Specify the details, and click **OK**.

For information about default and minimal scan frequency, see: [Frequency of scans and uploads of data \(on page cxcii\)](#).

[Schedule uploads of software scans results \(on page ccviii\)](#) to the BigFix server.

Types of software scans

The scanner can search for different types of information to determine whether the software is installed or to measure its usage. Generally, all types of scans should be run regularly. However, you can choose to run different types of scans at different times or distribute the scan schedule over the computers in your environment to improve the performance of the import process.

Software discovery

Catalog-based scan

In this type of scan, the BigFix server creates scanner catalogs that are sent to the endpoints. Based on those catalogs, the scanner discovers exact matches and sends its findings to the server. Scanner catalogs do not include signatures that can be found based on the list of file extensions nor entries that are irrelevant for a particular operating system.

File system scan

In this type of scan, the scanner uses a list of file extensions to check whether any files with those extensions exist on the endpoints. Then, it returns the findings to the BigFix server where the discovered files are compared with the software catalog. If a particular file matches an entry in the catalog, the software is discovered.

Collecting executable files based on application usage

9.2.8 **Linux** **Solaris** Starting from application update 9.2.8, the file system scan additionally reports files that are based on processes that have been run on the computers in your infrastructure, regardless of their extension. To view these files, you need to meet the following criteria:

- The BigFix server and client are in version 9.5.5 or higher.
- The operating system is Linux or Solaris.
- Application Usage Statistics analysis is activated.

The reported files, with supplementing information such as path and size, are available on the Scanned File Data report and can be used to create custom discovery and use signatures for software components. If the process has been active at least once since the feature was turned on and the associated file exists on the computer, this file is displayed on the report.

This feature is enabled by default. To stop collecting information about the files, run the **Disable Collecting Executable Files Based on Application Usage** fixlet. In case you need to reactivate this feature, run the **Enable Collecting Executable Files Based on Application Usage** fixlet.



Note: Checksums collection (MD5 and SHA-256) for these files is not supported.

Outputs of the file system scan

The scan generates two outputs: full file system scan and delta file system scan. The former contains information about all files that were discovered on the endpoint. The latter contains information only about files that changed between the last two full file system scans. Both outputs are generated during every scan and uploaded to the BigFix server. However, only one of them is imported to BigFix Inventory.

Delta file system scan is imported to BigFix Inventory in the first place to improve the import performance. Full file system scan is imported instead of the delta scan in the following cases:

- The file system scan is run on the endpoint for the first time
- The BigFix client was reinstalled on the endpoint
- The delta file system scan is larger than one third of the full file system scan
- The endpoint does not meet the prerequisites for generating the delta scan. If any of the prerequisites are not met on an endpoint, check the `BES_Client\LMT\CIT\delta.log` file for more details.
 - **UNIX** `sed, diff, wc, tar, gzip, expr`
 - **Windows** VBScript interpreter

Package data scan

In this type of scan, the scanner searches the system registry to gather information about Windows™ and UNIX™ packages that are installed on the endpoints. Then, it returns the findings to the server where the discovered packages are compared with the software catalog. If a particular package matches an entry in the catalog, the software is discovered.

Software identification tags scan

In this type of scan, the scanner searches for software identification tags that are delivered with software products. Then, it returns the findings to the BigFix server where the tags are processed. Based on the information that they contain, the software is discovered.

Additional scan capabilities

Application usage statistics

In this type of scan, the scanner gathers information about processes that are running on the endpoints. Then, it returns the findings to the BigFix server where the data is translated into usage statistics.

The usage data is first collected when the BigFix client is installed on an endpoint, and an application usage statistics is enabled. The statistics are displayed on the Metering Data report. When the processes are matched against usage signatures from the software catalog, or custom usage signatures, the statistics are available on the Software Classification panel.

9.2.11 Starting from application update 9.2.11, BigFix Inventory additionally collects application usage data in a new format.



Remember: By default, the usage scan is scheduled to run weekly to avoid performance issues. If you want to collect software usage on a daily basis, run the usage scan daily.

Resource utilization scan

In this type of scan, the scanner searches for software license metric tags that contain information about types of licenses that can be used by a product and their usage. The scanner returns its findings to the BigFix server where the tags are processed. Based on the information that they contain, the maximum usage of license metrics over the last 30 days and its trend value are calculated. For more information, see: [Raw utilization of license metrics \(on page dcxi\)](#).



Tip: The scan collects information about license metrics that are reported by products which implemented the ISO/IEC 19770-4:2017 standard. Because the volume of the collected data might be large, do not run this scan if you do not want to monitor these metrics.

9.2.13 User information

In this type of scan, the scanner collects information about users of the specific software products that help calculate the usage of some license-based products. The results are shown on the All Metrics and Software Users reports. For more information, see: [Available reports \(on page dcxvi\)](#).

Uploading software scan results

When the software scan finishes, its results must be uploaded from the computers in your infrastructure to the BigFix server. To ensure that software inventory data is up-to-date, software scans and uploads of their results should run on a similar schedule.

If you enabled the default scan configuration, uploads of software scan results are scheduled automatically and this configuration is not required.

1. Log in to the BigFix console.
2. In the navigation bar, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. In the upper right pane, select **Upload Software Scan Results**, and then in the lower pane, click **Take Action**.



Note: The size of a single compressed scan result cannot exceed 1 MB.

Fixlets and Tasks

Name	Source Sev...	Applicab...	Category
Uninstall VM Manager Tool	Low	1 / 13	VM Managers
Unset DSD Mode	Low	0 / 13	Configuration
Upload Scanner Diagnostic Data	Low	0 / 13	Troubleshooting
Upload Software Scan Results	High	7 / 13	Scanner
Upload VM Manager Tool Scan Results	Low	1 / 13	VM Managers

Task: Upload Software Scan Results

 **Take Action**
 Edit
  Copy
  Export
  Hide Locally
  Hide Globally
  Remove

4. Select computers from which you want to upload software scan results.



Tip: To ensure that the action is applied on all computers that are added in the future, select **Dynamically target by property**.

Target Execution Users Messages Offer Post-Action Applicability Success Criteria Action Script

Target:

Select devices
 Dynamically target by property
 Enter device names

Applicable Computers (4)

Computer Na...	OS	CPU	Last Report Ti...	Lo
NC91281112...	Linux Red Hat ..	2400 MHz Xeon	2014-06-12 14:...	No
NC91431260...	Linux Red Hat ...	2400 MHz Xeon	2014-06-12 13:...	No
NC91431261...	Linux Red Hat ...	2400 MHz Xeon	2014-06-12 14:...	No
NC91431261...	Linux Red Hat ...	2400 MHz Xeon	2014-06-12 14:...	No

5. **Optional:** By default, software scan results are uploaded to the BigFix server whenever they are available on the target computer. If you want to specify the dates and frequency of the uploads, open the **Execution** tab. Specify the details, and click **OK**.

The screenshot shows the 'Execution' tab in the BigFix console. It is divided into two main sections: 'Constraints' and 'Behavior'.

Constraints:

- Starts on: 2014-06-16 at 09:55:14 (client local time)
- Ends on: 2014-06-18 at 09:55:14 (client local time)
- Run between: 01:00:00 and 02:59:00 (client local time)
- Run only on: Sun, Mon, Tue, Wed, Thu, Fri, Sat (client local time)
- Run only when: Active Directory Path matches

Behavior (highlighted in red):

- On failure, retry: 3 times
 - Wait: 1 hour between attempts
 - Wait until computer has rebooted
- Reapply this action
 - whenever it becomes relevant again
 - while relevant, waiting: 15 minutes between reapplications
 - Limit to: 3 reapplications
- Start client downloads before constraints are satisfied
- Stagger action start times over: 5 minutes to reduce network load

[Schedule capacity scans and uploads \(on page ccx\)](#) of their results.

Collecting the hardware information and capacity data

The hardware information and capacity data concern information about your processors, and the processor capacity that is assigned to each computer. You can review the information on the Hardware Inventory report. The hardware and capacity data is essential to calculate processor-based metrics, for example PVU.

Concept

[Discovering software in Docker containers \(on page ccxxxvi\)](#)

Initiating the capacity scan on all computers

The capacity scan is the essential part of collecting capacity data and should run on all computers, including Oracle Logical Domains. The scan runs every 30 minutes, and collects data about the physical processor, system type (physical or virtual), guest operating system, and logical partitions.

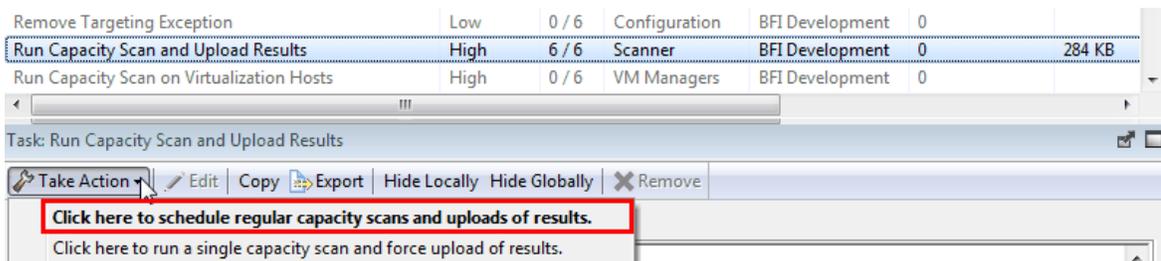
- **UNIX** You must have `tar` and `gzip` installed on each target computer.
- For KVM x86 hosts that are not controlled by RHEV-M and PowerKVM hosts, do not use the **Run Capacity Scan and Upload Results** task to collect the capacity data. Instead, use the **Run Capacity Scan on Virtualization Hosts** task. For more information, see: [Collecting capacity data from virtualization hosts for Xen and KVM \(on page ccxvi\)](#).

The scan runs every 30 minutes. However, it runs very quickly and has minimal impact on the processor usage. Moreover, scan results are uploaded to the BigFix server only if they changed since the last scan. It ensures that no unnecessary network traffic is generated.

Restriction: In case of subcapacity reporting, it is necessary to capture configuration of vCPU to count the maximum capacity of processor cores for each program concurrently. To ensure that dynamic changes of the vCPU configuration are captured, the capacity scan must run every 30 minutes. You can change the frequency of the scan only if you use BigFix Inventory for purposes other than subcapacity reporting.

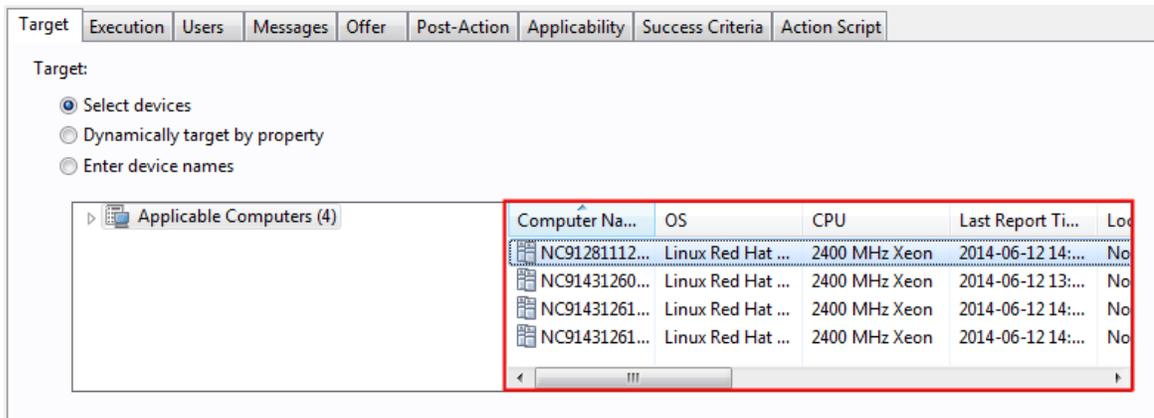
1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**, and select **Run Capacity Scan and Upload Results**.
3. In the lower pane, click **Take Action**, and choose **Click here to schedule regular capacity scans and uploads of results**.

Important: Single capacity scan runs only once. Results of this scan are uploaded to the BigFix server regardless of whether they differ from the results of the previous scan. Use this option only when you need to force the upload of scan results, for example when the import of scan results failed. Otherwise, run the regular capacity scan.



4. Open the **Target** tab, and select computers that you want to scan.

i **Tip:** To ensure that the action is applied on all computers that are added in the future, select **Dynamically target by property**.



When the scan completes successfully, scan results are automatically uploaded to the BigFix server.

Client Installation on Oracle Solaris

To ensure that the capacity data is properly collected from computers that are running on the Oracle Solaris virtualization, install the BigFix client on all required operating systems, zones and domains. Then, set up capacity scans on each endpoint.

Install the BigFix client in the following order:

1. Global zone in the Control Domain (if it exists)
2. Global zones
3. Local zones and Kernel zones

You might change the order and install the BigFix client on the computers in the local zone first. As a result, these computers have the No Host Scan Data status until you install the BigFix client on the computers in the global zone.

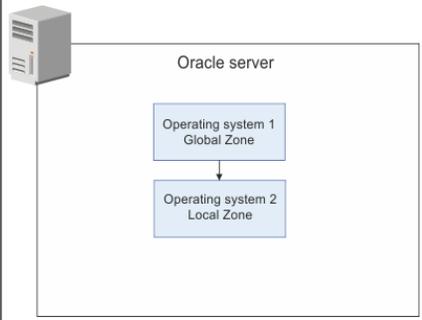
Examples

The following examples of possible installation scenarios describe in what order you might install the BigFix client on computers in the global and local zones.

Scenario	Applicable for	
	Solaris x86	SPARC

To discover an application that is installed on operating system 2 (local zone), install the client in the following order:

1. Operating system 1 (global zone)
2. Operating system 2 (local zone)

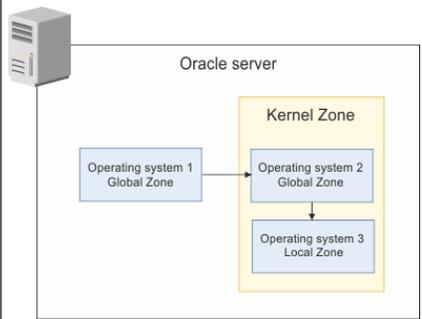


Yes

Yes

9.2.14 To discover an application that is installed on operating system 3 (local zone), install the client in the following order:

1. Operating system 1 (global zone)
2. Operating system 2 (global zone in the Kernel Zone)
3. Operating system 3 (local zone in the Kernel Zone)

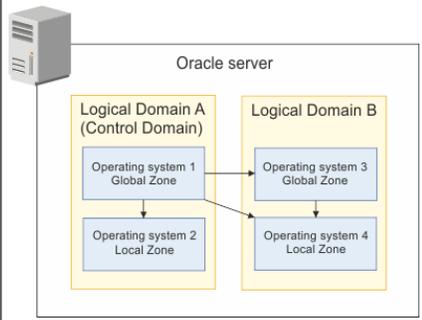


Yes

Yes

To discover an application that is installed on operating system 4 (local zone on a Logical Domain that is not the Control Domain), install the client in the following order:

1. Operating system 1 (global zone in the Control Domain)
2. Operating system 3 (global zone of operating system 4 in the Logical Domain)
3. Operating system 4

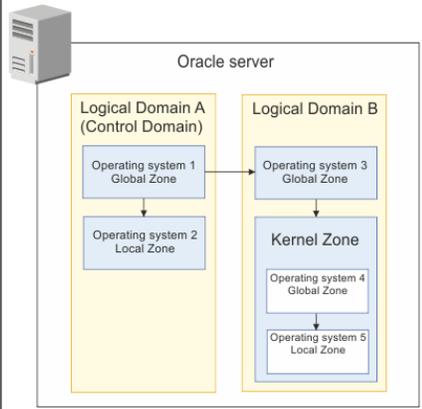


Yes

Yes

To discover an application that is installed on operating system 2, install the client in the following order:

1. Operating system 1 (global zone in the Control Domain)
2. Operating system 2

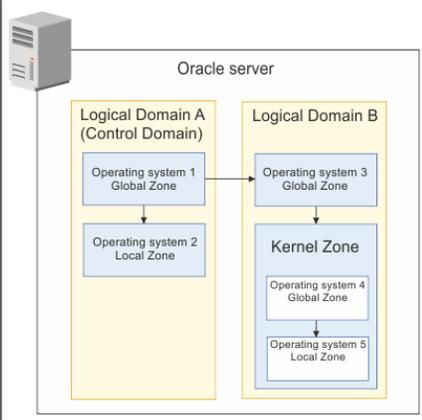


Yes

Yes

9.2.15 To discover an application that is installed on operating system 5 (local zone in the Kernel Zone), install the client in the following order:

1. Operating system 1 (global zone in the Control Domain)
2. Operating system 3 (global zone of operating system 4 in the Logical Domain)
3. Operating system 4 (global zone of operating system 5 in the Kernel Zone)
4. Operating system 5



Yes

Adding VM managers for VMware, Hyper-V, KVM with RHEV-M, Xen Server and Nutanix

After you configure the capacity scan, BigFix Inventory can evaluate capacity of virtual machines in your environment. However, to accurately calculate subcapacity consumption, it also needs data about physical hosts on which these VMs are running. The data can be collected only from VM managers. To ensure that the data is available to BigFix Inventory, configure connections to VM managers in your infrastructure.

- You must have the Manage VM Managers and Servers permission to perform this task.
- The computer where the main instance of the VM Manager Tool is installed must belong to the computer group to which you are assigned. Otherwise, the option to add VM managers is disabled.
- If the panel is blocked and displays a warning, see: [Troubleshooting: Enabling the VM Managers panel \(on page ccclxiv\)](#).

9.2.7 VM Manager Tool, starting from update 1.5.0.0, is enhanced to attempt to automatically fix the most common problems with the VM manager connection parameters. This feature is disabled by default. To enable this feature, go to **Management > Advanced Server Settings** and change the setting of **vmman_fix_parameters_enabled** parameter to `true`. While fixing connection problems the VM Manager Tool can connect using the HTTP protocol when the defined HTTPS connection is not available. This setting is enabled by default. To disable switching protocol from HTTPS to HTTP, go to **Management > Advanced Server Settings** and change the setting of **vmman_http_connection_allowed** parameter to `false`. For more information, see: [Advanced server settings \(on page cdii\)](#).

1. In the top navigation bar, click **Management > VM Managers**.
2. To add a VM manager, click **New**.
3. Select the virtualization type and provide the required details.

- To add Microsoft Hyper-V:
 - a. Select one of the available communication interfaces: PowerShell or WinRM. For more information about these interfaces, see: [Microsoft Hyper-V \(on page cccxxiv\)](#).
 -  **Important:** WinRM is the preferred communication interface. Before you add the VM Manager, configure WinRM on the Hyper-V host. For more information, see: [Configuring WinRM on Hyper-V hosts \(on page cccxxviii\)](#).
 - b. Provide the URL of the VM manager in the following format: `https://<Hyper-V_IP_address>/wsman`.
 - c. Optional: To share credentials with other hosts in the same cluster, select **Share credentials with other hosts in the same cluster**.
 - d. Provide the Administrator account credentials. Define the user as `user_name\domain` or `user_name@domain`. For example: `test\cluster.com` or `test@cluster.com`.
- To add KVM with RHEV-M:
 - a. Provide the URL of the VM manager in the following format:
 - For version 3.0: `https://<RHEV-M_IP_address>:8443/api`
 - For version 3.1 up to 3.6: `https://<RHEV-M_IP_address>/api`
 - For version 4.0: `https://<RHEV-M_IP_address>/ovirt-engine/api`
 - b. Provide the account credentials. Define the user as `user_name@domain`, for example: `test@cluster.com`.
- **9.2.17** To add Nutanix:
 - a. Provide the URL of the VM manager in the following format: `https://<NUTANIX_PRISM_IP_ADDRESS>:9440/PrismGateway/services/rest/v2.0`.
 - b. Provide the account credentials.
- To add VMware ESX, ESXi or vCenter:
 - a. Provide the URL of the VM manager in the following format:
 - For versions up to 5.1: `https://<vCenter_IP_address>/sdk`
 - For version 5.5 and higher: `https://<vCenter_IP_address>/sdk/vimService.wsdl`
 - b. Provide the account credentials. Define the user as `domain\user_name`, for example: `cluster.com\test`.
- **9.2.12** To add XenServer or Citrix XenServer:
 - a. Provide the URL of the VM manager in the following format: `https://xen-server`.
 - b. Provide the account credentials.

The following example shows a configured connection to VMware ESX, ESXi or vCenter.

Create VM Manager

VM Manager Type*

URL*

URL Template
 For versions lower than 5.5: <https://virtualcenter/sdk>
 For version 5.5 and higher: <https://virtualcenter/sdk/vimService.wsdl>

User Name*

Password*

4. Click **Create**. Connection to the VM manager is created and its status is Pending.

5. Select the VM manager, and click **Test Connection**.

The test might take a few minutes. Refresh the panel. If the VM manager connection is properly configured, the connection test status changes to Successful.

You added a VM manager. Its status remains Pending until the data is collected from the VM manager, uploaded to the BigFix server, and transferred to BigFix Inventory during the import of data.

The upload of data collected from VM managers is triggered shortly after you modify any VM manager, for example create a new connection or change the existing one. If you do not make any modifications, the data is uploaded according to the schedule, which is every 12 hours by default. For more information, see: [Capacity data flow \(on page cccxlv\)](#).

Collecting capacity data from virtualization hosts for Xen and KVM

For Citrix, Xen and KVM, connection to the VM manager does not have to be defined in BigFix Inventory to collect capacity data. The data can be collected directly from virtualization hosts by using the Run Capacity Scan on Virtualization Hosts task.

This method is applicable on the following virtualizations:

- **9.2.7** Xen Hypervisor
- **9.2.6** Citrix XenServer



Important: **9.2.12** Starting from application update 9.2.12, it is recommended to collect data from Citrix XenServer by configuring connections to VM managers. For more information, see: [Adding VM managers in central mode \(on page cccl\)](#).

- PowerKVM
- KVM x86

-  **Tip:** If your KVM x86 hosts are controlled by RHEV-M, you can configure the RHEV-M as a VM manager in BigFix Inventory. It allows for collecting capacity data from all controlled hosts without the need of using the Run Capacity Scan on Virtualization Hosts task on each of them. For more information, see: [Adding VM managers in central mode \(on page cccl\)](#).

For information about supported versions of these virtualizations, see: [System requirements \(on page c\)](#).

Ensure that virtualization hosts from which you want to collect capacity data meet the following prerequisites.

KVM hosts

Ensure that the following prerequisites are met on the KVM hosts:

- The host runs on Linux x86 or Linux on Power
- Bash shell is available
- `libvirt-client` library is installed (virsh command is available)
- `libxml2` library is installed (xmllint command is available)

Xen hosts

Ensure that the following prerequisites are met on the Xen hosts:

- The host runs on Linux x86
- Bash shell is available
- `xl` command is available
- `libxml2` library is installed (xmllint command is available)

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. In the upper right pane, select **Run Capacity Scan on Virtualization Hosts**.
4. **Optional:** To collect information about host names of virtualization hosts, select **Collect host names of virtualization hosts**.

This option affects the format of information about virtualization hosts that is displayed in the Server ID column on the reports. If you select this option, the column contains information about the vendor, type and host name of the host as specified in the **managedServerTagTemplate** parameter. If you do not select this option, the column contains the serial number of the host instead of its name. For more information about the **managedServerTagTemplate** parameter, see: [Advanced server settings \(on page cdii\)](#).

5. To start the scan, click **Take Action**.

-  **Important:** If you run this task on the host, you no longer need to run the **Run Capacity Scan and Upload Results** task on that host.

6. From the list of applicable computers, select the supported hosts.

 **Important:** The list is filtered out to the computers that meet the prerequisites. However, not all of them are virtualization hosts.

7. **Optional:** By default, the capacity scan is scheduled to run every 30 minutes. However, in an environment with many KVM or Xen hosts, consider lowering the frequency of the scan. To specify the frequency of the scan, open the **Execution** tab, specify the details, and click **OK**.

For environments with multiple hosts, run the scan once per one to six hours.

 **Restriction:** You can change the frequency after you obtain the acceptance of the BigFix Compliance Team.

8. Schedule the upload of scan results to the BigFix server.

a. In the upper right pane, select **Schedule VM Manager Tool Scan Results Upload**, specify the frequency, and click **Take Action**.

b. Select the hosts from which you collected the capacity data, and click **OK**.

 **Tip:** To check the status of the capacity scan on the endpoints, activate the analysis Status of Capacity Scan on Virtualization Hosts. If the status is other than OK, check the return code to learn what is the cause of the problem and how to solve it. For more information, see: [Return codes of capacity scans on virtualization hosts \(on page dcccx\)](#).

9. Wait for the scheduled import or run it manually to transfer the data to BigFix Inventory.

Related information

[Return codes of capacity scans on virtualization hosts \(on page dcccx\)](#)

Removing capacity scan data from the host

During the capacity scan, some files and folders are generated on the virtualization host. If you no longer collect capacity data from a particular host, remove the files and folders that were created by the capacity scan.

1. Stop the actions that were created by the **Run Capacity Scan on Virtualization Hosts** task on the specified endpoints. Otherwise, the scan continues running and the files are recreated.
 - a. In the navigation tree of the BigFix console, click **Sites > External Sites > IBM BigFix Inventory v9 > Actions**.
 - b. Select the actions, and in the lower pane, click **Stop**.
2. In the navigation tree, go to **Fixlets and Tasks**.

3. In the upper right pane, select **Remove Capacity Scan Data from Virtualization Hosts**, and click **Take Action**.
4. Select the computers from which you want to remove the data that was generated during the capacity scan, and click **OK**.

Identifying computers on public clouds

Available from 9.2.1. For all computers that run on public clouds, you must additionally identify them as such. This is required to apply the correct PVU rates, which might differ between the public clouds. It is enough to run a Fixlet against these computers and specify the type of the public cloud.

If a computer runs on a public cloud but is not identified as such, it has the *No VM Manager Data* or *No Scan Data* status on the Hardware Inventory report and its PVU value might be incorrectly calculated. To ensure that PVU is properly calculated, identify computers that run on public clouds and specify the types of clouds on which they run.

Public clouds that are currently supported by BigFix Inventory include:

- Amazon EC2
- IBM SoftLayer
- Microsoft Azure
- **9.2.8** Google Compute Engine

1. In the navigation tree of the BigFix console, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
2. In the upper right pane, select **Identify Computers on Public Clouds**.
3. Choose the type of the public cloud and click **Take Action**.



Note: If you incorrectly identified a computer as running on a public cloud, choose **Remove existing identification**.

4. Select all computers that run on this type of cloud and click **OK**.



Tip: To check which computers are identified as running on public clouds, activate the **Identified Public Cloud Computers** analysis and open the **Results** tab.

A computer that runs on a public cloud is assigned the number of PVUs that is specified in the [BigFix Eligible Public Cloud BYOSL policy](#). Information about the assigned number of PVUs is displayed on the Computers report after the next capacity scan and import of data. PVU consumption for the software that is installed on such a computer is properly calculated.

Advanced scan configuration

After you successfully schedule software and capacity scans and the scan data is displayed on the BigFix Inventory reports, you can further customize the scans. For example, you can exclude directories from software scans or scan remote shared file systems.

Detailed hardware scan

9.2.12 Available from 9.2.12. Detailed hardware scan collects hardware information related to memory, operating systems, storage, processors, partitions, network adapters, SMBIOS data, IP addresses and logical partition capacity data. The solution is supported on Windows, Linux x86 and AIX. It can be run also on other operating systems but the results might not be accurate.

In BigFix Inventory, there are two types of hardware scans: capacity scan and detailed hardware scan.

- Capacity scan collects the main information about the type and the capacity of processors in your infrastructure. BigFix Inventory uses this data to calculate software license usage for metrics that are based on processor capacity, such as PVU. BigFix Inventory displays the selected hardware attributes of scanned processors on the Hardware Inventory and Software Classification panels. You can also use API to retrieve the capacity details.
- Detailed hardware scan collects additional information about your hardware. The scan uses the same scanner technology as capacity scan. However, detailed hardware scan has a much broader spectrum and retrieves more hardware attributes. These details are not required to report license metrics. They can be used to monitor your hardware inventory or manage your assets. You can only retrieve the results of the detailed hardware scan through API.

For more information, see: [Software and hardware discovery \(on page lxiv\)](#).

Collecting detailed hardware scan information

9.2.12 Available from 9.2.12. **Linux** **Windows** You can collect the detailed hardware information by running the Collect Detailed Hardware Information task in the BigFix console. To retrieve the data use the relevant REST API.

Limitations

- Detailed hardware scan cannot be collected if you use the disconnected scan.
- The output of the detailed hardware scan cannot exceed 100 KB. This limit should be enough for most endpoints. However, if it is exceeded, the results will not be imported to BigFix Inventory server. Information about exceeding the limit of 100 KB is returned by the Detailed Hardware Information analysis.

Considerations

- If a piece of information is not applicable to your system, or cannot be retrieved, the detailed hardware scan returns no value, or the value equals "-1".
- In case of virtual systems, such as, VMware, Hyper-V or KVM guests, the scanner might not be able to retrieve the accurate value of some physical machine attributes. The physical system information is emulated by the hypervisor and the scanner cannot access the host system properties directly. The level of similarity between the virtual devices and the actual hardware depends on the virtualization technology used.

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > BigFix Inventory v9 > Fixlets and Tasks**.
3. Select **Collect Detailed Hardware Information**, and then in the lower pane, click **Take Action**.
4. Select computers from which you want to collect the detailed hardware information.
5. **Optional:** By default, the scan repeats every week. However, you can customize the frequency of the scan according to your needs. Select **Execution** tab and modify the **Reapply this action** section.
6. Click **OK**.
7. To populate the scan results to BigFix, activate the Detailed Hardware Information analysis.
8. Wait for the scheduled import, or run it manually.
9. To retrieve the detailed hardware information, use the relevant REST API call. For more information, see: [REST API for retrieving detailed hardware information \(v2\) \(on page mxliv\)](#).

The API call returns the detailed hardware information about the computers in your environment.

You can check the status and the date of the last scan in Detailed Hardware Information analysis.

Discovering software on shared disks

To discover software that is installed on shared disks, identify disks that exist in your infrastructure. Then, choose the method of scanning these disks that best suits your environment.

Software that is discovered on shared disks is reported on every computer on which the shared disk is mounted. It causes that license metric utilization that is calculated for the discovered software increases. When the software is reported on every computer on which the shared disk is mounted, it continues contributing to license metric utilization for the entire reporting period even if the shared disk is unmounted.

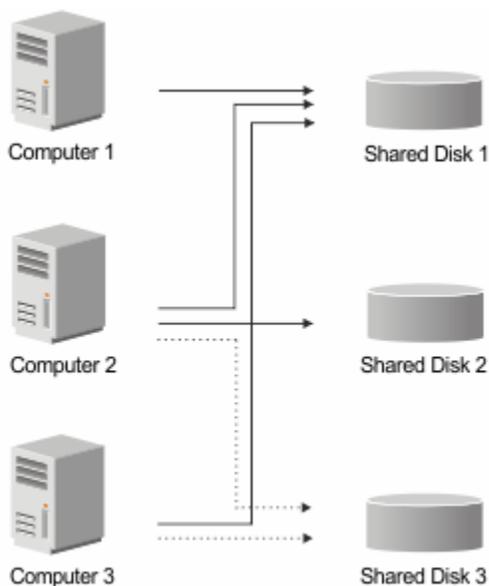
Scanned File Data

Requirements

- Scanning shared disks might significantly increase the number of scanned files per computer in your environment and, thus, affect the hardware requirements of your system. For more information, see: [Hardware requirements \(on page cvii\)](#).
- Discovery of shared disks is supported on Linux, AIX, Solaris, and HP-UX. It is not supported on Windows.
- Types of discovered shared disks include UNIX Network File System (NFS) and Windows Common Internet File System (CIFS).
- Types of software scans that are run on shared disks include catalog-based scan, file system scan, and software identification tags scan.
- If you have systems with multiple interfaces, ensure that you configure them correctly. For more information, see: [Configuring servers in separate networks \(on page clxxx\)](#).

Optimized mode

Optimized mode is advised for environments with a heavy use of shared disks in which a single shared disk is mounted on many computers.



There are two methods of scanning in this mode: automatic and manual. Automatic mode is recommended as it is easier to set up and maintain. However, this mode scans all shared disks that exist in your environment. If you want to scan only selected shared disks, use manual scanning.



Important: Do not use automatic and manual scanning of shared disks at the same time.

9.2.12 Automatic scanning

In this method, one of the computers on which a shared disk is mounted is automatically designated to scan the disk and discover the installed software. The data is then populated to all computers on which the same shared disk is mounted. Every disk is scanned starting from the top level through all the child-level directories. In case only a child-level directory is mounted on a computer, this directory is not scanned again. The data is populated based on the scan of the top-level directory. For detailed instructions, see: [Step 2a: Optimized mode - Automatic scanning of remote shared disks \(on page ccxxv\)](#).

9.2.8 Manual scanning

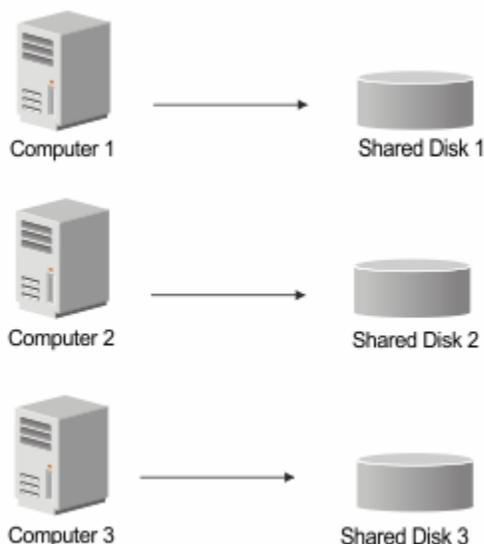
The outcome of this method is the same as of automatic scanning. However, it requires more manual setup. In this method, you create computer groups that contain only computers on which the same disk is mounted. Then, you scan one of the computers in that group to create a software template. Finally, you share the template between all computers on which the disk is mounted.

The choice of the scanned computer depends on whether some of the computer groups are overlapping and whether computers that access the same shared disk run on various operating systems. For

detailed instructions, see: [Step 2b: Optimized mode - Manual scanning of remote shared disks \(on page ccxxvii\)](#).

Basic mode

Basic mode is advised for environments in which a single shared disk is mounted on one or only few computers. In this mode, every discovered shared disk is scanned by every computer on which the shared disk is mounted. If many computers scan the same shared disk, performance of the disk might decrease. For detailed instructions, see: [Step 2c: Basic mode - Scanning remote shared disks \(on page ccxxxv\)](#).



Step 1: Discovering remote shared disks

To discover shared disks that are used in your infrastructure, run the Discover Remote Shared Disks task. Then, activate an analysis that gathers the retrieved information. After an import, the information is shown on the Shared Disks report. You can also retrieve it by using the `api/sam/v2/shared_disks` REST API.

1. Log in to the BigFix console.
2. Discover remote shared disks in your environment.
 - a. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
 - b. In the upper right pane, select **Discover Remote Shared Disks**, and then in the lower pane, click **Take Action**.
 - c. Select computers on which you want to discover shared disks, and click **OK**.

 **Tip:** Select all applicable computers to ensure that all shared disks are discovered.

3. Gather information about the discovered shared disks.

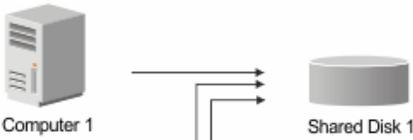
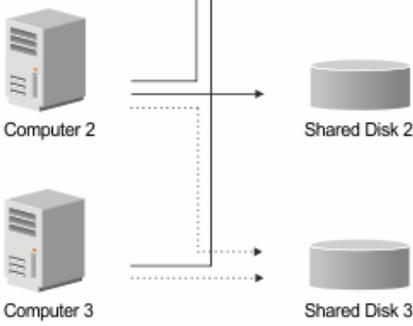
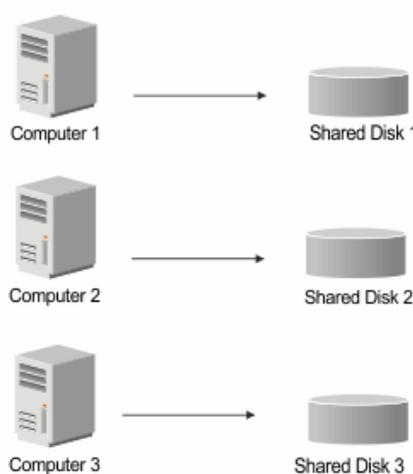
- a. In the navigation tree, click **Analyses**.
 - b. Right-click **Shared Disk Information**, and click **Activate**.
4. Wait for the scheduled import or run it manually.

Information about shared disks is shown on the Shared Disks report. You can also retrieve it by using REST API. For more information, see: [Retrieval of shared disks \(v2\) \(on page mxxxvii\)](#).

To view information about software that is installed on the discovered shared disks, scan the disks. Choose scanning method that best fits your use of shared disks.

Table 30. Methods of scanning remote shared disks

The table consists of three columns and three rows.

Method	Scenario	Diagram
9.2.12 Opti- mized mode - auto- matic scanning	A single shared disk is mounted on many computers. For more information, see: Step 2a: Optimized mode - Automatic scanning of remote shared disks (on page ccxxv) .	
9.2.8 Opti- mized mode - manual scanning	A single shared disk is mounted on many computers. For more information, see: Step 2b: Optimized mode - Manual scanning of remote shared disks (on page ccxxvii) . ! Important: Using automatic scanning is recommended in this scenario.	
Basic mode	A single shared disk is mounted on one or only few computers. For more information, see: Step 2c: Basic mode - Scanning remote shared disks (on page ccxxxv) .	

Step 2a: Optimized mode - Automatic scanning of remote shared disks

9.2.12 Available from 9.2.12. To discover software that is installed on shared disks in your infrastructure, enable automatic scanning of shared disks. As a result, a single computer is designated to scan a specific shared disk and discover the installed software. The data is then automatically populated to all computers on which the same shared disk is mounted. Use this mode when a single shared disk is mounted on many computers.

- This method of scanning is advised for environments with a heavy use of shared disks in which a single shared disk is mounted on many computers. For information, see: [Discovering software on shared disks \(on page ccxxi\)](#).
- This method is not supported in environments in which multiple instances of BigFix Inventory use a single BigFix server as the data source.
-  You must be an Administrator to perform this task.
- Ensure that the parameter **enable_automatic_task_deployment** is set to true on the **Advanced Server Settings** panel. It is the default setting.

1. [Discover shared disks that exist in your infrastructure. \(on page ccxxiii\)](#)
2. Enable optimized scan of shared disks.

- a. Log in to BigFix Inventory and go to **Management > Advanced Server Settings**.
- b. Change the value of the **autoscan_shared_disks_enable** parameter to true.



Note: If you disable this parameter, and then want to enable it back, you need to wait at least one day. Alternatively, you can log in to the BigFix console and remove the **LMT_Autoscan_Shared_Disks** property from all computers on which you ran the scan of shared disks.

3. Wait for the scheduled import or run it manually.

During the import, a single computer is designated to scan a specific shared disk. An action that is called Optimized Shared Disks Scan Update Resources List is created on the BigFix server. Each action represents designation of a single computer. Additionally, the software scan is triggered on every designated computer.



Important: Because these actions are triggered during the first import after you enable the scanning, it might take a couple of consecutive imports to discover software that is installed on shared disks.

To see software that was discovered on shared disks, go to **Reports > Software Classification**. Hover over the

Manage Report View icon , and click **Configure View**. Then, add the following filter: `From Software Template, equal to, yes`.

Publisher ...	Component Name	Version	Product Name	Metric	Computer N...	From Sc...	Installation Path	Details
IBM	IBM DB2 Enterprise Server...	10.5	IBM DB2 Enterprise Server...	RVU MAPC	NC050232	Yes	192.0.2.21:/file_server/shared/...	DETAILS >

When you click the link in the Details column, you can see on which shared disk the software is installed.

IBM Database Enterprise Developer Edition from 192.0.2.21:/file_server/shared

<p>IBM Database Enterprise Developer Edition</p> <p>Publisher IBM</p> <p>Product IBM DB2 Developer Edition</p> <p>Version IBM DB2 Developer Edition 10</p> <p>Release IBM DB2 Developer Edition 10.5</p> <p>Component IBM Database Enterprise Developer Edition 10.5</p> <p>GUID A02F1813-B524-41A7-894F-1336F6306D10</p> <p>IBM Database Enterprise Developer Edition was detected because the following condition was met:</p> <p>The signature was last modified at 12/18/2013 06:16 PM</p> <p>A computer has:</p> <ul style="list-style-type: none"> ✔ A signature that uses: File Signature <pre style="font-size: 8px; margin-top: 5px;"> <MultipleInstance><Iterator name="INSTALL_PATH"><FindFilePathEx name="\IBM_Database_Enterprise_Developer_Edition-10.5.0.swtag"/></Iterator><Instance><Variable name="IS_INSTALLED" export="true" /><FileExists absolutePath="\\$(INSTALL_PATH)IBM_Database_Enterprise_Developer_Edition-10.5.0.swtag"/></Variable><Condition withVariable="IS_INSTALLED"><Action do="SKIP" onValue="false" /></Condition></Instance></MultipleInstance></pre> ✔ Discovered matches: <ul style="list-style-type: none"> Installation Path 192.0.2.21:/file_server/shared/opt/ibm/db2/V10.5/properties/version 	<p>Application Usage</p> <p>Scan Template Name 192.0.2.21:/file_server/shared</p> <p>Usage data is not supported.</p>
--	--

Switching to automatic scanning of shared disks

To switch from manual scanning of shared disks to automatic scanning, set up the automatic scans. Then, wait for automatic scans to finish successfully to ensure data consistency. Finally, remove computer groups that were created to report software that is discovered on the shared disks.

1. [Enable automatic scanning of shared disks. \(on page ccxxv\)](#)

To ensure data consistency, wait for the results of the automatic scanning. To check the status of the scan, go to **Reports > Shared Disks**, and check the value in the Automatic Scan Status column. The value should be **OK** for every shared disk that was scanned by using the manual method.

2. In BigFix Inventory, remove computer groups that were created to report software that is discovered on shared disks.

- a. Go to **Management > Computer Groups**.
- b. Select the computer group that you want to remove, and click **Delete**.
- c. To confirm the removal, click **Delete**.

3. **Optional:** In the BigFix console, remove computer groups that contain computers on which a particular shared disk is mounted.

- a. In the navigation tree, go to **Computer Groups**.
- b. Right-click the computer group that you want to remove, and click **Remove**.
- c. To confirm the deletion, click **OK**.

Step 2b: Optimized mode - Manual scanning of remote shared disks

9.2.8 Available from 9.2.8. To discover software that is installed on shared disks in your infrastructure, create computer groups that contain only computers on which the same disk is mounted. Then, scan one of the computers in that group to create a software template. Finally, share the template between all computers on which the disk is mounted.

- This method of scanning is advised for environments with a heavy use of shared disks in which a single shared disk is mounted on many computers.



Important: **9.2.12** Starting from application updated 9.2.12, this method can be automated. Using the automated scanning is recommended as it is easier to set up and maintain. The difference is that automatic mode scans all shared disks that exist in your environment. If you want to scan only selected shared disks, use manual scanning. For more information, see: [Discovering software on shared disks \(on page ccxxi\)](#). If you already set up manual scanning, you can switch to automatic scanning. For more information, see: [Switching to automatic scanning of shared disks \(on page ccxxvi\)](#).

-  You must be an Administrator to perform this task.

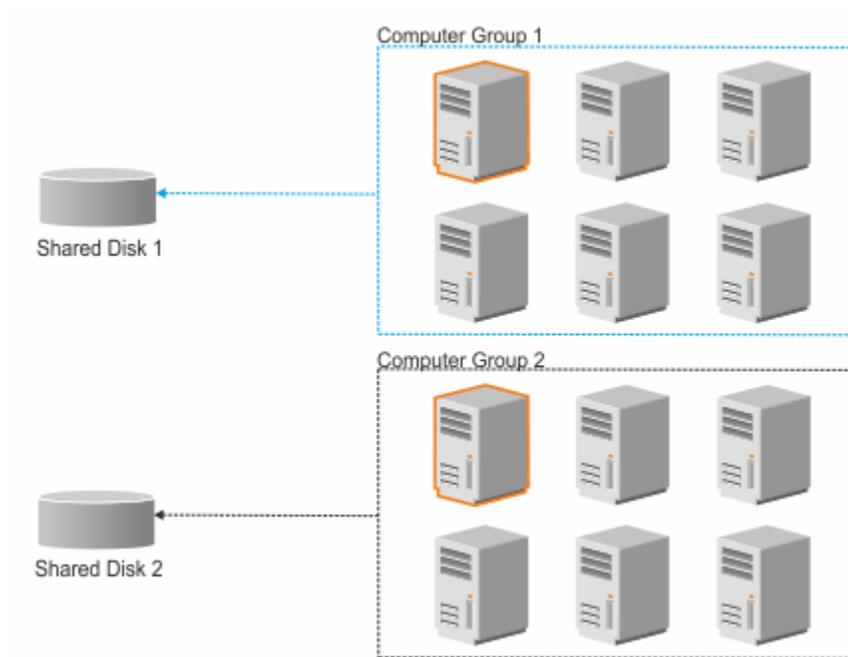
The choice of the scanned computer depends on whether some of the computer groups are overlapping and whether computers that access the same shared disk run on various operating systems.

Scenario 1: Each computer has access to one shared disk

There are two shared disks, each mounted on six computers. To properly report the discovered software, create two computer groups in BigFix Inventory:

- Computer Group 1 that contains computers on which Shared Disk 1 is mounted
- Computer Group 2 that contains computers on which Shared Disk 2 is mounted

From each group, select one computer to be scanned. Scan results from the designated computer are propagated to the remaining computers in the computer group. Thus, software that is installed on Shared Disk 1 is reported on all computers in Computer Group 1 without the need to scan all computers in that group. The same happens for computers in Computer Group 2.

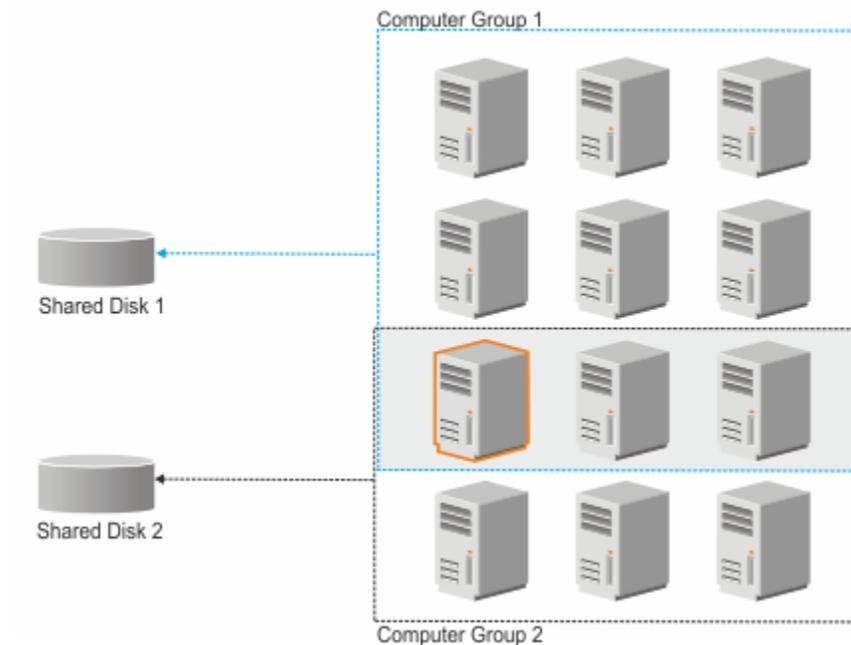


Scenario 2: Some computers have access to two shared disks

There are two shared disks. Shared Disk 1 is mounted on nine computers. Shared Disk 2 is mounted on six computers. To properly report the discovered software, create two computer groups in BigFix Inventory:

- Computer Group 1 that contains computers on which Shared Disk 1 is mounted
- Computer Group 2 that contains computers on which Shared Disk 2 is mounted

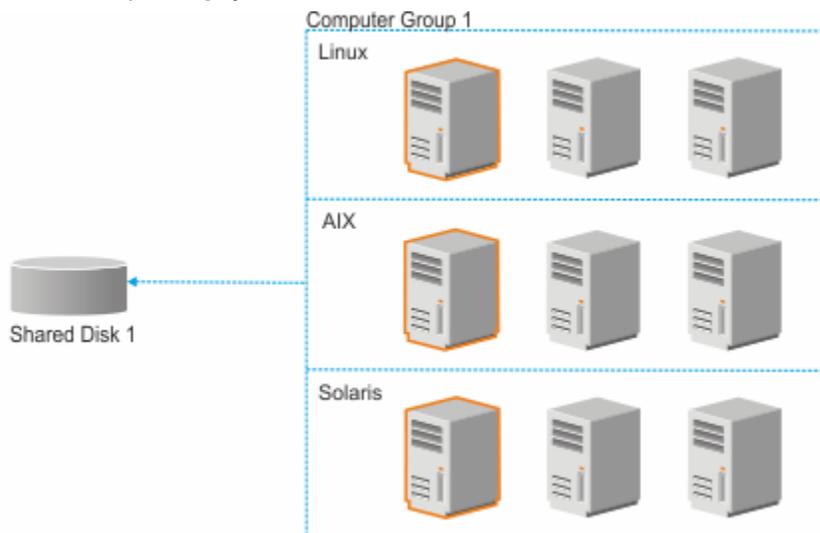
Three of the computers belong to both computer groups because they have both disks mounted. Select one of these computers to be scanned. This way, only one computer is scanned to discover software that is installed on both shared disks. Scan results from the designated computer are propagated to the rest of computers in both groups. Computers on which only one disk is mounted show software from that disk only. Computers on which both disks are mounted show software from both disks.



Scenario 3: Computers have access to one shared disk but run on different operating systems

There is one shared disk that is mounted on nine computers. Three of the computers run on Linux, three on AIX, and three on Solaris. To properly report the discovered software, create one computer group in BigFix Inventory. From this group, designate three computers to be scanned, one for every operating system.

It is necessary because the software catalog that is propagated to each computer monitored by BigFix Inventory contains software signatures specific to the operating system on which that computer runs. Thus, scan results from the shared disk must be matched against software catalog for every operating system. Scan results from each designated computer are propagated to the rest of computers that run on the same operating system.



1. [Discover shared disks that exist in your infrastructure. \(on page ccxxiii\)](#)
2. Create a list of unique access points of the discovered shared disks. The information is needed to group computers on which the same shared disk is mounted.

- a. In the navigation tree, click **Analyses**.
- b. In the upper-right pane, right-click **Shared Disks Information**, and click **Activate**.
- c. Open the **Results** tab. The Computer Name column lists computers on which shared disks are mounted. The Remote Shared Disks column provides information about the type of the mounted shared disk, the access point, and the path under which the disk is mounted on a particular computer.

Computer Name	Remote Shared Disks
 NC050232	network file system, 192.0.2.21:/file_server/shared, /mnt/nfs
 nc053026	<multiple results>

- d. Create a list of unique access points.
 - If every computer has only one disk mounted, right-click the results of the analysis, and click **Copy Text**. Then, paste the list to a text editor, and edit the file so that it contains only information about access points. For example:

```
192.0.2.21:/file_server/shared
```

- If a computer has multiple disks mounted, double-click the computer, and scroll down to the **Shared Disk Information** section.

▼ Shared Disk Information	
Remote Shared Disks	network file system, 192.0.2.21:/file_server/shared, /mnt/nfs1
	network file system, 192.0.2.22:/file_server/shared, /mnt/nfs2
	network file system, 192.0.2.23:/file_server/shared, /mnt/nfs3

Copy the text, paste it to a text editor, and edit the file so that it contains only information about access points. For example:

```
192.0.2.21:/file_server/shared
192.0.2.22:/file_server/shared
192.0.2.23:/file_server/shared
```

3. Create a BigFix computer group that contains computers on which a particular shared disk is mounted. The group is used to easily designate one computer that can be scanned to create a software template.

- a. To make grouping computers easier, create a computer property based on the information about the access point of the mounted shared disk.
 - i. In the top navigation, click **Tools > Manage Properties**.
 - ii. From the list of available properties, select **Remote Shared Disks**, and click **Make Custom Copy**. Then, click **OK**.

- b. Create a computer group that contains all computers on which the shared disk is mounted.
 - i. In the top navigation, click **Tools > Create New Automatic Computer Group**.
 - ii. Specify the name of the computer group. For example, Shared Disk 1.
 - iii. Specify the following conditions: `Remote Shared Disks, contains, <IP:shared_disk>`.

Include computers with the following property:

Remote Shared Disks	contains	192.0.2.21:/file_server/shared	-	+
---------------------	----------	--------------------------------	---	---

If a shared disk is mounted based on the IP address on some of the computers and based on the host name on other computers, BigFix Inventory does not unify addresses of access points. Thus, add the following condition to the definition of the computer group: `Remote Shared Disks, contains, <host_name:shared_disk>`. Alternatively, create two separate computer groups: one in which the access point based on the IP address and the other in which it is based on the host name.

Include computers with **any** of the following properties:

Remote Shared Disks	contains	192.0.2.21:/file_server/shared	-	+
Remote Shared Disks	contains	nc1920221:/file_server/shared	-	+

If the same shared disk is mounted on computers that run on different types of operating systems, create a computer group per each type. To create such a group, add the following condition to the definition of the computer group: `OS, contains, <operating_system>`, where `<operating_system>` can have the following values:

- linux
- aix
- hp-ux
- sunos

For more information, see scenario 3 in [Discovering software on shared disks \(on page ccxxi\)](#).

Include computers with **all** of the following properties:

Remote Shared Disks	contains	192.0.2.21:/file_server/shared	-	+
OS	contains	lin	-	+

4. From the computer group, choose one computer to be scanned. Information about software discovered on the shared disk that is mounted on this computer is used as a software template.

- a. In the navigation tree, click **Fixlets and Tasks**.
- b. In the upper right pane, select **Initiate Software Scan on Shared Disks**.
- c. **Optional:** If the computer has multiple shared disks mounted and you want to scan only some of them, specify URLs of the access points that you want to scan. Otherwise, all mounted shared disks are scanned.

You can enter each URL from a separate line or separate the URLs with a semicolon (;).

 **Tip:** Ensure that the URLs are exactly the same as the URLs returned by the Shared Disks Information analysis.

Disks to scan

Use the following field to enter access points of disks that you want to scan. For the list of remote disks check the output of the analysis: [Shared Disk Information](#). You can either enter each URL on a separate line, or separate them with a semicolon (;). If this text area is left empty, all disks currently mounted will be scanned.

192.0.2.21:/file_server/shared

- d. **Optional:** To limit the amount of processor resources that the scan consumes, select **Initiate the software scan with CPU threshold**. Specify the consumption limit that is in the range 5 - 100. The higher value you specify, the higher is the consumption limit.
For example, if you specify 75, scanner processes use the average of 75% of the processing power of the target computer.

 **Important:** Setting the threshold does not guarantee that the CPU consumption is always below the specified value. It fluctuates around that value, sometimes exceeding it and sometimes dropping below it. Temporary peaks are expected.

- e. To start the scan, click **Take Action**. Then, expand **Applicable Computers > By Group**, and select the group that you created in step 3 ([on page ccxxx](#)). From the list of computers that belong to this group, select one computer to be scanned, and click **OK**.

 **Tip:** If one of the computers has many shared disks mounted, scan this computer. This way, only one computer is scanned to discover software that is installed on multiple shared disks. For more information, see scenario 2 in [Discovering software on shared disks \(on page ccxxi\)](#).

- f. To see the status of the scan, activate the **Status of Shared Disks Software Scan** analysis, and open the **Results** tab.

If any of the columns shows `<multiple results>`, it means that multiple shared disks are mounted on the scanned computer. To see the status of the scan for each shared disk, hover over the results in the column specific to the type of the scan. The number next to the status corresponds to the number of the shared disk in the Scanned Disks column. You can also double-click the computer, and scroll down to the **Status of Shared Disks Software Scan** section.

Computer Na...	Mounted Disks	Scanned Disks	Time of Last S...	Status of Catalog-based Scan
 nc053026	<multiple results>	<multiple resu...	<multiple resu...	<multiple results>

1: OK
 2: OK
 3: OK

After the scan finishes, a software template is created.

5. If you have not previously scheduled upload of scan results from the scanned computer, schedule it now. Ensure that the upload is scheduled only once per computer.
 - a. In the navigation tree, click **Fixlets and Tasks**.
 - b. In the upper right pane, select **Upload Software Scan Results**, and click **Take Action**.
 - c. Select the computer that is scanned to create the software template, and click **OK**.
6. To make the software templates available in BigFix Inventory, wait for the scheduled import or run it manually.
7. Create a computer group in BigFix Inventory. The group is used to report software that is discovered on the shared disk on all computers that belong to this computer group.
 - a. Log in to BigFix Inventory, and click **Management > Computer Groups**.
 - b. To create a computer group, click **New**.
 - c. Provide the name and description of the computer group.

i

Tip: To make it easier to correlate computer groups that are created in the BigFix console and in BigFix Inventory, use the same name in both cases.
 - d. Specify filters according to which computers are assigned to the group, and click **Create**. To create a group that is based on a computer group from the BigFix console, choose `Data Source Groups, in set,` and select the group that you created in step 3 (on page ccxxx).
 - e. In the **Type** section, select **Software Template**.
 - f. Select the software template that you want to share among all computers that belong to this group. Then, click **Create**.



! **Important:** Ensure that you assign a software template to a group in which all computers have the shared disk mounted. Otherwise, software discovered on the shared disk is reported on computers that do not in fact have access to that software.

To see software that was discovered on shared disks, go to **Reports > Software Classification**. Hover over the **Manage Report View** icon , and click **Configure View**. Then, and add the following filter: From Software Template, equal to, yes.

Publisher ...	Component Name	Version	Product Name	Metric	Computer N...	From Sc...	Installation Path	Details
IBM	IBM DB2 Enterprise Server...	10.5	IBM DB2 Enterprise Server...	RVU MAPC	NC050232	Yes	192.0.2.21:/file_server/shared/...	DETAILS >

When you click the link in the Details column, you can see on which shared disk the software is installed.

IBM Database Enterprise Developer Edition from 192.0.2.21:/file_server/shared

<p>IBM Database Enterprise Developer Edition</p> <p>Publisher IBM</p> <p>Product IBM DB2 Developer Edition</p> <p>Version IBM DB2 Developer Edition 10</p> <p>Release IBM DB2 Developer Edition 10.5</p> <p>Component IBM Database Enterprise Developer Edition 10.5</p> <p>GUID A02F1813-B524-41A7-894F-1336F6306D10</p> <p>IBM Database Enterprise Developer Edition was detected because the following condition was met:</p> <p><u>The signature was last modified at 12/18/2013 06:16 PM</u></p> <p>A computer has:</p> <ul style="list-style-type: none"> ✔ A signature that uses: File Signature <pre style="font-family: monospace; font-size: 0.8em; margin-top: 5px;"> *<MultipleInstance><Iterator name="INSTALL_PATH"><FindFilePathEx name="IBM_Database_Enterprise_Developer_Edition-10.5.0.swtag"/></Iterator><Instance><Variable name="IS_INSTALLED" export="true" /><FileExists absolutePath="{(INSTALL_PATH)}IBM_Database_Enterprise_Developer_Edition-10.5.0.swtag"/><Variable><Condition withVariable="IS_INSTALLED"><Action do="SKIP" onValue="false"/></Condition></Instance></MultipleInstance></pre> ✔ Discovered matches: <ul style="list-style-type: none"> Installation Path: 192.0.2.21:/file_server/shared/opt/ibm/db2/V10.5/properties/version 	<p>Application Usage</p> <p>Scan Template Name 192.0.2.21:/file_server/shared</p> <p>Usage data is not supported.</p>
---	--

Changing the computer scanned to discover software on shared disks

9.2.8 Available from 9.2.8. If you designated a computer to be scanned to create a software template for shared disks, and you want to designate a different computer, start by scanning the new computer. After the computer is scanned, and scan results are uploaded to BigFix Inventory, stop the scan on the computer that you initially designated. Thanks to this approach, you ensure continuity of software discovery on shared disks.

1. Select the computer that you want to scan to create the new software template.
 - a. Log in to the BigFix console.
 - b. In the navigation tree, click **Sites > External Sites > BigFix Inventory v9 > Fixlets and Tasks**.
 - c. In the upper right pane, select **Initiate Software Scan on Shared Disks**, and click **Take Action**.
 - d. Select the new computer that you want to scan, and click **OK**.

 **Tip:** If you have the computers grouped, expand **Applicable Computers > By Group**, and select a group. Then, from the list of computers that belong to this group, select one computer to be scanned.

2. After the scan finishes, wait for the scheduled import of data or run it manually.
3. Select the new software template for the computer group that uses the template created from the initially designated computer.
 - a. Log in to BigFix Inventory, and go to **Management > Computer Groups**.
 - b. Open the computer group that used the old software template.
 - c. Select the new software template that you want to share among all computers that belong to this group. Then, click **Create**.
4. Stop the shared disks scan that is running on the computer that you initially designated.
 - a. Log in to the BigFix console.
 - b. In the navigation tree, click **Actions**.
 - c. Look for the action called **Initiate Software Scan on Shared Disks** that is running on the computer that you want to stop scanning.
 - d. Right-click the action, and click **Stop Action**. Then, click **OK** to confirm.

Step 2c: Basic mode - Scanning remote shared disks

To discover software that is installed on shared disks in your infrastructure, run the software scan with the **Scan remote shared disks** option enabled.

- This method of scanning is advised for environments in which a single shared disk is mounted on one or only few computers. For information, see: [Discovering software on shared disks \(on page ccxxi\)](#).
- If more computers are configured to scan a shared disk, performance of the disk might decrease drastically and the scanning might take a long time. It is recommended to schedule the clients to scan the shared disk at different times, for example to spread the scanning across many days of the week.
- If the shared disks to be scanned are large, scan timeout problems might occur.

1. [Discover shared disks that exist in your infrastructure. \(on page ccxxiii\)](#)
2. Exclude unsupported file systems from scanning.

- a. In the navigation tree, click **Analyses**.
 - b. Click **Shared Disk Information**, and open the **Results** tab to recognize file systems that must be excluded from scanning.
 - c. In the navigation tree, click **Fixlets and Tasks** and select **Add Excluded Directories**. Specify directories or mount points of unsupported file systems, and click **Take Action**. For more information, see: [Adding excluded directories \(on page ccxliv\)](#).
-  **Tip:** A mount point is a directory that is at the top of the file system hierarchy. You can verify whether a directory is a mount point by checking the output of the mount or mountpoint commands.
- d. Select computers on which you want to exclude the unsupported file systems and click **OK**.
3. Stop the current software scans.
 - a. In the navigation tree, click **Actions**.
 - b. In the upper right pane, select the **Initiate Software Scan** action.
 - c. In the lower pane, click **Stop**.
 4. [Initiate software scans \(on page cciii\)](#) with the option to scan remote shared disks selected.

 **Important:** Automounted remote disks will be scanned only if they are mounted during the scan.

Information about shared disks and the installed software is available on the following reports:

- Computers report (when you access it from the widget on the home page) contains information about the type of file system, shared disk IP address, and mount point
- Scanned File Data report contains information about the file path
- Software Classification report contains information about the installation path

9.2.6 Information about scanned files and directories in your file system is stored in the scanner cache folder. By knowing the hierarchy of files, the scanner can locate them quicker, which results in shorter scans. The amount of disk space that is needed for the cache depends on the number of files that are being scanned. If the current location of the cache folder cannot ensure sufficient disk space, you can change the location of the cache folder or optimize the cache. For more information, see: [Optimizing scanner cache configuration \(on page cclxvi\)](#).

Discovering software in Docker containers

9.2.5 Available from 9.2.5. Docker is a platform that allows for automating the deployment of applications inside software containers. BigFix Inventory discovers software that is installed inside Docker containers. It also measures license metric utilization of the discovered BigFix products.

Requirements

BigFix Inventory discovers software that is installed in Docker containers on condition that:

- Only one Docker engine is deployed on the host computer.
- The Docker container is deployed on one of the following platforms:
 - Red Hat Enterprise Linux 7 for x86
 - Red Hat Enterprise Linux 7 for BigFix 64-bit)
 - SUSE Linux 12 for x86
- The Docker container is running.
- The BigFix client is installed on the host computer.
- Scans and uploads of their results are enabled on the host computer.
- Software that is installed in the Docker container delivers software ID tags.
- To ensure proper discovery of software on Docker containers, the content of a Docker container cannot change throughout its lifecycle.

9.2.10

Starting from application update 9.2.10, BigFix Inventory additionally supports software discovery on Docker containers under the Red Hat® OpenShift container application platform.

Viewing software

Software that is installed in Docker containers can be viewed on the Software Installations report. It is presented under the host computer. To learn why the software was discovered, click **Details**.

IBM SPSS Statistics Base was detected because the following condition was met:

Software identification tag found

1

✔ Software identification tag `/var/opt/BESClient/LMT/CIT/docker/containers/0d4a4099ceb3/swidtags-python/2429/`

2 `ibm.com_IBM_SPSS_Statistics_Base-24.0.0.swidtag` with following data:

Publisher Name: IBM

Name: IBM SPSS Statistics Base

Version: 24.0.0

Unique ID: `ibm.com-1f1bf9d5371243de8cd0971c93dfc274-24.0.0`

Tag Creator Name: IBM

Licenser Name: IBM

Tag Version: 0

Detected on: 06/28/2016 04:20 PM

The details include, among other, information about:

1

Container on which the software was detected.

2

Software ID tag that caused the detection.

Measuring license metric utilization

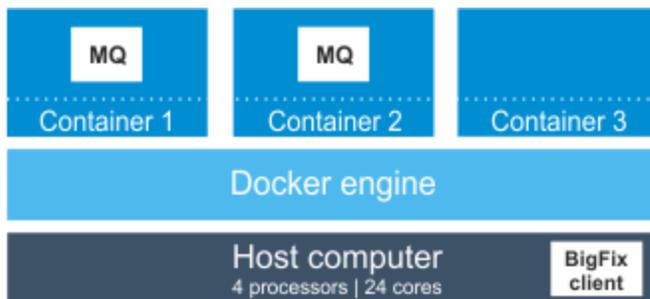
Apart from discovering software that is installed in Docker containers, BigFix Inventory also reports license metric utilization of the discovered BigFix products. When the Docker is deployed on a physical host, license metric utilization is calculated on the level of the host. When it is deployed on a virtual machine, utilization is calculated on the level of the virtual machine. For more details, see the following scenarios.

! **Important:** Docker is not a subcapacity eligible virtualization but it can be used in combination with a subcapacity eligible virtualization. The scenarios show how utilization of PVU and RVU MAPC is calculated. Utilization of other reported metrics is calculated in an analogical way.

Scenario 1: Docker deployed on a physical server

When the Docker engine is deployed directly on a physical server, PVU and RVU MAPC utilization is measured on the level of the host computer.

Example: Three containers are deployed on a physical server that has four Intel Xeon 3400 processors, each with six cores. It gives 24 cores in total. IBM MQ is installed in two out of three containers. BigFix Inventory counts PVU and RVU MAPC utilization on the level of the host computer.



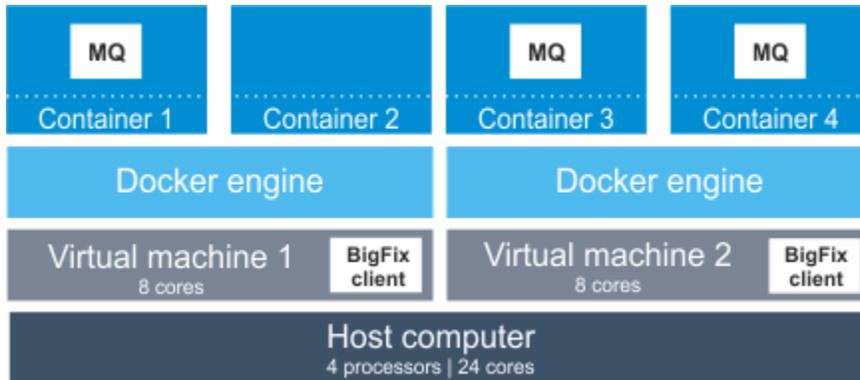
In this case, IBM MQ has access to 24 cores. According to the PVU table, when the server has four sockets, this processor model is assigned 100 PVUs per core. Thus, PVU utilization for IBM MQ equals 2400 PVUs. The value would be the same if another instance of IBM MQ was installed in the third container.

Scenario 2: Docker deployed on a virtual machine

When the Docker engine is deployed on a virtual machine, PVU and RVU MAPC utilization is counted as the highest number of PVUs that are available for the virtual machine.

Example: Two virtual machines are installed on a physical server that has four Intel Xeon 3400 processors, each with six cores. It gives 24 cores in total. Each virtual machine is assigned eight cores and has two containers deployed. IBM MQ is installed:

- In one container on the first virtual machine
- In two containers on the second virtual machine



In this case, IBM MQ that is installed on each of the virtual machines has access to eight cores. In total, it has access to 16 cores out of 24 cores that are available on the physical computer. According to the PVU table, when the server has four sockets, this processor model is assigned 100 PVUs per core. Thus, PVU utilization for IBM MQ equals 1600 PVUs. If the Docker engine was deployed directly on the physical server, IBM MQ would have access to 24 cores and its PVU utilization would equal 2400 PVUs.

Logs

To troubleshoot problems with discovery of software that is installed in Docker containers, see the [docker_scan.log](#) log. The log is stored in the BigFix client installation directory. By default, it is:

```
Linux var/opt/BesClient/LMT/CIT/docker_scan.log
Windows C:\Program Files (x86)\BigFix Enterprise\BESClient\LMT\CIT
\docker_scan.log
```

Concept

[Collecting the hardware information and capacity data \(on page ccx\)](#)

Configuring scans on Docker containers

9.2.5 Available from 9.2.5. In some Docker environments, you might need to perform additional steps to specify a non-default installation path, or to exclude directories from scanning.



Note: To check whether the Docker is installed in the default installation path, run the following command.

```
$ docker version
```

If the result of the command is a Docker version, the Docker is installed in the default installation path. Any other outcome indicates that the Docker is installed in a non-default path.

Specifying a non-default installation path for Docker

If Docker is installed in a non-default path, add this path as a setting of the BigFix client, so that the software can be successfully discovered.

1. Log in to the BigFix console, and click **Computer Management > Computers**.
2. Right-click on the computer that has the Docker installed, and click **Edit Computer Settings**.
3. Add a computer setting. Specify the name as `DOCKER_EXEC`, and provide an absolute path as the value, for example `/usr/bin/docker`.

Specifying additional command options

By default, the scan runs the Docker command without any options. If you want to use additional options provided by Docker, for example `-H` (daemon socket to connect to), add these options as a new setting of the BigFix client. Enter all options in one setting.

1. Log in to the BigFix console, and click **Computer Management > Computers**.
2. Right-click on the computer that has the Docker installed, and click **Edit Computer Settings**.
3. Add a computer setting. Specify the name as `DOCKER_OPTS`, and provide options as the value, for example `-H unix:///var/run/docker.sock`.

Excluding directories from scans

The default Docker file system directory `/var/lib/docker` is excluded from scanning.

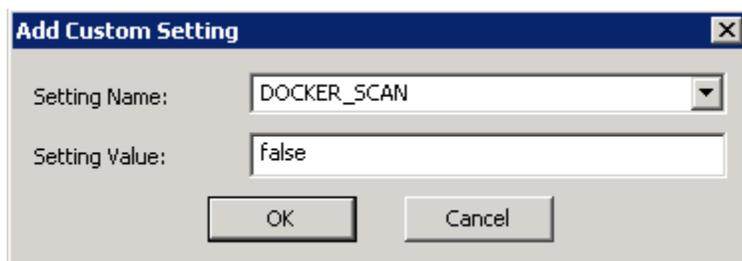
If you change the Docker file system directory to a custom directory, you need to manually exclude it from scanning because it might cause duplicated discoveries. For more information, see: [Excluding directories \(on page ccxli\)](#).

Disabling scans on Docker containers

9.2.5 Available from 9.2.5. By default, BigFix Inventory scans all Docker containers that are deployed on computers where the BigFix client is installed. If you do not want to scan the containers but still want to monitor the host computer, change the value of the **DOCKER_SCAN** parameter on the host computer.

You can disable software discovery on all containers that are deployed on a host computer. You cannot disable it on a subset of containers only.

1. Log in to the BigFix console, and click **Computers**.
2. Select the host computer on which Docker containers are deployed, and click **Edit Settings**.
3. Click **Add**. Specify **DOCKER_SCAN** as the setting name, and `false` as the setting value. Then, click **OK**.



Docker containers are no longer scanned. Scan results remain in the `<BES Client>/LMT/CIT/docker/containers` directory on the host computer but the directory itself is added to the list of excluded directories. Thus, the results are not transferred to BigFix Inventory.

! **Important:** The Docker file system directory `/var/lib/docker` might contain copies of software ID tags. When the Docker scan is enabled, the directory is excluded from scanning to avoid duplicated software discovery. When you disable the Docker scan, the directory is included back into regular scans.

If you want to re-enable the scans of Docker containers, change the value of the **DOCKER_SCAN** parameter to true.

Excluding directories from being scanned

Excluding some directories from scanning is useful if the directories are large and contain no information that is important to the software inventory. By excluding them, you can speed up the scanning process. Some directories are excluded from software scans by default. They must remain excluded to ensure accuracy of data. You can add or remove other directories from the list by using tasks in the BigFix console. You can also manually add them to the scanner files on particular endpoints.

i **Tip:** **9.2.5** Check how long the software scan lasts on each computer. If the scan is very long, it might indicate that some of the directories on this computer can be excluded from scanning. For example, because they contain software installation images or other file repositories not related to IBM software.

To check the scan length, activate the Software Scan Status analysis, and see the Total Time column.

Directories that are excluded by default

By default, the following directories are excluded from software scans. Ensure that you do not remove them from the list. It might lead to untested and unsupported results in BigFix Inventory.

AIX

```
/proc
*/tmp
*/eznim
*/unicore/lib/*
*/perl/lib/*
*/perl5/*
/usr/src/kernels/*
```

📝 **Note:** Starting from scanner version 2.8.0.5000, the `/usr/lpp` directory is no longer excluded from software scans. It might cause that more software components are discovered by BigFix Inventory.

Solaris

```
/proc
*/tmp
*/eznim
*/unicore/lib/*
*/perl/lib/*
*/perl5/*
/usr/src/kernels/*
/var/spool/pkg
```

UNIX

```
/proc
*/tmp
*/eznim
*/unicore/lib/*
*/perl/lib/*
*/perl5/*
/usr/src/kernels/*
```

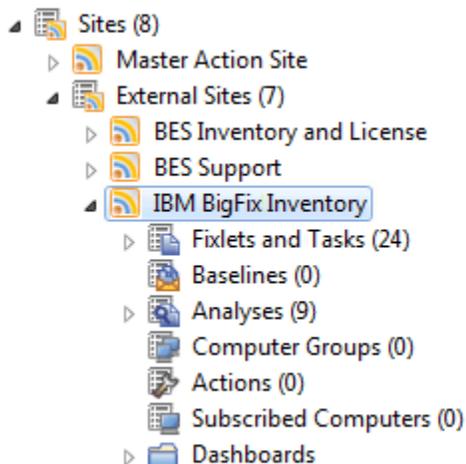
Windows

```
?:/System Volume Information
?:/$Recycle.Bin
?:/RECYCLER
%CSIDL_WINDOWS%/System32
%CSIDL_WINDOWS%/SysWOW64
%CSIDL_WINDOWS%/winsxs
%CSIDL_WINDOWS%/ServicePackFiles
%CSIDL_WINDOWS%/installer
%CSIDL_WINDOWS%/NtUninstall
%CSIDL_WINDOWS%/NtServicePackUninstall*$
%CSIDL_WINDOWS%/$hf_mig$
*/tmp
*/temp
*/cache/out-of-date
*/Temporary Internet Files
```

Retrieving the list of directories that are excluded from being scanned

To retrieve the list of all directories that are currently excluded from scanning, activate a relevant analysis in the BigFix console.

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9**, and then click **Analyses**.



3. Select the **Excluded Directories** analysis, right-click on it, and then click **Activate**. The analysis is activated and starts retrieving the list of directories.
4. To view the excluded directories, click the **Results** tab. Directories are divided according to endpoints. You can view the directories in three ways:
 - All directories printed on separate lines are listed in the **Excluded Directories** column. If an entry says `<multiple results>`, hover over it to view the complete list.
 - All directories printed on one line and separated with a semicolon (;) are listed in the **Excluded Directories (semicolon separated)** column.
 - You can also double-click on `<multiple results>` to view the summary. All directories are listed in the **Excluded Directories** entry.

Computer Name	Operating System	Excluded Directories	Excluded Directories (semicolon separated)
NC040203	Win2008R2 6.1.7600	<multiple results>	%CSIDL_WINDOWS%\\$NtServicePackUninst...
NC043019	Linux Red Hat Enterp...	<multiple results>	*/eznim;*/tmp;/proc
NC043024	Win2008R2 6.1.7600	<multiple results>	%CSIDL_WINDOWS%\\$NtServicePackUninst...
NC9128110031	Linux Red Hat Enterp...	<multiple results>	*/eznim;*/tmp;/proc;?;/tomek
NC9143126242	Win2008 6.0.6002	<multiple results>	%CSIDL_WINDOWS%\\$NtServicePackUninst...

You retrieved the list of all directories that are excluded from scanning. Whenever you add or remove a directory from the list, you can use this analysis to verify if the operation succeeded.

Excluding additional directories

You can exclude additional directories from being scanned by specifying them in the task and then running this task against the chosen endpoints. Excluding directories that do not contain any information relevant for software inventory can significantly speed up the scanning process.

Use ASCII characters to specify the names of directories that you want to exclude from scanning.

Some directories are excluded from being scanned by default. Ensure that you do not remove them from the list. It might lead to untested and unsupported results in BigFix Inventory. For the full list, see: [Directories that are excluded by default \(on page ccxli\)](#)

Using asterisks (*) and question marks (?) to specify the additional directories that should be excluded

You can use the asterisk (*) and question mark (?) wildcards to specify which directories to exclude from scanning. Remember about the following rules:

- Asterisks (*) and question marks (?) should be used according to their regular rules of usage while specifying excluded directories. An asterisk (*) might represent no characters, or a few consecutive characters. A question mark (?) represents a single character.
- The number of asterisks (*) and question marks (?) in a single path is not limited.
- Asterisks (*) and question marks (?) are the only wildcards that are supported. Other expressions, such as [a]* or [a-z] are not supported, and should not be used to specify the excluded directories.

Examples of properly specified directories.

Example 1. `?:/System Volume Information`

The `System Volume Information` directory located on any partition is excluded. For example, `c:/System Volume Information` and `d:/System Volume Information`.

Example 2. `*/backup*`

Any directory on a file system that begins with `backup` is excluded. For example, `/tmp/backup` and `/home/user1/backup20170101`.

Example 3.

```
/tmp
/tmp/
/tmp/*
```

The `/tmp` directory is excluded.



Note: It is important to place the wildcards in the right position. The following path patterns specify different directories.



`/tmp/*`

Only the `/tmp` directory is excluded.

`/tmp*`

The `/tmp` directory, as well as any directory starting with the `tmp` string are excluded. For example, `/tmp1`, `/tmp123`, and `/tmpABC`.

`*/tmp`

Any `/tmp` directory on file system is excluded. For example, `/tmp`, `/var/tmp`, and `/home/user/test/tmp`.

Example 4. `*per*lib6*all`

Any directory path that contains the `per` string which is followed by a string that contains `lib6` and ends with `all`. For example, `/usr/opt/per15/lib64/5.10.1/aix-thread-multi-64all` and `/usr/opt/per15/lib64/site_perl/5.10.1/aix-thread-multi-64all`.

However, according to the rules, this example does not cover the following path: `/usr/opt/per15/lib64/5.10.1/aix-thread-multi-64all2` as it does not end with the `all`.

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9**, and then click **Fixlets and Tasks**.
3. Select the **Add Excluded Directories** task.
4. Specify which directories are to be excluded from scanning. If you specify a particular directory, it is excluded along with all its sub-directories. **Windows** Use only slashes (/) in the directory path.

Excluded directories

Use the following field to enter directories that you want to exclude from scanning. You can either enter each directory on a separate line, or separate them with a semicolon (;).

```
*/eznim
*/tmp
*/proc
?:/System Volume Information
?:/$Recycle.Bin
```

5. Click **Take Action** and select endpoints for which you want to apply the changes.

You added new entries to the list of directories that are excluded from scanning.

Including the excluded directories back in scans

You can include a directory back in the software scan by specifying it in a task, and running this the task against the chosen endpoints.

Some directories are excluded from being scanned by default. Ensure that you do not remove them from the list. It might lead to untested and unsupported results in BigFix Inventory. For the full list, see: [Directories that are excluded by default \(on page ccxli\)](#).

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9**, and then click **Fixlets and Tasks**.
3. Select the **Remove Excluded Directories** task.
4. Specify which directories are to be removed from the list of excluded directories.

Excluded directories

Use the following field to enter directories that you want to remove from the list of excluded directories. You can either enter each directory on a separate line, or separate them with a semicolon (;).

```
*/eznim
*/proc
?:/System Volume Information
```

5. Click **Take Action** and select endpoints for which you want to apply the changes.

You removed the entries from the list of directories that are excluded from scanning. Those directories are now scanned during the software scan.

Manually excluding directories

After you install the scanner, you can specify which directories are to be excluded from scanning during the raw scan of the file system.



Important: Use ASCII characters to specify the names of directories that you want to exclude from scanning.

You specify those directories by adding paths to the `exclude_path.txt` file that is in the `<BES Client>LMT/CIT` directory. Each path must be added on a separate line. For more information about modifying the file, see: [Excluding additional directories \(on page ccxliv\)](#)

The file already contains some entries depending on the operating system. You can remove the content of the file which means that no paths are excluded from the scan. However, if you delete the whole file, it will be recreated with the default content before the next software scan.

Unless you exclude specific paths, all the following drives are included in the scan:

- **UNIX** All local drives and other drives, such as floppy disk, CD-ROM, and DVD.



Note: Remote drives are not scanned.

- **Windows** All local drives.

Specify paths according to the following syntax:

```
drive:path
```



Important: When you specify a path delimiter, you must use a forward slash (/) instead of a backslash (\). For example, `C:/Program Files`.

drive

Specifies the drive. Asterisks (*) and question marks (?) are supported. This variable is optional on UNIX.

path

Specifies the path. Asterisks (*) and question marks (?) are supported. This variable also supports the following CSIDL values on Windows:

```
%CSIDL_WINDOWS%
%CSIDL_PROGRAM_FILES%
%CSIDL_COMMON_DESKTOPDIRECTORY%
%CSIDL_COMMON_STARTMENU%
%CSIDL_COMMON_STARTMENU%
%CSIDL_COMMON_STARTUP%
%CSIDL_COMMON_PROGRAMS%
```



Important: The above CSIDL values already have a drive specified. If you use them, omit the *drive* variable.

You can refer to the following examples when specifying your paths.

- Excludes the `System Volume Information` folder on any local drive:

```
?:/System Volume Information
```

- Excludes the `System32` folder on the local drive that is specified in the CSIDL value:

```
%CSIDL_WINDOWS%/System32
```

Discovering software and hardware with disconnected scanner on Windows and UNIX

9.2.7 Available from 9.2.7. You can discover software and hardware inventory by using disconnected scans that do not require direct connection between the scanned computers and the BigFix server. Scripts that are provided in the disconnected scanner package initiate software and capacity scans, and prepare scan results that you later on upload to BigFix Inventory.

Disclaimer

 **Restriction:** If you use the disconnected scanner for BigFix Virtualization Capacity, also referred to as subcapacity licensing, you must comply with the following rules:

- You must obtain the approval from BigFix Compliance to use disconnected scans. To request for such an approval, contact your Sales Representative who will instruct you on how to contact BigFix Compliance.
- Disconnected scans can be applied when the BigFix client cannot be installed due to technical, legal, business, or security reasons, or in case of other valid justification.
- The disconnected scanner for BigFix Virtual Capacity can be deployed only to supported and eligible operating systems.

IBM i The BigFix client is not available for IBM iSeries, and thus it is the only method of software and hardware discovery on these systems. Unlike on other operating systems, approval from BigFix Compliance is not required to use disconnected scans on IBM i. For more information, see: [Discovering software and hardware on IBM i \(on page cclxxi\)](#).

Disconnected scans should be used with caution, especially when you use BigFix Inventory for subcapacity reporting. They require much more user control and manual maintenance including scanner and catalog updates, periodic transfer of data, and manual health checks to assure report correctness. The maintenance is customer responsibility because it is not automated as in case of the BigFix client.

Disconnected scanner package

To collect the software and hardware inventory data, you need an endpoint package that consists of:

- Scanner
- Configuration files
- Scripts to run the scans and prepare the scan result package

The scripts initiate software and capacity scans, gather scan results, and adjust them to the format that is compatible with BigFix Inventory. If the scripts are not appropriate for your environment, you can edit and customize them, or create new scripts that better fit your needs.

Scalability

The default number of the result packages that can be imported at a time is limited to 1000. However, the disconnected scanner provides support for more packages. It is recommended not to exceed 10 000. To increase the number of output capacity, configure the transaction logs size and increase java heap. For more information, see: [Tuning performance in medium and large environments \(on page dcccxliv\)](#).

Linux Additionally, for Linux systems, change the `ulimit -n` value to 4096.

Scan frequency

For information about default and minimal scan frequency as well as recommended frequency of importing scan results, see: [Frequency of scans and uploads of data \(on page cxci\)](#).

Multiple environments

When you have multiple environments, for example test and production, ensure that the following requirements are met:

- Every computer reports only one of the environments.
- Results of the disconnected scan from one environment are not uploaded to BigFix Inventory that monitors the other environment.

Limitations

- Resource utilization, and metering data are not supported. For more information, see: [Raw utilization of license metrics \(on page dcxi\)](#).
- Scanning remote shared file systems is not supported.
- Detailed hardware scan cannot be collected if you use the disconnected scan.
- The disconnected scanner always collects full scans, not delta scans.
- The directories that are excluded from software scans by default cannot be used as installation directories. To view the full list, see: [List of excluded directories \(on page ccxli\)](#).
- **Solaris** The Package Data report does not provide any information in the **Description** column.
- The information provided in the Operating System column might be slightly different for the computers that are scanned by the disconnected scan, and the computers scanned by a regular scan.
- The disconnected scan does not include additional computer properties that are defined by the user. To collect these details, you need to define additional entries in the `computer.yml` file.

Concept

[Discovering software and hardware on IBM i \(on page cclxxi\)](#)

Disconnected scanner requirements

9.2.7 Available from 9.2.7. The list of disk space requirements, and the supported operating systems compatible with the disconnected scanner.

Table 31. Disconnected scanner requirements

Summary for complex table

Operating system	Supported versions	Required disk space
Windows	For information about supported versions, see: BigFix 9.5 - System Requirements and choose BigFix Inventory 9.5 application update 9.2.16.	40 MB + disk space for the scanner cache
Red Hat Linux		24 MB + disk space for the scanner cache
9.2.8 Solaris		80 MB + disk space for the scanner cache.
9.2.9 AIX		80 MB + disk space for the scanner cache.
9.2.10 HP-UX		48 MB + disk space for the scanner cache.

Adding a new data source for disconnected scans

9.2.7 Available from 9.2.7. All scan results are imported to BigFix Inventory from a directory dedicated to disconnected scans. You must add this directory as a data source to allow BigFix Inventory to check its content and import all packages that are stored within. The packages are imported during each data import.

All scan results must be manually transferred from all endpoints into the new directory that works as a data source.

1. On the computer where BigFix Inventory server is installed, create a directory with read and write access, for example:

- **UNIX** `/disconnected`
- **Windows** `C:\disconnected`

2. Log in to BigFix Inventory, and in the top navigation bar click **Management > Data Sources**.
3. Click **New**, and specify the following values:

- a. Specify the name of this data source.
- b. Specify the Database Type as **Disconnected**.
- c. Specify the location of the directory that you created.



Note: The path is not validated. Ensure that you provide a correct one.

Create Data Source

Name*

Database for the BigFix Server*

Database Type*

Location of disconnected scan results*

Download catalog for:

[Windows](#) [AIX](#) [Linux](#) [Solaris](#) [HP-UX](#)

IBM i catalog is included in scanner package

Create

4. Download the catalog for the appropriate platform. It is needed for software discovery. You will copy it to your systems together with the scanner, and the configuration files.
5. Click **Create**.

[Prepare installation files for disconnected scans. \(on page cli\)](#)

Preparing installation files for disconnected scans

9.2.7 Available from 9.2.7. After you added the data source, prepare the installation package with the scanner and the configuration files that will be used to initiate scans, and move it to the disconnected system. The software scan also requires the software catalog. You can download the catalog from the BigFix Inventory server.

1. Download the scanner.
 - a. Log in to the BigFix console.
 - b. In the navigation bar, click **Sites > External Sites > BigFix Inventory v9 > Fixlets and Tasks**.
 - c. In the upper right pane, select **Download the Disconnected Scanner Package (version)**.
 - d. To run the task, click **Take Action**. Choose the disconnected scanner package for the appropriate platform from the drop-down menu.
2. Download the software catalog.
 - a. Log in to BigFix Inventory, and go to **Management > Data Sources**.
 - b. Select your *disconnected* data source.
 - c. Download the catalog for the appropriate platform.
3. Prepare the content of the disconnected scanner directory.

- a. Unpack the disconnected scanner package on the computer from which you intend to transfer the scanner to your endpoints.
 - b. Open the `config` directory in the unpacked package.
 - c. Place the catalog file `CIT_catalog_PLATFORM.xml` in the `config` directory.
4. **Optional:** If any modifications are required, customize the configuration files before distributing them to endpoints in your environment. Then, you do not need to configure each endpoint individually. You can change the default configuration; for example, limit the processor utilization, or specify the directories that need to be excluded from being scanned. For more information, see: [Troubleshooting the discovery by disconnected scanner \(on page cclviii\)](#).



Note: **Solaris** The `fs_config.xml` file contains the list of collected file extensions under the **FileMask** parameter. If it is required, you can narrow down this list by removing the unwanted extensions from the file. For other supported systems, the list of extensions is configured automatically based on the software catalog.

5. Upload the whole content of the disconnected scanner package to the endpoints. Do not change file names.



Note: After you install the scanner on endpoints, the scanner installation path cannot be changed. The installation path is a location where you upload the installation files. To change the installation path, back up the configuration files, scan results and logs. Then, uninstall the scanner, move the files to a different directory, and then install the scanner in a new location.



Restriction:

Linux

HP-UX

The installation path should not contain any white space characters.

Solaris

The installation path should be short, and should not contain any special characters or spaces.

Windows

The name of the directory where the installation files are copied cannot contain any special or national characters. Spaces are allowed. The entire path cannot exceed the total of 99 characters.

[Install the scanner and gather initial data. \(on page cclii\)](#)

Installing the scanner and gathering initial data

9.2.7

Available from 9.2.7. After you copied the installation files and the software catalog to the endpoints, start the installation and the initial scans.

- To perform this task, you must have the appropriate privileges.
 - **UNIX** root
 - **Windows** administrative privileges with elevation
- If you use License Metric Tool or Tivoli® Asset Discovery for Distributed 7.x, remember that the installation of the disconnected scanner stops the 7.x agent. You need to restart the 7.x agent after the setup.

1. Go to the directory where you copied the installation files.
2. Run the setup file.

- **UNIX** `setup.sh`
- **Windows** `setup.bat`

It automatically performs the following steps:

- Installs the scanner
- Runs the initial capacity scan, and creates the `computer.yml` file.

The initial capacity scan is needed to collect data about your system, and to create the `computer.yml` file that will identify this system. The file is added to each package with scan results and it helps to catalog and correctly display them in the user interface. It contains information about the operating system, such as host name, and a unique computer ID.



Note: **Linux** For older versions of Red Hat Linux on s390 mainframe, it can be required to manually provide machine type, processor type, shared pool capacity, and system active processors. The script will ask you to specify these parameters if the scanner cannot collect these details automatically.

Solaris For Solaris systems that run on SPARC, it is required to specify if the system is in Dynamic System Domains.

- Schedules regular capacity scans.

The capacity scan is scheduled to run every 30 minutes to gather current results for virtualized environments where the capacity can dynamically change depending on the allocated resources. If you want to collect the scan data only once, run the setup file with the `-noschedule` option.

The scanner is set up now, and ready to run the software scan and gather results. During installation, the `endpoint_id.txt` file was created in the `config` directory. The file contains the host name, and the time stamp of the first setup. The combination of these two elements becomes an endpoint identifier that can be modified by the user if needed. In case of scanner upgrade, or re-installation, the existing endpoint identifier from the file will be used.

[Run the software scan and gather scan results. \(on page ccliv\)](#)

Running software scans and gathering scan results

9.2.7 Available from 9.2.7. The script that is responsible for the software scan collects the results of the catalog-based, package data, and ISO tags scan as well as retrieves results of the capacity scan and the `computer.yml` file. All these results are combined and compressed into a common `.zip` or `tag.gz` package that can be imported to BigFix Inventory.

- Make sure that you have the appropriate privileges to perform this task.
 - **UNIX** root
 - **Windows** administrative privileges with elevation
- Each time you update the software catalog in BigFix Inventory, run an import of data or wait for the scheduled import. Next download the newest version of software catalog to an endpoint to run the software scan on it. For more information, see: [Downloading the software catalog \(on page ccli\)](#).
- The scan takes around 30 minutes. The exact time depends on the computer load, file system size, and the CPU threshold if you set it.
- For information about default and minimal scan frequency, see: [Frequency of scans and uploads of data \(on page cxciij\)](#).
- By default, no limit is set on the processor usage during the scan. To limit processor usage, set the CPU threshold. For more information, see: [Optimizing the processor utilization \(on page cclviii\)](#).

1. Go to the directory, where you uploaded the disconnected scan package.
2. Run the following script to initiate the scan.

- **UNIX** `run_sw_and_pack.sh`
- **Windows** `run_sw_and_pack.bat`

A compressed package with scan results was created in the output directory. It contains capacity scan results since the date when the last package was generated. The uncompressed scan results are deleted from their original location to ensure that they are not included in any future packages.

[Import the package to BigFix Inventory. \(on page ccliv\)](#) After it is successfully imported, delete it from the endpoint to save disk space.

Importing scan results to BigFix Inventory

9.2.7 Available from 9.2.7. To import new scan results to BigFix Inventory, copy the package to the disconnected scans directory that you added as a data source.

- For information about recommended frequency of importing scan results, see: [Frequency of scans and uploads of data \(on page cxciij\)](#).
- If you upload a package with older results, the current ones are overwritten.
- Discovered software is reported with the import time, not the package time stamp.

- Import only one package per endpoint. Otherwise, capacity results might be skipped.
 - Processed packages are not removed from your disconnected scans directory. The ones already imported are skipped.
1. Go to the directory where you uploaded the disconnected scan package. Ensure that a result package was created in the `output` directory after the software scan.
The package name uses the `<UTC_DATE>-<hostname>.<extension>` pattern. It contains the following files:
 - **UNIX** `capacity.tar.gz, catalog_scan.xml, computer.yml, isotag_scan.tar.gz, package_scan.xml, source.tar.gz, file_scan.xml`
 - **Windows** `capacity.zip, catalog_scan.xml, computer.yml, isotag_scan.zip, package_scan.xml, source.zip, file_scan.xml`
 2. Copy the package into the disconnected scans directory that you added as a data source in BigFix Inventory.
 3. Wait for the scheduled import or run it manually.

During the data import, BigFix Inventory checks the directory that you added as a data source and imports scan packages that meet the requirements. Scan results are processed and displayed on the reports in the user interface.

Files from processed packages are saved in the `installation_directory/wlp/usr/servers/server1/tmp/datasources/Datasource_ID` directory.

The processed packages are stored in the directory specified in the data source configuration. The packages are not automatically removed. It is recommended to create the audit snapshot once a quarter. Then, back up the packages in a different directory, and remove them from the directory specified in the data source configuration.

Disconnected scanner maintenance and troubleshooting

9.2.7 Available from 9.2.7. You can use the maintenance and troubleshooting tasks to monitor the scan progress, optimize the processor usage during the scan, understand and fix the most common issues, or learn more about the log files.

Updating the scanner and the scripts

9.2.7 Available from 9.2.7. Both the scanner and the scripts are updated on a regular basis to introduce improvements and new functions. To update the scanner, once again download all the required files, and copy them into the directory where the previous version is already located.

1. Log in to the BigFix console.
2. In the navigation bar, click **Sites > External Sites > BigFix Inventory v9 > Fixlets and Tasks**.
3. In the upper right pane, select **Download the Disconnected Scanner Package (version)**.
4. To run the task, click **Take Action**. Choose the disconnected scanner package for the appropriate platform from the drop-down menu.

5. **Optional:** If you made any changes to the software scan configuration files `sw_config.xml`, `fs_config.xml`, or `isotag_config.xml`, make sure that you back up the files before you extract the package. These files are overwritten.
6. Extract the contents to the disconnected scanner directory.
7. **Optional:** If you recently updated the software catalog in BigFix Inventory, update it also on the system with the disconnected scanner.
 - a. To download the software catalog for disconnected scanner, log in to BigFix Inventory and navigate to **Management > Data Source**. Choose the disconnected data source, and download the catalog for the appropriate platform.
 - b. Place the catalog in the `config` directory of the disconnected scanner package that you downloaded through fixlet.



Note: The catalog is available for download after uploading the main software catalog and running a data import.

8. **Optional:** If you made any changes to the software scan configuration files: `sw_config.xml`, `fs_config.xml` or `isotag_config.xml`, modify the new file versions to reflect the changes. Look at the back up files that you created in step 5, and rewrite the changes to modify the new files. Do not copy the old back up files as they might differ between the versions.
9. Run the setup file.
 - **UNIX** `setup.sh`
 - **Windows** `setup.bat`

Decommissioning computers with disconnected scanners

9.2.7

Available from 9.2.7. When you decommission a computer that is no longer in use, you can mark it as decommissioned on the Computers report, so that it is not displayed in BigFix Inventory.



You must have the Manage Computer Groups permission to perform this task.

1. In the top navigation bar, click **Reports > Computers**.
2. Click the computer name.
3. Click **Decommission computer**.

Decommission computer

Software on NC047118

3 Software Installations including IBM TCP/IP Connectivity Utilities for i, IBM i, and IBM HTTP Server for i

Computer Properties

DNS Name	Last Seen
NC047118	08/18/2016 05:34 PM
IP Address	Latest Scan Import
9.156.47.118	09/05/2016 01:24 PM
Operating System	
IBM i V7R3M0 PowerPC 64	

After you decommission the computer, this computer does not show up on the current reports. You can still manage the software on the Software Classification panel, but it is marked as uninstalled. The software is still displayed on the All Metrics report.

Uninstalling the disconnected scanner

9.2.7 Available from 9.2.7. If you no longer need the disconnected scanner to scan the computer, you can uninstall it from any endpoint.

1. Open the directory where the disconnected scanner is installed.
2. Create a back up of `endpoint_id.txt` file that is located in `config` directory. This file will be required to install disconnected scanner on the endpoint again in the future.
3. Run the following script.
 - **UNIX** `./uninstall.sh`
 - **Windows** `uninstall.bat`
4. Remove the whole directory where the disconnected scanner was installed.

The scanner is successfully uninstalled, the data is no longer collected and all software and hardware scans stop.

Understanding the Deployment Health and Scan Health

9.2.7 Available from 9.2.7. Deployment Health and Scan Health contain information about your clients and the scans that were initiated. You can view this information on the Deployment Health, and Scan Health widgets on the dashboard, or on the Computers report. Deployment Health excludes all the disconnected datasource endpoints. To properly understand values that are displayed for systems scanned with disconnected scans, see the explanation of particular columns on the Computers report.

Table 32. Explanation of Deployment Health and Scan Health information for systems with disconnected scans

Column Name	Description
Has Outdated Version	Always shows <code>No</code> .
Has Low Disk Space	Always shows <code>No</code> .
Is Out of Sync	Always shows <code>No</code> .
Has Missing Prerequisites	Always shows <code>No</code> .
Agent Version	Taken from the Agent Version property in the <code>computer.yml</code> file.
Catalog Scan Successful	Shows <code>Yes</code> if the imported package contains the <code>catalog_scan.xml</code> file.

Table 32. Explanation of Deployment Health and Scan Health information for systems with disconnected scans

(continued)

Column Name	Description
Software Tags Scan Successful	Shows <code>yes</code> if the imported package contains the <code>isotag_scan</code> package.
File System Scan Successful	Shows <code>yes</code> if the imported package contains the <code>file_scan.xml</code> file.
Package Scan Successful	Shows <code>yes</code> if the imported package contains the <code>package_scan.xml</code> file.
Catalog Version	Taken from the Catalog Version property in the <code>computer.yml</code> file.
Endpoint Scanner Catalog Version	Taken from the Catalog Version property in the <code>computer.yml</code> file. It displays only the second part of the whole catalog version, for example <code>1250110.1</code> .
Last Scan Attempt	Taken from the timestamp used in the final package name.
Scanning Shared Disks	Always shows <code>no</code> .
Mounted Shared Disks	Always shows <code>no</code> .

Troubleshooting the discovery by disconnected scanner

9.2.7 Available from 9.2.7. The troubleshooting section lists the most common issues you might encounter while discovering software and hardware inventory with the disconnected scanner. The main objective of troubleshooting is to determine why something does not work as expected and explain how to resolve the problem.

Log files

Scanner log files are in the following directory:

UNIX `Installation directory/logs`

Windows `Installation directory\logs`

Optimizing the processor utilization

By default, processor utilization is not limited during the scan. You can optimize it by setting up the CPU threshold. Setting the threshold might lengthen the time of the scan.



Note: Setting the threshold does not guarantee that the CPU consumption is always below the specified value. It fluctuates around that value, sometimes exceeding it and sometimes dropping below it. Temporary peaks are expected. Setting the threshold might lengthen the time of the scan.

1. Edit the configuration file or files that you downloaded with the scanner: `sw_config.xml`, `fs_config.xml`, `isotag_config.xml`.
2. Add new **cpuThreshold** attributes in the appropriate file.
 - For *FSScanner* add a corresponding parameter in the following files: `sw_config.xml`, `fs_config.xml`, `isotag_config.xml`
 - For *Signatures* add a corresponding parameter in the `sw_config.xml` file.

For example:

- `sw_config.xml`

```

<CIT>
<XSE version="1.0">
<Plugins>
<Plugin name="FSScanner">
(...)
<Attribute name="cpuThreshold" value="20:2:150:850"/>
<Attribute name="maxQueryTime" value="43200"/>
<Attribute name="interruptOnTimeout" value="true"/>
<Attribute name="assumeAutoFS" value="remote"/>
(...)
</Plugin>
</Plugins>
<Signatures>
<Attribute name="cpuThreshold" value="20:2:150:850"/>
(...)
</Signatures>
</XSE>
</CIT>
</IBM>

```

- `fs_config.xml` or `isotag_config.xml`

```

<IBM>
<CIT>
<FSScan version="1.0">
(...)
<Provider value="provider_cache"/>
<AssumeAutoFS value="remote"/>
<CpuThreshold value="20:2:150:850"/>
</FSScan>
</CIT>
</IBM>

```

3. Specify one of the following values:

Desired utilization	Value
5%	10:2:50:950
10%	20:2:100:900
15%	20:2:150:850
20%	30:2:200:800
25%	30:2:250:750
50%	50:2:300:500

4. Save the file.

Excluding directories from being scanned

You can exclude directories from being scanned. To do this, edit all configurations files: `sw_config.xml`, `fs_config.xml`, `isotag_config.xml`, and specify the directories as values of the `excludeDirectory` attributes.

- To exclude the `/tmp` directory:

in `sw_config.xml`: `<Attribute name="excludeDirectory" value="/tmp/" />`

in `fs_config.xml` and `isotag_config.xml`: `<ExcludeDirectory value="/tmp/">`

- To exclude all directories named `tmp`, for example `/tmp` and `/dir/tmp`:

in `sw_config.xml`: `<Attribute name="excludeDirectory" value="*/tmp/" />`

in `fs_config.xml` and `isotag_config.xml`: `<ExcludeDirectory value="*/tmp/">`

Common problems

The package is being skipped during the import despite meeting the requirements

Such a problem might occur if you are importing more than one package from an endpoint. BigFix Inventory always imports the latest package and omits all previous ones. The following message about skipping the package is related to irrelevant packages and can be ignored:

```
2015-11-05 09:35:22 WARN: Skipping file: package_scan.xml from: 201510101200-2inarow.zip.
The data was already processed.
2015-11-05 09:35:22 INFO: Skipping file: package_scan.xml.
Unsupported file in: 201510101200-2inarow.zip.
```

The detected software is duplicated in BigFix Inventory after reinstalling disconnected scanner

If you uninstall the disconnected scanner and remove its directory from a computer, and then reinstall the disconnected scanner, the endpoint ID of this computer is changed. As a result, the software inventory is duplicated in BigFix Inventory. To avoid duplication after you reinstall the disconnected scanner, perform the following actions.

1. Go to `config` directory, and substitute the `endpoint_id.txt` file with the backup file you created before uninstallation. For more information, see: [Uninstalling the disconnected scanner \(on page cclvii\)](#). Changing the endpoint ID allows you to keep the historical data, and collect the new data for this endpoint.
2. Check the DNS and IP address of the endpoint.
3. Log in to BigFix Inventory.
4. Navigate to **Reports > Computers**.
5. To identify the duplicated computer, add the following columns to the report DNS Name, IP Address and First Seen.
6. Compare the two computers with the same DNS Name and IP Address, and find the one that was first seen more recently.
7. Decommission this computer. For more information, see: [Decommissioning computers with disconnected scanners \(on page cclvi\)](#).
8. Wait for the scheduled import or run it manually.

The host name of an endpoint provided by the disconnected scanner is different than expected

The host name that is provided by the disconnected scanner and a regular scanner can be different. The regular scanner uses the host name that is provided by BigFix, while the disconnected scanner checks the host name by running the `hostname` command, or, if not available, through the scanner check.

Installation of the scanner and the capacity scan are slow on AIX

It happens when the host name cannot be resolved. To solve the problem, add the host name to DNS or to the `/etc/hosts` file. The default configuration of the name resolver on AIX first looks for the host name in the DNS then in the `/etc/hosts` file. If you add the host name to the hosts file, you might also need to make changes in the `/etc/netsvc.conf` file and restart the system.

Monitoring the import log

When you import a package with scan results to BigFix Inventory, the import log (**Management > Data Imports**) displays information and warning messages related to your package. If the package or its content does not meet the requirements, it might be skipped during the import, and the scan results will not be displayed in BigFix Inventory. While monitoring the import log, you can expect the following messages:

```
info: Processing scan package: package from: directory.
```

```
info: Invalid file name. The scan package must match the pattern YYYYMMDDHHMM-*.zip. Skipping package: package.
```

```
info: Invalid Catalog Version property in the computer.yml file: version. The property must match the pattern <Number>.<Number>. The catalog might be reported as outdated.
```

```
warn: The disconnected scans location does not exist: directory.
```

```
warn: Invalid file name. The scan date must be later than year 2000. Skipping package: package.
```

```
warn: Skipping corrupted scan package: package. The following error occurred: error.
```

```
warn: Missing computer.yml file. Skipping package: package.

warn: Missing property: property in the computer.yml file. Skipping package: package.

warn: Invalid computer.yml file. Skipping package: package. The following error occurred: error.

warn: The value of the endpointID property in the computer.yml file is longer than 512 characters. Skipping
package: package.

warn: The value of the endpointID property in the computer.yml file is empty. Skipping package: package.

warn: Missing property: endpointID in the computer.yml file. Skipping package: package.

warn: Importing older scan package from: date for computer: computer_id. This computer already has data from:
date.

warn: The file is larger than 2097150 bytes. Skipping file: scan_file from: package for computer: computer_id.

warn: The following error occurred when processing the package_scan.xml file: error.

warn: The insertion of the operating_system package data into the database failed on record content. The
following error occurred: error.
```

Scheduling imports of data

BigFix clients report data to the BigFix server that stores the data in its file system or database. BigFix Inventory server connects to the BigFix server and its database, downloads the stored data, and processes it. This process is called an import or Extract, Transform, Load (ETL). By default, the import runs once a day at midnight. You can schedule it for a period that is most suitable for your environment considering its size and specification.



You must have the Manage Imports permission to perform this task.

The import of data is an Extract, Transform, Load (ETL) process.

- During the Extract stage, data is extracted from the BigFix server. It includes information about the infrastructure, installed clients, and detected software. ETL also checks whether a new software catalog is available, gathers information about the software scan and files that are present on the computers. It also collects data from VM managers.
- During the Transform stage, the extracted data is transformed to a single format that can be loaded to the BigFix Inventory database. This stage also involves matching scan data with the software catalog, calculating processor value units (PVUs), processing the capacity scan, and converting information that is contained in the XML files.
- During the Load stage, the data that was extracted and transformed is loaded into the BigFix Inventory database and can be used by BigFix Inventory.

1. In the navigation bar, click **Management > Data Imports**.
2. To schedule regular imports, select **Enabled**, specify the number of daily imports and their hours, and click **Save**.

Import Settings

Enabled

Imports per day: (times specified in UTC +02:00)

Start Time:

Information about data imports and their status is displayed in the import history.

Status	Start Time	User Name	Duration	Download
✓	08/29/2017 08:12 AM	admin	00:02:56	↓
✓	08/28/2017 02:00 PM	Scheduled	00:01:59	↓
✓	08/28/2017 01:46 PM	admin	00:04:55	↓
✗	08/28/2017 12:32 PM	admin	01:00:38	↓

If you want to quickly discover recent changes to software inventory without having to wait for the complete data import, you can schedule imports of raw data. Such imports are much quicker but the imported data is not visible in the user interface. It can be retrieved only by using REST API. For more information, see: [Importing raw scan data \(on page cclxiii\)](#).

9.2.11 If the imports of data are failing, you can enable partial imports. For more information, see: [Enabling partial imports \(on page cclxiv\)](#).

Importing raw scan data

9.2.2 Available from 9.2.2. You can schedule an import of raw scan data. Such an import is much quicker, because it only imports raw scan results, and omits successive steps, such as aggregation and license calculations. The imported raw data, however, is not updated in the user interface, and can only be retrieved by using REST API. This type of import is used to quickly discover recent changes to software inventory without having to wait for the complete data import.

- The Raw Data Only import does not replace the complete data import, which must be run regularly to process and aggregate the data.
- The Raw Data Only import can be run only as a scheduled import. When you click **Import Now**, you always run a complete data import.
- To run the Raw Data Only import only once, use REST API. For more information, see: [Running data imports \(on page dcccclxxviii\)](#).

1. To enable this feature, go to <https://hostname:port/management/feature>, and select **Enable the Raw Data Only import mode**.
2. In the navigation bar, click **Management > Data Imports**.
3. To schedule regular imports, select **Enabled**, and specify the number of daily imports and their hours.
4. Select **Raw Data Only** next to each import that you want to use for importing raw scan results.

Import Settings

Enabled

Imports per day: (times specified in UTC +01:00)

06:00AM	<input type="checkbox"/> Raw Data Only
02:00PM	<input checked="" type="checkbox"/> Raw Data Only
04:00PM	<input checked="" type="checkbox"/> Raw Data Only

5. Click **Save**.

Information about completed data imports is displayed in the import history. Each entry has an icon that indicates the type of the import, either *Complete* or *Raw Data Only*.

Status	Type	Start Time	User Name	Duration	Download
✓		09/08/2015 07:3...	Administrator	1:27:29	
✓		09/08/2015 02:1...	Scheduled	0:00:22	

When the Raw Data Only import finishes, retrieve the imported raw data by using REST API. For more information, see [REST API for retrieving raw scan results \(on page dccc1xxvii\)](#).

Enabling partial imports

9.2.11 Available from 9.2.11. If the imports of data are failing, you can enable partial imports. Such imports process all steps and continue even if steps that are not crucial for the operation of BigFix Inventory fail. These imports allow for avoiding accumulation of scan data during the period when you troubleshoot the problem.

1. In the navigation bar, go to **Management > Advanced Server Settings**.
2. Change the value of the `enable_partial_imports` parameter to true.

Information about data imports and their status is displayed in the import history. Partial imports are indicated by a warning sign.

Status	Start Time	User Name	Duration	Download
	02/15/2018 10:16 AM	Administrator	00:05:47	
	02/15/2018 12:00 AM	Scheduled	00:01:50	
	02/14/2018 01:59 PM	Administrator	00:45:09	
	02/14/2018 11:43 AM	Administrator	00:00:00	

Partial imports might fail if any of the steps that are crucial for the operation of BigFix Inventory is disrupted, for example, if the database connection does not work. The imports that fail are displayed in red in the import history.

After the problem is solved, and imports complete successfully, disable partial imports.

Creating capacity configuration for Linux on z Systems

To properly calculate subcapacity values for products that are installed on Linux on z Systems, create capacity configuration for such computers. First, run a fixlet that creates a file with manually entered capacity values and places it on the target computer. Then, run regular software and capacity scans to discover the installed software and calculate its license metric utilization.



Note: Automatic capacity configuration is supported on Linux on z/KVM and on System z10 starting from model E64 (type 2097) mainframes with z/VM 6.3 and later that supports the Store Hypervisor Information (STHYI) instruction. If the automatic capacity configuration is supported, you do not have to perform any manual actions to calculate the capacity values.

The **Create Capacity Configuration for Linux on z Systems** fixlet is relevant on all computers where manual configuration is required. It is not relevant on computers where automatic capacity configuration is supported. If the fixlet is not relevant, you do not have to perform any manual actions to calculate the capacity values.

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. Select **Create Capacity Configuration for Linux on z Systems** and specify the following values:
 - Machine Type
 - Processor Type
 - Shared Pool Capacity
 - System Active Processors
4. To run the task, click **Take Action**.
5. Select computers for which you want to create the capacity configuration, and click **OK**.

The fixlet is relevant only on computers on which creating the manual configuration is required.

You created capacity configuration that will be used to calculate subcapacity values. If you want to edit or delete the capacity configuration, use the **Edit Capacity Configuration for Linux on z Systems** or **Delete Capacity Configuration for Linux on z Systems** task.

To discover software that installed on the computers and calculate its license metric utilization, schedule regular software and capacity scans. For more information, see: [Initiating software scans \(on page cciii\)](#), and [Initiating the capacity scan on all computers \(on page ccx\)](#).

Setting the DSD mode on Solaris

If a computer runs the Solaris operating system and is in the DSD domain, set the DSD mode to ensure that metric utilization is correctly calculated for software that is installed on this computer. If you do not set the DSD mode, Solaris machines in DSD might not be properly identified and metric utilization might be underestimated.

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. Select **Set DSD Mode** and then in the lower pane, click **Take Action**.
4. Select computers for which you want to set the DSD mode, and click **OK**.
5. **Optional:** If you want to remove the DSD mode, use the **Unset DSD Mode** task.

Optimizing scanner cache configuration

The scanner cache folder is used to store information about scanned files and directories in your file system. By knowing the hierarchy of files, the scanner can locate them quicker, which results in shorter scans. The amount of disk space that is needed for the cache depends on the number of files that are being scanned. If the current location of the cache folder cannot ensure sufficient disk space, you can change the location of the cache folder or optimize the cache.

- **9.2.2** Change the default location of the scanner cache folder.
 1. Log in to the BigFix console.
 2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
 3. Select **Change Scanner Cache Folder**.
 4. Enter an absolute path to the new cache folder, or restore the default `scanner_install_dir/cache_data` folder.
 5. Click **Take Action**, and select computers on which you want to change the scanner cache folder. Then, click **OK**.

After you change the cache folder or restore the default one, all current cache is removed. The new cache is created after you run the software scan. To check the current cache folder, see the results of the Scanner Information analysis.

- **9.2.4** Optimize the scanner cache. This option is recommended for infrastructures with a large number of files that need to be scanned. Optimization of the cache allows for reducing its size and shortening the time of the scan.

- **9.2.6** The option also optimizes remote shared file systems.

! **Important:** Starting from application update 9.2.6, fresh installations have the optimization enabled. However, scanner cache settings are not changed if you are upgrading. To check whether scan optimization is enabled on a particular computer, in the BigFix console go to **Computer Management > Computers**, and select the computer. Then, check the value of the **CIT_Scan_Provider** parameter. If the value is `provider_cache`, the optimization is not enabled. If it is `provider_cache2`, the optimization is enabled.

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. Select **Configure Scan Cache**.
4. To reduce the amount of disk space that is used by the cache, select **Enable cache optimization**.
5. Click **Take Action**, and select computers on which you want to enable the optimization. Then, click **OK**.

When you enable the scan optimization, the cache that is created during the software scan includes only files that are relevant for the types of software scans that you are running. The size of the cache is reduced.

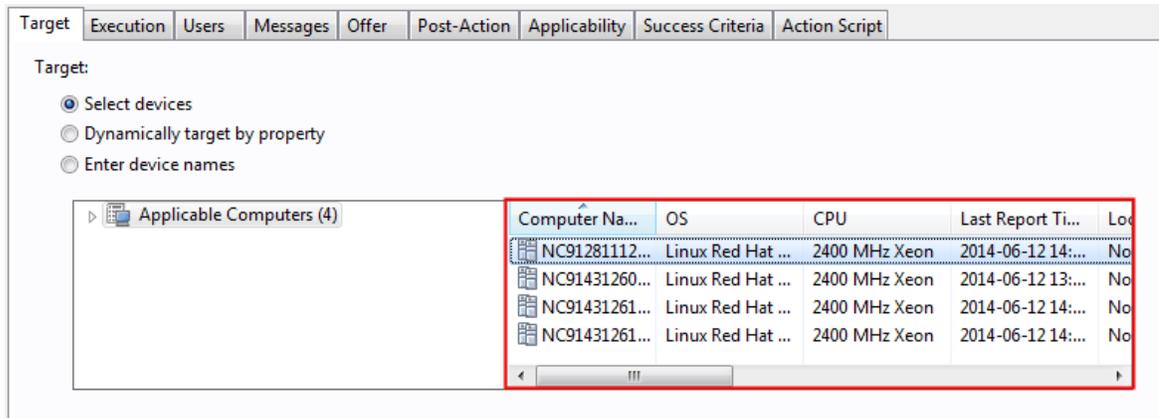
Updating scanner catalogs

Scanner catalogs are used to discover software on the endpoints. They are automatically created and distributed to the endpoints after every import that contains the BigFix software catalog. If the automatic distribution fails, endpoints on which the catalogs were not updated have the **Outdated Catalog** status on the **Scan Health** widget. You must manually update scanner catalogs on these endpoints.

- Before you update the scanner catalogs manually, try to determine why the automatic update failed. For more information, see [Server operation problems \(on page dcclxxxvii\)](#).
- Ensure that the BigFix Inventory server is visible to your BigFix server.
- If Secure Socket Layer (SSL) is enabled in BigFix Inventory, all updates are also downloaded through SSL. The BigFix server must recognize SSL certificates of BigFix Inventory as valid.

1. Log in to BigFix Inventory.
2. [Import \(on page cdxli\)](#) the software catalog.
3. In the top navigation bar, click **Management > Catalog Upload**.
4. To download the fixlet file to your computer, click the question mark sign. Then, click **Catalog Download Fixlet**. Choose the location where you want to save the `catalog_download.bes` file, and click **Save**.
5. Copy the file to the computer where the BigFix console is installed.
6. Log in to the BigFix console.
7. To import the fixlet, click **File > Import**.
8. Open the directory where you store the `catalog_download.bes` file, select the file, and click **Open**. The file is imported.

9. In the left pane, click **Sites > Master Action Site > Fixlets and Tasks**. A list of available fixlets opens in the upper right pane.
10. To run the fixlet on the endpoints, select **Catalog Download (Version: version)**, and click **Take Action**.
11. Select the computers on which you want to run the fixlet, and click **OK**.



You imported the scanner catalogs to the endpoints in your infrastructure. To gather inventory data from the endpoint, you must now [initiate software scans \(on page cciii\)](#).

Deprecated: Setting up additional analysis properties

Analysis properties are used to recognize software and gather information about its usage. Analysis properties are set by default in BigFix Inventory. You can also set up your own properties that you want to use to gather information from the endpoints.

9.2.12 Starting from application update 9.2.12, adding usage properties becomes deprecated.

 You must have the Manage Usage and Package Properties permission to perform this task.

1. In the top navigation, click **Management > Usage Properties**.
2. To add an application usage property, click **New**.
3. Specify the name of the property. From the list of available properties, choose the data source property that you want to use to discover software that is installed in your infrastructure, its use or other properties, and click **Create**.

Disabling the collection of software usage

Information about software usage is collected if you enabled the default scan configuration or manually activated the Application Usage Statistics analysis and ran the relevant scan. If you are in the initial deployment phase, or if you do not need information about software usage, you can disable the collection of this information to improve BigFix Inventory performance and shorten the import time.

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.

3. In the upper right pane, select **Disable Application Usage Statistics**, and then in the lower pane, click **Take Action**.
4. Select computers from which you no longer want to collect information about software usage, and click **OK**.

Optimizing the volume of scanned file data

The Scanned File Data report provides information about files with particular extensions that were discovered on the computers in your infrastructure, and the files that are associated with the processes detected on these computers regardless of the extension. The information is typically used for creating custom discovery and use signatures. To optimize the volume of the monitored data, narrow down the list of file extensions that are collected to only those extensions that are typically used for creating signatures. It reduces the workload on the BigFix Inventory server and improves the application performance.

If you are upgrading to version 9.0.1.2, a configuration panel on which you can automatically optimize the volume of scanned file data is displayed. The following procedure describes manual steps that you can perform if you skipped the automatic configuration.

Important:

- Changing the list of discovered extensions does not change the scope of the discovered software.
- Scanning shared disks significantly affects the number of scanned files that are reported by BigFix Inventory.
- **Windows** Changing the list of extensions of the files that are reported by BigFix Inventory might result in detailed versions not being reported.
- **9.2.8** **UNIX** Starting from application update 9.2.8, the file system scan reports files that are based on processes that have been run on the computers in your infrastructure, regardless of their extension. Thus, you can narrow down the list of file extensions to only those extensions that are typically used for creating signatures unless you use the external tools that retrieve the scanned file data.
- If you use disconnected scans, and you change the list of discovered extensions, distribute the new catalog to disconnected computers. For more information, see: [Preparing installation files for disconnected scans on Windows and UNIX \(on page ccli\)](#) or [Preparing installation files for disconnected scans for IBM i \(on page cclxxv\)](#).

1. To stop the BigFix Inventory server, run the following command.

- **Linux**

```
/etc/init.d/BFIserver stop
```

- **Windows**

```
installation_directory\cli\srvstop.bat
```

2. To optimize the volume of scanned file data, remove extensions that are not typically used for creating signatures from `file_names_all.txt`, `file_names_unix.txt`, and `file_names_win.txt` files.

The files are in one of the following directories:

- **Linux** `installation_directory/wlp/usr/servers/server1/apps/tema.war/WEB-INF/domains/sam/config`
- **Windows** `installation_directory\wlp\usr\servers\server1\apps\tema.war\WEB-INF\domains\sam\config`



Important: If you remove an extension which is a base for a signature that is currently used for software discovery, after the import of data the affected component is marked as uninstalled. The component will be reported back after you run the next software scan and import of data with a new software catalog. To check whether the extension is currently used to detect software, go to the Scanned File Data report and use the following filter:

```
File Name; ends with; <.extension> Recognized; equal to; Yes
```

- The following files are defined in the `file_names_all.txt` to be reported by default:

```
\.ear$
```

- The following files are defined in the `file_names_unix.txt` to be reported by default:

```
\.sh$
\.bin$
\.pl$
\.ear$
\.SH$
\.BIN$
\.PL$
\.EAR$
```

- The following files are defined in the `file_names_win.txt` to be reported by default:

```
\.exe$
\.sys$
\.com$
\.ear$
\.ocx$
```

3. To start the BigFix Inventory server, run the following command.

- **Linux**

```
/etc/init.d/BFIserver start
```

- **Windows**

```
installation_directory\cli\svrstart.bat
```

4. If a new software catalog is available, upload it. Otherwise, re-upload the current software catalog.
5. Wait for the scheduled import or run it manually.
 - During this import, performance might be lower because the software catalog is re-imported.
6. Wait for the scheduled software scan. Alternatively, if you have infrequent software scans, stop the current scan and start a new one. It will allow you for using the optimized list of file extensions in a shorter time.
 - a. Log in to the BigFix console and in the left navigation tree, click **Actions**.
 - b. In the upper-right pane, click **Initiate Software Scan** and then click **Stop**.
 - c. [Initiate a new software scan. \(on page cciii\)](#)
7. Wait for the scheduled import or run it manually. From now on, the optimized list of file extensions is used.

Discovering software and hardware on IBM i

9.2.5 Available from 9.2.5. You can discover software and hardware inventory on IBM i systems by using disconnected scans that do not require direct connection between the scanned computers and the BigFix server. Scripts that are provided in the disconnected scanner package initiate software and capacity scans, and prepare scan results that you later on upload to BigFix Inventory.

 **Important:** Because the BigFix client is not available for IBM i, the disconnected scanner is the only method of software and hardware discovery on IBM i systems. Unlike on other operating systems, approval from IBM Compliance is not required to use disconnected scans on IBM i.

Disconnected scanner package

To collect the software and hardware inventory data, you need an endpoint package that consists of:

- Scanner
- Configuration files
- Scripts to run the scans and prepare the scan result package

The scripts initiate software and capacity scans, gather scan results, and adjust them to the format that is compatible with BigFix Inventory. If the scripts are not appropriate for your environment, you can edit and customize them, or create new scripts that better fit your needs. You can also use a manual procedure. However, it is not recommended for inexperienced users.

Scalability

The default number of the result packages that can be imported at a time is limited to 1000. However, the disconnected scanner provides support for more packages. It is recommended not to exceed 10 000 at a time. To increase the output capacity, configure the transaction logs size and increase java heap. For more information, see: [Tuning performance in medium and large environments \(on page dcccxliv\)](#).

Scan frequency

The disconnected scanner triggers the following scans:

- Capacity scan - The scan runs via the `BFI_HW` jobs that are scheduled during the installation as QSECOFR user on its default queue. The capacity scan runs every 30 minutes.
- Software scan - You can schedule the software scan to run according to your preference and use any specified queue. The recommended scan frequency is once a month.

Each of these scans consists of a set of scripts and binaries that use up to 5 jobs. Thus, if you use the same queue, make sure that the job limit is set to 10 or more.

A Software scan scheduled / run on demand by end user using user/queue specified by end user. Recommended to be executed once a month. Each of these scans consists of a set of scripts and binaries that use up to 5 jobs. Thus, if you use the same queue, make sure that the job limit is set to 10 or more.

Multiple environments

When you have multiple environments, for example test and production, ensure that the following requirements are met:

- Every computer reports only one of the environments.
- Results of the disconnected scan from one environment are not uploaded to BigFix Inventory that monitors the other environment.

Limitations

- Resource utilization, and metering data are not supported. For more information, see: [Raw utilization of license metrics \(on page dcxi\)](#).
- Scanning remote shared file systems is not supported.
- Detailed hardware scan cannot be collected if you use the disconnected scan.
- The disconnected scanner always collects full scans, not delta scans.
- The disconnected scan does not include additional computer properties that are defined by the user. To collect these details, you need to define additional entries in the `computer.yml` file.

Concept

[Discovering software and hardware with disconnected scanner on Windows and UNIX \(on page ccxlviii\)](#)

Disconnected scanner requirements for IBM i

9.2.5 Available from 9.2.5. The list of disk space requirements, the supported IBM i version compatible with the disconnected scanner and other relevant requirements.

Supported operating systems, versions and required disk space

Table 33. Disconnected scanner requirements for IBM i systems

Operating system	Supported versions	Required disk space
IBM i	For information about supported versions, see: BigFix 9.5 - System Requirements and choose BigFix Inventory 9.5 application update 9.2.16.	Scanner installation directories: <ul style="list-style-type: none"> • <code>/etc/cit</code> - at least 20 MB. Also, provide disk space for the scanner cache. • <code>/cit</code> <ul style="list-style-type: none"> ◦ 50 MB for the script and installation media ◦ 50 MB for capacity scan results for 90 days • <code>/qsys.lib/qtivcit.lib/</code> - 200 MB

Other requirements

- The disconnected scanner consumes up to 10 jobs. Make sure that the job limit is sufficient before using the disconnected scanner. For more information, see: [Discovering software and hardware on IBM i \(on page cclxxi\)](#).

Other requirements

- The disconnected scanner consumes up to 10 jobs. Make sure that the job limit is sufficient before using the disconnected scanner. For more information, see: [Discovering software and hardware on IBM i \(on page cclxxi\)](#).

Other requirements

- The disconnected scanner consumes up to 10 jobs. Make sure that the job limit is sufficient before using the disconnected scanner. For more information, see: [Discovering software and hardware on IBM i \(on page cclxxi\)](#).

Adding a new data source for disconnected scans

9.2.5 Available from 9.2.5. All scan results are imported to BigFix Inventory from a directory dedicated to disconnected scans. You must add this directory as a data source to allow BigFix Inventory to check its content and import all packages that are stored within. The packages are imported during each data import.

All scan results must be manually transferred from all IBM i systems into the new directory that works as a data source. You can [automate this process by using NFS \(on page cclxxxii\)](#), or by other means that better suit your environment.

1. On the computer where BigFix Inventory server is installed, create a directory with read and write access, for example: `/disconnected`.
2. Log in to BigFix Inventory, and in the top navigation bar click **Management > Data Sources**.
3. Click **New**, and specify the following values:
 - a. Specify the name of this data source.
 - b. Specify the Database Type as **Disconnected**.
 - c. Specify the location of the directory that you created.



Note: The path is not validated. Ensure that you provide a correct one.

Create Data Source

Name*

Database for the BigFix Server*

Database Type*

Location of disconnected scan results*

Download catalog for:

[Windows](#) [AIX](#) [Linux](#) [Solaris](#) [HP-UX](#)

IBM i catalog is included in scanner package

Create

4. Download the catalog for IBM i. It is needed for software discovery. You will copy it to your IBM i system together with the scanner and the configuration files.
5. Click **Create**.

[Prepare installation files for disconnected scans. \(on page cclxxv\)](#)

Preparing installation files for disconnected scans for IBM i

9.2.5 Available from 9.2.5. After you add the data source, prepare the installation package and distribute it to IBM i endpoints. The package should contain the disconnected scanner and configuration files used to initiate the scans. These files are downloaded from the BigFix console. The package should also contain the software catalog that is downloaded from BigFix Inventory.

9.2.16 Starting from application update 9.2.16, each disconnected scanner package includes the software catalog. You no longer need to separately download the software catalog and place it in the disconnected scanner directory on IBM i computers.

1. Download the disconnected scanner.

- a. Log in to the BigFix console.
- b. In the navigation bar, click **Sites > External Sites > BigFix Inventory v9 > Fixlets and Tasks**.
- c. In the upper right pane, select **Download the Disconnected Scanner Package for IBMi (version)**.
- d. To run the task, click **Take Action**.

The `BFI-DisconnectedScanner-ibm_i-<version>-<timestamp>.zip` file is downloaded.

For versions before 9.2.16 file would be named: `BFI-DisconnectedScanner-ibm_i-<version>-<timestamp>.zip`

2. For versions before 9.2.16: Download the software catalog.

- a. Log in to BigFix Inventory, and go to **Management > Data Sources**.
- b. Select your *disconnected* data source.
- c. Download the catalog for IBM i.

The `CIT_catalog_I5OS.xml` file is downloaded.

3. **Optional:** To customize the disconnected scanner, edit configuration files.

You can change the scanner installation path, output directories for scan results, and some other settings, such as the processor utilization. Configuration files include:

- `configure_scan.sh` - The file contains the following parameters.
 - **CONFIG_DIR** parameter that specifies the installation path for the scanner as well as the location of configuration files and scripts.
 - **PACKAGE_OUTPUT_DIR** parameter that specifies the output directory for scan results.
 - **BASE_OUTPUT_DIR** parameter that specifies the temporary directory in which scan results are stored before they are included in the final package.

Each parameter is changed independently. Ensure that you provide correct paths.

- `schedule_scans.sh` - Modify the file and replace the default user profile name QSECOFR with the name of the profile that will be used to run hardware scans. The minimal requirements for the user profile are the class set to *USER and the special authority set to *NONE. You can use the same user

profile that is used to install the scanner and run the software scan. Then, set up the special authority to *ALLOBJ.

- `sw_config.xml` - The file contains a number of extra settings, such as the processor utilization and directories excluded from scanning. For more information, see: [Troubleshooting the discovery on IBM i \(on page cclxxxiv\)](#).

4. Upload the disconnected scanner package to IBM i computers. Place the files in correct directories.

- a. Extract the `LMT-DisconnectedScanner-ibm_i-<version>-<timestamp>.zip` file to the scanner installation directory. (For versions before 9.2.16: `BFI-DisconnectedScanner-ibm_i-<version>-<timestamp>.zip`)

By default, it is `/cit`. If you customized the directory, place the file in the directory that you specified in the `configure_scan.sh` file.

- b. For versions before 9.2.16: Place the software catalog file `CIT_catalog_I5OS.xml` in the `/cit/config` directory.

[Install the scanner and gather initial data. \(on page cclxxvi\)](#)

Installing the scanner and gathering initial data on IBM i systems

9.2.5 Available from 9.2.5. After you copied the installation files and the software catalog to your IBM i systems, start the installation and the initial scans.

- You must have the *ALLOBJ authority to perform this task.
- The `jar` tool must be available on the endpoint. To check whether the tool is available, type `jar` in the QSH.
- Set the `JAVA_HOME` variable on the computer on which you want to install the scanner and export it to your process.

1. To set the variable, run one of the following commands from the IBM i **Main Menu**.

- To add the variable to your current interactive job, run the following command:

```
ADDENVVAR ENVVAR(JAVA_HOME) VALUE(' /QOpenSys/QIBM/ProdData/JavaVM/jdk80/64bit')
```

- To set the variable at the system level, run the following command:

```
ADDENVVAR ENVVAR(JAVA_HOME) VALUE(' /QOpenSys/QIBM/ProdData/JavaVM/jdk80/64bit')
LEVEL(*SYS)
```

2. Export the `JAVA_HOME` variable to your process.

- a. Open a Qshell session by using the STRQSH command.
- b. Export the variable to your process by using the following command.

```
export -s JAVA_HOME=/QOpenSys/QIBM/ProdData/JavaVM/jdk80/64bit
```

1. Open the QSH and go to the directory where you copied the disconnected scanner package. By default, it is the `/cit` directory.
2. Run the `setup_all.sh` file.

It automatically performs the following steps:

- Installs the scanner.
- Runs the initial capacity scan, and creates the `computer.yml` file.

The initial capacity scan is needed to collect data about your IBM i system, and to create the `computer.yml` file that will identify this system. The file is added to each package with scan results and it helps to catalog and correctly display them in the user interface. It contains information about the operating system, such as host name, and a unique computer ID.

- Schedules regular capacity scans.

The capacity scan is scheduled to run every 30 minutes to gather current results for virtualized environments where the capacity can dynamically change depending on the allocated resources.



Note: After you install the scanner, its installation path cannot be changed. The installation path is a location where you upload the installation files. To change the installation path, back up the configuration files, scan results and logs. Then, uninstall the scanner, move the files to a different directory, and then install the scanner in the new location.

3. Go to the `/cit` or `/cit/config` directory, and verify that the created `computer.yml` file contains the correct data. For more information about the contents of this file and instructions about how to change them, see: [Creating the `computer.yml` file \(on page cclxxxix\)](#).

The scanner is set up now, and ready to run the software scan and gather results.

[Run the software scan and gather scan results. \(on page cclxxvii\)](#)

Running software scans and gathering scan results on IBM i systems

9.2.5 Available from 9.2.5. The script that is responsible for the software scan collects the results of the catalog-based, package data, and software ID tags scan as well as retrieves results of the capacity scan and the `computer.yml` file. All these results are combined and compressed into a common ZIP package that can be imported to BigFix Inventory.

- Make sure that you have the *ALLOBJ authority to perform this task.
- Each time you update the software catalog in BigFix Inventory, run an import of data or wait for the scheduled import. Next download the newest version of software catalog to an endpoint to run the software scan on it. For more information, see: [Downloading the software catalog \(on page cclxxv\)](#).
- The scan takes around 30 minutes. The exact time depends on the computer load, file system size, and the CPU threshold if you set it.
- For information about default and minimal scan frequency, see: [Frequency of scans and uploads of data \(on page cxci\)](#).
- By default, no limit is set on the processor usage during the scan. To limit processor usage, set the CPU threshold. For more information, see: [Optimizing the processor utilization \(on page cclxxxiv\)](#).

1. Run the following command to initiate the scan. Substitute the default `/cit` directory with the directory where you uploaded the scanner.

```
SBMJOB JOB(BFI_PKG) CMD(QSH CMD('/cit/run_sw_and_pack.sh'))
```

2. Monitor the progress of the scan. Check the `run_sw_and_pack_std.log` file created during the scan in the directory where you installed the disconnected scanner package, by default `/cit/logs`, or monitor the scanner logs. For more information, see: [Monitoring the scan progress \(on page cclxxxiii\)](#).

A ZIP package with scan results is created in the package output directory. The default directory is `/cit`. The package name is based on the following pattern: `<YYYYMMDDHHMM>-<endpoint_ID>.zip`.

[Import the package to BigFix Inventory. \(on page cclxxviii\)](#) After it is successfully imported, delete it from the endpoint to save disk space.

Importing scan results from IBM i systems to BigFix Inventory

9.2.5 Available from 9.2.5. To import new scan results to BigFix Inventory, copy the package to the disconnected scans directory that you added as a data source.

- For information about recommended frequency of importing scan results, see: [Frequency of scans and uploads of data \(on page cxcii\)](#).
- If you upload a package with older results, the current ones are overwritten.
- Discovered software is reported with the import time, not the package time stamp.
- Import only one package per endpoint. Otherwise, capacity results might be skipped.
- Processed packages are not removed from your disconnected scans directory. The ones already imported are skipped.

1. Go to the package output directory. The default directory is `/cit`. Verify that a ZIP package was created after running the software scan.

The package name is based on the following pattern: `<YYYYMMDDHHMM>-<endpoint_ID>.zip`, for example `201510191213-6552CFF-5.zip`. It contains the following files:

```
capacity.zip
catalog_scan.xml
computer.yml
isotag_scan.zip
package_scan.xml
source.zip
```

2. Copy the ZIP package into the disconnected scans directory that you added as a data source in BigFix Inventory.
3. To save disk space, remove the ZIP package from the IBM i endpoint.
4. Repeat steps 1, 2 and 3 for all the IBM i endpoints in your infrastructure.

5. Wait for the scheduled import or run it manually.
6. **Optional:** If you want to save disk space on the computer where BigFix Inventory is installed, go to Advanced Server Setting panel and enable the **delete_successfully_imported_scans** parameter. For more information, see: [Advanced server settings \(on page cdi\)](#).

During the data import, BigFix Inventory checks the directory that you added as a data source and imports scan packages that meet the requirements. Scan results are processed and displayed on the reports in the user interface.

Files from processed packages are saved in the `installation_directory/wlp/usr/servers/server1/tmp/datasources/Datasource_ID` directory.

The processed packages are stored in the directory specified in the data source configuration. The packages are not automatically removed. It is recommended to create the audit snapshot once a quarter. Then, back up the packages in a different directory, and remove them from the directory specified in the data source configuration.

Disconnected scanner maintenance and troubleshooting

9.2.5 Available from 9.2.5. You can use the maintenance and troubleshooting tasks related to the disconnected scanner to monitor the scan progress, optimize the processor usage during the scan, understand and fix the most common issues, or learn more about the log files.

Updating the scanner and the scripts on IBM i systems

9.2.5 Available from 9.2.5. Both the scanner and the scripts are updated on a regular basis to introduce improvements and new functions. To update the scanner, once again download all the required files, and copy them into the `/cit` directory in IFS on your IBM i system, or other directory where the previous version is already located.

1. Log in to the BigFix console.
2. In the navigation bar, click **Sites > External Sites > BigFix Inventory v9 > Fixlets and Tasks**.
3. In the upper right pane, select **Download the Disconnected Scanner Package for IBMi (version)**.
4. To run the task, click **Take Action**.
5. **Optional:** If you made any changes to the scan configuration file `configure_scan.sh`, or the software scan configuration file `sw_config.xml` make sure that you back up the files before you extract the package. Otherwise, the files are overwritten.
6. Extract the contents to the disconnected scanner directory.
7. **Optional:** If you recently updated the software catalog in BigFix Inventory, update it also on the system with the disconnected scanner.
 - a. To download the software catalog for disconnected scanner, log in to BigFix Inventory and go to **Management > Data Source**. Choose the disconnected data source, and download the catalog for IBM i.
 - b. Place the catalog in the `config` directory of the disconnected scanner package that you downloaded through fixlet.



Note: The catalog is available for download after uploading the main software catalog and running a data import.

8. **Optional:** If you made any changes to the scan configuration file `configure_scan.sh`, or the software scan configuration file `sw_config.xml`, modify the new file versions to reflect the changes. Look at the back up files that you created in step 5, and rewrite the changes to modify the new files. Do not copy the old back up files as they might differ between the versions.
9. Run the `update.sh` script.

Decommissioning computers with disconnected scanners

9.2.5

Available from 9.2.5. When you decommission a computer that is no longer in use, you can mark it as decommissioned on the Computers report, so that it is not displayed in BigFix Inventory.



You must have the Manage Computer Groups permission to perform this task.

1. In the top navigation bar, click **Reports > Computers**.
2. Click the computer name.
3. Click **Decommission computer**.

The screenshot shows a software report for computer NC047118. At the top right, there is a button labeled 'Decommission computer'. Below it, the text reads 'Software on NC047118' and '3 Software Installations including IBM TCP/IP Connectivity Utilities for i, IBM i, and IBM HTTP Server for i'. Under 'Computer Properties', there is a table with the following data:

DNS Name	Last Seen
NC047118	08/18/2016 05:34 PM
IP Address	Latest Scan Import
9.156.47.118	09/05/2016 01:24 PM
Operating System	
IBM i V7R3M0 PowerPC 64	

After you decommission the computer, this computer does not show up on the current reports. You can still manage the software on the Software Classification panel, but it is marked as uninstalled. The software is still displayed on the All Metrics report.

Uninstalling the disconnected scanner from IBM i systems

9.2.5

Available from 9.2.5. If you no longer need the disconnected scanner to scan the IBM i system, you can uninstall it from any endpoint.

1. Open the default `/cit` directory in IFS on your IBM i system, or other directory that you specified in `configure_scan.sh`.
2. Run the `uninstall.sh` script from QSH.
3. Remove the whole directory where the disconnected scanner was installed.

The scanner is successfully uninstalled, and the data is no longer collected.

Understanding the Deployment Health and Scan Health for IBM i systems

9.2.5 Available from 9.2.5. Deployment Health and Scan Health contain information about your clients and the scans that were initiated. You can view this information on the Deployment Health, and Scan Health widgets on the dashboard, or on the Computers report. Deployment Health excludes all the disconnected datasource endpoints. To properly understand values that are displayed for IBM i systems, see the explanation of particular columns on the Computers report.

Table 34. Explanation of Deployment Health and Scan Health information for IBM i systems

Column Name	Description
Has Outdated Version	Always shows <code>NO</code> .
Has Low Disk Space	Always shows <code>NO</code> .
Is Out of Sync	Always shows <code>NO</code> .
Has Missing Prerequisites	Always shows <code>NO</code> .
Agent Version	Taken from the Agent Version property in the <code>computer.yml</code> file.
Catalog Scan Successful	Shows <code>YES</code> if the imported package contains the <code>catalog_scan.xml</code> file.
Software Tags Scan Successful	Shows <code>YES</code> if the imported package contains the <code>isotag_scan</code> package.
File System Scan Successful	Always shows <code>YES</code> . IBM i endpoints do not run file system scan.
Package Scan Successful	Shows <code>YES</code> if the imported package contains the <code>package_scan.xml</code> file.
Catalog Version	Taken from the Catalog Version property in the <code>computer.yml</code> file.
Endpoint Scanner Catalog Version	Taken from the Catalog Version property in the <code>computer.yml</code> file. It displays only the second part of the whole catalog version, for example 1250110.1.
Last Scan Attempt	Taken from the timestamp used in the final package name.
Scanning Shared Disks	Always shows <code>NO</code> .
Mounted Shared Disks	Always shows <code>''</code> .

Using NFS to automate the transfer of scan results

9.2.5 Available from 9.2.5. You can use NFS to automate the transfer of scan results from your IBM i systems to the BigFix Inventory server. The results will be stored in a shared directory and imported by BigFix Inventory without any additional actions.

You can use NFS in several configurations:

- BigFix Inventory acts as an NFS server
- Each IBM i acts as an NFS server
- You use a separate NFS server

In this procedure, BigFix Inventory installed on Linux acts as an NFS server and shares a directory with an IBM i system that acts as an NFS client. The scan results saved into the shared directory on IBM i system are readily available to BigFix Inventory. The scenario can be adopted to be used in the Windows environment.

1. Configure the BigFix Inventory server to allow transfer by using NFS.

- Install the `nfs-utils` package.
- Configure the firewall to allow the NFS-related traffic from IBM i systems.
- Create a directory for the *disconnected scans* results, and change its ownership and rights.

```
mkdir /disconnected
chown nfsnobody:nfsnobody /disconnected
chmod 755 /disconnected
```

- Add this directory as a data source (on page [cclxxiii](#)).
- Edit `/etc/exports`, and add the directory to the list of shared directories. The following example shows a single IP address of an IBM i system, but you can specify a range of addresses.

```
/disconnected 172.16.0.31(rw, sync)
```

- Refresh the NFS configuration.

```
exportfs -a
```

2. On the IBM i system, create an output directory for scan results and mount the shared directory from the BigFix Inventory server.

This step connects these directories. Whenever you save files into `/bfi_server`, they will also be available in the `/disconnected` directory on the BigFix Inventory server.

```
MKDIR '/bfi_server'
ADDMFS TYPE(*NFS) MFS('IP_address_of_BFI_server:/disconnected') MNTOVRDIR('/bfi_server') OPTIONS(*DFT)
```

- Open the default `/cit` directory, or other directory where you store the disconnected scan scripts, and modify the following property in the `configure_scan.sh` script to save all scan results into the new output directory.

```
PACKAGE_OUTPUT_DIR="/bfi_server"
```

- Run the software scan (on page [cclxxvii](#)) to collect new scan results. When the scan is complete, the package with scan results can be imported to BigFix Inventory by running a data import.

Monitoring the scan progress

9.2.5 Available from 9.2.5. In case of any problems with the software scan, you can check whether the scan was started or completed, and view some extra information about the scanned directories.

Log files

The `run_hw.sh` and `run_sw_and_pack.sh` commands are the `_console.sh` wrappers, and they redirect the script output into the `run_hw_std.log` and `run_sw_and_pack_std.log` files. The default size limit of a single log file is 1024 bytes. The limit can be changed in the `configure_scan.sh` file. To change the size limit provide a new `LIMIT` value. If during the next scan the log files exceed the size limit, the existing files are renamed to `run_hw_std.log.1` and `run_sw_and_pack_std.log.1`. However, only up to two such files can exist per command. Therefore, if `run_hw_std.log.1` and `run_sw_and_pack_std.log.1` already exist, the oldest log files are removed.

Checking the scanner version

To monitor the scan progress, you must have the scanner in version 2.8.0.0001, or higher. To check the current version, open the `/etc/cit/cit.ini` file.

```
CIT_Version = 2.8.0.0001
```

Checking the initiated scans

Information about the started and completed scans is written in the `/etc/cit/logs/traceCIT.log` file.

```
16/07/04 15:18:36 E [main.C] seMain: 308 IMPORTANT: ---START--- Software scan started. CIT Version = '2.8.0.0001'
16/07/04 15:20:14 E [cache.C] 2800 IMPORTANT: File system access on '/' (drive: '/', type: 2, fs: 'os400fs')
16/07/04 15:22:02 E [cache.C] 2835 IMPORTANT: (OK) File system access on '/' (drive: '/', type: 2, fs: 'os400fs')
16/07/04 15:22:02 E [cache.C] 2800 IMPORTANT: File system access on '/QSYS.LIB' (drive: '/QSYS.LIB:', type: 2, fs: 'qsysfs')
16/07/04 15:25:24 E [cache.C] 2835 IMPORTANT: (OK) File system access on '/QSYS.LIB' (drive: '/QSYS.LIB:', type: 2, fs: 'qsysfs')
...
16/07/04 15:41:31 E [main.C] seMain: 483 IMPORTANT: ---STOP--- Software scan finished with return code=0
```

Checking the scanned directories

You can check which directories are currently being scanned in the `/etc/cit/logs/traceCIT_scanned_dirs.log` file.

```

2016-07-04T15:20:14 '/'
2016-07-04T15:20:14 '/iso'
2016-07-04T15:20:14 '/maestro'
...
2016-07-04T15:25:23 '/QSYS.LIB/QAUTOMON.LIB'

```

If the scan is taking too long to complete or hangs on a particular directory, you can stop it by aborting the job that is running the scan. Then, [exclude the directory from scanning \(on page cclxxv\)](#), and repeat the scan.

Troubleshooting the discovery on IBM i

9.2.5

Available from 9.2.5. The troubleshooting section lists the most common issues you might encounter while discovering software and hardware inventory with the disconnected scanner. The main objective of troubleshooting is to determine why something does not work as expected and explain how to resolve the problem.

Log files

Scanner log files are in the following directories:

- */Installation directory/logs/**.
- */etc/cit/logs/**

The default installation directory is */cit*.

Transferring files from IBM i systems

The internal file structure of IBM i systems displayed when connecting over FTP is very complex. However, you can use one of the following commands to make it more UNIX-like. Issue one of these commands after connecting over FTP:

```

cd /
site NAMEFMT 1

```



Tip: For more information, see [Using the FTP Subcommand: NAMEFMT](#).

Next, set the appropriate mode for transferring your files, either binary or text. Use the binary mode for transferring scan results, and text mode to obtain the scanner log files.

Optimizing the processor utilization

By default, processor utilization is not limited during the scan. You can optimize it by setting up the CPU threshold.



Note: Setting the threshold does not guarantee that the CPU consumption is always below the specified value. It fluctuates around that value, sometimes exceeding it and sometimes dropping below it. Temporary peaks are expected. Setting the threshold might lengthen the time of the scan.

1. Edit the `sw_config.xml` file that you downloaded with the scanner.
2. Add two new **cpuThreshold** attributes, one for *FSScanner* and another for *Signatures*. For example:

```
<BigFix>
<CIT>
<XSE version="1.0">
<Plugins>
<Plugin name="FSScanner">
<Attribute name="cpuThreshold" value="20:2:150:850"/>
(...)
<Attribute name="interruptOnTimeout" value="true"/>
</Plugin>
</Plugins>
<Signatures>
<Attribute name="cpuThreshold" value="20:2:150:850"/>
</Signatures>
</XSE>
</CIT>
</IBM>
```

3. Specify one of the following values:

Desired utilization	Value
5%	10:2:50:950
10%	20:2:100:900
15%	20:2:150:850
20%	30:2:200:800
25%	30:2:250:750
50%	50:2:300:500

4. Save the file.

Excluding directories from being scanned

You can exclude directories from being scanned during the catalog-based and ISO tags scan. To do this, edit the `sw_config.xml` file. Specify the directories as values of the **excludeDirectory** attributes.

- To exclude the `/tmp` directory:

```
<Attribute name="excludeDirectory" value="/tmp"/>
```

- To exclude all directories named `tmp`, for example `/tmp` and `/dir/tmp`:

```
<Attribute name="excludeDirectory" value="*/tmp/" />
```

Providing failing packages to BigFix Support

Import of a package is failing

If you see the following errors in the import log, provide the whole scan package.

```
2015-11-06 08:31:07 INFO: Processing scan package: 201511041200-9.128.0.223.zip from: /disconnected/.
2015-11-06 08:31:08 INFO: ETL from FileSystem - SAM::DisconnectedScan
(0x00000000563C63FF - 0x00000000563C6514): Failed
2015-11-06 08:31:08 ERROR: Sequel::UniqueConstraintViolation:
Java::ComIbmDb2JccAm::SqlIntegrityConstraintViolationException: One or more values in the INSERT statement,
UPDATE statement, or foreign key update caused by a DELETE statement are not valid because the primary key,
unique constraint or unique index identified by "1" constrains table "DBO.RAW_DATASOURCE_FILES" from having
duplicate values for the index key.. SQLCODE=-803, SQLSTATE=23505, DRIVER=3.64.104
```

Import of a single file is failing

The import of a single file from the scan package is failing if the import log contains the following message:

```
2015-11-06 10:48:07 WARN: An error occurred parsing XML file 2 in file isotagsearch_0_375.zip.
Install path value was not found.
```

Such a file must be provided with its long path. To determine the path, complete the following steps:

1. Go to the BigFix Inventory installation directory.
2. Search for the failing file:

```
find -name "isotagsearch_0_375.zip"
```

Command output that must be provided to Support:

```
./wlp/usr/servers/server1/tmp/datasources/2/UploadManager/sha1/75/375/isotagsearch_0_375.
zip
```

Common problems

The package is being skipped during the import despite meeting the requirements

Such a problem might occur if you are importing more than one package from an endpoint. BigFix Inventory always imports the latest package and omits all previous ones. The following message about skipping the package is related to irrelevant packages and can be ignored:

```
2015-11-05 09:35:22 WARN: Skipping file: package_scan.xml from: 201510101200-2inarow.zip.
The data was already processed.
```

```
2015-11-05 09:35:22 INFO: Skipping file: package_scan.xml.
Unsupported file in: 201510101200-2inarow.zip.
```

The import might fail if you use the same results of the capacity scan in different packages

This problem might occur if you reuse the capacity scan results for testing purposes. If more than one endpoint has the same scan results, the import fails and the following message is displayed:

```
ERROR: (ImportCapacity-Thread:3) doInsert com.ibm.db2.jcc.am.SqlIntegrityConstraintViolationException:
One or more values in the INSERT statement, UPDATE statement, or foreign key update caused
by a DELETE statement are not valid because the primary key, unique constraint or unique
index identified by "1" constrains table "ADM.LAST_IMPORTED_SCAN" from having duplicate values
for the index key.. SQLCODE=-803, SQLSTATE=23505, DRIVER=3.64.104
```

Ensure that you run the capacity scan for each of your endpoints and use appropriate results.

The software or ISO tags scan fails and no results are created

Run the following command in QSH, and then repeat the scan.

```
export QIBM_MULTI_THREADED=Y
```

The software is discovered and bundled, but there are no PVU calculations

The problem might be caused by an incorrect epoch, which points to the date in the future. The epoch is included in file names of the capacity scan results, for example `t1m_hw_201603021330_1456969683.xml`. Open the `capacity.zip` archive, and check whether the epochs are correct. You will have to convert the epoch into a normal date. On Linux, you can do it by running the `date -d @epoch` command.

The host name of an endpoint provided by the disconnected scanner is different than expected

The host name that is provided by the disconnected scanner and a regular scanner can be different. The regular scanner uses the host name that is provided by BigFix, while the disconnected scanner checks the host name by running the `hostname` command, or, if not available, through the scanner check.

Opening a support case

If none of the above solutions work for you, open a support case and provide the following files to BigFix Support:

- package with scan results
 - `/Installation directory/run_hw_std.log`
 - `/Installation directory/run_sw_and_pack_std.log`
- `/Installation directory/warning*.xml` file
- scanner logs from `/etc/cit/logs`
- the console output of a failed step in the debug mode, if possible. The example below generates the `/cit/failure.log` file with the output:

```
export QIBM_MULTI_THREADED=Y
/qsys.lib/qtivcit.lib/wscansw.pgm -s -c /cit/sw_config.xml -i /cit/CIT_catalog_I5OS.xml
```

```
-o /cit/catalog_scan_man.xml -e /cit/warning_man.xml -debug
>/cit/failure.log 2>&1
```

Monitoring the import log

When you import a package with scan results to BigFix Inventory, the import log (**Management > Data Imports**) displays information and warning messages related to your package. If the package or its content does not meet the requirements, it might be skipped during the import, and the scan results will not be displayed in BigFix Inventory. While monitoring the import log, you can expect the following messages:

```
info: Processing scan package: package from: directory.
info: Invalid file name. The scan package must match the pattern YYYYMMDDHHMM-*.zip. Skipping
package: package.
info: Invalid Catalog Version property in the computer.yml file: version. The property must match the pattern
<Number>.<Number>. The catalog might be reported as outdated.
warn: The disconnected scans location does not exist: directory.
warn: Invalid file name. The scan date must be later than year 2000. Skipping package: package.
warn: Skipping corrupted scan package: package. The following error occurred: error.
warn: Missing computer.yml file. Skipping package: package.
warn: Missing property: property in the computer.yml file. Skipping package: package.
warn: Invalid computer.yml file. Skipping package: package. The following error occurred: error.
warn: The value of the endpointID property in the computer.yml file is longer than 512 characters. Skipping
package: package.
warn: The value of the endpointID property in the computer.yml file is empty. Skipping package: package.
warn: Missing property: endpointID in the computer.yml file. Skipping package: package.
warn: Importing older scan package from: date for computer: computer_id. This computer already has data
from: date.
warn: The file is larger than 2097150 bytes. Skipping file: scan_file from: package for
computer: computer_id.
warn: The following error occurred when processing the package_scan.xml file: error.
warn: The insertion of the operating_system package data into the database failed on record content. The
following error occurred: error.
```

(Optional) Manually installing the scanner and running the scans on IBM i systems

9.2.5 Available from 9.2.5. scripts that automate the scans are not appropriate for your environment, and you need to customize the execution of scans without changing the required scan frequency, refer to this appendix. The appendix describes how to run each scan, how to collect the results, and finally how to create the final package that can be uploaded to BigFix Inventory manually. The automated scripts, however, remain a recommended way of discovering on IBM i systems.

(Optional) Manually installing the scanner

9.2.5 Available from 9.2.5. If the automated scripts are not appropriate for your environment, you can manually install the scanner. You will also need to run the initial capacity scan, create the `computer.yml` file, and schedule regular capacity scans. These tasks are optional and can be completed with the automated scripts.

Installing the scanner

9.2.5 Available from 9.2.5. Install the scanner on each of your IBM i systems to be able to initiate the software and capacity scans.



Note: This task is optional, because it can be completed by an automated script. Use it only for advanced scenarios.

Install the scanner by running the following commands from QSH:

```
cd /cit
./install_scanner.sh
```

Running the initial capacity scan

9.2.5 Available from 9.2.5. The initial capacity scan is needed only to gather detailed information about the scanned IBM i system. You will use this information to create the `computer.yml` file that describes your IBM i system, and is required for each import of new scan results. All subsequent capacity scans will be initiated from a script that schedules them to run automatically every 30 minutes.



Note: This task is optional, because it can be completed by an automated script. Use it only for advanced scenarios.

Go to `/cit`, and run the `run_hw.sh` script.

The results were saved in `/cit/output/hw_scan_all/last.xml`. You can now create the `computer.yml` file.

Creating the `computer.yml` file

9.2.5 Available from 9.2.5. The `computer.yml` file contains information about the scanned IBM i system, and is used to catalog scan results and correctly display them in the user interface. Each package with scan results must contain this file, but you can reuse the file as long as it is unique for each IBM i system.



Note: This task is optional, because it can be completed by an automated script. Use it only for advanced scenarios.

1. Go to `/cit`, or `/cit/config`, and edit the `computer.yml.sample` file. The file contains the following values:

```

endpointID: 6552CFF:5
Agent Version: IBM_doc_v.1.0
Catalog Version: 1250110.0
Operating System: IBM i V7R1M0 PowerPC 64
DNS Name: NC198511
Computer Name: NC198511
IP Address: 198.51.100.0

```

2. Take the values from the results of the capacity scan that are in the `/cit/output/hw_scan_all/last.xml` file. The table provides some details on how to retrieve them.

Attribute	Description	Value
endpoint-ID	Mandatory. Unique ID of the computer. It must contain the hardware serial number and LPAR ID in the <code><HWSerial>:<LparID></code> format, for example <code>6552CFF:5</code> .	In capacity results: <pre><HWSerial>6552CFF</HWSerial> <LparID>5</LparID></pre>
Agent Version	Optional. You can use this attribute to provide information that lets you differentiate between scan results. For example, if you used the manual procedure to gather them, you can call it <code>IBM_doc_v.1.0</code> , or choose other names for automated processes.	-
Catalog Version	Optional. Versions of the software catalog and its custom content, separated with a dot (<code>.</code>), for example <code>1250110.0</code> .	To view these versions, in BigFix Inventory go to Management > Catalog Upload .
Operating System	Mandatory. Operating system of the computer. It must start with <code>IBM i</code> , but you can add more information, and have it displayed in the user interface.	In capacity results: <pre><OSName>IBM i</OSName> <OSVersion>V7R1M0</OSVersion> <OSArch>PowerPC</OSArch> <OSKernelMode>64</OSKernelMode></pre>
Computer Name	Mandatory.	In capacity results:

Attribute	Description	Value
	Name of the computer, used for identification. It can be an IP address, or any name or ID that lets you recognize the computer.	<code><Hostname>NC046189</Hostname></code>
DNS Name	Optional. Domain name, used for identification. It can be the same as Computer Name.	-
IP Address	Optional. IP address of the computer, used for identification. For multiple IPs, place each address on a new line that starts with a space.	In capacity results: <code><Address IsKey="1">9.156.46.189</Address></code> Syntax for multiple IPs: <code>IP Address: 198.51.100.0 198.51.100.1 198.51.100.2</code>

3. Save the file as `computer.yml` in the same directory.

Scheduling regular capacity scans

9.2.5 Available from 9.2.5. The capacity scan must be initiated every 30 minutes to gather current results for virtualized environments, in which capacity can dynamically change depending on allocated resources. After you run the script, the scan will be run with such a frequency automatically. The results do not have to be imported to BigFix Inventory with such a high frequency. You can import only those that are already available, and include the rest in the next imported package.



Note: This task is optional, because it can be completed by an automated script. Use it only for advanced scenarios.

Go to `/cit`, and run the `schedule_scans.sh` script.

After you run the script, the capacity scan will be initiated every 30 minutes, each time creating two output directories:

`hw_scan_all`

All capacity scan files and their signatures. They are required for audit purposes to prove that the capacity scan was initiated every 30 minutes.

`hw_scan_unique`

Unique capacity scan files that contain only changes to the capacity.

Additionally, 48 instances of the BFI-HW job that run the `run_hw.sh` script are created. Since the capacity scan must be run every 30 minutes, two jobs are created for each hour.

All capacity results will be automatically copied from their directories and included in the common package with all scan results after you run the software scan.

(Optional) Manually running the scans and preparing the results

9.2.5 Available from 9.2.5. After you install the scanner, you use the QSH to initiate the software and capacity scans. Each type of scan creates a different output that must be copied to the final package that is imported to BigFix Inventory.

Before you start the scans, create two directories that will make it easier to manage the scan results.

/scans

In this directory, you will store all scan results that you want to include in a single import to BigFix Inventory. After you copy the results to this directory, you must compress and rename it, so that it has the right format. Such a final package is then imported to BigFix Inventory.

/source

In this directory, you will store signatures of all scan results that you copied to the `/scans` directory. In the case of the ISO tags scan, also the original output file. These files are required for audit purposes to ensure that scan results were not altered.

All details regarding the final package are described later in this document.

Running the capacity scan

9.2.5 Available from 9.2.5. In the capacity scan, the scanner reports the system type and details about the physical processor. If applicable, it also collects information about the guest operating system and logical partitions. This data is required to calculate the PVU and RVU MAPC consumption.



Note: This task is optional, because it can be completed by an automated script. Use it only for advanced scenarios.

The capacity scan must be run every 30 minutes to collect current results.

1. In QSH, run the following command to initiate the capacity scan:

```
/qsys.lib/qtivcit.lib/wscanhw.pgm -nolock -s -c tlm_hw_config.xml -o tlm_hw_<DATE>_<SECONDS>.xml
```

Where:

-s

Creates a signature for the output file. The signature must be stored for audit purposes.

-c tlm_hw_config.xml

Path to the `tlm_hw_config.xml` configuration file that you downloaded with CIT.

-o tlm_hw_<DATE>_<SECONDS>.xml

The output file.

DATE - Date in the YYYYMMDDHHMM format, for example 201511101248.

SECONDS - Number of seconds that passed since January 1st 1970, for example 1447177713.

 **Tip:** You can retrieve the date and the number of seconds by running the `date +%Y%m%d%H%M-%s` command on Linux.

2. Compress the output file into a `capacity.zip` package. The package can contain more than one output file.
3. Copy the `capacity.zip` package to the common `/scans` directory, in which you store all scan results.

Running the catalog-based scan

9.2.5 Available from 9.2.5. In the catalog-based scan, the scanner creates scanner catalogs that are based on the catalog that you attach while initiating the scan. Based on these catalogs, the scanner discovers exact software matches.



Note: This task is optional, because it can be completed by an automated script. Use it only for advanced scenarios.

1. Download the software catalog for IBM i systems from your BigFix Inventory server:
`https://hostname:port/sam/catalogs/CIT_catalog_I5OS.xml.bz2`. The catalog is available after uploading the main software catalog and running a data import. Next, extract it and upload to a directory in IFS on your IBM i system.
2. In QSH, run the following command to initiate the catalog-based scan:

```
/qsys.lib/qtivcit.lib/wscansw.pgm -s -c sw_config.xml -i CIT_catalog_I5OS.xml -o catalog_scan.xml
```

Where:

-s

Creates a signature for the output file. The signature must be stored for audit purposes.

-c sw_config.xml

Path to the `sw_config.xml` configuration file that you downloaded with the scanner.

-i CIT_catalog_I5OS.xml

Path to the software catalog for IBM i systems.

-o catalog_scan.xml

The output file.

3. Copy the output `catalog_scan.xml` file to the `/scans` directory, in which you store all scan results.
4. Copy the corresponding signature to the `/source` directory.

Running the ISO tags scan

9.2.5 Available from 9.2.5. In software identification tags scan, the scanner searches for ISO tags that are delivered with software products. After you upload the final package with scan results to BigFix Inventory, the tags are processed. Based on the information that they contain, the software is discovered.



Note: This task is optional, because it can be completed by an automated script. Use it only for advanced scenarios.

The ISO tags scan requires the software catalog for ISO tags. You can find it in the downloaded `ibm_i_cit_and_configs.zip` package.

1. In QSH, run the following command to initiate the ISO tag scan:

```
/qsys.lib/qtivcit.lib/wscansw.pgm -s -c sw_config.xml -i i5_iso_catalog.xml -o iso_scan.xml
```

Where:

-s

Creates a signature for the output file. The signature must be stored for audit purposes.

-c sw_config.xml

Path to the `sw_config.xml` configuration file that you downloaded with CIT.

-i i5_iso_catalog.xml

Path to the ISO software catalog for IBM i systems.

-o iso_scan.xml

The output file. It contains paths to all scanned ISO tags.

2. The scan creates the `iso_scan.xml` output file. Back up this file, because you will need to store the original output and its signature for audit purposes.
3. The `iso_scan.xml` output file lists all scanned ISO tags. You must now gather and copy these tags to a separate directory and create a text file that lists all of them. Create such a directory, for example `/isotag_scan`, and then complete the following steps:

- a. Open the `iso_scan.xml` file. Each entry in this file contains a path to the scanned ISO tag file, for example:

```
<MatchedSignature guid="ISO-TAGS">
<Variable name="INSTALL_PATH" value="/iso-swid/IBM_Capital_Regulatory_Reporting_Add-on-License-4.9.0.swidtag"/>
<Variable name="IS_INSTALLED" value="true"/>
</MatchedSignature>
```

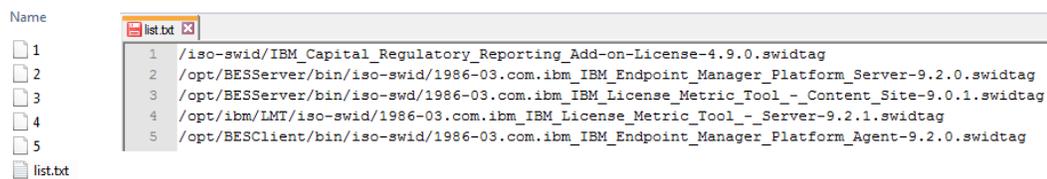
- b. Copy each path to a new file called `list.txt`. Place each path on a separate line, for example:

```
/iso-swid/IBM_Capital_Regulatory_Reporting_Add-on-License-4.9.0.swidtag
/iso-swid/1986-03.com.ibm_IBM_BigFix_Inventory_Server-9.2.1.swidtag
```

- c. Copy each of the scanned ISO tags to the `isotag_scan` directory. Rename the files so that they contain numeric values according to their order in the `list.txt` file.

For example, if the `/iso-swid/IBM_Capital_Regulatory_Reporting_Add-on-License-4.9.0.swidtag` file is written on the first line in the `list.txt` file, rename this ISO tag to `1`.

The following image shows sample ISO tag files and their paths in the `list.txt` file.



4. Compress the `isotag_scan` directory into a `[zip|tar.gz|tgz]` package, and copy it to the `/scans` directory, in which you store all scan results.

Warning: The `list.txt` file must be the first file in the archive. You can achieve it by compressing only this file, and then adding the rest of files to the archive.

5. Copy the backup `iso_scan.xml` file and its signature to the `/source` directory.

Running the package data scan

9.2.5 Available from 9.2.5. In the package data scan, the scanner searches the system registry to gather information about packages that are installed on your endpoints. After you upload the final package with scan results to BigFix Inventory, the discovered packages are compared with the software catalog resulting in software discoveries.



Note: This task is optional, because it can be completed by an automated script. Use it only for advanced scenarios.

1. In QSH, run the following command to initiate the package data scan:

```
/qsys.lib/qtivcit.lib/wscanvdp.pgm -s -o package_scan.xml -c vpd_config.xml
```

Where:

-s

Creates a signature for the output file. The signature must be stored for audit purposes.

-o package_scan.xml

The output file.

-c vpd_config.xml

Path to the configuration file that you downloaded with the scanner.

2. Copy the output `package_scan.xml` file to the `/scans` directory, in which you store all scan results.
3. Copy the corresponding signature to the `/source` directory.

9.2.5 IBM i Creating the final package

9.2.5 Available from 9.2.5. Gather all scan results and compress them into a ZIP package that can be imported to BigFix Inventory.



Note: This task is optional, because it can be completed by an automated script. Use it only for advanced scenarios.

BigFix Inventory does not import single scan results, but rather a whole package that has the right file name and format. The package must contain the following files: `computer.yml`, `catalog_scan.xml`, `capacity.zip`, `isotag_scan.[zip|tar.gz|tgz]`, `package_scan.xml` and `source.zip`. It must be copied to your disconnected scans directory that you added as a data source.

Considerations

- If you upload a package with older results, the current ones are overwritten.
- Discovered software is reported with the import time, not the package timestamp.
- Import only one package per endpoint. Otherwise, capacity results might be skipped.
- Processed packages are not removed from your disconnected scans directory. The ones already imported are skipped.

1. Compress the `/source` directory, in which you store signatures, into the `source.zip` package, and copy it to the `/scans` directory.
2. In the `/scans` directory, create a `computer.yml` file, and fill it in like in the following example:

```
endpointID: 6552CFF-5
Agent Version: IBM_doc_v.1.0
Catalog Version: 1250110.0
Operating System: IBM i V7R1M0 PowerPC 64
DNS Name: NC198511
Computer Name: NC198511
IP Address: 198.51.100.0
```

Where:

endpointID - (mandatory) Unique ID of your computer. It must contain the hardware serial number and LPAR ID in the `<HWSerial>-<LparID>` format, for example `6552CFF-5`. You can retrieve these values from the results of the capacity scan.

Agent Version - (optional) You can use this parameter to describe a way that you used to create the final package. For example, if you used the manual procedure, you can call it `IBM_doc_v.1.0`, or choose other names for automated processes. This information can be used for troubleshooting failing imports.

Catalog Version - (optional) Versions of the software catalog and its custom content, separated with a dot (`.`), for example `1250110.0`. You can view the versions in BigFix Inventory by going to **Management > Catalog Upload**.

Operating System - (mandatory) Operating system of your computer. It must start with `IBM i`, but you can add more information, and have it displayed in the user interface. You can retrieve information about the name, version, and architecture from the results of the capacity scan.

DNS Name - (optional) Domain name, used for identification. It can be the same as Computer Name.

Computer Name - (mandatory) Name of your computer, used for identification.

IP Address - (optional) IP address of your computer, used for identification.

3. Compress the `/scans` directory into a ZIP package, and rename it to `<UTC_DATE>-<endpointID>.zip`, for example `201510191213-6552CFF-5.zip`, where:

UTC_DATE - Date of scan in Coordinated Universal Time (UTC). It must have the following format: `YYYYMMDDHHMM`. The file name must start with this timestamp to ensure correct ordering.

endpointID - The same endpoint ID that you used in the `computer.yml` file. It is required to distinguish between various packages and computers.

4. Copy the ZIP package into your disconnected scans directory that you added as a data source in BigFix Inventory. The package is imported to BigFix Inventory after you run a data import.
5. Run a data import.

During the data import, BigFix Inventory checks the directory that you added as a data source and imports all scan packages that meet the requirements. Scan results are processed and displayed on reports in the user interface.

Files from processed packages are saved in `installation_directory/wlp/usr/servers/server1/tmp/datasources/Datasource_ID`.

Performing optional configuration

You can perform optional configuration tasks to further customize the application.

Data retention periods

Table 35. Data retention periods

Type of data	Period for which the data is retained	Changing the default period
Raw scan data: results of the file system scan, package data scan, and application usage statistics scan	7 days	To change the default period, change the value of the raw_data_api_history_keep_days parameter on the Advanced Server Settings panel. For more information, see: Advanced server settings (on page cdii) .
Raw utilization data that is displayed on the Resource Utilization report	90 days	To change the default period, go to the Server Settings panel and change the value in the Retention Period section. For more information, see: Configuring data retention period for raw utilization data (on page ccxcviii) .
Data that is needed to calculate utilization of license metrics that is displayed on the All Metrics report	Always	The period cannot be changed.
Data about an old complex signature that is substituted with a new complex signature	Until the software is detected based on the new signature	The period cannot be changed.

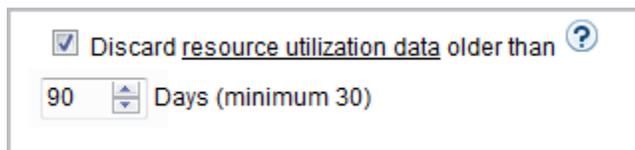
Configuring data retention period for raw utilization data

By default, raw data about utilization of license metrics is stored in the BigFix Inventory database for 90 days. You can change the period depending on the period for which you want to generate reports of raw utilization data.

Retention period is set by default in fresh installations of BigFix Inventory. If you upgraded from Software Use Analysis, your original settings were preserved.

Raw utilization data is related to license metrics that are reported by products, which implemented the ISO/IEC 19770-4:2017 standard. By default, retention period for raw utilization data is set to 90 days. After that period, the data is removed from the BigFix Inventory database. Software items for which raw license utilization is not reported for a period longer than the specified one are no longer displayed on the Resource Utilization report.

1. In the top navigation bar, click **Management > Server Settings**.
2. Specify the period after which you want the data to be removed from the database. If you want to disable the data retention period, clear the check box. Then, click **Save**.



The data that is older than the retention period is removed from the database during every import.

If the retention period was not configured or you shortened it considerably, removal of the large amounts of data might lengthen the first import after the change. During the consecutive imports, smaller amounts of data are removed and should not have considerable impact on the import time on condition that the import runs regularly.

Setting up computer properties

You can specify computer properties that are to be gathered from the computers in your infrastructure. You can then use those properties to filter data on the Computers report and to assign computers to computer groups.

 You must be an Administrator to perform this task.

9.2.5 **IBM i** You cannot set up additional computer properties for IBM i computers.

1. To view the properties that are specified for the computers in your infrastructure, click **Management > Computer Properties**.
2. To add a property, click **New**.
3. In the **Create Computer Property** pane, specify the name of the property to be displayed in BigFix Inventory. Select the property from the **Data Source Property** list and click **Create**.

 **Tip:** When you start typing in the property name, a list of possible values is displayed. It contains all properties whose names contain the letters that you entered in the specified order, regardless of whether the letters occur immediately one after another. For example, if you type `path`, the list might contain properties such as `Patches Applied - Solaris`, because letters `p, a, t, h` occur in this order in the property name.

To have the computer property displayed, wait for the next scheduled import or run it manually. Each computer property requires that a relevant analysis is activated. After the import finishes, check the import log for warning messages that indicate that an analysis related to a particular property is not activated. For example:

```
WARN: Analysis 'Software Scan Status', bound to Computer Property
'Status of catalog-based scan', is not activated and will not be imported.
```

If you find such a warning, activate the required analysis to have the computer property displayed.

Deprecated: Package and usage properties

In addition to computer properties, BigFix Inventory also uses package properties, UNIX package properties, and usage properties to retrieve data about your software and its usage. The data is originally gathered by analyses, and the properties are links between these analyses and BigFix Inventory. All properties are configured by default and do not require any action from you.

9.2.12 Starting from application update 9.2.12, editing package and usage properties becomes deprecated.

Package properties

Package properties retrieve data about software that is installed on the Windows operating systems. The data is gathered by the Installed Windows Applications analysis that retrieves the list of your software from the Windows Registry. The raw data collected by the analysis can also be viewed in the BigFix console by checking the analysis results.

To view the property in BigFix Inventory, go to the following URL: https://<bfi_host>:<port>/management/sam/package_properties, and select the property. The property must be linked to the relevant Fixlet site that contains the analysis, otherwise the results will not be collected.

Edit Package Property

Name*

Link to Datasource

Data Source Property*

- Installed UNIX Packages*
- Installed Windows Applications List
 - ActionSite*
 - Installed Windows Applications*
- Installed Windows Applications List
 - IBM BigFix Inventory*
 - Installed Windows Applications*
- Installed Windows Applications List
 - IBM License Reporting*
 - Installed Windows Applications*

Data retrieved from the analysis is displayed on the Package Data report. To access it, click **Reports > Package Data**.

UNIX package properties

UNIX package properties retrieve data about software that is installed on the UNIX operating systems. The data is gathered by the UNIX Installed Packages analysis that retrieves the list of your software from the package management system. The raw data collected by the analysis can also be viewed in the BigFix console by checking the analysis results.

To view the property in BigFix Inventory, go to the following URL: https://<bfi_host>:<port>/management/sam/unix_package_properties, and select the property. The property must be linked to the relevant Fixlet site that contains the analysis, otherwise the results will not be collected.

Edit UNIX Package Property

Name*

Link to Datasource

Data Source Property*

Installed Provisioning Profiles
MDM Beta
Installed Profiles - Apple iOS

Installed Provisioning Profiles
Mobile Device Management
Installed Profiles - Apple iOS

Installed Unix Packages List
IBM BigFix Inventory
Installed UNIX Packages

Installed Unix Packages List

Data retrieved from the analysis is displayed on the Package Data report. To access it, click **Reports > Package Data**.

Usage properties

Usage properties retrieve data about the use of your software. The data is gathered by the Application Usage Statistics analysis.

To view the property in BigFix Inventory, go to the following URL: https://<bfi_host>:<port>/management/sam/app_usage_properties, and select the property. The property must be linked to the relevant Fixlet site that contains the analysis, otherwise the results will not be collected.

Edit Application Usage Property

Name*

Link to Datasource

Data Source Property*

App Feedback
MDM Beta
Apps - Android / Apple iOS

App Feedback
Mobile Device Management
Apps - Android / Apple iOS

Application Usage
IBM Endpoint Manager for Software Use Analysis
Application Usage Statistics

Application Usage

Data retrieved from the analysis is displayed on the Metering Data report. To access it, click **Reports > Metering Data**.

Deprecated: Disabling subcapacity calculations

Subcapacity licenses are based on the number of processor value units (PVU) or resource value units (RVU) that can be used by the product not the number of all processor cores that are available on the physical server. Such a licensing scheme allows for reducing the cost of licenses for products that are eligible for PVU or RVU subcapacity. If no subcapacity products are installed in your environment, you can disable subcapacity calculations to improve the performance of the import.



Important: **9.2.5** Starting from application update 9.2.5, the option to entirely disable subcapacity calculations is no longer available. You can disable PVU and RVU MAPC calculations, both full capacity and subcapacity, per computer group. For more information, see: [Setting up computer groups \(on page cxc\)](#).

To disable subcapacity calculations, go to <https://hostname:port/management/feature> and clear the **Subcapacity** checkbox. Then, click **Save**.



Tip: You might restart the BigFix Inventory server to ensure that subcapacity calculations are disabled.

After you disable subcapacity calculations, subcapacity data is not uploaded to BigFix Inventory server during the import. Performance of the import is improved. Fixlets and tasks that are related to subcapacity configuration and VM managers remain available and relevant in the BigFix console. However, you do not have to run them. Additionally, the following elements of the BigFix Inventory user interface are hidden:

- Widgets:
 - IBM Capacity Data Completeness
 - IBM PVU Subcapacity
 - BigFix Software Classification
- Management panels:
 - BigFix Software Classifications
 - VM Managers
 - Metric Table Upload
 - Part Numbers Upload
- Reports:
 - IBM PVU Subcapacity
 - All Metrics
 - Audit Trail
 - Hardware Inventory

The following permissions are removed from the user roles:

- Manage Software Classification
- Manage VM Managers and Servers
- View Audit Trail

- View Hardware Inventory
- View License Metrics

Setting the home page

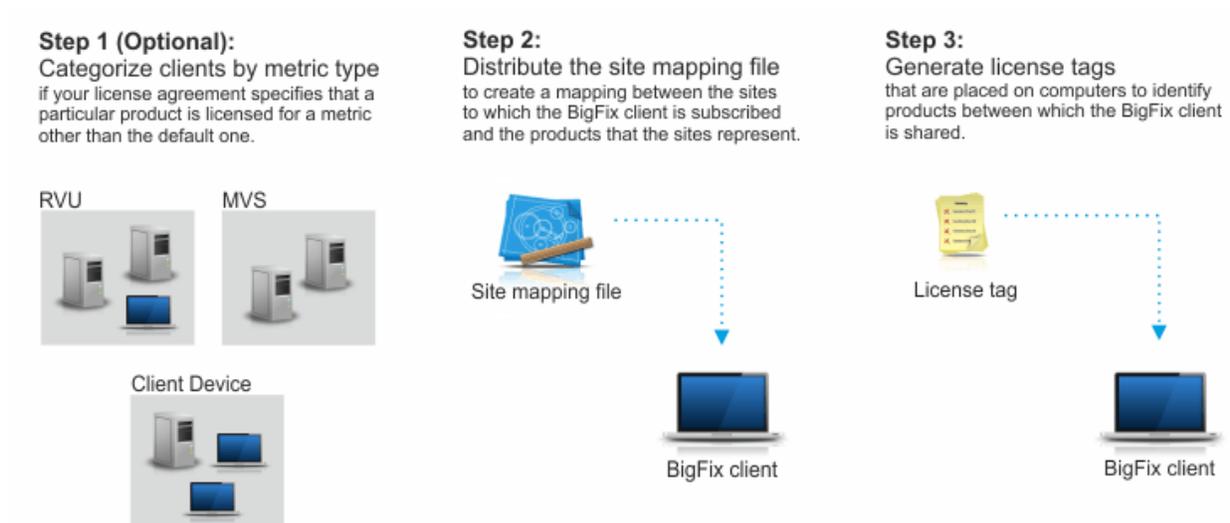
If you use a particular report or panel frequently, you can set it as BigFix Inventory home page.

- To set a new home page:
 1. Open the report or panel that you want to set as the home page.
 2. Hover over the **User** icon , and click **Set as home page**.
- To return to the default home page:
 1. Hover over the **User** icon , and click **Profile**.
 2. Under the **Home Page** field, click **Clear**.

Classifying the BigFix client used by multiple BigFix products

BigFix client is a component that is common for all products from the BigFix family. If you have a number of BigFix products installed, the client must be classified as used by these products. Because the client is usually installed on all computers in the enterprise, its manual classification can be time-consuming. To reduce the amount of time and resources that are needed to perform this task, you can generate license tags. The tags mark the client as used by the BigFix products that are installed in your infrastructure. The information is later on reflected in BigFix Inventory without the need of manual classification.

Below is the overview of the procedure. For detailed instructions and troubleshooting information, see: [Activating the license counting process](#).



Step 1: (Optional) Categorize clients by metric type

License metric that is assigned to the BigFix client depends on the operating system on which it is installed. For example, when the client is installed on a computer that runs on Windows 7, it is assigned the Client Device metric. However, when it is installed on a computer that runs on Red Hat Enterprise Linux, it is assigned the RVU MAPC metric. If your license agreement specifies that a particular product is licensed for a metric other than the default one, you can override the default metric.

For example, you might have BigFix Patch, BigFix Compliance and BigFix Protection deployed in your environment. Your license agreement specifies that the first two products are licensed for RVU MAPC but the last product is licensed for the Client Device metric. You can specify that every BigFix client that is used by these three applications is assigned the RVU MAPC metric for BigFix Patch and BigFix Compliance but the Client Device metric for BigFix Protection.

Step 2: Distribute the site mapping file

Create a mapping between the sites to which a BigFix client is subscribed and the products that the sites represent. The information is required to indicate products that use the client. Repeat the step every time you add a product from the BigFix family.

Step 3: Generate license tags

The tags are placed on the computers and contain information about the product name, license metric, and version of the BigFix client. After the software scan, information from the license tag is imported to BigFix Inventory.

Results

After you complete this procedure, wait for the scheduled software scan and the import of data. Then, BigFix client is listed as a component under each BigFix product. The name of the component consists of the product name, license metric, and version of the client.

For example, a BigFix client that is used by BigFix Compliance and BigFix Inventory, and is assigned the RVU MAPC metric for both products is listed as:

- BigFix Compliance RVU 9.2 under BigFix Compliance
- BigFix BigFix Inventory RVU 9.2 under BigFix Inventory



Tip: When you complete this procedure, you no longer have to assign the BigFix client to multiple BigFix products on the Software Classification panel.

Upgrading

A new version of BigFix Inventory is released periodically, typically at the end of each calendar quarter. Upgrade to the new version regularly to take full advantage of new features and application fixes.

Upgrading to BigFix Inventory 9.2.17

You can upgrade to BigFix Inventory 9.2.17 from all 9.x versions. Upgrade of the entire environment consists of upgrading the BigFix Inventory server, VM Manager Tool, and the scanner as well as restarting actions created by fixlets that were updated. You cannot upgrade between different operating systems.

 **Important:** It is strongly recommended, especially in big environments, to first perform the upgrade in the test environment. To do this, back up your production database, restore it on the test server, and perform the upgrade there. If it is successful, perform the upgrade on the production server.

Ensure that no files or directories in the BigFix Inventory installation directory are open in an external program. For example, that no log files are open in a text editor. Also, ensure that no PDF reports are being generated during the server upgrade. Otherwise, the upgrade might fail with error code 41.

For information about space requirements, see: [Hardware requirements for the server on Windows \(on page cvii\)](#) or [Hardware requirements for the server on Linux \(on page cxiii\)](#).

1. If the computer on which the BigFix server is installed does not have the Internet access, [update the fixlet site \(on page ccvii\)](#).
Otherwise, the site is updated automatically after its new version is released.



Note: To check whether you have the latest version of the fixlet site, see: [Checking the version of the fixlet site \(on page ccvii\)](#).

2. [Create a back up \(on page cccx\)](#) of the BigFix Inventory database, installation registry, and the server files to be able to restore the server if the upgrade fails.
3. Upgrade the server. The easiest and preferred method is by using a fixlet. To use this method on Linux, the server must be installed by a root user. Alternatively, you can upgrade the server in [interactive \(on page cccxi\)](#) or [silent \(on page cccxii\)](#) mode.
 - a. In the BigFix console, go to **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
 - b. Select **Upgrade to the latest version of BigFix Inventory**, and click **Take Action**.
 - c. Select the computer on which the server is installed, and click **OK**.



Tip: To check the status of the upgrade, in the BigFix console, go to **Actions** and look for the upgrade action. If the upgrade fails, [check the logs \(on page dccxviii\)](#).

During the upgrade of the server, new software catalog and PVU table are uploaded to BigFix Inventory. The versions that are uploaded are the newest ones that were available during the release of the application update to which you are upgrading. If the versions that are already available in BigFix Inventory are newer than the ones that are uploaded as part of the upgrade, the upload is skipped.

4. Upgrade the VM Manager Tool to the latest version.
 - a. In the BigFix console, go to **Fixlets and Tasks**.
 - b. Select **Update VM Manager Tool to version *version number***, and click **Take Action**.
 - c. Select the computer on which the VM Manager Tool is installed, and click **OK**.
5. If the new version of the scanner contains security fixes, fixes of problems that appear in your environment, or new features that you want to use, upgrade the scanner to the latest version. For information about changes introduced in each version of the scanner, see the [BigFix Inventory wiki](#).

To update the scanner, perform the following steps:

- a. In the BigFix console, go to **Fixlets and Tasks**.
 - b. Select **Install or Update Scanner**, and click **Take Action**.
 - c. Select all relevant computers, and click **OK**.
6. If some of the fixlets were updated on the action site, and you want to use features for which running the new version of the fixlet is required, restart actions created by these fixlets. Otherwise, the new features might not work properly.

To restart the actions, perform the following steps:

- a. In the BigFix console, go to **Actions**.
 - b. Look for actions started by the fixlets that were updated, right-click each action, and click **Stop Action**.
 - c. Run the updated version of the fixlets.
7. If the server database schema changed between the versions, update it after the upgrade of the server completes. For the MS SQL database, the user who is configured in this database must have the `db_owner` role. Log in to BigFix Inventory, and click **Update Schema**.

 **Important:**

The update can take from a couple of minutes to a couple of hours depending on your environment size and how long you have BigFix Inventory installed.

8. To ensure that all changes are processed, run an import. In BigFix Inventory, click **Management > Data Imports**, and click **Import Now**.

 **Important:** The first import after the upgrade might last longer because it includes the new software catalog.

After the upgrade, the following reports and panels are removed from the menu as they become deprecated.

- Usage Properties
- Package Properties

- UNIX Package Properties
- Current State Values

Table 36. Considerations after the upgrade

Version from which you are upgrading	Considerations
Earlier than 9.2.11	<p>Information about version of scanned files is re-imported to show more details. The data is displayed on the Scanned File Data report in the File Version column.</p> <p>The data is re-imported during a number of consecutive imports to avoid performance issues. Each consecutive import contains data from 10000 computers. The number of imports during which the data is re-imported is specified in the resync_imports parameter that is available on the Management > Advanced Server Settings panel.</p>
Earlier than 9.2.1	<p>If you had computer groups defined, subcapacity calculations are disabled for those groups after the upgrade. As a result, some of the panels and reports are not visible for users who are assigned to those computer groups. To re-enable subcapacity calculations, go to Management > Computer Groups, open a computer group, and select PVU, RVU MAPC and VPC. Then, click Save. For more information, see: Setting up computer groups (on page cxc).</p>

Updating the fixlet site

The content of the BigFix Inventory fixlet site can be periodically modified. New fixlets, tasks, and analyses can be added. The existing ones can be changed or might become obsolete due to functionality changes. If the BigFix server is installed on a computer with the Internet access, the BigFix Inventory fixlet site is updated automatically whenever the updates are available. However, if the server is installed on a computer without the Internet access, you must update the fixlet site manually. Start by checking whether the fixlet site that you are currently using is up-to-date. If a newer version of the fixlet site exists, download the site content by using the Airgap tool. Then, cache the files on the BigFix server by using the BES Download Cacher.



Note: The procedures in this section describe the extraction mode for upgrading the fixlet site. However, you can use non-extraction mode as an alternative. To explore the Airgap tool capabilities, see: [Using the Airgap tool](#).

Checking the version of the fixlet site

Compare the current version of the fixlet site with the latest version that was published. If a newer version is available, updated the content of your fixlet site.

1. To check what is the latest version of the fixlet site, open the <http://sync.bigfix.com/cgi-bin/bfgather/ibmforsua> site, and look for the `version` line.
2. To check what is the current version of the fixlet site that you are using, open the BigFix console and click the name of the fixlet site. The version of the site is displayed on the **Details** tab.

If a newer version of the fixlet site is available, update the content of your fixlet site.

Updating the content of the fixlet site on Windows

If the BigFix server is installed on a Windows computer without the Internet access, use the Airgap tool to download the content of the fixlet site to a Windows computer with the Internet access.

You need a Windows computer with the Internet access.

1. Open the directory where the BigFix server is installed and run the `BESAirgapTool.exe` file. When prompted, save the file to a new folder, for example `Airgap`. An airgap request file is created.
2. Copy all the created files to a Windows computer with the Internet access.
3. On the computer with the Internet access, run the `BESAirgapTool.exe`. The airgap request file is changed into a response file.
4. Copy the `AirgapResponse` file to the BigFix server and place it in the directory that you created in step 1.
5. Run `BESAirgapTool.exe`. The airgap response is loaded to the BigFix server.

[Cache the files \(on page cccix\)](#) and move them to the BigFix server.

Updating the content of the fixlet site on Linux

If the BigFix server is installed on a Linux computer without the Internet access, use the Airgap tool to download the content of the fixlet site to a Windows computer with the Internet access.

- You need a Windows computer with the Internet access.
- Download the Airgap tool to the Windows computer with the Internet access.

For more information, see: [Using the Airgap tool in the extraction mode](#).

1. Open the command line and enter the following commands to run the Airgap tool:
 - For BigFix versions up to 9.5.4

```
cd /opt/BESServer/bin
./Airgap.sh -run
```

The `airgap.tar` file is created in the current directory.

- For BigFix version 9.5.5 and higher

```
cd /opt/BESServer/bin
./Airgap.sh -remotedir directory
```

Where:

directory

The specified directory where the `airgap.tar` file is going to be created.

2. The `airgap.tar` file contains the air gap request file. Extract the file by using the following command:

```
tar xvf airgap.tar
```

3. Copy the extracted `AirgapRequest.xml` file to the Windows computer and place it in the folder that contains the downloaded `BESAirgapTool.exe` file.
4. On the Windows computer, run the `BESAirgapTool.exe`. The air gap request file is changed into the air gap response file.
5. Copy the `AirgapResponse` file to the BigFix server and place it in the `/opt/BESServer/bin` directory.
6. Run the Airgap tool again to upload the `AirgapResponse` file to BigFix.

```
cd /opt/BESServer/bin
./Airgap.sh -run
```

Wait a few minutes for the BigFix console to refresh.

[Cache the files \(on page cccix\)](#) and move them to the BigFix server.

Caching the files

Typically, fixlets and tasks download files that they require from the Internet. In separated networks, the required files must first be cached on a Windows computer with the Internet access and then moved to the BigFix server.

You need a Windows computer with the Internet access.

1. **Optional:** The default size of the BigFix server cache is enough if you use only the **IBM BigFix Inventory v9** fixlet site. If you plan to run fixlets from other sites, such as **BES Support**, increase the cache size so that the BigFix server does not try to delete any files. To increase the cache size, perform the following steps.
 - a. Log in to the BigFix console.
 - b. In the navigation tree, click **Computers** and right-click the computer on which the BigFix server is installed. Then, click **Edit Computer Settings**.
 - c. Increase the value of the `_BESGather_Download_CacheLimitMB` setting. If the setting is not on the list, add it and specify the value in MB.
The size depends on each fixlet site. You might need to increase it to at least a couple of gigabytes.
2. Copy the `BigFix Inventory.efxm` file from the BigFix server to the Windows computer with the Internet access.

- a. Log in to the computer where the BigFix server is installed and go to the `install_dir\BES Server\wwwrootbes\bfsites` directory.
 - b. Copy the `BigFix Inventory.efxm` file and move it to the computer with the Internet access. Place the file in the `C:\BigFix` directory.
3. Use the BES Download Cacher to download the required files.
- a. Download the BES Download Cacher to the computer with the Internet access.
 - b. On the computer with the Internet access, go to the `C:\BigFix` directory, and create a folder that is called `downloads`.



Tip: Do not clean the contents of this folder. Next time, when you run the Download Cacher only the files that were changed since the last download will be updated. The process will last shorter than if you downloaded the entire content every time.

- c. Run the BES Download Cacher.

```
BESDownloadCacher.exe -m "C:\BigFix\BigFix Inventory.efxm"
-x C:\BigFix\downloads
```

4. Copy the contents of the `downloads` folder from the computer with the Internet access to the following directory on the BigFix server.

```
install_dir\BES Server\wwwrootbes\bfmirror\downloads\shal
```

The cached files are automatically delivered to the BigFix relays and clients every time you run a fixlet that requires those files. Use both the Airgap tool and the BES Download Cacher periodically to ensure that the content of your fixlet site is always up-to-date.

When you run a fixlet, an action is created to indicate the status of its execution. If an action started by a fixlet is running and a newer version of the fixlet is delivered, stop the current action. Then, run the new version of the fixlet. To do this, perform the following steps.

1. Log in to the BigFix console and go to Actions.
2. Look for actions started by fixlets that were updated, right-click each action, and click **Stop Action**.
3. Run the updated version of the fixlets.

Backing up the server

Before you start the upgrade, back up the database, installation registry, and the server files to be able to restore the server if the upgrade fails.

1. [Stop the BigFix Inventory server](#).
2. Back up the database.

Linux Back up the DB2 database ([on page cdxxv](#)).

Windows Back up the SQL Server database ([on page cdxxx](#)).

3. Back up the installation registry.

Linux `/var/.com.zerog.registry.xml`

This file is hidden. To view it, run the `ls -la` command.

If you installed the server as a non-root user, the registry is in `$HOME/.com.zerog.registry.xml`.

Windows `C:\Program Files\Zero G Registry\.com.zerog.registry.xml`

This file is hidden. To view it, change the Folder Options in Control Panel to display hidden files, or type the complete path in the address bar of any open window.

4. Back up the installation directory. The default paths are:

Linux `/opt/ibm/BFI`

Windows `C:\Program Files\ibm\BFI`

In earlier versions, the default path is `C:\Program Files\IBM\SUA` or `/opt/ibm/SUA`.

5. If you did not start the BigFix Inventory server during the database backup, [start the server \(on page cccxcix\)](#) now.

You completed all steps that are required to restore the server after a failed upgrade. You can proceed to upgrading the server.

Upgrading the server in interactive mode

To upgrade the server in interactive mode, start the installation of a newer version of BigFix Inventory on the same computer. In interactive mode, you upgrade the server through a wizard.

To see the automated procedure for upgrading the server to the newest version, see: [Upgrading to BigFix Inventory 9.2.17 \(on page cccv\)](#)

[Create a back up \(on page cccx\)](#) of the database, installation registry, and the server files to be able to restore the server if the upgrade fails.

Ensure that no files or directories in the BigFix Inventory installation directory are open in an external program. For example, that no log files are open in a text editor. Also, ensure that no PDF reports are being generated during the server upgrade. Otherwise, the upgrade might fail with error code 41.

1. Download the latest installer and extract it.
 - a. Log in to the BigFix console and go to **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
 - b. Select **Download BigFix Inventory**, and click **Take Action**.
 - c. Select the computer on which the BigFix Inventory server is installed, and click **OK**. When the download completes, extract the installer.
2. Run one of the installation files:

```
Linux setup-server-linux-x86_64.sh
```

```
Windows setup-server-windows-x86_64.bat
```

3. Follow the steps in the wizard.
4. When the upgrade is complete, click **Done**.

 **Tip:** If the upgrade fails, check the upgrade logs. For more information, see: [Server installation and upgrade logs \(on page dccxcviii\)](#).

[Continue with the upgrade of the remaining components. \(on page ccv\)](#)

Related information

[Server installation and upgrade logs \(on page dccxcviii\)](#)

[Installation and upgrade problems \(on page dcclxxi\)](#)

[Restoring the server after a failed upgrade \(on page cccxiv\)](#)

Upgrading the server in silent mode

To upgrade the server in silent mode, start the installation of a newer version of BigFix Inventory on the same computer. Silent installation runs in the background.

To see the automated procedure for upgrading the server to the newest version, see: [Upgrading to BigFix Inventory 9.2.17 \(on page ccv\)](#)

[Create a back up \(on page ccx\)](#) of the database, installation registry, and the server files to be able to restore the server if the upgrade fails.

Ensure that no files or directories in the BigFix Inventory installation directory are open in an external program. For example, that no log files are open in a text editor. Also, ensure that no PDF reports are being generated during the server upgrade. Otherwise, the upgrade might fail with error code 41.

1. Download the latest installer and extract it.
 - a. Log in to the BigFix console and go to **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
 - b. Select **Download BigFix Inventory**, and click **Take Action**.
 - c. Select the computer on which the BigFix Inventory server is installed, and click **OK**. When the download completes, extract the installer.
2. Read the license agreement in the `license.txt` file. The file is in the `license/LA_language` directory.
3. Edit the `upgrade_response.txt` file.
4. In the response file, set the **RSP_LICENSE_ACCEPTED** parameter to `true`.
5. To start the upgrade, run one of the following commands:

```
Linux ./setup-server-linux-x86_64.sh -i silent -f /
response_file_path/upgrade_response.txt
```

```
Windows setup-server-windows-x86_64.bat -i silent -f
response_file_path\upgrade_response.txt
```

For `response_file_path`, enter the absolute path to the response file. For example:

```
setup-server-windows-x86_64.bat -i silent -f C:\images\SUA\upgrade_response.txt
```

 **Tip:** If the upgrade fails, check the upgrade logs. For more information, see: [Server installation and upgrade logs \(on page dccxcviii\)](#).

[Continue with the upgrade of the remaining components. \(on page cccv\)](#)

Related information

[Server installation and upgrade logs \(on page dccxcviii\)](#)

[Installation and upgrade problems \(on page dcclxxi\)](#)

[Restoring the server after a failed upgrade \(on page cccxiv\)](#)

Configuring the server after upgrading from version earlier than 9.0.1.2

If you upgraded from versions earlier than application update 9.0.1.2, you might be required to perform additional configuration. The configuration panel opens automatically. You can also complete the configuration later at

`https://hostname:9081`.

Optimizing the volume of scanned file data

You can optimize the volume of data that you monitor and use for creating custom signatures by narrowing down the list of collected file extensions to those that are typically used. This reduces the workload on the application server and improves the performance.

1. On the configuration panel, select the **Optimize the volume of scanned file data** check box.
2. Check whether the software catalog that is uploaded to BigFix Inventory is in the canonical format. Click **Management > Catalog Upload** and check the **Catalog Format** column.

Upload and Import History

Uploaded By	Uploaded At	Content	Provider	Version	Catalog Format	Date
Initial upload	02/26/2015 03...	Catalog XML file	IBM	1222224	Canonical	02/06/2015 12...

If the catalog is in the native format, upload a new catalog. If the catalog is in the canonical format, but a new version is available, upload the new catalog. Otherwise, proceed to the next step.

3. Run an import.

During this import, performance might be lower because the software catalog is imported.



Important: After the import, some software items might not be visible in the reports. This is an expected behavior. Complete the remaining steps for the software inventory to be properly reported.

4. Wait for the scheduled software scan. Alternatively, if you have infrequent software scans, stop the current scan and start a new one. This allows you to use the optimized list of file extensions in a shorter time.
 - a. Log in to the BigFix console and, in the left navigation tree, click **Actions**.
 - b. In the upper-right pane, click **Initiate Software Scan** and then click **Stop**.
 - c. [Initiate a new software scan. \(on page cciii\)](#)
5. Wait for the scheduled import or run it manually. From now on, the optimized list of file extensions is used.
6. If you upgraded from Software Use Analysis 9.0, the default Java heap size is not sufficient for the current version. Increase the Java heap size to 1536 megabytes (`Xmx:1536m`).

Configuring without optimizing the volume of the scanned file data

If you do not want to optimize the volume of the scanned file data, clear this option. You can narrow down the list of collected file extensions manually at any later time.

1. Clear the **Optimize the volume of scanned file data** check box.
2. Run the import of data to be able to see complete infrastructure information.
3. If you upgraded from Software Use Analysis 9.0, the default Java heap size is not sufficient for the current version. Increase the Java heap size to 1536 megabytes (`Xmx:1536m`).

Restoring the server after a failed upgrade

After a failed upgrade, you can restore the server to its original state by restoring the database, installation registry, and the server files from backups.

If you did not back up the installation directory and the registry, you can install the new server and restore only the database. You can keep the data that is already in your database and avoid running the initial import.

1. [Stop the server \(on page cd\)](#).
2. Restore the database from a backup.

Linux Restore the DB2 database (on page [cdxxviii](#)).

Windows Restore the SQL Server database (on page [cdxxxii](#)).

3. Replace the installation registry with the backup registry. Copy the `.com.zerog.registry.xml` backup file to one of the following directories:

Linux `/var/`

Windows `C:\Program Files\Zero G Registry\`

If you installed the server as a non-root user, copy the backup file to `$HOME/`.

4. Replace the installation directory with the backup directory.
5. [Start the server](#).

Migrating from License Metric Tool 9.2.17

It is possible to migrate from IBM License Metric Tool to HCL BigFix Inventory. Please contact [support](#) for additional instructions providing your current IBM License Metric Tool version and database details. Migration process is manual and it is recommended to get assistance from HCL services.

Managing the infrastructure

After you complete the initial configuration of BigFix Inventory, learn how to manage components of its infrastructure: VM managers, server, database, and data sources.

Computer statuses

Computer statuses provide information about the condition of the computers that are monitored by BigFix Inventory. To properly discover the installed software and measure its license metric utilization, ensure that computers in your environment work correctly.

OK

There are no problems with the computer. No actions are required.

No Scan Data

Results of the capacity scan are missing. To solve the problem, schedule the capacity scan, and upload its results. For more information, see: [Initiating the capacity scan on all computers \(on page ccx\)](#).

No Host Scan Data

The BigFix client is not installed on the host operating system. Depending on the virtualization technology, perform the following actions:

- For HP Integrity VM, install the BigFix client on the host computer and wait until it uploads the capacity scan results or upload them manually. For more information, see: [Initiating the capacity scan on all computers \(on page ccx\)](#).
- For Solaris Container/Zones or Logical Domains (LDOM), install the BigFix client in the following zones:
 - Global zone on the control domain
 - Global zones of all local zones in which software is installed

Wait until the client uploads the capacity scan results or upload them manually. For more information, see: [Client Installation on Oracle Solaris \(on page ccxii\)](#).

- For other virtualization technologies, refer to [Virtualization Capacity License Counting Rules](#) for more information.

Outdated VM Manager Data

Connection to the VM manager is configured, but hardware inventory data that was retrieved from the computer is older than 3 days. License metric utilization is not affected yet. If you do not resolve the issue within 30 days, the computer status changes to *No VM Manager Data* and default PVU counting is applied to the computer. As a result, license metric utilization that is reported can be higher. For more information, see: [Default PVU counting on x86 processors \(on page cccix\)](#).



Note: The period after which hardware inventory data is considered outdated is configured by the **vmManagerDetachmentNotificationPeriod** parameter. The period after which the computer status changes to *No VM Manager Data* is configured by the **vmManagerDetachmentPeriod**. For more information, see: [Advanced server settings \(on page cdi\)](#).

To solve the problem, perform the following actions:

- Ensure that the VM manager is connected.
- Ensure that all parameters of the VM manager connection are correctly configured.



Important: When you troubleshoot the VM manager connection, ensure that the following conditions are fulfilled:

- You have sufficient rights to collect information about the complete virtualization topology including virtual machines, hosts, clusters, and relations among them.
 - The UUIDs of virtual machines and hosts are valid and in scope of the particular virtualization technology.
 - Information about the number of processors, total number of cores, and processor description is available.
- If the problem persists on KVM although the VM manager connection is correctly configured, go to **Reports > Hardware Inventory**, and check the value in the **Server ID** column. If the value has the `TLM_VM_UUID_of_the_VM` format, check whether the UUID is correctly set on the virtual machine. If two or more virtual machines have the same UUID, manually set unique UUIDs for these machines. Note that virtual machines can operate on different hosts.
 - Ensure that each cluster in your infrastructure has a unique name. In the VMware environment, it is not possible to create two clusters with the same name in the same data center. However, vCenter can manage several data centers at the same time. As a result, one vCenter can control clusters that have the same name but are in different data centers. If multiple clusters share a name, rename them to keep each name unique.
 - Ensure that the value of the **vmManagerDetachmentPeriod** parameter is higher than the frequency with which data is retrieved from VM managers. The interval between consecutive retrievals of data is set by the **vmm_polling_time_interval** parameter and is 30 minutes by default. If the idle time after which a computer that is managed by a VM manager is considered detached is lower, change the value of the **vmManagerDetachmentPeriod** parameter.



No VM Manager Data

Hardware inventory data that was retrieved from the computer is older than a specific period or has never been retrieved. Default PVU counting is applied to the computer. As a result, reported license metric utilization can be higher than actual. For more information, see: [Default PVU counting on x86 processors \(on page cccix\)](#).

Impact of the peak of metric utilization on the reports depends on why the status appeared. The status can appear in the following situations:

- The VM manager is not specified for the computer and hardware inventory data has never been retrieved. If you configure the VM manager within a specific period, the peak of metric utilization is ignored and does not influence the reports. The period is:
 - 1 day for version up to 9.2.9.
 - 30 days for version 9.2.10 and higher. The period can be configured by the **vmManagerDetachmentPeriod** parameter.
- The VM manager is specified for the computer but the retrieved hardware inventory data is older than a specific period. The period can be configured by the **vmManagerDetachmentPeriod** parameter and is by default:
 - 7 days for version up to 9.2.9
 - 30 days for version 9.2.10 and higher

The peak of metric utilization is not removed from the reports that are generated for the period when the computer had the *No VM Manager Data* status even after you resolve the issue.

To solve the problem, perform the same actions as in case of the [Outdated VM Manager Data \(on page cccxvi\)](#) status.

Managing VM managers

VM managers are pieces of software that create, manage, and monitor virtual machines. You configure connections to VM managers so that BigFix Inventory can gather data that is required to calculate utilization of license metrics in virtual environments.

Collecting information from VM managers

Resources that are available to virtual machines are dynamically allocated based on current needs. When a virtual machine is not performing processor-intensive tasks, processor cores that are allocated to that machine are assigned to other virtual machines that need them to handle their workload. Because of these dynamics, BigFix Inventory cannot scan virtual machines to collect information about processor type and the number of available cores.

To collect this information, BigFix Inventory needs access to VM managers, which control resources that are available to virtual machines. With such access, BigFix Inventory can check the number of cores on the physical computer that hosts the virtual machines and see the allocation and usage of these resources.

You must define connections to VM managers for the following x86 virtualization types:

- [VMware vSphere \(on page cccxxii\)](#)
- [Microsoft Hyper-V \(on page cccxxiv\)](#)
- [Kernel-based Virtual Machine \(on page cccxli\)](#) except for the hosts that are not controlled by RHEV-M and the PowerKVM hosts that deliver capacity data with the Run Capacity Scan on Virtualization Hosts fixlet.
- [Oracle VM Server for x86 \(on page cccxliv\)](#)

Lack of information from VM managers

You should configure connections to VM managers within 24 hours, after you run the first successful hardware scan on new virtual machines. Otherwise BigFix Inventory is not able to retrieve information about the virtualization

hierarchy and physical processors. Thus, it is not able to properly calculate subcapacity values. In such a case, BigFix Inventory applies default PVU counting on virtual machines which causes that PVU utilization might be over-counted. For more information, see: [Default PVU counting on x86 processors \(on page cccix\)](#).

IBM accepts audit reports that contain PVU values calculated based on the default PVU counting instead of considering the client to be ineligible for subcapacity or liable for full capacity.

When you configure connections to VM managers later on, the reported PVU values are lower than or equal to the values that were reported when the connections were missing. However, peaks that resulted from missing connections remain on the reports that are generated for the period when the peaks occurred.

VM Manager Tool

To collect data from VM managers, BigFix Inventory uses the VM Manager Tool. The tool connects to VM managers and gathers data about their capacity, focusing on processors, their type, and usage. The data is then imported to BigFix Inventory and utilization of license metrics is calculated.

Approaches to managing VM managers in BigFix Inventory

You can choose between two approaches to managing VM managers: basic (central) or advanced (distributed). You can also combine these two approaches. The choice depends on the specifics of your environment.

Default PVU counting on x86 processors

To properly calculate PVU subcapacity, BigFix Inventory requires information about the number of processor cores that are available to virtual machines and on the physical computers that host these machines. If the information is not available, license metric utilization cannot be properly calculated and reported results might be over-counted for x86 processors.

When hypervisor data is not available, reported results might be over-counted due to the following reasons.

- **The applied PVU rate might be higher than the actual rate**

PVU rate that is applied to a processor depends on the model and type of the processor as well as the number of processors that can be installed on a physical host (the number of sockets). When hypervisor data is not available, BigFix Inventory cannot properly identify the PVU rate for a processor. It applies the PVU rate for the configuration with the highest number of sockets that is possible for the particular processor.



Important: **9.2.11** Starting from application update 9.2.11, BigFix Inventory always applies 120 PVUs per core when the PVU rate for a processor cannot be properly identified. In earlier versions of BigFix Inventory, the number of applied PVUs was specified in the PVU table.

- **The number of processor cores that are taken into account might be higher than the actual number of available cores**

If CPU is over-committed to virtual machines and the total virtual capacity exceeds the physical capacity, virtual capacity must be capped to physical capacity according to the pricing rules. When hypervisor data is not available, BigFix Inventory cannot limit virtualization capacity to physical capacity.

IBM accepts audit reports that contain PVU values calculated based on the default PVU counting instead of considering the client to be ineligible for subcapacity or liable for full capacity.

For more information about the results of applying default PVU counting, see: [Scenario: Capacity in virtual environments with no VM managers defined \(on page dlxxviii\)](#).

i **Tip:** If the hypervisor data is missing from a virtual machine and default PVU counting is applied, ID of the server on which the machine runs begins with `TLM_VM`.

Supported virtualization types

BigFix Inventory requires that connections to VM managers are defined for the following virtualization types: Kernel-based Virtual Machine, Microsoft Hyper-V, VMware vSphere, and Xen.

Table 37. Supported virtualization types and versions

The table consists of three columns and three body rows.

Virtualization Type		Version	Comments
Kernel-based	Managed by	3.0, 3.1, 3.5 or 3.6	
Virtual Machine	RHEV-M		
	Managed by	4.0	<ul style="list-style-type: none"> • BigFix Inventory supports instances that use HTTP Basic Authentication and version 3 REST API. • You must use VM Manager Tool in version 1.6.0.1 or higher. • Starting from VM Manager Tool update 9.2.17, BigFix Inventory supports instances that use version 4 API and OAuth2 Authentication.
	Not managed by		
	RHEV-M		<ul style="list-style-type: none"> • Capacity data from x86 KVM hosts that are not controlled by RHEV-M is collected directly from the host by using the Run Capacity

Table 37. Supported virtualization types and versions

The table consists of three columns and three body rows.

(continued)

Virtualization Type	Version	Comments
9.2.17 Nutanix AVH Managed by Nutanix Prism	AVH 5.11 +/-PRISM	Scan on Virtualization Hosts fixlet. For more information, see: Collecting capacity data directly from KVM hosts (on page cccxvi) .
Microsoft Hyper-V	<ul style="list-style-type: none"> • Windows Server 2008, 2008 R2 (stand-alone or direct mode) • Windows Server 2012, 2012 R2 (stand-alone or direct mode) • Windows Server 2016, 2016 R2 	<ul style="list-style-type: none"> • To use WinRM communication interface, configure WinRM on all hosts. For more information, see: Configuring WinRM (on page cccxxviii).
VMware vSphere	ESXi	<ul style="list-style-type: none"> • 5.0, 5.1, 5.5 • 6.0, 6.5, 6.7 <ul style="list-style-type: none"> • You need at least read-only rights to all VMs. For more information, see: Verifying permissions (on page cccxxiv).
	vCenter	<ul style="list-style-type: none"> • 5.0, 5.1, 5.5 • 6.0, 6.5, 6.7 <ul style="list-style-type: none"> • You need at least read-only rights to all VMs. For more information, see: Verifying permissions (on page cccxxiv). • BigFix Inventory supports configuration with one single sign-on domain, one single sign-on site, and vCenter Server with Platform Services Controller on the same machine (Embedded Deployment Model).

Table 37. Supported virtualization types and versions

The table consists of three columns and three body rows.

(continued)

Virtualization Type	Version	Comments
<p>9.2.12 Xen Server or Citrix</p> <p>XenServer</p>	<ul style="list-style-type: none"> • 6.5 • 7.0 	<ul style="list-style-type: none"> • BigFix Inventory supports vCenter Server installation on a Windows computer and as a vCenter Server Appliance. • 9.2.15 BigFix Inventory supports the configuration with external Platform Services Controller (PSC) with one PSC on a different machine. In case of a complex environment with multiple PSCs, contact IBM Support. • XenServer Tools must be installed on all VMs.
<p>9.2.14 Managed by the Oracle VM Manager</p> <p>Oracle VM Server for x86</p>	<ul style="list-style-type: none"> • 3.4.x 	<ul style="list-style-type: none"> • Oracle VM Server for SPARC is supported by BigFix Inventory on different terms. For more information, see: Client Installation on Oracle Solaris (on page ccxii).

VMware vSphere

VMware vSphere is one of the virtualization technologies supported by BigFix Inventory.

Purpose

This solution consists of two products: a VMware ESX (or ESXi) hypervisor and a VMware vCenter server. Both of these components provide the API that can be used to extract information concerning virtual machines in your infrastructure. The section below shows the differences between these two varieties.

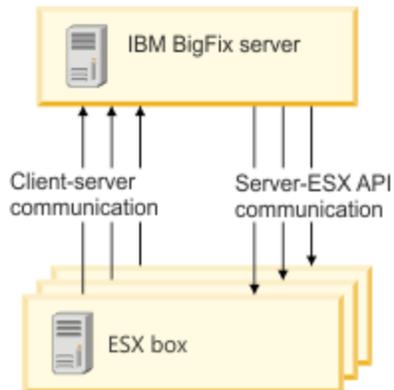
ESX or ESXi

It is an operating system that hosts virtual machines. It is enough to connect to a specific ESX if you do not use any software deployed in a clustered environment. However, if an ESX is controlled by a vCenter, it is recommended to connect a given VM manager to it via the vCenter. This is because BigFix Inventory needs data from every ESX box (every box that hosts a virtual machine where the CIT scanner is installed). If you choose a connection via ESX, you have to define all the ESX boxes separately in the **VM Managers** panel. If you have a vCenter deployed, you can manage all the ESX boxes via this server (it saves time and decreases the network load).

 **Note:** Changing the Universally Unique Identifier (UUID) on ESX virtual machines may lead to overcharging because when the identifier is changed, BigFix Inventory recognizes it as a brand new virtual machine.

Supported versions:

For a full list of supported versions, see: [Supported virtualization types \(on page cccxx\)](#).

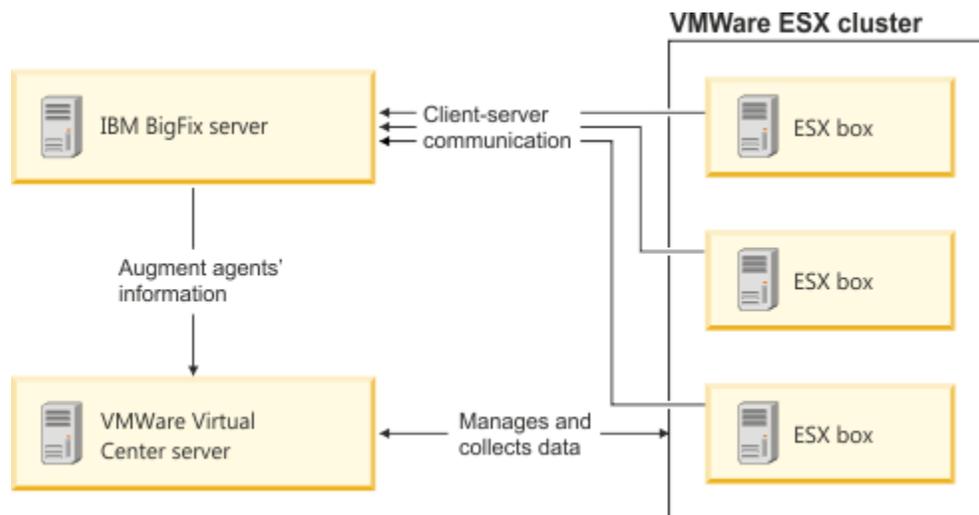


vCenter

It is used for managing computer systems on which virtual machines are installed. You have to use it if your software is installed in a clustered environment. Instead of connecting to each ESX box separately, you connect to all ESX boxes via a vCenter, whose main role is to retrieve data from them. Therefore, it is beneficial to you to connect to ESX boxes through a vCenter even if you do not have a clustered environment because in this way you can reduce the network traffic.

Supported versions:

For a full list of supported versions, see: [Supported virtualization types \(on page cccxx\)](#).



The default URL for ESX or ESXi, and vCenter

The default URL that is to be used (it may vary in your environment if the administrator of the ESX box, or vCenter has changed the configuration of the API, for example, if the HTTP protocol is used instead of the HTTPS one):

- For ESX or ESXi: `https://ESX_IP_address/sdk/vimService.wsdl`
- For vCenter: `https://vCenter_IP_address/sdk/vimService.wsdl`

Make sure that the vSphere Web Services SDK is set up and verify if the given URL is valid. Paste the URL into the Web browser and make sure that it can be accessed. If not, check the setup of the URL in vCenter settings. To do this, go to **vCenter Server Settings > Advanced Settings** and see the setting of **VirtualCenter.VimApiUrl**. If the URL cannot be reached, consult your VMware administrator.

Supported topologies

BigFix Inventory supports configuration with one single sign-on domain, one single sign-on site, and vCenter Server with Platform Services Controller on the same machine (Embedded Deployment Model). To set up other topologies, including external Platform Services or multiple single sign-on sites, you must connect to each ESX box separately.

Verifying permissions for VMware communication

Users must have sufficient privileges to collect all the data from the VM managers for VMware. The user must have at least read-only rights to all VMs on which the agents are running.

You verify whether users have sufficient privileges in the VMware Infrastructure Client. The user must have the correct access privileges for VMs on which the endpoints are running, and for the hosts of the VMs. If a user has insufficient privileges, agents return a No VM Manager Data status.

This procedure describes how to enable read-only rights for all elements in a virtual topology. Read-access is required only to the VMs on which the endpoints are running and to hosts of these VMs. However, the easiest way to set access permissions is to configure read-only access for all elements in a virtual topology.

1. To extend the privileges for the user, log in to VMware Infrastructure Client with a user that has administrator rights.
2. Right-click on the left bar and choose **Hosts and Clusters**.
3. Go to **Permissions** tab, right-click anywhere in the section, and then click **Add Permission**.
4. In the **Assign Permissions** panel, click **Add**, choose the user, and then click **OK**.
5. Choose Read-Only as an Assigned Role.
6. Select the **Propagate to Child Objects** check box, and click **OK**.

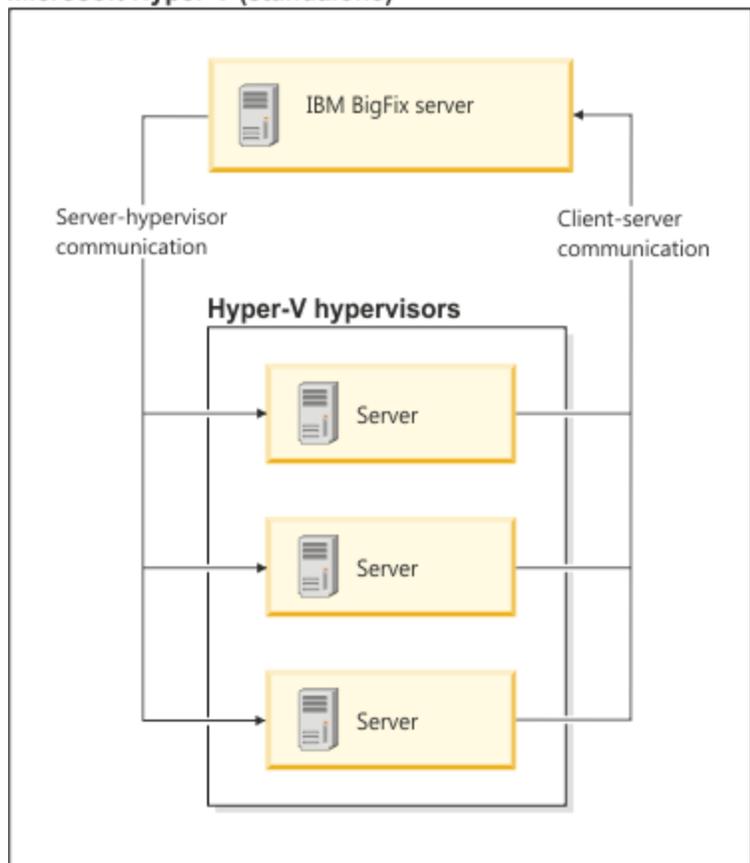
Microsoft™ Hyper-V

Microsoft™ Hyper-V is one of the virtualization technologies that are supported by BigFix Inventory. It is the successor of Microsoft™ Virtual Server.

Purpose

To retrieve information about the measures and virtualization structure of virtual machines, define the Hyper-V hypervisors as VM managers on the BigFix Inventory server. They are able to expose required data via the WS-MAN protocol.

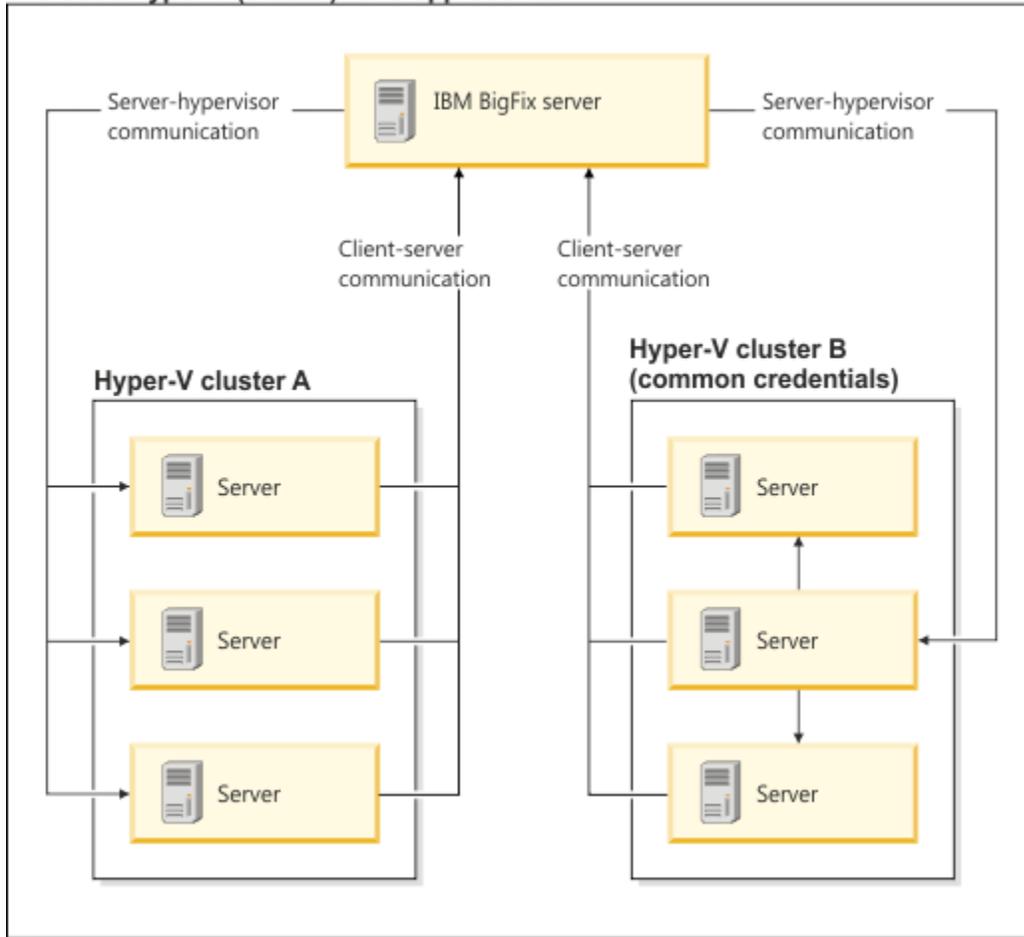
Microsoft Hyper-V (standalone)



In the case of clusters, two approaches exist:

1. Define all Hyper-V servers that are a part of the cluster in the **VM Managers** panel. If any servers are missing, they are in the *Incomplete definition* status.
2. If all members of the cluster have at least one common set of credentials that has privileges to access the WS-MAN interface (for example, the domain user), define one computer system from the cluster as a VM manager. Then, select the option to share credentials. The BigFix Inventory server uses the credentials to connect to the Hyper-V that you defined in the user interface. Then, it extracts addresses of the remaining cluster members and connects to them by using the same credentials.

Microsoft Hyper-V (cluster) - two approaches



Communication interface

You can choose to use PowerShell or WinRM as the communication interface for Hyper-V. For both communication interfaces, make sure that you provide the user name in one of the following formats:

- *user_name@domain*, for example: `test@cluster.com`
- *user_name\domain*, for example: `test\cluster.com`

9.2.16 Starting from application update 9.2.16, you can also provide the administrator user name in one of the following formats:

- *domain@user_name*, for example: `cluster.com@test`
- *domain\user_name*, for example: `cluster.com\test`



Note: Make sure that you upgrade the VM Manager Tool to version 9.2.16, not only the BigFix Inventory server.

If the target server is not in a domain, specify its host name as the domain. For example:

- `test\hostname` OR `test@hostname`
- **9.2.16** `hostname\test` OR `hostname@test`

PowerShell

PowerShell is a framework for managing the automation and configuration of tasks and can use the RPC protocol. In order for BigFix Inventory to use PowerShell for retrieving the information from VM managers, the VM Manager Tool must be installed on a Windows™ computer. The computer must also meet the following requirements:

- Microsoft™ RPC communication with Microsoft™ Hyper-V must be allowed by Windows firewall
- ICMP Echo Reply must be allowed by Windows firewall
- Microsoft™ PowerShell 2.0 or higher must be installed
- Microsoft™ .NET Framework 3.5 or higher must be installed

! **Important:** If you are using .NET Framework 4.0 or higher, use PowerShell version 3.0 or higher.

The use of PowerShell depends on the operating system of your BigFix server, because the main VM Manager Tool is installed along with the server:

Windows You configure a Hyper-V VM manager in the BigFix Inventory UI and choose PowerShell as the communication interface.

VM Manager Type*

Communication Interface*

Linux You use advanced VM management to install additional VM Manager Tool on any Windows computer, and then specify a Hyper-V VM manager in a configuration file. Such a VM manager is not visible in the UI, but is managed through the VM Manager Tool command line. For more information, see [Advanced VM management \(on page ccclxvii\)](#).

```
vmm_url=https://hyper-v/wsman
vmm_type=MICROSOFT_HYPER_V
vmm_communication_interface=POWERSHELL
```

WinRM

Windows Remote Management (WinRM) is the Microsoft implementation of the WS-Management protocol. It uses Simple Object Access Protocol (SOAP) over HTTP and HTTPS. In order for BigFix Inventory to use WinRM for retrieving the information from VM managers:

- You must configure Windows™ Remote Management Service (WinRM) for all members of the Microsoft™ Hyper-V cluster. For more information, see: [Configuring WinRM on Hyper-V hosts \(on page cccxxviii\)](#).
- The default URL that is to be used is `https://<HYPER-V_IP_address>:<port>/wsman`, where `<port>` is the port of the listener that was created for the WinRM service. The default value for the HTTP listeners is 5985 and for HTTPS listeners is 5986 (WinRM 2.0). To verify all the defined WinRM listeners, use the `winrm enumerate winrm/config/listener` command.

Supported versions

For a full list of supported versions, see: [Supported virtualization types \(on page cccxx\)](#).

Configuring WinRM on Hyper-V hosts

Configure Windows™ Remote Management to allow the BigFix Inventory server to gather data about virtualization topology of virtual machines installed in your infrastructure.

- To retrieve the data that is required to properly calculate PVU, you must be logged in as a local administrator on the Hyper-V host. It is necessary because the Windows Management Instrumentation call that accesses MsCluster namespace requires an administrative account.
- Hardcoded and select-only statements are run over WinRM. The obtained data is stored in a database schema. BigFix Inventory does not modify the Hyper-V settings and does not affect it any other way.

The WinRM service is an implementation of WS-Management specification that enables cooperation between hardware and operating systems that come from different vendors. The BigFix Inventory server connects to this service defined as a VM manager and collects data regarding virtualization hierarchy. Therefore, you must perform the following procedure on each Hyper-V host in your infrastructure, including those that are part of a cluster, to ensure the WinRM service is running and configured to enable communication with the server.

1. **Defining HTTP and HTTPS listeners.** By default, communication with the WinRM service is disabled because there are no listeners defined. To check whether there are any listeners that are currently defined, type the following command: `winrm enumerate winrm/config/listener`. If there is no output returned, there are no listeners defined.

- a. To define a default HTTP listener, type:

```
winrm quickconfig
```

The command starts the WinRM service and sets it to start automatically with the system start. It also creates an HTTP listener on the default port (accepting requests from any IP), it defines Internet Connection Firewall exceptions for the service, and it opens the HTTP port. Depending on the version of the WinRM service, the default HTTP port might be 80 or 5985. For more information, see [Installation and Configuration for Windows™ Remote Management](#).

- b. To define a listener for secure connection (HTTPS), you must have a valid certificate on the Hyper-V host with a CN that matches the host name that you are using to connect to Hyper-V. You must

also create a listener with the CertificateThumbprint of that certificate. For more information, see the Microsoft documentation: <http://support.microsoft.com/kb/2019527>.. You might be able to create a self-signed certificate for testing purposes, however, you should consult your certificate administrator.



Note: If an appropriate certificate was not found on the machine, the above command will not work and the following output will be returned *"The certificate must have a CN matching the host name, be appropriate for Server Authentication, and not be expired, revoked, or self-signed."* If there is a need to use a self-signed certificate, you can manually generate it and create the listener by starting the following command:

```
winrm create winrm/config/listener?Address=*&Transport=HTTPS @{Hostname="<the name of your server>";CertificateThumbprint="<certificate thumbprint>"}
```

In this case you have to configure the firewall settings manually.

2. **Enabling WinRM *Negotiate* authentication scheme.** The WinRM service offers several authentication schemes to be used to authenticate the client side. The BigFix Inventory server uses *Negotiate* authentication scheme, which is enabled by default.

a. To check the current setting of this property, type:

```
winrm get winrm/config/service/auth
```

b. To set the required value of this property, enter:

```
winrm set winrm/config/service/auth @{Negotiate="true"}
```

3. **Setting WinRM *AllowUnencrypted* property.** The server requires the property to be set to "true".

a. To check the current settings, type:

```
winrm get winrm/config/service
```

b. To set the required value of this property, type:

```
winrm set winrm/config/service @{AllowUnencrypted="true"}
```



Note: Setting this value to "true" does not mean that the sensitive data, such as user names or passwords, will be passed in an unencrypted form over the network. Only the content of the SOAP messages will be sent as a plain text. If this cannot be accepted because of security reasons, define the HTTPS listener and use the secured transport (HTTPS) while defining a VM manager in the BigFix Inventory server so that the TLS protocol will be used to encrypt all the network traffic.

4. **Verifying the listener.** After you define the HTTP or HTTPS listener, verify that you can remotely connect to the Hyper-V server.

- a. On the Hyper-V server, determine the port on which the Windows Remote Management client for the HTTP or HTTPS transport listens. Type the following command in the Windows command line:

```
winrm enumerate winrm/config/listener
```

- If the port number is listed in the `Port` line, the listener was properly created.
- If you receive an error or there is no information for the transport, the listener was not created properly. Go back to step one, and define the listener again.

- b. To verify the listener, type:

```
winrm enumerate winrm/config/listener /r:<transport>://
<Hyper-V_server_name>:<port>/wsman /u:<user_id> /p:<password> /a:Negotiate
```

Where

<transport>

Is either HTTP or HTTPS.

<Hyper-V_server_name>

Is the host name of the Hyper-V server. If you are using HTTPS, the host name must match the CN in the certificate.

<port>

Is the port number that you obtained in the previous step.

<user_id>

Is the user ID that is used to connect to the Hyper-V server.

<password>

Is the password that is used to connect to the Hyper-V server.

For example:

```
winrm enumerate winrm/config/listener /r:https://
myhyperv.ibm.com:5986/wsman /u:administrator /p:abc /a:Negotiate
```

5. **Verifying whether the Virtual System Management service is running.** To verify that the service that provides Hyper-V management is running, go to **Administrator Tools > Services** on the Hyper-V server. Look for the service called Hyper-V Virtual Machine Management
- If the service exists, but is not running, start the service.
 - If the service does not exist, the Hyper-V host was not configured properly.
6. **Verifying the MsCluster resource.** If the server is clustered, verify that you can access the MsCluster namespace. On the Hyper-V server, type the following command into the Windows command line:

```
winrm enumerate wmi/root/MsCluster/*

-dialect:"http://schemas.microsoft.com/wbem/wsman/1/WQL"

-filter:"SELECT PrivateProperties, Type FROM MsCluster_Resource WHERE Type='Network Name' AND Flags='1'"
```

If this command fails, refer to Microsoft documentation about WMI for MsCluster.

7. **Verifying remote connectivity and the server certificate.** To verify remote connectivity and the server certificate, type the following command into the Windows command line:

 **Restriction:** Enter the following command on the Windows command line of the BigFix Inventory server. If the server is not installed on a computer that runs on a Windows operating system, use a computer that is not the Hyper-V host and runs on Windows 2008 or higher.

```
winrm set winrm/config/client @TrustedHosts={"<Hyper-V_server_name>"}
winrm get winrm/config/client /r:<transport>://
<Hyper-V_server_name>:<port>/wsman /u:<user_id> /p:<password> /a:Negotiate
```

Where

<transport>

Is either HTTP or HTTPS.

<Hyper-V_server_name>

Is the host name of the Hyper-V server. If you are using HTTPS, the host name must match the CN in the certificate.

<port>

Is the port number on which the Windows Remote Management client for the HTTP or HTTPS transport listens.

<user_id>

Is the user ID that is used to connect to the Hyper-V server.

<password>

Is the password that is used to connect to the Hyper-V server.

For example:

```
winrm set winrm/config/client @TrustedHosts={"myhyperv.ibm.com"}
winrm get winrm/config/client /r:https://
myhyperv.ibm.com:5986/wsman /u:administrator /p:abc /a:Negotiate
```

The following error is often returned when a self-signed certificate is used is:

```
WSManFault
Message = The server certificate on the destination computer (myhyperv.ibm.com:5986)
has the following errors: The SSL certificate is signed by an unknown certificate authority.
```

If you receive this error, export the self-signed certificate from the Hyper-V host, and import it on the BigFix Inventory host. For other errors, refer to Microsoft documentation for the returned error code.

 **Tip:** For more information about Hyper-V configuration, see the following document: [Hyper-V connection fails CODVM0005E \(on page dcclxxv\)](#).

Troubleshooting Hyper-V connection issues

Troubleshoot, and understand the connection issues between Microsoft Hyper-V and BigFix Inventory. Analyze the listed solutions to resolve connection problems.

Glossary

The connection between BigFix Inventory and Hyper-V is established through the Windows Management Instrumentation (WMI) queries. These queries use either PowerShell over the DCE-RPC mechanism, or WinRM that uses SOAP, or XML over HTTP protocol. The BigFix Inventory client communicates with the Hyper-V host through the VM Manager interface, VM Manager Tool. The following list describes the basic components of the virtual environment:

Server

The Hyper-V host where the WinRM service is located.

Client

BigFix Inventory host that can be either a Windows or Unix computer.

The client authentication on both Windows and Unix systems uses the same authentication sequence: NTLM, NTLMV2, or HTTP Basic. The Windows domain members undergo Kerberos network authentication when you use PowerShell as the communication interface.

Troubleshooting multiple Hyper-V VM Manager Connections

When there are many Hyper-V VM Manager connections that are configured with WinRM interface, VM Manager Tool might not be able to correctly establish the connections simultaneously in multiple threads. In this case, you must set the **vmm_thread_pool_size** parameter to 1 to lower the number of connection threads. Additionally consider adding more VM Manager Tools to distribute the load between them.

Troubleshooting Hyper-V with the server diagnostic script

1. To detect potential problems with adding, or using the Hyper-V VM Manager, use the diagnostic script that checks the Hyper-V and BigFix Inventory servers.

- **UNIX** `NTLM_debug.jar` is a script that is dedicated for non-Windows systems. To download the script, see: [Troubleshooting of No VM Manager Data status](#). This script only checks the authentication schema, and tries to send out a SOAP, or XML queries to Hyper-V host. Therefore, it is recommended to additionally use a Windows script to troubleshoot the server from a Windows computer.
- **Windows** `HyperV_precheck.wsf` is a Windows script that uses WinRM commands. To download the script, see: [Troubleshooting of No VM Manager Data status](#). You need to set up WinRM on Hyper-V hosts to use the script. For more information, see: [Configuring WinRM on Hyper-V hosts \(on page cccxxviii\)](#).

The script checks:

- Whether the Virtual Machine Management Service (VMMS) is running.
- Whether the WinRM service is running.
- Whether the listeners for WinRM interface are defined.
- Whether the Certificate Thumbprint for HTTPS communication is in place.
- Whether the **Negotiate** authentication scheme is set as a valid protocol.
- Whether **AllowUnencrypted** is set for HTTP protocol. However, it does not check whether the HTTP Basic authentication is enabled in WinRM communication.
- The remote connectivity and server certificates by checking the trusted hosts in the WinRM configuration.
- The Local Security Policy settings that are used for NTLM, or NTLMV2 authentication.

Ensure that you are an administrator with sufficient permissions to run WinRM commands and run the script.

```
cscript hyperv_precheck.wsf /r:<transport>://<Hyper-V_server_name>:<port>/wsman
/u:<user_id> /p:<password> /s:<server> [/v:true] [/d:true]
```

Where:

<transport>

Is either HTTP or HTTPS.

<Hyper-V_server_name>

Is the host name of the Hyper-V server. If you are using HTTPS, the host name must match the CN in the certificate.

<port>

Is the port number on which the Windows Remote Management client for the HTTP or HTTPS transport listens.

<user_id>

Is the user ID that is used to connect to the Hyper-V server.

<password>

Is the password that is used to connect to the Hyper-V server.

<server>

Is either the Hyper-V or BigFix Inventory server where the script is run.

/v

Verbose mode.

/d

Debug mode.

- a. Run the script on the Hyper-V server. Use the `/s:hyperV` command.
- b. Run the script on the BigFix Inventory server. Use the `/s:LMT` command.
- c. If any of the checks fail, rerun the commands in the verbose mode to obtain detailed information about errors. Use the `/v:true /d:true` command.

Windows Manual Hyper-V troubleshooting

If the `hyperv_precheck.wsf` script was not able to diagnose the Hyper-V issues with satisfactory results, proceed with the manual troubleshooting.

1. Verify whether the WinRM service is running. Run the following command in command prompt.

```
sc query WinRM
```

The query results should include the following information.

```
SERVICE_NAME: WinRM
STATE        : 4  RUNNING
```



Note: The Windows User Account Control (UAC) can affect user access to the WinRM service. If you use the **Negotiate** authentication scheme, only the Administrator account can access the service. To allow all accounts to access the service, set the following registry value.

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\
LocalAccountTokenFilterPolicy to '1'
```

2. Verify the WinRM configuration properties by using command prompt.

- a. Run the following commands on the Hyper-V server:

- `winrm get winrm/config/service`
- `winrm enumerate winrm/config/listener`

- b. Run the following commands on the client:

```
winrm get winrm/config/client
```



Note: If Hyper-V and BigFix Inventory are installed on the same host, you can use a single command to retrieve the required information: `winrm get winrm/config`.

3. Set up a simple test case by changing some of the WinRM settings.

- a. On the BigFix Inventory server, add an asterisk to the **Trustedhosts** list.

```
winrm set winrm/config/client @{TrustedHosts="*"}
```

`TrustedHosts="*"` forces the client to abandon authentication of the remote end. However, the remote end still requires the client authentication. Typically, `TrustedHost` is set on the client to the Hyper-V server name.

- b. Allow Basic authentication and unencrypted traffic on a Hyper-V server.

```
winrm set winrm/config/service/auth @{Basic="true"}
winrm set winrm/config/service @{AllowUnencrypted="true"}
```

`Basic="true"` enables HTTP Basic authentication through user and password in clear text.

`AllowUnencrypted= true` allows the transfer of the unencrypted data through HTTP between the server and the client.

- c. Check the WinRM configuration setup.

```
'winrm get winrm/config/client':
Client:
    NetworkDelays = 5000
    URLPrefix = wsman
    AllowUnencrypted = true
    Auth
        Basic = true
        Digest = true
        Kerberos = true
        Negotiate = true
        Certificate = true
        CredSSP = false
    DefaultPorts
        HTTP = 5985
        HTTPS = 5986
    TrustedHosts = *

'winrm get winrm/config/service' on the Hyper-V server:
Service:
    RootSDDL = O:NSG:BAD:P(A;;GA;;;BA)(A;;GA;;;S-1-5-21-3273298017-2363932476
-3643925056-1633)S:P(AU;FA;GA;;;WD)(AU;SA;GWGX;;;WD)
```

```

MaxConcurrentOperations = 4294967295
MaxConcurrentOperationsPerUser = 15
EnumerationTimeoutms = 600000
MaxConnections = 15
MaxPacketRetrievalTimeSeconds = 120
AllowUnencrypted = true
Auth
    Basic = true
    Kerberos = true
    Negotiate = true
    Certificate = true
    CredSSP = false
    CbtHardeningLevel = Relaxed
DefaultPorts
    HTTP = 5985
    HTTPS = 5986
IPv4Filter = *
IPv6Filter = *
EnableCompatibilityHttpListener = true
EnableCompatibilityHttpsListener = true
CertificateThumbprint

```

If the previous change to `Basic = true` resolved the authentication issue, apparently the client and the server are not able to negotiate the authentication protocol. The Basic authentication scheme is not recommended, unless WinRM is set up with HTTPS. It might cause security exposure by sending a user name, a password and the message body in clear text. Microsoft uses three protocols during the **Negotiate** scheme: Kerberos, NTLMV2, and NTLM. Usually, the client and the server choose the best authentication mechanisms they agree upon. BigFix Inventory does not allow the Kerberos protocol. To check whether the client and the server allow the NTLMV2 or NTLM protocol, use the following registry query.

```
reg query HKEY_LOCAL_MACHINE\System\CurrentControlSet\control\LSA\MSV1_0
```

The following parameters under this registry key control the authentication schema behavior:

- **NtlmMinClientSec**
- **NtlmMinServerSec**

To configure these parameters, see: [How to disable LM authentication on Windows NT](#). Alternatively, the parameters can be modified through the appropriate Group Policy: **Computer Configuration > Windows Settings > Local Policies > Security Options: LAN Manager authentication level**

4. To manually identify the appropriate authentication mechanism for WinRM service on the remote host, use the following commands:

- Run the command from Windows PowerShell.

```
Test-WSMan -ComputerName http://<Hyper-V_server_name>:<port>
-Auth <authentication_scheme> -Cred <user_id>
```

Where:

<Hyper-V_server_name>

Is the host name of the Hyper-V server. If you are using HTTPS, the host name must match the CN in the certificate.

<port>

Is the port number on which the Windows Remote Management client for the HTTP or HTTPS transport listens.

<authentication_scheme>

Is the authentication scheme: Basic or Negotiate

<user_id>

Is the user ID to Windows PowerShell.

- Run the command from the command prompt.

```
winrm identify -r:http://<Hyper-V_server_name>:<port>
-auth <authentication_scheme> -u:<user_id> -p:<password>
```

Where:

- **<Hyper-V_server_name>**

Is the host name of the Hyper-V server. If you are using HTTPS, the host name must match the CN in the certificate.

<port>

Is the port number on which the Windows Remote Management client for the HTTP or HTTPS transport listens.

<authentication_scheme>

Is the authentication scheme: Basic or Negotiate.

<user_id>

Is the user ID to WinRM.

<password>

Is the password to WinRM.

Event tracing

If BigFix Inventory cannot connect to the Hyper-V host, use the custom Windows Event Tracing. When the settings of the Client and the server are conflicting, the Security event log generates the following error message: `Unknown user`

name or bad password. You cannot determine that the actual issue is the unsupported authentication protocols. The Windows event tracing obtains the diagnostic data about the WMI and WinRM queries.

1. To start Event Tracing on the Hyper-V server, use the following commands:

- `logman.exe start winrmtrace -p Microsoft-Windows-Winrm -o winrmtrace.etl -ets`
- `logman.exe start wmitrace -p Microsoft-Windows-WMI-Activity -o wmitrace.etl -ets`

2. Run the VM Manager connection test.

- **Windows** `vmman.bat -testconnection`
- **Linux** `vmman.sh -testconnection`

The formatted WinRM Event Trace Log file, `winrmtrace.etl`, contains the information about the user authentication issue.

```
<Task>User authentication </Task>
<Message>Sending HTTP 401 response to the client and disconnect the connection after
sending the response</Message>
```

Additionally, the file consists of the SOAP, or XML queries that are sent from the Client and provide essential data that is useful during troubleshooting the Hyper-V problems.

Troubleshooting Hyper-V connection status with Wireshark



Note: **UNIX** To analyze the network traffic stream on Unix with no GUI environment, use `tcpdump`. To dump all the traffic to the external file use `tcpdump -vvv -XX -w tcpdump.out`. The `tcpdump.out` file can be viewed with Wireshark.

Successful connection test between BigFix Inventory and Hyper-V with PowerShell interface.

Parameters:

- **Windows** `vmman.bat -testconnection`
- **Linux** `vmman.sh -testconnection`

```
vmm_communication_interface=POWERSHELL
```

PowerShell uses the NTLM authentication protocol.

Protocol	Length	Info	port src	port dst
DCERPC	174	Bind: call_id: 3, Fragment: Single, 1 context items: ISystemActivator V0.0 (32bit NDR), NTLMSSP_NEGOTIATE	51370	135
DCERPC	330	Bind_ack: call_id: 3, Fragment: Single, max_xmit: 5840 max_recv: 5840, 1 results: Acceptance, NTLMSSP_CHALLENGE	135	51370
DCERPC	300	AUTH3: call_id: 3, Fragment: Single, NTLMSSP_AUTH, User: ADMINIB-K14GP2L\ala	51370	135
ISystemActivator	870	RemoteCreateInstance request	51370	135
TCP	54	135 → 51370 [ACK] Seq=277 Ack=1183 Win=65024 Len=0	135	51370
TCP	1434	[TCP segment of a reassembled PDU]	135	51370
ISystemActivator	58	RemoteCreateInstance response	135	51370

It does not connect to the WinRM ports: 5985 and 80, but it uses a DCE-RPC mechanism to connect to the Endpoint Mapper service via port 135.

1. Select the `RemoteCreateInstance` response entry.
2. To check whether the authentication finished successfully, expand the parameters, and check the value of `HResult`. If authentication was success, `HResult` is marked as `S_OK`.

```

Distributed Computing Environment / Remote Procedure Call (DCE/RPC) Response,
ISystemActivator ISystemActivator Resolver, RemoteCreateInstance
  Operation: RemoteCreateInstance (4)
  [Request in frame: 34]
  DCOM, ORPCThat
  IActProperties
  HResult: S_OK (0x00000000)

```

PowerShell v2 by default uses NTLMv1 for negotiation. To update NTLM, use the following registry setting:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\LSA\LMCompatibilityLevel
```

For more information, see: [How to enable NTLM 2 authentication](#).

Unsuccessful connection test between BigFix Inventory and Hyper-V with NTLM interface.

Parameters:

- `vmm_communication_interface=NTLM`
- `vmm_url=http://.../wsman`
- The default port: 80

Protocol	Length	Info	port src	port dst
TCP	66	52905 → 80 [SYN] Seq=0 Win=8192 Len=0 MSS=1460 WS=4 SACK_PERM=1	52905	80
TCP	54	80 → 52905 [RST, ACK] Seq=1 Ack=1 Win=0 Len=0	80	52905

The connection test failed, and the `trace.log` contains the following error messages.

```
(...)com.ibm.license.mgmt.vmmmanager.hyperv.net.HttpConnector::send::
An error occurred while trying to send request to http://.../wsman
```

```
(...)com.ibm.license.mgmt.vmmmanager.hyperv.net.HttpConnector::send:java.net.
ConnectException: Connection refused: connect
```

While the errors from the `trace.log` file do not provide enough details, the Wireshark capture proves that the issue is related to TCP. Use the following command to check for listener on port 80.

```
winrm enumerate winrm/config/listener
```

If no listeners are detected, run the following command to set up listeners.

```
winrm set winrm/config/service @{EnableCompatibilityHttpListener="true" }
```



Note: The setup of listeners on port 80 can collide with the other HTTP services on this computer. To avoid the issues, specify a dedicated WinRM port, such as 5985 or 5986, on the `vmm_url` connection string.

Unsuccessful connection test between BigFix Inventory and Hyper-V with NTLM interface. HTTP Error 401.

1. The connection test failed, and the `trace.log` contains the following error messages.

```
(...)com.ibm.license.mgmt.vmmanger.hyperv.net.HttpConnector::retrieve::  
Response Code is: 401
```

```
(...)com.ibm.license.mgmt.vmmanger.hyperv.net.HttpConnector::retrieve::  
The following response code was returned while connecting to VM Manager  
http://.../wsman: responseCode = 401
```

The following Wireshark capture lists the NTLM negotiation.

Protocol	Length	Info	port src	port dst
HTTP	399	POST /wsman HTTP/1.1 , NTLMSSP_NEGOTIATE	54461	80
HTTP	468	HTTP/1.1 401 , NTLMSSP_CHALLENGE	80	54461
HTTP	767	POST /wsman HTTP/1.1 , NTLMSSP_AUTH, User: \ala	54461	80
HTTP	158	HTTP/1.1 200	80	54461
TCP	279	[TCP segment of a reassembled PDU]	54461	80
HTTP/XML	1239	POST /wsman HTTP/1.1	54461	80
HTTP	206	HTTP/1.1 401	80	54461

The first four rows show the successful NTLM negotiation. The fourth row ends the so-called four-way NTLM handshake. The last line detects HTTP Error 401: Unauthorized, and expands into the following HTTP response.

```
Hypertext Transfer Protocol  
HTTP/1.1 401 \r\n  
Server: Microsoft-HTTPAPI/2.0\r\n  
WWW-Authenticate: Negotiate\r\n  
Date: Thu, 01 Sep 2016 13:15:06 GMT\r\n  
Connection: close\r\n
```

The `WWW-Authenticate: Negotiate` response indicates that the server is ready to use the NTLM, NTLMV2, or HTTP Basic authentication scheme. However, the server already successfully responded to the negotiation sequence. Ensure that the WinRM **AllowUnencrypted = true**.



Note: VM Managers do not provide the HTTP body encryption. WinRM accepts only the HTTP body that is encrypted with **Negotiate**, or Kerberos protocol.

2. The Wireshark capture shows the unsuccessful NTLM negotiation handshake.

Protocol	Length	Info	port src	port dst
HTTP	399	POST /wsman HTTP/1.1 , NTLMSSP_NEGOTIATE	55445	80
HTTP	468	HTTP/1.1 401 , NTLMSSP_CHALLENGE	80	55445
HTTP	767	POST /wsman HTTP/1.1 , NTLMSSP_AUTH, User: \ala	55445	80
HTTP	206	HTTP/1.1 401	80	55445

The four-way NTLM handshake ends with a 401 error. The most common reason behind this error is the lack of sufficient access rights. To check whether the accounts assigned to the Administrators group are allowed to access WinRM, configure the following Windows registry entry.

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\LocalAccountTokenFilterPolicy
```

Unsuccessful connection test between BigFix Inventory and Hyper-V caused by the bug related to `vmm_login` domain prefix interpretation.

When you add a VM Manager to BigFix Inventory, make sure that you provide the administrator user name in one of the following formats:

- `user_name@domain`, for example: `test@cluster.com`
- `user_name\domain`, for example: `test\cluster.com`

9.2.16

Starting from application update 9.2.16, you can also provide the administrator user name in one of the following formats:

- `domain@user_name`, for example: `cluster.com@test`
- `domain\user_name`, for example: `cluster.com\test`



Note: Make sure that you upgrade the VM Manager Tool to version 9.2.16, not only the BigFix Inventory server.

If you do not follow this pattern, and provide the user name in the incorrect format, an error is generated during the VM Manager authentication.

Protocol	Length	Info	port src	port dst
HTTP	407	POST /wsman HTTP/1.1 , NTLMSSP_NEGOTIATE	55878	80
HTTP	468	HTTP/1.1 401 , NTLMSSP_CHALLENGE	80	55878
HTTP	799	POST /wsman HTTP/1.1 , NTLMSSP_AUTH, User: ala\NC9143126002	55878	80
HTTP	206	HTTP/1.1 401	80	55878

```
[truncated]Authorization: Negotiate T1RMTVNTUAAI
NTLM Secure Service Provider
  NTLMSSP identifier: NTLMSSP
  NTLM Message Type: NTLMSSP_AUTH (0x00000003)
  Lan Manager Response: 8fa95512dff721edd3c8a1
  NTLM Response: 0605a955cb3cc106b028874a67c97
  Domain name: ala
  User name: NC9143126002
  Host name: ADMINIB-K14GP2L
```

When the error occurs, the user name is assigned to the incorrect field.

Kernel-based Virtual Machine

Kernel-based Virtual Machine is one of the virtualization technologies supported by BigFix Inventory.

Purpose

Kernel-based Virtual Machine (KVM) is a full virtualization solution for Linux™. In the KVM architecture, each guest (virtual machine) is implemented as a regular Linux™ process. After you install KVM, you can run multiple guests, with each of them running a different operating system image. Each of these virtual machines has private, virtualized hardware, which includes memory, storage, graphics adapter, and a network card. This allows KVM to benefit from all the features of the Linux™ kernel.

Collecting capacity data

Capacity data is collected from virtual machines that are managed by KVM hosts in two ways:

- Capacity data from x86 KVM hosts that are controlled by RHEV-M is collected by the VM Manager Tool.
- Capacity data from x86 KVM hosts that are not controlled by RHEV-M and from PowerKVM hosts is collected directly from the host by using the **Run Capacity Scan on Virtualization Hosts** fixlet. For more information, see: [Collecting capacity data directly from KVM hosts \(on page ccxvi\)](#).

Red Hat Enterprise Virtualization Manager (RHEV-M)

Red Hat Enterprise Virtualization (RHEV) is an enterprise virtualization product based on the KVM hypervisor. RHEV-M is a service running on a Red Hat Enterprise Linux™ server that provides interfaces for controlling the virtualization platform. It manages provisioning, connection protocols, user sessions logins and logoffs, virtual desktop pools, virtual machine images, and the high availability and clustering systems. RHEV-M provides the REST API that is used by BigFix Inventory to collect information about the whole infrastructure that is managed by RHEV-M.

The default URL that is to be used:

- For RHEV-M 3.0

```
https://<RHEV-M_IP_address>:8443/api
```

- For RHEV-M 3.1 up to 3.6

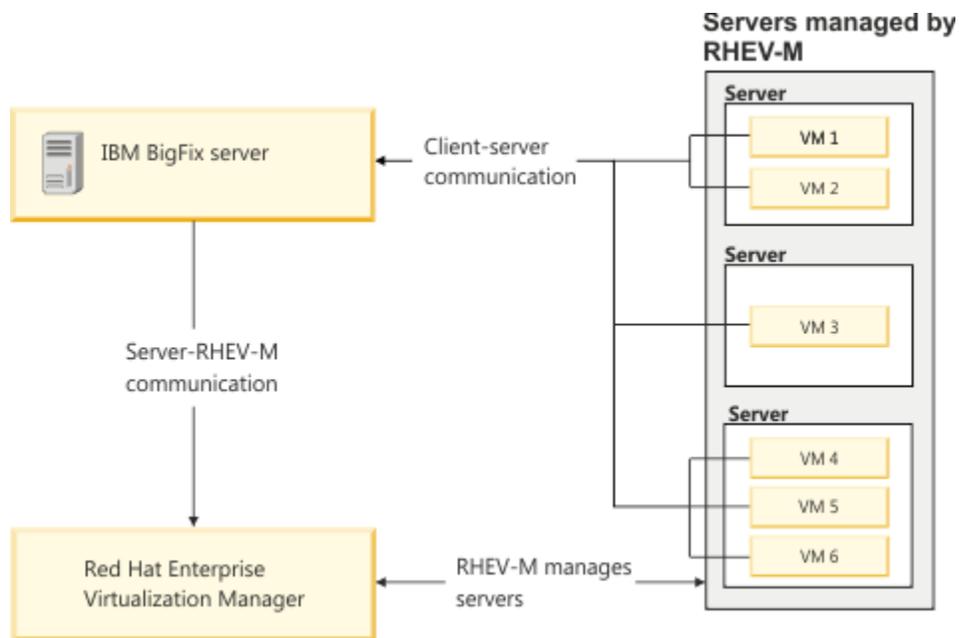
```
https://<RHEV-M_IP_address>/api
```

- For RHEV-M 4.0

```
https://<RHEV-M_IP_address>/ovirt-engine/api
```

Supported versions:

For a full list of supported versions, see: [Supported virtualization types \(on page cccxx\)](#).



Nutanix

Available from BigFix Inventory 9.2.17. Nutanix is one of the virtualization technologies supported by BigFix Inventory.

Purpose

Nutanix solutions are built on the hyperconverged infrastructure (HCI) technology. Nutanix Acropolis is a foundation of Nutanix Enterprise Cloud platform that provides built-in virtualization, storage services, networking and cross-hypervisor application mobility. Nutanix supports all of the leading hypervisors, including virtualization solutions from VMware, Microsoft, Citrix, and the native Nutanix hypervisor, AHV (Acropolis Hypervisor) which is a native virtualization solution that supports IBM Power and x86 based hosts depending on version.

Prism

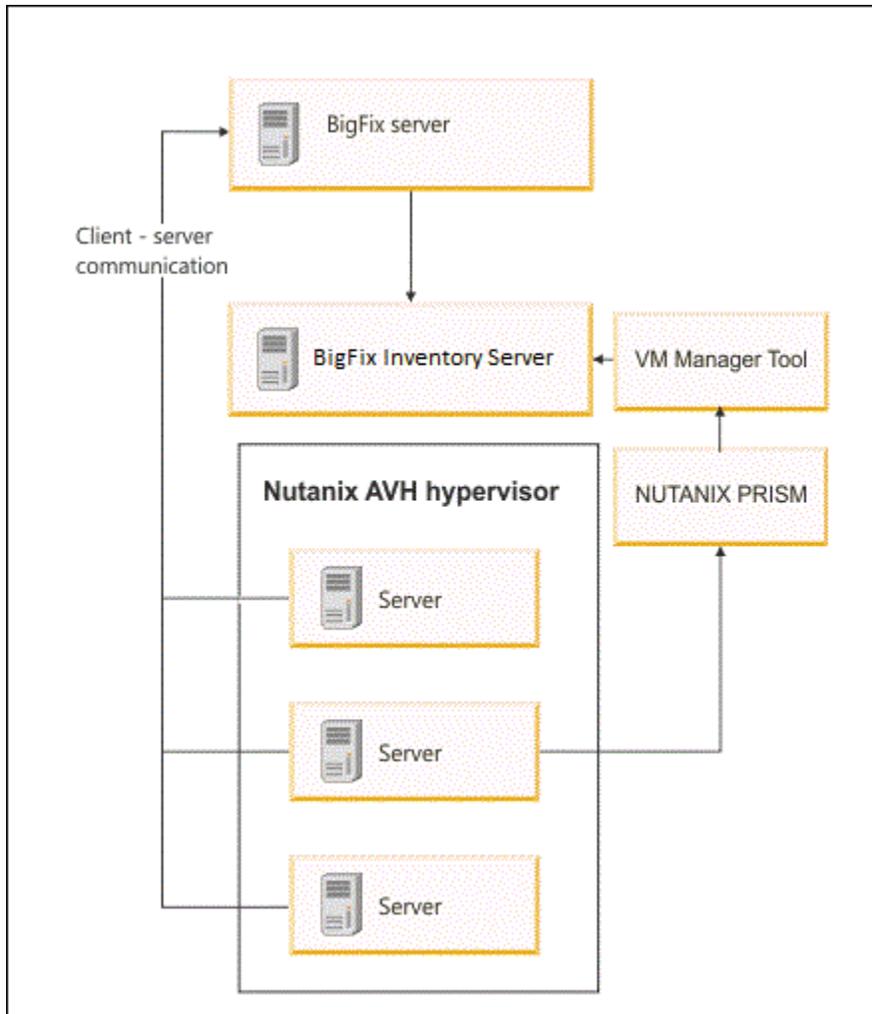
Nutanix Prism is a management pane that provides unified interface for Nutanix clusters. It is used to monitor and manage entities across Nutanix clusters through a web console. It allows to run a mixed environment, or easily switch hypervisors to meet changing needs.

The default URL is:

```
https://<NUTANIX_PRISM_IP_ADDRESS>:9440/PrismGateway/services/rest/v2.0
```

Supported technologies:

currently supports Nutanix technology limited to Nutanix AVH hypervisor running on IBM Power Hyperconverged Systems (CS Appliances) via Prism. Nutanix AVH hypervisor on x86 is currently not supported.



Oracle VM Server for x86

9.2.14 Available from 9.2.14. Oracle VM Server for x86 is one of the virtualization technologies supported by BigFix Inventory.

Purpose

Oracle VM Server for x86 is a full virtualization solution that manages the Oracle VM environment. BigFix Inventory supports the Oracle VM Servers for x86 which are managed by Oracle VM Managers. BigFix Inventory connects to Oracle VM Manager and collects the software and capacity data. The solution does not support the direct connection to the Oracle VM Server for x86.

! **Important:** The Oracle VM Manager to which BigFix Inventory connects can manage both Oracle VM Servers for x86 and SPARC. However, this solution does not retrieve the data from SPARC Servers. Oracle VM Server for SPARC is supported by BigFix Inventory on different terms. For more information, see: [Client Installation on Oracle Solaris \(on page ccxii\)](#).

Oracle VM Manager

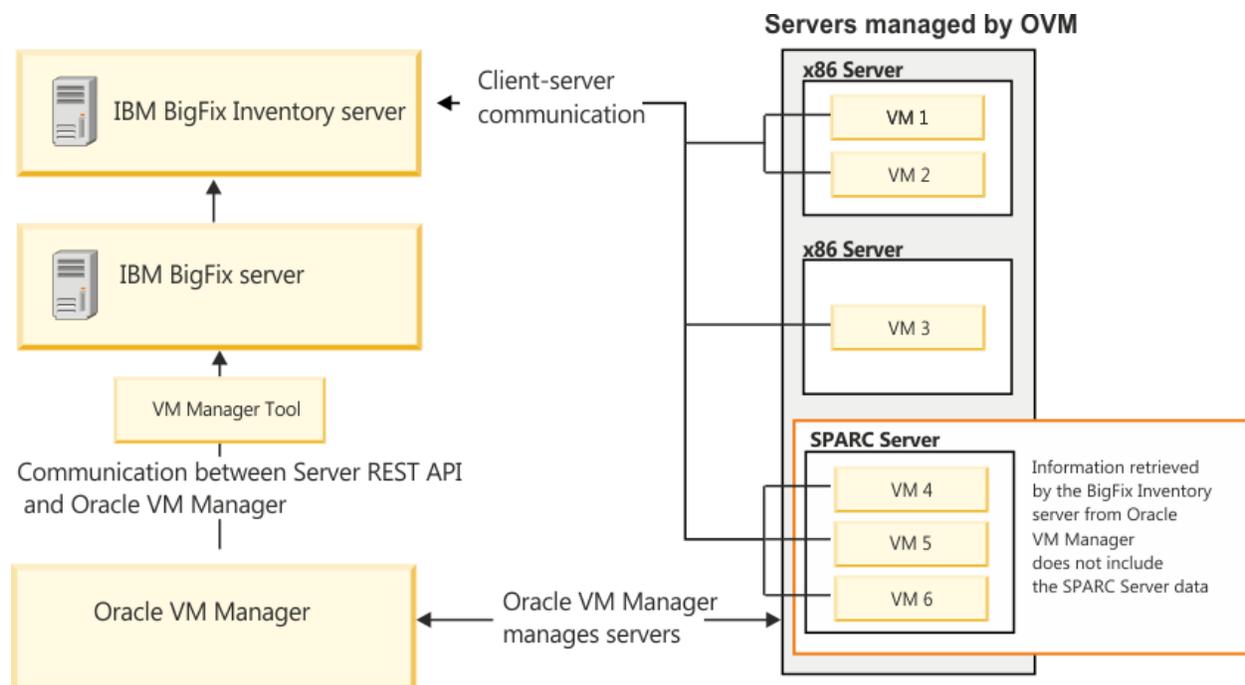
Oracle VM Manager provides interfaces for controlling the virtualization platform. It manages provisioning, connection protocols, user sessions logins and logoffs, virtual desktop pools, virtual machine images, and the high availability and clustering systems. Oracle VM Manager provides the REST API that is used by BigFix Inventory to collect information about the infrastructure that is managed by Oracle VM Manager.

The default URL that is to be used:

```
https://<OVM-Manager_IP_address>:7002/ovm/core/wsapi/rest
```

Supported versions:

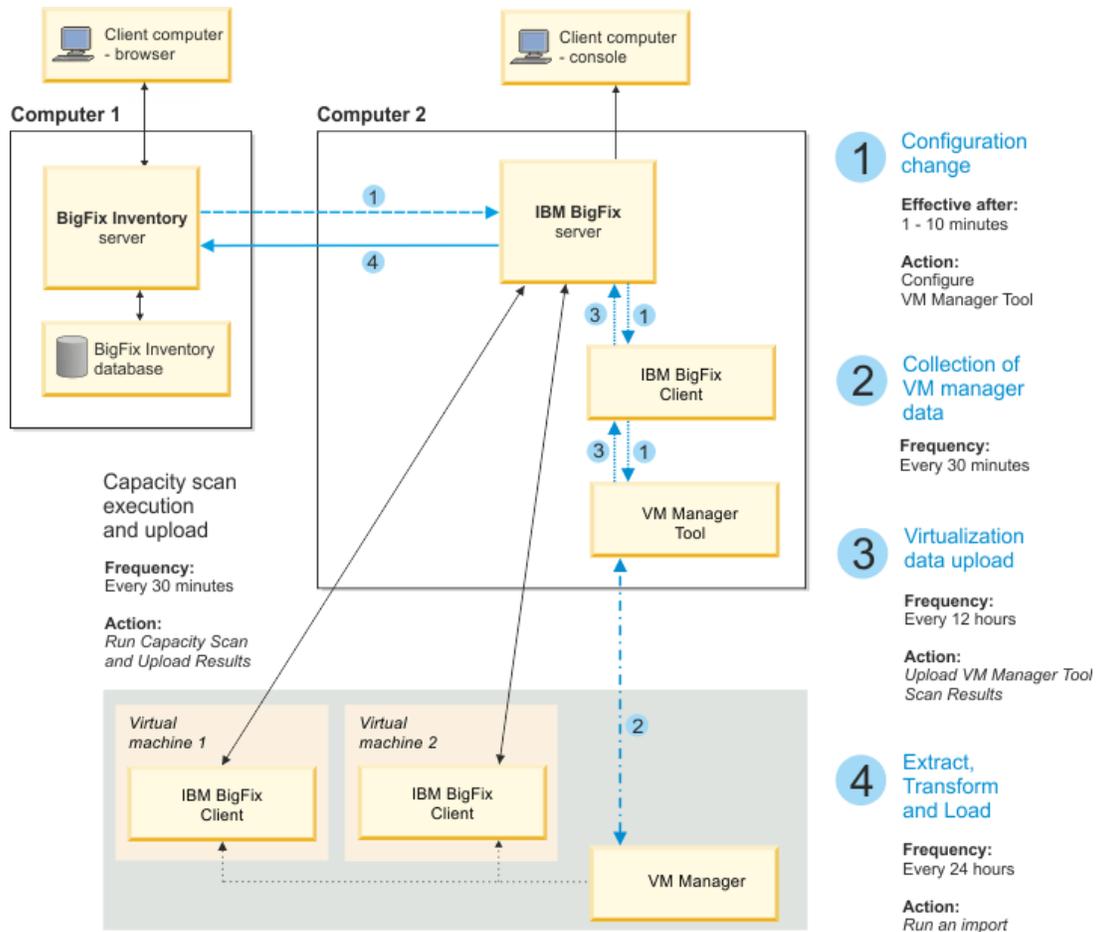
For a full list of supported versions, see: [Supported virtualization types \(on page cccxx\)](#).



Flow of capacity data from VM managers

To understand how capacity data is gathered in virtual environments, see what data flows between virtual machines, VM Manager Tool, and the BigFix Inventory server, and how frequently it is transferred.

The following diagram shows the flow of data between virtual machines, VM Manager Tool, and the BigFix Inventory server in case of centralized VM management where the VM Manager Tool is installed on the same computer as the BigFix server.



It might happen that the **VM Managers** panel does not display up-to-date information about your virtual systems. The delay might be caused by the fact that the data about your virtual environment and the configuration data goes through many infrastructure elements and depends on various scheduled tasks. As a result, the time that is needed for the complete virtualization data to reach BigFix Inventory and be displayed in the interface might last up to two days depending on your specific configurations.

Table 38. Frequency of steps of collecting capacity data

The table consists of three columns and four body rows.

Number	Step	Frequency	Comments
1.	Changes in configuration of the VM Manager Tool and VM manager connections	After every change of configuration	are sent to the endpoints.

Table 38. Frequency of steps of collecting capacity data

The table consists of three columns and four body rows.

(continued)

Number	Step	Frequency	Comments
2.	VM manager data is collected from VM managers by the VM Manager Tool.	Every 30 minutes	 Restriction: In case of subcapacity reporting, it is necessary to capture mobility of VMs. To ensure that it is captured, data from VM managers must be collected every 30 minutes. You can change the frequency of collecting data only if you use BigFix Inventory for purposes other than subcapacity reporting.
3.	VM manager data is uploaded from the VM Manager Tool to the BigFix server.	Every 12 hours	
4.	The current infrastructure data from the clients and the VM Manager Tool is imported from the BigFix server to the BigFix Inventory server.	Every 24 hours	

Basic VM management (central)

Use basic VM management if the BigFix server can connect to VM managers in your infrastructure. It is the default method of configuring connections to VM managers through the BigFix Inventory user interface.

Data about the capacity of VM managers is needed to calculate consumption of license metrics in the virtual environment. In basic VM management, the data is collected by the main instance of the VM Manager Tool that is by default installed on the BigFix server. You must only ensure that it can connect to your VM managers. If some of your VM managers are in separated networks with no connection from the BigFix server, manage these VM managers through advanced VM management. For more information, see: [Advanced VM management \(distributed\) \(on page cccxlviij\)](#).

Understanding the **VM Managers** panel

Before you start adding VM managers, see the details of the **VM Managers** panel to understand the available options and columns.

To open the **VM Managers** panel, click **Management > VM Managers**. Starting from application update 9.2.4, the panel lists not only VM managers that are managed in the central mode, but also those from the distributed mode. However, options such as editing, deleting, or testing connection are applicable only for VM managers that are managed in the central mode. For information how to perform these tasks for VM managers that are managed in the distributed mode, see: [Advanced VM management \(distributed\) \(on page ccclxvii\)](#).

The number of VM managers that are shown on the panel depends on your computer group assignment. You can see only VM managers from which data is collected by the VM Manager Tool that belongs to the computer group to which you are assigned.

Operation Status	Connection Test...	Type	URL	Data Import Time	Data Collection ...	Computer ID of V...	Deployment Type
OK	Successful (2018-...	VMware ESX, ESX...	http://9.174.39.77:7...	05/10/2018 04:33 ...	05/10/2018 04:33 AM	11448319	Central
Connection fail...	Not Tested	VMware ESX, ESX...	https://198.51.100/...	05/10/2018 10:25 AM	05/10/2018 10:25 AM	11448319	Distributed

1 New

Click this button to add a VM manager. Use this option to add VM managers to which the BigFix server can connect. If the VM manager is in a separated network to which the BigFix server cannot connect, create the VM manager configuration file. For more information, see: [Advanced VM management \(distributed\) \(on page ccclxvii\)](#).

! **Important:** The button is disabled when the computer where the main instance of the VM Manager Tool is installed does NOT belong to the computer group to which you are assigned.

2 Delete

Click this button to delete the selected VM manager. The option is applicable only to VM managers that are managed in the central mode. If the selected VM manager is managed in the distributed mode, the button is disabled. To delete such a VM manager, remove its configuration file from the computer that is listed in the **Computer ID of VM Manager Tool** column.

! **Important:** When you delete a configuration file of a VM manager that is managed in the distributed mode, the VM manager is removed from the BigFix Inventory user interface after the maximum visibility period is exceeded. The period is set by the **maxVMManagerVisibility** parameter and is set to 90 days by default.

3 Test Connection

Click this button to test the connection to the selected VM manager. The option is applicable only to VM managers that are managed in the central mode. If the selected VM manager is managed in the distributed mode, the button is disabled. The status of the connection test is shown in the **Connection Test Status** column.

4 9.2.12 Resume

Click this button to resume a connection to a VM manager that is in the *Invalid credentials - suspended* status. When the number of failed login attempts is exceeded, connection to the VM manager is suspended to prevent the account from locking. After you correct the underlying issue, you can resume the connection.

5 More information

A link to a documentation topic that describes how to manage VM managers.

6 Operation Status

The status of the VM manager that reflects its condition. The status changes when connection or configuration problems occur. For more information, see: [VM manager statuses \(on page ccclii\)](#).

7 Connection Test Status

The status of the connection test. The initial value is `Not Tested` and it changes to `Testing`, `Successful`, or `Failed` after you click **Test Connection**. The latter two statuses are always followed by a time stamp. For more information, see: [Connection test statuses \(on page ccclvi\)](#).

The option is applicable only to VM managers that are managed in the central mode. If the selected VM manager is managed in the distributed mode, the connection status is `Not Applicable`.

8 Type

Type of the VM manager.

9 URL

Web address of the VM manager. Each virtualization type uses a different address format. The template is always provided when you add VM managers. For more information, see: [Supported virtualization types \(on page cccxx\)](#).

10 Data Import Time

Date and time when the capacity data was imported from the BigFix server to BigFix Inventory.

11 Data Collection Time

Date and time when the capacity data was collected from the VM manager by VM Manager Tool. The time that is shown in this column might differ from the time in the Data Import Time column because the collected data might still need to be uploaded to the BigFix server and then imported to BigFix Inventory.

12 9.2.4 Computer ID of VM Manager Tool

Identifier of the computer on which the VM Manager Tool that collects data from the particular VM manager is installed. Click the link to display the details of the computer, for example its IP address. You can use the information to learn which computer to access to change the configuration file of the VM Manager Tool in the distributed mode.

13 9.2.4 Deployment Type

Indicates whether the VM manager is deployed in the central or distributed mode. If the deployment type is central, you can edit the VM manager on the BigFix Inventory user interface.

If the deployment type is distributed, you can view information about the VM manager on the BigFix Inventory user interface, but you cannot edit nor delete it. To change the settings of such a VM manager, edit its configuration file. The file is on the computer that is listed in the Computer ID of VM Manager Tool column. For more information, see: [VM manager configuration parameters \(on page ccclxxvi\)](#).

Deprecated: User Name

In earlier versions, this column shows the user name that is used to connect to the VM manager.

Adding VM managers in central mode

Configure connections to VM managers in your infrastructure to ensure that the VM Manager Tool can gather data that is needed to calculate consumption of license metrics in your virtual environment. To configure a connection to VM managers through the BigFix Inventory user interface, open the **VM Managers** panel, and specify details of VM managers, such as the web address, virtualization type, and credentials.

-  You must have the Manage VM Managers and Servers permission to perform this task.
- The computer where the main instance of the VM Manager Tool is installed must belong to the computer group to which you are assigned. Otherwise, the option to add VM managers is disabled.
- If the panel is blocked and displays a warning, see: [Troubleshooting: Enabling the VM Managers panel \(on page ccclxiv\)](#).

9.2.7 VM Manager Tool, starting from update 1.5.0.0, is enhanced to attempt to automatically fix the most common problems with the VM manager connection parameters. This feature is disabled by default. To enable this feature, go to **Management > Advanced Server Settings** and change the setting of **vmman_fix_parameters_enabled** parameter to `true`. While fixing connection problems the VM Manager Tool can connect using the HTTP protocol when the defined HTTPS connection is not available. This setting is enabled by default. To disable switching protocol from HTTPS to HTTP, go to **Management > Advanced Server Settings** and change the setting of **vmman_http_connection_allowed** parameter to `false`. For more information, see: [Advanced server settings \(on page cdii\)](#).

1. In the top navigation bar, click **Management > VM Managers**.
2. To add a VM manager, click **New**.
3. Select the virtualization type and provide the required details.

- To add Microsoft Hyper-V:
 - a. Select one of the available communication interfaces: PowerShell or WinRM. For more information about these interfaces, see: [Microsoft Hyper-V \(on page cccxxiv\)](#).
 -  **Important:** WinRM is the preferred communication interface. Before you add the VM Manager, configure WinRM on the Hyper-V host. For more information, see: [Configuring WinRM on Hyper-V hosts \(on page cccxxviii\)](#).
 - b. Provide the URL of the VM manager in the following format: `https://<Hyper-V_IP_address>/wsman`.
 - c. Optional: To share credentials with other hosts in the same cluster, select **Share credentials with other hosts in the same cluster**.
 - d. Provide the Administrator account credentials. Define the user as `user_name\domain` or `user_name@domain`. For example: `test\cluster.com` or `test@cluster.com`.
- To add KVM with RHEV-M:
 - a. Provide the URL of the VM manager in the following format:
 - For version 3.0: `https://<RHEV-M_IP_address>:8443/api`
 - For version 3.1 up to 3.6: `https://<RHEV-M_IP_address>/api`
 - For version 4.0: `https://<RHEV-M_IP_address>/ovirt-engine/api`
 - b. Provide the account credentials. Define the user as `user_name@domain`, for example: `test@cluster.com`.
- **9.2.17** To add Nutanix:
 - a. Provide the URL of the VM manager in the following format: `https://<NUTANIX_PRISM_IP_ADDRESS>:9440/PrismGateway/services/rest/v2.0`.
 - b. Provide the account credentials.
- To add VMware ESX, ESXi or vCenter:
 - a. Provide the URL of the VM manager in the following format:
 - For versions up to 5.1: `https://<vCenter_IP_address>/sdk`
 - For version 5.5 and higher: `https://<vCenter_IP_address>/sdk/vimService.wsdl`
 - b. Provide the account credentials. Define the user as `domain\user_name`, for example: `cluster.com\test`.
- **9.2.12** To add XenServer or Citrix XenServer:
 - a. Provide the URL of the VM manager in the following format: `https://xen-server`.
 - b. Provide the account credentials.

The following example shows a configured connection to VMware ESX, ESXi or vCenter.

Create VM Manager

VM Manager Type*

URL*

URL Template
For versions lower than 5.5: <https://virtualcenter/sdk>
For version 5.5 and higher: <https://virtualcenter/sdk/vimService.wsdl>

User Name*

Password*

4. Click **Create**. Connection to the VM manager is created and its status is Pending.
5. Select the VM manager, and click **Test Connection**.
The test might take a few minutes. Refresh the panel. If the VM manager connection is properly configured, the connection test status changes to Successful.

You added a VM manager. Its status remains Pending until the data is collected from the VM manager, uploaded to the BigFix server, and transferred to BigFix Inventory during the import of data.

The upload of data collected from VM managers is triggered shortly after you modify any VM manager, for example create a new connection or change the existing one. If you do not make any modifications, the data is uploaded according to the schedule, which is every 12 hours by default. For more information, see: [Capacity data flow \(on page cccx/v\)](#).

VM manager statuses

After you add a VM manager, see the explanation of its status to ensure that the VM manager works properly.

VM manager statuses

The statuses apply to VM managers that are managed both in central and distributed mode.

Connection failed

Connection to the VM manager failed and the data was not retrieved.

Action:

- To determine the cause of the problem, open the VM Manager Tool log files that are in the *BES Client\LMT\VMMAN\logs* directory, and look for the following exception.

```
VM Manager is not able to communicate with ESX 4.x due to:
javax.net.ssl.SSLHandshakeException: java.security.cert.CertificateException:
Certificates does not conform to algorithm constraints
```

For more information, see: [RQMCopyUtility error: SSL_TLS SSLContext not available. Certificate does not conform to algorithm constraints.](#)



Note: ESX 4 is no longer supported by VMWare. For more information, see: [End of Availability and End of Support Life for ESX 4.x and VMware Management Assistant.](#)

- In order to debug log in process modify the following file.
 - **Linux** `BES_Client/LMT/VMMAN/vmman.sh`
 - **Windows** `BES_Client\LMT\VMMAN\vmman.bat`

Update the Properties Definitions section by adding the debug dump option to **VMM_PROPERTIES_DEFS** parameter. Insert the following line:

- **Linux** `VMM_PROPERTIES_DEFS="-Dcom.sun.xml.internal.ws.transport.http.client.HttpTransportPipe.dump=true -Dhttps.protocols=TLSv1.2,TLSv1.1,TLSv1 -Dcom.ibm.jsse2.disableSSLv3=false -Dcom.ibm.tools.attach.enable=no`
- **Windows** `SET VMM_PROPERTIES_DEFS=%VMM_PROPERTIES_DEFS% -Dcom.sun.xml.internal.ws.transport.http.client.HttpTransportPipe.dump=true`

Run the following VMMAN tool command:

- **Linux** `vmman.sh -testconnection > vmman.out 2>&1`
- **Windows** `vmman.bat -testconnection > vmman.out 2>&1`

Verify the command output by checking credentials and VM manager response.



Data partially collected

Data was collected only from some of the hosts.

Action:

- If the problem occurs on Hyper-V, it indicates that some of the host in the cluster have connectivity issues. To identify the problematic hosts, open the VM Manager Tool log files that are in the `BES_Client\LMT\VMMAN\logs` directory and search for the following entries:
 - `Processing of the host: <HOST_UUID> ip: <HOST_IP_ADDRESS> from hyper-v cluster was skipped because of invalid connection status (<STATUS>).`
 - `A communication error occurred when processing one of the Hyper-V hosts in the cluster. Host name: <HOSTNAME>. Skip and continue.`

Fix connectivity issues for the identified hosts. If the problem persists, contact IBM Support.

- If the problem occurs on XenServer, it indicates that XenTools are not installed on some of the virtual machines. To identify the problematic virtual machines, open the VM Manager Tool log files that are in the `BES_Client\LMT\VMMAN\logs` directory and search for the following entries:

```
◦ It was not possible to retrieve the number of cores assigned to XEN VM with UUID:  
<VM_UUID>. Most likely the XenServer Tools are not installed on this particular VM.  
Skipping VM and continue.
```

Install XenTools on the identified virtual machines.

Duplicated address

The address of the VM manager is duplicated. The first VM manager with this address remains active and the status is displayed for the remaining ones.

Action: Delete duplicated VM managers to ensure that data is gathered only once. The duplicated addresses might occur only within a particular VM manager type.

Hard timeout - suspended

Connection to the VM manager failed and no further attempts are made. The status indicates that either the provided URL or configuration is incorrect, or that there are problems with the network.

Action: Check the connection between VM managers and the VM Manager Tool. Also, correct the URL and the credentials, and save the changes. Ensure that the `vmm_communication_locked` property is set to `false` in the corresponding `vmmconf_*.properties` file.

Inactive

Data from the VM manager was not retrieved for a period longer than the maximum VM manager inactivity. The status indicates that either the provided URL or configuration is incorrect.

Action: Correct the URL and the credentials, and save the changes.

Invalid credentials - attempting

Connection to the VM manager failed because the credentials are incorrect or the password expired. More attempts are made until the limit of failed login attempts is exceeded. Then, the status changes to **Invalid credentials - suspended** to prevent the account from locking.

Action: Edit the VM manager and correct the credentials.

Invalid credentials - suspended

Connection to the VM manager is suspended because the number of failed login attempts was exceeded. The status is used to prevent the account from locking if invalid credentials were provided.

Action:

- Ensure that the VM manager credentials are correct. Then, click **Resume**.
- Go to `vmmconf_*.properties` file and make sure that the `vmm_communication_locked` parameter is set to `false`.
- By default, VM Manager Tool is installed on the same computer as the BigFix server. Check the connection between VM Manager Tool and VM manager. Make sure that the firewall, proxy and NAT settings are correct.

9.2.5



No data (previously Insufficient rights)

No data was collected from the VM manager. The most common reason is that there is no data or that the user that is defined for this VM manager does not have sufficient rights.

Action: If you run the **Run Capacity Scan on Virtualization Hosts** fixlet, ensure that the virtualization hosts from which you want to collect capacity data have at least one virtual machine that is correctly defined and running.

Make sure that the user has read access to these virtual machines that are managed by this VM manager and their host.



OK

Data is collected from the VM manager and processed according to the schedule. No issues were detected.



OK - duplicated UUIDs discarded

Each virtual machine must have a unique UUID. When a UUID identifies more than one VM, both VMs are discarded from the results. The problem occurs most frequently on VMWare and is usually caused by one of the following reasons:

- Virtual machines were deployed from catalog templates in vCloud Director. By default, such machines are assigned the same BIOS UUID. For more information, see: [BIOS UUIDs in vCloud Director are not unique when virtual machines are deployed from catalog templates](#).
- Virtual machines were cloned without changing their UUIDs. For more information, see: [Changing or keeping a UUID for a moved virtual machine](#).

Action: Open the VM Manager Tool log files that are in the `BES Client\LMT\VMMAN\logs` directory and identify discarded virtual machines. Then, refer to the documentation for your virtualization type to learn how to assign unique UUIDs to each virtual machine. For information about VMWare, see: [Editing a virtual machine with a duplicate UUID.bios](#).

After the duplicates are corrected, the VM manager status changes to **OK**.



Pending

The status is applicable only to VM managers that are managed in the central mode. It is displayed after a VM manager is created or after its definition is modified. It indicates that changes are still

being propagated to the VM Manager Tool. The status changes to **OK** when the following steps are completed:

1. VM Manager Tool gathers data from VM managers.
2. The data is uploaded to the BigFix server.
3. The data is transferred to BigFix Inventory with a data import.

If a problem occurs during any of these steps, the status changes to indicate the problem.

Action: Wait until the VM Manager Tool completes the steps that are required to upload the data. For more information about each step, see: [Capacity data flow \(on page cccxlv\)](#).

If the status of the VM manager remains **Pending** for a longer period of time, it might indicate that the Schedule VM Manager Tool Scan Results Upload task was not scheduled or was stopped. To check whether it is the case, perform the following steps.

1. Log in to the BigFix console, and go to **Actions**.
2. Look for the Schedule VM Manager Tool Scan Results Upload action and check its status.
3. If the action does not exist, or its status is Stopped or Expired, start the action.
 - a. Go to **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
 - b. In the upper right pane, select **Schedule VM Manager Tool Scan Results Upload**, and click **Take Action**.
 - c. Select the computer on which the BigFix server is installed, and click **OK**.

After you schedule the action, wait for the collected data to be uploaded to the BigFix server and then imported to BigFix Inventory. The status of the VM manager should change to **OK** after the import.

Unknown problem

The problem does not fall into any other category. Check the [VM Manager Tool log \(on page cccxcj\)](#) to determine the cause.

Action: If you are not able to determine the cause of the problem from the VM Manager Tool log, contact IBM Support.

Connection test statuses

The following statuses are displayed in the **Connection Test Status** column after you select a VM manager and click **Test Connection**. The statuses are applicable only to VM managers that are managed in central mode. If the VM manager is managed in the distributed mode, the connection test status is `Not Applicable`.

Not Tested

The connection test was never started. Select a VM manager and click **Test Connection**.

Testing

The connection is being tested. It might take a few minutes until the status changes. Refresh the panel.

Successful

The connection to a VM manager was established. The VM manager is configured correctly.

Failed

The connection to a VM manager could not be established. Check the VM Manager Tool log files that are in the `BES Client\LMT\VMMAN\logs` to determine the cause.

Troubleshooting No VM Manager data

Troubleshoot, and understand why the VM Manager data cannot be retrieved by BigFix Inventory. Apply the listed solutions for the VM Manager connections with  **No VM Manager data**.

Background

The basic virtual environment is configured with one central VM Manager Tool that is installed on the same computer as the BigFix server. However, you can use a more complex environment, with many VM Manager Tools. Each tool must be connected with every hypervisor that is configured for this VM Manager, and the results from every such hypervisor must reach BigFix Inventory.

Before you start troubleshooting the issue, collect the following information:

- The number of VM Manager Tools that are installed in your environment.
- The number of hypervisors to be connected.
- Valid credentials with sufficient rights to log in to hypervisors.

9.2.12 Starting from application update 9.2.12, you do not need to go to the computer where VM Manager Tool is installed, to get the data that are needed for troubleshooting purposes. When cannot retrieve VM manager data from a computer, go to the **Computer Support Data** panel and download the logs from the related computer to investigate the issue. For more information about how to download the log package, see: [Collecting logs for troubleshooting purposes \(on page dccliv\)](#).

Reasons behind the data retrieval failure

There are several possible reasons why the VM Manager data cannot be retrieved by BigFix Inventory:

- The hypervisor connection is not defined.
- VM Manager Tool cannot connect to the hypervisor to collect the data.
- VM Manager Tool does not generate results, or does not properly communicate them to BigFix Inventory.
- VM Manager Tool collects invalid or incomplete results.

Troubleshooting

Check whether all VM Managers are added to BigFix Inventory.

1. Log in to BigFix Inventory.
2. Go to **Management > VM Managers**, and check whether all the VM Manager connections are defined. If not, see: [Adding VM Managers in central mode \(on page cccl\)](#) or [Adding VM Managers in distributed mode \(on page ccclxxii\)](#).



Note: Make sure that the VM Manager URL is in the correct format. The format differs between technologies, and versions.

Check whether all VM Manager connections are in the OK status.

If any of the VM Manager connections is in status other than  OK, refer to: [VM manager statuses \(on page ccclii\)](#).

Check whether VM Manager Tool is properly configured.

1. Check whether the VM Manager Tool is properly installed and configured.

The VM Manager Tool is installed in the `BESClient/LMT/VMMAN` directory.

- a. To ensure that the tool was installed successfully, log in to the BigFix console.
- b. Click **Actions** in the navigation tree, and select **Install VM Manager Tool**.
- c. Switch to the **Computers** tab, and check whether the status is set to Completed.
- d. If the installation failed, check which line of the action script caused the failure. Double-click the listed computer, and analyze the **View Action Info** to find the relevant line.

If you cannot find the **Install VM Manager Tool** action, see: [Troubleshooting: Enabling the VM Managers panel \(on page ccclxiv\)](#).

2. Run the following command and check whether the VM Manager Tool is working.

```
UNIX BESClient/LMT/VMMAN/vmman.sh -run
```

```
Windows BESClient\LMT\VMMAN\vmman.bat -run
```

3. Test the connection to VM Managers by running the following command.

```
UNIX BESClient/LMT/VMMAN/vmman.sh -testconnection
```

```
Windows BESClient\LMT\VMMAN\vmman.bat -testconnection
```

4. Check the statuses of the VM Managers by running the following command.

```
UNIX BESClient/LMT/VMMAN/vmman.sh -status
```

```
Windows BESClient\LMT\VMMAN\vmman.bat -status
```

For more information, see: [Running VM Manager Tool \(on page ccclxxiv\)](#) and [VM Manager Tool command-line options \(on page ccclxxix\)](#).

Check whether VM Manager Tool results are generated and properly communicated to the License Metric Tool server.

1. Check whether the scan results files are packed for upload.
 - a. Go to **Computer Support Data** panel and collect the logs from the computer that you need to investigate. For more information about how to download the log package, see: [Collecting logs for troubleshooting purposes \(on page dccliv\)](#).
 - b. Go to the following directory.
 - **UNIX** \sha1_files
 \logcollector_0_computer_id_logCollectorAgent.tar
 \computer_id_logCollectorAgent\var\opt\BESClient\LMT\VMAN
 \computer_id
 - **Windows** \sha1_files
 \logcollector_0_computer_id_logCollectorAgent\Program
 Files (x86)\BigFix Enterprise\BES Client\LMT\VMAN
 \computer_id_vmman.zip
2. If the result files are not packed for upload, check whether the upload is scheduled.
 - a. Log in to the BigFix console.
 - b. Click **Actions** in the navigation tree, and select **Schedule VM Manager Tool Scan Results Upload**.
 - c. Check the details and make sure that the state is marked as Open. If not, see: [Uploading collected data \(on page ccclxxi\)](#).
 - d. If the action failed, check which line of the action script caused failure. Go to the **Computer** tab, double-click the listed computer, and analyze the **View Action Info** to find the relevant line. You can also try [forcing the upload of collected data \(on page cccxciii\)](#).
3. Check whether the results exist in the BES server sha1 directory in the log package that you downloaded from the affected computer.
 - **UNIX** /sha1/vmman_scan_*_computer_id*_vmman.tar.gz
 - **Windows** \sha1\vmman_scan_*_computer_id*_vmman.zip

If the **Schedule VM Manager Tool Scan Results Upload** action was completed successfully, but the files are still not uploaded, go to the log package that was downloaded from this computer and in the sha1 directory check the **Index.txt** file. If the file contains the following error `MaxArchiveSize: Exceeded`, see: [Configuring VM manager for subcapacity reporting. \(on page cccxcvii\)](#)



Note: **9.2.14** Starting from application update 9.2.14, you can check whether the value of `_BESClient_ArchiveManager_MaxArchiveSize` parameter is exceeded on the **Computer Support Data** panel. For more information, see: [Checking whether the maximum archive size is exceeded \(on page dccliii\)](#).

4. To check whether the result file is available in the BigFix database, go to the log package that was downloaded from Computer Support Data panel and check whether the `BIGFIX_Uploads` and check whether the `BIGFIX_Uploads` and `BIGFIX_uploads_availability` files contain the data. If you cannot find the results, restart the **FillDB** service on the BigFix server.

- **UNIX** Run the following command: `/etc/init.d/besfilldb restart`.
- **Windows** Find the **FillDB** service on the list of services, and restart it.

It takes several minutes for the database to update after the restart.

- If the issue occurs more than once, upgrade the BigFix server. For more information, see: [Software/Hardware data may not appear in the LMT/BFI UI after successful ETL import because of the BigFix FillDB APAR IV83671](#).
5. Check whether you can find the result data in `DatasourceFile` in the log package that was collected from this computer.
 - If the result file exists in the BigFix database, but it is not found in `DatasourceFile` file, contact HCLsupport.
 6. Check whether you can find the results in `adm.last_imported_scan` file in the [log package \(on page dccliv\)](#) from this computer.

- If the results data exists in all databases (listed in point 5 and 6), but cannot be found in `adm.last_imported_scan` file, investigate why the file import fails.
 - Check the BigFix Inventory import logs, which are located in the following directory.

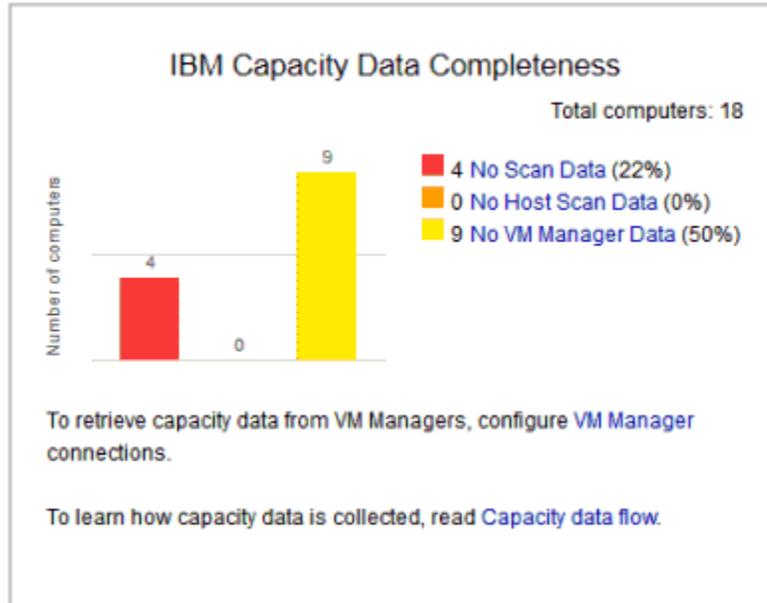
`Installation_directory/wlp/usr/servers/server1/logs/imports`

You can find also the last import log attached to the log package from this computer.

Check whether the VM Manager Tool collects valid and complete results.

Continue with this solution if the result file was imported without errors, and at least one of your VM Manager connections has the  **No VM Manager data** flag.

1. Check whether the results are complete.
 - a. Collect UUIDs of at least a few endpoints that cannot connect due to no VM Manager data.
 - b. Log in to BigFix Inventory, and view the IBM Capacity Data Completeness widget. It shows the number of endpoints with the No VM Manager Data status.



- c. Click the **No VM Manager Data** to see the list of computers with no VM Manager data on the Hardware Inventory report.

Status	Computer Name	IP Address	Partition Cores	Server ID	Server Cores	Processor Brand String
No VM Manager Data	NC040221	9.167.40.221, 192.168.122.1		2_TLM_VM_4234d9ad-380-55a8-a204-34473e535eea	2	Intel(R) Xeon(R) CPU X7560 @ 2.27GHz
No VM Manager Data	NC9128111123	9.128.111.123		2_TLM_VM_4236c85a-d07d-c111-c8cb-b8e27eb4b617	2	Intel(R) Xeon(R) CPU E7- 4870 @ 2.40GHz
No VM Manager Data	NC046182.krakib.pl.ibm.com	9.156.46.182, 192.168.122.1, 172.17.0.1		1_TLM_VM_421f7d1d-e374-da93-73fd-83edff959904	1	Intel(R) Xeon(R) CPU E5620 @ 2.40GHz
No VM Manager Data	NC040191	9.167.40.191		2_TLM_VM_4233b91c-d896-5f3-a0be-b10070b8c38	2	Intel(R) Xeon(R) CPU E7520 @ 1.87GHz

- d. Export the report to CSV file to view the information about the server ID.
- e. Get the result file. The result file is located in the `sha1` directory on the BigFix server. `sha1` directory is also attached to the log package that is collected from the affected computer and downloaded from [Computer Support Data \(on page dccliv\)](#) panel.
- f. If the result file is not in that directory, go to log package and check the `debugData.zip` file. The `debugData.zip` file contains the result files and additional information that you might need to troubleshoot the problem.

- g. Choose a server ID, also referred to as UUID, from the list of customer endpoints, and check the result files against this ID. The UUID does not include the `TLM_VM` phrase. See the example of the result file for a virtual machine.

```

12 <status>OK - duplicated UUIDs discarded</status>
13 <url>https://srvvcspr01/sdk/vimService.wsdl</url>
14 <cluster_layer>
15   <id>CP3/09. ROCKHAMPTON</id>
16   <name>CP3/09. ROCKHAMPTON</name>
17   <host_layer>
18     <uuid>564daf6b-4002-5ac4-7f4b-85703efe392c</uuid>
19     <name>srvhstpr41.gso.internal</name>
20     <cores_count>2</cores_count>
21     <active_sockets_count>1</active_sockets_count>
22     <available_sockets_count>1</available_sockets_count>
23     <cpu_model>Intel(R) Xeon(R) CPU E31220 @ 3.10GHz</cpu_model>
24     <hardware_vendor>VMware, Inc.</hardware_vendor>
25     <hardware_model>VMware Virtual Platform</hardware_model>
26     <connection_state>connected</connection_state>
27     <power_state>powered_on</power_state>
28     <virtual_technology>VSphere</virtual_technology>
29     <guest_layer>
30       <uuid>42362841-6b4e-aa26-9755-07b28dc0fd41</uuid>
31       <vp_count>1</vp_count>
32       <state>default_not_initialized</state>
33       <virtual_technology>VSphere</virtual_technology>
34     </guest_layer>
35   </host_layer>
36 </cluster_layer>

```

2. What to do, when you cannot find the UUID in the result files.

Repeat the actions from step one to ensure the UUID is not in the result files. Use only the final parts of the UUID in search.

The UUID returned by the hardware scan can differ from the UUID returned by the VM Manager Tool. Both values consist of 32 identical characters. However, the first 16 characters are arranged in different order. The final 16 characters are the same. For example, the Hardware Inventory server ID is `TLM_VM_e2bc2742-2329-caeb-2b5c-d1a004b2758a`. However, the result file shows the following value.

```

1189 <guest_layer>
1190   <uuid>4227bce2-2923-ebca-2b5c-d1a004b2758a</uuid>
1191   <vp_count>2</vp_count>
1192   <state>default_not_initialized</state>
1193   <virtual_technology>VSphere</virtual_technology>
1194 </guest_layer>

```

This issue was fixed in version 9.2.3. For permanent resolution upgrade BigFix Inventory. After the upgrade, upload the new capacity data to produce new scan results. To do that, run a single capacity scan and force upload of results. For more information, see: [Initiating the capacity scan on all computers \(on page ccx\)](#).

The capacity data is updated in up to two days, when the capacity scans and the VM Manager data are imported.

3. If the UUID is not in the result files, check the possible causes.

- The UUID is duplicated.
 - a. To verify whether the UUID is duplicated, check all the `trace.log` files that are located in the log package in the following directory.

```

UNIX \sha1_files
\logcollector_0_computer_id_logCollectorAgent.tar
\computer_id_logCollectorAgent\var\opt\BESClient\LMT\VMMAN
\logs

```

```

Windows \sha1_files
\logcollector_0_computer_id_logCollectorAgent\Program Files
(x86)\BigFix Enterprise\BES Client\LMT\VMMAN\logs

```

- b. Search for the following extract, and remove the duplicated UUID.

```

2016-05-30 15:42:40 : (pool-1-thread-1)
com.ibm.license.mgmt.datacollector.tasks.VMCollectorTask::
updateManagerSuccess(): (vmmconf_4.properties):
Duplicates of UUIDs are found on VM Manager,
URL: https://srvvcspr01/sdk/vimService.wsdl.
Discarded UUIDs for guests: 4227e866-7121-d731-235a-343077d7ee93.

```

- The user who is connecting with the VM Manager on VMware has no permission to collect the data about the virtual machine.

Check the VM Manager settings, and permissions. For more information, see: [Verifying permissions for VMware communication \(on page cccxxiv\)](#).

- The VM Manager is not valid for the virtual machine.

4. What to do when the UUID is found in the result files, and the VM Manager connection is still flagged as  No VM Manager data in BigFix Inventory.

- a. To find the reason behind the connection status, check whether the debug data is collected. Open the log package and go to the debug directory.

```

UNIX \sha1_files
\logcollector_0_computer_id_logCollectorAgent.tar
\computer_id_logCollectorAgent\var\opt\BESClient\LMT\VMMAN\debug

Windows \sha1_files\logcollector_0_computer_id_logCollectorAgent
\Program Files (x86)\BigFix Enterprise\BES Client\LMT\VMMAN\debug

```

- b. Open the `login.xml` file.

```

1 <?xml version="1.0" encoding="UTF-8"?>
2 <soapenv:Envelope xmlns:soapenc="http://schemas.xmlsoap.org/soap/encoding/"
3   xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/"
4   xmlns:xsd="http://www.w3.org/2001/XMLSchema"
5   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
6   <soapenv:Body>
7     <LoginResponse xmlns="urn:vim25">
8       <returnval>
9         <key>52d1785e-500f-1596-98c9-cc06729d0082</key>
10        <userName>KRAKLAB.PL.IBM.COM\lmtuser</userName>
11        <fullName>LMT User</fullName>
12        <loginTime>2016-06-07T09:24:50.965995Z</loginTime>
13        <lastActiveTime>2016-06-07T09:24:50.965995Z</lastActiveTime>
14        <locale>en</locale>
15        <messageLocale>en</messageLocale>
16      </returnval>
17    </LoginResponse>
18  </soapenv:Body>
19 </soapenv:Envelope>

```

- c. Check whether the `login.xml` file is correct, and error-free.

The `login.xml` file can contain the authentication error.

```

<faultcode>ServerFaultCode</faultcode>
<faultstring>Cannot complete login due to an incorrect user name
or password.</faultstring>
<detail>
<InvalidLoginFault xmlns="urn:vim25" xsi:type="InvalidLogin"/>

```

If the file contains the authentication error, and the VM Manager connection test was successful, the user login, or password probably contains \$, or \ sign. Upgrade the VM Manager Tool to the latest version to resolve the problem. For more information, see: [Checking the VM Manager Tool version \(on page cccxciv\)](#) and [Updating VM Manager Tool \(on page cccxciv\)](#).



Note: After you upgrade the VM Manager Tool, download the fresh log package and check the debug data.

- d. Open the `retrieveProperties.xml` file from the `debug` directory. The file contains information about all virtual machines, including the duplicate UUIDs, and the `<HostCpuPackage>` data. If the `<HostCpuPackage>` tag is not in the file, the user that connects to the VM Manager is underprivileged. Ensure that the user can collect the endpoints and the host data. Check user's permissions. For more information, see: [Verifying permissions for VMware communication \(on page cccxxiv\)](#) and [How to set the correct permissions for LMT - VMware communication](#).

Troubleshooting: Enabling the VM Managers panel

If the VM Managers panel is disabled and you cannot add VM managers, complete additional steps to enable the panel. This issue often occurs if your BigFix server does not have Web Reports configured, or if the automatic installation of the VM Manager Tool failed. It can also occur when the BigFix user that is used to connect to the BigFix console does not have sufficient rights.

The issue can be recognized by a disabled panel and a message similar to the following example:

```
Before you can configure VM managers, you must install and run the BigFix services,
including Web Reports. For more information, see the product documentation.
```

Required permissions

Before you perform the following steps, ensure that the BigFix user that is used to connect to the BigFix console is a Master Operator. If you do not want to use the Master Operator, you can create a dedicated user that fulfills the following requirements:

- Is assigned the BigFix Inventory v9 site
- Is assigned computers that you are going to monitor, and the computer where the BigFix server is installed
- Has the following permissions: *Can use REST API, Can use Console, Custom Content, Can Create Actions*

The option is supported starting from BigFix 9.5.

Installing Web Reports

The Web Reports component is typically installed together with the BigFix server. You might have chosen, however, to omit this component during the installation. To install it, complete one of the following steps depending on the operating system of the computer, where BigFix is installed.

- For Windows systems, see [Installing Web Reports](#).
- For Linux systems, start the installation of BigFix. Choose the Production installation and when prompted for the installation of components, choose Web Reports only.

```
Select the BigFix Features you want to install:
[1] All Components (Server, Client, and WebReports)
[2] Server and Client Only
[3] WebReports Only
Choose one of the options above or press <Enter> to accept the default: [1]
```

You installed Web Reports on the BigFix server.

[Install VM Manager Tool \(on page ccclxv\)](#).

(Optional) Installing VM Manager Tool

If the automatic installation of VM Manager Tool failed, you can install it manually on the BigFix server.

- [Review software requirements and other considerations \(on page cccxci\)](#)
- Install an BigFix client on the target computer.
- Install and start Web Reports on the BigFix server.

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.

3. Select **Install VM Manager Tool *version number***, and then click **Take Action**.
4. Select the BigFix server as target computer and click **OK**.

Schedule the upload of collected data (on page ccclxvi).

(Optional) Uploading collected data

The uploading of collected data is scheduled automatically. You can use this procedure if the automatic installation of VM Manager Tool failed.

1. Log in to BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. Select **Schedule VM Manager Tool Scan Results Upload**.

Fixlets and Tasks			
Name	Source Sev...	Applicab...	Category
Upload Scanner Diagnostic Data	Low	0 / 5	Troubleshooting
Install Additional VM Manager Tool	Low	2 / 5	VM Managers
Install VM Manager Tool	Low	1 / 5	VM Managers
Schedule VM Manager Tool Scan Results Upload	Low	2 / 5	VM Managers
Uninstall VM Manager Tool	Low	2 / 5	VM Managers
Update VM Manager Tool	Low	0 / 5	VM Managers

4. In Description, enter the upload frequency in hours.

Description

This task schedules regular uploads of capacity data collected by VM Manager Tool or by the following task: [Run Capacity Scan on Virtualization Hosts](#). The data is uploaded to IBM BigFix and can be transferred to IBM BigFix Inventory by running a data import.

Upload data every hours.

5. Select the BigFix server as the target computer and click **OK**.

The screenshot shows the 'Target' tab in the BigFix console. It has several sub-tabs: Execution, Users, Messages, Offer, Post-Action, Applicability, Success Criteria, and Action Script. Under the 'Target' section, there are three radio buttons: 'Select devices' (selected), 'Dynamically target by property', and 'Enter device names'. Below this is a table titled 'Applicable Computers (1)'. The table has four columns: Computer Name, OS, CPU, and Last Report Ti... (truncated). The first row contains the following data: NC9143126012, Linux Red Hat ..., 2400 MHz Xeon, and 2014-06-03 10:...

Computer Name	OS	CPU	Last Report Ti...
NC9143126012	Linux Red Hat ...	2400 MHz Xeon	2014-06-03 10:...



Note: All options in the Execution tab are disabled. The frequency can be specified only in the Description field.

You scheduled to upload capacity data from VM Manager Tool to BigFix. The data is transferred to BigFix Inventory with each data import. If no modifications are made to VM managers, the data is uploaded according to the schedule. If you modify a VM manager, the upload is triggered shortly afterward.

[Add VM managers \(on page cccl\)](#) in the user interface.

Advanced configuration of basic VM management

When you use the basic VM management, you can perform advanced configuration on the Advanced Server Settings panel. However, it is not usually needed.

The basic VM management does not usually require reconfiguring. However, you can change the configuration parameters on the Advanced Server Settings panel. For more information, see: [Advanced server settings \(on page cdii\)](#). Any customization can significantly impact the server and its performance. Thus, ensure that every configuration is carefully considered and needed.

Advanced VM management (distributed)

Use advanced VM management if the BigFix server cannot connect to some of the VM managers because they are in separated networks. This approach requires that you install an additional VM Manager Tool, and configure VM managers through the command line.

You can use advanced VM management as an extension of the basic configuration. VM managers to which the BigFix server can connect can still be managed in the BigFix Inventory user interface. Additional VM Manager Tool would manage only VM managers for which such a connection cannot be established. You can connect to such VM managers by creating configuration files and managing the connections in the VM Manager Tool command line.

Although VM managers are managed through the command line, their capacity data is uploaded to BigFix Inventory and updated on reports in the same way as for the basic configuration. In contradiction to the main VM Manager Tool, additional instances can be installed on any computer that can connect to separated VM managers, and you can deploy as many of them as you need. VM managers configured in the command line are visible on the [VM Managers](#)

[panel \(on page cccxlvi\)](#) in the BigFix Inventory user interface. However, their configuration cannot be changed through that panel. It can be change only through the VM manager configuration files.

The advanced VM management also gives you the following benefits:

- Balancing the network traffic by distributing it between multiple VM Manager Tools
- Collecting capacity data only for selected virtual machines (UUID-based filtering)

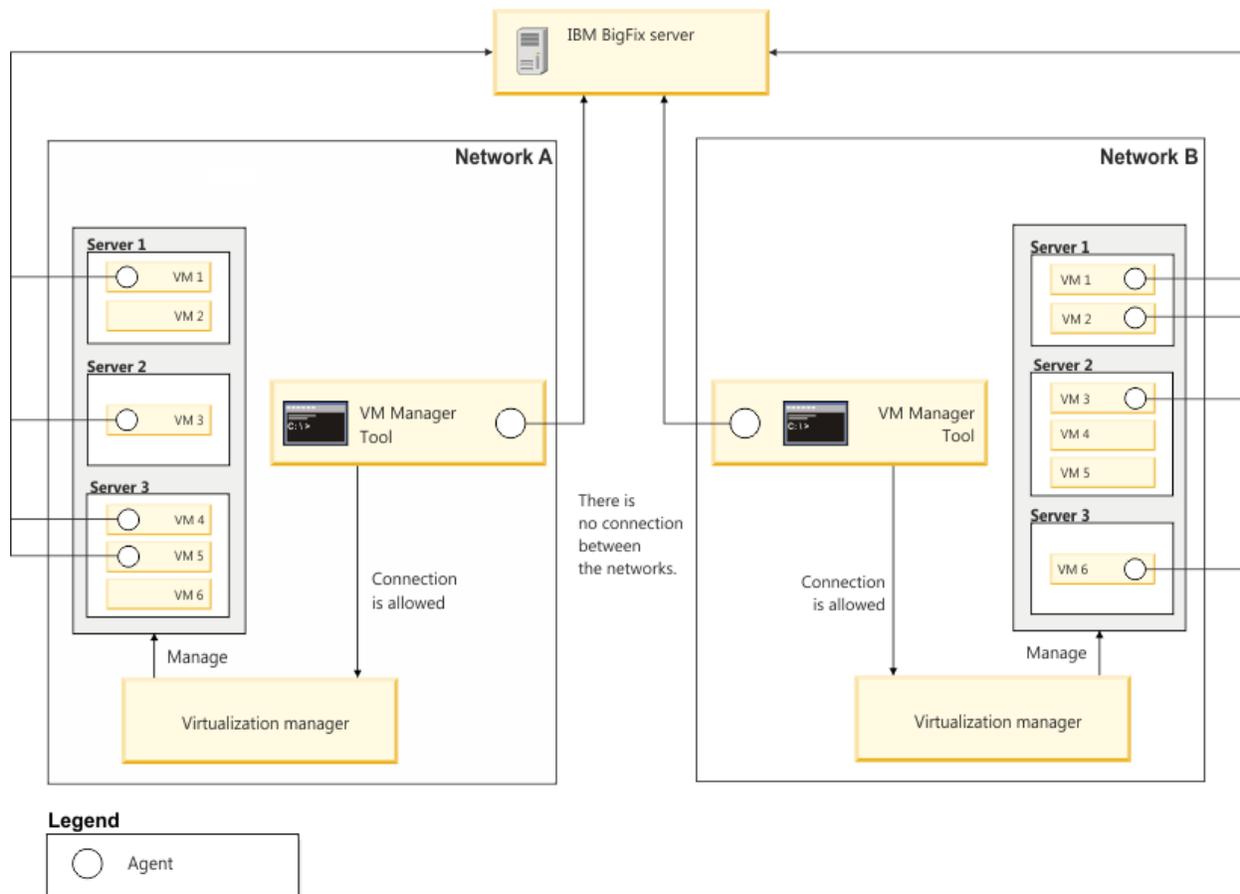
Deployment scenarios

Before you proceed to advanced VM management, review the available scenarios for which you can use the additional instance of VM Manager Tool. The main purpose is to transfer the capacity data from VM managers that are in separated networks, but you can also balance the network traffic, or filter your virtual machines based on their UUIDs.

Scenario 1: A VM manager is in a separated network and cannot connect to BigFix

Without a direct connection to BigFix, data about capacity of VM Managers cannot be collected, which conflicts with license compliance. If you have a separated network in your environment, you can use VM Manager Tool that serves as a connector between VM Managers and the server. The tool can be installed on any computer that can connect to VM Managers. It gathers data about their capacity and then uploads it to the server.

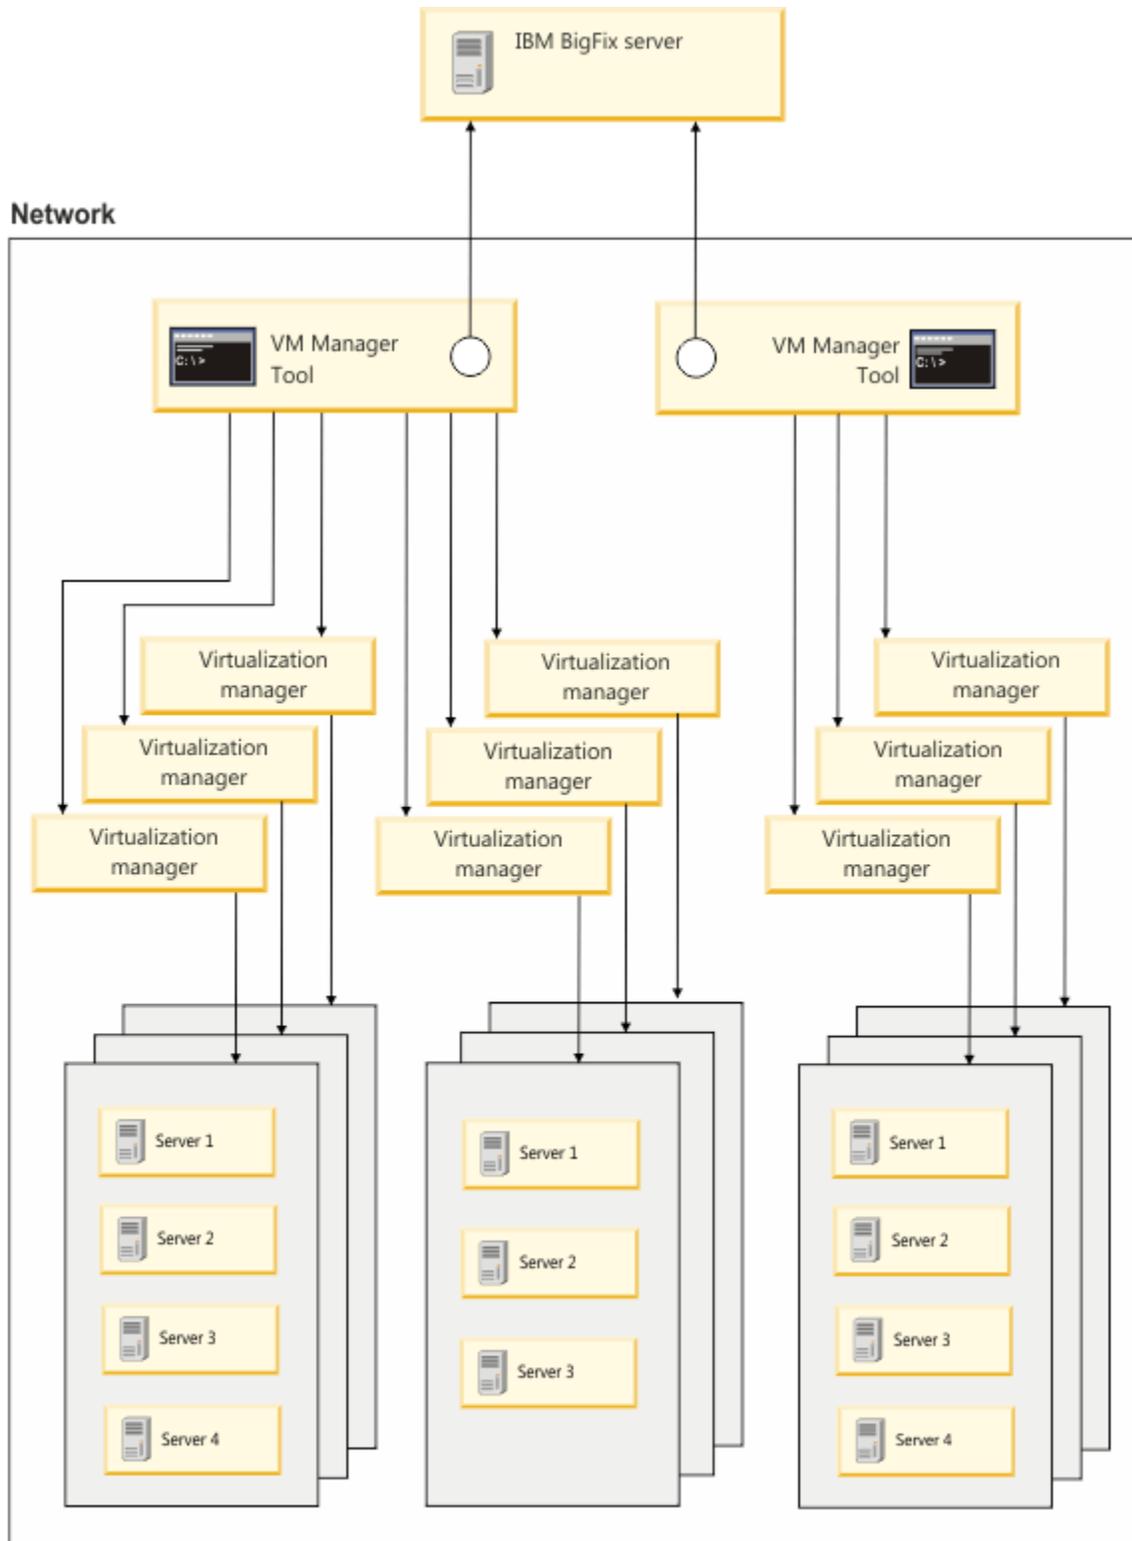
The following image shows two separated networks, each represented by a different VM Manager Tool.



Scenario 2: You want to balance the network traffic coming from multiple VM Managers

You can balance the network traffic by distributing it between multiple VM Manager Tools. Successful load balancing optimizes resource use, maximizes throughput, minimizes response time, and avoids overload.

The following image shows two VM Manager Tools connected to the BigFix server. Each of them is responsible for gathering capacity data from different VM Managers.



Legend



Scenario 3: UUID-based virtual machine filtering

The purpose of filtering based on universally unique identifier (UUID) is to select only those virtual machines whose capacity and topology data is to be included in a final report. The VM manager has the function of matching UUIDs of selected guests. To stay compliant, all the necessary virtual machines units must be included in the report. When you enable UUID-based filtering all of the non-selected guests, empty hosts and clusters are removed.

Installing additional VM Manager Tool

To connect to VM managers that are in separated networks, you must first install additional VM Manager Tool. You can install it on any computer that can connect to the separated networks.

- [Review software requirements and other considerations \(on page cccxci\)](#)
 - Install an BigFix client on the target computer.
1. Log in to the BigFix console.
 2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
 3. Select **Install Additional VM Manager Tool (OPTIONAL) version number** and click **Take Action**.
 4. Select the target computer and click **OK**. The target computer must be able to connect to separated VM managers.

You installed additional instance of VM Manager Tool. You can use its command-line interface to manage VM managers that are in separated networks.

[Schedule the upload of collected data \(on page ccclxxi\)](#).

Uploading collected data

After you install additional VM Manager Tool, schedule the regular uploads of collected data. This action is required so that the capacity data that is collected by VM Manager Tool is uploaded to the BigFix server.

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. Select **Schedule VM Manager Tool Scan Results Upload**.

Fixlets and Tasks			
Name	Source Sev...	Applicab...	Category
Upload Scanner Diagnostic Data	Low	0 / 5	Troubleshooting
Install Additional VM Manager Tool	Low	2 / 5	VM Managers
Install VM Manager Tool	Low	1 / 5	VM Managers
Schedule VM Manager Tool Scan Results Upload	Low	2 / 5	VM Managers
Uninstall VM Manager Tool	Low	2 / 5	VM Managers
Update VM Manager Tool	Low	0 / 5	VM Managers

4. In Description, enter the upload frequency in hours.

Description
<p>This task schedules regular uploads of capacity data collected by VM Manager Tool or by the following task: Run Capacity Scan on Virtualization Hosts. The data is uploaded to IBM BigFix and can be transferred to IBM BigFix Inventory by running a data import.</p> <p>Upload data every <input type="text" value="12"/> hours.</p>

5. Click **Take Action**.

6. Select the target computer and click **OK**. The computer must have VM Manager Tool installed.

Target	Execution	Users	Messages	Offer	Post-Action	Applicability	Success Criteria	Action Script										
<p>Target:</p> <p><input checked="" type="radio"/> Select devices <input type="radio"/> Dynamically target by property <input type="radio"/> Enter device names</p> <table border="1"> <thead> <tr> <th>Applicable Computers (1)</th> <th>Computer Name</th> <th>OS</th> <th>CPU</th> <th>Last Report Ti...</th> </tr> </thead> <tbody> <tr> <td></td> <td>NC9143126012</td> <td>Linux Red Hat ...</td> <td>2400 MHz Xeon</td> <td>2014-06-03 10:...</td> </tr> </tbody> </table>									Applicable Computers (1)	Computer Name	OS	CPU	Last Report Ti...		NC9143126012	Linux Red Hat ...	2400 MHz Xeon	2014-06-03 10:...
Applicable Computers (1)	Computer Name	OS	CPU	Last Report Ti...														
	NC9143126012	Linux Red Hat ...	2400 MHz Xeon	2014-06-03 10:...														



Note: All options in the Execution tab are disabled. The frequency can be specified only in the Description field.

You scheduled the upload of capacity data from VM Manager Tool to BigFix. The data is transferred to BigFix Inventory with each data import. If no modifications are made to VM managers, the data is uploaded according to the schedule. If you modify a VM manager, the upload is triggered shortly afterward.

[Configure VM managers \(on page ccclxxii\)](#).

Adding VM managers in distributed mode

In advanced VM management, you add VM managers by creating configuration files, where you specify web address, type, and credentials that are required to access a VM manager. Each VM manager uses a separate configuration file.

VM managers for which configuration files are created are displayed in the BigFix Inventory user interface. However, they can be managed only in the VM Manager Tool command-line interface.

9.2.7 VM Manager Tool, starting from update 1.5.0.0, is enhanced to attempt to automatically fix the most common problems with the VM manager connection parameters. This feature is disabled by default. To enable this feature, go to the `BES Client\LMT\VMAN\config` directory, open the `vmmmainconf.properties` file and change the setting of the `vmm_fix_parameters_enabled` parameter to `true`. While fixing connection problems the VM Manager Tool can connect using the HTTP protocol when the defined HTTPS connection is not available. This setting is enabled by default. To disable switching protocol from HTTPS to HTTP, go to the `BES Client\LMT\VMAN\config` directory, open the `vmmmainconf.properties` file and change the setting of the

`vmm_http_connection_allowed` parameter to `false`. For more information, see: [VM Manager Tool settings \(on page ccclxxxii\)](#).

1. Go to the `BES Client\LMT\VMMAN\config` directory.
2. Copy the `vmmconf_template.properties` file and rename it to `vmmconf_name.properties`. This file now represents a new VM manager.

! **Attention:** Each time you create a VM manager configuration file, use the `vmmconf_template.properties` file as a template. Do not copy and edit configuration files that you previously created for a different VM manager. After you load the configuration file for the first time, an ID is generated for a VM manager. Each ID must be unique. When you copy and edit an existing configuration file, the ID is duplicated.

3. Edit the file and specify the following parameters:

vmm_url

Specify the web address of the VM manager. You can specify either a full URL or only the host name or IP address.

For example, `vmm_url=http://192.0.2.0/wsman`.

! **Important:** The default URL differs depending on the virtualization type.

- vCenter
 - For versions up to 5.1: `https://<vCenter_IP_address>/sdk`
 - For version 5.5 and higher: `https://<vCenter_IP_address>/sdk/vimService.wsdl`
- RHEV-M
 - For version 3.0: `https://<RHEV-M_IP_address>:8443/api`
 - For version 3.1 and higher: `https://<RHEV-M_IP_address>/api`
 - For version 4.0: `https://<RHEV-M_IP_address>/ovirt-engine/api`
- Hyper-V
 - `https://<Hyper-V_IP_address>/wsman`
- **9.2.12** XenServer or Citrix XenServer
 - `https://<xen-server>`
- **9.2.14** Oracle VM Server for x86
 - `https://<OVM-Manager_IP_address>:7002/ovm/core/wsapi/rest`
- **9.2.17** Nutanix
 - `https://<NUTANIX_PRISM_IP_ADDRESS>:9440/PrismGateway/services/rest/v2.0`

vmm_type

Specify the type of the VM manager. The possible values are `VMWARE_V_SPHERE`, `MICROSOFT_HYPER_V`, `KVM_RHEV_M`, `XEN_MANAGER`, `ORACLE_VM` or `NUTANIX`.

For example, `vmm_type=MICROSOFT_HYPER_V`.

vmm_login

Specify the user name that is used to access the VM manager.

For example, `vmm_login=administrator\cluster.com`.



Important: VM managers use different definitions of users:

- For Microsoft Hyper-V, you must use the Administrator account. The user is defined as `user_name\domain` or `user_name@domain`. For example: `test\cluster.com` or `test@cluster.com`.
- For VMware, the user is defined as `domain\user_name`, for example: `cluster.com\test`.
- For RHEV-M, the user is defined as `user_name@domain`, for example: `test@cluster.com`.
- **9.2.12** For XenServer and Citrix XenServer, the user is defined as `user_name`, for example `root`.
- **9.2.14** For Oracle VM Server for x86, the user is defined as `user_name`, for example: `test`.
- **9.2.17** For Nutanix, the user is defined as `user_name`, for example: `test`.

vmm_password

Specify the password that is used to access the VM manager. The password is encrypted and saved when you load the configuration files.

For more information about optional parameters, see [Configuration parameters \(on page ccclxxvi\)](#).

4. Save the configuration file. To establish the connection, you must run the VM Manager Tool and load the configuration file.

[Run the VM Manager Tool \(on page ccclxxiv\)](#).

Running VM Manager Tool

After creating configuration files for your VM managers, you must run the VM Manager Tool to load the configuration, test connections, and check the status of VM managers.

For a complete list of commands that can be used in VM Manager Tool, see [VM Manager Tool command-line options \(on page ccclxxix\)](#).

1. Open the command-line interface and go to *BES Client\LMT\VMAN*.
2. **Optional:** VM Manager Tool is started automatically after installation, however you can use the following commands to start it, if needed.

Linux `./vmman.sh -run`

Windows `vmman.bat -run`

3. Load the new or updated configuration files to update the connection parameters in VM Manager Tool. During this step, the passwords provided in the files are encrypted. You must repeat this step after modifying or creating a configuration file.

Linux `./vmman.sh -reloadconfig`

Windows `vmman.bat -reloadconfig`

4. Test connections to VM managers.

Linux `./vmman.sh -testconnection`

Windows `vmman.bat -testconnection`

```
c:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\VMAN>vmman.bat -testconnection
CODUI00461: [ vmnconf_1.properties ] Test connection operation has been successfully completed.
CODUI00461: [ vmnconf_2.properties ] Test connection operation has been successfully completed.
```

5. Check the status of VM managers.

Linux `./vmman.sh -status`

Windows `vmman.bat -status`

 **Tip:** For more information about the available statuses, see [VM manager statuses \(on page ccclii\)](#).

```
c:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\VMAN>vmman.bat -status

Virtual manager file: vmnconf_1.properties
Status: OK
URL: http://-.---.--.203/usman
User name: administrator@cluster.com
Last successful scan date: 2013-12-18 14:58:10 CET
Last scan attempt date: 2013-12-18 14:58:10 CET
Last data transfer date: 2013-12-18 14:58:10 CET
Comment: There are no problems with the UM manager.

Virtual manager file: vmnconf_2.properties
Status: OK
URL: http://-.---.--.146:8080/api
User name: admin@internal
Last successful scan date: 2013-12-18 14:58:10 CET
Last scan attempt date: 2013-12-18 14:58:10 CET
Last data transfer date: 2013-12-18 14:58:10 CET
Comment: There are no problems with the UM manager.
```

If the VM managers report the **OK** status, they are working properly, and the capacity data is being collected. The upload of collected data is triggered shortly after making any modifications to VM managers, such as adding a

new connection or changing the existing one. If no modifications are made, the data is uploaded according to the schedule. For more information, see [Capacity data flow \(on page cccxlv\)](#).

9.2.4

Starting from application update 9.2.4, VM managers from the distributed mode are also visible in the user interface after the data about them is imported to the BigFix Inventory server. However, they can be modified only through configuration files.

Additional configuration

Additional configuration tasks are intended to give you more control over VM Manager Tool. You can use them to customize the VM Manager Tool, or to enable filtering of virtual machines based on their UUIDs. All of these tasks are optional.

VM manager configuration parameters

In advanced VM management, connection to each VM manager is created based on the information provided in the configuration files. Apart from specifying mandatory parameters, such as web address, virtualization type, user name, and password, you can also use additional ones to disable VM managers, set the number of login attempts, or allow Hyper-V to share credentials with hosts in the same cluster.



Important: The following settings cannot be changed for the main instance of VM Manager Tool that is installed on the BigFix server. Any changes are overwritten during the next action performed on VM managers. To configure VM managers on the main instance of VM Manager Tool, see: [Adding VM managers in central mode \(on page cccl\)](#).

Table 39. VM manager configuration parameters

The table describes global configuration parameters. For each parameter, there are 2 levels. The first level consists of four columns and provides the unit as well as the default, minimum, and maximum values of the parameter. The second level is common for all four columns and provides a description of the parameter.

Parameter	Unit	Default	Minimum	Maximum
vmm_url	Web address (URL)			
Specifies the web address of the VM manager. You can provide either a full URL, a partial URL, or only a host name or IP address. In the second case, the full address of the VM manager is built based on the selected type of the VM manager and protocol (if specified). HTTPS protocol is used by default.				
If you specify only the addresses, the defaults are used:				
<ul style="list-style-type: none"> • https://virtualcenter/sdk for VMWARE_V_SPHERE • https://hyper-v/wsman for MICROSOFT_HYPER_V 				

Table 39. VM manager configuration parameters

The table describes global configuration parameters. For each parameter, there are 2 levels. The first level consists of four columns and provides the unit as well as the default, minimum, and maximum values of the parameter. The second level is common for all four columns and provides a description of the parameter.

(continued)

Parameter	Unit	Default	Minimum	Maximum	
		Description <ul style="list-style-type: none"> • <code>https://rhev-m:8443/api</code> for <code>KVM_RHEV_M</code> • 9.2.12 <code>https://xen-server</code> for <code>XEN_MANAGER</code> • 9.2.14 <code>https://ovm-m:7002/ovm/core/wsapi/rest</code> for <code>ORACLE_VM</code> 			

If you do not specify the complete URL but only a protocol, or a port, or a context path, the URL is built based on the following defaults:

- `VMWARE_V_SPHERE` - default protocol `https`, port `443` (for `https`) or `80` (for `http`), context path - `sdk`
- `KVM_RHEV_M` - default protocol `https`, port `8443` (for `https`) or `8080` (for `http`), context path - `api`
- `MICROSOFT_HYPER_V` - default protocol `https`, port `443` (for `https`) or `80` (for `http`), context path - `wsman`
- **9.2.12** `XEN_MANAGER` - default protocol `https`, port `443` (for `https`) or `80` (for `http`), context path - `<empty>`
- **9.2.14** `ORACLE_VM` - default protocol `https`, port `7002` (for `https`) or `80` (for `http`), context path - `/ovm/core/wsapi/rest`

If the URL contains the name of a VM manager, the name is resolved to an IP address. However, the full URL, including a port number, is used by the server to identify the VM manager.

Each VM manager must have a different web address, that is, only one entry is allowed for a particular URL. If two or more configuration files duplicate the URL address, only the first file is treated as valid. The remaining files are ignored.

vmm_type

Characters

Specifies the type of VM manager. The possible values are:

- `VMWARE_V_SPHERE`
- `MICROSOFT_HYPER_V`
- `KVM_RHEV_M`

Table 39. VM manager configuration parameters

The table describes global configuration parameters. For each parameter, there are 2 levels. The first level consists of four columns and provides the unit as well as the default, minimum, and maximum values of the parameter. The second level is common for all four columns and provides a description of the parameter.

(continued)

Parameter	Unit	Default	Minimum	Maximum	Description
					<ul style="list-style-type: none"> • 9.2.12 XEN_MANAGER • 9.2.14 ORACLE_VM
vmm_communication_interface	Characters	POWERSHELL			<p>Specifies the communication interface that is used for communication with Microsoft™ Hyper-V. This option is supported only when the VM Manager Tool is installed on Windows. On Linux, the value of this parameter is ignored and the NTLM communication is always used. The possible values are:</p> <ul style="list-style-type: none"> • NTLM • POWERSHELL
vmm_login	Characters				<p>Specifies the user name that is used to access the VM manager.</p> <p> Important: For Microsoft™ Hyper-V, you must use the Administrator account.</p>
vmm_password	Characters				<p>Specifies the password that is used to access the VM manager. A password that is entered in plain text is immediately encrypted and saved while the configuration files are loaded.</p> <p> Important: For Microsoft™ Hyper-V, you must use the Administrator account.</p>
vmm_communication_locked	True/false	False			<p>Indicates whether the connection to the VM manager is locked. The possible values are:</p> <p>true</p> <p>Disables the connection to the VM manager but keeps the configuration file. If the number of failed logging attempts that is specified in the vmm_max_subsequent_login_failures parameter is exceeded, the value</p>

Table 39. VM manager configuration parameters

The table describes global configuration parameters. For each parameter, there are 2 levels. The first level consists of four columns and provides the unit as well as the default, minimum, and maximum values of the parameter. The second level is common for all four columns and provides a description of the parameter.

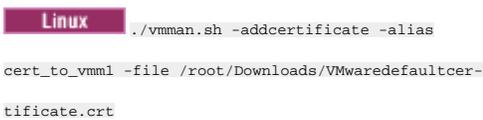
(continued)

Parameter	Unit	Default	Minimum	Maximum
<code>vmm_get_cluster_info_with_shared_credentials</code>	True/false	False		

VM Manager Tool command-line options

Apart from basic commands that allow you to create connections to VM managers and check their status, you can use additional commands to run or remove VM Manager Tool, encrypt passwords, or collect debug information for troubleshooting purposes.

The following table provides a list of command-line options that you can use when you run the VM Manager Tool. If you run the application without specifying any options, the help screen is displayed by default.

Option	Description	Example
<code>-addcertificate -alias unique_alias -file vm_manager_certificate</code>	Adds a certificate to the keystore. Use this option when the VM Manager Tool is set not to trust SSL certificates from all defined VM managers (<code>vmmman_trust_all_vm_managers_certificates</code> (on page cdii) parameter is set to <code>false</code>).	

Option	Description	Example
<p>9.2.9</p> <p><code>-changepassword -file file_path</code></p>	<p>Changes the default password to the VM Manager Tool keystore. For more information, see: Improving security of storing VM manager passwords (on page dccx).</p>	<p>Windows</p> <pre>vmman.bat -addcertificate -alias cert_to_vmml -file C:\Users\Administrator\Downloads\VMwaredefaultcertificate.crt</pre>
<p><code>-help</code></p>	<p>Displays the help screen. It is the default option when no other option is specified.</p>	<p>Linux</p> <pre>./vmman.sh -changepassword -file /var/opt/BESClient/LMT/VMMAN/config/keystore_password.txt</pre>
<p><code>-install</code></p>	<p>Installs VM Manager Tool as a system service.</p> <p> Important: You must have the administrative or root privileges to use this option.</p>	<p>Windows</p> <pre>vmman.bat -changepassword -file "C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\config\keystore_password.txt"</pre>
<p><code>-passwd passwordString -config file_path</code></p>	<p>Encrypts a password for the VM manager from a configuration file that is specified in the config parameter.</p>	<p>Linux</p> <pre>./vmman.sh -install</pre>
<p>9.2.9</p> <p><code>-regenerateencryptionkey</code></p>	<p>Overwrites the default key that is used to encrypt passwords to VM managers. Run the command to increase security of storing the passwords. For more information, see: Improving security of storing VM manager passwords (on page dccx).</p>	<p>Windows</p> <pre>vmman.bat -install</pre>
<p><code>-reloadconfig</code></p>	<p>Reloads all configuration files and updates the parameters in the memory of VM Manager Tool. You can use this option only if VM Manager Tool runs as a system service.</p>	<p>Windows</p> <p>The installation logs are in the directory: <i>BES Client\LMT\VMMAN\logs\install</i>.</p>
<p><code>-passwd passwordString -config file_path</code></p>	<p>Encrypts a password for the VM manager from a configuration file that is specified in the config parameter.</p>	<p>Linux</p> <pre>./vmman.sh -passwd newPassword -config ./config/vmmconf_1.properties</pre>
<p>Windows</p> <pre>vmman.bat -passwd newPassword -config config\vmmconf_1.properties</pre>	<p>Linux</p> <pre>./vmman.sh -regenerateencryptionkey</pre>	<p>Windows</p> <pre>vmman.bat -regenerateencryptionkey</pre>
<p>9.2.9</p> <p><code>-regenerateencryptionkey</code></p>	<p>Overwrites the default key that is used to encrypt passwords to VM managers. Run the command to increase security of storing the passwords. For more information, see: Improving security of storing VM manager passwords (on page dccx).</p>	<p>Linux</p> <pre>./vmman.sh -regenerateencryptionkey</pre>
<p>Windows</p> <pre>vmman.bat -regenerateencryptionkey</pre>	<p>Reloads all configuration files and updates the parameters in the memory of VM Manager Tool. You can use this option only if VM Manager Tool runs as a system service.</p>	<p>Windows</p> <pre>vmman.bat -regenerateencryptionkey</pre>
<p><code>-reloadconfig</code></p>	<p>Reloads all configuration files and updates the parameters in the memory of VM Manager Tool. You can use this option only if VM Manager Tool runs as a system service.</p>	<p>Linux</p> <pre>./vmman.sh -reloadconfig</pre>
<p>Windows</p> <pre>vmman.bat -reloadconfig</pre>		<p>Windows</p> <pre>vmman.bat -reloadconfig</pre>

Option	Description	Example
<code>-remove</code>	<p> Tip: Use this option each time a new VM manager connection is defined to load the newly created configuration.</p> <p>Removes VM Manager Tool from the service registry.</p>	<p>Linux <code>./vmman.sh -remove</code></p>
	<p> Important: You must have the administrative or root privileges to use this option.</p> <p>Windows The removal logs are in the directory: <code>BES Client\LMT\VMMAN\logs\install</code>.</p>	<p>Windows <code>vmman.bat -remove</code></p>
<code>-retrievebugdata</code>	<p>Collects debug information from all defined VM managers and stores it in the <code>debugData.zip</code> file that is in the main the VM Manager Tool directory. The collected information includes:</p> <ul style="list-style-type: none"> • Configuration files • Log files • Network communication log files • Data collected from VM managers • Status of all VM managers 	<p>Linux <code>./vmman.sh -retrievebugdata</code></p> <p>Windows <code>vmman.bat -retrievebugdata</code></p>
<code>-run</code>	<p>Starts the VM Manager Tool in service mode. In this mode, data is collected for all defined VM managers repeatedly at the interval that is set in the <code>vmm_polling_time_interval</code> (on page ccclxxxii) parameter. To use this option, you must install VM Manager Tool as a system service.</p> <p> Important: Windows You must have the administrative privileges to use this option.</p>	<p>Linux <code>./vmman.sh -run</code></p> <p>Windows <code>vmman.bat -run</code></p>
<code>-runonce [-config file_path]</code>	<p>Collects data from all defined VM managers once and exits the VM Manager Tool. To collect data from a particular VM manager, use the <code>-config file_path</code> option, where <code>file_path</code> is a full or relative path to the configuration file of the VM manager that you want to collect data from.</p>	<p>Linux <code>./vmman.sh -runonce -config ./config/vmmconf_1.properties</code></p> <p>Windows <code>vmman.bat -runonce -config config\vmconf_1.properties</code></p>

Option	Description	Example
<code>-status [-config file_path]</code>	<p>Displays the operation status for all VM managers. If the file path is specified, the operation status for the particular VM manager is displayed.</p> <p>To see the information about all the operation statuses, refer to the topic VM Manager Tool statuses (on page ccclij).</p>	<p>Linux <code>./vmman.sh -status -config ./config/vmmconf_1.properties</code></p> <p>Windows <code>vmman.bat -status -config config\vmmanconf_1.properties</code></p>
<code>-stop</code>	Stops VM Manager Tool that was started as a system service.	<p>Linux <code>./vmman.sh -stop</code></p> <p>Windows <code>vmman.bat -stop</code></p>
<code>-testconnection [-config file_path]</code>	<p>Tests connections to all defined VM managers. To collect data from a particular VM manager, use the <code>-config file_path</code> option, where <code>file_path</code> is a full or relative path to the configuration file of the VM manager that you want to collect data from.</p>	<p>Linux <code>./vmman.sh -testconnection -config ./config/vmmconf_1.properties</code></p> <p>Windows <code>vmman.bat -testconnection -config config\vmmanconf_1.properties</code></p>

VM Manager Tool settings

You can configure settings of VM Manager Tool by editing the `vmmmainconf.properties` file that is stored in the `BES Client\LMT\VMAN\config` directory.



Important: The following settings cannot be changed for the main instance of VM Manager Tool that is installed on the BigFix server. Any changes are overwritten during the next action performed on VM managers. To configure VM managers on the main instance of VM Manager Tool, see: [Adding VM managers in central mode \(on page cccl\)](#).

Table 40. Global configuration parameters

The table describes global configuration parameters. The first column includes names and definitions of parameters. Next four columns provide the unit as well as the default, minimum, and maximum values of the parameter.

Parameter	Unit	Default	Minimum	Maximum
<code>check_vm_managers_uniqueness</code>	Boolean (true/false)	True		
Distinguishes unique VM managers from duplicates.				
<code>uuid_filtering_enabled</code>	Boolean (true/false)	False		
Enables UUID filtering.				

Table 40. Global configuration parameters

The table describes global configuration parameters. The first column includes names and definitions of parameters. Next four columns provide the unit as well as the default, minimum, and maximum values of the parameter.

(continued)

Parameter	Unit	Default	Minimum	Maximum
9.2.2 <code>vmm_collecting_hostnames_enabled</code>	Boolean (true/false)	False		
<p>Specifies whether host names should be collected and saved in the log files. Available for VMware and Oracle VM Server for x86 only.</p> <p>You can enable this parameter to solve problems with duplicated UUIDs. In the case of found duplicates, the host names are written in the log files.</p>				
9.2.3 <code>vmm_collecting_host_serials_enabled</code>	Boolean (true/false)	False		
<p>Specifies whether serial numbers of VM manager hosts are collected for VM managers from which data is collected by the VM Manager Tool. Available for Hyper-V, VMware 5.0 and higher, and Oracle VM Server for x86.</p> <p>To display information about serial numbers in the Server ID column on the reports, adjust the Server ID template that is specified by the <code>managedServerTagTemplate</code> parameter. For more information, see: Advanced server settings (on page cdii).</p>				
<code>vmm_connection_timeout</code>	Seconds	90	10	3600 (1 hour)
<p>Specifies the time after which the connection with a VM manager is ended.</p>				
<code>vmm_data_transfer_period</code>	Minutes	720	0	10080
<p>Determines how often the scan data is transferred to the agent to be uploaded to the server if subsequent scans have the same results.</p>				
9.2.7 <code>vmm_fix_parameters_enabled</code>	Boolean (true/false)	False		
<p>Attempts to fix the most common problems with VM manager connection parameters, such as adding default</p>				

Table 40. Global configuration parameters

The table describes global configuration parameters. The first column includes names and definitions of parameters. Next four columns provide the unit as well as the default, minimum, and maximum values of the parameter.

(continued)

Parameter	Unit	Default	Minimum	Maximum
<p>URL suffix, correcting the user name format, and testing the WinRM and PowerShell protocols for Hyper-V VM managers.</p> <p>9.2.7 <code>vmm_http_connection_allowed</code></p> <p>Boolean (true/false) True</p>				
<p>Allows the VM manager to connect by using the HTTP protocol when the defined HTTPS connection is not available.</p> <p>9.2.9 <code>vmm_keystore_password_do_not_change_it</code></p> <p>Contains the encrypted password to the VM Manager Tool keystore. The parameter is specified only if you changed the default password by using the <code>regeneratekeystorepassword</code> command. For more information, see: Improving security of storing VM manager passwords (on page dccxl).</p> <p>Important: Do not change the value of this parameter.</p>				
<p><code>vmm_max_subsequent_login_failures</code></p> <p>Specifies the maximum number of failed attempts of logging in to the VM manager.</p>	Number	3	0	100
<p><code>vmm_pooling_time_interval</code></p> <p>Specifies the interval between the consecutive retrievals of data from VM managers.</p> <p>Restriction: In case of subcapacity reporting, it is necessary to capture mobility of VMs. To ensure that it is captured, data from VM managers must be collected every 30 minutes. You can change</p>	Minutes	30	30	10080 (1 week)

Table 40. Global configuration parameters

The table describes global configuration parameters. The first column includes names and definitions of parameters. Next four columns provide the unit as well as the default, minimum, and maximum values of the parameter.

(continued)

Parameter	Unit	Default	Minimum	Maximum
 the frequency of collecting data only if you use BigFix Inventory for purposes other than subcapacity reporting.				
9.2.8 <code>vmm_rmi_protocol</code>	String	TLSv1.2		
Specifies the protocol that is used for internal communication through the port that is specified in the <code>vmm_rmi_ssl_port</code> parameter when you run the <code>reloadconfig</code> or <code>stop</code> commands.				
Possible values are:				
<ul style="list-style-type: none"> • SSL • SSLv3 • TLSv1 • TLSv1.1 • TLSv1.2 				
<code>vmm_rmi_ssl_port</code>	Number	25001		
Specifies the port that is used for internal communication when you run the <code>reloadconfig</code> or <code>stop</code> commands.				
9.2.7 <code>vmm_slm_tags_file_size_limit</code>	Bytes	10240	10240	1048576
The maximum size of each <code>.slmtag</code> file that is used to collect VMware license metric utilization.				
9.2.7 <code>vmm_slm_tags_files_per_software</code>	Number	10	1	100
The maximum number of <code>.slmtag</code> files that are stored on the computer for each license key under VMware instance.				
9.2.7 <code>vmm_slm_tags_enabled</code>	Boolean (true/false)	False		

Table 40. Global configuration parameters

The table describes global configuration parameters. The first column includes names and definitions of parameters. Next four columns provide the unit as well as the default, minimum, and maximum values of the parameter.

(continued)

Parameter	Unit	Default	Minimum	Maximum
Enables collection of license metric utilization for VMware VM managers. If set to true, <code>.slmtag</code> files are created locally and collected by the resource utilization scan.				
vmm_thread_pool_size	Number	10	1	50
Specifies the number of threads in thread pool that is used for connections to VM managers.				
9.2.4 vmm_transfer_credentials_to_server	Boolean (true/false)	True		
Specifies whether VM manager credentials are transferred to the server. The option is enabled by default. If you disable the option, the credentials are not transferred and thus they are not displayed on the VM Managers panel when you edit the VM manager.				
9.2.4 vmm_trust_all_vm_managers_certificates	Boolean (true/false)	True		
Specifies whether SSL certificates from all defined VM managers are trusted. If the option is disabled, certificates that are trusted must be placed in the truststore on the computer where the VM Manager Tool is installed. For more information, see: Configuring VM Manager Tool to accept trusted VM manager certificates (on page dccxlv) .				

UUID-based virtual machine filtering

The purpose of filtering based on universally unique identifier (UUID) is to collect the capacity data only for chosen virtual machines. The remaining ones are omitted and not updated in BigFix Inventory.

Enabling UUID-based virtual machine filtering

The `uuid_filtering_enabled` parameter is set by default to false. You must set it to true before you can select virtual machines to be included in a final report.



Note: To change this parameter for basic (central) VM management, log in to BigFix Inventory, go to **Management > Advanced Server Settings**, and change the **vmman_uuid_filtering_enabled** value to true.

1. Go to the computer where you installed an additional instance of VM Manager Tool.
2. Go to `BES Client\LMT\VMMAN\config`.
3. Open the `vmmmainconf.properties` file.
4. Set the `uuid_filtering_enabled` property to true.

You enabled filtering of virtual machines based on their UUIDs. If you want to later reverse the changes, set the property to false.

Selecting virtual machines for UUID-based filtering

You can select virtual machines for whom the capacity data is to be collected by putting their UUIDs in the `vmmfilterconf.properties` file.

1. Obtain the list of virtual machines UUIDs by running the `select distinct uuid from ADM.VIRTUAL_VM_UUID` SQL query. For more information, see [Obtaining the list of virtual machine UUIDs \(on page ccclxxxvii\)](#).
2. In the `BES Client\LMT\VMMAN\config` directory, open the `vmmfilterconf.properties` file.
3. Add the UUIDs of chosen virtual machines
During filter matching, the VM manager removes any white space characters and hyphens (-). All letters are also converted to uppercase.

All of these entries are valid:

- 5030a6eb-485a-35b5-0fa0-a8bc4a459c9d
- 564da050-7b20-8754-b578-e8437da8653e
- 564D1E0d4C0CE65B8A54203D7E032D2B

Run the following command to load the list of UUIDs that you configured:

- **Linux** `./vmman.sh -reloadconfig`
- **Windows** `vmman.bat -reloadconfig`

Obtaining the list of virtual machine UUIDs

You can use the list of virtual machine UUIDs that are currently connected to the BigFix Inventory server to select those for whom the capacity data is to be collected.

You can find all of the virtual machine UUIDs that are currently connected to the BigFix Inventory server in the `ADM.VIRTUAL_VM_UUID` table in the TEMADB database. You can extract the complete list of UUIDs and then remove the chosen UUIDs to omit certain virtual machines.

The following procedure explains how to retrieve the list of UUIDs from the DB2 database. If you are using MS SQL, use a method of retrieving information from database tables that is suitable for this database.

1. Log in to the database server as `db2inst1` or other user with `DBADM` rights.
2. To obtain the list of virtual machines UUIDs, enter the following commands in the system command-line interface:

```
db2 connect to TEMADB
db2 "select distinct uuid from ADM.VIRTUAL_VM_UUID" >uuids.txt
```



Note: `TEMADB` is the default database name. If you are unsure whether it applies to your database, see: [Checking the database name \(on page dcclxx\)](#).

The list is now stored in the `uuids.txt` file.

Remove the unnecessary UUIDs from the `uuids.txt` file and copy the remaining UUIDs to the `vmmfilterconf.properties` file.

Collecting capacity data from virtualization hosts for Xen and KVM

For Citrix, Xen and KVM, connection to the VM manager does not have to be defined in BigFix Inventory to collect capacity data. The data can be collected directly from virtualization hosts by using the Run Capacity Scan on Virtualization Hosts task.

This method is applicable on the following virtualizations:

- **9.2.7** Xen Hypervisor
- **9.2.6** Citrix XenServer



Important: **9.2.12** Starting from application update 9.2.12, it is recommended to collect data from Citrix XenServer by configuring connections to VM managers. For more information, see: [Adding VM managers in central mode \(on page cccl\)](#).

- PowerKVM
- KVM x86



Tip: If your KVM x86 hosts are controlled by RHEV-M, you can configure the RHEV-M as a VM manager in BigFix Inventory. It allows for collecting capacity data from all controlled hosts without the need of using the Run Capacity Scan on Virtualization Hosts task on each of them. For more information, see: [Adding VM managers in central mode \(on page cccl\)](#).

For information about supported versions of these virtualizations, see: [System requirements \(on page c\)](#).

Ensure that virtualization hosts from which you want to collect capacity data meet the following prerequisites.

KVM hosts

Ensure that the following prerequisites are met on the KVM hosts:

- The host runs on Linux x86 or Linux on Power
- Bash shell is available
- `libvirt-client` library is installed (virsh command is available)
- `libxml2` library is installed (xmllint command is available)

Xen hosts

Ensure that the following prerequisites are met on the Xen hosts:

- The host runs on Linux x86
- Bash shell is available
- `xl` command is available
- `libxml2` library is installed (xmllint command is available)

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. In the upper right pane, select **Run Capacity Scan on Virtualization Hosts**.
4. **Optional:** To collect information about host names of virtualization hosts, select **Collect host names of virtualization hosts**.

This option affects the format of information about virtualization hosts that is displayed in the Server ID column on the reports. If you select this option, the column contains information about the vendor, type and host name of the host as specified in the **managedServerTagTemplate** parameter. If you do not select this option, the column contains the serial number of the host instead of its name. For more information about the **managedServerTagTemplate** parameter, see: [Advanced server settings \(on page cdii\)](#).

5. To start the scan, click **Take Action**.



Important: If you run this task on the host, you no longer need to run the **Run Capacity Scan and Upload Results** task on that host.

6. From the list of applicable computers, select the supported hosts.



Important: The list is filtered out to the computers that meet the prerequisites. However, not all of them are virtualization hosts.

7. **Optional:** By default, the capacity scan is scheduled to run every 30 minutes. However, in an environment with many KVM or Xen hosts, consider lowering the frequency of the scan. To specify the frequency of the scan, open the **Execution** tab, specify the details, and click **OK**.

For environments with multiple hosts, run the scan once per one to six hours.

 **Restriction:** You can change the frequency after you obtain the acceptance of the BigFix Compliance Team.

8. Schedule the upload of scan results to the BigFix server.
 - a. In the upper right pane, select **Schedule VM Manager Tool Scan Results Upload**, specify the frequency, and click **Take Action**.
 - b. Select the hosts from which you collected the capacity data, and click **OK**.

 **Tip:** To check the status of the capacity scan on the endpoints, activate the analysis Status of Capacity Scan on Virtualization Hosts. If the status is other than OK, check the return code to learn what is the cause of the problem and how to solve it. For more information, see: [Return codes of capacity scans on virtualization hosts \(on page dcccx\)](#).

9. Wait for the scheduled import or run it manually to transfer the data to BigFix Inventory.

Related information

[Return codes of capacity scans on virtualization hosts \(on page dcccx\)](#)

Removing capacity scan data from the host

During the capacity scan, some files and folders are generated on the virtualization host. If you no longer collect capacity data from a particular host, remove the files and folders that were created by the capacity scan.

1. Stop the actions that were created by the **Run Capacity Scan on Virtualization Hosts** task on the specified endpoints. Otherwise, the scan continues running and the files are recreated.
 - a. In the navigation tree of the BigFix console, click **Sites > External Sites > IBM BigFix Inventory v9 > Actions**.
 - b. Select the actions, and in the lower pane, click **Stop**.
2. In the navigation tree, go to **Fixlets and Tasks**.
3. In the upper right pane, select **Remove Capacity Scan Data from Virtualization Hosts**, and click **Take Action**.
4. Select the computers from which you want to remove the data that was generated during the capacity scan, and click **OK**.

Maintenance tasks

You can perform additional maintenance and troubleshooting tasks on VM Manager Tool. These tasks are not limited to a specific approach to VM management, but can be applied to any instance of VM Manager Tool.

VM Manager Tool installation requirements

Ensure that you fulfill all software requirements to install the VM Manager Tool.

Supported operating systems

Table 41. Requirements for installing VM Manager Tool.

The table consists of three columns and three rows

Operating system	Version	Software requirements
Linux™	Red Hat Enterprise Linux for Intel/AMD x86 (64-bit) versions 7	unzip
	Red Hat Enterprise Linux for Intel/AMD x86 (32 and 64-bit) versions 5 and 6	glibc library 2.7 or higher
	Red Hat Linux Desktop for Intel/AMD x86 (32 and 64-bit) versions 5 and 6	
	SUSE Linux Enterprise Server for Intel/AMD x86 (32 and 64-bit) versions 10 and 11	
	SUSE Linux Enterprise Desktop for Intel/AMD x86 (32 and 64-bit) versions 10 and 11	
Windows™	Server 2012 Standard and Datacenter (64-bit)	
	Server 2008 R2 Standard, Enterprise and Datacenter (64-bit)	
	Server 2008 Standard and Enterprise (32 and 64-bit)	
	8 Ultimate, Professional (32 and 64-bit)	
	7 Ultimate, Enterprise and Professional (32 and 64-bit)	
	Vista Ultimate, Enterprise and Business (32 and 64-bit)	

Required free disk space

The VM Manager Tool requires **250 MB** of disk space.

Log files

You can gather the log files to determine the problems that are related to the VM Manager Tool. You can also change the log settings.

The log files for the VM Manager Tool are in the following directories:

- Trace log files: *BES Client\LMT\VMMAN\logs*
- Installation log files: *BES Client\LMT\VMMAN\logs\install*

You can also gather the complete set of logs by running the `-retrievedebugdata` (on page [ccclxxxi](#)) command in VM Manager Tool.

Log settings

You can change the log settings by editing the `BES Client\LMT\VMMAN\config\log4j.properties` file. The following parameters are the most useful:

- **log4j.appender.mylogger.maxFileSize** specifies the maximum size of a log file

Default value = 1000 KB

- **log4j.appender.mylogger.MaxBackupIndex** specifies the maximum number of log files

Default value = 10

- **log4j.rootLogger** specifies the logging level

Default value = INFO. You can change it to DEBUG to enable debug logging

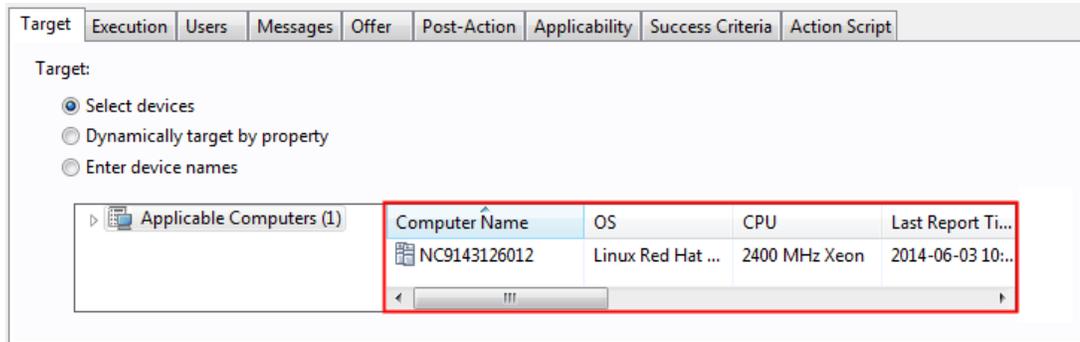
Changing the upload schedule

Capacity data collected by VM Manager Tool is uploaded to the BigFix server every 12 hours. You can change this schedule by removing the default schedule and creating a new one.

1. Stop and remove the default schedule.
 - a. In the navigation tree of the BigFix console, click **Actions**.
 - b. Locate the action **Schedule VM Manager Tool Scan Results Upload**.
 - c. Right-click the action and select **Stop Action**. The schedule is stopped.
 - d. Again, right-click the action and select **Delete Action**. The schedule is removed.
2. Create a new schedule.
 - a. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
 - b. Select **Schedule VM Manager Tool Scan Results Upload**.
 - c. In Description, enter the new frequency in hours.

Description
<p>This task schedules regular uploads of capacity data collected by VM Manager Tool or by the following task: Run Capacity Scan on Virtualization Hosts. The data is uploaded to IBM BigFix and can be transferred to IBM BigFix Inventory by running a data import.</p> <p>Upload data every <input type="text" value="12"/> hours.</p>

- d. Click **Take Action**.
- e. Select the target computer that has VM Manager Tool installed and click **OK**. The action is started under new schedule.



Target

Execution Users Messages Offer Post-Action Applicability Success Criteria Action Script

Target:

Select devices

Dynamically target by property

Enter device names

Applicable Computers (1)

Computer Name	OS	CPU	Last Report Ti...
NC9143126012	Linux Red Hat ...	2400 MHz Xeon	2014-06-03 10:...



Note: All options in the Execution tab are disabled. The frequency can be specified only in the Description field.

You set up the new schedule. The data collected by VM Manager Tool is now uploaded to BigFix with the frequency that you specified.

Forcing the upload of collected data

You can force the collected data to be immediately uploaded to the BigFix server. The data is uploaded once and no schedule is set.

1. Log in to BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. Select **Force VM Manager Tool Scan Results Upload**, and then click **Take Action**.

Name	Source Sev...	Applicab...	Category
Upload Software Scan Results	High	5 / 5	Scanner
Edit Scanner Trace Settings	Low	5 / 5	Troubleshooting
Force VM Manager Tool Scan Results Upload	Low	2 / 5	Troubleshooting
Gather Environment Information	Low	1 / 5	Troubleshooting
Initiate Scanner Diagnostic Tool	Low	5 / 5	Troubleshooting
Upload Scanner Diagnostic Data	Low	0 / 5	Troubleshooting

Task: Force VM Manager Tool Scan Results Upload

Take Action | Edit | Copy | Export | Hide Locally | Hide Globally | Remove

4. Select the target computer and click **OK**.

You uploaded the collected data to the BigFix server.

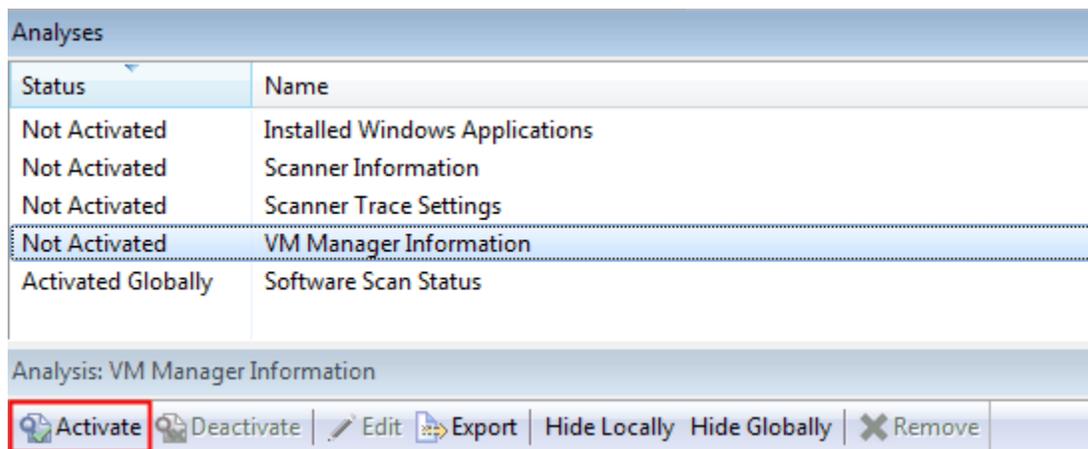
Run a data import in BigFix Inventory to update the data in the user interface.

Checking the VM Manager Tool version

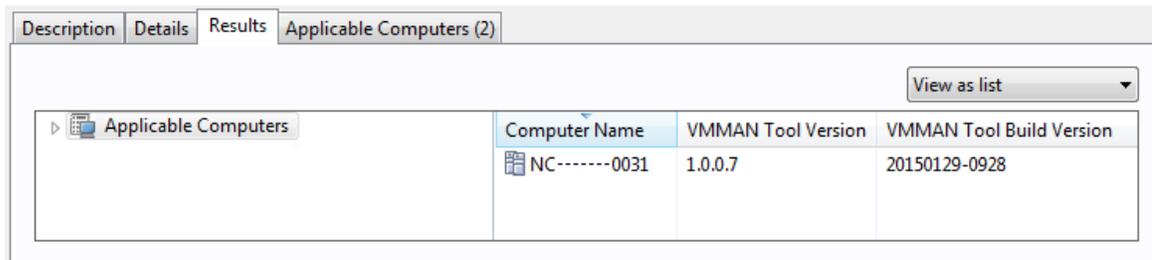
Activate the **VM Manager Information** analysis to retrieve data about the installed version of VM Manager Tool. You can use the analysis to quickly confirm the successful installation of the tool.

The **VM Manager Information** analysis returns information for two properties, `VMMAN_Tool_Version`, which specifies the installed version of the VM Manager Tool, and `Build_Version`, which shows the details of the build.

1. Log in to the BigFix console.
2. In the navigation bar, click **Sites > External Sites > IBM BigFix Inventory v9 > Analyses**.
3. Select **VM Manager Information** and click **Activate**.



4. After you activate the analysis, open the **Results** tab. The information about the version and the build is available next to the name of the applicable computer.



Updating VM Manager Tool

You can update VM Manager Tool to the newest version.

The **Update VM Manager Tool to version *version number*** task can be activated only if there are any older versions of VM Manager Tool.

1. Log in to BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. Select **Update VM Manager Tool to version *version number*** and click **Take Action**.
4. Select the target computer and click **OK**.

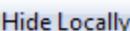
You updated VM Manager Tool to the newest version.

Uninstalling VM Manager Tool

Follow the procedure to uninstall the VM Manager Tool from the endpoints and to delete all the related settings and folders. Uninstallation of the VM Manager Tool is also required if you want to reinstall the tool.

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. Select **Uninstall VM Manager Tool** and click **Take Action**.

Fixlets and Tasks				
Name	Source Sev...	Applicab...	Category	Open
Initiate Scanner Diagnostic Tool	Low	9 / 13	Troubleshooting	0
Install VM Manager Tool	Low	9 / 13	VM Managers	0
Uninstall Scanner	Low	13 / 13	Scanner	0
Uninstall VM Manager Tool	Low	1 / 13	VM Managers	0
Upload VM Manager Tool Scan Results	Low	1 / 13	VM Managers	2

Task: Uninstall VM Manager Tool				
 Take Action	 Edit	 Copy	 Export	 Hide Locally
	 Hide Globally	 Remove		

4. Select the target computer and click **OK**.

Target	Execution	Users	Messages	Offer	Post-Action	Applicability	Success Criteria	Action Script								
Target:																
<input checked="" type="radio"/> Select devices <input type="radio"/> Dynamically target by property <input type="radio"/> Enter device names																
<div style="border: 1px solid gray; padding: 5px;"> Applicable Computers (1) <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th>Computer Name</th> <th>OS</th> <th>CPU</th> <th>Last Report Ti...</th> </tr> </thead> <tbody> <tr> <td>NC9143126012</td> <td>Linux Red Hat ...</td> <td>2400 MHz Xeon</td> <td>2014-06-03 10:...</td> </tr> </tbody> </table> </div>									Computer Name	OS	CPU	Last Report Ti...	NC9143126012	Linux Red Hat ...	2400 MHz Xeon	2014-06-03 10:...
Computer Name	OS	CPU	Last Report Ti...													
NC9143126012	Linux Red Hat ...	2400 MHz Xeon	2014-06-03 10:...													

You uninstalled VM Manager Tool.

Backing up and restoring the VM Manager Tool

To avoid losing configuration of VM manager connections due to corruption of VM Manager Tool files, back up the `/LMT/VMMAN/config/` and `/LMT/VMMAN/keydb/` directories. If the VM Manager Tool files are corrupted, reinstall the VM Manager Tool and restore the two directories from the backup location.

Backing up the VM Manager Tool

To back up the VM Manager Tool, copy the following directories to a backup location:

- `<client_install_dir>/LMT/VMMAN/config/`
- `<client_install_dir>/LMT/VMMAN/keydb/`

Restoring the VM Manager Tool

To restore the VM Manager Tool, reinstall it and copy the backed-up directories to the VM Manager Tool installation directory.

1. Reinstall the VM Manager Tool.
 - a. Log in to the BigFix console.
 - b. In the navigation tree, click **Sites > External Sites > BigFix Inventory v9 > Fixlets and Tasks**.
 - c. To uninstall the existing VM Manager Tool, select **Uninstall VM Manager Tool**, and click **Take Action**. Then, select the computer on which the VM Manager Tool is installed, and click **OK**.
 - d. To install a new VM Manager Tool, select **Install VM Manager Tool version number**, and click **Take Action**. Then, select the computer on which you want to install the VM Manager Tool, and click **OK**.
2. Restore the `/LMT/VMMAN/config/` and `/LMT/VMMAN/keydb/` directories from the backup location.

Best practices for configuring VM managers

Apply the best practices when configuring VM managers to ensure that your infrastructure works efficiently.

Optimal configuration strategies

Follow the optimal configuration strategies to collect all data from your VM managers.

Pooling interval

Pooling interval is a period during which VM Manager Tool is allowed to connect to all configured VM managers, and collect and upload the data from these VM managers. If pooling interval is exceeded, the VM Manager Tool times out and the next data collection cycle is delayed. If you want to count and report subcapacity, it is necessary to capture mobility of virtual machines. Therefore, data from VM managers must be collected every 30 minutes or more frequently. You can change the frequency of data collection only if you use BigFix Inventory for purposes other than subcapacity reporting. To specify the pooling interval, set the value of the `vmm_pooling_time_interval` parameter on the **Advanced Server Settings** panel or in the configuration file for VM Manager. For more information, see: [Advanced Server Settings](#) and [VM Manager Tool Settings](#).

Number of virtual machines

The number of virtual machines that reside on VM managers directly relates to the amount of data that needs to be gathered. If you configure multiple connections from one VM Manager Tool to VM managers that host vast number of virtual machines, problems might occur. If it takes too long to collect data from all configured connections, VM Manager Tool times out as it exceeds the pooling interval time. `_BESClient_ArchiveManager_MaxArchiveSize` parameter limits the total size of files that are uploaded to the BigFix Inventory server. The limit might be exceeded by data that is collected from all configured connections and as a result, the data is not uploaded. For more information about the `_BESClient_ArchiveManager_MaxArchiveSize` parameter, see: [Configuring VM manager for subcapacity reporting \(on page cccxcvii\)](#)

VM managers and network performance

Performance of VM Manager Tool depends on the network typology of your environment as well as the hardware specifications and the load of hypervisors. If network connection is slow or the hypervisors take long time to respond and provide the requested data, performance of VM Manager Tool can be significantly impacted.

Configuring VM manager for subcapacity reporting

If you use BigFix Inventory for subcapacity reporting purposes, perform additional configuration tasks to balance the workload of your VM Manager Tool.

Make sure that you define enough VM Manager Tool connections. Every VM Manager Tool should handle the workload and complete its tasks within the time that is determined by the pooling interval.

1. Set the value of **_BESClient_ArchiveManager_MaxArchiveSize** to at least 50 MB on the endpoint where the VM Manager Tool is installed. In this way, you eliminate the risk of reaching the maximum size of files that can be uploaded to the BigFix server during the configuration.
 - a. Log in to the BigFix console.
 - b. In the left navigation panel and, select **Computers**. Right-click the name of the appropriate computer, and click **Edit Computer Settings**.
 - c. Edit the value of the **_BESClient_ArchiveManager_MaxArchiveSize** parameter to increase the maximum archive size.



Note: **9.2.14** Starting from application update 9.2.14, you can check on the **Computer Support Data** panel whether the value of **_BESClient_ArchiveManager_MaxArchiveSize** is exceeded. For more information, see: [Troubleshooting problems with computers \(on page dccliiv\)](#).

2. Set the value of pooling interval to 30 minutes or less. It is a mandatory setting if you use BigFix Inventory to count subcapacity.
 - For the central VM Manager Tool, go to **Management > Advanced Server Settings** and set the value of **vmm_pooling_time_interval** parameter.
 - For distributed VM Manager Tool, go to the following directory and set the value of the **vmm_pooling_time_interval** parameter in the `vmmmainconf.properties` file.
 - **UNIX** `BESClient/LMT/VMMAN/config/`
 - **Windows** `BESClient\LMT\VMMAN\config`
3. To check whether you defined enough VM Manager Tool connections, perform the following steps.
 - a. Stop the VM Manager Tool service.
 - i. Go to the following directory
 - **UNIX** `BESClient/LMT/VMMAN`
 - **Windows** `BESClient\LMT\VMMAN`
 - ii. Run the following command.

```
./vmmman -stop
```

b. Define the VM manager connection. For more information, see: [Adding VM managers for VMware, Hyper-V, KVM with RHEV-M, Xen Server and Nutanix \(on page ccxiv\)](#).

c. Clear the following directory

- **UNIX** `BESClient/LMT/VMMAN/logs`
- **Windows** `BESClient\LMT\VMMAN\logs`

d. Check whether all connections are correctly defined and gather the data. Go to the following directory and run the command: `./vmmman -runonce`.

e. To make sure that the results of the command `./vmmman -runonce` are correct, check whether you see the following message for every defined connection.

```
CODSV0001I: [ vmmconf_<number>.properties ] Data was successfully retrieved
from the VM manager at https://<URL>
```

f. Go to the following directory.

- **UNIX** `BESClient/LMT/VMMAN/logs`
- **Windows** `BESClient\LMT\VMMAN\logs`

Check the `trace.log` file to see how long it took to complete the `./vmmman -runonce` command.

g. Note the time between log messages, for example:

```
2018-07-25 11:35:38 : (main) com.ibm.license.mgmt.datacollector.Runner:
:initialiseLogger::Logger initialised successfully.
...
2018-07-25 11:37:49 : (main) com.ibm.license.mgmt.datacollector.Runner::runCollector
(Configuration)::Finished 'run once' command.
```

- If the time between log messages stays under 20 minutes, it means that you defined enough VM Manager Tool connections and that every VM Manager Tool completes its tasks within the acceptable time. You can define more connections if needed.
- If the time between log messages goes above 20 minutes, distribute connections to another VM Manager Tool so that every VM Manager Tool completes its tasks within that time.
 - To limit the amount of data collected by a single connection, set users with limited access to a subset of computers and split the load between multiple VM Manager Tool connections.



Tip: Even though the pooling interval is 30 minutes, keep the time between log messages under 20 minutes. You need to leave some breathing room for VM migrations, higher network traffic or higher hypervisor workload periods.

- h. Run the `./vmman -run` command to start the VM Manager Tool service.
 - i. Check the size of the data set that was collected from this VM Manager Tool and whether it exceeds the `_BESClient_ArchiveManager_MaxArchiveSize` value.
4. After you complete the configuration of VM Manager Tool connections, results are stored in the following directory. To send the results to the BigFix server, run the **Schedule VM Manager Tool Scan Results Upload** task.

- **UNIX** `BESClient/LMT/VMMAN/upload`
- **Windows** `BESClient\LMT\VMMAN\upload`

You can run **Force VM Manager Tool Scan Results Upload** to upload the result files to the BigFix server faster.

5. When the task completes, the result file or files are stored in the following directory. Result file names follow the specific naming convention: `1323215_<number>_vmman.tar.gz`.

- **UNIX** `BESClient/LMT/VMMAN/`
- **Windows** `BESClient\LMT\VMMAN\`

- Your VM manager configuration is complete. If at any point your previously configured VM Manager Tool stops collecting data, go to the following directory.

- **UNIX** `BESClient/LMT/VMMAN/logs`
- **Windows** `BESClient\LMT\VMMAN\logs`

Check if the log file contains the following message.

```
"[WARNING] Data collection from some defined vm managers did not finish within defined pooling interval 30. Please consider increasing the vmm_pooling_time_interval setting in vmmmainconf.properties file or divide your vm manager connections to more VM Manager Tools."
```

If that happens, you need to configure new VM Manager Tool and split the workload.

- You can check whether `_BESClient_ArchiveManager_MaxArchiveSize` is exceeded by checking the **Software Scan Status** analysis. The **Is archive size exceeded?** column provides the information for computers with VM Manager Tool installed.

Administering the server

Perform administering tasks on your server to configure advanced settings, or restart its services.

Starting the server

You can start the BigFix Inventory server by running the `srvstart.bat` or `srvstart.sh` script.

Linux Starting the server on Linux

To start the server, run the `srvstart.sh` script.

cd

Only the user who installed BigFix Inventory or has root privileges can perform this task.

1. **Optional:** Check if your DB2® instance is running.
 - a. Log on to the computer where DB2® is installed. You must be a DB2® instance owner.
 - b. Type `db2start` at the command line.
2. Go to the BigFix Inventory installation directory.
3. Open the `cli` directory.
4. Run `srvstart.sh`.

Windows Starting the server on Windows

To start the server, run the `srvstart.bat` script.

1. **Optional:** Check if your SQL Server instance is running.
 - a. Log on to the computer where the SQL Server® is installed. You must have administrative privileges or be an SQL Server service owner.
 - b. Start the SQL Server instance. For more information, see: [Start, Stop, Pause, Resume, Restart SQL Server Services](#).
2. Go to the BigFix Inventory installation directory.
3. Open the `cli` directory.
4. Run `srvstart.bat`.

Stopping the server

You can stop the BigFix Inventory server by running the `srvstop.bat` or `srvstop.sh` script.

Linux Stopping the server on Linux

To stop the server, run the `srvstop.sh` script.

Only the user who installed BigFix Inventory or has root privileges can perform this task.

1. Go to the BigFix Inventory installation directory.
2. Open the `cli` directory.
3. Run `srvstop.sh`.
4. **Optional:** After stopping the server, check if the DB2® instance is stopped:
 - a. Log on to the computer where DB2® is installed. You must be a DB2® instance owner.
 - b. Type `db2stop` at the command line. The DB2® instance stops.

Windows Stopping the server on Windows

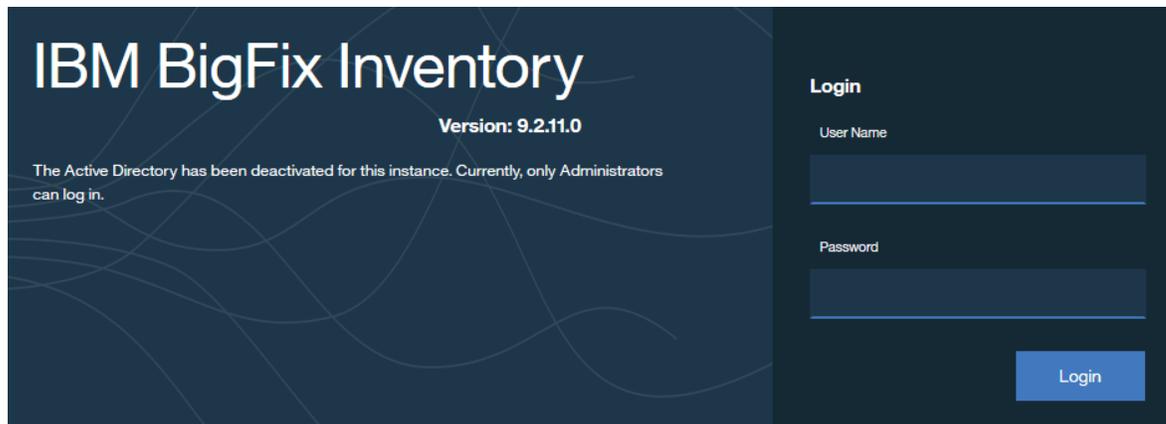
To stop the server, run the `srvstop.bat` file.

1. Go to the BigFix Inventory installation directory.
2. Open the `cli` directory.
3. Run `srvstop.bat`.

Customizing the login page

You can enter a custom message that will appear on the login page to differentiate between multiple instances of BigFix Inventory, or to provide extra information related to this instance. The custom message is commonly used to notify other users about specific conditions related to the instance, such as important configuration changes.

1. Log in to BigFix Inventory, and go to **Management > Session Settings**.
2. In the **Custom Message on Login Page** field, enter your message. After you click **Save**, the message will be displayed on the login page for all users of BigFix Inventory.



Setting session timeout

The session timeout specifies the period of user inactivity after which the session ends, and the user is required to log in again. Before it happens, you are notified that the session is about to end.

1. In BigFix Inventory, click **Management > Session Settings**.
2. Enter the timeout value in seconds. By default, it is set to 3600 seconds.
3. Click **Save**.

9.2.3 Configuring advanced server settings

9.2.3 Available from 9.2.3. You can change advanced settings of the BigFix Inventory server, for example settings related to the configuration of the central VM Manager Tool, on the **Advanced Server Settings** panel.



You must be an Administrator to perform this task.



Important: Configuration of advanced server settings is intended for users with in-depth knowledge of BigFix Inventory. If you are not sure about the impact of a particular parameter on your environment, do not change the parameter.

The following procedure describes how to change advanced server settings on the user interface. This option is available starting from version 9.2.3. For earlier versions, you can change the settings by using REST API calls. For more information, see: [REST API for changing advanced server settings \(on page cmlvi\)](#).

1. To open the panel, go to **Management > Advanced Server Settings**. If the panel is not available in the menu, go to the following URL: https://hostname:port/management/sam/control_values.
2. Click a parameter that you want to change and provide its new value.

For more information about each parameter, see: [Advanced server settings \(on page cdii\)](#).



Restriction: You cannot change advanced server settings when the import is in progress.

3. To save the changes, click **Save**.

9.2.3 Advanced server settings

9.2.3 Available from 9.2.3. Review to the list of parameters that you can use to configure the BigFix Inventory server. You can change the parameters on the **Advanced Server Settings** panel or through REST API.

To open the panel, go to **Management > Advanced Server Settings**. If the panel is not available in the menu, go to the following URL: https://hostname:port/management/sam/control_values. For information how to change the settings through REST API, see: [Configuring advanced server settings \(on page cdi\)](#).

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
9.2.8 <ul style="list-style-type: none"> • Name: Account lockout: enable account lockout • Parameter: <code>user_lockout_enabled</code> 	Boolean (true/false)	True		
Specifies whether exceeding the specified number of failed login attempts locks the user account.				
9.2.8 <ul style="list-style-type: none"> • Name: Account lockout: lockout period • Parameter: <code>user_lockout_length</code> 	Minutes	5	1	60
Specifies the period for which a user cannot log in to the application after exceeding the maximal number of failed login attempts.				
9.2.8 <ul style="list-style-type: none"> • Name: Account lockout: login attempts period • Parameter: <code>user_retry_time_period</code> 	Minutes	5	1	60
Specifies the period during which the user can try to log in to the application. When the maximal number of failed login attempts is exceeded within this period, the account is locked.				

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
Description				
<p>9.2.8</p> <ul style="list-style-type: none"> • Name: Account lockout: maximal number of failed login attempts • Parameter: <code>user_max_login_retries</code> 	Integer	10	2	128
<p>Specifies the maximum number of failed login attempts after which the user account is locked.</p>				
<p>9.2.8</p> <ul style="list-style-type: none"> • Name: Automatically decommission inactive computers • Parameter: <code>decommission_inactive_computers</code> 	Boolean	False		
<p>Enables the option to automatically remove inactive computers from BigFix Inventory after the period defined in the max_data_visibility_for_inactive_computers (on page <i>cdxi</i>) parameter. This option is disabled by default.</p>				
<p>9.2.17</p> <ul style="list-style-type: none"> • Name: Automatically delete the scan results after the successful import from the disconnected data source • Parameter: <code>delete_successfully_imported_scans</code> 	Boolean	False		
<p>Enables the option to automatically remove the scan results after they are successfully imported from the disconnected data source. After you set this option to true, automatic removal is applied after each subsequent import of data. You can manually remove remaining files, such as older scan results or files that are skipped. The option is useful especially when you handle many scan results packages. The removal does not affect the report data and the scan results are included in the reports.</p>				
<ul style="list-style-type: none"> • Name: CSV report separator • Parameter: <code>csvReportSeparator</code> 	Character	Comma (,)		
<p>Specifies the character that is used as a line separator during the creation of CSV files.</p>				
<ul style="list-style-type: none"> • Name: Calculate metric consumption for incomplete computers • Parameter: <code>calculateLicenseUsageForIncompleteComputers</code> 	Boolean (true/false)	True		
<p>Specifies how PVU consumption on x86 virtual machines that have incomplete data is counted. If set to true, PVU consumption is counted based on the number of PVUs on the host. In this case, the reported PVU consumption might be higher than the real value, but</p>				

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
	configuration of VM managers is not required. If set to false, the reported PVU consumption is 0, and VM managers must be configured.			
<p>9.2.10</p> <ul style="list-style-type: none"> • Name: Calculating usage per computer • Parameter: <code>calculate_during_import/product_usage</code> <p>Previously: <code>on_demand_sam/product_usage</code></p>	Boolean (true/false)	True		
<p>Specifies when the values that are displayed on the Usage per Computer report are calculated. By default, usage per computer is calculated during every import. If you change the setting to false, the usage is calculated on demand. To update the report on demand, open it and click Calculate.</p> <p>After you change the value of the parameter, restart the BigFix Inventory server for the change to take effect.</p>				
<p>9.2.3</p> <ul style="list-style-type: none"> • Name: Catalog distribution waiting time • Parameter: <code>catalog_download_task_stagger_interval</code> 	Minutes	0	0	10080 (1 week)
<p>Specifies the maximum time for which the distribution of the software catalog to the endpoints is postponed to reduce network traffic.</p> <p>The distribution of the catalog to each endpoint is postponed for a random amount of time that is within the specified waiting time. For example, if you set the waiting time to 1 week, the first computer can download the catalog after an hour since the import of the catalog, the second computer after a day, and the third computer after a week. The amount of time is randomly chosen by each computer.</p>				
<p>9.2.7</p> <ul style="list-style-type: none"> • Name: Collect VMware license metric utilization • Parameter: <code>vmman_slm_tags_enabled</code> 	Boolean (true/false)	False		
<p>Enables collection of VMware license metric utilization. The option is disabled by default. The setting applies only to the central VM.</p>				
<p>9.2.10</p>	Days	3	1	90

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
Description				
<ul style="list-style-type: none"> • Name: Computer approaching detachment from VM manager • Parameter: <code>vmManagerDetachmentNotificationPeriod</code> 	<p>Specifies the maximum idle time after which data from the VM manager is considered to be outdated. After that time, the computer status changes to <i>Outdated VM Manager Data</i>. The value should be lower than the value of the vmManagerDetachmentPeriod (on page cdv) parameter.</p>			
<ul style="list-style-type: none"> • Name: Computer detachment from VM manager • Parameter: <ul style="list-style-type: none"> ◦ 9.2.8 <code>vmManagerDetachmentPeriod</code> ◦ Previously: <code>computerVmManagerDetachmentPeriod</code> 	Days 30 Previously: 7	9.2.10 30	1	90
<p>Specifies the maximum idle time after which a guest computer or a host computer is considered to be detached from the VM manager. After that time, the computer status changes to <i>No VM Manager Data</i>. The data that is sent by an agent is not augmented by the data that is retrieved from the VM manager.</p> <p>9.2.8 Starting from application updated 9.2.8, if the host is a part of a cluster, it is also disconnected from that cluster.</p> <p>The value of this parameter should be greater than the value of the vmManagerDetachmentNotificationPeriod (on page cdv) parameter.</p>				
<p>9.2.3</p> <ul style="list-style-type: none"> • Name: Deploy tasks automatically • Parameter: <code>enable_automatic_task_deployment</code> 	Boolean (true/false)	True		
<p>Specifies whether tasks initiated by the BigFix Inventory server are automatically deployed to the BigFix server. The tasks include:</p> <ul style="list-style-type: none"> • Import of the new software catalog to the BigFix server. • Configuration of scan frequency and schedule on the Scan Configurations panel. 				

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
Description				
<p>9.2.11</p> <ul style="list-style-type: none"> • Name: Display data from partial imports • Parameter: <code>allow_partial_data_view</code> 	Boolean (true/false)	False		
<p>• Automatic configuration of the fixlet site on the BigFix server that is added as a new data source.</p> <p>• Installation and configuration of the main VM Manager Tool on the BigFix server that is added as a new data source.</p> <p>The option is enabled by default which is the recommended setting.</p> <p> Warning: If you disable the automatic deployment of tasks, you need to perform a number of manual tasks to ensure that the data needed by BigFix Inventory is collected and that proper values are displayed on the reports.</p> <p> Important: Use this parameter only when you are requested by IBM Support.</p> <p>Enables displaying data from the last import regardless of whether the import was successful or partial. It might cause inconsistency or lack of results on some reports. By default, the parameter is disabled and only data from successful imports is displayed. To enable this parameter, first enable the enable partial imports (on page cdvii) parameter.</p>				
<p>9.2.3</p> <ul style="list-style-type: none"> • Name: Download software catalog automatically • Parameter: <code>enable_automatic_catalog_download</code> 	Boolean (true/false)	True		
<p>Specifies whether the new software catalog is automatically distributed to the endpoints. If you disable this option, you must distribute the catalog manually. For more information, see: Updating scanner catalogs (on page cclxvii).</p>				

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
<p>9.2.3</p> <ul style="list-style-type: none"> • Name: Enable debug level logging during data imports • Parameter: <code>debug_logging_for_imports</code> 	Boolean (true/false)	False		
	Specifies whether import logs contain debug level information. When the parameter is set to false, the default logging level is used.			
<p>9.2.9</p> <ul style="list-style-type: none"> • Name: Enable information about files and packages that cause software detection • Parameter: <code>enable_caused_detection_data</code> 	Boolean (true/false)	False		
	Enables the data that specifies whether a file or a package contained enough information to cause detection of the related software. After enabling the data, you can display the Cause Detection column on the Scanned File Data and Package Data reports. This option is disabled by default.			
<p>9.2.11</p> <ul style="list-style-type: none"> • Name: Enable partial imports of data • Parameter: <code>enable_partial_imports</code> 	Boolean (true/false)	False		
	Enables skipping certain steps during the import of data. The import continues even if these steps fail. When the import finishes, it is marked as partial because only data from successfully completed steps is processed.			
<p>9.2.10</p> <ul style="list-style-type: none"> • Name: Enabled schema next feature • Parameter: <code>schema_next</code> 	Boolean (true/false)	False		
	Optimizes the process of importing data from BigFix. The option is available for BigFix 9.5.5 and higher. If you upgraded the BigFix server from earlier versions, run at least one import of data before you enable the optimization to ensure data consistency.			

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
<ul style="list-style-type: none"> • Name: Location of temporary application files • Parameter: <code>tempPathForGeneratedFiles</code> 	Path	Period (.)		
	Specifies the path to the folder where temporary application files, such as the CSV reports during generation, are stored. The default value indicates that the files are created in the following location: <ul style="list-style-type: none"> • Linux <code>install_dir/wlp/usr/servers/server1</code> • Windows <code>install_dir\wlp\usr\servers\server1</code> 			
9.2.12 <ul style="list-style-type: none"> • Name: Location of CVE files upload • Parameter: <code>tempPathForCVEFiles</code> 	Path	<code>../../../../cve_data</code>		
	Specifies the path to the folder where temporary CVE files should be stored for an upload. The default value indicates that the files are stored in the <code>cve_data</code> folder that is in the installation directory of the BigFix Inventory server.			
<ul style="list-style-type: none"> • Name: Maximum VM manager inactivity • Parameter: <code>maxVMManagerInactivity</code> 	Days	3	1	90
	Specifies the maximum time after which a VM manager is considered inactive if the server receives no new data from this VM manager. After that time, the status of the VM manager changes to <i>In-active</i> .			
9.2.4 <ul style="list-style-type: none"> • Name: Maximum VM manager visibility • Parameter: <code>maxVMManagerVisibility</code> 	Days	90	1	2147483647
	Specifies the number of days after which a VM manager is removed from the VM Managers panel if no data is collected from that VM manager.			
9.2.7 <ul style="list-style-type: none"> • Name: Maximum number of VMware <code>.slmtag</code> files • Parameter: <code>vmman_slm_tags_files_per_software</code> 	Number	10	1	10
	The maximum number of <code>.slmtag</code> files that are stored on the computer for each license key under VMWare instance.			

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
<p>9.2.2</p> <ul style="list-style-type: none"> • Name: Maximum number of aggregation threads • Parameter: <code>maxAggregationThreads</code> 	Number	2	1	64
<p> Important: Change the setting only if you created computer groups to comply with IBM subcapacity regions or if you are a service provider and created computer groups that represent your customers. Otherwise, changing this parameter has no impact on aggregation performance.</p> <p>Specifies the maximum number of threads that can be run during the aggregation. If you increase their number, performance of aggregation improves. However, memory usage and database locking increase. You should provide two processor cores on the database server for each thread.</p>				
<ul style="list-style-type: none"> • Name: Maximum number of import threads • Parameter: <code>numberOfImportThreads</code> 	Number	0	0	32
<p>Specifies the maximum number of threads that are used to process capacity scan data and VM manager scan data during the import of data from the BigFix server.</p> <p> Tip: In environments with over 10000 endpoints that are running on high-performance machines, it can be beneficial to increase the number of threads to speed up the import.</p>				
<ul style="list-style-type: none"> • Name: Maximum number of processed license metric files • Parameter: <code>maximumNumberOfLicenseUseMetricFilesToProcessPerImport</code> 	Number	-1	-1	2147483647
<p>Specifies the maximum number of license metric tag files that can be processed during an import. The oldest unprocessed files are processed in the first place. Then, the specified number of newer files is consecutively processed with every import. The value -1 indicates that all files are processed in a single import.</p>				
<p>9.2.2</p>	Number	4	1	64

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
Description				
<ul style="list-style-type: none"> • Name: Maximum number of reagggregation threads • Parameter: <code>maxReagggregationThreads</code> 	<p> Important: Change the setting only if you created computer groups to comply with IBM subcapacity regions or if you are a service provider and created computer groups that represent your customers. Otherwise, changing this parameter has no impact on reagggregation performance.</p> <p>Specifies the maximum number of threads that can be run during reagggregation. If you increase their number, performance of reagggregation improves. However, memory usage and database locking increase. You should provide two processor cores on the database server for each thread.</p>			
<p style="background-color: #0056b3; color: white; padding: 2px;">9.2.4</p> <ul style="list-style-type: none"> • Name: Maximum number of rows exported to PDF • Parameter: <code>pdfReportRowLimit</code> 	Number	100 000	0	2147483647
<p>Specifies the maximum number of rows that can be exported to a PDF file. If the number of rows in the report is larger than the specified limit, appropriate information is provided in the generated PDF file.</p>				
<p style="background-color: #0056b3; color: white; padding: 2px;">9.2.3</p> <ul style="list-style-type: none"> • Name: Maximum number of rows in REST API for raw scan results • Parameter: <code>raw_data_api_default_limit</code> 	Number	100 000	0	2147483647
<p>Specifies the maximum number of rows to be retrieved in a single request when using REST API for retrieving raw scan results. This parameter applies only if the REST API limit parameter, which always takes precedence, is omitted in the request. The default value is 100 000. 0 indicates that all rows are retrieved.</p>				
<p style="background-color: #0056b3; color: white; padding: 2px;">9.2.7</p> <ul style="list-style-type: none"> • Name: Maximum number of rows to select on the Software Classification report • Parameter: <code>selectAllLimit</code> 	Number	250000	0	500000
<p>Specifies the maximum number of rows that can be selected at once on the Software Classification report.</p>				
<p style="background-color: #0056b3; color: white; padding: 2px;">9.2.7</p>	Bytes	10240	10240	1048676

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
Description				
<ul style="list-style-type: none"> • Name: Maximum size of VMware .slm-tag files • Parameter: <code>vmman_slm_tags_file_size_limit</code> 	<p>The maximum size of each <code>.slmtag</code> file that is used to collect VMware license metric utilization.</p>			
<p>9.2.8</p> <ul style="list-style-type: none"> • Name: Maximum visibility of data for inactive computers • Parameter: <code>max_data_visibility_for_inactive_computers</code> 	Days	90	1	1000
<p>Specifies the number of days after which the data from the inactive computer no longer influences metric calculation and is not visible in BigFix Inventory. To enable this feature change the decommission_inactive_computers (on page cdiii) parameter to true.</p>				
<p>9.2.8</p> <ul style="list-style-type: none"> • Name: Number of days to count Microsoft Windows Server metrics after VM migration • Parameter: <code>historical_period_for_microsoft_metrics</code> 	Days	90	0	180
<p>Specifies the number of days for which data from a VM contributes to metric calculation of Microsoft Windows Server after the VM is moved to a different host. Run an import after you change the value of this parameter to recalculate the data.</p>				
<p>9.2.10</p> <ul style="list-style-type: none"> • Name: Number of days to keep the removed raw scan results • Parameter: <code>raw_data_api_history_keep_days</code> 	Days	7	-1	30
<p>Specifies the number of days after which historical raw scan data is removed from the database. The data was used to discover software but cannot be used for this purpose again. It is not displayed on any reports. It can only be retrieved by using REST API. The data is no longer needed and can be safely removed from the database without any impact on the reports.</p> <p>Set the value of the parameter to 0 to remove the data after every successful import. You can set it to -1 to never remove the data. However, such a setting is not recommended as it can generate lots of unnecessary data that impacts application performance.</p>				

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
<p>9.2.10</p> <ul style="list-style-type: none"> • Name: (Currently not used) Number of days to keep internal application data • Parameter: <code>pruning_days_to_keep</code> 	Days	0	-1	2147483647
	<p>Specifies the number of days after which internal data that is used by the application is removed from the database. The data is not displayed on any reports and can be safely removed from the database.</p> <p>Set the value of the parameter to 0 to remove the data after every successful import. You can set it to -1 to never remove the data. However, such a setting is not recommended as it can generate lots of unnecessary data that impacts application performance.</p>			
<p>9.2.7</p> <ul style="list-style-type: none"> • Name: Optimize migration from 7.x for large environments • Parameter: <code>new_migration_mode_enabled</code> 	Boolean (true/false)	True		
	<p>Specifies whether the migration from version 7.x to 9.x is optimized for large environments.</p>			
<p>9.2.12</p> <ul style="list-style-type: none"> • Name: Optimized shared disks scan: CPU threshold • Parameter: <code>autoscan_shared_disks_cpu_threshold</code> 	Integer	0	0	100
	<p>Limits the amount of processor resources that the scanner consumes during the optimized scan of shared disks. By default, the option is set to 0 which means that the limit is disabled.</p>			
<p>9.2.12</p> <ul style="list-style-type: none"> • Name: Optimized shared disks scan: enable • Parameter: <code>autoscan_shared_disks_enable</code> 	Boolean (true/false)	False		
	<p>Enables optimized scan of shared disks. When the option is enabled, scanning of shared disks is automated. As a result, the following actions are performed automatically:</p>			

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
				<ul style="list-style-type: none"> • Computers are designated to scan shared disks that exist in your infrastructure • Software scans are scheduled on the designated computers • New computers are designated if problems occur with scanning the initially designated computers <p>The option is disabled by default. For more information, see: Discovering software on shared disks (on page ccxxi).</p>
<p>9.2.12</p> <ul style="list-style-type: none"> • Name: Optimized shared disks scan: maximum inactivity time • Parameter: <code>max_inactivity_of_designated_shared_disk_scanner</code> 	Days	14	1	365
<p>9.2.12</p> <ul style="list-style-type: none"> • Name: Optimized shared disks scan: maximum waiting time • Parameter: <code>max_waiting_for_share_disk_scan_for_new_designated_endpoint</code> 	Days	7	1	90
<p>9.2.12</p> <ul style="list-style-type: none"> • Name: Optimized shared disks scan: scan interval • Parameter: <code>autoscan_shared_disks_scan_interval_days</code> 	Days	7	1	30
<p>9.2.3</p>	Boolean (true/false)	False		

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
Description				
<ul style="list-style-type: none"> • Name: Preserve Initiate Software Scan actions on the BigFix server • Parameter: <code>scanConfig_preserve_InitiateScanAction</code> 	Specifies whether Initiate Software Scan actions that are created on the BigFix server are preserved when the scan configuration is deleted on the Scan Configurations panel. If you set the value of this parameter to true, the actions are stopped but are not removed from the BigFix server.			
<div style="background-color: #0056b3; color: white; padding: 2px; margin-bottom: 5px;">9.2.13</div> <ul style="list-style-type: none"> • Name: Remove information about current software users during next import • Parameter: <code>remove_user_information</code> 	Boolean(true/false)	False		
Specifies whether information about current software users is removed during the next import. Historical aggregation data for Registered User metric will also be removed from the All Metrics report. This option is disabled by default.				
<div style="background-color: #0056b3; color: white; padding: 2px; margin-bottom: 5px;">9.2.4</div> <ul style="list-style-type: none"> • Name: Re-synchronize datasources after database restoration • Parameter: <code>resynchronize_datasources_once</code> 	Boolean (true/false)	False		
Enables an import that retrieves all software scan data from the restored BigFix database. Such an import is time-consuming and should be run only when problems with incomplete data occur after the database is restored. When you enable the option, the import runs once and returns to the normal operation mode.				
<div style="background-color: #0056b3; color: white; padding: 2px; margin-bottom: 5px;">9.2.11</div> <ul style="list-style-type: none"> • Name: Re-synchronize Scan Data • Parameter: <code>resync_imports</code> 	Integer	0	0	60
Defines the number of consecutive imports during which scan data is re-synchronized. Resynchronization is needed if you want to display detailed file version for software components that were discovered before you upgraded BigFix Inventory to version 9.2.11. To define the scope of the re-synchronized data, change the value of the resync_mode (on page cdxv) parameter.				
<div style="background-color: #0056b3; color: white; padding: 2px; margin-bottom: 5px;">9.2.11</div>	Enum	File Scan Only (0)		
Defines the scope of scan data that is re-synchronized.				

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
Description				
<ul style="list-style-type: none"> • Name: Re-synchronize Scan Data Mode • Parameter: <code>resync_mode</code> 	<p>Possible values:</p> <ul style="list-style-type: none"> • File Scan Only (0) - Only results of the file system scan are resynchronized. • All Scans (1) - Results of the file system scan, catalog-based scan, and software identification tags scan are re-synchronized. 			
<p>9.2.9</p> <ul style="list-style-type: none"> • Name: Recompute unrecognized scan data aggregates • Parameter: <code>recompute_unrecognized_scan_data</code> 	Enum	Never (0)		
<p>Specifies how often unrecognized scan data is calculated. The data allows for determining the most effective ways of improving signature coverage, but slows down imports, particularly in larger deployments. The aggregated data is visible on the Unrecognized Files report.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • Never (0) - Scan data aggregates are not calculated. • Next Import Only (1) - Scan data aggregates are calculated only during the next import. Then, the value of the parameter changes to Never (0). • Every Import (2) - Scan data aggregates are calculated during every import. 				
<ul style="list-style-type: none"> • Name: Reporting Missing Software Scan status <p>Previously: Software scan interval</p> <ul style="list-style-type: none"> • Parameter: <code>scanHealthLastScanAttempt</code> 	Days	30	1	365
<p>Specifies the number of days from the last software scan after which the computer is reported on the Scan Health widget as missing software scan.</p>				

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
<ul style="list-style-type: none"> • Name: Scan files import waiting time • Parameter: <code>scanFilesDownloadImportTimeout</code> 	Seconds	10	1	600
<p>Specifies the maximum number of seconds to wait for a file with scan results to be read from the BigFix server during import.</p>				
<p>9.2.3</p> <ul style="list-style-type: none"> • Name: Server ID template for managed hosts • Parameter: <code>managedServerTagTemplate</code> 	String	<code>%VENDOR %TYPE</code> <code>%HOSTNAME</code>		
<p>Specifies the format of information about VM manager hosts that is displayed in the Server ID column on the reports. The parameter affects hosts from which data is collected by the VM Manager Tool or by the Run Capacity Scan on Virtualization Hosts fixlet.</p> <p>By default, the column contains information about the vendor, type and host name of the VM manager host. For example: IBM NC123456.IBM.COM. You can choose any combination of the following characteristics of a VM manager host to be displayed on the reports:</p> <ul style="list-style-type: none"> • Host name (REST API value: <code>%HOSTNAME</code>) <p>! Important: To display this information for hosts from which data is collected by the Run Capacity Scan on Virtualization Hosts fixlet, you must first enable the collection of host names. To do this, select the Collect host names of virtualization hosts check box in the fixlet. Otherwise, a serial number is displayed instead of the host name for such VM manager hosts.</p> <ul style="list-style-type: none"> • Identifier (REST API value: <code>%ID</code>) • Model (REST API value: <code>%MODEL</code>) • Serial number (REST API value: <code>%SERIAL_NUMBER</code>) 				

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
<p>9.2.3</p> <ul style="list-style-type: none"> • Name: Server ID template for unmanaged hosts • Parameter: <code>standaloneServerTagTemplate</code> 	String	%VENDOR %TYPE %SERIAL_NUMBER		
<p> Important: To display this information for hosts from which data is collected by the VM Manager Tool, you must first enable the collection of serial numbers. To do this, set the value of the <code>vmman_collecting_host_serials_enabled</code> (on page cdxxii) parameter to true.</p> <ul style="list-style-type: none"> • Type (REST API value: %TYPE) • Vendor (REST API value: %VENDOR) <p> Note: Information is displayed in the new format after the next import.</p> <p>To change the format of information through REST API, list the characteristics of the VM manager host that you want to display in the value parameter. The value must be URL-encoded. For example:</p> <pre>PUT http://localhost:9081/api/sam/configs?token=7adc3efb175e2bc0f4484bdd2efca54a8fa&name=managedServerTagTemplate&value=%25VENDOR%20%25TYPE%20%25SERIAL_NUMBER</pre>				
<p>Specifies the format of information about the host that is displayed in the Server ID column on the reports. The parameter affects hosts that are not affected by the <code>managedServerTagTemplate</code> (on page cdxvi) parameter.</p> <p>By default, the column contains information about the vendor, type and serial number of the host. For example: IBM Corp. 7946 99B7166. You can choose any combination of the following characteristics of a host to be displayed on the reports:</p>				

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
Description				
	<ul style="list-style-type: none"> • Identifier (REST API value: %ID) • Model (REST API value: %MODEL) • Serial number (REST API value: %SERIAL_NUMBER) • Type (REST API value: %TYPE) • Vendor (REST API value: %VENDOR) <p> Note: Information is displayed in the new format after the next import.</p> <p>To change the format of information via REST API, list the characteristics of the host that you want to display in the value parameter. The value must be URL-encoded. For example:</p> <pre data-bbox="722 1045 1052 1192">PUT http://localhost:9081/api/sam/configs? token=7adc3efb175e2bc0f4484bdd2efca54a8fa &name=standaloneServerTagTemplate&value =%25VENDOR%20%25MODEL%20%25SERIAL_NUMBER</pre>			
<p>9.2.8</p> <ul style="list-style-type: none"> • Name: Size limit of Scan File to process • Parameter: scanFileSizeLimit 	MB	20	0	100
<p>Specifies the maximum size of the scan file that can be processed by BigFix Inventory. Files that exceed the limit are rejected. When you set the parameter to 0, the limit is disabled and all files are process regardless of their size.</p>				
<ul style="list-style-type: none"> • Name: Sort bundling options alphabetically • Parameter: blockUiBundlingComputations 	Boolean (true/false)	False		
<p>Specifies whether bundling options that are displayed when you reassign a software component are sorted by confidence or alphabetically. By default, the options are sorted by confidence. If you set the value of the parameter to true, the options are sorted alphabetically and thus are displayed more quickly on the user interface.</p>				
<p>9.2.8</p>	Boolean (true/false)	False		

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
Description				
<ul style="list-style-type: none"> • Name: Split SQL query for inserting software facts • Parameter: <code>split_ms_merge</code> 	<p> Important: Use this parameter only when you are requested by IBM Support.</p> <p>Splits the SQL query that is used for inserting software facts into multiple queries to reduce usage of TEMPDB in a large-scale environment. Splitting the query lengthens the import. The parameter is applicable only on Windows.</p>			
<ul style="list-style-type: none"> • Name: Store hardware data for VM manager nodes • Parameter: <code>storeHwDataForAllVM-ManagerNodes</code> 	Boolean (true/false)	False		
<p>Specifies whether information about nodes and clusters that is retrieved from VM managers is stored in the database, regardless of whether an agent is running on any virtual machine on such nodes or clusters.</p>				
<p>9.2.7</p> <ul style="list-style-type: none"> • Name: The default value of the Automatically confirm the affected components check box on the Software Classification panel • Parameter: <code>automaticConfirm</code> 	Boolean (true/false)	True		
<p>Specifies the default value of the Automatically confirm the affected components check box on the Software Classification. If the parameter is set to true, the check box is selected by default and components affected by an action are confirmed. Otherwise, affected component are not confirmed.</p>				
<p>9.2.13</p> <ul style="list-style-type: none"> • Name: To Do: Show that VM manager connection is missing • Parameter: <code>todo_display_vm_managers_need</code> 	Boolean (true/false)	True		
<p>Specifies whether the To Do List shows information about missing connections to VM managers. By default, the parameter is set to true which means that the item is displayed.</p>				
<p>9.2.7</p>	Integer	0	0	1024
<p>Specifies the maximum length of the user password. By default, the parameter is set to 0 which means that the requirement is disabled.</p>				

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
Description				
<ul style="list-style-type: none"> • Name: User password: maximal number of characters • Parameter: <code>user_password_max_length</code> 				
<p>9.2.7</p> <ul style="list-style-type: none"> • Name: User password: maximal number of the same characters in a row • Parameter: <code>user_password_max_of_identical_characters_in_row</code> 	Number	0	0	32
<p>Specifies the maximum number of the same characters that can be used in a row in the user password. By default, the parameter is set to 0 which means that the requirement is disabled.</p>				
<p>9.2.7</p> <ul style="list-style-type: none"> • Name: User password: minimal number of characters • Parameter: <code>user_password_min_length</code> 	Number	8	0	128
<p>Specifies the minimum length of the user password. Set the value to 0 to disable the requirement.</p>				
<p>9.2.7</p> <ul style="list-style-type: none"> • Name: User password: regular expression • Parameter: <code>user_password_regular_expression</code> 	Regular Expression	<code>^(?=.*[a-z])(?=.*[A-Z])(?=.*\d).+\$</code>		
<p>Specifies the regular expression that must be matched by the user password. By default, the password must contain at least one uppercase character, one lowercase character, and one digit. Set the value to an empty string to disable this requirement.</p>				
<p>9.2.7</p> <ul style="list-style-type: none"> • Name: User password: required special characters • Parameter: <code>user_password_special_characters</code> 	String		0	
<p>Specifies the list of special characters from which at least one must be used in the user password. By default, no special characters are required.</p>				

Table 42. Server settings

The table consists of 6 columns and 31 rows. For each row, there are two levels. The first level contains a column with units, default value, minimum value, and maximum value. The second level contains a description of the parameter.

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
	 Note: Provide the characters one after another without any separators. For example: ?!&%. If you add a space, it is also treated as a special character.			
9.2.3 <ul style="list-style-type: none"> • Name: Validate VM manager data • Parameter: <code>validateVMManagersData</code> 	Boolean (true/false)	True		
	Specifies whether the XML schema of files with VM manager data is validated by the BigFix Inventory server.			

Settings of the main VM Manager Tool

The following parameters can be configured from the server only if you are using [basic VM management \(on page cccxlvii\)](#). If you are using [advanced VM management \(on page ccclxvii\)](#), you must edit the `vmmmainconf.properties` configuration file to change the parameters. For more information about changing the settings of an additional VM Manager Tool, see: [VM Manager Tool settings \(on page ccclxxxii\)](#).

Table 43. Settings of the main VM Manager Tool

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
9.2.7 <ul style="list-style-type: none"> • Name: Allow VM manager HTTP connection fallback • Parameter: <code>vmman_http_connection_allowed</code> 	Boolean (true/false)	True		
	Allows the VM manager to connect by using the HTTP protocol when the defined HTTPS connection is not available.			
9.2.7	Boolean (true/false)	False		

Table 43. Settings of the main VM Manager Tool

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
<ul style="list-style-type: none"> • Name: Attempt to fix VM manager connection parameters • Parameter: <code>vmman_fix_parameters_enabled</code> 	Attempts to fix the most common problems with VM manager connection parameters, such as adding default URL suffix, correcting the user name format, and testing the WinRM and PowerShell protocols for Hyper-V VM managers.			
<ul style="list-style-type: none"> • Name: Check VM managers for uniqueness • Parameter: <code>vmman_check_uniqueness_enabled</code> 	Boolean (true/false)	False		
Specifies whether unique VM managers are distinguished from duplicates.				
9.2.3 <ul style="list-style-type: none"> • Name: Collect serial numbers of VM manager hosts • Parameter: <code>vmman_collecting_host_serials_enabled</code> 	Boolean (true/false)	False		
Specifies whether serial numbers of VM manager hosts are collected for VM managers from which data is collected by the VM Manager Tool. Available for Hyper-V and VMware 5.0 and higher.				
To display information about serial numbers in the Server ID column on the reports, adjust the Server ID template that is specified by the managedServerTagTemplate (on page cdxvi) parameter.				
9.2.2 <ul style="list-style-type: none"> • Name: Collect virtual machines host names • Parameter: <code>vmman_collecting_host_names_enabled</code> 	Boolean (true/false)	False		
Specifies whether host names of virtual machines are collected by the VM Manager Tool and saved in its log files. Available for VMware only.				
You can enable this parameter to solve problems with duplicated UUIDs. If duplicates are found, the host names are written in the log files.				
<ul style="list-style-type: none"> • Name: Filter VM manager UUIDs • Parameter: <code>vmman_uuid_filtering_enabled</code> 	Boolean (true/false)	False		
Enables UUID filtering of VM managers. For more information, see: UUID-based virtual machine filtering (on page ccclxxxvi).				

Table 43. Settings of the main VM Manager Tool

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
<ul style="list-style-type: none"> • Name: Maximum number of subsequent login failures • Parameter: <code>vmman_max_subsequent_login_failures</code> 	Number	3	0	100
	Specifies the maximum number of failed attempts of logging in to the VM manager.			
9.2.4 <ul style="list-style-type: none"> • Name: Transfer VM manager credentials • Parameter: <code>vmman_transfer_credentials_to_server</code> 	Boolean (true/false)	True		
	Specifies whether VM manager credentials are transferred to the server. The option is enabled by default. If you disable the option, the credentials are not transferred and thus they are not displayed on the VM Managers panel when you edit the VM manager.			
9.2.4 <ul style="list-style-type: none"> • Name: Trust SSL certificates from all defined VM managers • Parameter: <code>vmman_trust_all_vm_managers_certificates</code> 	Boolean (true/false)	True		
	Specifies whether SSL certificates from all defined VM managers are trusted. If the option is disabled, certificates that are trusted must be placed in the truststore on the computer where the VM Manager Tool is installed. For more information, see: Configuring VM Manager Tool to accept trusted VM manager certificates (on page dccxlv) .			
<ul style="list-style-type: none"> • Name: VM manager connection timeout • Parameter: <code>vmman_connection_time_out</code> 	Seconds	90	10	3600 (1 hour)
	Specifies the time after which the connection with a VM manager is ended.			
<ul style="list-style-type: none"> • Name: VM manager data transfer period • Parameter: <code>vmman_transfer_period</code> 	Minutes	720	30	10080 (1 week)
	Determines how often scan data is transferred to the agent to be uploaded to the server if subsequent scans have the same results.			
<ul style="list-style-type: none"> • Name: VM manager pooling interval • Parameter: <code>vmman_pooling_time_interval</code> 	Minutes	30	30	10080 (1 week)

Table 43. Settings of the main VM Manager Tool

(continued)

Name & Parameter	Type	Default	Minimum	Maximum
	Description			
<ul style="list-style-type: none"> • Name: VM manager thread pool size • Parameter: <code>vmman_thread_pool_size</code> 	Number	10	1	50
Specifies the number of threads in the thread pool that is used for connections to VM managers.				



Restriction: In case of subcapacity reporting, it is necessary to capture mobility of VMs. To ensure that it is captured, data from VM managers must be collected every 30 minutes. You can change the frequency of collecting data only if you use BigFix Inventory for purposes other than subcapacity reporting.

Administering the database

Perform administering tasks on your database to keep it in good condition, protect your data, or configure additional settings.

Backing up and restoring the database

Perform regular backups of the data that is stored in the database, and restore it when needed to prevent data loss.

To maintain data consistency during data recovery, it is recommended to perform backup at least once per week. When you plan regular backups, you must consider the dynamics of your environment, for example, infrastructure changes and software installations. It is a good practice to back up the database before you update the software catalog, upgrade the server, or introduce significant changes to the middleware or the operating system. The number of backups that you keep depends on the storage space available. It is recommended to ensure that the backup that you store allows the restoration of your environment to the last functional state. Thus, if you introduce several consecutive changes in the application, the database, or the operating system, ensure that you perform backup after each of those change.

Consider the start time and duration of the import when you plan the backup schedule to avoid performance impact on the data import.

Backing up the DB2® database

You should back up the database on a regular basis to protect your data. Each backup is the copy of the entire database that can be restored in the case of corruption or malfunction.

- Users

To perform the backup, log in as the database instance owner.

- Space requirements

Ensure that enough storage space is available in the target directory. The backup file size depends on the number of computers that report to BigFix Inventory, and the amount of historical data that is stored in the database. Therefore, you cannot estimate the exact size of the backup file.

The backup file size can reach:

- 2GB or more in the environments that consist of less than 100 computers.
- 50GB for 30000 computers.
- 75GB for 100000 computers.

You can use compression to significantly reduce the backup size.

Related information

[DB2 data recovery](#)

Creating offline backups

You should back up the database on a regular basis to protect your data. Each backup is the copy of the entire database that can be restored in the case of corruption or malfunction. Choose offline backups if you can afford to deactivate the database during each backup.

During an offline backup, the entire database is copied to a file. Since the database must be deactivated during each backup (no active transactions), the backup file contains all transactions performed on the database so far. Therefore, the backup file is not dependent on transaction logs, and the database can use the default circular logging, in which transaction logs are overwritten with new transactions, which saves the disk space. To restore the database, you only need the backup file.

1. Stop the BigFix Inventory server.
2. Terminate the connections and deactivate the database.



Note: `TEMADB` is the default database name. If you are unsure whether it applies to your database, see:

[Checking the database name \(on page dcclxx\)](#).

```
db2 terminate
db2 deactivate db TEMADB
```

```
DB20000I The TERMINATE command completed successfully.
DB20000I The DEACTIVATE DATABASE command completed successfully.
```

3. Make a full offline backup. The entire database is copied to a backup file in the specified location.

```
db2 backup database TEMADB to <location>
```

```
Backup successful. The timestamp for this backup image is : xxxxxxxxxxxx
```

4. Verify that the backup file was correctly created.

```
db2ckbkp <backup file>
```

```
Image Verification Complete - successful.
```

5. Activate the database and restart the BigFix Inventory server to restore its connection to the database.

```
db2 activate db TEMADB
/etc/init.d/BFIserver restart
```

6. If DB2 is shared with the BigFix server, restart the BigFix server to restore its connection to the database.

For more information, see: [Starting and stopping the BigFix server](#).

Creating online backups

Choose online backups if you back up your database frequently and cannot afford to deactivate it for each backup. Online backups allow you to maintain all connections to the database, but they require more disk space to store the history of transaction logs.

During an online backup, the entire database is copied to a file, but it might not contain transactions that were being performed on the database while the backup was being created. Therefore, the backup file must always be complemented with transaction logs to restore the database to the consistent state. Transaction logs are backed up automatically after enabling archive logging, and must always be stored in a safe location together with the backup file. After you restore the database from the backup file, you add the archive logs, and the database is complete. Enabling the archive logging for the first time requires that you deactivate the database.

1. Set up archive logging.



Note: TEMADB is the default database name. If you are unsure whether it applies to your database, see:

[Checking the database name \(on page dcclxx\)](#).

- a. Check the status of archive logging in your database. If the first archive method contains a value other than OFF, you can omit the following steps and make an online backup.

```
db2 get db cfg for TEMADB | grep LOGARCHMETH
```

```
First log archive method          (LOGARCHMETH1) = OFF
Second log archive method         (LOGARCHMETH2) = OFF
```

- b. Create a backup directory that will be used to store archive logs and online backups, and give it proper permissions.

```
mkdir /var/online_backup
chown db2inst1:db2iadml /var/online_backup
```

- c. Update the database configuration to enable archive logging and point to the new directory.

```
db2 update database configuration for TEMADB using LOGARCHMETH1 'disk:/var/online_backup'
```

```
DB20000I  The UPDATE DATABASE CONFIGURATION command completed successfully.
```

The status of archive logging settings also changed to a specific path:

```
First log archive method          (LOGARCHMETH1) = DISK:/var/online_backup/
```

- d. Terminate the connections and deactivate the database.

```
db2 force application all
db2 deactivate db TEMADB
```

- e. Make a full offline backup.

```
db2 backup database TEMADB to <location>
```

Creating an offline backup is a requirement after enabling the archive logging. Although you will not need to create offline backups later, it is a good practice to create them before major changes to the database. The archive logs can be applied in the same way both to offline and online backups to restore them to the most recent state. You can store this offline backup in a safe location as an extra precaution.

- f. Activate the database and restart the BigFix Inventory server to restore its connection to the database.

```
db2 activate db TEMADB
/etc/init.d/BFIServer restart
```

- g. If DB2 is shared with the BigFix server, restart the BigFix server to restore its connection to the database.

For more information, see: [Starting and stopping the BigFix server](#).

2. Make an online backup.

```
db2 backup database TEMADB online to /var/online_backup/ compress include logs
```

```
Backup successful. The timestamp for this backup image is : xxxxxxxxxxxx
```

Whenever you create an online backup, all archive logs created up to this point are no longer needed. To restore a complete database, you need a backup file and subsequent archive logs.

3. Verify that the backup file was correctly created.

```
db2ckbkp <online backup file>
```

```
Image Verification Complete - successful.
```

Restoring the DB2® database

If your database has been corrupted or you want to revert some changes, you can restore it to its previous state from one of the backup files. Depending on the backup procedure that you chose, use either offline or online backup files. In the latter case, you must also restore the archive logs to make the database consistent.

Restoring database from offline backup files

An offline backup file is the copy of the entire database and all transactions performed on it at the time when the backup was created. To restore the database, you only need this backup file.

1. Terminate the connections and deactivate the database.



Note: TEMADB is the default database name. If you are unsure whether it applies to your database, see: [Checking the database name \(on page dcclxx\)](#).

```
db2 terminate
```

```
db2 deactivate db TEMADB
```

2. Restore the database from the offline backup file.

```
db2 restore db TEMADB from <location> taken at <timestamp> replace existing
```

3. [Upload the software catalog \(on page cdxliv\)](#) to avoid problems with the data import. At this point, there might be some discrepancies between the server and the contents of the database, which will be resolved after the upload.

Restoring database from online backup files

An online backup file is also the copy of the entire database, but it might not include transactions that were being performed on the database while the backup was being created. Therefore, the backup file must be complemented with archive logs that contain the latest transactions. By choosing the right archive logs, you can either restore the database to the most recent state, or to one of the previous states.

Restoring to the most recent state

To restore the database to the most recent state, use one of your online backup files (preferably, the most recent one), and then apply to it all subsequent archive logs. If you store all archive logs in one directory, the correct ones will be applied automatically.

1. Terminate the connections and deactivate the database.



Note: `TEMADB` is the default database name. If you are unsure whether it applies to your database, see: [Checking the database name \(on page dcclxx\)](#).

```
db2 terminate
db2 deactivate db TEMADB
```

2. Restore the database from the online backup file.

```
db2 restore database TEMADB from <location> taken at <timestamp> replace existing
```

3. Restore all archive logs that were saved after the backup file was created.

```
db2 rollforward db TEMADB to end of logs and stop log path (/var/archive_logs)
```

4. Activate the database and restart the BigFix Inventory server to restore its connection to the database.

```
db2 activate db TEMADB
/etc/init.d/BFIserver restart
```

5. [Upload the software catalog \(on page cdxliv\)](#) to avoid problems with the data import. At this point, there might be some discrepancies between the server and the contents of the database, which will be resolved after the upload.

Restoring to one of the previous states

You can also restore the database to one of the previous states to avoid all subsequent transactions that might have corrupted it. In such a case, you do not apply all archive logs that you store, but only those that are bundled in the backup file. These bundled logs include only transactions from the exact state when the backup was created. You can extract them from the backup file.

1. Terminate the connections and deactivate the database.



Note: `TEMADB` is the default database name. If you are unsure whether it applies to your database, see: [Checking the database name \(on page dcclxx\)](#).

```
db2 terminate
db2 deactivate db TEMADB
```

2. Extract the transaction logs from the online backup file. This step is required, because you will not be able to activate the database unless some logs are applied to it.

```
db2 restore database TEMADB logs from <online backup file> LOGTARGET /var/extracted_logs/
```



Note: Ensure that the location for the extracted log files is different from the one in which you store all other archive logs.

3. Restore the database from the online backup file.

```
db2 restore database TEMADB from <location> taken at <timestamp> replace existing
```

4. Apply the extracted archive logs.

```
db2 rollforward database TEMADB to end of logs overflow log path (/var/extracted_logs/)
```

5. Activate the database and restart the BigFix Inventory server to restore its connection to the database.

```
db2 activate db TEMADB
/etc/init.d/BFIserver restart
```

6. [Upload the software catalog \(on page cdxliv\)](#) to avoid problems with the data import. At this point, there might be some discrepancies between the server and the contents of the database, which will be resolved after the upload.

Backing up the SQL Server database

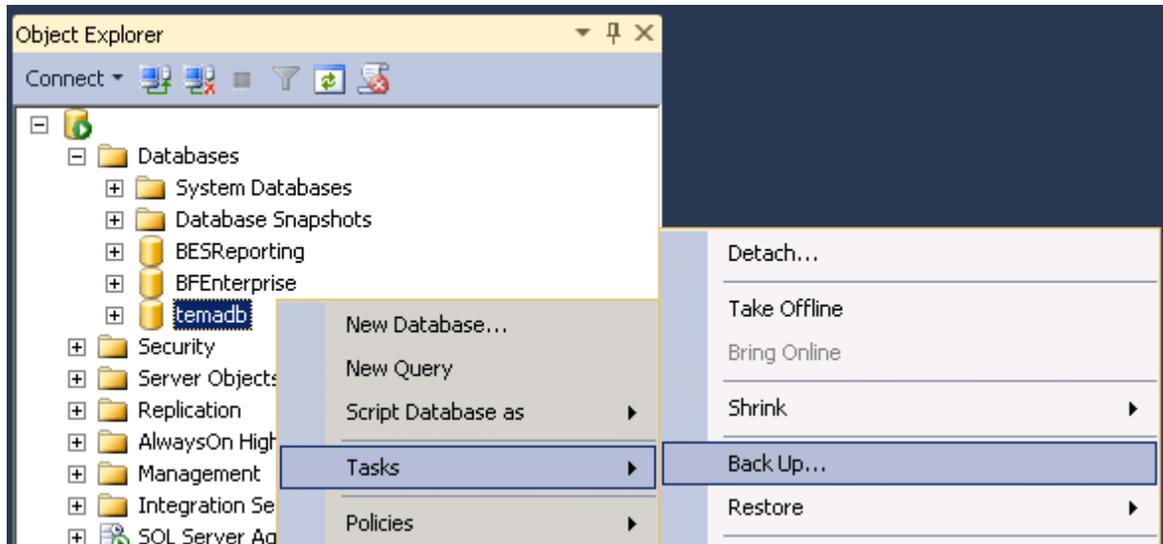
You can make a copy of your database by saving it to a backup file. If you want, you can then move the backup to another computer and restore it in a different BigFix Inventory instance.

- You can back up and restore the database only within one version of BigFix Inventory.
- BigFix Inventory and Microsoft™ SQL Server Management Studio must be installed.
- Ensure that enough storage space is available in the target directory. The backup file size depends on the number of computers that report to BigFix Inventory, and the amount of historical data that is stored in the database. Therefore, you cannot estimate the exact size of the backup file. In small environments the backup file size starts at 1 GB, and for 30000 computers the file reaches 40 GB. You can use compression to significantly reduce the backup file size.

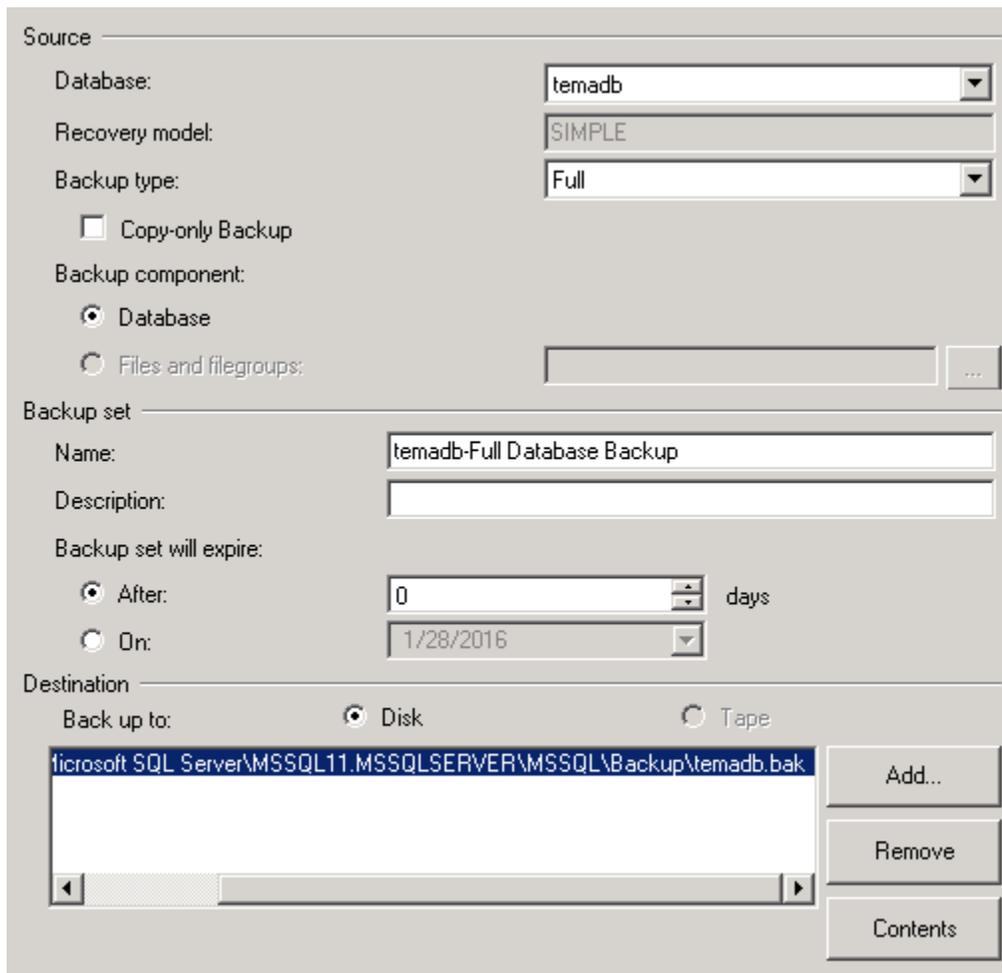
1. Log in to the computer that hosts the database that you want to back up.
2. Open Microsoft™ SQL Server Management Studio.
3. In the left navigation bar, expand **Databases**.
4. Right-click on the database that you want to back up and then click **Tasks > Back Up**.



Note: `TEMADB` is the default database name. If you are unsure whether it applies to your database, see: [Checking the database name \(on page dcclxx\)](#).



5. Review the details of the backup and then click **OK** to create the backup.



6. Click **OK**.

If the database was backed up successfully, you can find the **bak** file in the location that you specified in step 5.

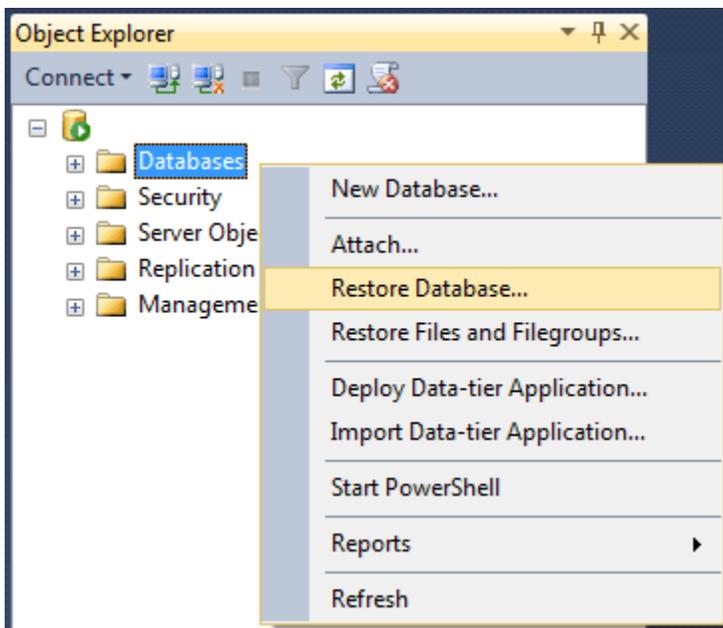
If you want to move the database to a different BigFix Inventory instance, copy the backup file to the target computer and then [restore the database \(on page cdxxxii\)](#).

Restoring the SQL Server database

If you encounter any problems with your database or if you want to move it between different instances of BigFix Inventory, you can use a backup file to restore the database.

- You can back up and restore the database only within one version of BigFix Inventory.
- Ensure that you are logged in to Microsoft™ SQL Server Management Studio as the user who created the **temadb** database. If you log in as a different user, the restoring will fail.
- BigFix Inventory and Microsoft™ SQL Server Management Studio must be installed.
- Stop the `BFIserver` service. Open the command prompt and run `net stop BFIserver`.

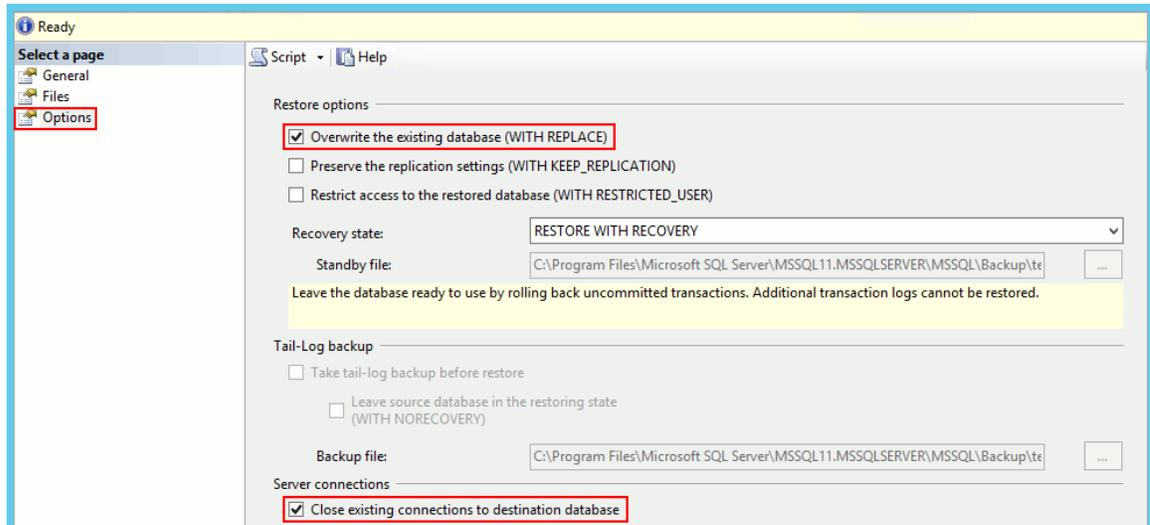
1. Log in to the computer on which you want to restore the database.
2. Open Microsoft™ SQL Server Management Studio.
3. In the left navigation bar, right-click on **Databases** and then click **Restore Database**.



4. In the **Source** section, select **Device** and click the button with three dots.



5. In the pop up window that opens, click **Add** and browse for your backup file. Click **OK**.
6. In the left navigation menu, click Options.
7. In the pane on the right select **Overwrite the existing database (WITH REPLACE)** and **Close existing connections to destination database**.



8. Click **OK**.

Upload the software catalog ([on page cdxliv](#)) to avoid problems with the data import. At this point, there might be some discrepancies between the server and the contents of the database, which will be resolved after the upload.

Creating a new database

To create a new database without reinstalling BigFix Inventory, you must modify the configuration file that stores information about the database. After you do so, you will be able to complete the initial configuration again, and create a new database.

1. [Stop the BigFix Inventory server \(on page cd\)](#).
2. Delete the current database.



Note: `TEMADB` is the default database name. If you are unsure whether it applies to your database, see: [Checking the database name \(on page dcclxx\)](#).

DB2:

- a. Log in as the DB2 instance owner, which by default is `db2inst1`.
- b. Run the following command to check the details of the existing database.

```
db2 list db directory
```

- c. Run the following commands to close the existing connections, deactivate the database, and delete it.

```
db2 deactivate db TEMADB
db2 drop db TEMADB
```

SQL Server:

- a. Log in to SQL Server Management Studio.
 - b. In the navigation tree, expand **Databases**.
 - c. Right-click on the *TEMADB* database, and then click **Delete**.
3. Go to `installation_dir/wlp/usr/servers/server1/config/`, and rename the `database.yml` file to `database.yml.bak`. This step resets the initial configuration and allows you to create a new database. You will also have to create a new administrative user and specify connections to the BigFix platform.
 4. [Start the BigFix Inventory server \(on page cccxcix\)](#).
 5. Go to `https://hostname:port/setup/database` to perform the initial configuration and create a new database. For more information, see: [Performing initial configuration on Windows \(on page cxlvi\)](#) or [Performing initial configuration on Linux \(on page clxxiii\)](#).
 6. After the configuration is completed, [upload the software catalog \(on page cdxliv\)](#) to avoid problems with the data import. At this point, there might be some discrepancies between the server and the contents of the database, which will be resolved after the upload.

Reinstalling BigFix Inventory with an existing database

You can reinstall BigFix Inventory and connect it to an existing database. This allows you to preserve all the data collected so far and avoid changing configuration settings.

1. Back up the following files from the BigFix Inventory installation directory.


```
installation_dir/wlp/usr/servers/server1/logs/imports/
installation_dir/wlp/usr/servers/server1/data/sam/
installation_dir/wlp/usr/servers/server1/apps/tema.war/WEB-INF/domains/sam/
config/file_names*.txt
```
2. Uninstall the BigFix Inventory server. For more information, see: [Uninstalling on Windows \(on page cli\)](#) or [Uninstalling on Linux \(on page clxxvii\)](#).
3. Install the BigFix Inventory server. For more information, see: [Installing the server on Windows \(on page cxli\)](#) or [Installing the server on Linux \(on page clxxviii\)](#).
4. Before you complete the initial configuration, copy the files that you backed up in step 1 into the new installation directory.
5. Perform the initial configuration and point to the existing database instance. It is enough to specify the right host name and leave the same database name. BigFix Inventory will connect to your existing database. For more information, see: [Performing initial configuration on Windows \(on page cxlvi\)](#) or [Performing initial configuration on Linux \(on page clxxiii\)](#).
6. [Upload the software catalog \(on page cdxliv\)](#) to avoid problems with the data import. At this point, there might be some discrepancies between the server and the contents of the database, which will be resolved after the upload.

Moving the BigFix Inventory database

To move the BigFix Inventory database to a different computer, back up the database, move it to the target computer, and point to that computer in the `server.xml` and `database.yml` files.

1. Stop the BigFix Inventory server (*on page cd*).
2. Back up the BigFix Inventory database.
 - Backing up the DB2 database (*on page cdxxxv*).
 - Backing up the SQL Server database (*on page cdxxx*).
3. Move the database to the target computer and make it operational.
4. Back up the following files from the installation directory of the BigFix Inventory server.
 - `installation_directory/wlp/usr/servers/server1/server.xml`
 - `installation_directory/wlp/usr/servers/server1/config/database.yml`
5. Update the database location in the `server.xml` file. Provide the IP address or domain name of the computer on which the database is located.
 - DB2 database

```
<properties.db2.jcc
password="{aes}xxxxXXXxxxxXXXxxxxXXXxxxxXXXxxxxXXXxxxxXXX"
user="db2inst1"
serverName="localhost"
databaseName="temadb" />
```

- MSSQL database

```
<properties.microsoft.sqlserver
password="{aes}xxxxXXXxxxxXXXxxxxXXXxxxxXXXxxxxXXXxxxxXXX"
user="sa"
serverName="localhost"
databaseName="temadb" />
```

6. Update the database location in the `database.yml` file. Provide the IP address or domain name of the computer on which the database is located.

```
host: localhost
database: temadb
database_type: mssql
windows_authenticated: true
```

7. Start the BigFix Inventory server (*on page cd*).

Updating the database password

You can update the database user password when needed, for example if the password is changed and users cannot log in to BigFix Inventory.

Before you can run the security utility you must set the JAVA_HOME environment variable, for example `export`

```
JAVA_HOME=install_dir/jre/jre.
```

Important: Create a backup of **server.xml** and **database.yml** files before you make any change.

1. Log in to the server where BigFix Inventory is installed.
2. Run the following command and enter the new password for the database user.

```
install_dir/wlp/bin/securityUtility encode
```



Note: If you use AES encryption, add the following parameter to the command above: `--encoding=aes`.

The password is returned as an encrypted string.

3. Edit the following configuration files, and enter the new encrypted password.

- `install_dir/wlp/usr/servers/server1/server.xml`

Search the database entry "**properties.db2.jcc**" for **DB2** and "**properties.microsoft.sqlserver**" for **SQL Server** and then change the password for the database entry.

Copy the new encrypted password to **password** attribute, for example:

- in case of 'xor' encoding : `password="{xor}fgspGzAWPTApFzk="`
- in case of 'aes' encoding : `password="{aes}fgspGzAWPTApFzk="`

- `install_dir/wlp/usr/servers/server1/config/database.yml`

Copy the new encrypted password to **encrypted_password**.

4. If you use the same user for the BigFix database, you must also enter the new password in the following locations:
 - a. Update the database password in the BigFix server. For more information, see: [Changing the database password](#).
 - b. In BigFix Inventory, click **Management > Data Sources**. Enter the unencrypted password for the database user and click **Save**.
5. Restart the BigFix Inventory server.

Related information

[Starting the server \(on page cccxcix\)](#)

Removing inactive computers

When you decommission a computer in your infrastructure, it continues to report to BigFix Inventory. The data is still collected from this computer, and displayed on the metric reports. To ensure that all infrastructure changes are reflected in BigFix Inventory, you should remove each decommissioned computer from BigFix. If you cannot remove this computer from BigFix for a valid reason, you can use an alternative solution, and decommission the computer in BigFix Inventory.

Removing inactive computers from BigFix

When you decommission a computer in your infrastructure, you need to remove this computer from BigFix to reflect the change in the BigFix Inventory reports.

The computer that is decommissioned in your infrastructure continues to report to BigFix Inventory. The software is listed in the metric reports, and included in license calculation. To make sure that the components that are installed on this computer are not counted, remove the computer from BigFix.

1. Remove the decommissioned computers from the database.
 - Log in to BigFix console. To remove the computers manually, follow the [Removing Computers from the database](#) instruction.
 - To automatically detect, and remove the inactive computers from the database, use the BigFix Administration Tool. For more information, see: [Removing computers with BES Computer Remover](#) and [BES Computer Remover utility](#).
2. Run a data import.

The computer that you removed from the database is no longer listed in the Computers report. The metric reports that are generated for the period before the removal, still list the software that was installed on the computer. The reports that collect the data for the period after the computer was removed, do not include the software that was installed on the computer.

9.2.8 Automatically decommissioning inactive computers

9.2.8 Available from 9.2.8. When infrastructure management policy does not allow the deletion of decommissioned BigFix clients in BigFix, you can enable automatic decommissioning of inactive computers directly in BigFix Inventory. By decommissioning the inactive computers you ensure that they do not have negative impact on metric calculation.



Important: To ensure consistency between the BigFix platform and the BigFix applications, it is recommended to manage the entire infrastructure in BigFix. Decommissioning computers in BigFix Inventory is an alternative solution allowed only when decommissioning in BigFix is not possible.

If a computer stops reporting to BigFix, it is considered inactive. Configure BigFix Inventory to automatically decommission computers that are inactive for a specified period.



Note: This solution is not supported on the computers with disconnected scanner. For more information, see: [Decommissioning computers with disconnected scanners \(on page cclvi\)](#).

1. Log in to BigFix Inventory.
2. Go to **Management > Advanced Server Settings**.
3. To enable the feature, set the `decommission_inactive_computers` parameter to `true`, and click **Save**.
4. To set the inactivity period after which the computers are automatically decommissioned, configure the value of `max_data_visibility_for_inactive_computers` parameter to the required number of days, and click **Save**. By default, the inactivity period is set to 90 days.

The reports display only the data for the period before the date when the computer was decommissioned.



Note: When a decommissioned computer reactivates and starts reporting back to BigFix, the report data is updated.

To disable automatic decommissioning of inactive computers, go to **Advanced Server Settings**, select the `decommission_inactive_computers` parameter, and set it to `False`.

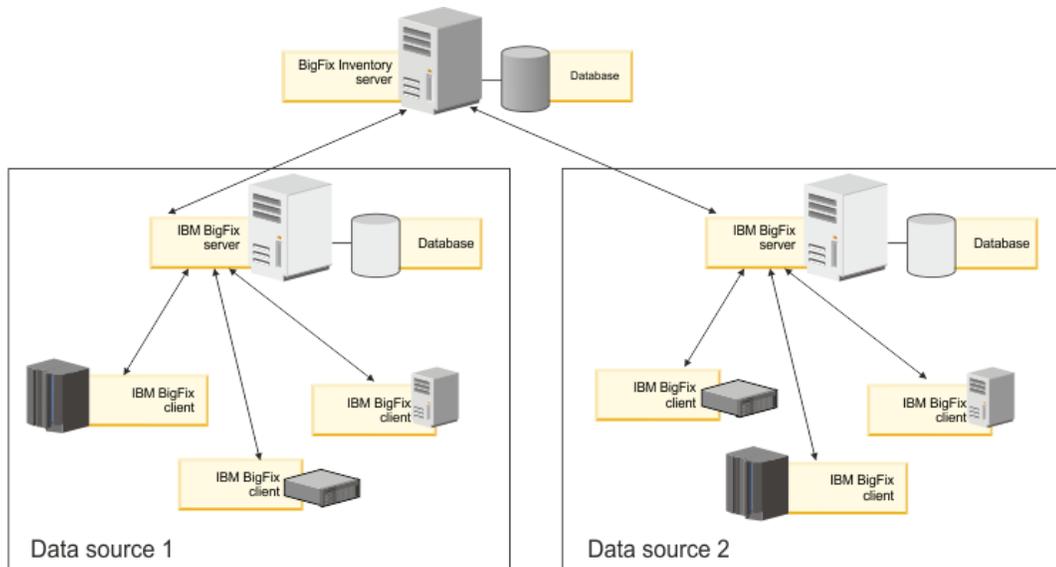
Managing data sources

Computers in an organization can be divided into organizational units, and can be monitored by multiple BigFix infrastructures. Each infrastructure is referred to as a data source. Data from multiple data sources can be imported to one instance of BigFix Inventory.

Definition of a data source

A *data source* can be viewed as a BigFix server and the clients that report to that server. Thus, a data source is the infrastructure from which BigFix Inventory imports the scan data. Organizations can define multiple data sources that mirror various types of organizational units such as departments or countries in which the organizations have their branches.

A *disconnected data source* is specific for environments that are monitored by disconnected scanners. It can be defined as a directory on the BigFix Inventory server that provides the scan data during the import. The disconnected data source cannot be the only data source that is set up in BigFix Inventory.



Hardware requirements

Sum up the number of the BigFix clients that report to BigFix Inventory from all data sources, and ensure that the computer where BigFix Inventory is installed meets the hardware requirements specific to the environment size. You can designate a physical server or virtual machine that is capable of handling large data imports (ETL) and has the capability to process a large amount of data. For more information, see:

- BigFix Inventory [hardware requirements \(on page cvii\)](#).
- [Tuning performance \(on page dcccxxxi\)](#)

Import of data from multiple data sources

During the import of data to BigFix Inventory, connectivity to all data sources is checked. If any of the data sources is not reachable, the entire import fails. Ensure that all data sources are reachable during the import.

Adding a data source

To collect data from multiple BigFix infrastructures and report it to a single instance of BigFix Inventory, add multiple data sources to BigFix Inventory.

1. In the top navigation bar, click **Management > Data Sources**, and then click **New**.
2. Provide a unique name for the data source.
3. To automatically enable scans that collect data from the computers in your infrastructure, select **Enable default scan schedule for this data source**.

If you enable the default scan schedule, actions that are needed to collect data from the computers in your infrastructure are automatically started on the BigFix server. This option is advised for environments with up to a few thousand computers. For larger environments, divide the computers into groups, and then manually

set up scan schedule for each group to avoid performance issues. For more information about the default and manual scan schedule, see: [Setting up scans to discover software and hardware inventory \(on page cxciii\)](#).

4. Select the database type.

- If you choose DB2, specify the host, port, database name, and credentials of the user that can access the BigFix server database.



Important: Ensure that the DB2 user has the following permissions. These permissions apply only if the databases were installed with default settings, and all customizations and hardening configurations were consulted with BigFix support.

- For the BigFix database (BFENT): **DBAUTH**
- For the Web Reports database (BESREPOR): **DATAACCESS**

- If you choose SQL Server, specify the host, database name, and credentials of the user that can access the BigFix server database.



Important:

- Ensure that the MS SQL Server user has the following permissions. These permissions apply only if the databases were installed with default settings and all customizations and hardening configurations were consulted with BigFix support.
 - For the BigFix database (BFEnterprise): `CREATE FUNCTION, CREATE SCHEMA, CREATE TABLE, CREATE VIEW, EXECUTE, SELECT`
 - For the Web Reports database (BESReporting): `SELECT`
- Ensure that the MS SQL Server user has the appropriate role to create the BigFix Inventory database.
 - If you create a new database during initial configuration in BigFix Inventory, the user must have the `sysadmin` role in MS SQL Server.
 - If you create an MS SQL database manually before the configuration, make sure that the database is empty and use the `SQL_Latin1_General_CP1_CS_AS` collation. The user for such database must have the `db_owner` role in MS SQL Server.
 - If you are using MS SQL Server 2012 and you chose Local System account as the service owner during the installation, give the `dbcreator` or `sysadmin` role to the `NT AUTHORITY\SYSTEM` user in MS SQL Server.

- If you choose Disconnected data source, see: [Adding a disconnected data source. \(on page cclxxiii\)](#)

5. Provide credentials of the Console Operator that you created while installing BigFix (by default, *IEMAdmin*).

6. **Optional:** If the BigFix and BigFix Inventory servers are in separated networks, the automatic address lookup might return incorrect address. To disable the lookup, select **Disable automatic address lookup**, and specify the address manually. Then, configure additional environment variables on the BigFix Inventory server. For more information, see: [Configuring servers in separate networks \(on page clxxx\)](#).

7. **Optional:** You can additionally configure connection between your BigFix Inventory data source and the Web Reports database to allow the Web Reports users to access BigFix Inventory. Specify the database type,

host name, database name, and credentials of the Web Reports database user. For more information, see: [Integrating users with Web Reports \(on page dccx\)](#).

8. Click **Create**.

Editing a data source

If you move the BigFix database to a different computer, or change credentials of the Console Operator that is used for connecting with BigFix, edit the parameters to maintain data source connectivity.

1. In the top navigation bar, click **Management > Data Sources**.
2. Click the data source that you want to edit, and change connection parameters.
 - To edit connection parameters to the BigFix database, edit fields in the first column.
 - To edit credentials of the Console Operator that is used for connecting with BigFix, edit fields in the second column.
 - To edit connection parameters to the Web Reports database, edit fields in the last column.
3. During editing, passwords for the BigFix database, Console Operator, and the Web Reports database are cleared. Re-enter the passwords, and click **Save**.

Deleting a data source

If you no longer want to import data from a particular BigFix infrastructure to BigFix Inventory, remove this data source.

1. In the top navigation bar, click **Management > Data Sources**.
2. Click the data source that you want to delete, and click **Delete**.

Data for computers that reported to this data source is deleted from BigFix Inventory.

Importing software scan data

The inventory results are stored on your BigFix server. To import software scan data, the software catalog and other settings that changed since the last update, you must extract the data from BigFix server and load it into BigFix Inventory.



You must have the Manage Imports permission to perform this task.

The import of data is an Extract, Transform, Load (ETL) process.

- During the Extract stage, data is extracted from the BigFix server. It includes information about the infrastructure, installed clients, and detected software. ETL also checks whether a new software catalog is available, gathers information about the software scan and files that are present on the computers. It also collects data from VM managers.
- During the Transform stage, the extracted data is transformed to a single format that can be loaded to the BigFix Inventory database. This stage also involves matching scan data with the software catalog, calculating

processor value units (PVUs), processing the capacity scan, and converting information that is contained in the XML files.

- During the Load stage, the data that was extracted and transformed is loaded into the BigFix Inventory database and can be used by BigFix Inventory.

1. In the navigation bar, click **Management > Data Imports**.
2. To schedule regular imports, select **Enabled**, specify the number of daily imports and their hours, and click **Save**.

Import Settings

Enabled

Imports per day (times specified in UTC +02:00)

Information about data imports and their status is displayed in the import history.

Status	Start Time	User Name	Duration	Download
✓	08/29/2017 08:12 AM	admin	00:02:56	↓
✓	08/28/2017 02:00 PM	Scheduled	00:01:59	↓
✓	08/28/2017 01:46 PM	admin	00:04:55	↓
✗	08/28/2017 12:32 PM	admin	01:00:38	↓

If you want to quickly discover recent changes to software inventory without having to wait for the complete data import, you can schedule imports of raw data. Such imports are much quicker but the imported data is not visible in the user interface. It can be retrieved only by using REST API. For more information, see: [Importing raw scan data \(on page cclxiii\)](#).

9.2.11 If the imports of data are failing, you can enable partial imports. For more information, see: [Enabling partial imports \(on page cclxiv\)](#).

Importing raw scan data

9.2.2 Available from 9.2.2. You can schedule an import of raw scan data. Such an import is much quicker, because it only imports raw scan results, and omits successive steps, such as aggregation and license calculations. The imported raw data, however, is not updated in the user interface, and can only be retrieved by using REST API. This type of import is used to quickly discover recent changes to software inventory without having to wait for the complete data import.

- The Raw Data Only import does not replace the complete data import, which must be run regularly to process and aggregate the data.
- The Raw Data Only import can be run only as a scheduled import. When you click **Import Now**, you always run a complete data import.
- To run the Raw Data Only import only once, use REST API. For more information, see: [Running data imports \(on page dcccclxxviii\)](#).

1. To enable this feature, go to <https://hostname:port/management/feature>, and select **Enable the Raw Data Only import mode**.
2. In the navigation bar, click **Management > Data Imports**.
3. To schedule regular imports, select **Enabled**, and specify the number of daily imports and their hours.
4. Select **Raw Data Only** next to each import that you want to use for importing raw scan results.

Import Settings

Enabled

Imports per day: (times specified in UTC +01:00)

06:00AM	<input type="checkbox"/>	Raw Data Only
02:00PM	<input checked="" type="checkbox"/>	Raw Data Only
04:00PM	<input checked="" type="checkbox"/>	Raw Data Only

5. Click **Save**.

Information about completed data imports is displayed in the import history. Each entry has an icon that indicates the type of the import, either *Complete* or *Raw Data Only*.

Status	Type	Start Time	User Name	Duration	Download
✓		09/08/2015 07:3...	Administrator	1:27:29	
✓		09/08/2015 02:1...	Scheduled	0:00:22	

When the Raw Data Only import finishes, retrieve the imported raw data by using REST API. For more information, see [REST API for retrieving raw scan results \(on page dcccclxxvii\)](#).

Managing software inventory and metric utilization

Learn how to classify the discovered software so that reports in BigFix Inventory reflect your entitlements and properly show utilization of license metrics by particular products.

Software catalogs

To correctly identify components of software products in your infrastructure, ensure that your software catalog is always up-to-date. Periodically import a new catalog that contains the most recent software products. Also, manually add products that are installed in your infrastructure but do not have corresponding entries in the catalog.

There are two types of the software catalog:

- Catalog provided by BigFix that contains signatures for BigFix products as well as signatures for a range of products from vendors such as Microsoft or Oracle.
- Custom catalog that contains signatures created manually to discover products that are not contained in the catalog provided by BigFix.

These two catalogs are fully independent. If you create a custom signature to discover a product and then a signature for that product is delivered in the BigFix catalog, the two signatures are not merged. If the same product is discovered by both signatures, the number of discovered instances might be higher than the number of instances that are installed. In such case, delete the custom signature. For more information, see: [Catalog problems \(on page dcclx\)](#).

Updating the software catalog to ensure accuracy of software discovery and reporting

The software catalog is automatically updated during every upgrade of . However, as an exception, you can manually update the software catalog for troubleshooting purposes.

- **9.2.11** Starting from application update 9.2.11, the Software Catalog Update task automatically triggers the following actions:
 - The catalog is downloaded to the following directory:
 - **Linux** `/opt/ibm/BFI/wlp/usr/servers/server1/data/sam/catalog/BFI`
 - **Windows** `C:\Program Files\IBM\BFI\wlp\usr\servers\server1\data\sam\catalog\BFI`
 - The catalog is automatically uploaded to the BigFix Inventory server during the next import of data. No additional actions are required.
- For versions before 9.2.11, the software catalog is downloaded with the Software Catalog Update task but it still needs to be manually uploaded to the BigFix Inventory server according to this procedure.  You must have the Manage Uploads permission to perform this task.
- If the computer where the BigFix server is installed does not have the Internet access, ensure that the content of the BigFix Inventory site is up-to-date before you download the catalog. For more information, see: [Updating the fixlet site \(on page cccvii\)](#).

The software catalog is called `BFI_Catalog_version-timestamp.zip` and consists of the following files:

- Software catalog in canonical 2.0 format: `IBMSoftwareCatalog_canonical_2.0_form_date.xml`.
- **9.2.5** Software catalog in canonical 3.0 format: `IBMExtendedCatalog_canonical_3.0_form_date.xml`. The file consists of low priority applications; for example, open source software, freeware, and games. The content of the file will be expanded to include high priority business applications with time. The file is included in each software catalog, however, it is ignored by BigFix Inventory versions lower than 9.2.5.
- Charge unit data file that contains information about charge unit definitions, their relations with products, and some additional parameters: `charge_unit_parameters_dataversion_version.csv`.
- Part numbers file that contains information about part numbers that are used for software licensing purposes: `part_numbers_dataversion_version.csv`.
- Catalog metadata file that describes the software catalog: `catalog_package.properties`.
- **9.2.13** FlexPoint bundles file that contains information about products that can be assigned to each of the available FlexPoint Offerings: `flexpoint_bundles_dataversion_version.csv`
- **9.2.11** For versions 9.2.11 and higher
 1. Log in to the BigFix console.
 2. In the navigation bar, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
 3. In the upper-right pane, select **Software Catalog Update**, and then click **Take Action**. The action is applicable only on the computer on which the BigFix Inventory server is installed. Select the computer, and click **OK**.
If the task is not applicable on that computer, see: [Catalog problems \(on page dcllviii\)](#).
 4. Wait for the scheduled import of data or run it manually.
- For versions before 9.2.11
 1. Download the software catalog
 - a. Log in to the BigFix console.
 - b. In the navigation bar, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
 - c. In the upper-right pane, select **Software Catalog Update**, and choose the option to download the software catalog for BigFix Inventory version before 9.2.11.0.
 - d. Optional: Copy the file to the computer from which you access the BigFix Inventory web user interface.
 2. Upload the software catalog to BigFix Inventory.
 - a. Log in to BigFix Inventory.
 - b. In the navigation bar, click **Management > Catalog Upload**.
 - c. Click **Browse** and select the software catalog file.

- d. To upload the file, click **Upload**. The software catalog file and the charge unit data are listed in the **Upload and Import History** table. Their status is **Pending**.
- e. Wait for the scheduled import of data or run it manually.

During the import, scanner catalogs that are used for software discovery are created and automatically distributed to the computers in your infrastructure.

If the automatic distribution of scanner catalogs fails, computers on which the catalog was not updated have the **Outdated Catalog** status on the **Scan Health** widget. You must [manually update scanner catalogs \(on page cclxvii\)](#) on these computers. If you use the disconnected scanner, [distribute the downloaded software catalog \(on page cclv\)](#) to each computer.

Checking product components and their signatures in the software catalog

You can browse the software catalog to check what components belong to a particular product and what signatures cause that these components are detected.

1. In the navigation bar, click **Reports > Products & Metrics**.
2. Select a product. To narrow down the report, filter it by a product name or a part number.
3. Drill down through the version name and release name to the component level.



Note: If the Component Version column contains an asterisk (*), the listed software component is detected regardless of its version. The reported version of each instance of the component is based on the generic signature that pulls the component version from the file or package that caused detection.

4. To check the signature that allows the discovery of a given component click **Signatures**.

Expanding the custom software catalog

To ensure that all software items that are installed in your infrastructure are discovered, create custom signatures. You can create them from scratch or directly from the Package Data and Scanned File Data reports. You can also create extended signatures to discover software that is not detected by using file or package signatures.

The custom catalog is fully independent from the catalog that is provided by IBM. If you create a custom signature to discover a product and then a signature for that product is delivered in the BigFix catalog, the two signatures are not merged. If the same product is discovered by both signatures, the number of discovered instances might be higher than the number of instances that are installed. In such case, delete the custom signature. For more information, see: [Catalog problems \(on page dcclx\)](#).

Before you create custom signatures, check the [signature community](#). It is a place where you can find custom signatures that are created by members of the community. You can download such signatures and import them to BigFix Inventory. After you create your custom signatures, consider contributing the community.

Catalog customization process

The process of adding custom signatures to the software catalog requires careful planning and consideration. It starts with identifying products that you expect to discover but are not reported or are missing from the software catalog. Next, you should organize your work to ensure that the most important products are given the highest priority. Then, you determine what files or packages can be used to detect the software and report its usage. Finally, you can create the signatures.



Note: BigFix Inventory is intended for software inventory and license management. It is suggested that you extend the software catalog for those purposes only.

Process input

The input to the catalog customization process is a list of software products that you want to be discovered. The list can be based on the general knowledge of your environment and the procurement data. As the discovery, monitoring, and license management of every software product require extra effort, target products on the basis of license expenses. Focus on a few selected products and adopt an iterative approach instead of trying to work on all products simultaneously.

Required skills

A set of skills that are helpful during the creation of the custom catalog content includes:

- General knowledge of the software products that are used in your company, their architecture, and licensing models
- General knowledge of the operating systems on which the software is installed

Step 1: Prioritize the work

Before you start adding custom signatures to the software catalog, prioritize your work. Start with products that are critical from the business perspective. Then, iteratively continue with the remaining products until all software that is used in your company is discovered.

To optimize your work on the custom catalog content, adopt an iterative approach. Start by grouping all software products that are used in your company into software families and focus on one software family per iteration. Analyze and create software signatures for all products that belong to the family on which you are working. If a software family is too extensive, split it into smaller groups, for example based on editions or distributions. The iterative approach allows for continuously extending the catalog content and verifying its correctness at the same time. Thus, you can gain more benefits in less time.

Step 2: Analyze your software

Extend the software catalog only with information that is related to products that are used in your company but are not reported or do not exist in the catalog. Consult the product administrators and available documentation to gather information that is related to the product architecture and licensing model.

To gather basic information about a software product, work with the administrator who is responsible for its installation and has the knowledge of its architecture and licensing model. Try to establish a reference installation. Investigate the architecture of the product to obtain the following details:

- Separately installable pieces (components) that constitute the product and are important for discovery or licensing purposes
- Version number of each component
- Ways to check whether the component is installed
- Platforms on which the component can be installed
- Licensing model

If the product administrator is unable to provide the information necessary for creating the licensing model, consult the available documentation. It can be found on the installation media, in the product installation path, license agreements, installation instructions, or online resources.

Step 3: Review information about files and packages

To identify candidates for software signatures, review information about files and packages that exist on the computer where a particular software product is installed.

Candidates for file and package signatures

Most efficient signatures are file-based signatures with constant size that is different for each release or signatures that are based on package data where a wildcard (*) can be used. In both cases, the file and package data should be removed during the product uninstallation or changed during the product upgrade or downgrade.

When you are looking for candidates for such signatures, always check:

- Whether any product documentation describes methods for determining the product edition and release based on files or package data
- Whether the product name and version can be unambiguously determined, for example, by the file name
- Whether there are any component-specific files:
 - Executable files whose version, part of version, or size is specific to the particular release
 - Files whose name or its part is specific to the particular release
 - Files whose content defines the product name and version
 - Other files with constant size, for example, a graphic that contains the product release number
 - Libraries with version or constant size
- Whether there are application-specific packages

Do not create signatures that are based on:

- Shared or external libraries
- Files that can be used by another product
- File names that are commonly used, for example, `readme.txt`

Available reports

Use one of the following reports to identify candidates for signatures.

Package Data

The report provides information about the packages that are installed on the computers in your infrastructure. To find a package that can be used as a signature, filter the report to the data from the computer on which the particular software product is installed.

Scanned File Data

The report provides information about files that were detected on the computers in your infrastructure. The files with any file extensions can be used to create software signatures.



Note: To find a file or a package that can be used as a signature, filter the report to the data from the computer on which the particular software product is installed. To further narrow down the results, specify the whole or part of the path to the directory where the software is installed. Then, look at the following columns.

- **Recognized** - the column shows whether the file was recognized as part of an existing signature.
- **Caused Detection** - the column shows whether the file contained enough information to cause detection of the related software. This column is not enabled by default. For more information, see: [Report columns \(on page dcxxx\)](#).

If the value in both columns is No, the file is not used for software detection and you can use it to create a new software signature.

Unrecognized Files

The report creates a ranking of files that are most commonly encountered in your computer infrastructure but do not produce matches for any signature. For more information, see: [Available reports \(on page dcxvi\)](#).

Step 4: Create component signatures

Use the following guidelines to avoid common mistakes and to create the most accurate signatures that better recognize software.

To obtain the best results, try to create signatures that meet the following requirements:

- The signature is based on the main executable file of the software.
- The signature combines the main executable file of the software with package data if it is available.
- The signature discovers only one release. If it discovers release 7.2, it cannot be valid for releases 7.1 or 7.3.
- The signature is generic and discovers a particular release and all its mod-releases and fix packs. To ensure that the signature is generic:

- Use a wildcard (*) in the package version, for example: 7.2.*

If you use a dot (.) in the package version, the dot is also matched with a dash (-). For example, if you specify the package version as 4.9.3.1*, version 4.9.3-1.el1 is also matched.

If you specify a version without the wildcard, the signature also matches a package that contains a comma in the version number. For example, if the package version is **11.10.1,REV=2005.01.21.15.53**, signatures that have the following package version specified match it:

- 11.10.1
- 11.10.1,REV=2005.01.21.15.53

Signatures that have the following package version specified do not match it:

- 11.10
 - 11.10.1,REV=2005.02.14.12.32
- Specify the version of the file limited to the release number. For example, 11.10. Avoid using a complete file version such as 10.11.2 or file size. These values are likely to change with every patch.
 - If a signature combines a file and package data and neither of them has a version, accept a file with any size.

 **Tip:** If you are unable to cover all fix packs and mod-releases, combine separate conditions for each of them.

Version of the software that is provided in the package and in the file might sometimes be inaccurate and might not match the version that is displayed in BigFix Inventory. If such a discrepancy occurs, ensure that the proper version is defined in the software catalog. Consider creating more specific rules for handling issues that are related to software versioning.

Creating signatures

You can use the built-in functionality to extend your custom catalog with the software products that are installed in your infrastructure but are not present in the HCL® catalog.

 You must have the Manage Catalogs permission to perform this task.

1. In the top navigation bar, click **Management > Catalog Customization**, and then click **New**.
2. In the **Create Catalog Entry** window, provide the publisher name, product name, and release number.

Create Catalog Entry

Publisher Name*	<input type="text" value="Software Publisher"/>
Product Name*	<input type="text" value="Product"/>
Release*	<input type="text" value="2.1"/>

A component will be created based on the release information. It represents a discoverable software item that is displayed on reports.

3. Decide whether the product should be detected with a package or file signature, or an extended signature.

- For a package or file signature, select a condition or conditions that must be fulfilled for the software item to be discovered.
 - To create an installation package signature, select the first condition, and specify the name of the package, its version, and vendor.
 - To create a file signature, select the second condition, and specify the name of the file and its size or version. If neither the file size nor the version is relevant, select **Any size or version**.

 **Tip:** If you select one condition and specify multiple files or installation packages, detection of any of them causes that the software item is discovered. If you select both conditions, at least one file and one installation package must be detected for the software item to be discovered.

Report release when the following conditions are fulfilled. If you specify two conditions, they both must be accomplished.

The following installation package is found:

Name*	Version	Vendor	
<input type="text" value="Package"/>	<input type="text" value="1"/>	<input type="text"/>	<input type="button" value="Remove"/>
Name*	Version	Vendor	
<input type="text" value="Package_1"/>	<input type="text" value="2"/>	<input type="text" value="Vendor"/>	<input type="button" value="Remove"/> <input type="button" value="Add"/>

The following file is found:

Name*	With	
<input type="text" value="file.exe"/>	<input type="text" value="Any size or version"/> ▼	<input type="button" value="Remove"/> <input type="button" value="Add"/>

If you select both conditions, at least one installation package and one file must be detected to report the release.

Important:

- If you add a rule with common executable files or other popular file types, it is enough to run the data import once and software inventory is available in the web user interface after the import. The popular file extension list includes the following file types:

`*.exe, *.sys, *.com, *.ear, *.ocx, *.sh, *.bin, *.pl, *.ear, *.SH, *.BIN, *.PL, and *.EAR.`

- If you add a rule with non-standard file extensions, wait until all the steps in the catalog data flow complete or perform those steps yourself. For more information, see: [Software discovery process after customizing the catalog \(on page cdxcviii\)](#).
- File rules for both common executable files and non-standard extensions are only supported on Windows™ endpoints if version is specified.

- Choose an extended signature to use various data elements such as file names, registry keys, CPU architectures, or specified strings to discover software that otherwise might not be discovered with file or package signatures.

- a. Open an XML editor or text editor that has support for the XML language, for example Notepad ++, and create the extended signature in the XML format. The signature length must not exceed 4000 characters.



Note: The supported version of the XML signature schema is 2.7.

- b. In the XML editor, select and copy the elements of the extended XML signature file.
- c. In the **Signature Content** text area, paste the contents of the extended signature that you created in step a.



4. To save the catalog entry, click **Submit**.

You added a software signature to your custom catalog, and it can now be used to detect software or its usage. Although you added a new signature, the custom content version of the software catalog does not have to change. The version is updated only for signatures that are based on non-standard file extensions. For more information, see [Software discovery process \(on page cdxcviii\)](#).

You can edit and delete a custom catalog entry on the **Catalog Customization** panel.

To edit a custom catalog entry, select it, edit a chosen element and confirm the change by clicking **Save**. When you edit the name of a publisher, product, version or release you cannot change it to a name that already exists in the custom catalog.

To delete a custom entry, select it and click **Delete**. You can delete an entire hierarchy, that is a publisher, software product, version, release, component, and signature simultaneously. However, you cannot delete multiple entries, for example multiple publishers, at the same time.

9.2.2 Extended signature templates

9.2.2 Available from 9.2.2. Extended signatures are flexible signatures that are used for software recognition.

They consist of a number of exported variables and the declaration of the operations that are required to assign a value to these variables.

Operands and operators

Every extended signature can be viewed as a set of processing instructions that are used by software scans to detect if a software item is installed on the target computer system. These processing expressions consist of two types of components: *operands* and *operators*. Operands are the objects that are manipulated, and the operators as the

symbols that represent specific actions. Operands and operators can be either scalar (consisting of only one value) or vector (consisting of multiple values).

The schema version that is supported by BigFix Inventory is 2.7.

Extended signature templates

The following section provides extended signature templates to help you create a new extended signature and lists the necessary operators and operands that explain the values that they have. The table lists the templates of extended signatures as they are created in BigFix Inventory.



Note: To prevent the XML validation from failing, ensure that the following XML element is present in the extended signature that you are creating. Otherwise, the XML validation fails.

```
<Variable name="IS_INSTALLED">
...
</Variable>
```

Identification or detection of	Sample template
A file name	<p>The following template checks the identification or detection of a file name.</p> <pre data-bbox="391 1066 808 1367"> <Variable name="IS_INSTALLED" export="true"> <Compare type="string" relation="ne"> <VectorToScalar> <FindFilePathEx name="A_FILENAME.EXE" /> </VectorToScalar> <ValueOf value="" /> </Compare> </Variable> </pre> <p>The size of the file need not be provided.</p>
Existence of a number of Windows™ registry keys	<p>The following template checks the identification or detection of the existence of a number of Windows™ registry keys.</p> <p>The number of keys to be provided depends on the application.</p> <pre data-bbox="391 1619 1222 1793"> <Variable name="IS_INSTALLED"> <And> <RegKeyValueEqual key="KEY NAME A" entry="ENTRY NAME" value="VALUE" section="64or32" /> <RegKeyValueEqual key="KEY NAME B" entry="ENTRY NAME" value="VALUE" section="64or32" /> <RegKeyValueEqual key="KEY NAME N" entry="ENTRY NAME" value="VALUE" section="64or32" /> </And> </Variable> </pre>

Identification or detection of	Sample template
Existence of at least one of the registry keys	<p>The following template checks the identification or detection of the existence of at least one of the registry keys.</p> <pre data-bbox="391 296 1419 380" style="background-color: #f0f0f0;"> </And> </Variable> </pre> <pre data-bbox="391 470 1419 764" style="background-color: #f0f0f0;"> <Variable name="IS_INSTALLED" export="true"> <Or> <RegKeyValueEqual key="KEY NAME A" entry="ENTRY NAME" value="VALUE" section="64or32" /> <RegKeyValueEqual key="KEY NAME B" entry="ENTRY NAME" value="VALUE" section="64or32" /> <RegKeyValueEqual key="KEY NAME N" entry="ENTRY NAME" value="VALUE" section="64or32" /> </Or> </Variable> </pre>
File that is in a particular directory that is defined by a filter	<p>The following template checks the identification or detection of the file that is in a particular directory that is defined by a filter.</p> <p data-bbox="391 915 1138 957"> Note: Use "?" as the single-character wildcard and "*" to denote multiple characters.</p> <p data-bbox="391 999 1146 1041"> Restriction: Replace the forward or backslashes in the file path with the "?" wildcard.</p> <pre data-bbox="391 1100 1419 1499" style="background-color: #f0f0f0;"> <Variable name="IS_INSTALLED" export="true"> <Compare type="string" relation="ne"> <VectorToScalar> <Vector filter="?opt?symantec*"> <FindFilePathEx name="FILENAME" /> </Vector> </VectorToScalar> <ValueOf value="" /> </Compare> </Variable> </pre>
Key name, key entry, and its value for a registry key	<p>The following template checks the identification or detection of the key name, key entry, and its value for a registry key.</p> <pre data-bbox="391 1604 1419 1730" style="background-color: #f0f0f0;"> <Variable name="IS_INSTALLED" export="true"> <RegKeyValueEqual key="KEY NAME" entry="ENTRY NAME" value="VALUE" section="64or32" /> </Variable> </pre>
CPU architecture	<p>The following template checks the identification or detection of CPU architecture.</p>

Identification or detection of

Sample template

```
<Variable name="IS_INSTALLED" export="true">
  <Compare relation="eq" type="string">
    <ValueOf value="$(B_PROCESSOR_ARCH)" />
    <ValueOf value="ARCHITECTURE_CODE" />
  </Compare>
</Variable>
```

Where **ARCHITECTURE_CODES**:

- AIX, IBM i, and Linux PPC: **PPC**
- Linux x86, Mac OS X, and Solaris x86: **IX86**
- Linux on System z: **S390**
- HP-UX Itanium: **ITANIUM**
- HP-UX PA-RISC: **PA-RISC**
- Solaris SPARC: **SPARC**

File contents and
matching specified
strings

The following template checks the identification or detection of the file contents and matching specified strings.



Note: The scanner reads only text files. It cannot read binary files and archives, such as **.jar**, **.ear**, and **.zip**.

```
<MultipleInstance>
  <Iterator name="file_list" export="false">
    <FindFilePathEx name="MYFILE.txt" />
  </Iterator>
  <Instance>
    <Variable name="file_read" export="false">
      <ReadFile path="$(file_list)MYFILE.txt" />
    </Variable>
    <Variable name="IS_INSTALLED">
      <Compare type="string" relation="ne">
        <VectorToScalar>
          <Vector filter="*MYAPP_version=MYAPP-1.3*">
            <ValueOf value="$(file_read)" />
          </Vector>
        </VectorToScalar>
        <ValueOf value="" />
      </Compare>
    </Variable>
    <Condition withVariable="IS_INSTALLED">
      <Action onValue="false" do="SKIP" />
    </Condition>
  </Instance>
</MultipleInstance>
```

Identification or detection of	Sample template
Existence of two files in a particular directory where <code>someprog.jar</code> file has a particular size	<pre data-bbox="391 296 1409 415"></Condition> </Instance> </MultipleInstance></pre> <p data-bbox="375 436 1273 457">where the name of the file to be detected is <code>MYFILE.txt</code> and the string value is <code>MYAPP_version=MYAPP-1.3</code>.</p> <p data-bbox="375 506 1386 569">The following template checks the identification or detection of the existence of two files in a particular directory where <code>someprog.jar</code> file has a particular size.</p> <pre data-bbox="391 590 1409 1209"><MultipleInstance> <Iterator name="INSTALL_PATH"> <Vector filter="*SOMEPROG.ear?SOMEPROG.war?"> <FindFilePathEx name="essbase.jar" /> </Vector> </Iterator> <Instance> <Variable name="IS_INSTALLED"> <FileInfoCompare absolutePath="\$(INSTALL_PATH)someprog.jar" field="size" value="649432" relation="eq" /> </Variable> <Condition withVariable="IS_INSTALLED"> <Action onValue="false" do="SKIP"/> </Condition> </Instance> </MultipleInstance></pre>
<code>httpd</code> file with a size in a specific directory (but not in other paths)	<p data-bbox="375 1241 1409 1325">The following template checks the identification or detection of the <code>httpd</code> file with a size in a specific directory (but not in other paths).</p> <pre data-bbox="391 1346 1409 1871"><MultipleInstance> <Iterator name="DIR" export="false"> <Vector filter="*Apache?Apache?bin?"> <FindFilePathEx name="httpd" excludeDirectories="*backup*,*bak*,*bkup*" /> </Vector> </Iterator> <Instance> <Variable name="IS_INSTALLED" export="true"> <FileInfoCompare absolutePath="\$(DIR)httpd" field="size" value="629868" relation="eq"/> </Variable> <Condition withVariable="IS_INSTALLED"> <Action onValue="false" do="SKIP"/> </Condition> </Instance> </MultipleInstance></pre>

**Identification
or detection of**

Sample template

```
</Instance>
</MultipleInstance>
```

Component information on AIX® platform

The following template checks the identification or detection of the component information on AIX® platform.

```
<Variable name="IS_INSTALLED">
  <And>
    <Compare type="string" relation="eq">
      <AixProductInfo code="WSBAA@5.0.0.0" field="version"/>
      <ValueOf value="5.0.0.0"/>
    </Compare>
    <Compare type="string" relation="eq">
      <AixProductInfo code="WSBAA@5.0.0.0" field="description"/>
      <ValueOf value="BigFix"/>
    </Compare>
  </And>
</Variable>
```

Provide three values for each AIX® native component:

Code: WSBAA@5.0.0.0

Version: 5.0.0.0

Description: IBM WebSphere® Application Server

Component information on HP-UX platform

The following template checks the identification or detection of the component information on HP-UX platform.

```
<Variable name="IS_INSTALLED">
  <And>
    <Compare type="string" relation="eq">
      <HpProductInfo code="MQSERIES" field="version"/>
      <ValueOf value="6.0.1.0"/>
    </Compare>
    <Compare type="string" relation="eq">
      <HpProductInfo code="MQSERIES" field="name"/>
      <ValueOf value="WebSphere MQ for HP-UX"/>
    </Compare>
  </And>
</Variable>
```

Provide three values for each HP-UX native component:

**Identification
or detection of**

Sample template

Code: MQSERIES

Version: 6.0.1.0

Name: WebSphere® MQ for HP-UX

The existence of a
particular file

The following template checks the identification or detection of the existence of a particular file.

```
<MultipleInstance>
  <Iterator name="INSTALL_PATH">
    <FindFilePathEx name="SWTAG_NAME.swtag"/>
  </Iterator>
  <Instance>
    <Variable name="IS_INSTALLED">
      <Compare type="string" relation="ne">
        <ValueOf value="$(INSTALL_PATH)" />
        <ValueOf value="" />
      </Compare>
    </Variable>
    <Condition withVariable="IS_INSTALLED">
      <Action onValue="false" do="SKIP"/>
    </Condition>
  </Instance>
</MultipleInstance>
```

File type with only
a part of the name
given

The following template checks the identification or detection of the file type with only a part of the name given.

```
<MultipleInstance>
  <Iterator name="INSTALL_PATH">
    <FindFilePathEx name="*.FILETYPE" appendFileName="true"/>
  </Iterator>
  <Instance>
    <Variable name="IS_INSTALLED">
      <Compare type="string" relation="ne">
        <ValueOf value="$(INSTALL_PATH)" />
        <ValueOf value="" />
      </Compare>
    </Variable>
    <Condition withVariable="IS_INSTALLED">
      <Action onValue="false" do="SKIP"/>
    </Condition>
```

Identification or detection of	Sample template
The existence of a specific file, another file does not exist	<pre data-bbox="391 296 1409 380" style="background-color: #f0f0f0;"> </Instance> </MultipleInstance> </pre> <p data-bbox="375 422 1409 453">The following template checks the identification or detection of the existence of a specific file, another file does not exist.</p> <pre data-bbox="391 474 1409 1409" style="background-color: #f0f0f0;"> <MultipleInstance> <Iterator name="INSTALL_PATH"> <FindFilePathEx name="file1"/> </Iterator> <Instance> <Variable name="IS_INSTALLED"> <And> <FileExists absolutePath="\$(INSTALL_PATH)file1"/> <Not> <Compare type="string" relation="ne"> <VectorToScalar> <FindFilePathEx name="file2"/> </VectorToScalar> <ValueOf value=""/> </Compare> </Not> </And> </Variable> <Condition withVariable="IS_INSTALLED"> <Action onValue="false" do="SKIP"/> </Condition> </Instance> </MultipleInstance> </pre>
Software items when the registry key identifies the product, Windows™ systems	<p data-bbox="375 1461 1409 1524">The following template checks the identification or detection of the Software items when the registry key identifies the product, Windows™ systems.</p> <p data-bbox="375 1566 1409 1629">Compares a part of the key string. To address this case, two signatures must be created, one for 32-bit and one for 64-bit architecture.</p> <pre data-bbox="391 1650 1409 1890" style="background-color: #f0f0f0;"> <MultipleInstance> <Iterator name="reg"> <Vector filter="*000000FF1CE}" ignoreCase="false"> <RegKeyList root="HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Uninstall*" section="32" subkeys="false"/> </Vector> </Iterator> </pre>

**Identification
or detection of**
Sample template

```

<Instance>
  <Variable name="IS_INSTALLED">
    <And>
      <Compare relation="eq" type="string">
        <Clip endIndex="76" inputString="$(reg)" startIndex="74"/>
        <ValueOf value="140"/>
      </Compare>
      <Compare relation="eq" type="string">
        <Clip endIndex="84" inputString="$(reg)" startIndex="81"/>
        <ValueOf value="0018"/>
      </Compare>
    </And>
  </Variable>
  <Condition withVariable="IS_INSTALLED">
    <Action do="SKIP" onValue="false"/>
  </Condition>
</Instance>
</MultipleInstance>

```

```

<MultipleInstance>
  <Iterator name="reg">
    <Vector filter="*000000FF1CE}" ignoreCase="false">
      <RegKeyList root="HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Uninstall\*" section="64"
subkeys="false"/>
    </Vector>
  </Iterator>
  <Instance>
    <Variable name="IS_INSTALLED">
      <And>
        <Compare relation="eq" type="string">
          <Clip endIndex="76" inputString="$(reg)" startIndex="74"/>
          <ValueOf value="140"/>
        </Compare>
        <Compare relation="eq" type="string">
          <Clip endIndex="84" inputString="$(reg)" startIndex="81"/>
          <ValueOf value="0018"/>
        </Compare>
      </And>
    </Variable>

```

Identification or detection of

Sample template

```
<Condition withVariable="IS_INSTALLED">
  <Action do="SKIP" onValue="false" />
</Condition>
</Instance>
</MultipleInstance>
```

Registry keys,
where edition and
version are stored

The following template checks the identification or detection of the registry keys, where edition and version are stored.

This signature matches a specific version range.

```
<Variable name="IS_INSTALLED">
  <And>
    <RegKeyValueContains entry="Edition" key="HKEY_LOCAL_MACHINE\SOFTWARE\Vendor\SoftwareName\Data"
value="Enterprise" section="64or32" />
    <RegKeyVersionCompare key="HKEY_LOCAL_MACHINE\SOFTWARE\Vendor\SoftwareName\CurrVer" entry="Version"
relation="ge" value="3.00.00.00" section="64or32" />
    <RegKeyVersionCompare key="HKEY_LOCAL_MACHINE\SOFTWARE\Vendor\SoftwareName\CurrVer" entry="Version"
relation="le" value="3.9999.9999.9999" section="64or32" />
  </And>
</Variable>
```

Software instances
by identifying their
registry keys.

The following template checks the identification or detection of the software instances by identifying their registry keys.

Instance name is reflected by the asterisk character in registry path: `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\SoftwareName*\Info\CurrVer`

This signature matches a specific edition and version range. To address this case, two signatures must be created, one for 32-bit and one for 64-bit architecture.

```
<MultipleInstance>
  <Iterator export="false" name="rg">
    <Vector>
      <RegKeyList root="HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\SoftwareName\*" section="64" subkeys="false" />
    </Vector>
  </Iterator>
  <Instance>
    <Variable export="false" name="vr">
      <RegKeyValue entry="Version" key="$(rg)\Info\CurrVer" mustExists="true" section="64" />
    </Variable>
    <Variable name="IS_INSTALLED">
      <And>
        <RegKeyValueContains entry="Edition" key="$(rg)\Data" section="64" value="Enterprise" />
```

**Identification
or detection of**
Sample template

```

    <Contains inputString="\${vr}" position="start" searchedString="3."/>

    </And>

    </Variable>

    <Condition withVariable="IS_INSTALLED">

        <Action do="SKIP" onValue="false"/>

    </Condition>

    </Instance>

</MultipleInstance>

```

```

<MultipleInstance>

    <Iterator export="false" name="rg">

        <Vector>

            <RegKeyList root="HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\SoftwareName\*" section="32" subkeys="false"/>

        </Vector>

    </Iterator>

    <Instance>

        <Variable export="false" name="vr">

            <RegKeyValue entry="Version" key="\${rg}\Info\CurrVer" mustExists="true" section="32"/>

        </Variable>

        <Variable name="IS_INSTALLED">

            <And>

                <RegKeyValueContains entry="Edition" key="\${rg}\Data" section="32" value="Enterprise"/>

                <Contains inputString="\${vr}" position="start" searchedString="3."/>

            </And>

        </Variable>

        <Condition withVariable="IS_INSTALLED">

            <Action do="SKIP" onValue="false"/>

        </Condition>

    </Instance>

</MultipleInstance>

```

Uninstall registry
keys, Windows™
systems

The following template checks the identification or detection of the uninstall registry keys, Windows™ systems.

This signature iterates through all keys and tries to match specific patterns in DisplayName and DisplayVersion values. To address this case, two signatures must be created, one for the 32-bit and one for the 64-bit architecture.

```

<MultipleInstance>

    <Iterator name="reg" export="true">

        <RegKeyList root="HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall\*"

            section="32" subkeys="false"/>

    </Iterator>

```

**Identification
or detection of**

Sample template

```

<Instance>
  <Variable name="valueName" export="true">
    <RegKeyValue entry="DisplayName" key="$(reg)" section="32" />
  </Variable>
  <Variable name="valueVersion" export="true">
    <RegKeyValue entry="DisplayVersion" key="$(reg)" section="32" />
  </Variable>
  <Variable name="IS_INSTALLED">
    <And>
      <Contains inputString="$(valueName)" searchedString="TestProduct" position="start"/>
      <Contains inputString="$(valueVersion)" searchedString="12.0" position="start"/>
    </And>
  </Variable>
  <Condition withVariable="IS_INSTALLED">
    <Action onValue="false" do="SKIP"/>
  </Condition>
</Instance>
</MultipleInstance>

<MultipleInstance>
  <Iterator name="reg" export="true">
    <RegKeyList root="HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall\*"
    section="64" subkeys="false"/>
  </Iterator>
  <Instance>
    <Variable name="valueName" export="true">
      <RegKeyValue entry="DisplayName" key="$(reg)" section="64" />
    </Variable>
    <Variable name="valueVersion" export="true">
      <RegKeyValue entry="DisplayVersion" key="$(reg)" section="64" />
    </Variable>
    <Variable name="IS_INSTALLED">
      <And>
        <Contains inputString="$(valueName)" searchedString="TestProduct" position="start"/>
        <Contains inputString="$(valueVersion)" searchedString="12.0" position="start"/>
      </And>
    </Variable>
    <Condition withVariable="IS_INSTALLED">
      <Action onValue="false" do="SKIP"/>
    </Condition>
  </Instance>
</MultipleInstance>

```

Identification or detection of	Sample template
	<pre data-bbox="391 296 1425 401"> </Condition> </Instance> </MultipleInstance> </pre>

9.2.2 Extended signature operators

9.2.2 Available from 9.2.2. The scan operations during which extended XML signature definitions are used are performed based on the processing instructions that are contained in the signature catalog. In computer languages, expressions consist of two types of components: operands and operators. Operands are the objects that are manipulated and operators are the symbols that represent specific actions. Operands and operators can be either scalar or vector.

The following definitions apply to operators and operands in the context of the signature catalog:

scalar operand

Returns a value in string format. An operand cannot have children. For example, the **RegKeyValue** scalar operand returns the value of the specified registry key.

scalar operator

Performs an operation on other operators and operands and returns results in string format. An operator can have children. For example, the **AND** operator, returns true or false as a result.

vector operand

Returns results as a set of strings. An operand cannot have children. For example, the **FindFilePathEx** operand, returns a set of file paths.

vector operator

Performs an operation on other operators and operands and returns results as a set of strings. An operator can have children. For example, the **Vector** operator, returns a set of strings that contain the results of all its children.

Contents

[The structure of the signature catalog \(on page cdlxv\)](#)

[Core expressions \(on page cdlxvi\)](#)

[File system expressions \(on page cdlxxiii\)](#)

[Native registry expressions \(on page cdlxxvi\)](#)

[Registry operands \(on page cdlxxxiii\)](#)

[Windows™ registry expressions \(on page cdlxxxiv\)](#)

The structure of the signature catalog

Table 44. The structure of the signature catalog

Operands and operators	Function and supported attributes
Action	<p>Specifies the action to be performed depending on the result of the Condition element.</p> <p>Attributes:</p> <p>onValue</p> <p>The value that causes the action to be performed. This attribute is required.</p> <p>do</p> <p>Specifies the operation to be performed. This attribute is required. The supported value is <i>SKIP</i>, which interrupts the signature or the iterator instance evaluation and prevents it from being returned in the output.</p>
Condition	<p>Verifies if the specified condition is met. It contains one or more Action elements, which specify the action to be performed depending on the result of the Condition element.</p> <p>Attributes:</p> <p>withVariable</p> <p>Specifies the name of the variable whose condition is to be verified. This attribute is required.</p>
MultipleInstance	<p>Groups information to identify a software product that might be installed in several instances on the same workstation. It contains 1 Iterator and 1 Instance element.</p> <p>Attributes:</p> <p>None</p> <p>Multiple instance - iterator on INSTALL_PATH</p> <pre data-bbox="483 1377 1365 1873"> <MultipleInstance> <Iterator name="INSTALL_PATH"> <FindFilePathEx name=" "/> </Iterator> <Instance> <Variable name="IS_INSTALLED"><!-- RULE --> </Variable> <Condition withVariable="IS_INSTALLED"> <Action onValue="false" do="SKIP"/> </Condition> </Instance> </MultipleInstance> </pre>

Table 44. The structure of the signature catalog (continued)

Operands and operators	Function and supported attributes
Variable	<p>Multiple instance - iterator on a vector with calculation (options) of INSTALL PATH later on</p> <pre data-bbox="483 407 1365 852"> <MultipleInstance> <Iterator name="ITER_VAR"> </Iterator> <Instance> <Variable name="IS_INSTALLED"><!-- RULE --> </Variable> <Variable name="INSTALL_PATH"><!-- RULE --> </Variable> <Condition withVariable="IS_INSTALLED"> <Action onValue="false" do="SKIP"/> </Condition> </Instance> </MultipleInstance> </pre>
	<p>Is a container for a single value that is expressed as a string whose return value is given by the result of a set of operators and operands.</p>
Attributes:	
name	<p>Specifies the name of the variable. It is unique in the containing signature scope. This attribute is required.</p>
export	<p>Specifies if the output of the variable is to be exported. Supported values are <i>true</i> and <i>false</i>. The default value is <i>true</i>. This attribute is optional.</p>

Core expressions

Operands and operators	Function and supported attributes
And	<p>Performs a logical AND operation between the results of all its children. If one of the operations returns false, the whole operation is set to false without performing further evaluation on the remaining children. At least one child is required.</p>
Properties:	
	<ul style="list-style-type: none"> • Result type: scalar [<i>true</i> <i>false</i>] • Expression type: operator
Example:	

Operands and operators

Function and supported attributes

```
<And>
...scalar operators and operands returning a Boolean value
</And>
```

Clip

Returns a substring that is a range of consecutive characters from a string, starting with the character whose index is specified as **startIndex** and ending with the character whose index is specified as **endIndex**.

Properties:

- Result type: scalar
- Expression type: operand

Example:

For example: the substring 234 is returned when:

```
<Clip inputString="123456789" startIndex="1" endIndex="3"/>
```

Attributes:

inputString

Specifies the string to be cropped.

startIndex

Specifies the index of the character in the input string that the substring should start with. An index of 0 refers to the first character of the input string.

endIndex

Specifies the index of the character in the input string that the substring should end with. An index of 0 refers to the last character of the input string.



Note: It is recommended to use the new operator **StringSlice** whose range is wider than that of **Clip** because **StringSlice** supports negative indexes in Python syntax.

Compare

Compares two expressions based on the value of the relation attribute. The result is *true* if the expression is satisfied, *false* if the expression is not satisfied.

Properties:

- Result type: scalar
- Expression type: operator

Example:

```
<Compare relation="relation" type="type" >
<Expression 1/>
```

Operands and operators

Function and supported attributes

```
<Expression 2/>
</Compare>
```

Attributes:

relation

Specifies the relation to be evaluated. The supported values are as follows:

- *eq* - equal to
- *ne* - not equal to
- *gt* - greater than
- *ge* - greater than or equal to
- *lt* - less than
- *le* - less than or equal to

type

Specifies the format of the expressions to be compared. The supported values are as follows:

- *version* - Compares two versions. The version is expressed by no more than four numbers that are separated by periods or commas.
- *integer* - Compares two integers.
- *string* - Compares two strings.
- *cistring* - Compares two case-insensitive strings.
- *boolean* - The comparison is to be performed between two Boolean values.
- *hex* - Compares two unsigned 4-byte hexadecimal integers. The 0x prefix is optional.

Concat

Returns a concatenation of the expressions.

Properties:

- Result type: scalar
- Expression type: operator

Example:

```
<Concat [separator="separator"] [ifEmpty="{add|skip}"] >
  <Expression 1/>
  <Expression 2/>
  ...
</Concat>
```

Attributes:

Operands and operators

Function and supported attributes

separator

Specifies the separator to insert between two expressions. This attribute is optional. The default value is an empty string.

ifEmpty

Specifies if the separator must be added if the expression to concatenate is empty. This attribute is optional. Supported values are as follows:

skip

If the string is empty, the separator is not added. This is the default value.

add

If the string is empty, the separator is added.

Contains

Checks if the string contains a string.

Properties:

- Result type: scalar [{true|false}]
- Expression type: operand

Example:

```
<Contains inputString="example_string" searchedString="example_string_to_search" position = ["start" | "end" | "contains"]/>
```

Attributes:

inputString

Specifies string to search in.

searchedString

Specifies string to search.

position

Specifies where to search.

start

Checks if the searched string is at the beginning.

end

Checks if searched string is at the end.

contains

Checks if searched string is somewhere in the original string.

**Operands and
operators****Function and supported attributes****FileInfoCompare**

Returns the information that is specified in the **field** attribute.

Properties:

- Result type: scalar [{true|false}]
- Expression type: operand

Example:

```
<FileInfoCompare
  absolutePath="absolute_path"
  field="field"
  value="value"
  relation="relation"
  [mustExist="{true|false}"]
/>
```

Attributes:**absolutePath**

Specifies the absolute path to the file to evaluate.

field

Compares two files by one of the values that are supported, which are as follows:

- *size*
- *inode* (UNIX™ only)
- *atime* - the date the file was accessed in UNIX™ time stamp format.
- *ctime* - the date the file was created in UNIX™ time stamp format.
- *mtime* - the date the file was modified in UNIX™ time stamp format.

relation

Specifies the relation to be evaluated.

value

Specifies the value to compare with.

mustExist

Indicates that the specified file must exist. This attribute is optional. The default value is *false*. If this attribute is set to *true* and the specified file does not exist, the current signature is not evaluated.

FileExists

Returns a Boolean value that indicates if the specified file exists.

Properties:

Operands and operators	Function and supported attributes
Not	<ul style="list-style-type: none"> • Result type: scalar [<i>true</i> <i>false</i>] • Expression type: operand <p>Example:</p> <pre data-bbox="483 464 1365 506"><FileExists absolutePath="absolute_path" /></pre> <p>Attributes:</p> <p>absolutePath</p> <p>Specifies the absolute path to the file to be searched.</p> <p>Performs a logical NOT operation on the result of its one child. Only one child is allowed.</p> <p>Properties:</p> <ul style="list-style-type: none"> • Result type: scalar [<i>true</i> <i>false</i>] • Expression type: operator <p>Example:</p> <pre data-bbox="483 957 1365 1094"><Not> ... a single scalar operator or operand returning a Boolean value </Not></pre>
PathExists	<p>Returns a Boolean value that indicates if the specified path exists.</p> <p>Properties:</p> <ul style="list-style-type: none"> • Result type: scalar • Expression type: operand <p>Example:</p> <pre data-bbox="483 1398 1365 1440"><PathExists absolutePath="absolute_path" /></pre> <p>Attributes:</p> <p>absolutePath</p> <p>Specifies the absolute path to be searched.</p>
Or	<p>Performs a logical OR operation between the results of all its children. If one of the operations returns "true", the whole operation is set to "true" without performing further evaluation on the remaining children. At least one child is required.</p>
Substring	<p>Searches through a string that contains multiple lines of text and returns the lines of a string that contain the search text.</p> <p>Properties:</p>

Operands and operators	Function and supported attributes
Vector	<ul style="list-style-type: none"> • Result type: vector • Expression type: operand <p>Example:</p> <pre data-bbox="483 464 1365 543"><Substring inputString="input_string_to_search" searchedString="string_to_search_for" /></pre> <p>Attributes:</p> <p>inputString Input string that contains multiple lines.</p> <p>searchedString The string that is searched for.</p> <p>Returns a vector that contains the union of all its children, duplicates are included. If you specify a filter, the result contains only the strings that match the pattern.</p> <ul style="list-style-type: none"> • filter: Specifies the filter criteria to be applied to the search. The asterisk (*) and question mark (?) wildcards are supported.
VectorAt	<p>Returns the vector element for the specified position.</p> <p>Properties:</p> <ul style="list-style-type: none"> • Result type: scalar • Expression type: operator <p>Example:</p> <pre data-bbox="483 1318 1365 1451"><VectorAt index="index_number"> ... any number of scalar/vector operator/operand </VectorAt></pre> <p>Attributes:</p> <p>index Specifies the wanted vector position. The first vector element has position "1".</p>
VectorToScalar	<p>Converts a vector to its scalar representation. You must convert a vector to scalar because only scalars can be compared.</p>

File system expressions

Operands and

operators

Function and supported attributes

FindFilePathEx

Searches for the specified file in the specified locations. The result is a vector of all the paths where at least one entry that matches the specified file is found.

The operand searches either the cache or the file system depending on the settings that are defined for the file system scanner. To have the operand search the file system, set the **maxDataAge** attribute to zero in the configuration file. An extra filter can be set to refine the file search. Such a filter can contain a file system scanner query, which is applied to the found file. If the filter returns a true condition, the directory for the file is returned, otherwise it is not returned.

Properties:

- Result type: vector
- Expression type: operand

Example:

```
< FindFilePathEx name="file_name" [filter="query_filter"] [recursive="{true|false}"]
[appendFileName="{true|false}"]/>
```

Attributes:

name

Specifies the file mask to be searched. This attribute is required.

appendFileName

Specifies if the results contain the file name. This attribute is optional. The supported values are as follows:

- *true*
- *false*

filter

Specifies a file system query.

FileInfoMatch

Searches for a file with the specified name and attributes in the specified locations. The result is *true* if at least one matching file is found.

Properties:

- Result type: scalar
- Expression type: operand

Example:

Operands and operators

Function and supported attributes

```
<FileInfoMatch
  [rootPath="root_path"]
  [recursive="{true|false}"]
  name="file_name"
  field="field"
  values="values"
/>
```

Attributes:

rootPath

Specifies the path where the search is to be performed. The default value is /. Wildcards (*) (?) are supported. The path is indicated according to the following syntax:

```
media_type::]drive:path
```

where

media_type

Specifies the media type and can assume the following values:

\$local

Includes local disks. This is the default value.

\$remote

Includes remote disks

\$other

On Windows™ systems, includes other disk types, such as CDROM, DVD. On UNIX™ systems, this value is not supported and is expressed with the `$local` value.

\$all

Includes all disks that are listed above.

If the `media_type` variable is not specified, the `$local` value is assumed, therefore, if you want to scan disks other than local, specify a value for the `media_type` variable.

drive

Specifies the drives to be excluded. You can indicate one or more specific drives or you can use the `$local`, `$remote`, `$other`, and `$all` values. In this case, all drives in the specified category are excluded from the search. This variable is optional on UNIX™ operating systems.

path

Operands and**operators****Function and supported attributes**

Specifies the path to be excluded. Wildcards (*) and (?) are supported.

This attribute is optional.

recursive

Specifies if the search is to be performed in the specified path and in all its subdirectories.

Supported values are *true* and *false*. The default value is *true*. This attribute is optional.

name

Specifies the file name to be searched. This attribute is required. Wildcards (*) and (?) are supported.

field

Compares two files by one of the values that are supported, which are as follows:

- *size*
- *inode* (UNIX™ only)
- *atime* - the date the file was accessed in UNIX™ time stamp format.
- *ctime* - the date the file was created in UNIX™ time stamp format.
- *mtime* - the date the file was modified in UNIX™ time stamp format.

values

Specifies one or more values for the attribute that is specified in the **field** attribute. If more than one value is specified, the items are concatenated by using the Unicode control character "Private Use 1" with ASCII code 0x0091.

date

Returns the creation, access, and modification dates of the file.

ReadFile

Returns in string form the text from a file on the system.

Properties:

- Result type: scalar
- Expression type: operand

Example:

```
<ReadFile path="path_to_file" [trim="{true|false}"][fileType="{native|utf8}"]/>
```

Attributes:**Path**

Specifies the path to the file.

trim

Operands and operators	Function and supported attributes
	Indicates if new line special characters must be removed from the content of the file. The default is <i>false</i> .
	fileType
	Indicates the type of file. The default is <i>native</i> .

Native registry expressions

Operands and operators	Function and supported attributes
AixProductInfo	Returns any information about the supported AIX® operating systems.
	Properties:
	<ul style="list-style-type: none"> • Result type: scalar • Expression type: operand
	Example:
	<pre><AixProductInfo code="product_code" field="{name version vendor description installDir}" [mustExist="{true false}*"] /></pre>
	Attributes:
	code
	Specifies the product code that uniquely identifies an AIX® ODM package. The product code can be retrieved only by using the AixProductList operand. For more information about this operand, see AixProductList (on page cdlxxvii) .
	field
	Specifies the information to be retrieved. Supported values are as follows:
	<ul style="list-style-type: none"> • <i>name</i> • <i>version</i> - The version is expressed by no more than four numbers that are separated by periods or commas. • <i>vendor</i> • <i>description</i> • <i>installDir</i>
	mustExist
	Specifies whether to stop the operand evaluation if the specified file does not exist. Supported values are <i>true</i> and <i>false</i> . The default value is <i>false</i> . This attribute is optional. If this attribute is set to <i>true</i> and the specified file does not exist, the current signature is not evaluat-

Operands and operators

Function and supported attributes

ed. If the attribute is *false* and the specified file does not exist, the return value is an empty string.

AixProductList

Returns the set of AIX® ODM package codes. The codes that are returned can be used with the **AixProductExists** and **AixProductInfo** operands.

Properties:

- Result type: vector
- Expression type: operand

Example:

```
<AixProductList [filter="filter_expression"]
[type="type"] />
```

Attributes:

type

Applies a filter based on the product type. This attribute is optional. Supported values are as follows:

any

Returns both products and any subproducts. This is the default value.

product

Returns only products.

filter

Specifies the filter criteria to be applied to the search. The asterisk (*) and question mark (?) wildcards are supported. This attribute is optional.

HpProductInfo

Returns the requested information, if available.

Properties:

- Result type: scalar
- Expression type: operand

Example:

```
<HpProductInfo code="product_code" field="{name|version|
vendor|installDir}"/>
```

Attributes:

code

Operands and operators

Function and supported attributes

Specifies the product and subproduct code that uniquely identifies an HP-UX package. The product and subproduct code can be retrieved only by using the **HpProductList** operand. For more information about this operand, see [HpProductList \(on page cdlxxviii\)](#). This attribute is optional.

field

Specifies the information to be retrieved. Supported values are as follows:

- *name*
- *version* - The version is expressed by no more than four numbers that are separated by periods or commas.
- *vendor*
- *installDir*

HpProductList

Returns the set of HP-UX package codes. The codes that are returned can be used with the HpProductExists and HpProductInfo operands.

Properties:

- Result type: vector
- Expression type: operand

Example:

```
<HpProductList [filter="filter_expression"]
  [type="type" ] />
```

Attributes:

type

Applies a filter based on the product type. This attribute is optional. Supported values are as follows:

any

Returns both products and any subproducts. This is the default value.

product

Returns only products.

subproduct

Returns only subproducts.

filter

Specifies the filter criteria to be applied to the search. The asterisk (*) and question mark (?) wildcards are supported. This attribute is optional.

Operands and**operators****Function and supported attributes****RpmProductExists**

Indicates if the specified package is installed.

Properties:

- Result type: scalar [true | false]
- Expression type: operand
- Operating system: supported on Linux™ and AIX

Example:

```
<RpmProductExists code="product_code" />
```

Attributes:**code**

Specifies the product code that uniquely identifies an RPM package. The product code can be retrieved by using the **RPMProductList** operand. For more information about this operand, see [RpmProductList \(on page cdlxxix\)](#).

The product code can also be listed with the `rpm -qa` command. The command returns all available codes that identify RPM packages. The codes can either define the package architecture or not; for example: `kernel-2.6.32-358.el6.x86_64` or `kernel-2.6.32-358.el6`.

RpmProductList

Returns the set of RPM package codes. The codes that are returned can be used with the **RpmProductExists** and **RpmProductInfo** operands.

Properties:

- Result type: vector
- Expression type: operand
- Operating system: supported on Linux™ and AIX

Example:

```
<RpmProductList [filter="filter_expression"]
[type="type"] />
```

Attributes:**filter**

Specifies the filter criteria to be applied to the search. The asterisk (*) and question mark (?) wildcards are supported. This attribute is optional.

type

Applies a filter based on the product type. This attribute is optional. Supported values are as follows:

**Operands and
operators**

Function and supported attributes

any

Returns both products and any subproducts. This is the default value.

product

Returns only products.

SunProductInfo

Returns the requested information, if available.

Properties:

- Result type: scalar
- Expression type: operand
- Operating system: supported Oracle Solaris operating environments

Example:

```
<SunProductInfo code="product_code" field="{name|version|displayVersion|vendor|installDir}"/>
```

Attributes:

code

Specifies the product code that uniquely identifies an Oracle Solaris package. The product code can be retrieved only by using the **SunProductList** operand. For more information about this operand, see [SunProductList \(on page cdlxxx\)](#).

field

Specifies the information to be retrieved. Supported values are as follows:

name

Returns the name of the product.

version

Returns the product version and revision number, for example 3.7.2101.

displayVersion

Returns the product version and revision number, for example 3.7.2101,
REV=0.98.08.26.

vendor

Returns the name of the vendor.

installDir

Returns the topmost installation directories.

SunProductList

Returns the set of Sun Solaris package codes. The codes that are returned can be used with the **SunProductExists** and operands.

Operands and operators

Function and supported attributes

Properties:

- Result type: vector
- Expression type: operand
- Operating system: supported Sun Solaris operating environments

Example:

```
<SunProductList [filter="filter_expression"]
  [type="type"] />
```

Attributes:

type

Applies a filter based on the product type. This attribute is optional. Supported values are as follows:

any

Returns both products and any subproducts. This is the default value.

product

Returns only products.

filter

Specifies the filter criteria to be applied to the search. The asterisk (*) and question mark (?) wildcards are supported. This attribute is optional.

Win32ProgramInfo

Returns the requested information, if available.

Properties:

- Result type: scalar
- Expression type: operand
- Operating systems: all supported Windows™ operating systems

Example:

```
<Win32ProgramInfo code="product_code"
  field="{name|description|version|vendor|installDir|uninstData}"/>
```

Attributes:

code

Specifies the product code that uniquely identifies a Win32 program. The program code can be retrieved only by using the **Win32ProgramList** operand. For more information about this operand, see [Win32ProgramList \(on page cdlxxxii\)](#).

**Operands and
operators**

Function and supported attributes

field

Specifies the information to be retrieved. Supported values are as follows:

name

Returns the patch name.

description

Returns the patch description.

version

Returns the product version. The version is expressed by no more than four numbers that are separated by periods or commas.

vendor

Returns the product vendor.

installDir

Returns the product installation directory.

uninstData

Returns the product uninstallation command.

Win32ProgramList

Returns the set of packages installed. The codes that are returned can be used with the **Win32ProgramExists** and **Win32ProgramInfo** operands.

Properties:

- Result type: vector
- Expression type: operand
- Operating systems: all supported Windows™ operating systems

Example:

```
<Win32ProgramList [filter="filter_expression"] [type="type"] />
```

Attributes:

type

Applies a filter based on the product type. This attribute is optional. Supported values are as follows:

any

Returns both products and any subproducts. This is the default value.

product

Returns only products.

Operands and operators	Function and supported attributes
filter	Specifies the filter criteria to be applied to the search. The asterisk (*) and question mark (?) wildcards are supported. This attribute is optional.

Registry operands

Operands and operators	Function and supported attributes
PackageInfoEqual	Returns true if an item with the same name, version, and vendor exists in the specified registry.
Properties:	<ul style="list-style-type: none"> • Result type: scalar [true false] • Expression type: operand
Example:	<pre data-bbox="483 869 1365 1087"><PackageInfoEqual provider="{any ismp os}" code="product_code" value="name_version_vendor" /></pre>
Attributes:	provider
	Specifies the installation database to be searched. Supported values are as follows:
any	Specifies that the search is performed on the operating system installation database and then on the ISMP installation database.
os	Specifies that the search is performed on the operating system installation database.
ismp	Specifies that the search is performed on the ISMP installation database.
code	Specifies the product code that uniquely identifies the product.
value	Comprises the concatenation of the name, version, and vendor strings for the product. Separate each string with a blank space; empty strings can be omitted.

Operands and operators	Function and supported attributes
PackageInfoMatch	Returns true if an item with the same name, version, and vendor exists in any registry.
	Properties:
	<ul style="list-style-type: none"> • Result type: scalar [true false] • Expression type: operand • Operating systems: all supported operating systems
	Example:
	<pre data-bbox="483 590 1365 804"><PackageInfoMatch vendor="product vendor" name="product name" version="product version" /></pre>
	Attributes:
	vendor
	Specifies the product vendor.
	name
	Specifies the name of the product.
	version
	Specifies the product version number.
	The wildcard (*) is supported.

Windows™ registry expressions

Operands and operators	Function and supported attributes
RegKeyExists	Indicates whether the specified key exists.
	Properties:
	<ul style="list-style-type: none"> • Result type: scalar [true false] • Expression type: operand
	Example:
	<pre data-bbox="483 1667 1365 1724"><RegKeyExists key="key" [section="{32 64 32or64 64or32}"]/></pre>
	Attributes:
	key
	Specifies the full registry key name in <code>registry_key\registry_subkey</code> format.

Operands and operators**Function and supported attributes****section**

Specifies the attribute that indicates which section of the registry the search is to be performed on. This is the parameter introduced to make sure that the scanner functions correctly on 64-bit systems. Its values are:

32

The search is only performed on the 32-bit part of the registry.

64

The search is only performed on the 64-bit part of the registry.

32or64

The search is performed on the 32-bit part first, then on the 64-bit part of the registry.

64or32

The search is performed on the 64-bit part first, then on the 32-bit part of the registry.

Default value is 32.

RegKeyList

Searches the Windows™ registry under the key that is specified in the root attribute, and returns a list of keys that are in the registry under the root key. If the **subKeys** attribute is set to *false*, only the values in the root key are returned.

Properties:

- Result type: vector
- Expression type: operand

Example:

```
<RegKeyList
  root="root_key"
  [section="{32|64|32or64|64or32}"]
  [subKeys="{true|false}"]/>
```

Attributes:**root**

Specifies the full registry key name in `registry_key\registry_subkey` format.

section

Specifies the attribute that indicates which section of the registry the search is to be performed on. (This is the parameter introduced to make sure that the scanner runs correctly under 64-bit systems. Its values are:

32

Operands and operators	Function and supported attributes
	<p>The search is only performed on the 32-bit part of the registry.</p> <p>64</p> <p>The search is only performed on the 32-bit part of the registry.</p> <p>32or64</p> <p>The search is performed on the 32-bit part first, then on the 64-bit part of the registry.</p> <p>64or32</p> <p>The search is performed on the 64-bit part first, then on the 32-bit part of the registry.</p> <p>Default value is 32.</p> <p>subKeys</p> <p>Specifies if subkeys must be returned. The default value is <i>false</i>.</p>
<p>RegKeyValue</p>	<p>Returns the current value for the specified registry key. When the registry key is a REG_DWORD entry, the key attribute returns a decimal value.</p> <p>Properties:</p> <ul style="list-style-type: none"> • Result type: scalar • Expression type: operand <p>Example:</p> <pre data-bbox="483 1165 1369 1304"><RegKeyValue entry="entry" key="key" [section="{32 64 32or64 64or32}"] [mustExist="{true false}"]/></pre> <p>Attributes:</p> <p>key</p> <p>Specifies the full registry key name in registry_key\registry_subkey format.</p> <p>entry</p> <p>Specifies the name of the value that is contained in the key. To retrieve the value of registry entries that are named "(Default)", you must specify a null value for the entry attribute, for example <i>entry=""</i>.</p> <p>section</p> <p>Specifies the attribute that indicates which section of the registry the search is to be performed on. This is the parameter introduced to make sure that the scanner runs correctly on 64-bit systems. Its values are:</p> <p>32</p>

Operands and operators**Function and supported attributes**

The search is only performed on the 32-bit part of the registry.

64

The search is only performed on the 32-bit part of the registry.

32or64

The search is performed on the 32-bit part first, then on the 64-bit part of the registry.

64or32

The search is performed on the 64-bit part first, then on the 32-bit part of the registry.

Default value is 32.

mustExist

Specifies whether to stop the operand evaluation if the specified key does not exist. Supported values are *true* and *false*. The default value is *false*. This attribute is optional. If this attribute is set to *true* and the specified key does not exist, the current signature is not evaluated. If this attribute is set to *false* and the specified key does not exist, an empty string is returned.

RegKeyValueContains

Specifies whether the specified registry key contains the specified value.

Properties:

- Result type: scalar [true | false]
- Expression type: operand

Example:

```
<RegKeyValueContains
  entry="entry"
  key="key"
  value="value"
  [section="{32|64|32or64|64or32}"]
  [mustExist="{true|false}"]/>
```

Attributes:**key**

Specifies the full registry key name in `registry_key\registry_subkey` format.

entry

Specifies the name of the value that is contained in the key.

value

Specifies the value that must be contained in the entry.

Operands and operators	Function and supported attributes
	<p>section</p> <p>Specifies the attribute indicating which section of the registry the search should be performed on. This is the parameter introduced to make sure that the scanner runs correctly on 64-bit systems. Its values are:</p> <p>32</p> <p>The search is only performed on the 32-bit part of the registry.</p> <p>64</p> <p>The search is only performed on the 32-bit part of the registry.</p> <p>32or64</p> <p>The search is performed on the 32-bit part first, then on the 64-bit part of the registry.</p> <p>64or32</p> <p>The search is performed on the 64-bit part first, then on the 32-bit part of the registry.</p> <p>Default value is 32.</p> <p>mustExist</p> <p>Specifies whether to stop the operand evaluation if the specified key does not exist. Supported values are <i>true</i> and <i>false</i>. The default value is <i>false</i>. This attribute is optional. If this attribute is set to <i>true</i> and the specified key does not exist, the current signature is not evaluated. If this attribute is set to <i>false</i> and the specified key does not exist, an empty string is returned.</p>
<p>RegKeyValueEqual</p>	<p>Indicates if the specified registry key value is equal to the specified key.</p> <p>key</p> <p>Specifies the full registry key name in <code>registry_key/registry_subkey</code> format.</p> <p>entry</p> <p>Specifies the name of the value that is contained in the key.</p> <p>value</p> <p>Specifies the value to be compared. When the registry key that is specified in the entry attribute is a REG_DWORD entry, enter a decimal value in the value attribute.</p>
<p>RegKeyVersionCompare</p>	<p>Compares the specified registry key value with the specified version based on the value of the relation attribute.</p> <p>Properties:</p> <ul style="list-style-type: none"> • Result type: scalar [<i>true</i> <i>false</i>] • Expression type: operand

Operands and operators

Function and supported attributes

Example:

```
<RegKeyVersionCompare
  key="key"
  entry="entry"
  relation="relation"
  value="value"
  [section="{32|64|32or64|64or32}"]
  [mustExist="{true|false}"]/>
```

Attributes:

key

Specifies the full registry key name in *registry_key\registry_subkey* format.

entry

Specifies the name of the value that is contained in the key.

relation

Specifies the relation to be evaluated. Supported values are as follows:

- *eq* - equal to
- *ne* - not equal to
- *gt* - greater than
- *ge* - greater than or equal to
- *lt* - less than
- *le* - less than or equal to

value

Specifies the version to be compared.



Note: When the registry key that is specified in the entry attribute is a **REG_DWORD** entry, enter a decimal value in the value attribute. The supported value is **version**. The version is expressed by up to four numbers that are separated by periods or commas.



Note: When the registry key that is specified in the entry attribute is a **REG_BINARY** entry in the value field, enter a string that looks exactly like the "data" field when you are viewing the field in the registry editor. This string must be of a series of two-digit numbers that are separated by spaces, for example, "01 01" "10 11 10".

section

Operands and operators**Function and supported attributes**

Specifies the attribute that indicates which section of the registry the search is to be performed on. This is the parameter introduced to make sure that the scanner runs correctly on 64-bit systems. Its values are:

32

The search is only performed on the 32-bit part of the registry.

64

The search is only performed on the 64-bit part of the registry.

32or64

The search is performed on the 32-bit part first, then on the 64-bit part of the registry.

64or32

The search is performed on the 64-bit part first, then on the 32-bit part of the registry.

The default value is 32.

mustExist

Specifies whether to stop the operand evaluation if the specified key does not exist. Supported values are *true* and *false*. The default value is *false*. This attribute is optional. If this attribute is set to *true* and the specified key does not exist, the current signature is not evaluated.

9.2.2 Testing extended signatures

9.2.2 Available from 9.2.2. You can test the extended signatures that you create by using a BigFix client that is installed on the platform for which a signature is defined.

1. Prepare a test catalog file that contains only the signature that you want to verify and the name the file `extended_signature_test.xml`.

The following example shows a test catalog file. It consists of three main parts: the header, the body and the footer. The body, which is in bold, is formatted basing on the available signature templates. For the full list, see: [Extended signature templates \(on page cdlii\)](#).

```
<?xml version="1.0" encoding="UTF-8"?>
<Signatures>
  <Signature guid="EXTENDED SIGNATURE TEST" name="EXTENDED SIGNATURE TEST" type="OtherSignature">

    <!-- SIGANTURE BODY START -->

    <Variable name="IS_INSTALLED" export="true">

      <RegKeyExists key="HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\" section="64or32"/>

    </Variable>
  </Signature>
</Signatures>
```

```

<!-- SIGANTURE BODY END -->

</Signature>
</Signatures>

```

2. Add a scanner location to the environment variables by running the following command:

o **AIX**

```

LIBPATH="/opt/tivoli/cit/bin:$LIBPATH"
export LIBPATH

```

o **HP-UX**

```

SHLIB_PATH="/opt/tivoli/cit/bin:$SHLIB_PATH"
export SHLIB_PATH

```

o **Linux** **Solaris**

```

LD_LIBRARY_PATH="/opt/tivoli/cit/bin:$LD_LIBRARY_PATH"
export LD_LIBRARY_PATH

```

o **Windows**

```

set Path=C:\Program Files\Tivoli\cit\bin;%Path%

```

3. Reset the file system scanner by running the following command:

o **UNIX** /opt/tivoli/cit/bin/wscanfs -reset

o **Windows** "C:\Program Files\Tivoli\cit\bin\wscanfs.exe" -reset

4. Run a scan by using the test catalog file.

o **UNIX** /opt/tivoli/cit/bin/wscansw -i extended_signature_test.xml -o output.xml -e warning.xml

o **Windows** "C:\Program Files\Tivoli\cit\bin\wscansw.exe" -i extended_signature_test.xml -o output.xml -e warning.xml

5. Review the `output.xml` and `warning.xml` files.

a. In `output.xml`, ensure that the **MatchedSignature** element is returned for every matching instance. If no instance is found, there are no **MatchedSignature** elements in the file.

```

<?xml version="1.0" encoding="UTF-8"?>
<MatchedSignatures>
  <MatchedSignature guid="EXTENDED SIGNATURE TEST" name="EXTENDED SIGNATURE TEST"
    type="OtherSignature">
    <Variable name="IS_INSTALLED" value="true"/>
  </MatchedSignature>
</MatchedSignatures>

```

b. Ensure that there are no error messages in the `warning.xml` file.

Creating signatures from package data

You can use package data to expand your custom catalog. You can add signatures for software products that are installed in your infrastructure but are not present in the HCL® catalog.

 You must have the View Raw Data, View Endpoints and Manage Catalogs permissions to perform this task.

- **9.2.11** For versions 9.2.11 and higher

1. Log in to BigFix Inventory, and go to **Reports > Package Summary**.
2. Set up the relevant filters, and identify the packages that you want to create signatures for.
3. Check one or more packages. You can select up to twenty packages at a time.
4. Hover over the **Signature** icon  and click **Add Signature**.



Note: You cannot create a signature on a Package Data report, if you already added either a file or an extended signature for the component with the same details. If you create a package signature for a component that is a part of a basic catalog, the component and the detection will be duplicated.

- For versions before 9.2.11

1. Log in to BigFix Inventory, and go to **Reports > Package Data**.
2. To filter the report data, hover over the **Manage Report View** icon , and click **Configure View**. Then, specify the appropriate filtering options.



Note: To find a file or a package that can be used as a signature, filter the report to the data from the computer on which the particular software product is installed. To further narrow down the results, specify the whole or part of the path to the directory where the software is installed. Then, look at the following columns.

- **Recognized** - the column shows whether the file was recognized as part of an existing signature.
- **Caused Detection** - the column shows whether the file contained enough information to cause detection of the related software. This column is not enabled by default. For more information, see: [Report columns \(on page dcxxx\)](#).

If the value in both columns is No, the file is not used for software detection and you can use it to create a new software signature.

3. In the **Name** column, look for the location where the missing software title is installed. Click the arrow on the right from the **Name** of the file that you want to use as a signature, and click **Create Signature**.

Computer Name	Last Seen	Name	Version
NC046213	about 4 hours ago	IBM 32-bit SDK for Java 2, v5.0	5.0
NC046213	about 4 hours ago	IBM 32-bit SDK for Java 2, v5.0	5.0
VMW009128109094	about 4 hours ago	IBM 64-bit Runtime	
VMW009128109094	about 4 hours ago	IBM 64-bit Runtime Environment for 6	

4. In the **Create Catalog Entry** window, provide the publisher name, product name, and release number.

Create Catalog Entry

Publisher Name*

Product Name*

Release*

A component will be created based on the release information. It represents a discoverable software item that is displayed on reports.

Report release when the following conditions are fulfilled. If you specify two conditions, they both must be accomplished.

The following installation package is found:

Name*	Version	Vendor	
<input type="text" value="IBM 32-bit SDK for Jav"/>	<input type="text" value="5.0"/>	<input type="text" value="IBM"/>	<input type="button" value="Remove"/> <input type="button" value="Add"/>

The following file is found:

Name*	With	
<input type="text"/>	<input type="text" value="Any size or version"/>	<input type="button" value="Remove"/> <input type="button" value="Add"/>

If you select both conditions, at least one installation package and one file must be detected to report the release.

 **Tip:** The most efficient signatures are file-based signatures with a constant size that is different for each release, or signatures that are based on registry entries where a wildcard (*) can be used. You can use a wildcard in the **Name**, **Version**, and **Vendor** fields for a package. A wildcard represents any string and can be combined with other characters. For example, you might use IBM * SDK for Java™, 5.*, * for the name, version, and vendor.

5. To save the catalog entry, click **Submit**.

 **Important:** If you create an entry that exists in the software catalog an informational message is displayed, and both signatures are saved in the catalog.

You added a software signature to your custom catalog, and it can now be used to detect software. When you create signatures from package data, the custom content version of the software catalog remains the same, because the signatures are matched on the server side, and no new catalog needs to be created.

You can view the audit trail of the changes on the **Catalog Audit** report.

You can edit the details of the signature and delete a custom catalog entry on the **Catalog Customization** panel.

Related information

[Software discovery process after customizing the catalog \(on page cdxcviii\)](#)

Creating signatures from scanned file data

You can use scanned file data to expand your custom catalog. You can add signatures for software products that are installed in your infrastructure but are not present in the IBM catalog. You identify the file name, or another item of information that can be used to identify the software and add the signature.

 You must have the View Endpoints and Manage Catalogs permissions to perform this task.

1. Log in to BigFix Inventory.
2. In the top navigation bar, click **Reports > Scanned File Data**.
3. To filter the report data, hover over the **Manage Report View** icon , and click **Configure View**. Then, specify the appropriate filtering options.
 - a. **Optional:** Select the **Recognized** and the **Caused Detection** check boxes to add these two columns to the view of scan data.



Note: To find a file or a package that can be used as a signature, filter the report to the data from the computer on which the particular software product is installed. To further narrow down the results, specify the whole or part of the path to the directory where the software is installed. Then, look at the following columns.

- **Recognized** - the column shows whether the file was recognized as part of an existing signature.
- **Caused Detection** - the column shows whether the file contained enough information to cause detection of the related software. This column is not enabled by default. For more information, see: [Report columns \(on page dcxxx\)](#).

If the value in both columns is No, the file is not used for software detection and you can use it to create a new software signature.

4. **Optional:** In the **File Path** column, look for the location where the missing software title is installed.
5. Click the arrow on the right from the **File Name** that you want to use as a signature, and click **Create Signature**.

Computer Name	Last Seen	File Path	File Name	Size
VMW009128109094	about 5 hours ago	C:\airgap	BESAirgapTool.exe	-
VMW009128109094	about 5 hours ago	C:\cygwin\bin	a2p.exe	-
VMW009128109094	about 5 hours ago	C:\cygwin\bin	addffinfo.exe	-

6. In the **Create Catalog Entry** window, provide the publisher name, product name, and release number.

9.2.11 Starting from application update 9.2.11, the Scanned File Data report provides full versions of discovered files. However, best practices of creating software signatures recommend that file version that is specified in the signature is limited to the release number, for example, 9.0. When you specify a more detailed version of the file such as 9.0.7, the signature does not detect all patches or mod-releases of the software component. When you create a signature from the scanned file data, the prefilled version of the file is limited to the release number according to the best practices.

Create Catalog Entry

Publisher Name*

Product Name*

Release*
A component will be created based on the release information. It represents a discoverable software item that is displayed on reports.

Report release when the following conditions are fulfilled. If you specify two conditions, they both must be accomplished.

The following installation package is found:

Name*	Version	Vendor	
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="button" value="Remove"/> <input type="button" value="Add"/>

The following file is found:

Name*	With	Size	
<input type="text" value="BESAirgapTool.exe"/>	<input type="text" value="Size (in bytes)"/> <input type="button" value="v"/>	<input type="text" value="1850304"/>	<input type="button" value="Remove"/>
Name*	With	Version	
<input type="text" value="BESAirgapTool.exe"/>	<input type="text" value="Version"/> <input type="button" value="v"/>	<input type="text" value="9.0"/>	<input type="button" value="Remove"/> <input type="button" value="Add"/>

If you select both conditions, at least one installation package and one file must be detected to report the release.

7. To save the catalog entry, click **Submit**.

! **Important:** If you create an entry that exists in the software catalog, an information message is displayed, and both signatures are saved in the catalog.

8. To make the signatures available for software detection, click **Reports > Import Now**, or wait for a scheduled import.

You added a software signature to your custom catalog, and it can now be used to detect software or its usage. When you create signatures from scanned file data, the custom content version of the software catalog remains the same, because the signatures are matched on the server side, and no new catalog needs to be created.

You can view the audit trail of the changes on the **Catalog Audit** report.

You can edit the details of the signature and delete a custom catalog entry on the **Catalog Customization** panel.

Related information

[Software discovery process after customizing the catalog \(on page cdxcviii\)](#)

Adding a new product to software catalog

The software catalog consists of a list of standard software products. Use the Add Product feature to define a custom product with the accompanying information, such as a publisher name, license metric and release.

 You must have the View Software Catalog and Signatures and Manage Licenses permissions to perform this task.

You can add a custom product to the software catalog to ensure it contains all your business applications. You can later use the new product during software classification.

1. Log in to **Products & Metrics**, and go to **Reports > Products & Metrics**.
2. Hover over the **Add** icon , and select **Add Product**.
3. Provide the publisher name. If the publisher already exists in the software catalog, you can search for it, and select it from the drop-down list.
4. Provide the product name. The product name must be unique for the publisher that you provided in step 3.
5. Select a metric from the list. To search for a metric, type the metric name in the search bar.
6. Provide the release of the product. By default, this field is set to 1.0.
7. Click **OK**. You will be notified about product creation.
8. Wait for the scheduled import of data, or run it manually. You cannot add new products during the import.

Thus, if you plan to create more than one product, first create all products and then run the import of data.

A new line is added on the Products & Metrics report. You can now use the newly created product to classify your software. The newly created product is available on all reports after the import of data is successfully completed.

To assign components to the new product, go to **Management > Software Classification**, select the components, and reassign them. For more information, see: [Assigning components to products \(on page dxlvi\)](#).

To delete the catalog entry that you create for the product, see: [Deleting custom catalog entries](#).

9.2.8 Assigning an additional metric to a product

9.2.8 Available from 9.2.8. The pricing calculations are based on the relation between a product and a metric.

The software catalog consists of an inventory of standard relations. To define a custom relation, assign one of available metrics to an existing product.

 You must have the View Software Catalog and Signatures and Manage Licenses permissions to perform this task.

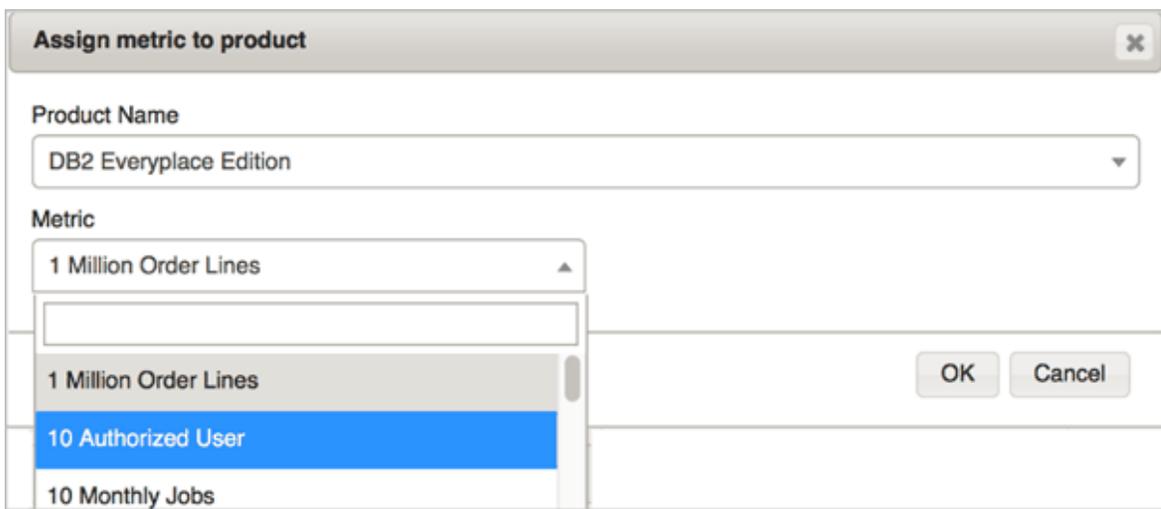
You can assign one of available metrics to a product from the software catalog, and thus, define a custom relation.

You cannot duplicate an existing relation. You can only assign a metric that is not yet assigned to a certain product.

The custom relation can be used while classifying your software inventory. It allows you to make your pricing calculations accurate and tailor-made.

 **Note:** If you use BigFix Inventory version 9.2.8 or higher, go to Software Catalog report to define a custom relation. Starting from application update 9.2.14, the Software Catalog report is renamed to Products & Metrics report.

1. Log in to BigFix Inventory, and go to **Reports > Products & Metrics**.
2. Select a product for which you want to assign a metric. You can select any product from the software catalog.
3. Hover over the **Add** icon , and select **Assign Metric to Product**.
4. Select a metric from the list. To search for a metric, type the metric name in the search bar.



5. Click **OK** to confirm.

After you confirm the assignment of a new metric to a product, a new line is added on the Products & Metrics report. You can now use the newly created custom relation to classify your software.

To assign components to the product with the new metric, go to **Reports > Software Classification**, select the components, and reassign them. For more information, see: [Assigning components to products \(on page dxlvi\)](#).

9.2.11 Changing end of support date

9.2.11 Available from 9.2.11. Set or change the end of support date for your software components to have a comprehensive overview of your software inventory.



You must have the Manage Catalogs permission to perform this task.



You must have the View Software Catalog and Signatures permission to perform this task.

After you install or upgrade the BigFix Inventory server, the end of support date is automatically populated for the following components after the initial import of data.

- BigFix components announced via https://support.hcltechsw.com/csm?id=bigfix_support
- **9.2.14** Microsoft software that according to information provided on 20 November 2018 is announced to be out of support between 30 November 2015 and 1 November 2021. For more information, see: [Microsoft Lifecycle Policy](#).

. When you manually set the end of support date for a component, the imported date is permanently overwritten.

1. Go to **Reports > Software Components**.
2. Select one or more components for which you want to set the end of support date. You can select up to a 100 rows at a time.
3. Hover over the **Edit Component** icon , and click **Change End of Support**.
4. Choose the end of support date, and click **Change**.
5. **Optional:** To clear the end of support date, remove the date, and click **Change**.

You set the end of support date for the selected components.

9.2.12 Starting from application update 9.2.12, end of support date is also displayed on the Software Classification report.

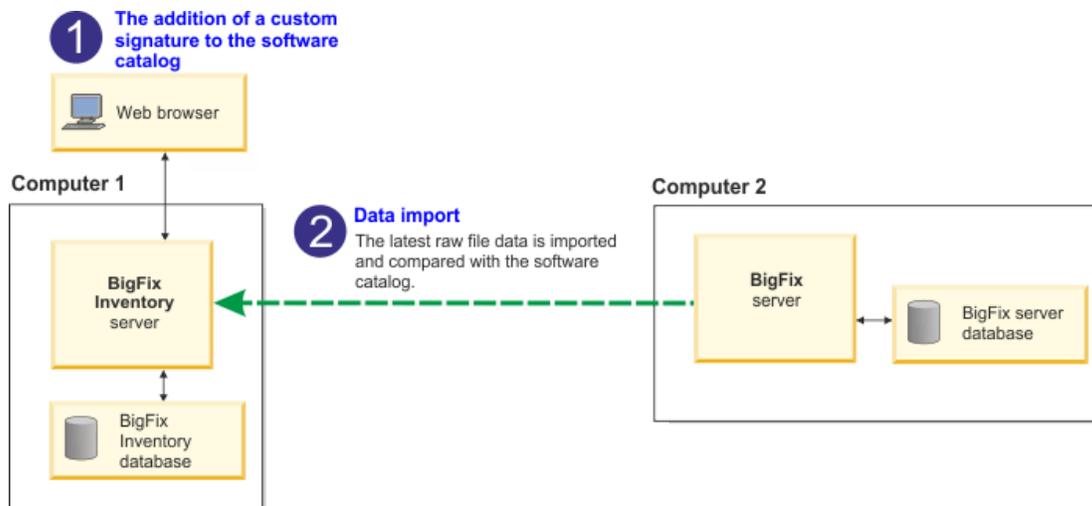
You can use end of support dates to filter and sort the report. You can also create a saved report view that is filtered by the specified information, and enable notifications that are triggered when specific criteria are met.

Software discovery process after customizing the catalog

Software matching your catalog customizations can be identified in one of two ways depending on the custom signature.

Case 1: Common executable files or other popular file types

If you created a custom signature, which is based on package data or file data such as common executable files, then you need to run a data import to pull the latest scan data from the BigFix server. The new data is compared with the latest catalog and the updated software inventory is displayed in the web user interface.

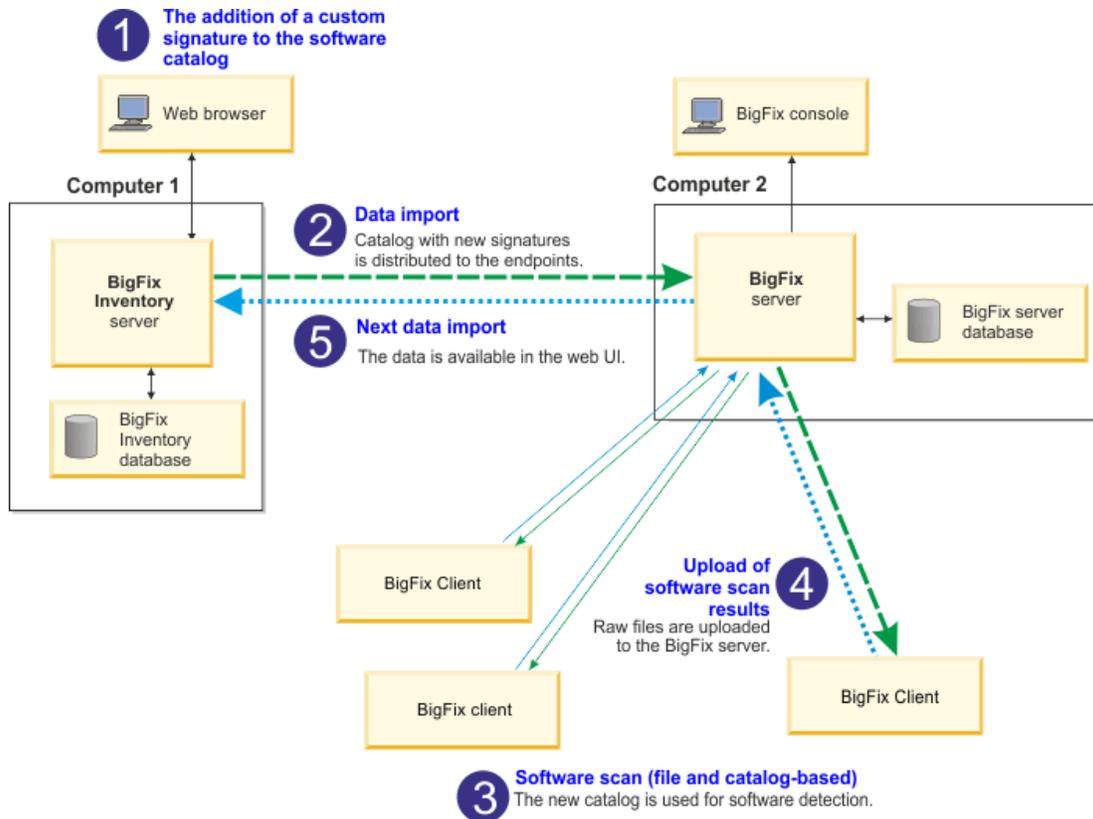


The popular file extension list includes the following file types: `*.exe`, `*.sys`, `*.com`, `*.ear`, `*.ocx`, `*.sh`, `*.bin`, `*.pl`, `*.ear`, `*.SH`, `*.BIN`, `*.PL`, and `*.EAR`.

When you create signatures based on these file types, the custom content version of the software catalog remains the same, because no new catalog is created. The data is pulled from your endpoints and matched against the signatures on the server side.

Case 2: Non-standard file types

If you added new software signatures to the software catalog and those signatures are based on files other than common executable files, you must either wait until all the following steps in the catalog data flow complete or perform those steps yourself.



When you create signatures based on non-standard file types, the custom content version of the software catalog changes, because the new catalog is created and distributed to your endpoints. The catalog with the new signatures is matched against the data on your endpoints.

Exporting and importing custom signatures

You can export custom signatures to an XML file and import them to a different instance of BigFix Inventory to augment your software catalog. It is useful if you maintain separate instances of BigFix Inventory for test and production sites. In such case, you can reuse the custom catalog content instead of creating the new one. You can also import signatures that are published on the signature community or export your custom signatures and contribute them to the community.

 You must have the Manage Catalogs permission to perform this task.

The [signature community](#) is a place where you can find custom signatures that are created by members of the community. You can download such signatures and import them to BigFix Inventory. You can also export your custom signatures and contribute them to the community for other members to download.

1. Export custom signatures to an XML format.
 - a. In the top navigation bar, click **Management > Catalog Customizations**.
 - b. Click **Export Mode**.

- To export only a selection of signatures, check the appropriate check boxes and click **Export Selected**. You can select individual signatures, or all signatures for a publisher. Each signature is exported to a separate XML file, and compressed. When you select multiple signatures, all XML files are compressed into one ZIP file.
 - To export the entire custom catalog, click **Export All**. You export all signatures, regardless of your selection, or filters, in a form of one ZIP file that contains individual XML files.
2. Import custom signatures in an XML format to the other instance of BigFix Inventory.
 - a. In the top navigation bar, click **Management > Catalog Customizations**.
 - b. Click **Import** to open the **Import Signatures** window.
 - c. Browse to the location of the file that you want to import.
 - d. Select the file and click **Import**. You can either import a single XML file or a ZIP package that contains one or more XML files.
 - e. Click **Done**.

The imported signatures are available on the **Catalog Customizations** panel. They are now used by software scans to discover software in your environment.

Deprecated: Synchronizing software catalog content

You can export your custom catalog content from one instance of BigFix Inventory and import it into another to merge the custom data. It is useful if you maintain separate instances of BigFix Inventory for test and production sites. In such case, you can reuse the custom catalog content instead of creating the new one.



Important: **9.2.11** Starting from application update 9.2.11.0, you can no longer export the software catalog with your custom content from the Catalog Upload panel. As an alternative, you can export your custom signatures and import them onto the other instance of BigFix Inventory using the Catalog Customization panel. For more information, see: [Exporting and importing custom signatures \(on page d\)](#).



You must have the Manage Catalogs permission to perform this task.



Restriction:

- The catalog is exported in to the native format. This format overwrites the existing catalog, so it can be imported only into a fresh installation of BigFix Inventory before the initial import.
- You can transfer custom catalog content only between two instances of BigFix Inventory if they have the same version.

When you export the custom catalog content from BigFix Inventory, an XML file that contains all catalog entries is saved to your computer. Then, you must log in to the instance of BigFix Inventory to which you want to import the custom catalog content and upload the XML file.

1. Log in to the instance of BigFix Inventory from which you want to export the custom catalog content.
2. In the top navigation bar, click **Management > Catalog Upload**.
3. In the **Catalog Export** section, click **Export**. An XML file with your custom catalog content is saved to your computer.
4. Log in to the instance of BigFix Inventory to which you want to import the custom catalog content.
5. In the top navigation bar, click **Management > Catalog Upload**.
6. Click **Browse**, select the XML file with your custom catalog content, and click **Upload**.

The software catalog file is listed in the table. The status is **Pending** until you go to **Management > Imports**, and click **Import Now** to process new changes.

Enterprise applications inventory

Software components are discovered by software scans. The scans correlate data that is collected from the computers with the content of the software catalog, and thus determine whether a component is installed. However, software scans do not discover details of some enterprise applications, for example licensable options of Oracle Database or editions of Microsoft Exchange. The details are discovered by dedicated fixlets. If you are interested in discovering additional details of enterprise applications, run the dedicated fixlets apart from scheduling software scans. Otherwise, running the fixlets is not required.

Software identification tag files for enterprise applications are stored in the following directories:

- `BESClient/LMT/ORACLE`
- `BESClient/LMT/SAP`
- `BESClient/LMT/MS/exchange`

The software identification tag path that is provided in BigFix Inventory shows the product installation path, and the tag name.

9.2.9

Starting from application update 9.2.9, the software identification tag path and installation path are shown separately on condition that the paths are different and the tag contains information about the installation path. Paths to software identification tags discovered before version 9.2.9 might also be updated after the upgrade on condition that sufficient data is available.

Extended discovery of Oracle Database

You can retrieve granular information about the edition (Standard or Enterprise), options, Management Packs (including Oracle Real Application Cluster, RAC), and user concurrent sessions of Oracle Databases that are deployed in your environment. When you combine this information with the collected hardware inventory data, you can use it to calculate the demand for Oracle Database licenses. The functionality is supported for Oracle Database 11.1, 11.2, and 12.1.

Table 45. Discovered versions and editions of Oracle Database

The table consists of two columns and four body rows.

Oracle Database version	Edition	Supported operating systems
9.2.15 18c: 18.3	Enterprise	Oracle Database is discovered on the following operating systems. <ul style="list-style-type: none"> • AIX • Linux • Windows For detailed versions of the operating systems, see: BigFix 9.5 - System Requirements .
12c Release 1: 12.1.0.2.0	Enterprise	
11g Release 2: 11.1.0.1.0	Enterprise and Standard	
11g Release 1: 11.1.0.6.0	Enterprise and Standard	

Coverage of Oracle licensable options

Check which options can be discovered for a particular version of Oracle Database. The options are discovered by the official Oracle SQL auditing script.

Table 46. Coverage of Oracle licensable options

Licensable Options for the Enterprise Edition of Oracle Database					Oracle SQL script coverage
	9.2.15	18c	12c R1	11g R2	11g R1
Oracle Active Data Guard	Yes		Yes	Yes	Yes
Oracle Advanced Analytics	Yes				
Oracle Advanced Compression	Yes		Yes	Yes	Yes
Oracle Advanced Security	Yes		Yes	Yes	Yes
Oracle Database In-Memory	Yes		Yes		
Oracle Database Vault	Yes		Yes	Yes	Yes
Oracle Label Security	Yes		Yes	Yes	Yes
Oracle Multitenant	Yes				
Oracle On-Line Analytical Processing (OLAP)	Yes		Yes	Yes	Yes
Oracle Partitioning	Yes		Yes	Yes	Yes
Oracle Real Application Clusters (Oracle RAC)	Yes		Yes	Yes	Yes

Table 46. Coverage of Oracle licensable options (continued)

Licensable Options for the Enterprise Edition of Oracle Database	9.2.15	18c	12c R1	11g R2	11g R1	Oracle SQL script coverage
Oracle Real Application Clusters One Node						
Oracle Real Application Testing	Yes		Yes		Yes	Yes
Oracle Spatial and Graph Management Packs	Yes					
Total Recall						Yes
Data Mining					Yes	Yes
Spatial Warehouse Builder						Yes
Oracle Change Management Pack				Yes	Yes	Yes
Oracle Configuration Management Pack				Yes	Yes	Yes
Oracle Diagnostic Pack	Yes		Yes	Yes	Yes	Yes
Tuning Pack	Yes		Yes	Yes	Yes	Yes
Oracle Provisioning and Patch Automation Pack				Yes	Yes	Yes
Transparent Gateways						
Oracle Cloud Management Pack for Oracle Database						
Oracle Data Masking Pack for Oracle and Non-Oracle Databases			Yes			Yes
Oracle Database Lifecycle Management Pack for Oracle Database						
In-Memory Database Cache						
Oracle TimesTen Application-Tier Database Cache			Yes			

Discovering usage of Oracle Database features

To retrieve detailed information about Oracle Database editions, options and management packs that are used in your infrastructure, distribute the Oracle reporting script on relevant endpoints, and run the **Get Oracle Features** task. When the information is imported to BigFix Inventory, it is displayed on the Oracle Databases report.

The task establishes Oracle database connection as a user that has rights to run the BigFix Agent process. Ensure that this user has the Oracle `SYSDBA` privilege.

The discovery of Oracle Database editions, options and management packs is based on the use of one of the Oracle scripts: `options_packs_usage_statistics.sql` or `ReviewLite.sql`. The former should be used by default. Both scripts retrieve information about usage of Oracle features.



Note: BigFix Inventory supports `ReviewLite.sql` 17.1 and higher. Certification was performed on version 17.1 and 17.2.

The `ReviewLite.sql` script is used mostly in case of an audit, review, or contract renewal. It additionally gathers data that might be required by the Oracle License Management Services (LMS) and uploads it to the BigFix server. The `ReviewLite.sql` script produces the `options.csv` file which might contain data that is not entirely accurate and require confirmation from Oracle LMS. Based on the `options.csv` file, BigFix Inventory might detect usage or potential usage of Oracle features.

The following examples show the potential usage of Oracle features :

- Potential use of Tuning Pack or Real Application Testing might be reported based on the data from SQL Tuning Sets. If they are contained in the results, the Tuning Pack or Real Application Testing license is required.
- Potential use of Diagnostic Pack might be reported based on `control_management_pack_access` data if the value of `control_management_pack_access` is `DIAGNOSTIC` OR `DIAGNOSTIC+TUNING`.

1. Choose the type of the Oracle script that you want to use to discover features of the used Oracle databases, and download it.

You can choose `options_packs_usage_statistics.sql` or `ReviewLite.sql`.

- If you choose `options_packs_usage_statistics.sql`, download the script from the [Oracle support website](#). If the script is not available under this link, refer to Oracle documentation specific to your version of the database for more information:
 - [Oracle Database 12c](#)
 - [Oracle Database 11g](#)



Important: To download the script, you must have a valid Oracle account. It is recommended that you contact [Oracle License Management Services](#) to ensure good cooperation.

- If you choose `ReviewLite.sql`, contact the [Oracle License Management Services](#) to obtain the script.

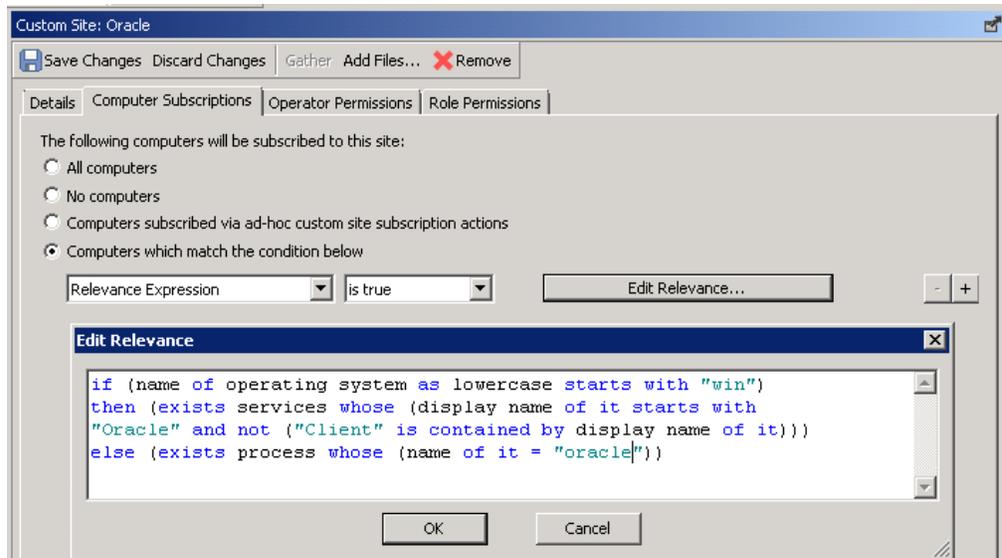
2. Propagate the downloaded script to all computers on which you want to discover usage of Oracle features.

You can propagate the script in one of the following ways.

- Use the method of propagating files that is typical in your enterprise. Place the script in the following directory: `<BES Client>\LMT`.
- Add the script to the Master Action Site and send it to all computers in your enterprise.
 - a. In the top navigation bar of the BigFix console, click **Tools > Add Files to Site**.
 - b. From the drop-down list, choose **Master Action Site**.

- c. Click **Browse** and go to the directory in which the script is located.
 - d. Select the script and click **Open**.
 - e. Select **Send to clients** and click **Add files**. The script is sent to all computers that are subscribed to the Master Action site.
- Create a custom site to which you subscribe computers on which you want to discover the used features. Then, add the script to this site and send it to the subscribed computers.
 - a. In the top navigation bar of the BigFix console, click **Tools > Create Custom Site**. Provide the name of the custom site that begins with "Oracle", and click **OK**.
 - b. Open the **Computer Subscriptions** tab, and subscribe computers on which you want to discover usage of Oracle features. For example, you can choose **Computers which match the condition below**, select **Relevance Expression, is true**, and provide the following relevance expression.

```
if (name of operating system as lowercase starts with "win")
then (exists services whose (display name of it starts with
"Oracle" and not ("Client" is contained by display name of it)))
else (exists process whose (name of it = "oracle"))
```



- c. In the top navigation bar of the BigFix console, click **Tools > Add Files to Site**.
 - d. From the drop-down list, choose the custom site that you created.
 - e. Click **Browse** and go to the directory in which the script is located.
 - f. Select the script and click **Open**.
 - g. Select **Send to clients** and click **Add files**. The script is sent to all computers that are subscribed to the custom site.
3. To retrieve information about features of the used Oracle databases, log in to the BigFix console, and go to **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**. Then, click **Get Oracle Features**.
 4. Choose the type of the Oracle script that you downloaded in step 1 (on page dv).
If you choose `ReviewLite.sql`, read and accept the license agreement related to the Oracle script.

You can find the license terms in the script or you can print them to the console by running the following script on the Oracle database.

```
sqlplus <user_name>/<password> @ReviewLite.sql
```

When you print the license terms, you are asked whether you accept them. You can ignore this question as you confirm that you accept the terms by selecting **Accept License Agreement** in the **Get Oracle Features** task. By selecting this option, you also agree that BigFix will accept the script license on your behalf on all computers on which you run the **Get Oracle Features** task.

1. Choose the Oracle script that you want to use to discover the used Oracle features.

options_packs_usage_statistics.sql

This script should be used by default. You can download it from the [Oracle support website](#).

ReviewLite.sql

This script is used mostly in case of an audit, review, or contract renewal. To obtain the script, contact the [Oracle License Management Services](#).

Read and accept the license agreement related to the script. To view the license, run the following script on the Oracle database:sqlplus <user_name>/<password> @ReviewLite.sql

Accept License Agreement

When you accept the license agreement, you indicate that you read and accept the license agreement related to the Oracle script. You also agree that BigFix will accept the script license on your behalf on all endpoints on which you run this task.

5. **Optional:** To enable debug logging, select **Debug mode**.

The information is logged in the `script_trace.txt` file. By default, the file is in the following location.

- **UNIX** /var/opt/BESClient/LMT/ORACLE

- **Windows** C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\ORACLE

6. To run the task, click **Take Action**. On the **Target** tab, select computers from which you want to retrieve the information, and click **OK**.



Tip: To check whether the task completed successfully or to troubleshoot a problem, go to **Analyses**, right-click **Oracle Features Health Check**, and then click **Activate**. The analysis provides information about the type and version of Oracle script that was initiated through the Get Oracle Features fixlet. It also provides information about the time of last execution of the fixlet, and its status: Success, Failure or Partial Success. If the status is Partial Success, the script was run successfully only on some Oracle instances on a given computer.

7. Ensure that software identification tags scan and upload of its results are scheduled. For more information, see: [Initiating software scans \(on page cciii\)](#) and [Uploading software scan results \(on page ccviii\)](#).
8. To make the data available in BigFix Inventory, wait for the scheduled import of data or run it manually.

Information about the used features is available on the Oracle Databases report. It is also included on the Software Installations report.

Computer Name	Name	Release	Edition	Feature	SID	Install Path
NC9128110208	Oracle Database	11.2	Oracle Database - Enterprise Edition	Oracle Database Licensable Option - TuningPack	ordc	/home/oracle/app/oracle/product/11.2.0.4
NC9128110167	Oracle Database	11.2	Oracle Database - Enterprise Edition	Oracle Database Licensable Option - TuningPack	ORCL	c:\app\administrator\product\11.2.0.4
NC9128110167	Oracle Database	11.2	Oracle Database - Enterprise Edition	Oracle Database Licensable Option - AdvancedCompr...	ORCL	c:\app\administrator\product\11.2.0.4
NC9128110167	Oracle Database	11.2	Oracle Database - Standard Edition	Oracle Database Licensable Option - AdvancedCompr...	ORCL3	c:\app\administrator\product\11.2.0.4

Additionally, raw output of the script is sent to the to BigFix server. It can be found in the following directory: `<BigFix Enterprise>\BES Server\UploadManagerData\BufferDir\shal\<last_2digits_of_computer_ID>\<computer_ID>\oracleresults_0_<computer_ID>(.zip/.tar.gz).`

9.2.6 Metering the number of Oracle Database concurrent sessions

To retrieve detailed information about the number of Oracle Database concurrent user sessions that run on the computers in your environment, use the Get Oracle Concurrent Sessions Number fixlet. When the data is imported to BigFix Inventory, it is displayed in the Resource Utilization report.

The task establishes Oracle database connection as the `LocalSystem` user on Windows, or as the database instance owner on Linux, AIX or Solaris. Ensure that this user has the Oracle `SYSDBA` privilege.

BigFix Inventory supports metering of Named User Plus license consumption. Run this task to retrieve information about the current number of concurrent sessions, and the high water mark value of concurrent sessions for Oracle Databases that are installed on the computers in your infrastructure. These values are individually retrieved for each database instance with the following SQL query.

```
SELECT sessions_highwater, sessions_current FROM v$license;
```

The database instance must be running when the task is executed. Otherwise, the values of high water mark, and current number of sessions equal 0 in the report. After you run this task, the data is collected on a daily basis by default, to preserve the high water mark value that is reset after each database instance restart.

For more information about the Oracle licenses, see: [Oracle Database Licensing](#).

1. Log in to the BigFix console.
2. In the navigation bar, click **Sites > External Sites > BigFix Inventory v9 > Fixlets and Tasks**.
3. In the upper right pane, select **Get Oracle Concurrent Sessions Number**.
4. **Optional:** To enable debug logging, select the **Debug mode** check box.

The information is logged in the `script_trace_sessions.txt` file that is, by default, in the following location.

UNIX `/var/opt/BESClient/LMT/ORACLE`

Windows `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\ORACLE`

5. To run the task, click **Take Action**. On the **Target** tab, select computers from which you want to retrieve the information, and click **OK**.
6. Ensure that the software scans are enabled, and run automatically. If not, [run the software scan \(on page cciiii\)](#) manually. Select the **Resource Utilization Scan**.

7. Ensure that the scan results are automatically uploaded. If not, [upload the results \(on page ccviii\)](#) to the BigFix server manually.
8. To make the data available in BigFix Inventory, wait for the scheduled import of data, or run it manually.

The retrieved data about the Oracle database concurrent sessions is listed in the Resource Utilization report. To view the data, log in to BigFix Inventory, and go to **Reports > Resource Utilization**.

Retention period for historical data is configured for 90 days.

Computer Name	Software Name	Metric Type	First Measured	Last Measured	Maximal Trend Value	Maximal Trend Value History 2016-10-17 - 2016-11-15
NC040211	Oracle Database - Standard Edition	SESSION_CURRENT	11/04/2016 01:00 AM	11/14/2016 04:29 PM	3	
NC040211	Oracle Database - Standard Edition	SESSION_HIGHWATERMARK	11/04/2016 01:00 AM	11/14/2016 04:29 PM	8	

9.2.6 Discovering Microsoft Exchange edition

9.2.6 Available from 9.2.6. To retrieve detailed information about Microsoft Exchange servers that are installed in your infrastructure, run the **Get Microsoft Exchange Edition** fixlet. The data is collected, imported to BigFix Inventory, and displayed in the Software Installations report.

The discovery of Microsoft Exchange edition is only applicable to the computers with the Microsoft Exchange server installed. The server does not need to be running during this task.

The following Microsoft Exchange versions are supported: 2007 (only 64bit version), 2010, 2013, 2016. The Microsoft Exchange edition is obtained through the cmdlet `Get-ExchangeServer` from Exchange Management Shell. For more information, see: [Exchange Server Updates: build numbers and release dates](#).

1. Log in to the BigFix console.
2. In the navigation bar, click **Sites > External Sites > BigFix Inventory v9 > Fixlets and Tasks**.
3. In the upper right pane, select **Get Microsoft Exchange Edition**.
4. To run the task, click **Take Action**. On the **Target** tab, select computers from which you want to retrieve the information, and click **OK**.
5. Ensure that the software scans are enabled, and run automatically. If not, [run the software scan \(on page cciii\)](#) manually. Select the **Software identification tags scan**.
6. Ensure that the scan results are automatically uploaded. If not, [upload the results \(on page ccviii\)](#) to the BigFix server manually.
7. To make the data available in BigFix Inventory, wait for the scheduled import of data, or run it manually.

The discovered Microsoft Exchange editions are listed on the Software Installations report. To view the data, log in to BigFix Inventory, and go to **Reports > Software Installations**.



Notice: After you activate the Get Microsoft Exchange Edition fixlet, the Microsoft Exchange instances are duplicated on the Software Installations report. The reason behind the duplication is the dual detection. The Microsoft Exchange is included in the software catalog, and thus, detected as a standard installation. During the edition discovery, a separate line is added to the report that provides information about the Microsoft Exchange edition. See the duplicates on the following example:

Publisher Name	Component Name	Component Version	Product Name	Metric	Computer Name	Installation Path	Details
Microsoft	Microsoft Exchange Server Standard Edition	2013.0	Microsoft Exchange Server Standard Edition	Unknown	NC912811123	<See details>	DETAILS >
Microsoft	Microsoft Exchange Server	2013.0	Microsoft Exchange Server	Unknown	NC912811123	<See details>	DETAILS >

You can filter the report to remove the redundant lines.

9.2.7 Measuring license metric utilization of VMware products

9.2.7 Available from 9.2.7. BigFix Inventory allows you to retrieve granular information about license metric utilization of VMware vCenter and vSphere in your infrastructure. The resource utilization data is collected by the VM Manager Tool and displayed on the Resource Utilization report.

9.2.7 Enabling metric calculation of VMware products (basic VM management)

9.2.7 Available from 9.2.7. You can calculate license metric utilization of the VMware products: vCenter and vSphere in standard environments where the VM managers are centralized and configured through user interface. Use basic (central) VM management to enable data retrieval.

To collect granular information about the license metric utilization of VMware vCenter and vSphere

- Update the VM Manager Tool to the newest version. Metric calculation of VMware products is supported starting from version 1.5.0.0.
- Ensure you have BigFix Inventory version 9.2.1 or higher.
- Ensure that the user defined for the VMware connection has Global.Licenses permission.



Note: The API for retrieving the license data is available starting from vSphere API 4.0. The resource utilization is collected only for the VMware installations that are defined as VM managers for BigFix Inventory.

When you enable license metric calculation of VMware vCenter and vSphere, the .slmtag files are created locally on the machine running VM Manager Tool and collected by the resource utilization scan.

1. Log in to BigFix Inventory.
2. Go to **Management > Advanced Server Settings**.
3. Set the **vmman_slm_tags_enabled** parameter to `true`, and click **Save**.
4. **Optional:** To specify the maximum size of each VMware .slmtag file, configure the **vmman_slm_tags_file_size_limit** parameter value, and click **Save**.
5. **Optional:** To specify the maximum number of VMware .slmtag files that are stored on the computer for each license key under VMware instance, configure the **vmman_slm_tags_files_per_software** parameter value, and click **Save**.
6. The data will be available in BigFix Inventory after the next resource utilization scan and data import.

To view license metric utilization, log in to BigFix Inventory and navigate to **Reports > Resource Utilization**.

The following columns are specific to the VMware vCenter and vSphere license metric utilization data:

- **Computer Name** always shows a name of the computer where the VM Manager Tool is installed.
- **Software Name** shows the name and the version of the software with additional information about license.
- **Metric Subtype** indicates if the count refers to the license metric utilization or overall license capacity. The column shows one of two values. USED refers to the license metric utilization of the specified license on a particular machine. CAPACITY reflects the limit agreed in this license. The number is specified in **Maximal Trend Value** column.

Computer Name	Software Name	Metric Type	Metric Subtype	Instance	Maximal Trend Value
NC9143126163	VMware vSphere 6 Enterprise Plus (unlimited cores per CPU)	SOCKET	USED	/P142Q-08L40-H8041-013RP-CHPQ1?VirtualCenter_198.51.100.0	12
NC9143126163	VMware vSphere 6 Enterprise Plus (unlimited cores per CPU)	SOCKET	CAPACITY	/P142Q-08L40-H8041-013RP-CHPQ1?VirtualCenter_198.51.100.0	999,999



Note: The maximal trend value for unlimited licenses is always set to 999,999.

- **Instance** shows the instance identifier that consists of:
 - License key, for example: *P142Q-08L40-H8041-013RP-CHPQ1*
 - Machine type: virtual center or ESX box.

Virtual center refers to the entire vSphere infrastructure for which VM manager connection is defined. The number of used licenses for virtual center shows the license metric utilization for all the boxes within this center. Virtual center can consist of more than one ESX, or ESXi box. VM manager can connect directly to the ESX, or ESXi box. Then, the displayed license metric utilization is calculated for each individual box.

License utilization is reported for the VMware connection that is defined in BigFix Inventory. Thus, if you define the virtual center connection, the .slmtag files are generated for this center, not individual ESX, or ESXi boxes.



Note: One license key can be shared between different virtual centers. Therefore, license metric utilization might not represent the entire license consumption.

- Host name, for example: *198.51.100.0*

To enable the **Instance** column, hover over the **Settings** icon , and click **Configure View**.

To disable the license metric calculation of the VMware products, change the **vmman_slm_tags_enabled** parameter value to false. The collected data will no longer be retrieved, however, the historical data will still be displayed on the **Resource Utilization** report.

9.2.7

Enabling metric calculation of VMware products (advanced VM management)

9.2.7

Available from 9.2.7. You can calculate license metric utilization of the VMware products: vCenter and vSphere in complex environments where the VM managers are distributed across separated networks. Use advanced (distributed) VM management to enable data retrieval.

To collect granular information about the license metric utilization of VMware vCenter and vSphere

- Update the VM Manager Tool to the newest version. Metric calculation of VMware products is supported starting from version 1.5.0.0.
- Ensure you have BigFix Inventory version 9.2.1 or higher.
- Ensure that the user defined for the VMware connection has Global.Licenses permission.



Note: The API for retrieving the license data is available starting from vSphere API 4.0. The resource utilization is collected only for the VMware installations that are defined as VM managers for BigFix Inventory.

When you enable the license metric calculation of VMware vCenter and vSphere, the .slmtag files are created locally on the machine running VM Manager Tool and collected by the resource utilization scan.

1. On the computer where VM Manager Tool is installed, go to the `BES Client\LMT\VMMAN\config` directory.
2. Open and edit the `vmmmainconf.properties` file.
3. Set the `vmm_slm_tags_enabled` parameter to `true`.

```
vmm_slm_tags_enabled=true
```

If the parameter does not appear in the properties file yet, add the line above to this file. For more information, see: [VM Manager Tool settings \(on page ccclxxxii\)](#).

4. **Optional:** To specify the maximum size of each VMware .slmtag file, configure the `vmm_slm_tags_file_size_limit` parameter value.
5. **Optional:** To specify the maximum number of VMware .slmtag files that are stored on the computer for each license key under VMware instance, configure the `vmm_slm_tags_files_per_software` parameter value.
6. Save the file.
7. Go to `BES Client\LMT\VMMAN`.
8. Run the following command in the command-line interface.

- **Linux** `./vmmman.sh -reloadconfig`
- **Windows** `vmmman.bat -reloadconfig`

For more information, see: [VM Manager Tool command-line options \(on page ccclxxxix\)](#).

9. The data will be available in BigFix Inventory after the next resource utilization scan and data import.

To view license metric utilization, log in to BigFix Inventory and navigate to **Reports > Resource Utilization**.

The following columns are specific to the VMware vCenter and vSphere license metric utilization data:

- **Computer Name** always shows a name of the computer where the VM Manager Tool is installed.
- **Software Name** shows the name and the version of the software with additional information about license.
- **Metric Subtype** indicates if the count refers to the license metric utilization or overall license capacity. The column shows one of two values. USED refers to the license metric utilization of the specified license on a particular machine. CAPACITY reflects the limit agreed in this license. The number is specified in **Maximal Trend Value** column.

Computer Name	Software Name	Metric Type	Metric Subtype	Instance	Maximal Trend Value
NC9143126163	VMware vSphere 6 Enterprise Plus (unlimited cores per CPU)	SOCKET	USED	/P142Q-08L40-H8041-013RP-CHPQ1?VirtualCenter_198.51.100.0	12
NC9143126163	VMware vSphere 6 Enterprise Plus (unlimited cores per CPU)	SOCKET	CAPACITY	/P142Q-08L40-H8041-013RP-CHPQ1?VirtualCenter_198.51.100.0	999,999



Note: The maximal trend value for unlimited licenses is always set to 999,999.

- **Instance** shows the instance identifier that consists of:
 - License key, for example: *P142Q-08L40-H8041-013RP-CHPQ1*
 - Machine type: virtual center or ESX box.

Virtual center refers to the entire vSphere infrastructure for which VM manager connection is defined. The number of used licenses for virtual center shows the license metric utilization for all the boxes within this center. Virtual center can consist of more than one ESX, or ESXi box. VM manager can connect directly to the ESX, or ESXi box. Then, the displayed license metric utilization is calculated for each individual box.

License utilization is reported for the VMware connection that is defined in BigFix Inventory. Thus, if you define the virtual center connection, the .slmtag files are generated for this center, not individual ESX, or ESXi boxes.



Note: One license key can be shared between different virtual centers. Therefore, license metric utilization might not represent the entire license consumption.

- Host name, for example: *198.51.100.0*

To enable the **Instance** column, hover over the **Settings** icon , and click **Configure View**.

To disable the license metric calculation of the VMware products, change the **vmm_slm_tags_enabled** parameter value to false. The collected data will no longer be retrieved, however, the historical data will still be displayed on the **Resource Utilization** report.

Discovering and measuring SAP

You can discover components of SAP Business Suite that are installed on Linux and AIX computers in your infrastructure. You can also collect license metric utilization data from SAP landscapes and integrate it with your other resources on the Resource Utilization report. Discovery and measuring are two separate processes that do not impact each other.

Discovering

BigFix Inventory discovers components of SAP Business Suite that are installed on Linux and AIX computers. These components are listed on the Software Classification report, and can be assigned to appropriate products. It allows you to manage SAP products that are installed in your environment.

Measuring

The dedicated SAP Metric Data Collector collects information about the license metric utilization of SAP ABAP-based products. The data is available on the Resource Utilization report. The report shows the maximum metric utilization and its trend value over the last 30 days.

9.2.7

AIX

Linux

Discovering SAP components

9.2.7

Available from 9.2.7. To discover SAP components on Linux and AIX computers in your infrastructure, run the **Get SAP Components** fixlet. The data is collected, imported to BigFix Inventory, and displayed on the Software Classification panel.

The discovery of SAP components is supported on Linux and AIX systems.

BigFix Inventory discovers SAP Business Suite components that are based on SAP NetWeaver 7.x; for example, the application server, cross-application components and applications. For more information about SAP component hierarchy, see: [Software Components in SAP ERP](#).

If multiple instances of SAP are installed on the same computer, all of these instances are discovered. The task is relevant only to the computers where SAP is installed.

1. Log in to BigFix console.
2. In the navigation bar, click **Sites > External Sites > BigFix Inventory > Fixlets and Tasks**.
3. In the upper right pane, select **Get SAP Components**.
4. **Optional:** Select the **Retrieve information about SAP subcomponents** check box to additionally discover SAP sub-components.
5. To run the task, click **Take Action**. On the **Target** tab, select computers from which you want to retrieve the information, and click **OK**.
6. Ensure that the software scans are enabled, and run automatically. If not, [run the software scan \(on page cciii\)](#) manually. Select the **Software identification tags scan**.
7. Ensure that the scan results are automatically uploaded. If not, [upload the results \(on page ccviii\)](#) to the BigFix server manually.
8. To make the data available in BigFix Inventory, wait for the scheduled import of data, or run it manually.

The discovered SAP components are listed on the Software Classification panel. To view the data, log in to BigFix Inventory, and go to **Reports > Software Classification**. The component name is extended with a short description. You can classify these components by manually assigning them to appropriate products.

Publisher Name	Component Name	Version	Product Name	Metric	Computer Name	Installation Path	Details
SAP	SAP_HR (Human Resources)	604	<not set>	<n/a>	nc9098038016	<See details>	DETAILS >
SAP	SAP_BW (SAP Business Warehouse)	731	<not set>	<n/a>	nc9098038016	<See details>	DETAILS >
SAP	SAP_BS_FND (SAP Business Suite Foundation)	731	<not set>	<n/a>	nc9098038016	<See details>	DETAILS >
SAP	SAP_BASIS (SAP Basis Component)	731	<not set>	<n/a>	nc9098038016	<See details>	DETAILS >
SAP	SAP_APPL (Logistics and Accounting)	606	<not set>	<n/a>	nc9098038016	<See details>	DETAILS >
SAP	SAP_AP (SAP Application Platform)	700	<not set>	<n/a>	nc9098038016	<See details>	DETAILS >
SAP	SAP_ABA (Cross-Application Component)	731	<not set>	<n/a>	nc9098038016	<See details>	DETAILS >
SAP	SAP kernel	720	<not set>	<n/a>	nc9098038016	<See details>	DETAILS >

In case of any inconsistencies or issues, log in to BigFix console, select **Actions**, and choose **Get SAP Components** from the list. Check the exit code:

- 0 - Completed. Action successful.
- 1 - Failed. Action failed.
- 2 - Completed. Action partially completed. Some SAP instances did not return any data.

For more information, see the following log file that is located on the computer where SAP is installed:

```
/var/opt/BESClient/LMT/SAP/discover_sap.log
```

9.2.9 Measuring and collecting the license metric utilization of SAP

9.2.9 Available from 9.2.9. BigFix Inventory allows you to retrieve granular information about license metric utilization of SAP ABAP-based products. The utilization data is collected by the SAP Metric Data Collector and displayed on the Resource Utilization report. SAP Metric Data Collector retrieves the license utilization data available in License Administration Workbench in SAP.

9.2.9 Collecting utilization of SAP license metrics

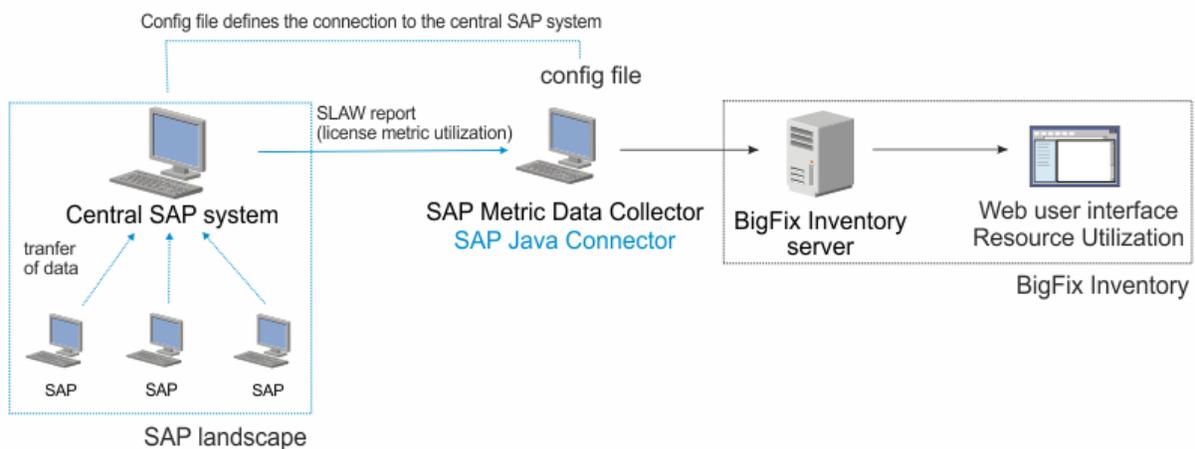
9.2.9 Available from 9.2.9. Thanks to SAP Metric Data Collector, you can view the license metric utilization of SAP ABAP-based products in BigFix Inventory. In this scenario, the utilization data is retrieved from the central SAP system, uploaded into BigFix Inventory, and displayed on the Resource Utilization report.

- You can install SAP Metric Data Collector on any computer, not necessarily the computer where the BigFix Inventory server or SAP is installed. However, the computer must meet the following requirements:
 - Runs on Windows x86 (64-bit) or Red Hat Linux x86 (64-bit). Linux version requires the unzip executable.
 - Has at least 250 MB of free disk space.
 - Can connect through RFC to the central SAP system that consolidates utilization data in SAP License Administration Workbench (LAW).
- Ensure that the SAP landscape meets the following criteria:

- The licensing data that is consolidated in the central SAP system is up-to-date.
- A user account with permissions to run end-to-end functions of LAW or LAW 2.0 exists on the central SAP system.
- License utilization data can be processed with either LAW or LAW 2.0.

The SAP ABAP-based products use the licensing model that consists of two types of licenses: software packages and named users.

Each SAP system measures utilization of SAP license metrics. The measurements from all SAP systems in the landscape are transferred to the central SAP and combined into one LAW report. To collect the report, SAP Metric Data Collector establishes an RFC connection with the SAP License Administration Workbench (LAW) that is located on the central SAP system. Information about license metric utilization from the report is displayed in BigFix Inventory on the Resource Utilization report.



9.2.12 Support for License Administration Workbench 2.0

LAW 2.0 allows you to manually declare engine usage information. SAP Metric Data Collector retrieves these self-declared details as well.

LAW 2.0 allows the creation and storage of multiple consolidations. SAP Metric Data Collector uses the current consolidation if it is in the Consolidated status. If no current consolidation is applicable, the tool retrieves the completed consolidation that has been modified last. The completed consolidation must be in Sent to SAP or Closed status.

SAP Metric Data Collector supports LAW and LAW 2.0. If both versions are installed on the central SAP system, SAP Metric Data Collector determines which instance of LAW to use by checking the target of RFC actions in LAW 2.0 settings.

1. Install SAP Metric Data Collector.
 - a. Log in to the BigFix console.
 - b. In the navigation bar, click **Sites > External Sites > BigFix Inventory > Fixlets and Tasks**.

- c. In the upper right pane, select **Install SAP Metric Data Collector**, and click **Take Action**.
 - d. Select the computer on which you want to install SAP Metric Data Collector, and click **OK**.
2. Download the 64-bit version of the SAP Java Connector from <http://support.sap.com/en/product/connectors.html>. Unpack it to the following directory on the computer where you installed SAP Metric Data Collector:

- **Linux** `/var/opt/BESClient/LMT/SAPTOOL/lib/sapjco3`
- **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\lib\sapjco3`

The following files, which are included in the package, are required to collect SAP utilization:

- **Linux** `sapjco3.jar, libsapjco3.so`
- **Windows** `sapjco3.jar, sapjco3.dll`

Make sure that these files are located directly in the `sapjco3` directory.



Note: Log in to <http://support.sap.com/en/product/connectors.html> with your SAP S User ID and password.

3. If you use LAW 2.0, install the ABAP plug-in that is delivered with SAP Metric Data Collector on a central SAP system.

- a. Copy the following file from the SAP Metric Data Collector installation directory:

- **Linux** `/var/opt/BESClient/LMT/SAPTOOL/trans/cofiles/K900041.BF1`
- **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\trans\cofiles\K900041.BF1`

- b. Paste the file to the `trans/cofiles` directory in the directory where SAP is installed on the central SAP system.

- c. Copy the following file from the SAP Metric Data Collector installation directory:

- **Linux** `/var/opt/BESClient/LMT/SAPTOOL/trans/data/R900041.BF1`
- **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\trans\data\R900041.BF1`

- d. Paste the file to the `trans/data` directory in the directory where SAP is installed on the central SAP system.

- e. Run the `STMS` SAP t-code on the central SAP system and import the incoming transport request. Select the Ignore Invalid Component Version option if needed. After the import, check the transport logs to ensure that it was successful.

4. Manually configure SAP Metric Data Collector. Use the `template.sapsystem` file to create the `<host_name>.sapsystem` file that defines the central SAP system in the landscape.



Important: You must not remove the `template.sapsystem` file from the `config` directory.

- a. Open the `template.sapsystem` file that is located in the following directory.
 - **Linux** `/var/opt/BESClient/LMT/SAPTOOL/config`
 - **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\config`
 - b. Save the file as `<host_name>.sapsystem` in the `config` directory. Do not overwrite the template.
 - c. Specify all parameters, and save the file. For more information, see: [SAP Metric Data Collector configuration parameters \(on page dxxiv\)](#).
5. To complete the configuration of SAP Metric Data Collector, edit the `saptool_config.properties` file that is located in the following directory:
- **Linux** `/var/opt/BESClient/LMT/SAPTOOL/config`
 - **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\config`
- a. Specify the **landscape_name** parameter. The landscape name is displayed on the Resource Utilization report in the Instance column.
 - b. To collect data without triggering measurements, set the value of the **skip_measurements** parameter to true and save the file. If you want to automate the measurements, see: [Measuring and collecting utilization of SAP license metrics \(on page dxx\)](#).
- For more information about configuration, see: [SAP Metric Data Collector configuration parameters \(on page dxxiv\)](#).
6. **Optional:** To test the connection to SAP and the correctness of data provided in the configuration file, run the following script with the **-testconnection** parameter.

- **Linux** `/var/opt/BESClient/LMT/SAPTOOL/saptool.sh -testconnection`
- **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\saptool.bat -testconnection`

Apart from testing the connection, the **-testconnection** parameter encrypts the password of the SAP user account. If you do not run the connection test, the password will be encrypted when you run SAP Metric Data Collector.

To verify whether the connection test was successful, check the return code in the `trace.log` file that is located in the following directory:

- **Linux** `/var/opt/BESClient/LMT/SAPTOOL/logs`
- **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\logs`

The return code that is provided in the log file can differ from the code in SAP Metric Data Collector Information analysis. The analysis does not provide a return code for the connection test. For more

information and the meaning of the return codes, see: [Troubleshooting SAP Metric Data Collector \(on page dxxv\)](#).

7. Initiate the collection of SAP metric data.
 - a. Log in to the BigFix console.
 - b. In the navigation bar, click **Sites > External Sites > BigFix Inventory > Fixlets and Tasks**.
 - c. In the upper right pane, select **Initiate Collection of SAP Metric Data**, and click **Take Action**.
 - d. Select the computer on which SAP Metric Data Collector is installed, and click **OK**.

 **Tip:** By default, the Initiate Collection of SAP Metric Data fixlet is scheduled to run every 30 days. The fixlet starts the collection of the utilization data, which can still be in progress after the fixlet is completed.

To schedule the fixlet to run periodically, select it, click **Take Action** and go to the **Execute** tab. Select the **Reapply this action** check box, and specify the frequency.

To check the status of SAP metric data collection at any point, activate the SAP Metric Data Collector Information analysis.

8. Ensure that the software scans are enabled, and run automatically. If not, [run the software scan \(on page cciii\)](#) manually. Select the **Resource utilization scan**.
9. Ensure that the scan results are automatically uploaded. If not, [upload the results \(on page ccviii\)](#) to the BigFix server manually.
10. To make the data available in BigFix Inventory, wait for the scheduled import of data, or run it manually.

License metric utilization data is displayed on the Resource Utilization report. The report shows the following metrics:

- Named users metric
- Other metrics that are related to the software packages license

On the SAP audit report, metrics are reported for a specified period, or not. BigFix Inventory reports these metrics as follows.

- If a metric is reported by SAP for a specified period, BigFix Inventory reports this metric for that period.
- If a metric is not reported by SAP for a specified period, BigFix Inventory reports this metric for the period since the previous collection of SAP metric data.
- If a metric is not reported by SAP for a specified period and you collect the SAP metric data for the first time, BigFix Inventory reports this metric for the last 30 days.

Computer Na...	Software Name	Metric Type	Metric Subtype	First Measured	Last Measured	Maximal Trend Value	Maximal Trend Value History 2017-06-26 - 2017-07-25
nc9143126222...	SAP - (0100) mySAP Human Resources	(0100) Master records PD Organizational Mgmt	<no data>	07/20/2017 09:52 AM	07/21/2017 09:47 AM	3,796	
nc9143126222...	SAP - (0100) mySAP Human Resources	(0170) Master records PA Benefits	<no data>	07/20/2017 09:52 AM	07/21/2017 09:47 AM	1,505	
nc9143126222...	SAP - (8010) SAP EHS: Product and Reach Compliance	(8010) SRC: Created Tasks	<no data>	07/20/2017 09:52 AM	07/24/2017 05:53 PM	1,144	
nc9143126222...	SAP - (3110) Information Lifecycle Management	(3111) ILM - Database Size in TB	<no data>	07/20/2017 09:52 AM	07/24/2017 05:53 PM	192	

Related information

[Changing the default secret key and password to the SAP Metric Data Collector keystore \(on page dccxlii\)](#)

9.2.10 Measuring and collecting utilization of SAP license metrics

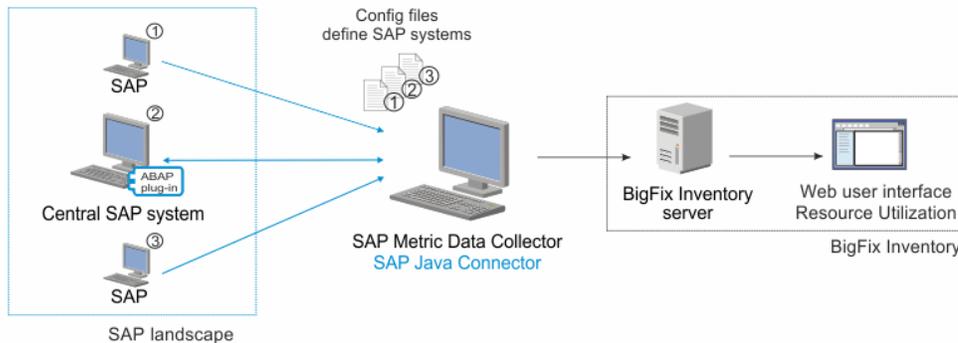
9.2.10 Available from 9.2.10. Thanks to SAP Metric Data Collector and the installation of ABAP plug-in, you can automatically measure the license metric utilization of SAP ABAP-based products, and then view the data on the Resource Utilization report in BigFix Inventory. In this scenario, SAP Metric Data Collector connects to each SAP system in the landscape and initiates measurements. Then, the tool manages the consolidation of data and integrates the results with BigFix Inventory.

- You can install SAP Metric Data Collector on any computer, not necessarily the computer where the BigFix Inventory server or SAP is installed. However, the computer must meet the following requirements:
 - Runs on Windows x86 (64-bit) or Red Hat Linux x86 (64-bit). Linux version requires the unzip executable.
 - Has at least 250 MB of free disk space.
 - Can connect through RFC to all SAP systems to collect the data required for calculating the license metric utilization.
- You should install the HCL-provided ABAP plug-in only on the central SAP system.
- Ensure that the SAP landscape meets the following criteria:
 - Configuration of Clients, Price Lists and User Types in System Measurement dialog, which can be accessed via `USMM` transaction code (t-code), is up-to-date on all the SAP systems.
 - A dedicated user account with permissions to run `USMM` transaction exists on each SAP system.
 - A user account with permissions to run `USMM` transaction and end-to-end functions of License Administration Workbench (LAW) exists on the central SAP system.
 - License utilization data must be processed with LAW. If you use LAW 2.0, see: [Collecting utilization of SAP license metrics \(on page dxv\)](#).

The SAP ABAP-based products use the licensing model that consists of two types of licenses: software packages and named users. BigFix Inventory uses a native SAP functionality for determining utilization of these license metrics called System Measurement and License Administration Workbench.

SAP Metric Data Collector connects to all SAP systems in the landscape through RFC, automatically initiates measurements, and retrieves the results. Next, SAP Metric Data Collector uses the ABAP plug-in to upload the results onto the central SAP system, and manage automatic user grouping and data consolidation in License Administration Workbench to produce a LAW report. SAP Metric Data Collector gathers the report data, and transfers it to BigFix

Inventory. Information about license metric utilization from the report is displayed in BigFix Inventory on the Resource Utilization report .



1. Install SAP Metric Data Collector.
 - a. Log in to the BigFix console.
 - b. In the navigation bar, click **Sites > External Sites > BigFix Inventory > Fixlets and Tasks**.
 - c. In the upper right pane, select **Install SAP Metric Data Collector**, and click **Take Action**.
 - d. Select the computer on which you want to install SAP Metric Data Collector, and click **OK**.
2. Download the 64-bit version of the SAP Java Connector from <http://service.sap.com/connectors>. Unpack it to the following directory on the computer where you installed SAP Metric Data Collector:
 - **Linux** `/var/opt/BESClient/LMT/SAPTOOL/lib/sapjco3`
 - **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\lib\sapjco3`

The following files, which are included in the package, are required to collect the SAP utilization data:

- **Linux** `sapjco3.jar, libsapjco3.so`
- **Windows** `sapjco3.jar, sapjco3.dll`

Make sure that these files are located directly in the `sapjco3` directory.



Note: Log in to <http://service.sap.com/connectors> with your SAP S User ID and password. You must comply with the license terms of the SAP Connectors.

3. Install the ABAP plug-in that is delivered with SAP Metric Data Collector on a central SAP system.
 - a. Copy the following file from the SAP Metric Data Collector installation directory:
 - **Linux** `/var/opt/BESClient/LMT/SAPTOOL/trans/cofiles/K900041.BF1`
 - **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\trans\cofiles\K900041.BF1`
 - b. Paste the file to the `trans/cofiles` directory in the directory where SAP is installed on the central SAP system.

c. Copy the following file from the SAP Metric Data Collector installation directory:

- **Linux** /var/opt/BESClient/LMT/SAPTOOL/trans/data/R900041.BF1
- **Windows** C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\trans\data\R900041.BF1

d. Paste the file to the `trans/data` directory in the directory where SAP is installed on the central SAP system.

e. Run the `STMS` SAP t-code on the central SAP system and import the incoming transport request. Select the Ignore Invalid Component Version option if needed. After the import, check the transport logs to ensure that it was successful.

4. Define the connection to each SAP system in the landscape by configuring a set of parameters. Use the `template.sapsystem` file to create the `<host_name>.sapsystem` file with all the required information for each SAP system in the landscape.



Important: You must not remove the `template.sapsystem` file from the `config` directory.

a. Open the `template.sapsystem` file that is located in the following directory:

- **Linux** /var/opt/BESClient/LMT/SAPTOOL/config
- **Windows** C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\config

b. Save the file as `<host_name>.sapsystem` in the `config` directory. Do not overwrite the template.

c. Specify all parameters, and save. For more information, see: [SAP Metric Data Collector configuration parameters \(on page dxxiv\)](#).

d. Create a separate file for each SAP system in the landscape.

5. To complete the configuration of SAP Metric Data Collector, edit the `saptool_config.properties` file that is located in the following directory:

- **Linux** /var/opt/BESClient/LMT/SAPTOOL/config
- **Windows** C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\config

a. Specify the `landscape_name` parameter and save the file. The landscape name is displayed on the Resource Utilization report in the Instance column.

b. Ensure that the value of the `skip_measurements` parameter is set to false.

6. **Optional:** To test the connection to SAP and the correctness of data provided in the configuration file, run the following script with the `-testconnection` parameter.

- **Linux** /var/opt/BESClient/LMT/SAPTOOL/saptool.sh -testconnection
- **Windows** C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\saptool.bat -testconnection

Apart from testing the connection, the **-testconnection** parameter encrypts the password of the SAP user account. If you do not run the connection test, the password will be encrypted when you run SAP Metric Data Collector.

To verify whether the connection test was successful, check the return code in the `trace.log` file that is located in the following directory:

- **Linux** `/var/opt/BESClient/LMT/SAPTOOL/logs`
- **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\logs`

The return code that is provided in the log file can differ from the code in the SAP Metric Data Collector Information analysis. The analysis does not provide a return code for the connection test. For more information and the meaning of the return codes, see: [Troubleshooting SAP Metric Data Collector \(on page dxv\)](#).

7. Initiate the measurement and collection of SAP metric data.

- a. Log in to the BigFix console.
- b. In the navigation bar, click **Sites > External Sites > BigFix Inventory > Fixlets and Tasks**.
- c. In the upper right pane, select **Initiate Collection of SAP Metric Data**, and click **Take Action**.
- d. Select the computer on which SAP Metric Data Collector is installed, and click **OK**.



Tip: By default, the Initiate Collection of SAP Metric Data fixlet is scheduled to run every 30 days. The fixlet starts the collection of the utilization data, which can still be in progress after the fixlet is completed.

To schedule the fixlet to run periodically, select it, click **Take Action** and go to the **Execute** tab. Select the **Reapply this action** check box, and specify the frequency.

To check the status of SAP metric data collection at any point, activate the SAP Metric Data Collector Information analysis.

8. Ensure that the software scans are enabled, and run automatically. If not, [run the software scan \(on page cciii\)](#) manually. Select the **Resource utilization scan**.
9. Ensure that the scan results are automatically uploaded. If not, [upload the results \(on page ccviii\)](#) to the BigFix server manually.
10. To make the data available in BigFix Inventory, wait for the scheduled import of data, or run it manually.

License metric utilization data is displayed on the Resource Utilization report. The report shows the following metrics:

- Named users metric
- Other metrics that are related to the software packages license

On the SAP audit report, metrics are reported for a specified period, or not. BigFix Inventory reports these metrics as follows.

- If a metric is reported by SAP for a specified period, BigFix Inventory reports this metric for that period.
- If a metric is not reported by SAP for a specified period, BigFix Inventory reports this metric for the period since the previous collection of SAP metric data.
- If a metric is not reported by SAP for a specified period and you collect the SAP metric data for the first time, BigFix Inventory reports this metric for the last 30 days.

Computer Na...	Software Name	Metric Type	Metric Subtype	First Measured	Last Measured	Maximal Trend Value	Maximal Trend Value History 2017-06-26 - 2017-07-25
nc9143126222...	SAP - (0100) mySAP Human Resources	(0100) Master records PD Organizational Mgmt	<no data>	07/20/2017 09:52 AM	07/21/2017 09:47 AM	3,796	
nc9143126222...	SAP - (0100) mySAP Human Resources	(0170) Master records PA Benefits	<no data>	07/20/2017 09:52 AM	07/21/2017 09:47 AM	1,505	
nc9143126222...	SAP - (8010) SAP EHS: Product and Reach Compliance	(8010) SRC: Created Tasks	<no data>	07/20/2017 09:52 AM	07/24/2017 05:53 PM	1,144	
nc9143126222...	SAP - (3110) Information Lifecycle Management	(3111) ILM - Database Size in TB	<no data>	07/20/2017 09:52 AM	07/24/2017 05:53 PM	192	

Related information

[Changing the default secret key and password to the SAP Metric Data Collector keystore \(on page dccxljii\)](#)

9.2.9 SAP Metric Data Collector configuration parameters

The configuration files contain important information that SAP Metric Data Collector requires to connect to the SAP systems from which the tool collects the license metric utilization data.

The configuration files are located in the following directory:

- **Linux** `/var/opt/BESClient/LMT/SAPTOOL/config`
- **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\config`

List of main configuration parameters that define SAP systems

The following parameters define the SAP systems in your environment. They are configured based on the `template.sapsystem` file. Do not overwrite the template file. Save the configuration in a new file named `<host_name>.sapsystem`.

Table 47. List of main configuration parameters

Parameter	Description
jco.client.ashost	Hostname or IP address of the SAP system.
jco.client.user	SAP user account name.
jco.client.passwd	SAP user account password.
jco.client.client	SAP client number.

Table 47. List of main configuration parameters (continued)

Parameter	Description
<code>jco.client.sysnr</code>	SAP system number.
<code>jco.client.lang</code>	Connection language; for example: EN or DE.
<code>is_law</code>	Confirm whether this is the central SAP system with consolidated LAW data. Only one system in the landscape must have this parameter set to true.

Other parameters

Other parameters that are required to collect SAP utilization data should be configured in a separate configuration file `saptool_config.properties`. After you set up the value of these parameters, overwrite the file by saving.

Table 48. List of other parameters

Parameter	Description
<code>landscape_name</code>	A custom name of the SAP landscape that is used to distinguish the SAP landscape in BigFix Inventory. The landscape name is displayed on the Resource Utilization report in a Instance column.
<code>skip_measurements</code>	<p>By default, this parameter is set to false. If you only need to collect SAP utilization data, without triggering measurements, set the value of the parameter to true and save the file. If you use the automated procedure, do not change the default setting of the parameter.</p> <p> Note: The <code>skip_measurements</code> parameter is relevant only if you use LAW. In case of LAW 2.0 measurements and consolidation are always skipped.</p>

9.2.9 Troubleshooting SAP Metric Data Collector

9.2.9 Available from 9.2.9. If you encounter problems with measuring utilization of SAP license metrics by using SAP Metric Data Collector, check the log and the return code to determine the possible cause.

Log file

The `trace.log` file is located in the following directory.

- **Linux** `/var/opt/BESClient/LMT/SAPTOOL/logs`
- **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\logs`

Return codes

If the **Initiate SAP Metric Data Collector** task fails, an error code that indicates why the scan failed is returned. The return code is included in the log file.

You can also check the status and the return code in the BigFix console. Go to **Sites > External Sites > BigFix Inventory v9 > Analyses**, and activate the **SAP Metric Data Collector Information** analysis. The analysis does not provide a return code for the connection test.

Table 49. SAP Metric Data Collector return codes

Return code	Possible cause and solutions
0	No errors.
1	An unknown error occurred. Contact BigFix Support.
2	Writing in to the <code>trace.log</code> file could not be initiated. Ensure that the file is not used by other processes and that it is not opened in any program, for example a text editor.
3	Encryption module could not be initiated. Ensure that the keystore and keystore password are correctly configured. For more information, see: Changing the default secret key and password to the SAP Metric Data Collector keystore (on page dccxlii) .
4	The specified command-line parameters are incorrect. Run the <code>saptool.bat</code> or <code>sap-tool.sh</code> script with the <code>-h</code> parameter to see the list of supported command-line parameters.
5	Configuration settings are invalid or missing. Ensure that all required parameters are specified and correct. For more information, see: SAP Metric Data Collector configuration parameters (on page dxxiv) .
6	Connection to SAP failed. Ensure that connection to SAP is possible and that all required parameters are specified and correct. For more information, see: SAP Metric Data Collector configuration parameters (on page dxxiv) .
7	The SAP Metric Data Collector process was interrupted. Run SAP Metric Data Collector again.
8	Generating <code>.slmtag</code> files failed. Ensure that files in the <code>slmtags</code> subdirectory are not locked by other processes.
9	Measuring utilization data in SAP exceeded the time limit.
10	Grouping SAP users in License Administration Workbench (LAW) failed because SAP user grouping mode is not set to automatic. Set the mode to automatic, and run SAP Metric Data Collector again.
11	The data collected from SAP is incomplete. Ensure that the consolidated results are available in LAW or LAW 2.0, and then run the SAP Metric Data Collector again. The measure-

Table 49. SAP Metric Data Collector return codes (continued)

Return code	Possible cause and solutions
	ment results can be retrieved and consolidated automatically in LAW when the skip_measurements parameter is set to false. If the problem persists, contact BigFix Support.
12	The status of the SAP Metric Data Collector completion could not be stored on disk. Ensure that the <code>saptool.status</code> file is not used by other processes and that it is not opened in any program, for example a text editor.
13	The configuration file could not be updated with the encrypted password. Ensure that none of the configuration files in the <code>config</code> subdirectory is used by other processes and that it is not opened in any program, for example a text editor.
14	The SAP Metric Data Collector configuration that you initiated is not supported. Contact BigFix Support.
15	SAP Java Connector is corrupted or missing. Download the SAP Java Connector again. For more information, see: Collecting utilization of SAP license metrics (on page dxv) .
16	Currently not used.
17	SAP Metric Data Collector was run with the <code>-encode</code> parameter by using a method other than the command line. Run the tool through the command line to encode keystore password. For more information see: Changing the default secret key and password to the SAP Metric Data Collector keystore (on page dccxlii) .
18	SAP Metric Data Collector tried to decrypt the stored password by using a key other than the one that was used to encrypt it. The encryption key was changed. Re-enter the password in plain text. The password will be encrypted the next time the tool is run. For more information see: Changing the default secret key and password to the SAP Metric Data Collector keystore (on page dccxlii) and Collecting utilization of SAP license metrics (on page dxv) .
19	SAP Metric Data Collector tried to encrypt the stored password by using a key other than AES or the key was stronger than the Java JCE jurisdiction policy allows. The encryption key was changed. Change the encryption key to AES with correct length. For more information see: Changing the default secret key and password to the SAP Metric Data Collector keystore (on page dccxlii) .
20	A given SAP system could not find the requested function module.
21	A function module on a given SAP system failed.
22	An error occurred when cleaning the last measurement data from the central SAP system.
23	The ABAP plug-in is not installed on the central SAP system in the landscape. For information about the installation of the code, see: Installing the ABAP plug-in on a central SAP system. (on page dxxi)

Table 49. SAP Metric Data Collector return codes (continued)

Return code	Possible cause and solutions
24	SAP Metric Data Collector could not find the consolidated SAP utilization data.

Assigning components to active BigFix products

With BigFix Inventory you can assign components to active BigFix products. After the first discovery of the component, you must manually assign the component to the respective product. When the same component is discovered again, it is automatically assigned to the same product that it was assigned earlier.

To apply automated component assignments, complete the following steps:

1. Categorize the clients by running one of following fixlets:
 - Categorize ClientDevice clients
 - Categorize MVS clients
 - Categorize RVU clients
2. Check "License Overrides" analysis and displaying the current license metrics.
3. Run the "Distributing the site mapping file" fixlet.
4. Run the "Generating the license tags" fixlet.

You must run the "Initiate Software Scan" and "Upload Software Scan Results" fixlets so that the newly discovered components are visible.

For information on activating license counting process, refer to [Activating the license counting process](#).

You can discover the new components by clicking **Software Classification > Current Installations**.

Software Classification											
Current Installations Assign Exclude Confirm Edit 1 / 11 rows (filtered) Export Configure											
Publisher Na...	Component Name	Compon...	Compon...	Product Name	Metric	FlexPoint or ...	Entitled	Computer Name	Installation Path	Details	
<input type="checkbox"/>	HCL	BigFix Platform Agent	10.0	10.0.1.41	IBM BigFix Inventory	Client Device	None	No	NC9174039241	C:\Program Files (x86)\BigFix Enterprise\BES Client\swidtag	DETAILS >
<input checked="" type="checkbox"/>	HCL	BigFix Lifecycle MVS	10.0	10.0.0	<not set>	<N/A>	<N/A>	<N/A>	NC9174039241	C:\Program Files (x86)\BigFix Enterprise\BES Client\swidtag	DETAILS >
<input type="checkbox"/>	HCL	BigFix Compliance ClientDevice	10.0	10.0.0	<not set>	<N/A>	<N/A>	<N/A>	NC9174039241	C:\Program Files (x86)\BigFix Enterprise\BES Client\swidtag	DETAILS >
<input type="checkbox"/>	HCL	BigFix Patch MVS	10.0	10.0.0	<not set>	<N/A>	<N/A>	<N/A>	NC9174039241	C:\Program Files (x86)\BigFix Enterprise\BES Client\swidtag	DETAILS >
<input type="checkbox"/>	HCL	BigFix Compliance Payment Card Industry Add-on ClientDevice	10.0	10.0.0	<not set>	<N/A>	<N/A>	<N/A>	NC9174039241	C:\Program Files (x86)\BigFix Enterprise\BES Client\swidtag	DETAILS >
<input type="checkbox"/>	HCL	BigFix Platform Agent	9.5	9.5.16.90	<not set>	<N/A>	<N/A>	<N/A>	NC9174038204	/opt/BESClient/bin/swidtag	DETAILS >
<input type="checkbox"/>	IBM	IBM BigFix Compliance Payment Card Industry Add-on RVU	9.5	9.5.0	IBM BigFix Compliance Payment Card Industry Add-on	RVU MAPC	None	No	NC9174038204	/opt/BESClient/bin/swidtag	DETAILS >
<input type="checkbox"/>	IBM	IBM BigFix Compliance RVU	9.5	9.5.0	IBM BigFix Compliance	RVU MAPC	None	No	NC9174038204	/opt/BESClient/bin/swidtag	DETAILS >
<input type="checkbox"/>	IBM	IBM BigFix Patch MVS	9.5	9.5.0	IBM BigFix Patch	RVU MAPC	None	No	NC9174038204	/opt/BESClient/bin/swidtag	DETAILS >
<input type="checkbox"/>	IBM	IBM BigFix Inventory RVU	9.5	9.5.0	IBM BigFix Inventory	RVU MAPC	None	No	NC9174038204	/opt/BESClient/bin/swidtag	DETAILS >
<input type="checkbox"/>	IBM	IBM BigFix Lifecycle MVS	9.5	9.5.0	IBM BigFix Lifecycle	RVU MAPC	None	No	NC9174038204	/opt/BESClient/bin/swidtag	DETAILS >

Software classification

After new components are discovered, they are assigned to products based on bundling rules. If the initial assignment does not reflect your entitlements, work with the software inventory until it matches the entitlements.

To make metric calculation accurate, assign each component to a product, exclude, or suppress the instances that should not be included in the calculation, and confirm your actions. After you confirm the actions, the classification is complete, and the metric calculation is accurate.

Click a selected step to view detailed instructions

Step 1: Assign

Assign components to products from catalog, or create custom relations.

Step 2: Tailor

Exclude products, or suppress components to remove them from metric calculation.

Step 3: Confirm

Confirm your actions and flag the items as completed.



Bundling Assistant

Bundling Assistant instances of discovered components to products based on five bundling rules: bundling tags, part numbers, partition collocation, infrastructure collocation, and stand-alone product discovery. Results of this assignment are shown on the Software Classification panel.

Bundling Assistance process

After a component is discovered, it undergoes the process of bundling assistance. During the process, bundling rules are applied to discover possible component-product relations. Each relation is evaluated and assigned an internal score. The component is assigned to the product with the highest score.

If the discovered component belongs to the product and license metric assigned during bundling assistance, you can confirm the assignment. Otherwise, assign the component to a different product or license metric so that the relation reflects your entitlement, and confirm it. If the assignment is not confirmed, it might be automatically changed. It might happen when a new software catalog is uploaded or when a different product is suggested as the best bundling option based on the new results of bundling assistance.

! **Important:** Instances that are a part of an excluded product cannot be rebundled automatically. Instances can be bundled with excluded products only if they are associated with the excluded products based on the software catalog or the stand-alone product discovery. Any new instance that is assigned to the excluded product is also excluded.

Bundling rules

The following table lists rules that are used during the process of bundling assistance.

Table 50. Bundling rules

Rule	Description
Custom bundling	This rule assigns a component to the product and license metric based on the custom bundling that was defined manually.
9.2.5 Bundling tags (on page dxxx)	This rule assigns a component to the product and license metric that are represented by the part number listed in the bundling tag. The tag is placed on the

Table 50. Bundling rules (continued)

Rule	Description
Part numbers (on page dxxxiv)	computer where the component is installed which allows for assigning the specific instance of the component to the right product and license metric.
Partition collocation (on page dxxxix)	This rule uses part numbers that you upload to BigFix Inventory to narrow down the list of products and license metrics to which a component can be assigned to the ones that you purchased.
Infrastructure collocation (on page dxl)	This rule discovers components that are installed on the same partition, and are related to the same product in the software catalog.
Stand-alone product discovery (on page dxli)	This rule discovers components that are installed in the same infrastructure, and are related to the same product in the software catalog. The search does not include the computer on which the component is installed.
Stand-alone product discovery (on page dxli)	This rule identifies all products to which a component can be assigned and chooses the product for which the fewest components are listed in the software catalog.

Related information

[Main background application tasks \(on page lxxix\)](#)

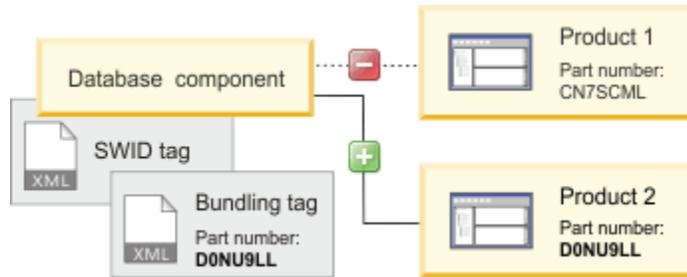
9.2.5 Bundling tags

9.2.5 Available from 9.2.5. Bundling tag is a type of the ISO/IEC 19770-2 supplemental tag. It allows for unambiguously assigning a component to a product and license metric based on the part number that is specified in the tag. When a bundling tag is available, it takes precedence over other bundling rules. Thus, you can create bundling tags to facilitate the process of software classification.

Example

A bundling tag was manually created for the database component. The tag contains a part number that represents the product and license metric to which the component should be assigned. The part number is correlated with the information from the software catalog. It is established that the part number represents Product 2.

The database component can be assigned to two products. Based on the part number from the bundling tag, the database component is assigned to Product 2. The assignment needs to be manually confirmed.



Details

- Bundling tags can be created only for IBM products.
- Bundling tags can be created only for products that are discovered by [SWID tags \(on page lxxxiii\)](#).
- The bundling tag must be placed in the same directory as the SWID tag.
- The bundling tag can contain only one part number.
- `PersistentId` in the bundling tag and SWID tag must be the same.

Supported schemas

BigFix Inventory supports bundling tags in the 19770-2:2014-CD1 and 19770-2:2015 schemas. However, the latter is recommended for creating bundling tags. For more information about each schema, see the following links:

- <http://standards.iso.org/iso/19770/-2/2014-CD1/schema.xsd>
- <http://standards.iso.org/iso/19770/-2/2015-current/schema.xsd>

Example of a bundling tag in the 19770-2:2015 schema

```
<?xml version="1.0" encoding="UTF-8"?>
<SoftwareIdentity
name="IBM DB2 Enterprise Server Edition"
tagId="ibm.com-b3414ccd1e364aa6b82f4e1c3adecdf6-D0NU9LL"
version="10.5.0"
versionScheme="multipartnumeric"
supplemental="true"
xmlns="http://standards.iso.org/iso/19770/-2/2015/schema.xsd"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://standards.iso.org/iso/19770/-2/2015-current/schema.xsd">
  <Meta productId="PN:D0NU9LL" />
  <Link rel="parent" href="swidpath://SoftwareIdentity
[Meta/@persistentId='b3414ccd1e364aa6b82f4e1c3adecdf6']" />
  <Entity name="BigFix" regid="ibm.com" role="licensor publisher" />
  <Entity name="customer" regid="customer.com" role="tagCreator"/>
</SoftwareIdentity>
```

Related information

[Support for ISO/IEC 19770-2 \(on page lxxxiii\)](#)

9.2.5 Manually creating bundling tags

9.2.5 Available from 9.2.5. You can facilitate the process of software classification by manually creating bundling tags. When a bundling tag is placed on the computer where a component is installed, the component is automatically assigned to the product that is specified in the bundling tag.

Creating bundling tags is most effective when your company uses a process of packaging software installers. In this case, you can add a bundling tag to the packaged installer, and automatically place it on the computers during the installation.

1. Find the SWID tag that is used to discover the component.
 - a. In the top navigation bar, go to **Reports > Software Installations**.
 - b. Find the component, and click the link in the **Details** column. Information about the SWID tag and its location is displayed at the bottom of the page.
2. Go to the computer where the component is installed and open the SWID tag.

To create a bundling tag, you need the following information from the SWID tag.

- `name`
- `version`
- `persistentId`

BigFix Inventory supports SWID tags in the 19770-2:2014-CD1 and 19770-2:2015 schemas. The information necessary to create the bundling tag is available in SWID tags in both schemas. The following example shows SWID tag in the 19770-2:2014-CD1 schema.

```
<?xml version="1.0" encoding="UTF-8"?>
<SoftwareIdentity name="IBM Control Desk"
  uniqueId="0ecd37302a7648a5aca7a26aec76731d-7.6.0"
  version="7.6.0"
  versionScheme="multipartnumeric"
  xmlns="http://standards.iso.org/iso/19770/-2/2014-CD1/schema.xsd"
  xmlns:ds="http://www.w3.org/2000/09/xmldsig#"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://standards.iso.org/iso/19770/-2/2014-CD1/schema.xsd
  swid.xsd">
  <Meta persistentId="0ecd37302a7648a5aca7a26aec76731d"/>
  <Meta taxonomyCode="SSWT9A"/>
  <Meta taggingProcess="4-1-20150603"/>
```

```
<Entity name="IBM" regid="regid.1986-03.com.ibm" role="licensor publisher tagCreator"/>
</SoftwareIdentity>
```

3. Open a text editor and create a bundling tag according to the following template.

```
<?xml version="1.0" encoding="UTF-8"?>
<SoftwareIdentity name="COMPONENT_NAME"
  tagId="ibm.com-PERSISTENT_ID-PART_NUMBER"
  version="VERSION"
  versionScheme="multipartnumeric"
  supplemental="true"
  xmlns="http://standards.iso.org/iso/19770/-2/2015/schema.xsd"
  xmlns:ds="http://www.w3.org/2000/09/xmldsig#"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://standards.iso.org/iso/19770/-2/2015-current/schema.xsd
  swid.xsd">
  <Meta productId="PN:PART_NUMBER" />
  <Link rel="parent" href="swidpath://SoftwareIdentity[Meta/@persistentId='PERSISTENT_ID']" />
  <Entity name="IBM" regid="ibm.com" role="licensor softwareCreator" />
  <Entity name="CUSTOMER" regid="CUSTOMER.COM" role="tagCreator" />
</SoftwareIdentity>
```

Where:

- `COMPONENT_NAME` is the name of the component as specified in the SWID tag
- `PERSISTENT_ID` is the persistent ID of the component as specified in the SWID tag
- `VERSION` is the version of the component as specified in the SWID tag
- `PART_NUMBER` is the part number that represents the product to which the component is to be assigned and its license metric
- `CUSTOMER` is the name of your organization
- `CUSTOMER.COM` is the registration identifier of your organization in the URI format. For more information, see: [Uniform Resource Identifier \(URI\): Generic Syntax](#).

A bundling tag that is created based on the SWID tag presented in step 2 looks as follows.

```
<?xml version="1.0" encoding="UTF-8"?>
<SoftwareIdentity name="BigFix Control Desk"
  tagId="ibm.com-0ecd37302a7648a5aca7a26aec76731d-D0LKPLL"
  version="7.6.0"
  versionScheme="multipartnumeric"
  supplemental="true"
  xmlns="http://standards.iso.org/iso/19770/-2/2015/schema.xsd"
  xmlns:ds="http://www.w3.org/2000/09/xmldsig#"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://standards.iso.org/iso/19770/-2/2015-current/schema.xsd
```

```

swid.xsd">
  <Meta productId="PN:D0LKPLL" />
  <Link rel="parent"
href="swidpath://SoftwareIdentity[Meta/@persistentId='0ecd37302a7648a5aca7a26aec76731d']" />
  <Entity name="IBM" regid="ibm.com" role="licensor softwareCreator" />
  <Entity name="MyCompany" regid="mycompany.com" role="tagCreator"/>
</SoftwareIdentity>

```

4. Save the bundling tag. Name it according to the following template:

```
<Tag_Creator_Regid>-<Component_Name>-<Part_Number>.swidtag
```

For example, `mycompany.com-IBM_Control_Desk-D0LKPLL.swidtag`.

5. Place the bundling tag in the same location as the SWID tag.
Place the bundling tag on all computers on which the component is to be assigned to the product specified in the bundling tag.
6. Wait for the next software scan and import of data.
7. Confirm the assignment of the components.



Note: You can easily assign all components that are bundled based on the bundling tag when you use the Software Classification panel.

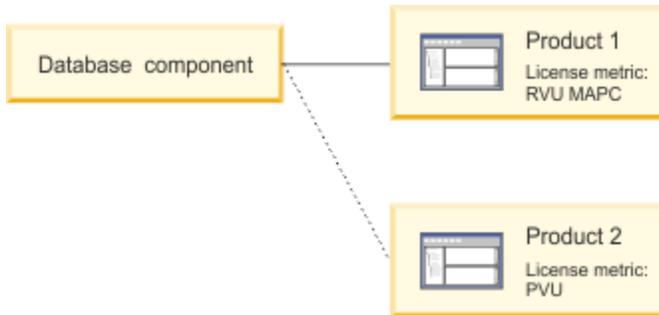
- a. In the top navigation bar, go to **Reports > Software Classification**.
- b. To filter out the report to components that were bundled by using bundling tags, hover over the **Manage Report View** icon , and click **Configure View**. Then, specify the following filter: `Bundling Tag Used, equal to, Yes`.
- c. Press **CTRL + A** to select all components.
- d. To confirm the assignment, hover over the **Confirm** icon , and click **Confirm**.

Part numbers

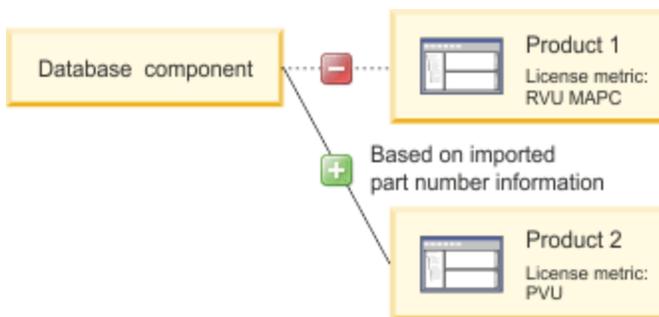
Part numbers represent software products that you purchased and license metrics that these products use. They can be used to increase the accuracy of automated bundling by narrowing down the number of possible products to which a discovered component can be assigned.

Example

Automated bundling assigns a database component to one of the products that were discovered in your infrastructure. However, the database component can also be potentially assigned to another product.



Before you confirm the assignment of the database component or change it manually, you can import the part numbers file to BigFix Inventory. Based on the file, and the component-product relations that are defined in the software catalog, the initial assignment of the database component is more accurate.



Types of part numbers

There are three types of part numbers:

- Catalog part numbers are provided by BigFix together with the software catalog. Each part number represents a product and a license metric that this product uses.
- **9.2.7** Custom part numbers are user-defined part numbers that are not included in the software catalog by default. They might be assigned to products with custom deals, or any additional software you wish to add to the software catalog.
- Uploaded part numbers are part numbers of products that you purchased. You can download the uploaded part numbers list from the Passport Advantage website or create it manually, and upload the file on the **Part Numbers Upload** panel. Part numbers that represents products that you purchased and their license metrics are matched against the part numbers that are provided with the software catalog to improve the results of automated bundling.

BigFix Inventory takes into consideration part numbers that start with the following letters:

- D - part numbers that represent License Entitlements
- E - part numbers that represent Support and Subscription

Limitations

The extent to which part numbers can improve automated bundling is limited by the following factors:

- Part numbers influence the assignment of only those components whose assignment is not confirmed
- Part numbers that were purchased before 2010 are not contained in the software catalog
- Part numbers that represent products that have limited availability are not contained in the software catalog
- Part numbers that represent products that are available only outside of the United States are not included in the catalog



Note: **9.2.7** You can add a custom part number to the software catalog for any existing product-license metric pair. This solution allows you to customize the software catalog, and ensure that products that were purchased before 2010, or are specific to your business and geographical location are additionally considered during automated bundling. For more information, see: [Adding custom part numbers \(on page dxxxvi\)](#).

9.2.7 Adding custom part numbers

9.2.7 Available from 9.2.7. Add custom part numbers to the software catalog to improve the results of automated bundling of software components.



You must have the View Software Catalog and Signatures and Manage Licenses permissions to perform this task.

You can add custom part numbers to any product-license metric pair existing in the software catalog. After adding custom part numbers to the catalog, make sure to include these custom part numbers in the part numbers file that is uploaded to BigFix Inventory. After the next data import, the custom part numbers are matched against the uploaded part numbers and used for automated bundling. For more information, see: [Uploading part numbers \(on page dxxxviii\)](#).

1. Log in to BigFix Inventory, and go to **Reports > Products & Metrics**.
2. Select the product for which you want to add a custom part number.
3. Hover over the **Edit Product** icon , and click **Edit Part Number**.
4. Provide a custom part number, or part numbers in the **Custom Part Number** field. Separate each number with a comma. The custom part number can contain letters, numbers and an underscore (_). Other characters are not allowed. Each part number needs to be unique. Thus, you cannot add a duplicate of an existing catalog part number.
5. Click **OK**.

The specified custom part numbers are added to the software catalog, and displayed in the **Catalog Part Numbers** column.

To remove custom part numbers, select the product, hover over the **Edit Product** icon , and click **Edit Part Number**. Then, remove custom part numbers from the list, and click **OK**.

Preparing the part numbers file

The part numbers file can be used to increase the accuracy of automated bundling by narrowing the number of possible products to which a discovered component can be assigned. You can download the file from the HCL Software Flexnet site or create it manually.

If you added custom part numbers to the software catalog, make sure that they are included in the part numbers file. Otherwise, the custom part numbers are not matched against the uploaded part numbers and are not used during automated bundling.

 **Restriction:** BigFix Inventory supports only CR LF and LF end of line characters. While editing the part number file do not use CR end of line characters specific for Mac systems.

- Download the file with part numbers from the HCL Software Flexnet site.
 1. Log in to the [Passport Advantage®](#).
 2. From the menu on the left, select **Reporting > Order history**.
 3. If applicable, select the relevant site or sites.
 4. In the **Type of order** section, select **All**.
 5. Specify the **Sales order date** and click **Submit detail report**. Your report is ready.
 6. To save the file, click **Download report** and then select **Comma delimited**. Ensure that the file is in the CSV or ZIP format.
- Prepare a custom CSV file that consists of two mandatory columns: `Part number` and `Order type`. Only the rows that contain the `Licenses` element in the `Order type` column are imported. The list of part numbers can include HCL Software Flexnet site part numbers, as well as your custom part numbers. You can add any other columns to the file, however, these columns are ignored during upload. Use a comma as a field separator and quotation marks (" ") as a text delimiter. For example:

```
"Part number", "Order type"
"D07UMLL", "Licenses"
"D08UMLL", "Licenses"
"Custom part number", "Licenses"
```

- Prepare a CSV file that contains a single column without a header. This file can only contain HCL Software Flexnet site part numbers, not your custom part numbers. List each part number in a new line. Do not use quotation marks (" ") to define part numbers. For example:

```
D07UMLL
D08UMLL
D09UMLL
```

Ensure that you use only valid part numbers that consist of seven characters. Upload of part numbers that are shorter or longer fails with the following message.

```
Upload failed. The uploaded csv file does not contain Part number or Order type column.
```

After you prepare the file, [upload the part numbers \(on page dxxxviii\)](#).

Uploading part numbers

Upload and import part numbers to increase the accuracy of automated bundling of software components.

 You must have the Manage Part Numbers Uploads permission to perform this task.

If you added custom part numbers to the software catalog, make sure that they are included in the part number file. Otherwise, the custom part numbers will not be matched against the uploaded part numbers and used during automated bundling.

1. [Prepare the part numbers file. \(on page dxxxvii\)](#)
2. In the top navigation bar, click **Management > Part Numbers Upload**.
3. **9.2.2** Choose the computer group for which you want to upload the part numbers.



Note: The part numbers are taken into account only for the software that is installed on the computers that belong to this group. They are also visible only to the user who is assigned to this computer group. For example, if you upload part numbers for a subgroup of the All Computers group, the part numbers are visible only to the user who is assigned to the subgroup and are not visible to the user who is assigned to the All Computers group.

4. Click **Browse** and choose the part numbers file to upload. You can upload a CSV or a ZIP file.
5. **Optional:** If you want to overwrite the existing part numbers, select **Overwrite existing part numbers**. Otherwise, the new part numbers are added to the existing ones.
The option is enabled only if you previously imported a part numbers file for the computer group that you selected in step 3.
6. Click **Upload**.
When you upload the file, a new entry is created in the Upload History table. The status is Pending until you run the import.



Important: If more than one entry has the Pending status, only the latest part numbers file is imported.

7. Wait for the scheduled import or run it manually.

The part numbers were imported to BigFix Inventory.

If the part numbers are not matched with products after you upload the part number file and run import, see: [Catalog problems. \(on page dcclx\)](#)

- **Removing the imported part numbers from BigFix Inventory server**

If you want to remove the part numbers from the server, go to **Management > Part Numbers Upload**, click **Remove All Part Numbers** and then run the import. You can remove the part numbers only if you previously imported the part numbers file to the server.

- **Checking what part numbers are currently imported to BigFix Inventory**

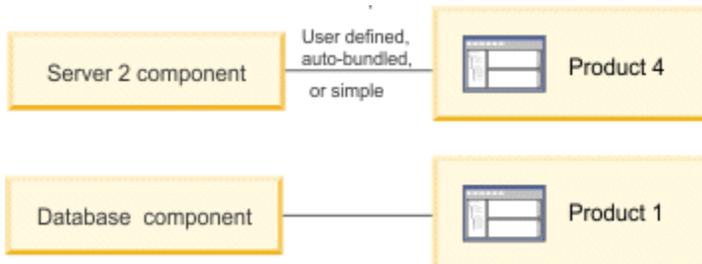
You can check what part numbers are currently imported to BigFix Inventory. To do it, go to **Reports > Products & Metrics** and configure the report view to display the **Imported Part Numbers** column. You can also narrow down the report to only see the products that are linked with the imported part numbers by setting the following filter: `Imported Part Numbers, is not empty.`

Partition collocation

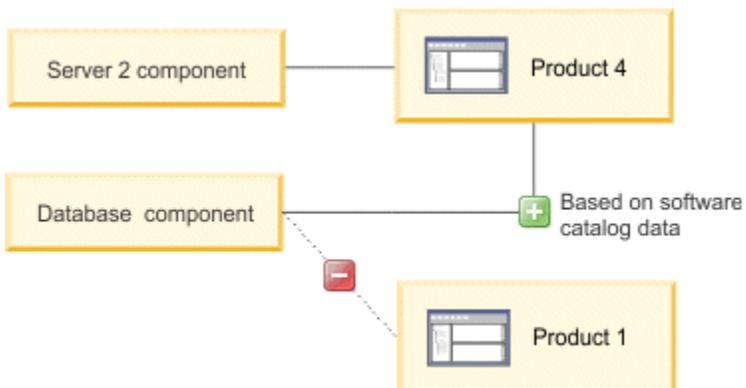
Partition collocation is a rule that discovers component instances that are on the same partition and are related to a product in the software catalog. If any of the discovered component instances has a confirmed, one-to-one, or auto-bundled relation to the product, the other instances are automatically bundled with the same product.

Example

The assignment of the Server 2 component to Product 4 is simple, auto-bundled, or manually confirmed. The database component is assigned to Product 1. Both components are on the same partition.



Based on the software catalog data, the database component is automatically reassigned to Product 4.



9.2.6 Components that are not evaluated

Some components are not evaluated as part of the partition collocation rule and do not affect other components that are installed on the same partition. The components include:

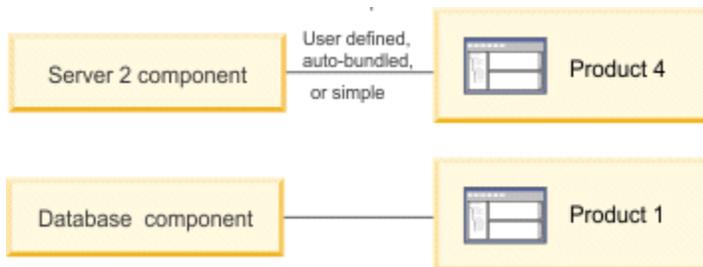
- Tivoli Monitoring for Databases - *<name>* Agent
- Tivoli Monitoring - *<name>* Agent

Infrastructure collocation

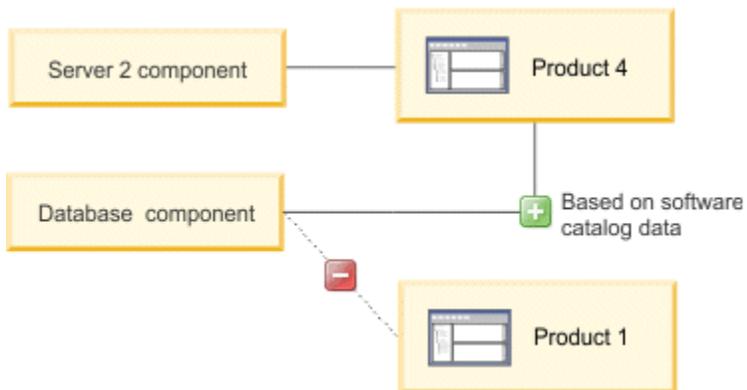
Infrastructure collocation is a rule that discovers component instances that are installed in the same infrastructure and that are related to a specific product in the software catalog. The search does not include the agent on which the product is installed. If any of the discovered component instances has a confirmed, one-to-one, or auto-bundled relationship with the product, the other instances are automatically bundled with the same product.

Example

The assignment of the Server 2 component to Product 4 is simple, auto-bundled, or manually confirmed. The database component is assigned to Product 1. Both components are in the same infrastructure.



Based on the software catalog data, the database component is automatically reassigned to Product 4.



9.2.6 Components that are not evaluated

Some components are not evaluated as part of the infrastructure collocation rule and do not affect other components that are installed in the same infrastructure. The components include:

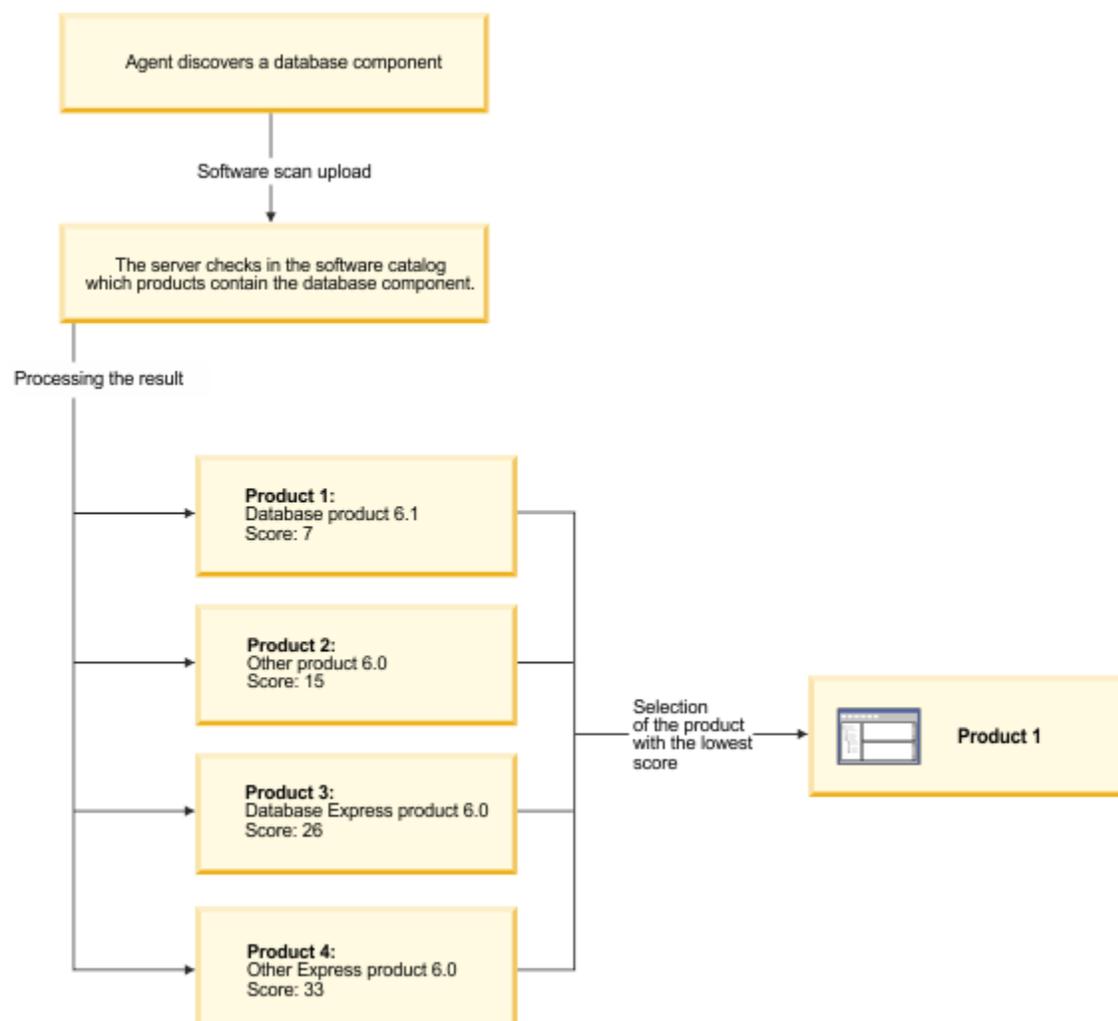
- Tivoli Monitoring for Databases - *<name>* Agent
- Tivoli Monitoring - *<name>* Agent

Stand-alone product discovery

Stand-alone product discovery is a rule that identifies components that have more than one default component to product relation. For each relation, it compares the name of the component with the name of the product and calculates the Levenshtein distance. During the calculation, production components are assigned to production products, and non-production components such as trial, demonstration, limited, or express versions are assigned to non-production products. The algorithm does not favor shorter names, but names that have more letters in common. Based on the calculations, the product-component relation with the shortest distance (represented by the lowest score) is selected. If a few relations have the same score, they are ordered alphabetically.

Example

The agent discovers the database component and the server checks which products contain this component in the software catalog. The database component is assigned to the product whose name produces a match with the lowest score, in this case, Product 1.



Software Classification panel

9.2.7 Available from 9.2.7. Software Classification panel provides an overview of your software inventory and the relations between the discovered software components and licensable products. It provides a flat structure with software installations broken into single components, which makes it easier to view your software assets and manage relations between them.



Note: The Software Classification panel replaced the IBM Software Classification report and is the only integrated place to manage the discovered software components and their relations with licensed products.

Permissions



You must have the Manage Software Classification permission to view the panel and classify software. To view some additional columns related to computers on which the components are installed, you also need the View Endpoints permission.

Panel overview

Before you start working with the discovered components, see the details of the Software Classification panel to understand the available options and columns.

● Current Installations							
Publisher...	Component Name	Version	Product Name	Metric	Computer Name	Installation Path	Details
IBM	IBM BigFix Inventory Server	9.2	IBM BigFix Inventory	RVU MAPC	NC042208	/varia	DETAILS >
IBM	IBM BigFix Inventory Server	9.2	IBM BigFix Inventory	RVU MAPC	NC042193	C:\varia	DETAILS >
IBM	IBM BigFix Inventory Server	9.2	IBM BigFix Inventory	RVU MAPC	NC042101	C:\varia	DETAILS >

● Current Installations							
Component Name	Version	Product Name	Metric	Computer Name	Installation Path	Details	
IBM BigFix Inventory Server	9.2	IBM BigFix Inventory	RVU MAPC	NC042208	/varia	DETAILS >	
IBM BigFix Inventory Server	9.2	IBM BigFix Inventory	RVU MAPC	NC042193	C:\varia	DETAILS >	
IBM BigFix Inventory Server	9.2	IBM BigFix Inventory	RVU MAPC	NC042101	C:\varia	DETAILS >	

Menu

- The menu consists of a set of options for software assignment, classification, and suppression. The menu also presents a preview of functions that are intended to be implemented in the future. When you hover over unavailable options, they are described as *Not yet implemented*.

Default columns

- **Component Name** shows components that were discovered on your computers.
- **Version** shows version of the component.
- **Product Name** shows products under which the components are licensed based on the software catalog.

- **Metric** shows license metric that is assigned to the product.
- **Computer Name** shows the computer where a component was discovered.
- **Installation Path** shows details about an installation path of IBM products. The installation path of non-IBM products is provided in Details.
- **Details** shows details about a component, such as signatures that were used to discover it.

Additional columns not enabled by default

- **Always Not Charged** indicates whether a component is always not charged.
- **Bundling Tag Used** shows whether the product assignment is based on a bundling tag.
- **Charged** indicates whether the component is charged, or not charged based on its license.
- **Comment** shows a comment that was entered while excluding, or suppressing the software instance from pricing calculations.
- **Computer Deletion Date** shows the date and time when the computer stopped being discovered, for example because the computer was removed from the infrastructure, the data source was removed, or the BigFix client was uninstalled.
- **Confirmed** indicates whether the software classification was manually confirmed.
- **Discovery End** indicates when the component instance was reported for the last time. If the component is still detected, the column shows <n/a>.
- **Discovery Start** indicates when the component instance was reported for the first time.
- **9.2.12 End of Support** indicates the date when the software component reaches end of support. For information about how to change the date, see: [Changing end of support date \(on page cdxcvii\)](#).
- **9.2.13 Vulnerability Risk (Preview)** shows Common Vulnerabilities and Exposures (CVEs) that were matched with the particular software component through its detailed version. The column can contain multiple values that are displayed after expanding. For more information about CVEs, see: [Preview: Checking Common Vulnerabilities and Exposures \(CVEs\) \(on page dclxv\)](#).
- **Excluded** indicates whether the product on a particular computer is currently excluded from pricing calculations.
- **9.2.13 FlexPoint Bundle** shows the FlexPoint bundle to which the product is assigned.
- **IBM-provided Bundling Option** indicates whether the product has the selected component listed as a bundling option in the software catalog provided by IBM.
- **Imported Part Numbers** shows the part numbers that you imported to BigFix Inventory.
- **One to One** indicates whether the product to which the component is currently assigned is the only bundling option that is available in the software catalog.
- **Present** indicates whether the component instance is currently installed.
- **Shared** indicates whether a component instance is shared between multiple products.
- **Suppressed** indicates whether the component is currently suppressed from the inventory.

Customization

Sorting

You can adjust the view by sorting the displayed data by as many columns as you need. To sort the data by multiple columns, hold **SHIFT**, and click the columns. The column you choose first takes precedence in sorting. Other columns are sorted in the order of the selection. In this way, if the data in a single column is the same (for example, BigFix Platform Agent), you can choose another column to further sort it by (for example, by version and product name).

<i>Component Name</i>	<i>Version</i>	<i>Product Name</i>
IBM BigFix Platform Agent	9.2	IBM BigFix Protection
IBM BigFix Platform Agent	9.5	IBM BigFix Protection
IBM BigFix Platform Agent	9.5	IBM BigFix Protection
IBM BigFix Platform Agent	9.5	IBM License Metric Tool

To clear a sorted column, hold **CTRL** or **CMD** on Mac, and click on this column. To clear all, click any of the columns.

Multiselect

To simplify your actions, you can select multiple items at once. To select multiple components:

- Hold **CTRL**, or **CMD** on Mac, to add another item to the current selection.
- Hold **SHIFT** and select two items to extend the selection to all items in between (any previous selection is canceled).
- Select all instances by using **CTRL + A**, or **CMD + A** on Mac.



Restriction: You can select up to 250 000 components at once.

Limitations

Currently, the Software Classification panel has the following limitations:

- When you exclude, or suppress instances, the notification informs you whether the action was successful or not. However, the notification does not provide the number of instances affected, or any further details.
- When you exclude multiple products or products on multiple computers, the exclusion window lists these items separately. The order does not reflect the relation between the computer and the product. The lists are limited to display up to three items.
- When you reassign the component to the different release of the same product, and this component is currently assigned to the excluded release of the product, the exclusion flag is not cleared.
- If you created and saved a custom report filtered by *Exclusion Status* in the versions before 9.2.6, the report is not replicated when you upgrade to the newest version of BigFix Inventory.
- Installation path is provided only for IBM products.
- When you reach the limit of 10 unique custom rules, and create an additional rule without applying a different filter, or refreshing the Software Classification panel view, the additional rule is added.

- When you sort components by metric, components with the same metrics are grouped together, but metrics are not sorted alphabetically.
- Going to the first and last row by using the `HOME` and `END` keys is not supported.

9.2.7 Benefits and limitations of the Software Classification panel

9.2.7 Available from 9.2.7. The Software Classification panel replaced the BigFix Software Classification panel. It displays all your software installations and provides a streamline software management, and classification.

Benefits

Main benefits of the Software Classification panel:

- Provides improved filtering capabilities, multiple filters can be applied to all available columns.
- Allows for flexible ordering and sorting of data by any column.
- Allows for selecting multiple components and assigning them to a product in one go.
- Provides extended search capabilities when you assign components to products. The list of target products can be searched by product name, version, metric, or by other attributes.
- Shows components that are discovered with software ID tags.
- Shows components that were installed in the past but are no longer detected, for example because they were uninstalled or the computer on which they were installed was decommissioned.

Limitations

Currently, the Software Classification panel has the following limitations:

- When you exclude, or suppress instances, the notification informs you whether the action was successful or not. However, the notification does not provide the number of instances affected, or any further details.
- When you exclude multiple products or products on multiple computers, the exclusion window lists these items separately. The order does not reflect the relation between the computer and the product. The lists are limited to display up to three items.
- When you reassign the component to the different release of the same product, and this component is currently assigned to the excluded release of the product, the exclusion flag is not cleared.
- If you created and saved a custom report filtered by *Exclusion Status* in the versions before 9.2.6, the report is not replicated when you upgrade to the newest version of BigFix Inventory.
- Installation path is provided only for IBM products.
- When you reach the limit of 10 unique custom rules, and create an additional rule without applying a different filter, or refreshing the Software Classification panel view, the additional rule is added.
- When you sort components by metric, components with the same metrics are grouped together, but metrics are not sorted alphabetically.
- Going to the first and last row by using the `HOME` and `END` keys is not supported.

Assigning components to products

9.2.7 Available from 9.2.7. After new components are discovered, they are associated with products based on bundling rules. As a result, most components are automatically assigned to products. Correct assignment of software components reduces the cost of software ownership and ensures audit readiness. Compare the automatically created bundles with information in your license agreement, confirm the accurate assignments, or reassign the components to right products.

 You must have the Manage Software Classification permission to view the Software Classification panel and classify software. To view some additional columns related to computers on which the components are installed, you also need the View Endpoints permission.

Assign components to products based on the following scenarios:

- Assign a single component to a single product.
- Assign multiple components to a single product. To properly use this option, set filters that limit the displayed components, select them, and assign to the product.
- Assign a single component more than once to share this component between multiple products. Each assignment associates the component with one selected product. To view all shared components, apply the following filter: `Shared, equal to, Yes`.

You can automate your software classification process by creating a set of custom rules that will automatically classify your newly discovered software. For more information, see: [Creating and managing custom rules \(on page dl\)](#).

1. Log in to BigFix Inventory, and go to **Reports > Software Classification**.
2. Select a component or components that you want to assign.
3. Hover over the **Assign** icon , and click **Choose from Catalog**.

 **Tip:** To narrow the list of displayed products, you can sort the data by columns or use the search bar that is available at the top of each column. After you type the search phrase, press `Enter`.

4. Select the product to which you want to assign the component and the license metric that is suitable for your pricing calculations model.
 - If you want to assign the component to a product according to a relation that is defined in the software catalog, simply select the product from the list.
 - If the product is not detected and the value of the Product Name column is `<not set>`, see: [Catalog Problems: The `<not set>` value appears under the Product Name on Software Classification panel \(on page dcclviii\)](#).
 - If none of the products that is listed by default matches your license agreement, create a custom relation known as *custom bundling*. You create a custom bundling in the following situations.

- The component was obtained as part of a product that is not listed as an existing relation.
- **9.2.12** The type of the relation that is listed in your license agreement is not available by default. For example, when the existing relation between the component and the product is charged, but your license agreement states that the component is not charged.

To create a custom bundling, perform the following steps.

- a. Clear the **List products for which the component is listed as a bundling option** check box to view all products that exist in the catalog, and select the appropriate product.
- b. Specify whether you want the relation between the component and the product to be charged or not charged. If you specify the relation as not charged, the component does not contribute to the license metric utilization of the product to which it is assigned.

Creating a custom bundling impacts [automated bundling](#). It increases the chance for newly discovered components to get automatically bundled according to custom bundling definition.

5. **Optional:** To extend the current component assignment with the selected product instead of reassigning the component, select **Share component**. The component will be shared between the product to which it is currently assigned and the selected product.

If you share the same component between multiple products, each assignment is listed as a separate instance on the report.



Note: To stop sharing the component between multiple products, select all instances of the component and assign them to a common product, but leave the **Share component** check box cleared.

6. **Optional:** If you do not want the assignment to be automatically confirmed, clear the **Automatically confirm the affected components** check box.



Note: To change the default setting of the check box, go to **Management > Advanced Server Settings**, and change the value of the **automaticConfirm** parameter according to your preferences.

7. To change the assignment, click **Assign**.

The component is reassigned to the specified product and metric. The change is reflected on the report.

To ensure that metric utilization reflects the change of the component assignment, go to **Reports > All Metrics**, and click **Recalculate**.

9.2.7 Excluding and suppressing software instances

9.2.7 Available from 9.2.7. If a software instance should not contribute to license metric utilization, you can exclude it on the product level or suppress it on the component level. Exclusion applies to all components that are assigned to a product and are installed on the selected computer. Use it if license information specifies that the

product should not contribute to metric utilization under specific circumstances. Suppression applies to a single component that is installed on the selected computer. Use it when a component is improperly discovered.

 You must have the Manage Software Classification permission to view the Software Classification panel and classify software. To view some additional columns related to computers on which the components are installed, you also need the View Endpoints permission.

Exclusion

Exclude a product when it is installed on a non-production, test, or back up computer and license information specifies that the product should not contribute to license metric utilization under these circumstances. When you exclude the product, all components that are assigned to this product and are installed on the selected computer are excluded from metric utilization.

Excluded products and their components are still displayed on the Software Classification report. They are also listed in the audit snapshot, together with exclusion comments.

Suppression

Suppress a component when it is improperly discovered. For example, when some files were not removed from the computer after the component was uninstalled. Suppression applies only to the selected component on the selected computer. A suppressed component is no longer assigned to any product, and thus does not contribute to metric utilization.

By default, suppressed components are not displayed on the Software Classification report. You can show them by removing one of the default report filters. Suppressed components are not included in the audit snapshot in order not to obscure the view. However, the components and suppression comments are included in the Audit Trail.

1. Log in to BigFix Inventory and go to **Reports > Software Classification**.
2. **Optional:** To filter down the report, hover over the **Manage Report View** icon , and click **Configure View**. Then, specify appropriate filters.
3. Select one or more rows and hover over the **Exclude or Suppress** icon .
 - To exclude the product on the selected computer, click **Exclude Product from Metric Calculation**.
 - To suppress the selected component on the selected computer, click **Suppress Component from Inventory**.
4. Provide an obligatory comment.

If a comment was provided for any selected item before, it is overwritten. Exclusion comment applies to all components that are assigned to the excluded product on the selected computer. Suppression comment applies to the suppressed component on the selected computer.
5. To automatically confirm the action, select **Automatically confirm the affected components**.



Note: To change the default setting of the check box, go to **Management > Advanced Server Settings**, and change the value of the **automaticConfirm** parameter according to your preferences.

6. To complete the action, click **Exclude** or **Suppress**.

The action completes even if any of the selected items has already been excluded or suppressed.

Excluded products and the suppressed components do not contribute to license metric utilization. You can include them back to the calculations by using the **Clear Exclusion**, or **Clear Suppression** option. The instances that you include back in the calculation are marked as unconfirmed.

To review exclusion and suppression actions for BigFix components and products, go to **Reports > Audit Trail**.

9.2.7 Confirming assignment of software components

9.2.7 Available from 9.2.7. When the assignment of a component to the product is correct, you can confirm it. The task is optional. However, components whose assignment is not confirmed can be automatically reassigned to different products. It might happen when a new software catalog is uploaded or when a different product is suggested as the best bundling option based on the new results of automated bundling. To avoid such a situation, it is recommended to confirm the assignment of all components.

 You must have the Manage Software Classification permission to view the Software Classification panel and classify software. To view some additional columns related to computers on which the components are installed, you also need the View Endpoints permission.

1. Log in to BigFix Inventory, and go to **Reports > Software Classification**.
2. Confirm all components that are marked as One to One.

If a component relates to a single product in the software catalog, the relation is marked as One to One. Previously, all BigFix components that had a single bundling option in the software catalog were automatically confirmed. Now, you can easily filter the report to review these components, and confirm them. It ensures that you are fully in charge of the software classification process.

- a. Hover over the **Manage Report View** icon , and click **Configure View**. Then, add the following filters to the filters that are already defined in this report:

- Publisher Name, equal to, IBM
- IBM-provided Bundling Option, equal to, Yes
- One to One, equal to, Yes
- Confirmed, equal to, No

Filters

Specify the report filter which matches all of the following conditions:

Publisher Name	equal to	IBM		
IBM-provided Bundling Option	equal to	Yes		
One to One	equal to	Yes		
Confirmed	equal to	No		

Submit
Cancel

b. Press **CTRL + A**, or **CMD + A** on Mac, to select all the filtered components.

c. Hover over the **Confirm Components** icon , and click **Confirm**.

3. Confirm all assignments that are correct.

a. Select a component or components that you want to confirm.

b. Hover over the **Confirm Components** icon , and click **Confirm**.

9.2.7 Creating and managing custom rules

9.2.7 Available from 9.2.7. You can create custom rules to automate the process of software classification.

The rules automatically perform the action that is defined by the rule criteria. For example, when you create a custom rule while assigning a component to a product, the rule is applied whenever the exact same component is discovered in the future.

 You must have the Manage Imports and Manage Software Classification permissions to perform this task.

Custom rules can be applied to:

- A single computer when they are created for this computer
- A computer group when they are created for this computer group
- Computer group to which the user is assigned. The user cannot choose a computer group, and is limited to the group or groups to which the user belongs. The rule is always applied to the current user group. When user's computer group changes, the rule scope is revised.

 **Restriction:** The number of unique custom rules you can create is currently limited to 10.

9.2.16 Starting from application update 9.2.16, you can create a custom rule for suppressing components.

1. Log in to BigFix Inventory, and go to **Reports > Software Classification**.
2. Filter the panel so that it shows only instances for which you want to create a custom rule.

 **Tip:** Specify the exact version of the component to avoid inconsistency.

3. A custom rule can be created only when you select all instances that are currently displayed. To select all filtered instances, press `CTRL + A`, or `COMMAND + A` on Mac.
4. [Assign a component \(on page dxlvi\)](#) to a product, or [exclude a product or suppress a component \(on page dxlvii\)](#) from software calculations.
5. Before you confirm the assignment, exclusion, or suppression select the **Create a custom rule** check box.

 **Restriction:** You cannot create a custom rule and share the component at the same time. Perform these actions separately.

6. Click **Assign**, **Exclude**, or **Suppress**.

The custom rule was created. To view your custom rules, navigate to **Management > Custom Rules**. The rules are added to the panel in the order of creation, and are applied from the top during import.

Deleting custom rules

All custom rules that you create are displayed, and can be deleted on the Custom Rules panel. Any user who has the Manage Imports and Manage Software Classification permissions can add, and delete any rules, including the rules created by other users. When you delete a user, the rules created by that user are not automatically removed.

 You must have the Manage Imports and Manage Software Classification permissions to perform this task.

1. Go to **Management > Custom Rules**.
2. Select a custom rule that you want to delete, and click **Delete**.

The rule is deleted and no longer applied. The software classification based on the rule remains unchanged.

9.2.13 Assigning products to FlexPoint bundles

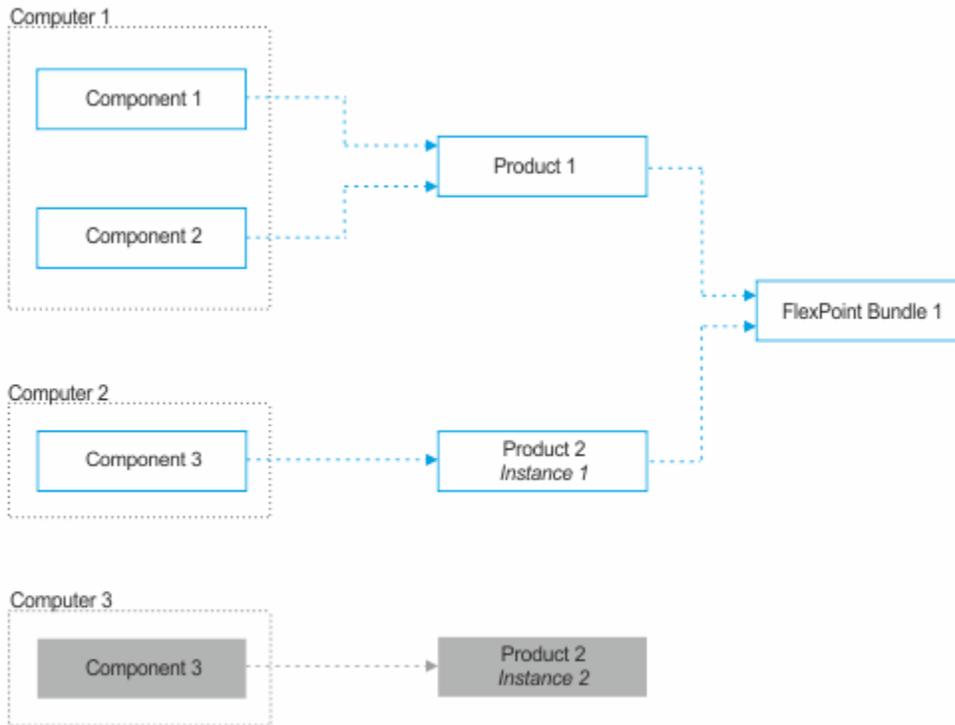
9.2.13 Available from 9.2.13. When you purchase a product as part of a FlexPoint bundle, assign instances of this product to that bundle for improved tracking and to calculate utilization of FlexPoints.

-  You must have the Manage Software Classification permission to view the Software Classification panel and classify software. To view some additional columns related to computers on which the components are installed, you also need the View Endpoints permission.
- Ensure that components that were discovered in your environment are assigned to software products in accordance with your license agreement. For more information, see: [Assigning components to products \(on page dxlvi\)](#).

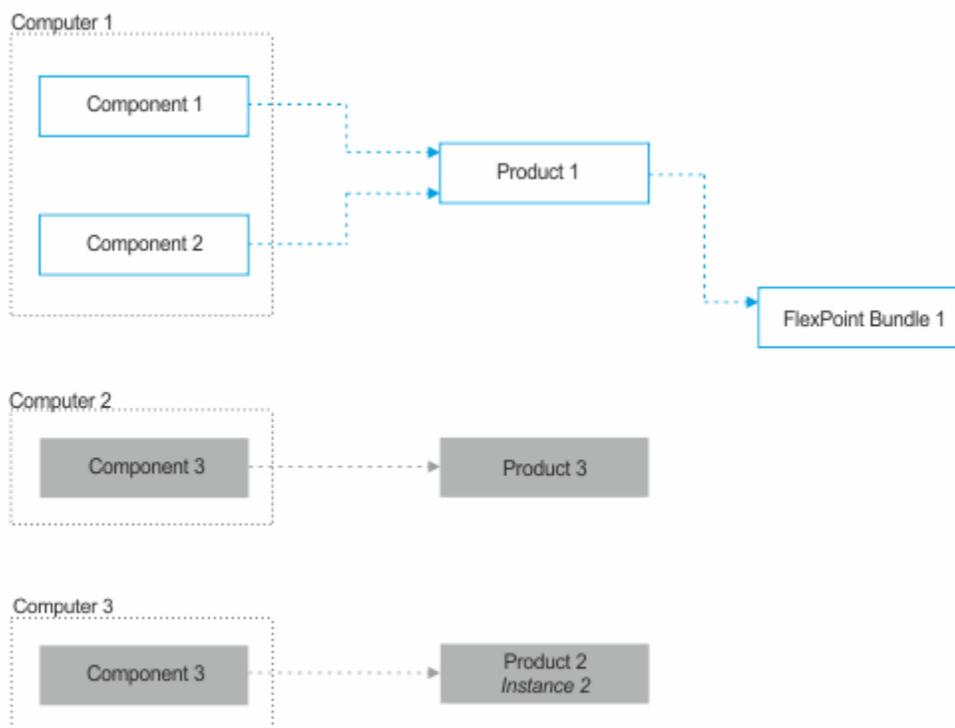
To properly classify the installed software, start by assigning the discovered components to products in accordance with your license agreement. In this example, the discovered components are assigned as follows.

- Component 1 and Component 2 installed on Computer 1 are assigned to Product 1.
- Component 3 installed on Computer 2 is assigned to the first instance of Product 2.
- Component 3 installed on Computer 3 is assigned to the second instance of Product 2.

Then, assign proper product instances to FlexPoint bundles. In this example, Product 1 and the first instance of Product 2 are a part of FlexPoint bundle 1. The second instance of Product 2 is licensed under a different license metric and thus is not assigned to any FlexPoint bundle.



If a component is assigned to a product that is a part of a FlexPoint bundle, and you reassign such a component, the component instance is no longer assigned to any FlexPoint bundle. In this example, Component 3 that is installed on Computer 2 is reassigned from the first instance of Product 2 to Product 3. As a result, the reassigned component is no longer a part of any FlexPoint bundle. The same applies to a situation in which you suppress a component that is assigned to a product that is a part of a FlexPoint bundle.



1. Log in to BigFix Inventory, and go to **Reports > Software Classification**.
2. Select the product instance that you want to assign to a FlexPoint bundle.
When you select the product, make sure that you choose the product instance that is assigned the license metric which is defined in your license agreement.
3. Hover over the **Assign** icon , and click **FlexPoint Bundle**.
4. Select the bundle to which you want to assign the product, and click **Assign**.
Only the product instance that is installed on the particular computer and its components are assigned to the FlexPoint bundle.

FlexPoint Bundle
✕

Selected products: IBM Cognos Controller Administrator

FlexPoint Bundle: Data Science & Business Analytics ▼

Assign
Cancel

 **Note:** If you previously assigned a product to a FlexPoint bundle and later on want to remove the product from that bundle, choose **None**.

Information about the FlexPoint bundle to which the product is assigned is displayed in the FlexPoint Bundle column.

To view utilization of FlexPoints, go to **Reports > IBM FlexPoints**, and click **Recalculate** to refresh the data on the report.

9.2.13 Setting the number of metric units convertible to FlexPoints

9.2.13 Available from 9.2.13. When a software product is assigned to a FlexPoint bundle, BigFix Inventory converts the number of metric units that the product used to FlexPoints. If the product uses a license metric whose utilization is not calculated by BigFix Inventory, you can manually set the number of metric units to be converted into FlexPoints.

 You must have the View License Metrics permission to view the IBM FlexPoints report, and Manage Software Classification permission to set the number of metric units to be converted into FlexPoints.

You can manually set the number of metric units convertible to FlexPoints only for metrics whose utilization is not calculated by BigFix Inventory. For the list of metrics whose utilization is calculated in BigFix Inventory, see: [Table 51: Metrics for which utilization is calculated \(on page dlvi\)](#).

1. Log in to BigFix Inventory, and go to **Reports > IBM FlexPoints**.
2. To display the list of products that are assigned to a FlexPoint bundle, click the name of that bundle.



Product Name	Metric	Metric Quantity	Metric Quantity History
Data Science & Business Analytics	FlexPoints	0 <n/a>	06/06/2018 - 09/04/2018

3. Select the product for which you want to specify the number of license metrics, hover over the **Metric Quantity** icon , and click **Set Value**.



Product Name	Metric	Metric Quantity	Metric Units Convertible t...	FlexPoints per Metric Unit	FlexPoints
IBM SPSS Modeler Premium	Authorized Us	<n/a>	<n/a>	30140	<n/a>

4. Specify the number of metric units that the product uses, and click **OK**.



Edit Metric Quantity ✕

Metric Quantity:

To refresh the data that is displayed on the report, click **Recalculate**. Metric units that you specified are converted into FlexPoints.

Report Date: 2018-09-12

FlexPoints: 0

Product needs to be recalculated. Use the Recalculate button to initiate the recalculation process.

Product Name	Metric	Metric Quantity	Metric Units Convertible t...	FlexPoints per Metric Unit	FlexPoints
IBM SPSS Modeler Premium	Authorized User	<n/a>	100	30,140	3,014,000

9.2.7 Viewing history of installed components

9.2.7 Available from 9.2.7. You can view information about components that were installed in your infrastructure during a specific period and their assignment. You can use this information to view details of components that contributed to the overall license utilization for a specific period. You can also reassign and exclude uninstalled components if their original assignment was incorrect and impacted the license metric utilization.

 You must have the View Endpoints and Manage Software Classification permissions to perform this task.

1. In the top navigation bar, click **Reports > Software Classification**.
2. To show components that were discovered during a particular period, hover over the **Manage Report View** icon , and click **Configure View**. Then, perform the following steps.

- a. Remove the current filter.
- b. Specify discovery start on or before the last day of the period for which you want to display the data, and discovery end on or after the first day of that period.
For example, to display all components that were discovered between 1 September and 30 September, use the following filters:
 - **Discovery Start, on or before**, and choose 30 September.
 - **Discovery End, on or after**, and choose 1 September.
- c. Click **Submit**.

The report shows components that were discovered during the specified period, including uninstalled components, and components that were installed on computers that are no longer discovered. Details of the computer on which the component was discovered, computer health, scan results, hardware inventory, and application usage represent the current state of the infrastructure, not historical data from the specified period. If a computer is no longer discovered, the data is not available.

Publisher ...	Component Name	Version	Product Name	Computer Na...	Installation Path	Details
IBM	IBM DB2 Workgroup Server Edition OEM Limite...	10.5	IBM BigFix Inventory	 NC91431271	/opt/ibm/db2/10.5/properties/version	<no data>
IBM	IBM DB2 Workgroup Server Edition OEM Limite...	10.5	IBM BigFix Inventory	NC9143127057	/opt/ibm/db2/10.5/properties/version	<no data>
IBM	IBM SDK 5.0 for Linux IA32, Java Technology	5.0	IBM SDK 5.0 for Linux IA32, Java Technology	NC9143127057	/opt/IBMts&am/sam/java/docs	<no data>
IBM	IBM SDK 5.0 for Linux IA32, Java Technology	5.0	IBM SDK 5.0 for Linux IA32, Java Technology	 NC91431270	/opt/IBMts&am/sam/java/docs	<no data>

If a computer is no longer discovered, it has the  icon displayed next to its name. It happens in the following situations:

- The computer was removed from the infrastructure. For more information, see [Removing inactive computers from BigFix \(on page cdxxxvii\)](#).
- The BigFix client was uninstalled from that computer
- The data source to which the computer belonged was removed

You can reassign components or exclude them from pricing calculations to obtain the correct reports for the specified period.

License metrics

BigFix Inventory provides two types of information about license metrics utilization: aggregated utilization and raw utilization. The former shows utilization of the assigned license metric by all instances of a particular product within the reporting period. However, it is limited to a subset of license metrics. The latter shows data that is collected from `.slmtag` files. It shows raw utilization of all license metrics that can be used by a particular product. The data is not aggregated and requires further processing.

Reported license metrics

BigFix Inventory reports utilization for a number of license metrics for IBM and non-IBM products. It also shows products that are assigned any BigFix metric even if calculating utilization is not yet supported for this metric. It gives you an overview of all IBM metrics for which you should have licenses.

Metrics for which utilization is calculated

License metric utilization is calculated for products that are assigned to one of the metrics that are listed in the following table. Their utilization is shown on the All Metrics report. For more information, see: [License metric utilization \(on page dlxxxvii\)](#).

For every of these metrics, BigFix Inventory calculates the so-called high-water mark. For more information, see: [High-water mark \(on page lxxvi\)](#).

Table 51. Metrics for which utilization is calculated

The table consists of two columns, a header row, and 12 rows.

Vendor	Product	Metric	Description
BigFix	<ul style="list-style-type: none"> • BigFix products that are sold as part of FlexPoint bundles 	9.2.13 FlexPoint	PVU and RVU software that is assigned to a FlexPoint bundle, must be licensed based on subcapacity rules. For more information, see: IBM FlexPoints (on page dlx) .

Table 51. Metrics for which utilization is calculated

The table consists of two columns, a header row, and 12 rows.

(continued)

Vendor	Product	Metric	Description
	<ul style="list-style-type: none"> BigFix products that are licensed based on the listed metrics 	9.2.7 Install Instances	A unit of measure that is based on the number of installed instances regardless of whether the instances are installed on one or many computers. For example, when three instances of a product are installed on one computer, the metric quantity is 3.
		Processor Value Unit (PVU)	For more information, see: IBM Processor value unit (PVU) (on page dlxxvi) .
		Resource Value Unit (RVU MAPC)	For more information, see: IBM Resource value unit (RVU MAPC) (on page dlxxxiv) .
		9.2.11 Virtual Processor Core (VPC)	For more information, see: IBM Virtual processor core (VPC) (on page dlxxxvi) .
		Other IBM metrics	For more information, see: Other IBM metrics (on page dlxxiv) .
SAP	<ul style="list-style-type: none"> SAP-ABAP based products 	Engines, Users	For more information, see: Discovering and measuring SAP (on page dxiii) and Raw utilization of license metrics (on page dcxi) .  Note: Information about SAP metrics is shown only on the Resource Utilization report.
Oracle	<ul style="list-style-type: none"> All discoverable versions of Oracle Database (Enterprise Edition) 	9.2.8 Oracle Processor Core	For more information, see: Oracle Processor Core (on page dlxx) .
		Oracle Concurrent Sessions, Oracle User Sessions	For more information, see: Extended discovery of Oracle Database (on page dii) and Raw utilization of license metrics (on page dcxi) .

Table 51. Metrics for which utilization is calculated

The table consists of two columns, a header row, and 12 rows.

(continued)

Vendor	Product	Metric	Description
			 Note: Information about Oracle Database concurrent user sessions metrics is shown only on the Resource Utilization report.
	<ul style="list-style-type: none"> • Programs discovered by BigFix Inventory and based on the Instal Seats metric 	9.2.5 Install Seats	A unit of measure that is based on the number of computers on which the software is installed. For example, when three instances of a product are installed on one computer, the metric quantity is 1.
Microsoft	<ul style="list-style-type: none"> • Windows Server 2008 R2 Datacenter 	9.2.8 Microsoft Single Physical Processor	For more information, see: Microsoft Single and Dual Physical Processor (on page dlxii) .
	<ul style="list-style-type: none"> • Windows Server 2012 Datacenter • Windows Server 2012 R2 Datacenter 	9.2.8 Microsoft Dual Physical Processor	
	<ul style="list-style-type: none"> • Microsoft SQL Server 2012, 2014, and 2016 (Standard and Enterprise) 	9.2.7 Microsoft Physical Core with SA	For more information, see: Microsoft Physical Core with SA (on page dlxiv) .
	<ul style="list-style-type: none"> • Microsoft SQL Server 2012, 2014, and 2016 (Standard and Enterprise). 	9.2.7 Microsoft Virtual Core with SA	For more information, see: Microsoft Virtual Core with SA (on page dlxvii) .
	<ul style="list-style-type: none"> • Microsoft Office 365 	9.2.13 Registered User	For more information, see: Registered User (on page dlxxxi) .

Table 51. Metrics for which utilization is calculated

The table consists of two columns, a header row, and 12 rows.

(continued)

Vendor	Product	Metric	Description
	<ul style="list-style-type: none"> • Programs discovered by BigFix Inventory and based on the Instal Seats metric 	9.2.5 Install Seats	A unit of measure that is based on the number of computers on which the software is installed. For example, when three instances of a product are installed on one computer, the metric quantity is 1.
VMware	<ul style="list-style-type: none"> • vSphere • vCenter 	9.2.7 VMware sockets	For more information, see: Measuring license metric utilization of VMware products (on page dx) and Raw utilization of license metrics (on page dcxi) .  Note: Information about VMware metrics is shown only on the Resource Utilization report.
	<ul style="list-style-type: none"> • Programs discovered by BigFix Inventory and based on the Instal Seats metric 	9.2.5 Install Seats	A unit of measure that is based on the number of computers on which the software is installed. For example, when three instances of a product are installed on one computer, the metric quantity is 1.
Adobe	<ul style="list-style-type: none"> • Adobe Illustrator CC • Adobe PhotoShop CC • Adobe InDesign CC • Adobe Premiere Pro CC • Adobe After Effects CC • Adobe Dreamweaver CC • Adobe Muse CC 	9.2.13 Registered User	For more information, see: Registered User (on page dlxxxi) .

Table 51. Metrics for which utilization is calculated

The table consists of two columns, a header row, and 12 rows.

(continued)

Vendor	Product	Metric	Description
	<ul style="list-style-type: none"> Programs discovered by BigFix Inventory and based on the Instal Seats metric 	9.2.5 Install Seats	A unit of measure that is based on the number of computers on which the software is installed. For example, when three instances of a product are installed on one computer, the metric quantity is 1.
Other	<ul style="list-style-type: none"> Non-BigFix products that use this metric. Non-BigFix products that are discovered by SWID tags and for which the software catalog does not define license details. 	9.2.5 Install Seats	A unit of measure that is based on the number of computers on which the software is installed. For example, when three instances of a product are installed on one computer, the metric quantity is 1.

9.2.13 IBM FlexPoints

9.2.13 Available from 9.2.13. FlexPoint is a license metric unit that can be used to determine the cost of IBM products that are purchased as part of FlexPoint bundles. Each product from a bundle is licensed based on a different license metric but all metrics are converted into IBM FlexPoints. The conversion is based on the specific calculations. PVU and RVU software that is assigned to a FlexPoint bundle, must be licensed based on subcapacity rules. Products that are currently available as part of FlexPoint bundles are products from the IBM Cloud and IBM Analytics.

Predicting the number of required IBM FlexPoints

To predict the number of IBM FlexPoints that are required to license products in your environment, you can use the [FlexPoint Calculator](#). For each product that you are planning to deploy as part of a FlexPoint bundle, specify the number of license metric units. The numbers are then converted into IBM FlexPoints and summed up.

Example

The following example describes how the number of IBM FlexPoints that are used in your environment is counted by BigFix Inventory. The calculations are based on sample values and should be treated only as an example.

You purchased the IBM Cloud DevOps offering and you deployed three products from this offering: UrbanCode Deploy Server Agent, Rational Test Workbench, and Rational Test Virtualization Server. Each of the products is licensed based on a different license metric as presented in [Table 52: IBM FlexPoints needed to license products deployed as](#)

part of the IBM Cloud DevOps offering (on page dlxi). To calculate the number of IBM FlexPoints that are used by these products, BigFix Inventory multiplies the number of license metric units that each product uses by the number of IBM FlexPoints that are assigned per metric unit for that product. Results for each product are then summed up to obtain the number of IBM FlexPoints used by the entire FlexPoint bundle.

Table 52. IBM FlexPoints needed to license products deployed as part of the IBM Cloud DevOps offering

The table consists of four columns and five rows. In the fifth row, the first three columns are joined.

Product	Used license metrics	IBM FlexPoints per license metric unit	Calculations
UrbanCode Deploy Server Agent	2 Virtual Servers	129,800	2 x 129,800 = 259,600
Rational Test Workbench	50 Authorized Users	15,950	50 x 15,950 = 797,500
Rational Test Virtualization Server	210 PVUs	2,893	210 x 2,893 = 607,530
Total FlexPoints Used			1,664,630

Reading metric utilization

Information about FlexPoint bundles

You can view information about FlexPoint bundles on the IBM FlexPoints report. The report lists the FlexPoint bundles based on the decisions that you make during software classification. The metric quantity represents the license metric utilization for an entire bundle over the specified period of time expressed in FlexPoints.

Product Name	Metric	FlexPoint Bundle	Metric Quantity	Metric Quantity History	
				06/01/2018	08/30/2018
Data Science & Business Analytics	FlexPoints	Data Science & Business Analytics	1019040	<n/a>	
IBM Cloud Integration - Monthly	FlexPoints	IBM Cloud Integration - Monthly	1595	<n/a>	
IBM Cloud Integration - Perpetual	FlexPoints	IBM Cloud Integration - Perpetual	9108	<n/a>	
Cloud DevOps - Monthly	FlexPoints	Cloud DevOps - Monthly	13440	<n/a>	

Information about the products that contribute to FlexPoint bundles

To view the detailed report with license metric utilization for individual products that are parts of a bundle, go to the IBM FlexPoints report, and click the FlexPoint bundle name in the Product Name column. The detailed information include the peak date of the entire FlexPoint bundle and the number of FlexPoints used by each individual product. Additionally, you can check the number of original metric units by which the products are licensed that are used over the reporting period.



Note: **9.2.16** You can change the metric quantity for the products that contribute to the bundle. Once you set the declared metric quantity for a product that is a part of a FlexPoint bundle, the number of the FlexPoints for the entire bundle is automatically adjusted. For more information, see: [Setting and removing the declared metric quantity \(on page dlxxxiii\)](#).

Cloud DevOps - Perpetual

Send Feedback

Report Date: 2018-08-30

FlexPoints: 325440

Product Name	Metric	Metric Quantity	Metric Units Convertible to FlexPoints	FlexPoints per Metric Unit	FlexPoints
IBM Rational Performance T...	PVU Subcapacity	960	960	339	325440

Information about FlexPoints in audit snapshot

9.2.16 When you generate the audit snapshot, information about FlexPoint bundles and the number of FlexPoints that they use is contained in the `flexpoints.csv` file. The file additionally contains information about the software products within the bundle and the number of FlexPoints used by them.

Table 53. A fragment of the `flexpoints.csv` file

Row No.	Publisher	Product Name	FlexPoint Bundle	Metric	Metric Quantity	Peak Date	FlexPoint Quantity	Recalculation Needed
1	IBM	Cloud Dev- Ops - Per- petual	FlexPoint	None	139,073	06/14/2019		No
2	IBM	BigFix Rational ClearCase	Floating User	Cloud Dev- Ops - Per- petual	20	06/14/2019	136,400	No
3	IBM	BigFix Rational ClearQuest	Authorized User	Cloud Dev- Ops - Per- petual	1	06/14/2019	2,673	No

Related information

[High-water mark \(on page lxxvi\)](#)

9.2.8 Microsoft Single and Dual Physical Processor

9.2.8 Available from 9.2.8. Microsoft Single Physical Processor and Dual Physical Processor metrics are used to determine the licensing cost of Microsoft products. They are based on the number of active physical processors

that are deployed on the host where the software is installed. When the software is installed on multiple hosts, all active processors from all hosts contribute to metric utilization.

Microsoft Single Physical Processor metric covers one physical processor. Microsoft Dual Physical Processor license covers up to two physical processors. BigFix Inventory reports the number of processors, not licenses. Thus, the number of processors is rounded up to an even number.

By default, calculation of both metrics takes into account data from the last 90 days because Microsoft licensing rules do not allow for license migration more often than once per 90 days. To change the period for which data is calculated, go to **Management > Advanced Server Settings** and change the value of the **historical_period_for_microsoft_metrics** parameter according to your needs.

Supported software

- Reporting of the Microsoft Single Physical Processor metric is supported for Windows Server 2008 R2 Datacenter.
- Reporting of the Microsoft Dual Physical Processor metric is supported for Windows Server 2012 Datacenter or Windows Server 2012 R2 Datacenter.

Requirements

VM managers should be defined for all virtual machines in your environment. If a VM manager is not defined for a particular host, the number of processors on the host is counted as 0. As a result, license utilization might be underestimated. If a component of a product with any of the two metrics is detected on a host for which a VM manager is not defined, a warning sign (⚠️) is displayed on the All Metrics report. It indicates calculation inaccuracy.



Note: The warning is displayed only on the user interface. It is not displayed in the CSV or PDF version of the All Metrics report.

Limitations

When a computer on which the software was discovered is deleted, report columns related to server capacity show `<no data>` in the user interface. Information about computers that contributed to the metric peak value is available in the audit snapshot regardless of whether these computers were removed.

Examples

Example 1: Software is deployed on two physical computers

Microsoft Windows Server 2012 Datacenter is installed on 2 physical computers, each with 4 processors. Metric utilization equals 8.

Example 2: Software is deployed on a virtual machine that was migrated between hosts

Microsoft Windows Server 2012 Datacenter is installed on a virtual machine that runs on a host with 2 processors. Within a period of 90 days from the software discovery, the virtual machine is migrated to

another host that also has 2 processors. The number of processors from the two hosts is added. Metric utilization equals 4.

Table 54. Utilization of the Dual Physical Processor metric for a virtual machine that was migrated

	Host 1	Host 2
Physical processors	2	2
Metric utilization	4	

Related information

[High-water mark \(on page lxxvi\)](#)

9.2.6 Microsoft Physical Core with SA

9.2.6 Available from 9.2.6. The Microsoft Physical Core with SA metric is used to determine the licensing cost of Microsoft products. It is based on the number of physical cores that are deployed on the host where the software is installed. When the software is installed on multiple hosts, all cores from all hosts contribute to metric utilization. Because this license type is based on physical cores, the same number of cores is assigned regardless of whether the software is installed on a physical host or on a virtual machine.

Depending on the processor type, Microsoft requires that an appropriate core factor is applied to calculate the number of required licenses. The number of cores that is reported by BigFix Inventory represents the actual number of physical cores without any core factor applied. The minimal number of licensable cores is four per physical processor. It means that each processor on a host with one or two cores is assigned four cores.

Supported software

Reporting of the metric is supported for Microsoft SQL Server 2012, 2014, and 2016 (Standard and Enterprise).

Requirements

To report the number of physical cores for the discovered Microsoft software, the licensed software must have Software Assurance.

Additionally, VM managers should be defined for all virtual machines in your environment. If a VM manager is not defined for a particular host, the number of cores on the host is counted as 0. As a result, license utilization might be underestimated. If a component of a product with the Microsoft Physical Core with SA metric is detected on a host for which a VM manager is not defined, a warning sign (⚠️) is displayed on the All Metrics report. It indicates calculation inaccuracy.



Note: The warning is displayed only on the user interface. It is not displayed in the CSV or PDF version of the All Metrics report.

Limitations

When a computer on which the software was discovered is deleted, report columns related to server capacity show `<no data>` in the user interface. Information about computers that contributed to the metric peak value is available in the audit snapshot regardless of whether these computers were removed.

Examples

The following examples describe how the number of required licenses is calculated. They are based on sample data.

Example 1: Software is deployed on a physical server

Microsoft SQL Server 2014 Enterprise is installed on a physical computer with four dual-core processors. According to the licensing terms, each dual-core processor is assigned four cores which gives 16 cores in total. Metric utilization for Microsoft SQL Server is 16.

Table 55. Metric utilization for software in a physical environment

	Host 1
Physical processors on the host	4
Physical cores on the host	8
Core licenses needed	16

Example 2: Software is deployed on two virtual machines

Microsoft SQL Server 2014 Enterprise is installed on two virtual machines that run on a host with 10-core processor. The cores are shared by the virtual machines. The number of virtual cores that are assigned to a virtual machine is ignored and BigFix Inventory reports the total number of physical cores on the host. Metric utilization for Microsoft SQL Server is 10.

Table 56. Metric utilization for software in a virtual environment

	Virtual Machine 1	Virtual Machine 2
Virtual cores	8	16
Physical cores on the host	10	10
Physical processors on the host	1	
Core licenses needed	10	

Example 3: Software is deployed on two virtual machines that were migrated between hosts

Microsoft SQL Server 2014 Enterprise is installed on two virtual machines that run on a host with 10-core processor. Both virtual machines are migrated on the same day to a host with a 12-core processor. Metric utilization is checked once per day at midnight (GMT). Taking advantage of License Mobility, all licenses from Host 1 are reassigned to Host 2. Metric utilization for Microsoft SQL Server is 10 for the period before the migration, and 12 starting from the day of the migration.

Example 4: Software instance uses a different license metric

Two instances of Microsoft SQL Server 2014 Enterprise are discovered. The first instance uses the Microsoft Physical Core with SA metric. The second instance uses a different metric, for example Microsoft Virtual Core with SA. To ensure that only the first instance contributes to the calculation of the Microsoft Physical Core with SA metric, reassign the second instance to the proper metric. If the metric is not available in the software catalog, exclude the instance from pricing calculations. For more information, see: [Assigning components to products \(on page dxlvi\)](#), and [Excluding and suppressing software instances \(on page dxlvii\)](#).

Example 5: Two software instances use the same metric and are installed on the same server

Microsoft SQL Server 2012 Enterprise and Microsoft SQL Server 2014 Enterprise are discovered on one server. Both instances use the Microsoft Physical Core with SA metric. Metric utilization is counted once for this server regardless of the number of installed instances.

Reading metric utilization

The first level of the All Metrics report shows a product and its metric quantity, which is the highest number of metric units that the product used during the specified period. The warning sign (⚠️) indicates that the VM manager is not defined for at least one of the hosts on which the product is installed.

Publisher	Product Name	Metric	Metric Quantity	Metric Quantity History	
				11/16/2016	02/14/2017
Microsoft	Microsoft SQL Server Enterprise E...	Microsoft Physical Core with SA	44	⚠️	

When you click the product name, the second level of the report shows all instances of product components that contributed to the metric quantity, including historical instances that are no longer installed. The value in the Server Cores column shows the number of physical cores on the host. For hosts without VM managers defined, the column shows the number of virtual cores. Thus, it is skipped in the calculation of the overall metric quantity.

Component Name	Version	Product Name	Metric	Excluded	Comput...	Status	Server ID	Server C...	Server A...	Details
Microsoft Windows Server 2...	6.1	Microsoft SQL Server Enterpr...	Microsoft P...	No	NC040202	OK	IBM NC14...	12		2 DETAILS >
Java Support	5.2	Microsoft SQL Server Enterpr...	Microsoft P...	No	NC0472...	OK	IBM 9117 ...	32		4 DETAILS >
Microsoft SQL Server 2012 f...	2012.0	Microsoft SQL Server Enterpr...	Microsoft P...	No	NC9143...	No VM ...	TLM_VM_...	1		1 DETAILS >

The fact that the value is skipped is reflected in the Core Metric column in the `microsoft_physical_core_with_sa.csv` file which is a part of the audit snapshot. Additionally, information about incomplete virtualization technology is provided in the Comment column.

Table 57. A fragment of the `microsoft_physical_core_with_sa.csv` file

Publisher	Product Name	Metric Quantity	Server Cores	Core Metric	Comment
Microsoft	Microsoft SQL Server Enterprise Edition	44	12	12	
Microsoft	Microsoft SQL Server Enterprise Edition	44	32	32	
Microsoft	Microsoft SQL Server Enterprise Edition	44	0	0	Incomplete virtualization hierarchy

Related information

[High-water mark \(on page lxxvi\)](#)

9.2.7 Microsoft Virtual Core with SA

9.2.7 Available from 9.2.7. The Microsoft Virtual Core with SA metric is used to determine the licensing cost of Microsoft products that have Software Assurance. It is based on the number of virtual cores that are available to the product. When the product is installed on multiple virtual machines, virtual cores from all machines contribute to metric utilization.

The number of virtual cores can be greater than the number of physical cores that are available on the hosts. The minimal number of licensable cores is four. It means that when a virtual machine is assigned one or two cores, four cores are reported. If the number of cores is not a multiple of 2, it is rounded up.

Supported software

Reporting of the metric is supported for Microsoft SQL Server 2012, 2014, and 2016 (Standard and Enterprise).

Requirements

To report the number of virtual cores for the discovered Microsoft software, the licensed software must have Software Assurance.

Limitations

Reporting of the metric has the following limitations:

- By default, the supported products are assigned to the Microsoft Physical Core with SA metric. To calculate utilization of the virtual core based metric, reassign the products to this metric. For more information, see: [Assigning components to products \(on page dxlvi\)](#).
- BigFix Inventory does not recognize when a virtual machine is migrated to a server farm located in a time zone that is more than four hours apart from the current time zone. Instead of reporting two instances of a product that is installed on such a virtual machine, BigFix Inventory reports one instance.
- When a computer on which the software was discovered is deleted, report columns related to server capacity show `<no data>` in the user interface. Information about computers that contributed to the metric peak value is available in the audit snapshot regardless of whether these computers were removed.

Examples

The following examples describe how the number of required licenses is calculated. They are based on sample data.

Example 1: Software is deployed on one virtual machine

Microsoft SQL Server 2014 Enterprise is installed on a virtual machine that is assigned two virtual cores. According to the licensing terms, the minimal number of licensable cores is four. Metric utilization for Microsoft SQL Server is 4.

Table 58. Metric utilization for software on one virtual machine

	Virtual machine 1
Virtual cores	2
Core licenses needed	4

Example 2: Software is deployed on two virtual machines

Microsoft SQL Server 2014 Enterprise is installed on two virtual machines that run on the same physical host. The host has 10 cores. The first virtual machine is assigned 8 cores. The second virtual machine is assigned 16 cores. Metric utilization is the sum of virtual cores that are available to Microsoft SQL Server and equals 24.

Table 59. Metric utilization for software on two virtual machines

Summary for complex table

	Virtual Machine 1	Virtual Machine 2
Virtual cores	8	16
Physical cores on the host	10	
Core licenses needed	24	

Example 3: Software is deployed on a virtual machine that was assigned additional cores

Microsoft SQL Server 2014 Enterprise is installed on a virtual machine that is assigned 10 cores. The virtual machine is assigned two additional cores. Metric utilization for Microsoft SQL Server is 10 until the day when additional cores were added. Starting from the next day, metric utilization is 12.

Example 4: Software instance uses a different license metric

Two instances of Microsoft SQL Server 2014 Enterprise are discovered. The first instance uses the Microsoft Virtual Core with SA metric. The second instance uses the Microsoft Physical Core with SA metric. To ensure that the first instance contributes to the calculation of the virtual core based metric, reassign this instance to the proper metric. For more information, see: [Assigning components to products \(on page dxlvi\)](#).

Example 5: Two software instances use the same metric and are installed on the same virtual machine

Microsoft SQL Server 2012 Enterprise and Microsoft SQL Server 2014 Enterprise are discovered on one virtual machine. Both instances use the Microsoft Virtual Core with SA metric. Metric utilization is counted once for this virtual machine regardless of the number of installed instances.

Reading metric utilization

The first level of the All Metrics report shows a product and its metric quantity, which is the highest number of metric units that the product used during the specified period.

Publisher	Product Name	Metric	Metric Quantity	Metric Quantity History
				11/15/2016 - 02/13/2017
Microsoft	Microsoft SQL Server Enterprise E...	Microsoft Virtual Core with SA	16	
Microsoft	Microsoft SQL Server Standard Edi...	Microsoft Virtual Core with SA	4	

When you click the product name, the second level of the report shows all instances of product components that contributed to the metric quantity, including historical instances that are no longer installed. The values in the Logical Processors column show the number of virtual cores that are reported on the computers where the component is installed. Because the minimal number of licensable cores is four, virtual machines with one and two cores are counted as having four cores. This requirement is not reflected in the Logical Processor column because the column shows the actual number of processor cores.

Component Name	Version	Product Name	Metric	Excluded	Comput...	Server ID	Logical Pr...	Details
Microsoft Windows Server 2008 R2 Ent...	6.1	Microsoft SQL Server Enterprise Edition	Microsoft Virtu...	No	NC040202	IBM NC142...	2	DETAILS >
Microsoft SQL Server 2012 for Microsof...	2012.0	Microsoft SQL Server Enterprise Edition	Microsoft Virtu...	No	NC9143...	TLM_VM_4...	1	DETAILS >
Java Support	5.2	Microsoft SQL Server Enterprise Edition	Microsoft Virtu...	No	NC0472...	IBM 9117 0...	8	DETAILS >

The requirement is reflected in the Core Metric column in the `microsoft_virtual_core_with_sa.csv` file which is a part of the audit snapshot.

Table 60. A fragment of the `microsoft_virtual_core_with_sa.csv` file

Publisher	Product Name	Metric Quantity	Logical Processors	Core Metric
Microsoft	Microsoft SQL Server Enterprise Edition	16	2	4
Microsoft	Microsoft SQL Server Enterprise Edition	16	1	4
Microsoft	Microsoft SQL Server Enterprise Edition	16	8	8

Related information

[High-water mark \(on page lxxvi\)](#)

9.2.8 Oracle Processor Core

9.2.8 Available from 9.2.8. Oracle Processor Core metric is used to determine the licensing cost of Oracle products. It is based on the number of physical cores that are deployed on the host where the software is installed multiplied by an appropriate processor core factor. When the software is installed on multiple hosts, all cores that require applying the same core factor are added. Then, the factor is applied to calculate metric utilization.

By default, the 0.5 factor is applied to all servers. If the number of cores multiplied by the core factor is not an integer, it is rounded up at the core factor level. If a processor requires a different core factor, the value can be adjusted [through the user interface \(on page dlxxii\)](#) or [by using REST API \(on page dlxxiii\)](#).

Supported software

Reporting of the metric is supported for all discoverable versions of Oracle Database (Enterprise Edition).

Requirements

VM managers should be defined for all virtual machines in your environment. If a VM manager is not defined for a particular host, the number of cores is counted at the level of the virtual machine. As a result, license utilization might be incorrectly calculated. If a component of a product with the Oracle Processor Core metric is detected on a host for which a VM manager is not defined, a warning sign (⚠️) is displayed on the All Metrics report. It indicates calculation inaccuracy.



Note: The warning is displayed only on the user interface. It is not displayed in the CSV or PDF version of the All Metrics report.

In case of VMware, if a virtual machine is a cluster, metric utilization is counted on the basis of the number of cores in a cluster, not the physical server. To obtain correct count of cluster cores, go to **Management > Advanced Server**

Settings and set the value of the **storeHwDataForAllVMMManagerNodes** parameter to true. Otherwise, only hosts on which the BigFix client is installed are included in the cluster core count.

Limitations

Only active sockets are taken into account during the calculations.

Examples

Example 1: Software is installed on multiple servers that require various core factors

Oracle Database Enterprise Edition is installed on the following six servers:

- One server with two cores that require applying the 0.75 core factor
- Four servers with one core each that requires applying the 0.75 core factor
- One server with one core that requires applying the 0.5 factor

All cores that require the same core factor are added and the factor is applied. Then, the values obtained at each core factor level are rounded up and added. Metric utilization equals 6.

Table 61. Utilization of the Oracle Processor Core metric on multiple servers that require various core factors

Summary for complex table

	Cores that require the 0.75 factor	Cores that require the 0.5 factor
Actual number of cores	$(1 \times 2 \text{ cores} + 4 \times 1 \text{ core}) = 6$	$1 \times 1 \text{ core} = 1$
Number of cores after applying the core factor	$(1 \times 2 \text{ cores} + 4 \times 1 \text{ core}) \times 0.75 = 4,5$	$1 \times 1 \text{ core} \times 0.5 = 0.5$
Number of cores after rounding up	5	1
Total	6	

Related information

[High-water mark \(on page lxxvi\)](#)

Changing the Oracle core factor

By default, all servers are assigned the 0.5 factor. If the processor that is deployed on a server requires applying a different core factor, you can change it through the user interface or, to perform the operation on multiple servers, by using REST API.

9.2.8 Changing the Oracle core factor through the user interface

9.2.8 Available from 9.2.8. To change the Oracle core factor through the user interface, open the Hardware Inventory report, and edit the value in the Oracle Core Factor column.

 You must have the Manage Hardware Inventory permissions to perform this task.

1. Log in to BigFix Inventory, and go to **Reports > Hardware Inventory**.
2. To adjust the report view, hover over the **Manage Report View** icon , and click **Configure View**.
 - a. **Optional:** To narrow the report to a server for which you want to change the core factor, specify appropriate filters.

For example: `Computer Name`, `contains`, and specify the name of the server.

Filters

Specify the report filter which matches all of the following conditions:

Computer Name	contains	NC040202	✖	+
---------------	----------	----------	---	---

- b. To add a column in which you can edit the core factor, select **Oracle Core Factor**. Then, click **Submit**.
3. To change the Oracle core factor, click the pencil icon () in the **Oracle Core Factor** column, and click **Change Oracle Core Factor**.
4. Specify the new core factor, and click **Change**.

Change Oracle Core Factor
✖

Change the Oracle Core Factor for the processor that was identified on server IBM NC142253.

The value will be changed for the server and all virtual machines that are hosted on that server. After you change the value, open the All Metrics report, and click Recalculate to reflect the changes on the reports.

New Oracle Core Factor*

5. To update utilization of the Oracle Processor Core metric, go to **Reports > All Metrics**, and click **Recalculate**.

The core factor value is updated for the selected server. Metric utilization is updated for products that are installed on that server.

9.2.8 Changing the Oracle core factor through REST API

9.2.8 Available from 9.2.8. To change the Oracle core factor through REST API, retrieve the list of servers that exist in your infrastructure. Then, choose IDs of the servers for which you want to change the core factor and update the value.



You must have the Manage Hardware Inventory permissions to use this API.

1. Retrieve information about servers that exist in your infrastructure and the Oracle core factor that is assigned to every server.

- If you retrieve the information through a web browser plug-in, use the following request:

```
GET https://hostname:port/sam/processors.json?columns[]=server_id
&columns[]=server_name&columns[]=oracle_core_factor
&token=cble554922312e0e4aa44a1ef2b9a426117efe39
Accept: application/json
Accept-Language: en-US
```

- If you retrieve the information by using cURL, run the following command:

```
curl --insecure -H "Content-type: application/json; charset=UTF-8"
-X GET "https://hostname:port/sam/processors?columns%5B%5D=server_id
&columns%5B%5D=server_name&columns%5B%5D=oracle_core_factor
&token=cble554922312e0e4aa44a1ef2b9a426117efe39" -o out.json
```

The request returns the list of server IDs and names together with core factors that are currently assigned to each server.

```
{
  "total":3,
  "rows":
    [
      {
        "server_id":1,
        "server_name":"TLM_VM_42362841-6b4e-ea26-9755-07b28dc0fd41",
        "oracle_core_factor":"0.5"
      },
      {
        "server_id":2,
        "server_name":"Oracle Corporation SPARC-Enterprise-T5220 2225894042",
        "oracle_core_factor":"0.5"
      },
      {
        "server_id":3,
        "server_name":"Microsoft Azure VMware-42 3b f0 3c fb d5 4c 97-4f bd ca 3f 83 3a 91 46",
```

```
"oracle_core_factor": "0.5"
}]}
```

2. Update the core factor that is assigned to a selected server.

- If you update the information through a web browser plug-in, use the following request:

```
POST api/v1/servers?token=cble554922312e0e4aa44a1ef2b9a426117efe39
{
  "oracle_core_factors":
  [
    {
      "server_id": 1,
      "oracle_core_factor": 1.0
    }
  ]
}
```

- If you update the information by using cURL, run the following command:

```
curl --insecure -H "Content-type: application/json; charset=UTF-8"
-X PUT https://hostname:port/api/v1/servers?token=cble554922312e0e4aa44a1ef2b9a426117efe39
-d '{"oracle_core_factors":[{"server_id": 1,"oracle_core_factor": 1.4}]}'
```

3. To update utilization of the Oracle Processor Core metric, log in to BigFix Inventory and go to **Reports > All Metrics**. Then, click **Recalculate**.

Other IBM metrics

BigFix Inventory reports utilization for a vast number of BigFix license metrics. It also shows products that are assigned any IBM metric even if calculating utilization is not yet supported for this metric. BigFix Inventory also shows utilization of license metrics for products that deliver `.slmtag` files.

Other IBM metrics retrieved by BigFix Inventory

A number of IBM products deliver `.slmtag` files. For more information, see: [IBM products that deliver SLM tags](#).

Information about the utilization of license metrics for these products is presented on the following reports:

Resource Utilization

The report shows only the raw license utilization data that is retrieved from `.slmtag` files. For more information, see: [Raw utilization of license metrics \(on page dcxi\)](#).

9.2.13 All IBM Metrics and All Metrics

Starting from application update 9.2.13, BigFix Inventory reports additional IBM metrics basing on the resource utilization delivered by `.slmtag` files. BigFix Inventory scans a computer to find the software components that are reported with `.slmtag` files and then aggregates the metric quantity.



Note: The license metrics are reported only for the products that deliver `.slmtag` files and are detected and reported on the Software Classification panel. The license metric that is assigned to a product on the Software Classification panel must match the license metric embedded in the `.slmtag` file. Thus, some records displayed on the Resource Utilization panel might not be reflected on the All Metrics and All IBM Metrics reports.

9.2.14 Starting from application update 9.2.14, BigFix Inventory reports license consumption data from `.slmtag` files that refer to products. When the license usage that is included in `.slmtag` files is generated on product level, an artificial component is created in the catalog for that product. The component name follows the specific naming convention:

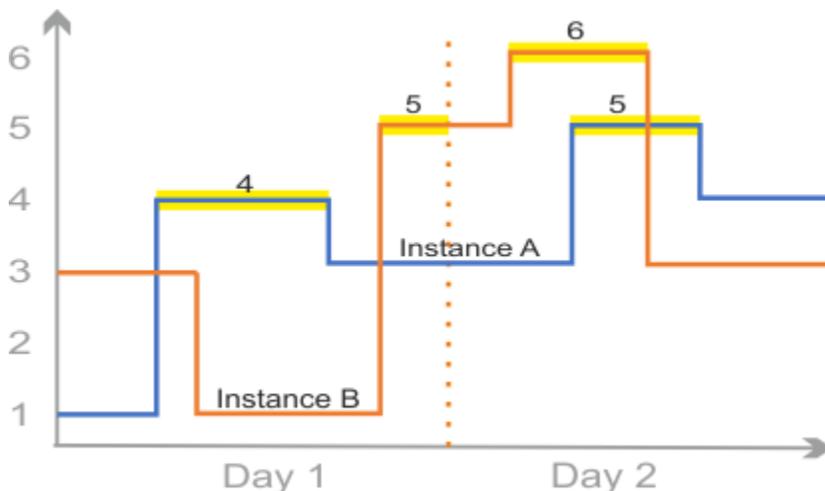
- License Reporter for `<product name>` with metric `<metric name>`; for example, License Reporter for BigFix Rational ClearCase with metric Floating User.

This component does not constitute a real part of the product and it is created only for internal reporting purposes.

Metric quantity that is shown on the report is counted in the following way:

- The aggregation algorithm sums the maximum daily values from all reporting software instances.
- The report shows the highest maximum daily value over a period of time.

The following chart shows the license metric utilization of two instances of a product over two days that is embedded in `.slmtag` files for two instances of the product over two days:



Day 1: $4 + 5 = 9$

Day 2: $6 + 5 = 11$

The highest reported maximum daily value over a period of time: **11**

According to the `.slmtag` data, both instances of the product are matched with the Authorized User metric. On the first day, the maximum daily utilization of instance A equaled 4 and of instance B - 5. Thus, the total metric quantity used on day one equaled 9. On the second day, the maximum daily utilization of instance A equaled 5 and of instance B - 6. Thus, the total metric quantity used on day two equaled 11. The highest maximum daily value over these two days is 11 and it is reported as such on the metric report.

If the overall metric quantity is not correct, the `.slmtag` data for the product is not proper. Contact the product support for a fix.

If you require an BigFix product to deliver `.slmtag` files, submit an idea for the product team. To submit an idea, see: [Submit an idea](#).

For the list of BigFix products that deliver `.slmtag` files, visit the [Resource Utilization Measurement Community](#).

9.2.7 IBM metrics for which calculating utilization is not yet supported

Apart from license metrics for which utilization is calculated, the All Metrics report also shows all other metrics that are used by BigFix products in your infrastructure. Even though utilization is not calculated for these metrics, it gives you an overview of all IBM metrics for which you should have licenses. For explanation of some of these license metrics, see: [Learn about Software licensing](#).

When no license metric is defined for a particular product in the software catalog, the metric that is displayed for this product is Unknown.

IBM Processor value unit (PVU)

A processor value unit (PVU) is a unit of measurement that is used to determine the licensing cost of BigFix middleware products. It is based on the type of processor that is deployed on the server where the software is installed. For each product, you must acquire the appropriate number of PVUs that is defined for the specific processor on which the software is installed.

The number of required PVUs is based on the processor technology and the number of processors that are available to the software product. For PVU-based licensing, IBM defines a processor as a processor core on a chip. For example, a dual-core processor chip has two processor cores. PVUs are assigned per core, not per processor. The number of PVUs that are assigned to a processor core is defined in the PVU table format. For more information, refer to [Processor Value Unit \[PVU\] licensing for Distributed Software](#).

Processor value unit license types

PVU licensing can be full capacity or subcapacity. Full capacity licensing is based on the overall number of processor cores on the physical server on which the product is installed. Subcapacity licensing is available only in virtualized environments. It is counted as the highest number of PVUs that are available for the VM on which the product is installed, not the total PVU count on the physical server. You acquire licenses for the lower value: subcapacity or full capacity.

Subcapacity licensing can significantly reduce the licensing cost, especially with a move to a newer processor core technology. However, virtual machines or partitions that restrict the available processor capacity must be created by using eligible virtualization technologies. For more information about subcapacity licensing requirements, see: [Passport Advantage® Virtualization Capacity \(Subcapacity\) Licensing](#).

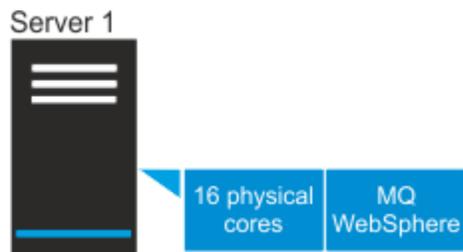
For information about rules of subcapacity counting on each virtualization type, see: [IBM Virtualization Capacity License Counting Rules](#).

Scenarios

The following scenarios describe how the cost of software license is calculated. The calculations are based on sample values and should be treated as examples.

Scenario 1: Full capacity on a physical server

IBM MQ software is installed on a server that has two Intel Xeon 3400 processors, each processor with eight cores (16 cores in total).



Because the environment is not virtualized, subcapacity license does not apply. Full capacity license is counted as the highest number of PVUs on the server where the software is installed. According to the PVU table, when the server has two sockets, this processor model is assigned 70 PVUs per core. The following table shows how the cost of full capacity license is calculated for IBM MQ software that is installed on this server.

Table 62. Cost of full capacity license for IBM MQ software

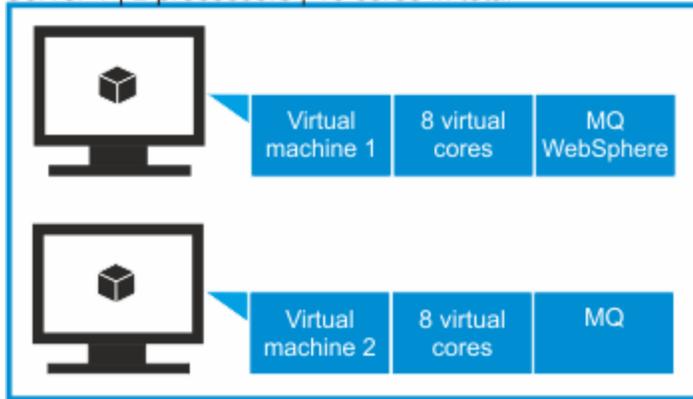
The table consists of two columns and five rows.

Description	Full capacity
Cores to license	16
PVU per core	70
Software cost per PVU	\$50
License cost	16 * 70 * \$50 = \$56,000

Scenario 2: Subcapacity on two virtual machines

Two virtual machines are deployed on a server that has two Intel Xeon 3400 processors, each processor with eight cores (16 cores in total). Each VM is assigned eight virtual cores. WebSphere software is installed only on the first virtual machine and has access to eight cores. IBM MQ software is installed on both virtual machines and thus has access to 16 cores.

Server 1 | 2 processors | 16 cores in total



According to the PVU table, when the server has two sockets, this processor model is assigned 70 PVUs per core. The following table shows how the cost of full capacity license and subcapacity license is calculated for IBM MQ software and WebSphere software that is installed on this server.

Table 63. Cost of full capacity and subcapacity license for IBM MQ software and WebSphere software

The table consists of two columns and five rows.

Description	Full capacity for IBM MQ and WebSphere software	Subcapacity for IBM MQ software	Subcapacity for WebSphere software
Cores to license	16	16	8
PVU per core	70	70	70
Software cost per PVU	\$50	\$50	\$50
License cost	16 * 70 * \$50 = \$56,000	16 * 70 * \$50 = \$56,000	8 * 70 * \$50 = \$28,000

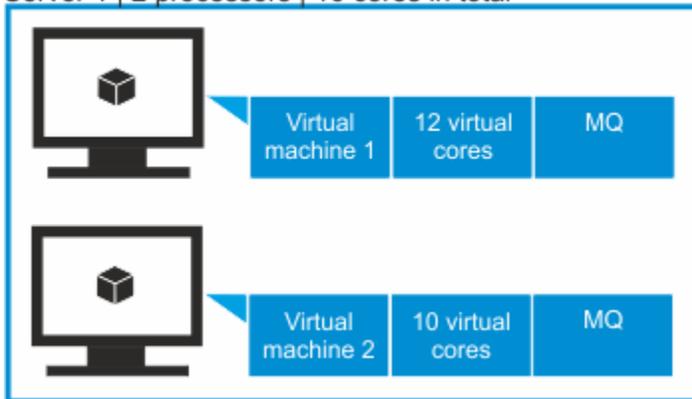
Scenario 3: Capacity in a virtual environment where connection to the VM manager is not configured

! **Important:** In this scenario, default PVU counting is applied due to the lack of connection to the VM manager. IBM accepts audit reports that contain PVU values calculated based on the default PVU counting instead of considering the client to be ineligible for subcapacity or liable

! for full capacity. For more information, see: [Default PVU counting on x86 processors \(on page cccxix\)](#) and [Managing VM managers \(on page cccxviii\)](#).

Two virtual machines are deployed on a server that has two Intel Xeon 3400 processors, each with eight cores (16 cores in total). The first VM is assigned 12 virtual cores. The second VM is assigned 10 virtual cores. IBM MQ software is installed on both VMs.

Server 1 | 2 processors | 16 cores in total



If a virtual machine manager is not defined, BigFix Inventory does not have access to information about the number of sockets and physical cores on the server on which the VMs are deployed. It has access only to information about the processor model and the number of virtual cores that each virtual machine is assigned. The sum of virtual cores to which IBM MQ software has access is greater than the number of physical cores that are available on the server.

i Tip: If the hypervisor data is missing from a virtual machine and default PVU counting is applied, ID of the server on which the machine runs begins with `TLM_VM`.

The number of PVUs per core that is assigned to a processor depends on the number of sockets on the host. Because BigFix Inventory does not have access to this information, it assigns the highest PVU per core value that is defined for the particular processor in the PVU table. According to the PVU table, when the server has two sockets, this processor model is assigned 70 PVUs per core. However, due to the lack of information about the number of sockets, software is charged 120 PVUs per core. It is the highest number of PVUs that is defined in the PVU table for this type of processor.

The following table shows the cost of license for the IBM MQ software depending on whether the VM manager is defined.

Table 64. Cost of license for IBM MQ software depending on whether the VM manager is defined

The table consists of two columns and five rows.

Description	Capacity for IBM MQ software when the VM manager is defined		Capacity for IBM MQ software when VM manager is not defined
	Full capacity	Subcapacity	
Cores to license	16	22 capped to 16	22
PVU per core	70	70	120
Software cost per PVU	\$50	\$50	\$50
License cost	16 * 70 * \$50 = \$56,000	16 * 70 * \$50 = \$56,000	22 * 120 * \$50 = \$132,000

Related information

[High-water mark \(on page lxxvi\)](#)

IBM Processor value unit table

A processor value unit table contains information about the conversion ratio that is used for available processor types.

The processor value unit table supports processor-based pricing models in which charges differ according to the type of processor on which the licensed product is installed and running. In the table, a number of units is assigned to each processor type for which this type of pricing model is available. It is necessary to ensure that the table is updated to reflect the ratios for all processors in use.

If a processor cannot be properly identified by an agent, it is shown in the processor table as OTHER. This case can occur because it is a new processor type and its discovery is not yet supported. In this situation, upload the newest processor value unit table.

Related information

[Uploading an IBM PVU table \(on page dlxxx\)](#)

Uploading an IBM PVU table

The PVU table supports processor-based pricing models in which charges differ according to the type of processor on which the licensed product is installed and running. In the table, a number of units are assigned to each processor type on which this type of pricing model is available. BigFix continues to update the PVU table to include new processor types and models. To ensure that the processors are properly detected and charged, upload a new version of the PVU table whenever it is available.

 You must have the Manage Uploads permission to perform this task.

9.2.6 A new PVU table is uploaded to BigFix Inventory during the upgrade. Use this procedure when a new table is published between application updates.

The PVU table file is authenticated with SHA-256 to comply with the FIPS (Federal Information Processing Standards) security standards. Previous versions of PVU tables that are signed with SHA1 are not supported.

1. In the top navigation bar, click **Management > Metric Table Upload**.
2. To download PVU table, use **Download PVU table** fixlet.



Note: **Download PVU table** fixlet is available since 10.0.4.

3. To upload PVU table, click **Choose File**, select the file that you downloaded, and click **Upload**.
A new entry is created in the **Upload and Import History** table. The status of the PVU table is Pending.
4. Wait for the scheduled import or run it manually.
For BigFix Inventory to be compliant with Subcapacity, you must upgrade the server.

Related information

[IBM Processor value unit table \(on page dlxxx\)](#)

9.2.13 Registered User

9.2.13 Available from 9.2.13. Registered User is a license metric unit that can be used to determine the licensing cost of Microsoft Office 365 and a number of Adobe Creative Cloud products.

A product can be licensed based on the number of registered users. Then, one product that is installed on multiple computers or servers can be used by a number of registered users. Only the number of registered users is taken into account when calculating license metric utilization.

Note that the Registered User metric differs from the Authorized User metric. Authorized User has a broader spectrum and relates to all users who are authorized to use a product. Registered User is more accurate and pertains only to those users who registered.

Supported Software

Windows Reporting of the Registered User metric is supported for the following software that is installed on Windows computers:

- Microsoft Office 365
- **9.2.14** Adobe Creative Cloud (CC) applications on Windows computers.
 - Adobe Illustrator CC
 - Adobe PhotoShop CC

- Adobe InDesign CC
- Adobe Premiere Pro CC
- Adobe After Effects CC
- Adobe Dreamweaver CC
- Adobe Muse CC

Limitations

- BigFix Inventory extracts users only of those Office 365 components in which the user logs in to Microsoft. If an Office 365 component is deployed but the user does not log in to Microsoft, the user is not reported. Due to this fact, users are reported only when they are logged in to at least one of the following components: Access, Excel, PowerPoint, Publisher, Word.
- Because BigFix Inventory extracts only users that log in to Microsoft, users are reported only for components that provide the option to log in to Microsoft. For example, users of Skype are reported only if they use another component of Office 365 and log in to Microsoft.
- If the software scan runs during a time when the user of Office 365 is not logged in to Windows, information about the user is not captured. If the user was reported in the past and Office 365 is still deployed on the scanned computer, the user is counted. If the user was reported in the past but Office 365 is no longer deployed nor activated, the user is not counted.
- Utilization of the Registered User metric is calculated only for products that are installed on Windows. Product instances that are installed on computers with other operating systems, on mobile devices, or are available online are not taken into account. Thus, the reported license metric utilization might differ from the number provided by the software manufacturer.
- BigFix Inventory incorrectly reports users of Office 365 that is installed in the Shared Client Activation (SCA) mode.
- To display the information about software users of Adobe CC applications on the Software Users report, you must have Adobe Creative Cloud Desktop App installed on all relevant computers. Otherwise, BigFix Inventory measures and reports only the usage of these components that is based on the Windows processes.

Examples

Publisher	Product Name	Metric	Metric Quantity	Metric Quantity History	
				07/30/2018	08/09/2018
Microsoft	Microsoft Office 365	Registered User	7		

Related information

[High-water mark \(on page lxxvi\)](#)

Setting and removing the declared metric quantity

9.2.16 Available from 9.2.16. On the All Metrics report, you can manually declare the metric quantity for selected license types for which BigFix Inventory does not calculate the utilization. You can also use this option to customize the metric quantity if needed. The declared value takes precedence over the value that is measured by the application.

 You must have the Manage License Metrics and the Manage Software Bundling permissions to set the license metric quantity.

You can check which of your license metric quantities can be customized. Look at the pencil icon  in the **Metric Quantity** column. If the icon is active, you can declare the custom metric quantity for the product.

By default, you set the declared value from the reporting start date onwards. However when you explicitly provide also the end date for the report, then the declared metric quantity is set only until this date. Using this mechanism, you can declare metric quantities for any period of time by defining the appropriate time range for the report.

1. Open the All Metrics report.
2. **Optional:** To change the period for which you want to set the declared metric quantity, hover over the **Manage Report View** icon , and click **Configure View**. Then, select the time range. If you specify a period from a date to a date, the license metric quantity is set only for this particular period of time.
3. To set the metric quantity, click the pencil icon () in the **Metric Quantity** column.

Publisher	Product Name	Metric	Metric Quantity	Metric Quantity History
IBM	IBM Virtual Desktop	Concurrent User		02/23/2019 - 05/24/2019
IBM	IBM i2 iBridge Designer	Concurrent User		



Note: In case of FlexPoint bundles, you can set the license metric quantities for particular products that contribute to the bundle but not for the bundle itself. Click on the name of the bundle to see the detailed report that lists the products within bundle.

4. Specify the metric quantity and click **Set**. The value is set only for the computer group to which you are assigned and is not inherited by its subgroups.

Set Metric Quantity ✕

Set the declared metric quantity. The value that you provide takes precedence over the value that was measured by the application. You can remove the declared value to restore the initial one at any time.

Time range: from 06/02/2019 ?

Product: IBM Virtual Desktop

Metric: Concurrent User

Computer group: All Computers

Metric quantity ?



Note: If you want to set a metric quantity for a subgroup of your computer group, go to **Reports > Computer Groups** and open the subgroup. In the upper left corner, select All Metrics and set the license metric quantity for a particular product.

The metric quantity for the product is set.

The declared metric quantity is independent of your software inventory or bundling and you should make sure that it is kept up-to-date. If you reassign software to another product or FlexPoint bundle, the declared metric quantity remains unchanged. If, as a result of reassignment, a product is left with no components, it is still visible on the All Metrics report with the declared metric quantity.

Removing the declared metric quantity

You can remove the declared value to restore the value that is measured by the application.

1. Open the report.
2. To remove the declared metric quantity, click the pencil icon () in the **Metric Quantity** column.
3. Clear the text field to remove the declared metric quantity for the specified period of time.
4. Click **Set**.

IBM Resource value unit (RVU MAPC)

Resource value unit (RVU MAPC) is a license metric unit that is based on the number of activated processor cores that are available to the product. An activated processor core is a core that is managed or used by a product, regardless of whether the capacity of the processor core can be limited through virtualization technologies.

Resource value unit license types

RVU MAPC licensing can be full capacity or subcapacity. Full capacity licensing is counted as the highest number of RVUs in the physical environment that are managed or used by the product.

Subcapacity licensing is available only in virtualized environments. It is counted as the highest number of RVUs that are available to the VM on which the product is installed, not the total RVU count on the physical server.

Resource value unit tier table

The RVU tier table provides factors that are specific to the number of activated processor cores. The greater the number of activated processors, the lower is the factor and thus the fewer RVUs are required per processor core.

Table 65. The RVU tier table for activated processor cores

Tier	Activated processor cores	Factor
1	0 - 2,500	1.00
2	2,501 - 10,000	0.80
3	10,001 - 50,000	0.60
4	50,001 - 150,000	0.40
5	Above 150,001	0.20

Example

A product has access to 45,000 activated processor cores. The number is divided into tiers based on the RVU tier table. The number of cores that fall into each tier is multiplied by the factor that is applicable for that tier. The values are summed up to obtain the total number of required RVUs. The following table shows how the number of required RVUs is obtained for this product.

Table 66. Activated processor cores divided into quantity tiers

Tier	Number of managed cores in the tier	Calculations
1	2,500	$2,500 \times 1.00 = 2,500$
2	7,500	$7,500 \times 0.80 = 6,000$
3	35,000	$35,000 \times 0.60 = 21,000$
Total	45,000	29,500

Related information

[High-water mark \(on page lxxvi\)](#)

IBM Virtual processor core (VPC)

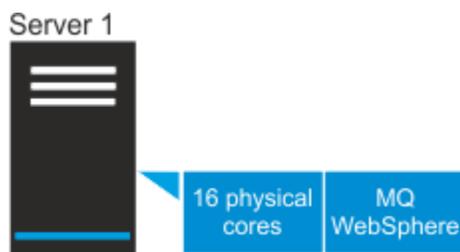
9.2.11 Available from 9.2.11. A virtual processor core (VPC) is a unit of measurement that is used to determine the licensing cost of BigFix products. It is based on the number of virtual cores (vCPUs) that are available to the product. A vCPU is a virtual core that is assigned to a virtual machine or a physical processor core if the server is not partitioned for virtual machines.

Scenarios

The following scenarios describe how the number of vCPUs is counted under different circumstances. The calculations are based on sample values and should be treated as examples.

Scenario 1: VPC on a physical server

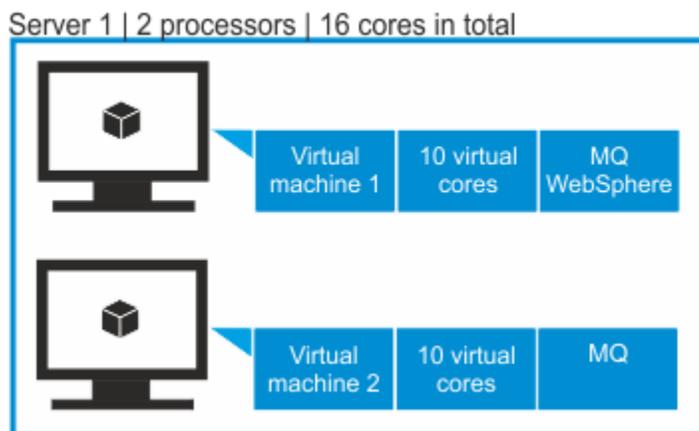
IBM MQ software is installed on a server that has two processors. Each processor has 8 physical cores which gives 16 cores in total. The number of vCPUs to license is 16.



Scenario 2: VPC on two virtual machines

Two VMs are deployed on a server that has two processors. Each processor has 8 physical cores which gives 16 cores in total. Each VM is assigned 10 vCPUs.

WebSphere software is installed only on the first VM and has access to 10 vCPUs. The number of vCPUs to license is 10. IBM MQ software is installed on both VMs and thus has access to 20 vCPUs. Because the number of physical cores is lower than the number of vCPUs that are assigned to the VMs, the number of vCPUs to license is capped to 16.

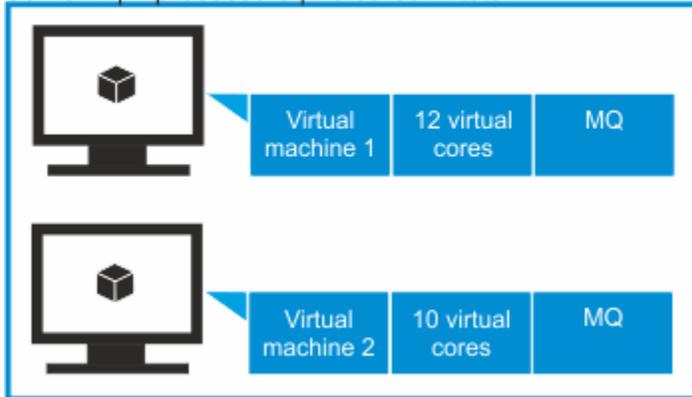


Scenario 3: VPC in a virtual environment where the VM manager is not defined

! **Important:** This scenario is not supported. It is provided to explain the consequences of not defining the VM manager that controls resources that are available to virtual machines. VM managers must be defined to properly report VPC consumption.

Two VMs are deployed on a server that has two processors. Each processor has 8 cores which gives 16 physical cores in total. The first VM is assigned 12 vCPUs. The second VM is assigned 10 vCPUs. It gives 22 vCPUs in total. IBM MQ software is installed on both VMs.

Server 1 | 2 processors | 16 cores in total



If a connection to the VM manager is not defined, BigFix Inventory does not have access to information about the number of sockets and physical cores on the server on which the VMs are deployed. It has access only to information about the number of vCPUs that each VM is assigned. The sum of vCPUs to which IBM MQ software has access is greater than the number of physical cores that are available on the server.

If a connection to the VM manager is not defined, the number of vCPUs to license is 22. If the connection is defined, 22 vCPUs are capped to 16 physical cores and thus the number of vCPUs to license is 16.

Related information

[High-water mark \(on page lxxvi\)](#)

License metric utilization

The primary source of information about license metric utilization is the All Metrics report. For every product, you can specify a metric threshold to verify whether metric utilization is above or below your expectations. You can also create a snapshot of the report, and store it for audit purposes.

Report overview

1 All Metrics Send Feedback

Time Range: 05/10/2018 - 08/08/2018 UTC
 Current Catalog Version: 9.2.12.1
 Endpoint Scanner Catalog: 0

(Base Report) Recalculate Audit Snapshot **2**

Publisher	Product Name	Metric	FlexPoint Bundle	Metric Quantity	Threshold	Threshold Delta	Imported Pa...	Metric Quantity History
IBM	IBM Db2 Advanced Wo...	PVU Full Capa...	(none)	960	1,000	40	E0HP3LL;D0...	05/10/2018 - 08/08/2018
IBM	IBM Db2 Connect Appli...	PVU Full Capa...	(none)	960	1,500	540	E022ILL;D55...	
IBM	IBM Homogeneous Re...	PVU Full Capa...	(none)	960	750	-210	D61P2LL;E04...	
IBM	IBM Db2 Advanced Ent...	PVU Full Capa...	(none)	480	<not set>	<no data>	D0ZVOLL;D0...	
IBM	IBM Db2 Direct Advanc...	Virtual Proces...	(none)	8	<not set>	<no data>	<no data>	
IBM	IBM Db2 Big SQL	Virtual Proces...	Hybrid Data Mana...	4	<not set>	<no data>	<no data>	
IBM	IBM InfoSphere Optim ...	Managed Virt...	(none)	<n/a>	<not set>	<no data>	D09UWLL;D0...	<n/a>
Microsoft	Microsoft Help Viewer	Install Seats	(none)	4	<not set>	<no data>	<no data>	
Microsoft	Microsoft SQL Server ...	Microsoft Phy...	(none)	0	<not set>	<no data>	<no data>	

3 **4** **5** **6** **7** **8** **9** **10**

1 9.2.4 Predefined report views

Predefined report views give you quick access to filtered information. You can schedule these reports to be sent and set them as default views. However, you cannot modify them.

2 9.2.9 Thresholds and Custom Fields

Custom fields allow you to extend information that is related to software products by adding details pertinent to your organization. You can then use the information to filter and sort the reports as well as group products based on flexible criteria. A value that is specified in a custom field applies to the selected product with the particular license metric. It also applies only within the computer group to which you are assigned. For more information, see: [Adding custom information to license metric reports \(on page dxciii\)](#).

The menu can also be used to set license metric thresholds in the same way as in the [Threshold \(on page dxc\)](#) column.

3 Product Name

Name of the licensed product. When you click the link for a product that uses the PVU or RVU MAPC metric, the next level of the report shows servers on which the product was installed when the metric quantity was the highest.

When you click the link for any other product, the next level shows all instances of product components, including historical instances that are no longer installed. To view the list of component instances that were installed when the metric quantity was the highest, generate the audit snapshot.

4 Metric

License metric that is assigned to the product. The metrics fall into three categories:

- Metrics for which utilization is calculated, including:
 - 9.2.7 Install Instances
 - 9.2.5 Install Seats
 - 9.2.8 Microsoft Dual Physical Processor
 - 9.2.8 Microsoft Single Physical Processor
 - 9.2.7 Microsoft Physical Core with SA
 - 9.2.7 Microsoft Virtual Core with SA
 - 9.2.8 Oracle Processor Core
 - Processor Value Unit (PVU)
 - 9.2.13 Registered User
 - Resource Value Unit (RVU MAPC)
 - 9.2.11 Virtual Processor Core (VPC)
- 9.2.7 Metrics for which calculating utilization is not yet supported. These metrics are available only for BigFix products.
- 9.2.7 Metric is unknown

For more information about the license metrics, see: [Reported license metrics \(on page dlvi\)](#).

 **Tip:** If you want to display only products for which metric utilization is calculated, hover over the **Manage Report View** icon , and click **Configure View**. Then, set the following filter: `Metric Quantity, greater than or equal to, 0`.

5 9.2.13 FlexPoint Bundle

FlexPoint bundle to which the product and its components are assigned. For more information, see: [Assigning products to FlexPoint bundles \(on page dli\)](#).

6 Metric Quantity

The highest number of metric units that the product used during the specified period. The value `<n/a>` indicates that the quantity is not calculated for this metric.

The following icons indicate that further actions are required to ensure that the metric quantity is up-to-date:

- Red square () indicates that the metric quantity for the product needs to be updated because an action that influences the value was performed. For example, a component was assigned to this product. To update the metric quantity, click **Recalculate**. The progress of recalculation is displayed below the top menu bar.
- Exclamation mark () indicates that metric quantity might not be calculated correctly because inventory data from at least one VM manager is not reported properly. To solve the problem, ensure that all VM managers are configured. For more information, see: [Adding VM managers in central mode \(on page cccl\)](#) or [Adding VM managers in distributed mode \(on page ccclxxii\)](#).

7 9.2.4 Threshold

The maximum number of metric units that the product can use within the computer group. The value is set manually. For more information, see: [Setting license metric threshold \(on page dxc\)](#).

8 9.2.4 Threshold Delta

The difference between the threshold and the metric quantity. It indicates whether the metric utilization is above or below the manually set threshold. The following values indicate that the delta cannot be calculated:

- The value `<no data>` indicates that the threshold is not set.
- The value `<n/a>` indicates that the threshold was set for a product for which metric quantity is not calculated.

9 Imported Part Numbers

Part numbers that represents products that you purchased. For more information, see: [Part numbers \(on page dxxxiv\)](#).

10 Metric Quantity History

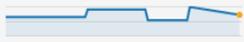
History of changes in the metric utilization over the specified period. The history is not available for products for which metric quantity is not calculated.

9.2.4 Setting license metric threshold

9.2.4 Available from 9.2.4. You can set a threshold to indicate the maximum number of metric units that a product can use within a computer group. The threshold is then used to calculate whether metric utilization is above or below your expectations.

 You must have the View License Metrics permission to view the threshold and the Manage Contracts permission to change it.

1. Open the All Metrics report and from the list of predefined reports, choose Metric Thresholds.
2. **Optional:** To filter down the report to the product for which you want to set the threshold, hover over the **Manage Report View** icon , and click **Configure View**. Then, specify appropriate filters. For example, choose Product Name, contains and specify the name of the product. Then, click **Submit**.
3. To set the threshold, click the pencil icon () in the **Threshold** column, and click **Set Threshold**.

Publisher	Product Name	Metric	Metric Quantity	Threshold	Threshold Delta	Metric Quantity History
IBM	IBM MQ	PVU Full Capacity	270	<not set> 	<no data>	
IBM	IBM MQ	PVU Subcapacity	270	Set Threshold	<no data>	

4. Specify the threshold and click **Set**. The value is set only for the computer group to which you are assigned and is not inherited by its subgroups.

Set Threshold ✕

Set the metric threshold to indicate the maximum number of metric units that the product is entitled to use.

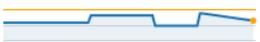
Product: IBM MQ
Metric: PVU Full Capacity
Computer group: All Computers

Metric threshold*

Set

 **Tip:** If you want to set a threshold for a subgroup of your computer group, go to **Reports > Computer Groups** and open the subgroup. In the upper left corner, select All Metrics. Then, from the list of predefined reports, choose Metric Thresholds and set the threshold value for a particular product.

When you set the threshold, its delta is automatically calculated by subtracting the metric quantity from the threshold. If the utilization is lower than or equal to the threshold, the delta is displayed in green. Otherwise, it is displayed in red. Additionally, a line that represents the threshold is displayed on the metric quantity history graph.

Publisher	Product Name	Metric	Metric Quantity	Threshold	Threshold Delta	Metric Quantity History
IBM	IBM MQ	PVU Full Capacity	270	300 	30	
IBM	IBM MQ	PVU Subcapacity	270	300 	30	

If you want to remove the threshold, click the pencil icon () and click **Clear Threshold**. Then, click **Clear**.

From the list of predefined reports, you can choose the Exceeded Thresholds report and schedule emails to be sent whenever a product appears on that report. It will ensure that you are automatically notified when metric utilization for any product exceeds the threshold. For more information, see: [Scheduling report emails \(on page dclv\)](#).

Setting and removing the declared metric quantity

9.2.16 Available from 9.2.16. On the All Metrics report, you can manually declare the metric quantity for selected license types for which BigFix Inventory does not calculate the utilization. You can also use this option to customize the metric quantity if needed. The declared value takes precedence over the value that is measured by the application.

 You must have the Manage License Metrics and the Manage Software Bundling permissions to set the license metric quantity.

You can check which of your license metric quantities can be customized. Look at the pencil icon  in the **Metric Quantity** column. If the icon is active, you can declare the custom metric quantity for the product.

By default, you set the declared value from the reporting start date onwards. However when you explicitly provide also the end date for the report, then the declared metric quantity is set only until this date. Using this mechanism, you can declare metric quantities for any period of time by defining the appropriate time range for the report.

1. Open the All Metrics report.
2. **Optional:** To change the period for which you want to set the declared metric quantity, hover over the **Manage Report View** icon , and click **Configure View**. Then, select the time range. If you specify a period from a date to a date, the license metric quantity is set only for this particular period of time.
3. To set the metric quantity, click the pencil icon () in the **Metric Quantity** column.

Publisher	Product Name	Metric	Metric Quantity	Metric Quantity History
IBM	IBM Virtual Desktop	Concurrent User		02/23/2019 - 05/24/2019
IBM	IBM i2 iBridge Designer	Concurrent User		



Note: In case of FlexPoint bundles, you can set the license metric quantities for particular products that contribute to the bundle but not for the bundle itself. Click on the name of the bundle to see the detailed report that lists the products within bundle.

4. Specify the metric quantity and click **Set**. The value is set only for the computer group to which you are assigned and is not inherited by its subgroups.

Set Metric Quantity ✕

Set the declared metric quantity. The value that you provide takes precedence over the value that was measured by the application. You can remove the declared value to restore the initial one at any time.

Time range: from 06/02/2019 ?

Product: IBM Virtual Desktop

Metric: Concurrent User

Computer group: All Computers

Metric quantity ?

Set



Note: If you want to set a metric quantity for a subgroup of your computer group, go to **Reports > Computer Groups** and open the subgroup. In the upper left corner, select All Metrics and set the license metric quantity for a particular product.

The metric quantity for the product is set.

The declared metric quantity is independent of your software inventory or bundling and you should make sure that it is kept up-to-date. If you reassign software to another product or FlexPoint bundle, the declared metric quantity remains unchanged. If, as a result of reassignment, a product is left with no components, it is still visible on the All Metrics report with the declared metric quantity.

Removing the declared metric quantity

You can remove the declared value to restore the value that is measured by the application.

1. Open the report.
2. To remove the declared metric quantity, click the pencil icon () in the **Metric Quantity** column.
3. Clear the text field to remove the declared metric quantity for the specified period of time.
4. Click **Set**.

9.2.9

Adding custom information to license metric reports

9.2.9

 Available from 9.2.9. You can extend information that is related to software products by adding custom fields on the All Metrics report and its derivatives. You can then use the information to group products based on flexible criteria such as contract name, placement, allocation, date, or other.


You must have the Manage Contracts permission to add and edit custom fields, and the View License Metrics permission to view them.

1. To add a custom field that is displayed on the All Metrics report and its derivatives, perform the following steps.

- a. In the top navigation bar, click **Management > Metric Custom Fields**.
- b. To add a custom field, click **New**.
- c. Provide the name of the field.
For example, Needs Renewal.
- d. Choose the type of data to be entered in the field.
For example, **Boolean**.

Create Metric Custom Field

Name*

Type* ▼



Note: After you create a custom field, you can edit only its name. You cannot change the type of data that is required in that field.

- e. Click **Create**.

2. To specify the value of a custom field, perform the following steps.

- a. In the top navigation bar, click **Reports > All Metrics**.
- b. Select the row for which you want to specify the value. The value applies only to the selected product with the particular license metric. It is also set only for the computer group to which you are assigned and is not inherited by its subgroups.



Tip: If you want to set the value for a subgroup of your computer group, go to **Reports > Computer Groups** and open the subgroup. Then, set the value for a particular product within the selected subgroup.

- c. Hover over the **Threshold and Custom Fields**  icon, and click **Set Value**.
- d. Specify the value, and click **OK**.

Set Threshold and Custom Fields ✕

Set the metric threshold to indicate the maximum number of metric units that the product is entitled to use.

Product: IBM Tivoli Storage Manager
Metric: PVU Full Capacity
Computer group: All Computers

Threshold

Needs Renewal

 **Tip:**

- If the type of the custom field is Boolean, the check box that is displayed on the reports has three possible values.
 - Not set
 - Yes
 - No

3. To display the custom field on the All Metrics report, perform the following steps.

- a. Hover over the **Manage Report View** icon , and click **Configure View**.
- b. Select the custom field, and click **Submit**.

You created a custom field that can be used to provide specific information about a product with a particular license metric.

Product Name	Metric	Metric Quantity	Contract Name	Contract Renewal Date	Needs Renewal	Contact Person	Metric Quantity History
							05/30/2017 - 08/28/2017
IBM Maximo Archivin...	PVU Full Capacity	480	Maximo	12/31/2018	No	John Doe	
IBM Maximo Archivin...	PVU Subcapacity	480	Maximo	12/31/2018	No	John Doe	
WebSphere Service ...	PVU Full Capacity	480	WebSphere	12/31/2017	Yes	Ian Smith	
WebSphere Service ...	PVU Subcapacity	480	WebSphere	12/31/2017	Yes	Ian Smith	

If you want to remove a value in the custom field, select a specific row, and hover over the **Threshold and Custom Fields**  icon. Then, click **Clear Value**. If you want to remove a custom field, go to **Management > Metric Custom Fields**, select the field that you want to remove, and click **Delete**.

You can use values in the custom fields to filter and sort the reports. You can also create saved report views that are filtered by the specified information, and enable notifications that are triggered when specific criteria are met.

Creating snapshots of license metric utilization for auditing purposes

Create the snapshot of the **All Metrics** report, and store it as a record of license metric utilization in your enterprise over a particular period. You can create the snapshot with any frequency that is suitable for your purposes. However, it is a good practice to create the snapshot every quarter. You can also regenerate the snapshot at any time as the historical data is always editable.

 You must have the following permissions to perform this task:

- View License Metric to open the **All Metrics** report.
- View Audit Trail if you want the snapshot to contain the audit trail file.

To ensure completeness and accuracy of data that is contained in the audit snapshot, ensure that:

- Software scan data is collected from the entire environment. To verify whether there are any problems, check the **Scan Health widget (on page liv)** on the **Overview** panel.
- Capacity data is collected from the entire environment. To verify whether there are any problems, check the **IBM Capacity Data Completeness widget (on page liv)** on the **Overview** panel.
- Imports of data are successful. To check the import history, go to **Management > Data Imports**.
- Software is correctly classified. For more information, see: [Software classification \(on page dxxviii\)](#).

1. Log in to BigFix Inventory, and go to **Reports > All Metrics**.
2. **Optional:** To configure the period for which you want to create the snapshot, hover over the **Manage Report View** icon , and click **Configure View**. Then, select the time range.
 - If you specify a period from a particular date until the last successful import, the reports contain entries from that date starting at 00:00:00.
 - If you specify a period from a date to a date, the reports contain entries from the first date starting at 00:00:00 to the second date ending at 23:59:59.
3. **Optional:** To limit the number of products that are included in the snapshot, specify the name of the product in the **Filters** section.

Time Range

90 before the last successful import
 06/18/2015 to the last successful import
 06/18/2015 to

Filters

Specify the report filter which matches of the following conditions:

4. To create the snapshot, in the top navigation bar, click **Audit Snapshot**.
5. **Optional:** Provide a comment in the comment text field.
6. **Optional:** To add an attachment, click **Browse** and select the file that you want to attach to the snapshot.

Generate Audit Snapshot ✕

Before you generate the Audit Snapshot, you can optionally add a comment or attach any file. They will be included in the generated package.

Provide your comment here

The attachment contains information why particular instances were excluded from pricing calculations.

Select a file that you want to attach

exclusion_reasons.xls

7. Click **Generate**.

The `audit_snapshot_date_time.zip` file is downloaded to your computer. The date that is provided as part of the file name is expressed in the GMT time zone. For information about the content of the file, see: [Content of the audit snapshot \(on page dxcvii\)](#).

Content of the audit snapshot

Audit snapshot is a record of license metric utilization in your enterprise over a period of time. It is generated on the All Metrics report and its derivatives. It is downloaded to your computer as the `audit_snapshot_date_time.zip` file that contains a number of other files.

! **Important:** If you use the audit snapshot for integration of BigFix Inventory with external tools, base the integration on column names, not their order. The order of columns might change if new columns are added.

Table 67. General files

File	Content
<code>audit_attachment.extension</code>	Optional. A file that was attached during the creation of the snapshot. The file is listed in the <code>checksums.txt</code> file.

Table 67. General files (continued)

File	Content
<code>audit_comments.txt</code>	Optional. Comments that were added during the creation of the snapshot. The file is listed in the <code>checksums.txt</code> file.
9.2.12 <code>audit_summary.csv</code>	<p>Summary of all software products that were installed in your environment during the period for which the snapshot is generated and which are charged. The file contains information about:</p> <ul style="list-style-type: none"> • License metric that is used by the product. • The highest number of metric units that the product used during the specified period. • Part numbers that represents products that you purchased. • The need to perform recalculation to update the metric quantity for the product.
<code>audit_trails.csv</code>	<p>Information about:</p> <ul style="list-style-type: none"> • Software classification actions • Uploads and imports of the PVU table, software catalog, and charge units • Changes to the VM managers <p>The file is not generated if the user who generates the snapshot does not have the permission to view audit trails. It is also not generated when the snapshot contains only products that use license metrics other than PVU and RVU MAPC. For example, when the snapshot contains only products that use the Install Seats metric.</p>
<code>bundling_definitions.csv</code>	<p>Information about:</p> <ul style="list-style-type: none"> • Software limited to the vendor that was specified on the user interface • 9.2.7 License metric that each product uses • Details of computers on which the software is installed
<code>checksums.txt</code>	<p>Checksums that are used to verify whether any files from the audit snapshot were tampered with.</p> <p>Linux To verify the report files on Linux, run the following command:</p> <pre>sha256sum -c pvu_sub_capacity.csv</pre> <p>Windows To verify the report files on Windows, run the following command:</p> <pre>CertUtil -hashfile pvu_sub_capacity.csv SHA256</pre>
<code>data_condition.txt</code>	<p>Information about:</p> <ul style="list-style-type: none"> • 9.2.9 Version of BigFix Inventory in which the snapshot was generated • Author of the snapshot • Date of the snapshot

Table 67. General files (continued)

File	Content
	<ul style="list-style-type: none"> • Period that is covered by the snapshot • Version of the PVU table • Version of the software catalog • Status of aggregation and recalculation • Filters that are used on the reports • Attached file (optional)
<code>part_numbers.csv</code>	Part numbers that are listed in the software catalog for software products included in any of the CSV report files.
<code>pub_key.pem</code>	Public key file that can be used to verify the <code>signature.rsa</code> file against the <code>checksums.txt</code> file.
<code>signature.rsa</code>	A digital signature that can be used to verify whether the <code>checksums.txt</code> file was tampered with. To verify the file, you can use OpenSSL.
	<p>Linux To verify the <code>checksums.txt</code> file on Linux, run the following command:</p> <pre>openssl dgst -sha256 -verify pub_key.pem -signature signature.rsa checksums.txt</pre>
	<p>Windows To verify the <code>checksums.txt</code> file on Windows, run the following command:</p> <pre>openssl -sha256 -verify pub_key.pem -signature file signature.rsa checksums.txt</pre>

The following files are included in each audit snapshot unless you filter the report to show only a specific metric type, for example PVU subcapacity. In this case, files with other metrics are not included.

Table 68. Files with license metric utilization

File	Content
9.2.16 <code>flexpoints.csv</code>	<p>Information about:</p> <ul style="list-style-type: none"> • Software licensed per number of FlexPoints • Details of software products that are licensed under the Flex Point bundles
9.2.7 <code>install_instances.csv</code>	Information about:

Table 68. Files with license metric utilization (continued)

File	Content
	<ul style="list-style-type: none"> • Software licensed per number of installed instances • 9.2.13 FlexPoint offerings to which products and their components are bundled • Details of computers on which the software was installed when the metric quantity was the highest
9.2.5 <code>install_seats.csv</code>	<p>Information about:</p> <ul style="list-style-type: none"> • Software licensed per number of computers on which it is installed • Details of computers on which the software was installed when the metric quantity was the highest
9.2.8 <code>microsoft_dual_physical_processor.csv</code>	<p>Information about:</p> <ul style="list-style-type: none"> • Microsoft software that is licensed for the number of active physical processors • Utilization of the metric • Details of computers on which the software was installed when the metric quantity was the highest
9.2.8 <code>microsoft_single_physical_processor.csv</code>	<p>Information about:</p> <ul style="list-style-type: none"> • Microsoft software that is licensed for the number of active physical processors • Utilization of the metric • Details of computers on which the software was installed when the metric quantity was the highest
9.2.7 <code>microsoft_physical_core_with_sa.csv</code>	<p>Information about:</p> <ul style="list-style-type: none"> • Microsoft software that has Software Assurance and is licensed for the number of physical cores • Utilization of the metric • Details of computers on which the software was installed when the metric quantity was the highest

Table 68. Files with license metric utilization (continued)

File	Content
9.2.7 microsoft_virtual_- core_with_sa.csv	Information about: <ul style="list-style-type: none"> • Microsoft software that has Software Assurance and is licensed for the number of virtual cores • Utilization of the metric • Details of computers on which the software was installed when the metric quantity was the highest
9.2.8 oracle_processor_- core.csv	Information about: <ul style="list-style-type: none"> • Oracle software that is licensed for the number of physical cores • Utilization of the metric • Details of computers on which the software was installed when the metric quantity was the highest
pvu_full_capacity.csv	Information about: <ul style="list-style-type: none"> • PVU software and its PVU utilization under full capacity license • 9.2.13 FlexPoint offerings to which products and their components are bundled • Details of computers on which the software was installed when the metric quantity was the highest
pvu_sub_capacity.csv	Information about: <ul style="list-style-type: none"> • PVU software and its PVU utilization under subcapacity license • 9.2.13 FlexPoint offerings to which products and their components are bundled • Details of computers on which the software was installed when the metric quantity was the highest
9.2.13 registered_user.csv	Information about: <ul style="list-style-type: none"> • Software that is licensed for the number of registered users • Utilization of the metric • Details of users who used or registered the software when the metric quantity was the highest

Table 68. Files with license metric utilization (continued)

File	Content
<code>rvu_full_capacity.csv</code>	Information about: <ul style="list-style-type: none"> • RVU software and its RVU MAPC utilization under full capacity license • 9.2.13 FlexPoint offerings to which products and their components are bundled • Details of computers on which the software was installed when the metric quantity was the highest
<code>rvu_sub_capacity.csv</code>	Information about: <ul style="list-style-type: none"> • RVU software and its RVU MAPC utilization under subcapacity license • 9.2.13 FlexPoint offerings to which products and their components are bundled • Details of computers on which the software was installed when the metric quantity was the highest
9.2.11 <code>vpc_capacity.csv</code>	Information about: <ul style="list-style-type: none"> • VPC software and its VPC utilization • 9.2.13 FlexPoint offerings to which products and their components are bundled • Details of computers on which the software was installed when the metric quantity was the highest

Snapshot columns

This glossary provides a brief description of every column that can be found in the audit snapshot generated in BigFix Inventory.

[C \(on page dcii\)](#) [D \(on page dciv\)](#) [E \(on page dciv\)](#) [F \(on page dciv\)](#) [I \(on page dciv\)](#) [L \(on page dcv\)](#) [M \(on page dcv\)](#) [O \(on page dcvi\)](#) [P \(on page dcvi\)](#) [R \(on page dcviii\)](#) [S \(on page dcviii\)](#) [U \(on page dcviii\)](#) [V \(on page dcix\)](#)

C

Changed PVU per Core

Specifies whether the number of processor value units assigned to a single core was manually changed for the server that is listed in the Server Name column. If the value is `Yes`, you can check the comment that was provided as the reason of the change. It is listed in the `audit_trails.csv` file, and has the following format.

```
Modification Type: PVU per core value changed
Details: The PVU per core value was changed to 70 on server server_name. The provided
reason: reason.
```

9.2.12 Charged

Indicates whether the relation between the component and the product is charged based on the license terms. If the relation is charged, the component contributes to the license metric utilization of the product to which it is assigned.

Cluster Cores

Total number of processor cores that were available within the cluster of servers when the snapshot was generated. The information is available in the `oracle_processor_core.csv` file.

Cluster Name

Name of the cluster to which the server listed in the Computer column belonged when the snapshot was generated. The information is available in the `oracle_processor_core.csv` file.

Comment

Extra technical information. Possible values include:

- Incomplete virtualization hierarchy
- All instances are excluded

Component

Software component that was discovered on the computer that is listed in the Computer column.

Computer

Host name of the endpoint on which the BigFix client is installed.

Computer Deleted

Specifies whether the computer that is listed in the Computer column existed in the environment when the snapshot was generated.

9.2.11 Computer Last Seen

Date and time when the BigFix client connected to the BigFix server the last time.

Computer Type

Type of the computer. Possible values include:

dciv

- Physical
- Virtual
- Public cloud

Current Server Name

Name of the server on which the computer that is listed in the Computer column resided when the snapshot was generated. Capacity of that server does not influence utilization of license metric. The server whose capacity influences metric utilization is listed in the Server Name column.

D

Details

Information about the type and date of action that is listed in the `audit_trails.csv` file as well as the user who performed the action.

Discovery End

Date and time when the component instance stopped being detected. If the component is still detected, the column is empty.

Discovery Start

Date and time when the component was first detected by BigFix Inventory.

E

Exclusion Comment

Comment that was provided as the reason for excluding the product instance from pricing calculations on the computer that is listed in the Computer column.

F

9.2.13 FlexPoint Bundle

FlexPoint bundle to which the product and its components are assigned.

I

Imported Part Numbers

Part number that was imported to BigFix Inventory. It represents the product that is listed in the Product Name column and the license metric that is assigned to the product. License metric is the same for all products that are listed in a single `csv` file. For example, all products in the `pvu_sub_capacity.csv` file are assigned the PVU subcapacity metric.

IP Address

IP address of the computer that is listed in the Computer column.

Is Imported

Specifies whether a part number that is listed in the Part Number column was also imported to BigFix Inventory by a user.

L

Last Seen

Date and time when the BigFix client connected to the BigFix server the last time.

9.2.11 Starting from application update 9.2.11, the name of the column is changed to Computer Last Seen.

License Use For Cluster

Total number of licenses needed for a cluster, based on number of processor cores, that were available within the cluster of servers when the snapshot was generated and multiplied by Oracle Core Factor.

The information is available in the `oracle_processor_core.csv` file.

9.2.13 License Use For Endpoint

Number of unique users who used the software on a specified computer within the period for which the snapshot was generated. License Use for Endpoint is measured at the metric peak value time.

License Use For Server

Total number of licenses needed for a server, based on number of processor cores, that were available on the server that is listed in the Server Name column when the snapshot was generated and multiplied by Oracle Core Factor. The information is available in the `oracle_processor_core.csv` file.

Logical Processors

Number of processor cores that were assigned to the computer that is listed in the Computer column when the peak of metric value occurred. When hyper-threading is used on the x86 architecture, each core is counted as two logical processors. The information is available in the `microsoft_virtual_core_with_sa.csv` file.

M

Metric

Type of software license metric that is used by the product.

Metric Peak Value Time

Date and time when utilization of the license metric was the highest during the period for which the snapshot was generated.

Metric Quantity

The highest number of metric units that the product used within the period for which the snapshot was generated. Metric quantity is measured at the metric peak value time. It represents metric utilization in the entire environment.

Modification Date

Date and time when the modification that is listed in the `audit_trails.csv` file was performed.

Modification Type

Type of modification.

Modified By

User who modified the item.

O

Oracle Core Factor

Factor that is applied to the server on which Oracle products are installed.

OS

Operating system of the computer that is listed in the Computer column.

P

Part Number

Part number that represents the product and its license metric as listed in the software catalog.

Partition Cores

Number of partition cores that were assigned to the computer that is listed in the Computer column when the peak of metric value occurred.

Path

Installation path of the software that was discovered on the computer or the path where the `.swidtag` file was discovered.

Physical Server CPU Core Full Capacity

Number of physical cores that were available on the server that is listed in the Server Name column when peak of metric value occurred. The information is available in the `pvu_full_capacity.csv` and `rvu_full_capacity.csv` files.

Physical Server CPU Core Subcapacity

Number of physical cores that were available to the product that is listed in the Product Name column when peak of metric value occurred. The information is available in the `pvu_sub_capacity.csv` file.

Physical Server CPU Core Subcapacity Limit

Number of physical cores that were available on the server that is listed in the Server Name column when peak of metric value occurred. The information is available in the following files:

- `pvu_sub_capacity.csv`
- `rvu_sub_capacity.csv`
- `9.2.11 vpc_capacity.csv`

Physical Server PVU Full Capacity

Number of cores (Physical Server CPU Core Full Capacity column) multiplied by the number of PVUs assigned to the processor core (PVU per Core column). It represents PVU full capacity for the server that is listed in the Server Name column. The information is available in the `pvu_full_capacity.csv` file.

Physical Server PVU Subcapacity

Number of cores (Physical Server CPU Core Subcapacity column) multiplied by the number of PVUs assigned to the processor core (PVU per Core column). It represents PVU subcapacity for the product that is listed in the Product Name column. The information is available in the `pvu_sub_capacity.csv` file.

Physical Server PVU Subcapacity Limit

Number of cores (Physical Server CPU Core Subcapacity Limit column) multiplied by the number of PVUs assigned to the processor core (PVU per Core column). It represents PVU full capacity for the server that is listed in the Server Name column. The information is available in the `pvu_sub_capacity.csv` file.

Processor

Name of the processor normalized to the name that is listed in the PVU table. If the discovered processor was not matched with any processor from the PVU table, the value is `Unknown processor`.

Processor Brand String

Full specification of the processor including its brand, model and speed as read from the operating system of the computer that is listed in the Computer column.

Processor Metric

Number of processors that contributed to calculating license metric utilization. The information is available in the `microsoft_single_physical_processor.csv` and `microsoft_dual_physical_processor.csv` files.

Product Name

Name of the software product.

Product Release

Release of the software product.

Publisher

Name of the company that published the software product.

PVU per Core

Number of processor value units that are assigned to a processor core on a server that is listed in the Server Name column. By default, the value is taken from the PVU table but it can also be set manually. If it was set manually, the value in the column Changed PVU per Core is `yes`.

R

Recalculation Needed

Specifies whether recalculation was needed when the snapshot was generated. If the value is `yes`, metric quantity for the product might not be correct. It happens because actions that influence metric quantity were performed but the value was not recalculated. For example, a component was reassigned to a different product. A correct snapshot contains only entries for which the value in the Recalculation Needed column is `no`.

S

Server Active Sockets

Number of physical processor chips that are mounted on the server that is listed in the Server Name column. The number also includes empty sockets without processors.

Server Cores

Number of processor cores that were available on the server that is listed in the Server Name column when peak of metric value occurred. The information is available in the `microsoft_physical_core_with_sa.csv` and `oracle_processor_core.csv` files.

Server Name

A unique system in the enterprise. The value depends on what type of computer is listed in the Computer column.

- If the computer is a physical machine, the Server Name column shows the hardware manufacturer, type, and a machine serial number.
- If the computer is a virtual machine, the Server Name column shows the manufacturer and host name of the physical host.
- If the computer is a virtual machine for which a VM manager is not configured, the Server Name column shows a UUID assigned by BigFix Inventory prefixed with `TLM_VM_`. A correct snapshot should not contain any entries with such a value unless you have the permission of the IBM Compliance Team.

If multiple virtual machines are running on the same physical server, the same server name is listed multiple times.

U

Unconfirmed Product Instance

Specifies whether the assignment of the component to the product was confirmed. Confirming the assignment is optional and does not impact metric utilization.

V

Virtualization Layer ID

Internal ID of the virtualization layer on which the software component is installed. The information is used by the IBM Support for troubleshooting purposes.

Understanding CSV files with metric utilization

Due to differences in the structure of the user interface reports and CSV reports, information about metric utilization is presented differently in each type of report. Learn more about these differences to better understand how to correlate the information between user interface reports and CSV reports.

Multilevel versus flat report structure

The All Metrics and IBM PVU subcapacity reports consist of the following levels for PVU and RVU MAPC products:

1. First level - a list of all products.

Product Name	Metric	Metric Quantit...	Metric Quantity History
			03/18/2016 - 06/17/2016
IBM MQ	PVU Full Capacity	1360	
IBM MQ	PVU Subcapacity	810	

2. Second level - when you click a product name, you go to information about the physical server on which the product is installed.

Server Name	Processor	CPU Core Subcapacit...	CPU Core Subcapacit...	PVU Subcapacity Lim...	PVU Subcapacity	Comment
X.XXX.XX.242	Intel(R) Xeon(R) Multi-c...	8	6	800	600	
X.XXX.XX.217	Intel(R) Xeon(R) Multi-c...	8	3	560	210	

3. Third level - when you click the server name, you go to the list of virtual machines that run on this server and have the product installed.

Computer Name	Operating System	IP Address	Product Version	Component Name	Path	Last Seen	Comment
Computer 1	Linux Red Hat Ente...	X.XXX.XX.151	7.5	IBM WebSphere M...	/utitle/MQ/	06/15/2016	
Computer 2	Win10 10.0.10240	X.XXX.XX.153	7.5	IBM WebSphere M...	C:\Uttitle\CompleteS...	06/15/2016	
Computer 3	Win10 10.0.10240	X.XXX.XX.154	7.5	IBM WebSphere M...	C:\Uttitle\CompleteS...	06/16/2016	

The audit snapshot is not multilevel. However, it must present the same information that the report that is available in the user interface. Below is an excerpt from the audit snapshot that corresponds to the PVU report presented on the screen captures above. Some columns are hidden in the snapshot to simplify it.

Table 69. Excerpt of the audit snapshot

D	E	F	H	K	M	N	O	P	S
Product Name	Metric Quantity	Metric Peak Value Time	Server Name	PVU Per Core	Physical Server CPU Core Subcapacity Limit	Physical Server CPU Core Subcapacity	Physical Server PVU Subcapacity Limit	Physical Server PVU Subcapacity	Computer
IBM MQ	810	2016-06-01 00:00:00.000	X.XXX.XX.217	70	8	3	560	210	Computer 1
IBM MQ	810	2016-06-01 00:00:00.000	X.XXX.XX.217	70	8	3	560	210	Computer 2
IBM MQ	810	2016-06-01 00:00:00.000	X.XXX.XX.217	70	8	3	560	210	Computer 3

When values in column O (Physical Server PVU Subcapacity Limit) are added, PVU consumption equals 1680 PVUs, which is greater than the value that is presented on the user interface. These values should not be added because they all refer to a single server (X.XXX.XX.217) on which the three VMs run, not an individual VM that is listed in column S (Computer). The value in column M is repeated due to the flat structure of the CSV report.

If you want to relate the values from the CSV report to the user interface report, the relation is as follows:

- Column O in the CSV report relates to the PVU Subcapacity Limit column on the second level of the report
- Column N in the CSV report relates to the CPU Core Subcapacity column on the second level of the report

Server Name	Processor	CPU Core Subcapacity Limit	CPU Core Subcapacity	PVU Subcapacity Limit	PVU Subcapacity	Comment
X.XXX.XX.242	Intel(R) Xeon(R) Multi-c...		8	6	800	600
X.XXX.XX.217	Intel(R) Xeon(R) Multi-c...		8	3	560	210

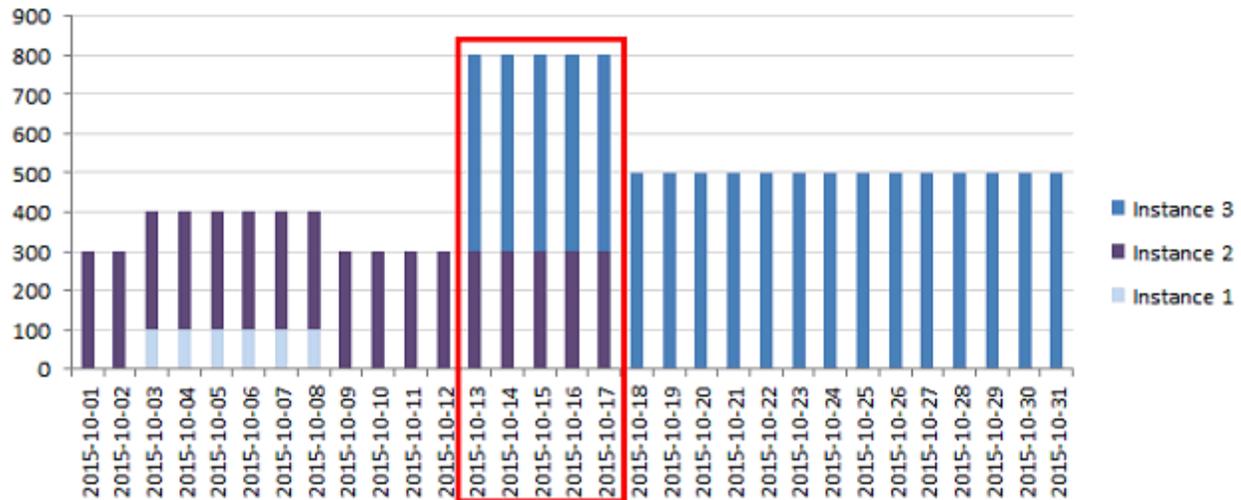
PVU and RVU MAPC subcapacity license terms specify license consumption only on the physical server level. The values cannot be calculated and presented on the virtual machine or LPAR level, and such data is not provided by BigFix Inventory. For more information, see: [Virtualization Capacity License Counting Rules](#).

CSV reports contain only software instances discovered during the license peak value time

CSV reports list only software instances that were discovered during the license peak value time. If an instance was installed in your environment during the reporting period but was not discovered during the license peak value time, the instance is not listed in the CSV reports.

For example, three instances of DB2 are installed. The reporting period is set to one month and DB2 license utilization is as follows:

- Instance 1 uses the maximum of 100 PVUs
- Instance 2 uses the maximum of 300 PVUs
- Instance 3 uses the maximum of 500 PVUs



The license peak value occurs when the sum of PVU values for all instances is the highest and equals 800 PVUs (300 PVUs for instance 2 + 500 PVUs for instance 3). The peak occurs during the period that is marked with a red square. It is the license peak value time. The peak usage for instance 1 is outside of the peak value time. It means that this instance was installed in the environment during the current reporting period but was not discovered during the license peak value time. This instance is not listed in the CSV report.

Raw utilization of license metrics

Available from 9.2.1. Raw data about utilization of license metrics is collected for products that deliver .slmtag files and is shown on the Resource Utilization report. The data that is shown on the report is not aggregated. It means that if a product can use five license metrics, five rows are shown on the report for that product. What is more, all metrics that are available for a product are displayed, not only metrics that are defined by your license agreement.

To collect raw utilization data, run the resource utilization scan. The scan is scheduled automatically if you enabled the default scan configuration. Otherwise, enable it manually. For more information, see: [Initiating software scans \(on page cciii\)](#).

For the list of IBM products that deliver .slmtag files, see: [Resource Utilization Measurement Community](#).

The screenshot shows the 'Resource Utilization' report interface. At the top, there is a 'Send Feedback' button. Below it, a message states: 'Retention period for historical data is configured for 90 days.' This message is highlighted with a red box and labeled '1'. Below the message is a dropdown menu set to '(Base Report)' and an 'Export History' button, which is also highlighted with a red box and labeled '2'. The main part of the interface is a table with columns: Computer Name, Software Name, Metric Type, Metric Subtype, First Measured, Last Measured, Maximal Trend Value, and Maximal Trend Value History. The table contains four rows of data. The first row shows 'NC042222' with 'IBM ILOG CPLE...' and 'FLOATING_USER' metric, showing '0' for the maximal trend value. The second row shows 'VMW009128109...' with 'IBM MobileFirst ...' and 'CLIENT_DEVICE' metric, showing '27'. The third row shows 'NC042222' with 'IBM Endpoint Ma...' and 'MANAGED_VIRT...' metric, showing '0'. The fourth row shows 'VMW009128109...' with 'IBM MobileFirst ...' and 'CLIENT_DEVICE' metric, showing '97'. Each row has a corresponding line graph in the 'Maximal Trend Value History' column. Red arrows point from the numbered callouts to these elements: '3' points to the 'Metric Type' and 'Metric Subtype' columns, '4' points to the 'First Measured' and 'Last Measured' columns, '5' points to the 'Maximal Trend Value' column, and '6' points to the 'Maximal Trend Value History' column.

1 Retention period setting

To reduce the amount of resource utilization data, set the retention period after which the data is removed from the database. For more information, see: [Configuring data retention period for raw utilization data \(on page ccxcviii\)](#).

2 Export History

History provides initially processed data about the average and maximum utilization of each license metric for each product for a specified period. It is different from export to CSV or PDF that includes only data that is visible on the user interface.

3 Metric Type and Metric Subtype

The type and subtype of license metrics that a product uses. The report does not show products that use the PVU, RVU MAPC or Install Seats metric. To view these products, open the All Metrics report.

4 First Measured and Last Measured

The first and last time when each metric was reported.

5 Maximal Trend Value

The maximum utilization of each license metric over the last 30 days. If a particular product is unused, no resource utilization data is available, or the data is older than 30 days, the value is 0.

6 Maximal Trend Value History

Trend value of the maximum utilization over the last 30 days. When the graph shows 0 for a specific period, it means that there is no data, or there was no license utilization during this period. To recognize which is the case, click **Export History** to save the report to a file. Then, open the file, and check if the entries that show 0 in the user interface also show 0 in the file. If they do, it means that the license utilization during this period equals 0. If there are no such entries, the data from this period was not reported.

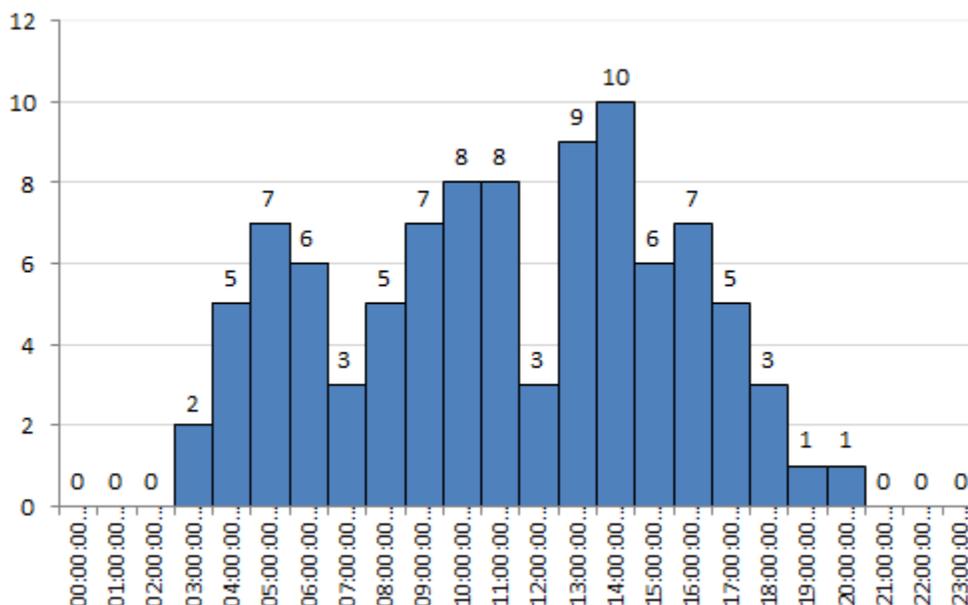
Aggregating the data

If you want to aggregate the data according to your license agreement, export the history of license metric utilization and aggregate it outside of BigFix Inventory. Apply all necessary filters on the Resource Utilization report, and click **Export History**. The `csvhistory.zip` file is downloaded to your computer. It contains the `export_history.csv` file that provides initially processed data about the average and maximum utilization of each license metric for each product.

Example

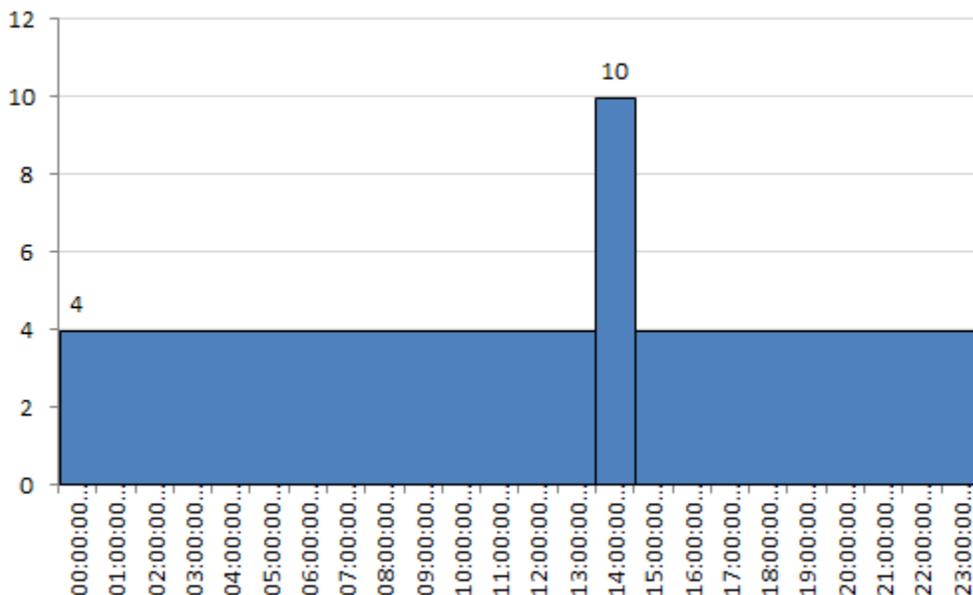
IBM Product is an application that is licensed based on the number of users who simultaneously log in to that application per hour. [Figure 6: License metric utilization data retrieved from the .slmtag file for BigFix Product \(on page dcxiii\)](#) is a visualization of the license metric utilization that is retrieved from the .slmtag file specific for BigFix Product.

Figure 6. License metric utilization data retrieved from the .slmtag file for BigFix Product



The data undergoes initial processing during which the average daily utilization of the license metric is calculated. Information about the maximum utilization, including its value and time, is preserved. [Figure 7: License metric utilization data for BigFix Product after initial processing \(on page dcxiv\)](#) is a visualization of the license metric utilization data for BigFix Product after initial processing. It shows that on average, four users were simultaneously logged in to the application. The maximum utilization occurred between 2:00 PM and 3:00 PM when 10 users were simultaneously logged in to BigFix Product.

Figure 7. License metric utilization data for BigFix Product after initial processing



The `export_history.csv` file contains license metric utilization data after initial processing. It does not contain all data that was collected from the endpoints in your infrastructure. [Table 70: A fragment of the `export_history.csv` file \(on page dcxiv\)](#) shows a fragment of the `export_history.csv` file that corresponds to the data that is presented in [Figure 7: License metric utilization data for BigFix Product after initial processing \(on page dcxiv\)](#). The first row corresponds to the average license metric utilization before the peak. The second row corresponds to the peak value. The third row corresponds to the average utilization after the peak value.

Table 70. A fragment of the `export_history.csv` file

A simple table with 4 rows and 7 columns.

Computer Name	Software Name	Metric Type	Metric Subtype	Start Time (UTC 0)	End Time (UTC 0)	Value
VMW00912	IBM Product	CONCURRENT_	USER	2015-06-20 00:00:00.0	2015-06-20 14:00:00.0	4
VMW00912	IBM Product	CONCURRENT_	USER	2015-06-20 14:00:00.0	2015-06-20 15:00:00.0	10
VMW00912	IBM Product	CONCURRENT_	USER	2015-06-20 15:00:00.0	2015-06-21 00:00:00.0	4

Disabling collection of raw utilization data

Available from 9.2.1. Raw utilization data is collected by default. However, you can disable the option if you are not interested in this data.

 You must be an Administrator to perform this task.

1. Stop all actions that were created by the **Initiate Software Scan** task.
 - a. Log in to the BigFix console.
 - b. In the left navigation tree, click **Actions**.
 - c. Select all open actions that were created by the **Initiate Software Scan** task, right-click them, and click **Stop Action**.
The status of the actions changes to *Stopped*.
2. [Run the software scan \(on page cciii\)](#). Clear the **Resource utilization** check box.
3. Log in to BigFix Inventory, go to <https://hostname:port/management/feature>, and clear the **Measure resource utilization** check box. Then, click **Save**.

Raw utilization data is no longer collected from the computers in your infrastructure. The Resource Utilization report is hidden from the BigFix Inventory menu.

Software license optimization

BigFix Inventory helps you optimize your software licenses. You can analyze the software inventory in your environment and maximize the utilization of your assets. Learn how to minimize the costs of acquiring and maintaining licenses and track license compliance for the software products that are installed in your environment. You can verify the needs and requirements of your environment and then apply the most appropriate license model.

Monitoring and reporting software usage

Software usage statistics contain details about the usage of specific software that is installed in your environment, such as how many times it was used and when it was used for the last time. Based on this information, you can quickly identify used and unused software. This analysis can help you reduce the license costs when planning license renewals or upgrades.

BigFix Inventory collects information about the running processes and, based on that, measures software usage for all software instances that are discovered with usage signatures.



Important: **9.2.15** Starting from application update 9.2.15, BigFix Inventory introduces a new method of reporting usage for Java. Paths to Java processes that run on a computer are matched with paths in which the software is discovered. Based on this information, BigFix Inventory can determine exactly which Java instance is used. It is possible to distinguish between usage of BigFix and Oracle Java that is installed on a single computer.

Software catalog in version 9.2.15 does not contain usage signatures for Java any more. To report usage of Java, ensure that the BigFix client and BigFix server are in version 9.5.10 or higher, and that you upgrade the BigFix Inventory server to version 9.2.15.

Monitoring and reporting of software usage is enabled by default if during the initial configuration you choose the default scan settings. For more information, see: [Default scan configuration \(on page cxcvi\)](#). Otherwise, follow the instruction for manual scan configuration to enable the collecting of software usage statistics.

1. Log in to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Analyses**.
3. To measure software usage, activate the **Application Usage Statistics** analysis. Right-click the analysis, and click **Activate**.
4. Run the **Initiate Software Scan** fixlet. For more information, see: [Initiating software scans \(on page cciii\)](#).
5. Run the **Upload Software Scan Results** fixlet. For more information, see: [Uploading software scan results \(on page ccviii\)](#).
6. Wait for the scheduled import or run it manually.

The results of the analysis are reflected on the following reports:

- Metering Data shows the aggregated usage statistics from all instances that are discovered on each endpoint based on given binary file.
- Software Classification panel lists the discovered software components with the details about usage statistics.
- Usage per Computer panel provides an overview of the software usage statistics that are aggregated for all instances on product level and presented in the report for each computer.

For more information about the reports, see: [Available reports \(on page dcxvi\)](#).



Tip: To view information about usage on the reports, configure the report view to display the columns that collect software usage statistics. To do that, hover over the **Manage Report View** icon , select **Configure View** and choose the applicable columns.

If you are in the initial deployment phase, or if you do not need information about software usage, you can disable the collection of this information to improve the BigFix Inventory performance and shorten the import time. For more information, see: [Disabling the collection of software usage \(on page cclxviii\)](#).

Reports

Reports provide detailed information about the computers in your infrastructure and the software items that are installed on these computers. You can decide what type of information you want to display by choosing the appropriate type of report. You can also customize the type and amount of information that is displayed in a report and save the report settings to reuse them.

Available reports

Reports are the primary source of information about computers in your infrastructure and software that is installed on these computers.

The inventory data is collected from all computers in your infrastructure. When you decommission a computer it continues to report to BigFix Inventory. For report accuracy, ensure that all infrastructure changes are reflected in BigFix Inventory. For more information, see: [Removing inactive computers \(on page cdxxxvii\)](#).



Note: The data that is displayed on the reports is limited to the computer group to which the user belongs.

General

Report	Description
Saved Reports	<p>The report provides information about saved reports, for example:</p> <ul style="list-style-type: none"> • The base report from which the saved report was created. • The user who created it. Some of the reports are delivered with BigFix Inventory. They are marked as created by the SYSTEM. Other reports are created by individual users. • Date and time when the saved report will be sent to the recipients that you specified in the schedule. <p>For more information, see: Creating saved report views (on page dclliii).</p>

Software Inventory

Report	Description
Software Summary (Preview)	<p> Required permission: View Endpoints</p> <p>9.2.15 The report provides an overview of software inventory in your environment. Each row in the report represents a set of software instances that are grouped by the selected columns. You can limit the instances in the groups by using the report filters. By default, you can see only the instances that are currently installed and not suppressed. To change the way how the data is displayed, select different report columns.</p> <p>For example, you can filter the group by selecting a publisher. In the column Count you get the number of software instances only from the selected publisher.</p> <p> Note: You cannot filter the report by selecting columns that aggregate the values, for example: Count, Used and Unused.</p> <p>Drill down</p> <p>You can drill down the Software Summary report to view the Software Classification panel with the detailed information about all software instances from the related group, including component name, component version, and installation path.</p>

Report	Description
	<p data-bbox="404 264 1422 422">9.2.16 Starting from application update 9.2.16, you can additionally use drill down to see the detailed information about the software instances that contribute to software usage through the related columns, for example, Used, Unused, Used in the Last Month and other summary columns.</p> <p data-bbox="404 453 753 485">Summary of the usage statistics</p> <p data-bbox="404 516 532 548">9.2.16</p> <p data-bbox="404 590 1422 779">This report also provides software usage statistics that contain details about the usage of specific software on the listed instances. You can check how many instances of a product within the group were used and how many were not. You can see an overview of software usage statistics from the last month, the last quarter, and the last year. You can track the number of software instances within the group where software usage is monitored and where it is not.</p> <p data-bbox="404 810 526 842">Limitations</p> <ul data-bbox="469 884 1390 1041" style="list-style-type: none"> • Software Summary (Preview) report does not work on MS SQL 2008 and 2008 R2. The required MS SQL Server version is MS SQL 2012 and higher. • You might experience performance issues on SQL Server, especially when you select Product Name column or usage related columns, for example Used, for grouping. <p data-bbox="404 1083 1338 1115">In case of problems with opening links in Internet Explorer, see: User interface problems.</p>
Inventory Exploration	<p data-bbox="404 1167 837 1209"> Required permission: View Endpoints</p> <p data-bbox="404 1251 1414 1398">The report provides information about software items that are installed in your infrastructure. The items are structured into hierarchies that consist of the publisher, software titles, their versions, releases, and components. You can drill down through the hierarchy to gather detailed information about software usage including its time and frequency.</p> <p data-bbox="404 1440 1422 1629">For each product, its version and release, you can also view information about contracts to which this product is assigned. The report shows the number of software instances that are covered by contracts and the cost of acquisition and maintenance of the licenses for that product. You can use the report to efficiently manage your software inventory by monitoring the cost of software licenses and the use of the particular software products.</p>
Software Installations	<p data-bbox="404 1682 837 1724"> Required permission: View Endpoints</p> <p data-bbox="404 1755 1414 1875">9.2.15 The Software Installations report has been changed. The old report is based on the old data model that does not take shared disks and multiple component instances into account. Now, it is deprecated, moved to the bottom of the Reports menu and will be entirely removed</p>

Report	Description
	<p>in the future. The old report is substituted with a new report view that is based on the new data model and consistent with other reports. Thanks to the new data model, the report criteria are extended to cover additional component instances and the reported number of installed software might be affected. The new report is based on the Software Classification panel. For more information, see: Software Classification (on page dcxix).</p> <p>The report provides information about computers in your infrastructure and software items that are installed on these computers. You can see a detailed report for a single software item, when you click the link in the Details column.</p> <p>The detailed view shows information about software signatures that were tried to detect the software or software identification tags that were discovered. The information is accompanied by the following icons:</p> <ul style="list-style-type: none"> •  - The signature matched the software item, or the software ID tag was discovered on the computer and caused software detection. •  - The signature did not match the software item. <p>If you believe that software was erroneously detected, you can use this information to examine how the signature or the software ID tag was evaluated. You can also view the installation path of software that was discovered by a file-based signature. If software was discovered by a complex signature, you can view the hierarchical structure of the signature.</p> <p>In this view, you can also see information about software usage. When a usage signature exists and matches a process listed on the Metering Data report, statistics about usage time and frequency are shown. If the information cannot be gathered for the particular software, one of the following messages is displayed:</p> <ul style="list-style-type: none"> • <code>Usage data is not supported</code> - Usage data cannot be gathered because the usage signature for this software item does not exist. • <code>Usage data was not found</code> - At least one usage signature exists for the particular software but the usage data was not gathered. Thus, report columns with statistics about usage time and frequency show 0. Usage data might not be found because the software was never used and thus has no usage information. It might also happen because the option to run the Application Usage Statistics scan was cleared in the Initiate Software Scan fixlet.
Software Classification	<p> Required permission: View Endpoints or Manage Software Classification</p> <p>The report provides an overview of your software inventory and the relations between the discovered software components and licensable products. It provides a flat structure with software installations broken into single components, which makes it easier to view your software as-</p>

Report	Description
	<p>sets and manage relations between them. The Software Classification panel allows streamline management of your software inventory. To make your metric calculation accurate, assign each component to a product, exclude, or suppress the instances that should not be included in the calculation, and confirm your actions. After you classify all your components, and confirm all actions, software classification is complete.</p> <p>9.2.15 Starting from application update 9.2.15, BigFix Inventory collects detailed information about the usage of Java processes on each instance separately. The usage statistics for each instance is available on Software Classification report.</p> <p>The report provides the following predefined views:</p> <ul style="list-style-type: none"> • BigFix Software Only that shows only BigFix software that is currently installed. • First Detected Last Week - BigFix Software Only shows the BigFix software that was detected within the last week and is currently installed. It helps you to review the recently installed software and make an appropriate action, such as changing the component assignment. • Software Installations shows all components that are installed in your environment. • Current Installations shows software that is currently installed in your environment and its classification. • 9.2.13 Out of Support Software shows installations for which support already ended. Starting from application 9.2.14, the report additionally shows the installations of old software components that do not have the end of support date set but were created more than 6 years ago. • 9.2.13 Vulnerable Software (Preview) shows software installations for which at least one CVE was matched. • 9.2.13 Software Approaching End of Support shows software installations for which support ends within the next three months. Starting from application update 9.2.14, the report additionally shows installations of the components that do not have the end of support set but were created between 5 years 9 months and 6 years ago. These installations will be reported as out of support within the next three months. • 9.2.13 Vulnerable Software - Recent (Preview) shows software installations that were discovered within the last two weeks and for which at least one CVE was matched. • 9.2.14 Software Installations - Recent shows new software installations that were discovered within the last two weeks. • 9.2.14 IBM Software - Pending Classification shows the IBM software components that are not yet classified. • 9.2.14 Blacklist Software shows the discovered software components that are marked with a <code>blacklist</code> tag that represents unwanted software. For more information, see: Adding and removing tags (on page dcli).

Report	Description
<p data-bbox="196 268 323 296">9.2.11</p> <p data-bbox="196 310 383 380">Package Summary</p>	<p data-bbox="402 289 831 327"> Required permission: View Raw Data</p> <p data-bbox="402 363 1421 554">The report lists and counts the packages that are installed on the computers in your infrastructure. The report is calculated on demand only. You can find out what is the name of every package, its version, vendor, description, and type. You can use this information to create custom software signatures. For more information, see: Creating signatures from package data (on page cdxci).</p> <p data-bbox="402 590 1421 701">To refresh the report, open it, and click Calculate. The data is calculated immediately, unless the import is in progress. In this case, the report is calculated after the import completes. If new data arrives, information that the report is not up-to-date is displayed.</p> <p data-bbox="402 737 1071 764">The report might not be up-to-date it in the following situations:</p> <ul data-bbox="467 810 1390 919" style="list-style-type: none"> • The import of data was completed and the new package data might be available. • A new package signature was created and it might influence the Package Data report. • Package signatures were changed or deleted on the Catalog Customizations panel.
<p data-bbox="196 947 345 974">Package Data</p>	<p data-bbox="402 968 831 1005"> Required permission: View Raw Data</p> <p data-bbox="402 1041 1406 1232">The report provides information about packages that are installed on each computer in your infrastructure. The information is retrieved from the Windows registry and UNIX RPM database. You can find out what is the name of every package, its version, vendor, description, and type. You can use this information to create custom software signatures. For more information, see: Creating signatures from package data (on page cdxcii).</p>
<p data-bbox="196 1272 386 1299">Scanned File Data</p>	<p data-bbox="402 1293 831 1331"> Required permission: View Raw Data</p> <p data-bbox="402 1367 1414 1600">The report provides information about files that were found on the computers in your infrastructure. You can find out what is the location, size, and version of each file. The files that are displayed on the report are detected basing on the discovery of particular file extensions, or the processes that run on the computers in your environment. You can use this information to create custom software signatures. For more information, see: Creating signatures from scanned file data (on page cdxci).</p> <p data-bbox="402 1635 946 1663">The report provides the following predefined views:</p> <ul data-bbox="467 1709 1406 1864" style="list-style-type: none"> • 9.2.14 Malware Threats (Preview) lists files that are found in the environment by BigFix Inventory and classified as malware threats. • 9.2.14 Malware Threats - Recent (Preview) lists files that are found in the environment BigFix Inventory within the last two weeks and classified as malware threats.

Report	Description
	<p>The tool and the fixlets related to getting file reputation are not available anymore as this preview feature has been withdrawn. However, the historical data is still displayed, and the predefined views are still available.</p>
<p>Unrecognized Files</p>	<p> Required permission: View Raw Data</p> <p>The report provides a ranking of 1000 files that are most commonly encountered in your computer infrastructure but do not produce matches for any signature. The ranking is separate for each computer group. It is based on the number of computers on which the files were detected, the number of copies of a particular file, and the alphabetical order. The files are listed by name with no regard to their size or version. You can use this report to investigate unrecognized files that can be used to create custom signatures.</p> <p>Because aggregation of unrecognized scan data slows down imports, particularly in large deployments, it is disabled by default. To enable it, perform the following steps:</p> <ul style="list-style-type: none"> • For versions up to 9.2.8, go to Management > Scanned File Data, and choose the frequency of the aggregation. • Starting from version 9.2.9, go to Management > Advanced Server Settings, and change the value of the recompute_unrecognized_scan_data parameter to specify frequency of the aggregation.
<p>Metering Data</p>	<p> Required permission: View Raw Data</p> <p>The report provides information about processes that run on the computers in your infrastructure. You can track when the process was used for the first and the last time as well as what is the frequency and the total time of its usage. You can also see what is the average run time of each process and the average number of its runs per day. This report provides aggregated usage statistics from all instances that are discovered on each endpoint based on given binary file. In this report you cannot track individual installations on one endpoint.</p> <p>9.2.15 Starting from application update 9.2.15, BigFix Inventory collects detailed information about the usage of Java processes and provides the aggregated usage statistics on Metering Data report.</p>
<p>9.2.10 Usage per Computer</p>	<p> Required permission: View License Metrics</p> <p>The report provides information about the frequency and time of using software products that are installed on the computers in your infrastructure. For every product you can see what was the frequency and the total time of its use per computer as well as when it was used for the first</p>

Report	Description
	<p>and last time. You can also see what is the average run time of each product and the average number of its runs per day.</p> <p>By default, the report is calculated after every import. To change this behavior so that the report is calculated on demand, go to Management > Advanced Server Settings, and change the value of the calculate_during_import/product_usage parameter.</p> <p>After you change the parameter, the report must be refreshed on demand. To refresh the report, open it, and click Calculate. The data is calculated immediately, unless the import is in progress. In this case, the report is calculated after the import completes. If new data arrives, information about the need of recalculating the data is displayed on the report.</p> <p>9.2.15 Starting from application update 9.2.15, BigFix Inventory collects detailed information about the usage of Java processes on each instance separately. The aggregated usage for all instances of the same product is available on Usage per Computer report.</p>

Business Applications

Report	Description
Oracle Databases	<p> Required permission: View Endpoints</p> <p>The report provides information about the edition (Standard or Enterprise), options, and Management Packs (including Oracle Real Application Cluster, RAC) of Oracle Databases that are deployed in your environment. When you combine this information with the collected hardware inventory data, you can use it to calculate the demand for Oracle Database licenses. For more information, see: Extended discovery of Oracle Database (on page dii).</p>

License Metrics

Report	Description
All Metrics	<p> Required permission: View License Metrics</p> <p>The report provides information about the utilization of license metrics for products that are installed on the computers in your infrastructure. For every product, you can see the maximum metric utilization and its history over the specified period. You can also specify a metric threshold to verify whether metric utilization is above or below your expectations. You can create a snapshot of the report, and store it for audit purposes. For more information, see: License metric utilization (on page dlxxxvii).</p>

Report	Description
	<p> Important: If the report does not show products with a particular license metric, calculation of this metric might be disabled for the computer group to which you are assigned. For information about enabling the calculations, see: Setting up computer groups (on page cxc).</p> <p>9.2.16 Starting from application update 9.2.16, the report provides information about the top level of metric utilization, which means that stand-alone products and flexpoint bundles are displayed. You can also see products that contribute to a FlexPoint bundle by applying the appropriate report filter.</p> <p>9.2.4 The report provides the following predefined views:</p> <ul style="list-style-type: none"> • IBM PVU Subcapacity that shows only products that are licensed for PVU subcapacity. • 9.2.13 IBM FlexPoints that shows FlexPoint bundles. • All IBM Metrics that shows only BigFix products. • 9.2.5 Install Seats Metric that shows only products that are licensed based on the number of computers on which the product is installed. • Metric Thresholds that includes a column in which you can specify a metric threshold. The threshold indicates the maximum number of metric units that a product can use within a computer group. It is then used to calculate whether metric utilization is above or below your expectations. For more information, see: Setting license metric threshold (on page dxc). • Exceeded Thresholds that shows only products for which metric utilization is above the specified threshold. • 9.2.7 Microsoft Metrics (previously Microsoft Core Metric) that shows Microsoft products for which metrics are calculated in BigFix Inventory. • 9.2.8 Oracle Metrics that shows Oracle products for which metrics are calculated in BigFix Inventory.
Audit Trail	<p> Required permission: View Audit Trail</p> <p>The report provides information about actions that have influence on the information that is displayed on other reports. You can view details of each action, its type, date, and the user who performed it.</p> <p> Important: Audit trail shows actions that are performed only within the computer group to which the user who views the report has access.</p> <p>The actions that are shown on the report include:</p>

Report	Description
	<ul style="list-style-type: none"> • Modifications to the IBM software classification: confirming and changing the default classification, sharing an instance between products, including an instance in and excluding it from pricing calculations <p> Note: Audit trail does not contain information about software classification when either component, product or both are non-IBM.</p> <ul style="list-style-type: none"> • Uploads: successful or failed upload of the software catalog, charge unit data, and PVU table • Imports: successful or failed import of the software catalog, charge unit data, and PVU table • VM managers: adding, modifying, and deleting VM managers • Updates to the scan configuration on the Scan Configurations panel • Changes of server parameters and advanced server settings • Changes of the PVU per core value
9.2.1 Re-source Utilization	<p> Required permission: View License Metrics</p> <p>The report provides information about the utilization of license metrics for products that deliver .slmtags. The report shows the maximum metric utilization and its trend value over the last 30 days. For every product, you can see the type and subtype of license metrics that a product uses as well as the first and last time when each metric was reported. You can also see the path to the .slmtag file from which the data was retrieved.</p> <p>The report shows only products that deliver resource utilization data. It does not show products for which metric utilization is calculated and shown on the All Metrics report. For more information, see: Raw utilization of license metrics (on page dcxi).</p>

Users

Report	Description
9.2.13 Software Users	<p> Required permission: View User Details</p> <p>Scope of the report</p> <ul style="list-style-type: none"> • 9.2.13 Microsoft Office 365 • 9.2.14 Windows The following Adobe Creative Cloud (CC) applications that are installed on Windows computers. <ul style="list-style-type: none"> ◦ Adobe Illustrator CC ◦ Adobe Photoshop CC

Report	Description
	<ul style="list-style-type: none"> ◦ Adobe InDesign CC ◦ Adobe Premiere Pro CC ◦ Adobe After Effects CC ◦ Adobe Dreamweaver CC ◦ Adobe Muse CC <p>Overview</p> <p>The report provides information about the current users of software and their details. Users are matched with the instance of the product to which they are registered with the given ID. By clicking the product name, you are redirected to the Software Classification panel.</p> <p>The report shows a list of users who use or registered Microsoft Office 365 and a set of Adobe CC components on the computers in your environment. They have a unique user ID, that might be an e-mail address, or an alternative user ID, which is an identifier that is specific to the particular software.</p> <p>User details</p> <p>The user information that is collected might include the personal data. By default, the user personal data, such as an e-mail address, is obfuscated with SHA-1 and is not shown on the report or stored in the database. If you would like to collect and display it, select the Store and show user names on reports when initiating software scans. For more information, see: Initiating software scans (on page cciii).</p> <p>If you no longer need to view information about software users on the Software Users report, go to Management > Advanced Server Settings, and change the value of the remove_user_information parameter to True. After the next import of data, the user information is removed and the report is emptied. To repopulate the report, run the Initiate Software Scan fixlet. Removing user information also removes all historical data related to the Registered User metric on the All Metrics report.</p> <p> Important: BigFix Inventory is GDPR compliant. By default, BigFix Inventory does not collect any personal data about the users. Gathering personal data is optional and should be enabled consciously. Personal data is stored and used by BigFix Inventory only. It is not passed or used outside BigFix Inventory.</p>

Infrastructure

Report	Description
Computers	 Required permission: View Endpoints

Report	Description
	<p>The report provides information about computers in your infrastructure, such as their operating systems, IP addresses, and other information that can be used to identify them. Additionally, you can display information about the computer health, for example:</p> <ul style="list-style-type: none"> • When the computer last connected to BigFix • What version of the client is installed on the computer • What version of the catalog is available on the computer • When was the last attempt to run the scan • Whether each type of the software scan was successful • Whether the computer misses scan results • Whether it is running out of disk space • Whether it misses some prerequisites <p>When you click a link in the Name column, detailed information about a particular computer is displayed. You can now view reports narrowed down to this computer. To do this, select the report from the upper left corner.</p> <p>9.2.5 IBM i You can decommission computers that are imported from the disconnected data source (on page cclxxi) to remove them from the report. For more information, see: Decommissioning computers with disconnected scanners (on page cclxxx).</p> <p>The report provides the following predefined views:</p> <ul style="list-style-type: none"> • 9.2.4 Outdated Agents shows computers on which the BigFix client is not updated to the latest available version. <p> Important: For some operating systems, support is withdrawn with newer versions of the BigFix client and only earlier versions are supported. Such a computer is reported as having an outdated version of the client even though a newer version is not available.</p> <ul style="list-style-type: none"> • 9.2.14 Computers lists all computers in your infrastructure. • 9.2.14 Computers - Recent shows the computers that were added to BigFix Inventory within the last two weeks. • 9.2.14 Computers without Scan (Preview) shows computers that are visible in BigFix Inventory for more than a week, and on which the scan was not initiated. • 9.2.14 Computers without Recent Scan (Preview) shows computers with scan that is older than two weeks.
Computer Groups	 Required permission: Manage Computer Groups

Report	Description
	<p>The report provides information about computer groups that are defined in your infrastructure. The information includes the ID of the computer group, its name, and description. You can find out how many computers and subgroups are in the computer group, and what software items are installed in that group. The report also provides information about the contracts that are assigned to a computer group, with the acquisition and maintenance cost for each contract.</p> <p>When you click a link in the Name column, detailed information about a particular computer group is displayed. You can now view the reports narrowed down to the computers that belong to the group. To do this, from the upper left corner, select the report that you want to view for this computer group.</p> <p>9.2.8 The report shows two types of computer groups: reporting and software template. Both types are visible to all entitled users. However, only an Administrator can create software template groups and modify software templates of existing groups of this type. For more information, see: Setting up computer groups (on page cxc).</p>
<p>9.2.12 Shared Disks</p>	<p> Required permission: View Shared Disks</p> <p>The report provides information about shared disks that are used in your infrastructure. For each shared disk, you can view information about its type, the exported directory, and the number of computers on which the disk is mounted. You can drill down the report to view detailed information about the shared directory or the computers on which the disk is mounted. If you use automatic scanning of shared disks, the report also shows the status of the scan and the computer that is designated to scan the shared disk.</p> <p>To view information on the report, discover shared disks in your infrastructure. For more information, see: Discovering software on shared disks (on page ccxxi).</p>
<p>Hardware Inventory</p>	<p> Required permission: View Hardware Inventory</p> <p>The report provides information about the state of BigFix clients in your infrastructure. For every computer on which a client is installed, you can view processor details, and the number of PVUs assigned to each core.</p> <p>If the PVU per core value that is assigned to a processor in your environment is incorrect, you can manually change it on the Hardware Inventory report to ensure that subcapacity licenses are properly calculated. For more information, see: Changing the PVU per core value (on page dcccxxiii).</p>
<p>9.2.3 Current State Values</p>	<p>9.2.12 Starting from application update 9.2.12, the report is deprecated.</p>

Report	Description
	 Required permission: Edit Server Configuration The report provides information that is intended to help the IBM Support team troubleshoot problems with BigFix Inventory. When you open a support case, generate a CSV or PDF version of this report and provide it to the HCL Support team.

Catalog

Report	Description
9.2.14 Products & Metrics previously Software Catalog	 Required permission: View Software Catalog and Signatures The report provides information about the content of the software catalog including software products and their publishers, part numbers that represent the products, and license metrics that the product uses. 9.2.11 When you drill down to the component level, you can also view the date when the software component reaches end of support.
Signatures	 Required permission: View Software Catalog and Signatures The report provides information about signatures that are available in the software catalog. You can find out what signature is used to detect each software item that exists in the catalog.
Catalog Audit	 Required permission: View Catalog Audit The report provides information about modifications that were made to the content of your custom software catalog. For each modification, you can see detailed information about the current status of the software entity, as well as its status before and after modification.

Contracts

Report	Description
Contract Usage Data	 Required permission: View Contracts The report provides information about contracts that you created for the software items. You can see the relationship between the number of software products that you are entitled to use and the actual number of software instances that are used. You can use the report to eliminate unused licenses and track computers that use unlicensed instances of a software product.

Report columns

This glossary provides a brief description of every column that can be found in the reports that are available in BigFix InventoryBigFix Inventory.

[A \(on page dcxxx\)](#) [B \(on page dcxxxi\)](#) [C \(on page dcxxxi\)](#) [D \(on page dcxxxv\)](#) [E \(on page dcxxxvi\)](#) [F \(on page dcxxxvi\)](#) [G \(on page dcxxxvii\)](#) [H I \(on page dcxxxviii\)](#) [L \(on page dcxxxix\)](#) [M \(on page dcxl\)](#) [N \(on page dcxli\)](#) [O \(on page dcxli\)](#) [P \(on page dcxlii\)](#) [Q R \(on page dcxliv\)](#) [S \(on page dcxlv\)](#) [T \(on page dcxlvii\)](#) [U \(on page dcxlix\)](#) [V \(on page dcl\)](#) [W \(on page dcl\)](#) [X Y \(on page dcl\)](#) [Z](#)



Note: When the value of the column cannot be populated because there is no data, one of the following values is displayed.

- `<no data>` when the data does not exist.
- `<not set>` when the value is not yet set. You can set it using the dedicated option.

If the `<not set>` value appears under the Product Name column, see: [The <not set> value appears under the Product Name on Software Classification panel \(on page dcclxi\)](#).

- `<n/a>` when the data is not relevant for the particular report line.
- The cell is empty when the value does not exist and cannot be populated.

A

Acquisition Cost

Overall cost of purchasing the license for the software product. It is the value that you specify while creating the contract.

Agent Version

Version of the BigFix agent that is deployed on the computer.

Always Not Charged

Indicates whether a component is always not charged.

Alternative User ID

An identifier of a user who currently uses or registered a software program. It is a number or a string specific to this software program.

Automatic Scan Status

Indicates whether the shared disk is scanned by using the automatic version of the optimized scanning mode. For more information, see: [Step 2a: Optimized mode - Automatic scanning of remote shared disks \(on page ccxxv\)](#). The column can have one of the following values:

- Disabled
- OK
- Pending

When you hover over the Pending status, you can see more details why the action is pending.

Average Runs per Day

Average number of times a process or software item is started per day.

Average Run Time

Average amount of time during which a process or software was used. It is calculated by dividing the total run time by the number of total runs. The average run time is the same as the total run time when the total number of runs equals one.

B

Brand

Brand of the processor that was matched in the PVU table. If the brand is "Other", the processor is not listed in the PVU table. For details of the processor read from the computer operating system, see: [Processor Brand String \(on page dcxlvi\)](#).

Bundling Tag Used

Indicates whether the product assignment is based on a bundling tag.

C

Catalog Part Numbers

Part numbers that represent the product that are listed in the software catalog.

Catalog-provided

- On the Products & Metrics report, it indicates whether the product and metric relation is provided in the software catalog. If the column value is NO, the relation comes from the SWID tag, or is defined by a user.
- On the Package Data report, it indicates whether the component from the particular vendor and in the specific version already exists in the software catalog.

Catalog Provider

Name of the catalog content originator.

Catalog Scan Results Missing

Indicates whether the catalog scan results are missing. The scan might be missing because the computer or a relay is offline, there is a network outage, or the last scan attempt was more than 30 days ago.

Catalog Scan Successful

Indicates whether the catalog scan was successful. The scan might be unsuccessful because the computer is out of space, it is misconfigured, or the scan was stopped.

Catalog Version

Version of the software catalog that is available on the computer. If different versions of the catalog are available on various computers, [update the catalog \(on page dcccxxiv\)](#).

Caused Detection

Indicates whether a file or a package contained enough information to cause detection of the related software. This column is not available by default. The data that is required to populate this column uses many resources, and the time to load and filter the reports will be extended, especially for huge environments.

To enable this column on the Scanned File Data and Package Data reports, go to **Management > Advanced Server Settings** and change the value of `enable_caused_detection_data` parameter to true. To display the column on the Scanned File Data or Package Data reports hover over the **Manage Report**

View icon , and click **Configure View**. Select the appropriate check box.

In versions earlier than 9.2.9, the Caused Detection column can be enabled only on the Scanned File Data Report. To display it, enable the aggregation of unrecognized scan data. Go to **Management > Unrecognized Scan Data** and choose whether you want to aggregate the data after every import or after the next import only. Next, configure the report view.

Changed PVU per Core Value

Specifies whether the number of processor value units for a specific server was manually changed.

Charged

Indicates whether the relation between the component and the product is charged based on the license terms. If the relation is charged, the component contributes to the license metric utilization of the product to which it is assigned.

Children

Number of nested computer groups.

Cluster Cores

The total number of processor cores available within the cluster of host computer systems.

Cluster Name

Name of the cluster to which a particular host computer system belongs.

Comment

Extra technical information. Possible values on the server level include: Incomplete virtualization hierarchy and All instances are excluded. Possible values on the software instance level include reasons for excluding the instance from pricing calculations, and suppressing the product from the inventory.

Component Definition Source

Specifies the entity that provided the component definition. From application update 9.2.9, this column replaced the [Software Catalog column \(on page dcxlvj\)](#).

Component Detailed Version

Detailed version of the software component that is installed on the computer. This information is provided for components that are discovered by the following means:

- Software identification (SWID) tags. Supplemental SWID tags are not supported. Thus, if the fix pack version is embedded in the supplemental tag, it is not shown.
- **9.2.11** Windows files in which the file version attribute is set and which is used as a software signature

For components that are discovered with other detection methods the values in Component Detailed Version and Component Version columns are duplicated.

Starting from application update 9.2.11, the column presents the highest reported version of the component. Versions are sorted in numerical order. If the version consists of digits and letters, it is sorted in numerical order and then the alphabetical order. Precision of sorting is up to four places per position, and up to four positions.

Component Name

Product component that was discovered on the computer.

Component Tags

Tags defined on the Software Components panel. These tags are common for all discovered instances of the same component.

Component Version

Version of the software component that is installed on the computer.

Computer Count

Number of computers on which the software is installed or the shared disk is mounted.

Computer Deletion Date

The date and time when the computer stopped being discovered, for example because it was removed from the infrastructure, the data source was removed, or the BigFix client was uninstalled.

Computer Name

Host name of the computer on which the data specific to the report was discovered. For VMware products displayed on the Resource Utilization report, the name of the computer where VM Manager Tool is installed.

Computer Type

Type of the computer: physical, virtual, or public cloud.

Computers

Number of computers.

Confidence

Reliability of the signature that is defined in the software catalog. The lowest confidence equals 0, the highest confidence equals 3.

Confirmed

Indicates whether the software classification was manually confirmed.

Contact Information

The column is not used.

Contracts

On the Computer Groups report, it is the number of contracts that are assigned to the computer group. On the Inventory Exploration report, it is the number of contracts that are defined for the software item.

Copies

Number of copies of the unrecognized file that are discovered on the computers in your infrastructure.

Cost Delta

Difference between the cost of purchasing the license for the particular software product, and the cost of used or unused instances of that product. If the cost delta is a positive number that is displayed in green, the number of product instances that are used does not exceed the license entitlement. The number indicates that you have some spare licenses that you can distribute in your enterprise or that you can reduce the license-related cost by not extending the excessive licenses. If the cost delta is a negative number that is displayed in red, the number of software product instances that are used exceeds the license entitlement. You must purchase licenses for more instances of the software product.

Cost per Seat

Cost of the license for an instance of the software product.

Count

The number of software instances in the group that is filtered based on the selected columns.

CPU Core Full Capacity

Total count of physical processor cores of the server on which the product is installed.

CPU Core Subcapacity

Number of processor cores that are used by a software product.

CPU Core Subcapacity Limit

Total count of physical cores of the server on which the product is installed. It is counted at the time when the subcapacity is maximum which means when the number of cores that are used by the product is the highest.

Creation

Date and time when the signature or software catalog entry was created.

Current Status

Details of the signature after modification.

D

Data Source Computer ID

An identifier of a computer as defined in BigFix.

Data Source Name

Name of the BigFix server from which the data is uploaded to BigFix Inventory.

Default Report

Indicates whether the saved report view is set as the default view for a particular report.

Default PVU Value

Indicates whether a default number of processor value units is assigned to the computer. The default number is assigned when the processor is not listed in the PVU table. The default value is 100 PVUs per core.

Designated Computer

Provides information about the computer that is designated to scan the particular share disk.

Description

Additional description of an item.

Details

On the Audit Trail, it is information about the type and date of action that influences data on other reports as well as the user who performed the action. On other reports, it is a link to detailed information about the software product.

Discovery

Indicates whether the signature is used for software discovery.

Discovery End

Date and time when the component instance was reported for the last time. If the component is still detected, the column shows <n/a>.

Discovery Start

Date and time when the component was first detected by BigFix Inventory. The Discovery Start date might be later than the [First Used \(on page dcxxxvii\)](#) date.

DNS Name

Name of the computer that is specified in the Domain Name System.

E

Edition

Edition of the Oracle Database.

End of Support

Date when the software component reaches end of support.

Endpoint Scanner Catalog Version

Version of the custom software catalog that is available on the computer.

Entitled Computers

Number of computers that are entitled to use the software.

Entitlement End

Date of the end of entitlement for the product that was provided during the creation of the contract.

Entitlement Start

Date of the start of entitlement for the product that was provided during the creation of the contract.

Excluded

Indicates whether the product on a computer is currently excluded from the pricing calculations.

Exported Directory

Exported directory of the shared disk.

F

Feature

Feature of the Oracle Database.

File Name

Name of the file that is used by the signature to discover the software.

File Path

Directory in which the file is located.

File System Scan Results Missing

Indicates whether the file system scan results are missing. Scan results might be missing because the computer or a relay is offline, there is a network outage, or the last scan attempt was more than 30 days ago.

File System Scan Successful

Indicates whether the file system scan was successful. The scan might be unsuccessful because the computer is out of space, it is misconfigured, or the scan was stopped.

File Version

Version of the file.

Windows **9.2.11** Starting from application update 9.2.11, the column returns the full version of the file, not only version limited to the release number. If you upgrade from earlier versions of BigFix Inventory, reupload results of the file system scan to show the full version. Otherwise, it is shown only for newly discovered or updated files. To reupload the result, run the Force Reupload of Software Scan Results fixlet.

First Detected

Date and time when the software component was detected by the scanner for the first time.

First Measured

Date and time when the license metric was first measured.

First Seen

Date and time when the information about the computer was imported from BigFix for the first time.

First Used

Date and time when the information about the process was first collected by the BigFix client. The First Used date might be prior to the [Discovery Start \(on page dcxxxv\)](#) date.

FlexPoint Bundle

FlexPoint bundle to which the product and its components are assigned. For more information, see: [Assigning products to FlexPoint bundles \(on page dli\)](#).

FlexPoints

Total utilization of FlexPoints by the FlexPoint bundle.

FlexPoints per Metric Unit

The factor by which the number of metric units is multiplied to obtain the number of FlexPoints that the products uses.

FlexPoint Quantity

Indicates the number of FlexPoints to which the product and related metric are assigned as a part of a FlexPoint bundle.

From Software Template

Indicates whether the software was discovered on a particular computer or was generated from a software template. In the latter case, the software is installed on a shared disk that is mounted on the particular computer. For more information, see: [Discovering software on shared disks \(on page ccxxi\)](#).

G**Global Default Report**

Indicates whether the saved report view is set as a default view of the particular report for all users.

H

Has Low Disk Space

Indicates whether the computer has less than 100 MB of free disk space.

Has Missing Prerequisites

Indicates whether any scanner prerequisites are missing on the computer.

Has Outdated Version

Indicates whether the latest version of the BigFix agent is installed on the computer. For some operating systems, no newer version of the agent is available.

Has Usage

Indicates whether a usage signature exists for the particular software.

I

IBM-provided Bundling Option

Indicates whether the product has the selected component listed as a bundling option in the software catalog provided by BigFix.

ID

Identifier of a computer group.

Imported Part Numbers

Part numbers that you imported to BigFix Inventory. They represent BigFix products that you purchased.

Install Path

Installation path of the Oracle Database.

Installed Software

Number of software items that are installed on the computer.

Instance

Identifier of the software instance. For VMware products displayed on the Resource Utilization report, it specifies the license key, machine type and the host name.

Instance Tags

Tags for individually discovered instances of a software. These tags are independent from component tags in catalog. Discovery tags are set by using the dedicated option.

Instances Covered by Contracts

Number of software items by a particular vendor for which contracts are defined.

Internal Computer ID

Identifier of a computer as defined in BigFix Inventory.

IP Address

IP address of the computer.

Is Capped

Indicates whether the partition is capped, that is whether its access to processor resources is limited.

Available only for AIX. For other operating systems, the column shows *no data*.

Is Out of Sync

Indicates whether the time that is set on the computer is at least an hour from the time that is set on the BigFix Inventory server.

L

Last Login

Date and time when the user last logged in to BigFix Inventory.

Last Scan Attempt

Date and time when the last scan was initiated.

Last Seen

Date and time when the agent connected to the BigFix server the last time.

Last Used

Date and time when the software item was last used or a process was last run.

Latest Scan Import

Date and time when scan data from a particular computer was last imported to BigFix Inventory. The date is only updated if there are changes in the content of the scan file. If scan data is unchanged after any new scan, the **Latest Scan Import** column will not be updated.

Licensed Computers

Number of computers that are entitled to use the software and have this software installed.

License Delta

Difference between the number of software product instances that you are entitled to under a particular license and the actual number of software instances that are used. If the license delta is a positive number that is displayed in green, the number of product instances that are used does not exceed the license entitlement. If the license delta is a negative number that is displayed in red, the number of product instances that are used exceeds the license entitlement.

Logical Processors

Number of processor cores that are assigned to the partition on which the BigFix agent is installed. The number takes into account hyper-threading on the x86 architecture.

LPAR Entitlement

Number of processor cores that are assigned to the logical partition. Available only for AIX. For other operating systems, the column shows *no data*.

M

Maintenance End

Date of the end of maintenance of the software that was provided during the creation of the contract.

Maintenance Start

Date of the start of maintenance of the software that was provided during the creation of the contract.

Maximal Trend Value

Maximum value of the license metric utilization over the last 30 days.

Maximal Trend Value History

Graphic representation of the history of license metric utilization over the last 30 days.

MD5

MD5 checksum calculated for a file. Checksums work as fingerprints of files, and can be compared to checksums of known virus signatures to recognize whether files were altered or tampered with. They are used to ensure file integrity and safety in advanced security scenarios.

Metric

Type of the metric used by the software item.

Metric Quantity

The highest number of metric units that the product used over the specified period.

Metric Quantity History

Graphic representation of the history of metric consumption over the specified time period. If you set the license threshold, the graph also includes a line that represents the threshold. For more information, see: [Setting license metric threshold \(on page dxc\)](#).

Metric Subtype

Subtype of the license metric. For VMware products displayed on the Resource Utilization report, metric subtype indicates if the count refers to the metric utilization or overall license capacity.

Metric Type

Type of the license metric.

Metric Units Convertible to FlexPoints

The number of metric units that are converted into FlexPoints. The number of convertible metric units can be lower than the overall number of units that are used by the product. It happens when some instance of the product are not a part of the FlexPoint bundle.

Model

Model of the processor.

Modification

Date and time when the item was modified.

Modification Date

Date and time when the item was modified.

Modification Type

Type of the modification.

Modified By

User who modified the item.

Modified Software

Software item that was modified.

Mounted Shared Disks

List of directories of all shared disks that are mounted on the computer. To see the data, run the Discover Remote Shared Disks task on the computer.

N**Name**

Name of an item. It depends on the report on which the column is displayed.

Normalized Component Name

The component name that is suggested after the analysis of package details and current content of software catalog.

Normalized Component Vendor

The vendor that is suggested after the analysis of package details and current content of software catalog.

Normalized Component Version

The component version that is based on what is embedded in package details and rounded up to two decimal places.

Number of Packages

Number of specific packages that are identified in your environment.

O**One to One**

Indicates whether the product to which the component is currently assigned is the only bundling option that is available in the software catalog.

Operating System

Operating system of the computer on which the BigFix agent is installed.

P**Package Scan Successful**

Indicates whether the latest package scan was successful. The scan might be unsuccessful because the computer is out of space, it is misconfigured, or the scan was stopped.

Parent Organization

Organization that owns the content of the catalog.

Partition Cores

Specifies the number of processor cores that are assigned to the partition on which the BigFix agent is installed.

On the All Metrics report, when you click a product that is licensed under the PVU subcapacity or RVU MAPC subcapacity, and you drill down to information about a single computer, the column shows the number of partition cores measured when the peak of the metric value occurred. When you sum up partition cores from VMs running on a single server, the value might not be equal to the number of cores that contribute to metric utilization on the server level. It is in line with subcapacity pricing rules. For more information, see: [Virtualization Capacity License Counting Rules](#).

Partition Virtual Processors

Number of virtual processors that are assigned to the partition on which the BigFix agent is installed. The value is used in further calculations when the endpoint works in the uncapped mode. Available only for AIX. For other operating systems, the column shows *no data*.

Path

Installation path of the software that was discovered on the computer or where the `.slmtag` file was discovered.

Physical Shared Pool Cores

Number of processors that work in the shared mode.

Peak Date

Indicates the date of license metric peak value for the reporting period.

Platform Name

Name of the platform on which the software is discovered by means of the particular signature.

Predefined End of Support

Indicates whether the end of support date was provided by BigFix or was specified manually by a user.

Present

Indicates whether the component instance is currently installed.

Private

Indicates whether the saved report view can be used only by its creator or by other users as well.

Process Name

Name of the process on which the calculation of software usage is based.

Processor

Details of the processor.

Processor Brand String

Full specification of the processor including its brand, model and speed as read from the computer operating system.

Product Definition Source

Specifies the entity that provided the product definition. From application update 9.2.9, this column replaced the [Software Catalog column \(on page dcxlvj\)](#).

Product Name

Name of the product.

Product Version

Version of the software product.

Publisher

Name of the company that published the software item.

Publisher Definition Source

Specifies the entity that provided the publisher definition. From application update 9.2.9, this column replaced the [Software Catalog column \(on page dcxlvj\)](#).

Publisher Name

Name of the company that published the software item.

PVU per Core

Number of processor value units that are assigned to a processor core on a computer system. By default, the value is taken from the PVU table. It can also be manually set on the Hardware Inventory report. For more information, see: [Changing the PVU per core value \(on page dcccxxiii\)](#). See also: [Changed PVU per Core Value \(on page \)](#).

PVU Full Capacity

Number of processor value units that are used by a product that is installed on the server.

PVU Subcapacity

Number of processor value units that are used by a product that is installed on the server partition.

PVU Subcapacity Limit

Total count of processor value units for the server on which the product is deployed, counted at the time when subcapacity is maximum. Maximum subcapacity is when the number of cores that are used by the product is the highest.

R

Rank

The position of a file in the ranking of 1000 files that are most commonly encountered in your computer infrastructure but do not produce matches for any signature.

Recognized

Indicates whether the file was recognized as part of an existing signature.

Release

Release of the software product.

Release Definition Source

Specifies the entity that provided the release definition. From application update 9.2.9, this column replaced the [Software Catalog column \(on page dcxlvj\)](#).

Reported Since

Date and time when the package or file was discovered for the first time.

Reputation

Indicates whether the file is a malware threat. Possible values:

- `<n/a>` is when SHA-256 checksum is not collected.
- `<no data>` is when the Tool for Reporting File Reputation is not installed.
- `<No threat>` is when the file is not classified as a security threat.
- `<No reputation>` is when the classification was done but the type of reputation is unknown.
- `<Threat>` is when the file is classified as a security threat. This value is a link which directs you to BigFix website that contains details about the hash reputation. To see the information you must log in using your IBMid.

Reputation Details

Indicates the type of a threat. Possible values:

- Trojan
- Downloader
- Spyware
- Backdoor
- Rootkit
- Ransomware
- Bot

- Adware
- Worm
- Unwanted
- Virus
- Miner

Reputation Import

The date when the reputation was updated in BigFix Inventory.

Reputation Update

The date when the reputation was updated in the reputation data source.

S

Scanning Shared Disks

Indicates whether shared disks were scanned during the last software scan.

Server Active Sockets

Number of physical processor chips that are mounted on the server.

Server Cores

Number of processor cores that are used by the server.

Server ID

Identifier of the server. The format of the identifier is specified by the **managedServerTagTemplate** parameter for hosts from which data is collected by the VM Manager Tool or by the Run Capacity Scan on Virtualization Hosts fixlet or by the **standaloneServerTagTemplate** parameter for the remaining hosts. For more information about the parameters, see: [Advanced server settings \(on page cdii\)](#).

Server Name

A unique system in the enterprise. For a physical machine, it is the hardware manufacturer, type, and a machine serial number. For a virtual machine, it is the manufacturer and host name. For a virtual machine with an incomplete definition, it is a UUID that is prefixed with `TLM_VM_`.

SHA-256

SHA-256 checksum calculated for a file. Checksums work as fingerprints of files, and can be compared to checksums of known virus signatures to recognize whether files were altered or tampered with. They are used to ensure file integrity and safety in advanced security scenarios.

Shared

Indicates whether a component instance is shared between multiple products.

Shared Disk Type

Type of the shared disk.

Shared Mode

Indicates whether the logical partition works in the shared mode. Available only for AIX. For other operating systems, the column shows *no data*

Shared Pool ID

Identifier of the shared processor pool.

SID

Unique name of the Oracle Database instance.

Signature Creation

Date and time when the signature was created.

Signature Definition Source

Specifies the entity that provided the signature. From application update 9.2.9, this column replaced the [Software Catalog column \(on page dcxlvj\)](#).

Signature ID

Identifier of the signature.

Signature Modification

Date and time when the signature was last modified.

Size

Size of the file in bytes.

Software Catalog

Software catalog that contains the signature for the software item. Starting from application update 9.2.9, the Software Catalog column is substituted with the following columns:

- [Component Definition Source \(on page dcxxxii\)](#)
- [Product Definition Source \(on page dcxliii\)](#)
- [Publisher Definition Source \(on page dcxliv\)](#)
- [Release Definition Source \(on page dcxlv\)](#)
- [Signature Definition Source \(on page dcxlvj\)](#)
- [Version Definition Source \(on page dcl\)](#)

Software Component

Name of the software component.

Software ID

Identifier of the software to which the modification applies.

Software Level

Specifies the version, release, and modification number of a software item.

Software Name

Name of the software product.

Software Product

Name of the software product.

Software Release

Release of the software product.

Software Tags Scan Successful

Indicates whether the last software tags scan was successful.

Software Version

Version of the software product.

Status

Indicates whether data from a hypervisor is available.

System Model

Name of the hardware system or virtualization technology that is discovered by the agent.

Suppressed

Indicates whether the component is currently suppressed from the inventory.

T

Threshold

The maximum number of metric units that a product is entitled to use within a computer group. The value is set manually and is used to calculate the metric threshold delta.

Threshold Delta

It is calculated by subtracting the metric quantity from the threshold. If license utilization is lower than or equal to the threshold, the delta is displayed in green. Otherwise, it is displayed in red.

Top-level Directory

Indicates whether the shared disk is a top-level directory in the shared disk structure.

Total Acquisition Cost

On the Computer Groups report, it is the total cost of purchasing licenses for software items in the computer group. On the Inventory Exploration report, it is the cost of purchasing licenses for software items by a particular vendor.

Total Components

Number of components of software products that are listed in the catalog.

Total Licenses

Number of licenses for the particular software product that are available. If the product is licensed per computer, the number of computers that are entitled to use the product is provided. If an unlimited

number of products can be deployed under the particular license, the column includes the abbreviation ELA (Enterprise Licensing Agreement). Higher-level users can define a contract that applies to a broad group. If lower-level users view a contract usage report that is defined in such a way, they see the larger total number of licenses for the product. However, when the lower-level users view the list of instances, they see only those instances of the product for which they have permissions.

For example, a global administrator defines a contract for 5000 instances of Lotus Notes, out of which 3500 instances are used. A lower-level user has an infrastructure that consists of 100 computers. The user is entitled to 100 instances of Lotus Notes, but only 70 instances of the product are used. If the global administrator makes the contract viewable to the lower-level user, the user sees that 3500 instances of Lotus Notes out of 5000 available instances are used. However, when lower-level users view the list of instances that are pertinent to their computer group, they see only the numbers that reflect the use of licenses in their computer group. The report shows 70 used instances of Lotus Notes out of 100 instances available under the particular license.

Total Maintenance Cost

On the Computer Groups report, it is the cost of maintaining all software items in the computer group. On the Inventory Exploration report, it is the cost of maintaining software items by a particular vendor.

Total Publishers

Number of software publishers that are listed in the catalog.

Total Releases

Number of releases of software products that are listed in the catalog.

Total Runs

Number of times a software item or process was started.

Total Run Time

Indicates how long the software item or process was running.

Total Signatures

Number of signatures that are listed in the catalog.

Total Software Products

Number of software products that are listed in the catalog.

Total Versions

Number of versions of software products that are listed in the catalog.

Type

- On the Inventory Exploration report, it indicates whether the displayed item is a publisher, software product, version, release, or component.
- On the Package Data report, it is the type of the package management system.

- On the Computer Groups report, it indicates whether the type of the group is Reporting or Software Template. For more information, see: [Setting up computer groups \(on page cxc\)](#).
- On other reports, it is the type of the processor according to number of cores.

U

Unique User ID

An identifier of a user who currently uses or registered a software program. It is a specific ID that is assigned to one user only, and can either be an e-mail, first name, last name, initials or a random string.

Unknown Use

Number of software instances for which use is not monitored. These instances are grouped based on the columns defined or selected by the user in the report view.

Unused

Number of software instances that were not used and that are grouped based on the columns defined or selected by the user in the report view.

Unused in the Last Month

Number of software instances that were not used within the last month and that are grouped based on the columns defined or selected by the user in the report view.

Unused in the Last Quarter

Number of software instances that were not used within the last three months and that are grouped based on the columns defined or selected by the user in the report view.

Unused in the Last Week

Number of software instances that were not used within the last seven days and that are grouped based on the columns defined or selected by the user in the report view.

Unused in the Last Year

Number of software instances that were not used within the last year and that are grouped based on the columns defined or selected by the user in the report view.

Used

Number of software instances that were used and that are grouped based on the columns defined or selected by the user in the report view.

Used in the Last Month

Number of software instances that were used within the last month and that are grouped based on the columns defined or selected by the user in the report view.

Used in the Last Quarter

Number of software instances that were used within the last three months and that are grouped based on the columns defined or selected by the user in the report view.

Used in the Last Week

Number of software instances that were used within the last seven days and that are grouped based on the columns defined or selected by the user in the report view.

Used in the Last Year

Number of software instances that were used within the last year and that are grouped based on the columns defined or selected by the user in the report view.

User Group Name

Specifies the source of the information related to software users.

User Name

Name of the user who created the custom report view.

V

Vendor

Vendor of the processor or of the software that is installed on the computer.

Version

Version of an item. It depends on the report on which the column is displayed.

Version Definition Source

Specifies the entity that provided the version definition. From application update 9.2.9, this column replaced the [Software Catalog column \(on page dcxlvj\)](#).

Virtual Shared Pool Cores

Number of active processor cores in the virtual shared pool.

Vulnerability Risk (Preview)

Lists Common Vulnerabilities and Exposures matched with a given software component. The column can contain multiple values that are displayed after expanding.

W

Website

Website of the software publisher.

Y

Your Next Scheduled Export

Date and time when the saved report will be sent to the recipients that you specified in the schedule. The column does not contain sending dates scheduled by other users.

Adding and removing tags

9.2.14 Available from 9.2.14. BigFix Inventory allows you to tag certain elements on the reports. You can tag software components, and whole report rows and thus mark them to meet your individual business needs.

 To perform this task on the Software Components report, you must have the View Software Catalog and Signatures and Manage Catalogs permissions.

 To perform this task on the Software Classification panel, you must have the View Endpoints and Manage Software Classification permissions.

Tagging software components, and whole report rows might be used to differentiate the specific kind of software components, or the discovered software. It can help you group items according to their importance, specificity, type or use. You can tag the following report elements:

- Software components on the Software Components report.
- Individually discovered instances of software on the Software Classification panel.

Adding tags

1. Log in to BigFix Inventory.
2. Go to **Reports** and choose the relevant report.
3. Select one or more rows for which you want to add tags.



Note: On the Software Components report, you can select up to a 100 rows at a time.

4. Hover over the  icon, and click **Add Tags**.
5. List tags that you want to add for the selected components. You can select an existing tag or add a new one.
6. Click **Add Tags**.

The tags are visible on a relevant report.

Removing tags

To remove tags, hover over the  icon, and click **Remove Tags**. Select the tags and confirm by clicking **Remove Tags**.

Applying report filters

The type and amount of information that is displayed on the reports is set by default. You can apply filters to narrow down the scope of the report to information that meets specific criteria.

1. In the top navigation bar, click **Reports** and open the report that you want to view.
2. To view the filtering options, in the upper right corner, hover over the **Manage Report View** icon , and click **Configure View**.
3. **Optional:** If you want to add multiple filters, choose **all** or **any** to specify whether the items must match all filtering criteria or any of them.
4. Select the column by which you want to filter the report and the filter operator. Then, enter the value that you want to use as the filter.

Filter operators that are available depend on the type of column by which you are filtering the report.

Filters

Specify the report filter which matches all of the following conditions:

Component Name	equal to		🗑️	+
	<ul style="list-style-type: none"> equal to not equal to contains does not contain begins with ends with does not end with does not begin with 			

 **Restriction:** When you use the **IP Address** operator and choose the **ends with** operator, no computers are displayed even if some computers match the criteria. To work around this issue, you can choose other relations, such as **contains** or **begins with**.

5. **Optional:** To add another filter, click the plus sign.

 **Tip:** Click the trash can to delete a filter.

6. Click **Submit**.

Exporting reports to a file

If you need a hardcopy of a report that is generated in BigFix Inventory you can export the report to a **CSV** or **PDF** file. You can also use the **CSV** file to create charts and statistics that are based on the information from the report.

1. In the top navigation bar, click **Reports** and select the report that you want to view.
2. **Optional:** You can customize the report before export. To specify the required filters and columns, hover over the **Manage Report View** icon , and click **Configure View**. To display all columns on the report, select the **Select All** check box in the Columns section.
3. To export the report to a file, hover over the **Export Report** icon , and click **PDF** or **CSV**.

 **Note:** If your reports are not displayed correctly, see: [User interface problems \(on page dccxciii\)](#).

9.2.15 Report files names, including report files that are scheduled to be emailed, follow new naming convention. Starting from application update 9.2.15, the file name consists of the predefined or custom report name, the default report name, and the time stamp. For example: `My_Custom_Report-software-classification-20190215085933.csv`.

Creating saved report views

Saved report views are derived from base reports and provide quick access to information that you use frequently. Some saved report views are delivered with BigFix Inventory and are available for you immediately. You can also specify filters and parameters of any report and create your own saved report views.

Three types of saved reports are available.

The table consists of five columns and four rows.

Type	Description	Can be sent to specified recipients	Can be set as a default view	Can be modified
9.2.4 Predefined	Reports that are delivered with BigFix Inventory.	Yes	Yes	
My Reports	Reports that you created.	Yes	Yes	Yes
Other Reports	Reports that were created by other users and marked as public.	Yes		

1. Open the report for which you want to create a custom view.
2. To customize the report view, hover over the **Manage Report View** icon , and click **Configure View**. Add or remove columns and apply filters, then click **Submit**. Resize and sort the report columns by clicking and dragging.
3. To save the report, hover over the **Manage Report View** icon , and click **Save As**. Specify the name of the report. You can also choose one of the following options:
 - If you do not want other users of BigFix Inventory to view your custom report, select **Private**.



Note: The Administrator can view and delete reports that are marked as private.

- If you want the custom report to be your default view for the particular report, select **Set as default**.



Tip: If the Administrator set a customized view as global default for a particular report, you can overwrite this setting by creating your default view. Your settings take precedence over the global settings.

- If you want the custom view to be the default report view for all users of BigFix Inventory, select **Set as global default**. You must be an Administrator to perform this task.

All types of saved reports can be accessed from the **Saved Reports** panel or from a drop-down list on the base report from which the saved report was created.



Deleting saved report views

To delete a saved report view, select the report from the list of all saved reports, and click **Delete**.

1. To open the list of all saved reports, click **Reports > Saved Reports**.
2. Select the report that you want to delete.
3. In the upper left corner, click **Delete**.

Changing the default report view

To change the default view that you set for a particular report, you must first clear the settings and then set the new custom view.

1. To open the list of all saved reports, click **Reports > Saved Reports**.
2. Open the report for which you want to edit or change the default customized view.
3. Clear the **Set as default** or **Set as global default** check box, and click **Save**.
4. Open the report for which you want to change the default view and customize it according to your needs.
5. Hover over the **Manage Report View** icon , and click **Save As**. Specify the report name, and select **Set as default** or **Set as global default**. Click **Create**.

Scheduling report notifications

You can configure report notifications to be sent to a list of recipients regularly or when a specific number of items appears on a report. For example, when the number of computers that have not reported to BigFix Inventory within the last seven days exceeds five. Thanks to the automatic notifications, you are always up-to-date with changes in your environment.

Configuring the server to send mail notifications

To be able to schedule report notifications, you must first configure the BigFix Inventory server to send mail notifications.

 You must be an Administrator to perform this task.

1. Log in to BigFix Inventory.
2. In the top navigation bar, click **Management > Mail Settings**.
3. Specify the SMTP server to which you want to have the email notifications sent.

 **Important:** The SMTP port must be open for communication with BigFix Inventory.

4. Choose the port through which you want to have the email notifications sent.
5. **Optional:** To have the email encrypted, select **Use STARTTLS**.
6. Specify the domain through which you access the BigFix server.
7. Choose the authentication method.
8. In the **From address** field, specify the address that is displayed as the sender of the email.
9. **Optional:** To check whether you correctly configured mail settings, click **Send Test Email**.
10. To save the configuration, click **Save**.

You can now schedule report notifications to be sent automatically.

Scheduling report emails

You can schedule reports to be emailed to the specified recipients. This option is especially useful if a person does not work with BigFix Inventory, or is not familiar with the application, yet must have access to the reports.

Ensure that the BigFix Inventory server is configured to send mail notifications.

1. Open the report that you want to email, hover over the **Export Report** icon , and click **Schedule Export**.

 **Important:** Reports that you want to email must be saved before they can be sent. You can either save the report without changing its default view or create a customized view. You can set only one schedule per a saved report. If you try to set a next schedule, it replaces the previous one. For more information, see: [Creating saved report views \(on page dcliii\)](#).

2. Select the file format and adjust the page settings.
3. In the **Email** field, enter a comma-separated list of email addresses to which you want to send the notifications.
4. Select the language of the report.
5. Select at least one option for emailing the report:

- **Email on a regular basis**

The report is emailed according to a daily, weekly, or monthly schedule.



Note: The start time for each report must differ by at least 5 minutes.

- **9.2.3 Email if the number of rows exceeds**

The report is emailed only if a specific number of rows is exceeded. You can customize a report to show only products for which PVU exceeds a particular value, or contracts whose entitlement expired.

You will be notified whenever the specified number of items appears on the report. This option is an additional notification and does not affect the regular schedule.



Note: Reports based on conditions are emailed during data imports.

Schedule Export ✕

Report Subscription

Format* PDF ▼

Page Size Letter ▼

Orientation

Portrait
 Landscape

Email* name@example.com

Separate multiple recipients with a comma.

Language* English ▼

Email on a regular basis:

Start Time 02/06/2019 03:56PM

Frequency Daily ▼

Every 1 day(s)

Email if the number of rows exceeds: 3

Save

6. Click **Save**.

9.2.15 Report file names, including report files that are scheduled to be emailed, follow new naming convention. For more information, see: [Exporting reports to a file \(on page dclii\)](#). Additionally, starting from application update 9.2.15, the email subject contains the name of the predefined or custom report. For example: *Scheduled Report 'My Custom Report' from BigFix Inventory*.

To disable the export schedule, clear the **Report Subscription** check box and click **Save**.

Contracts

You use contracts to track license compliance for the software products that are installed in your environment. When you purchase a software license, you create a contract that contains information about the cost of acquiring and

maintaining the license, and the period of entitlement. You use that information to track spending that is associated with software licenses and manage costs of extending or downsizing licenses for the software products. You can also avoid costs that are related to license non-compliance by ensuring that enough software licenses are purchased.

Creating contract custom fields

You can customize your contracts by creating custom fields that store information of your choice about the software licenses in your environment. Choose a meaningful name for each custom field to reflect its purpose. After you create a custom field, it is available in the **Create Contract** pane, among other fields that are required to create a contract.

 You must have the Manage Contracts permission to perform this task.

1. In the top navigation bar, click **Management > Contract Custom Fields**.
2. To create a custom field for a contract, in the upper left corner of the **Contract Custom Fields** window, click **New**.
3. In the **Create Contract Custom Field** pane, enter the name of the field, specify whether the field is required for a contract, and choose the type of the requested input.

Create Contract Custom Field

Name*

Required

Type* ▼

4. To save the custom field, click **Create**.
5. To make the custom field available for the reports, click **Reports > Import Now**.

You created a custom field for a contract. You can now create a contract or edit the existing ones to add the values for the new fields.

Creating contracts

You create contracts to store information about your software licenses, their cost, and the period for which they entitle you to use the software. The contracts are then used to generate a contract usage report. It shows the number of software instances to which you are entitled and the actual number of software instances that are used.

 You must have the Manage Contracts permission to perform this task.

You can create a contract for a software product, its version, or a release. If you have an entitlement that applies to all versions of a product, create a contract on the level of the product. It will cover all versions and releases. If you have an entitlement for a specific version or release, create a contract on the appropriate level to ensure that it covers the right software version or release. You can also create a contract that covers multiple products, versions, and releases according to your entitlements.

1. In the top navigation bar, click **Management > Contracts**.
2. To create a contract, in the upper left corner click **New**.
3. Enter the name of the contract and specify the software product, version, or release that you want to assign to that contract. When you start typing, a list of suggestions is displayed below the search field. It is based on the content of your software catalog. Choose a product from the list.

! **Important:** The software product for which you want to create a contract must be in the software catalog. If it is not available in the BigFix catalog, create a custom catalog entry.

Create Contract

Name* IBM Contracts

Software Software Version IBM Endpoint M

License Type Seats ELA

Computer Group All Computers

Acquisition Cost*

Maintenance Cost*

IBM Endpoint Manager for Core Protection 9
 IBM Endpoint Manager for Core Protection Data Protection Add-on 8
 IBM Endpoint Manager for Core Protection Data Protection Add-on 9
 IBM Endpoint Manager for Lifecycle Management - Content Site

4. Choose the type of licensing that is used for the product that you assigned to the contract:
 - Choose **Seats** if a limited number of product instances can be deployed. Additionally, you must specify the number of software instances that you are entitled to under the license.
 - Choose **ELA** if an unlimited number of product instances can be deployed.

Create Contract

Name* IBM Contracts

Software IBM Endpoint Manager for Core Protection 9 x

Software Version Enter a software version

License Type Seats ELA

Seats* 200

Computer Group All Computers

Acquisition Cost* 20000

Maintenance Cost* 1000

Entitlement Start* 04/30/2015

Entitlement End 04/30/2016

Maintenance Start

Maintenance End

Create

5. Choose the computer group to which you want to assign the contract. Provide the acquisition and maintenance cost, and the entitlement start date. You can also specify the entitlement end date, and the maintenance start and end dates.



Important: Acquisition cost is the total cost of the license for a particular product. If you create a contract for a specific number of seats, provide the total cost of acquisition for all seats.

6. To save the contract, click **Create**.

7. To make the contract available for generating a contract usage report, run an import.

You created a contract that reflects your entitlement to use a particular software product under the license that you have for that product. You can now view the contract usage report to find out whether you conform to the license regulations.

Managing security threats

Learn how to use BigFix Inventory to manage security threats in your environment. You can view whether any of the installed components is prone to any Common Vulnerability and Exposure (CVE).

Collecting file checksums

9.2.3 Available from 9.2.3. Checksums are long strings that describe the content of files and act as their fingerprints. You can enable the calculation of checksums for files on your computers to check their integrity and to ensure that they were not altered or tampered with.

In general, checksums are used to ensure safe transmission of files between devices, or across the Internet. In the latter case, after downloading a file, you can calculate its checksum and then compare it to the original checksum that was published together with the download. If both checksums match, you know that your download is exactly the same as the published source file, and that it is secure. Otherwise, there was some data loss or alteration, which means that the file was corrupted, either by accident or intentionally.

In BigFix Inventory, checksums are calculated during the file system scan and are created for each file that is discovered by the scan. In other words, apart from collecting usual information about files, such as their version or size, the scanner also collects information about their checksums. The checksums can then be viewed in the user interface or retrieved by using REST API, and compared to the checksums of known virus signatures or malware to ensure that none of such files is present in your environment. If any of the collected checksums finds a match in the database of corrupted checksums, there is a high probability that the file associated with this checksum is not secure, and presents a potential security breach.

Checksums in BigFix Inventory can be calculated by using two hash functions, MD5 and SHA-256. The choice of either of these should depend on the database of corrupted checksums that you own, and that you can compare with checksums collected by the scan.

Enabling the collection of checksums

9.2.3 Available from 9.2.3. You enable the collection of checksums by running a fixlet that changes the configuration of the software scan for BigFix Inventory. When you select checksums and run the fixlet against a chosen endpoint, this endpoint is given new properties that represent either MD5, SHA-256, or both. Based on these properties, the software scan can recognize which checksums must be collected from an endpoint. If you choose just one type of checksum and want to add another one later, you need to run the fixlet again.

- Upgrade the BigFix Inventory server to version 9.2.3, or higher.
- If you upgraded from previous versions, stop all *Initiate Software Scan* actions, and rerun the software scan against your endpoints. The software scan was changed to collect checksums, and must be started from the updated fixlet.
- Divide your environment into scan groups to distribute the load of the imported data:

- If the extended data import (three times longer) is acceptable, enable the collection of file hashes for all scan groups, and collect the data according to schedule.
- If the extended data import does not meet your expectations, rearrange your scan groups into smaller ones with fewer endpoints to lower the amount of data included in a single data import. After the first import is completed for all scan groups, you can go back to the previous setup.

It is important that you enable checksums for a small group of endpoints. After data from the first group is imported to BigFix Inventory, proceed to the next group. Because of the checksums, each file is detected as changed. By default, BigFix Inventory imports the delta file system scan, so only data that changed since the last scan. However, when it detects too many changes, it always chooses the full file system scan over the delta file system scan. Importing so many results might overload your data import. In such a case, you must recover from accumulated scans. This applies not only to enabling the checksums for the first time, but also to each change, such as adding a new type of checksum to be collected, or removing it.

1. Log in to the BigFix console.
2. In the navigation bar, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. In the upper right pane, select **Configure File Checksums Collection (MD5/SHA-256)**.
4. Select the types of checksums that you want to collect, either MD5, SHA-256, or both.

Description

Collected File Checksums

MD5 checksum

SHA-256 checksum

 **Important:** Always select all checksums that you want to collect, especially if you are running the fixlet for the next time. New properties always overwrite the previous ones.

5. Click **Take Action**.
6. Select the computers from which you want to collect checksums, and click **OK**.
7. **Optional:** In the navigation bar, go to **Analyses**, select **File Checksums Collection Settings (MD5/SHA-256)**, and click **Activate**. The analysis shows which checksums are collected from your endpoints.

Computer Name	Minimal Scanner version installed	MD5 checksum calculation enabled	SHA-256 checksum calculation enabled
 NC007143	True	True	True
 NC126312	False	False	False

Checksums are displayed in BigFix Inventory after the next software scan finishes and its results are imported from BigFix during the import of data.

In an environment with 60 000 endpoints divided into 6 scan groups (each with 10 000 endpoints), where each scan group is scanned on a different day, the file hashes will be collected in 6 days. The initial import for each scan group after enabling the collection might be three times longer. Next imports will take about 10% longer.

Impact of file hashes on the BigFix Inventory database size

For both DB2 and SQL Server databases, the collection of file hashes (MD5 and SHA-256) is expected to result in a 20% growth of the disk space consumption.

Impact of file hashes on the BigFix client

File hashes are calculated during the software scan and the results are gathered on the endpoint. The size of scan results will increase by about 5%. For an average endpoint with 30 matched and 800 unmatched raw data files, an additional 0.5 MB of disk space might be consumed.

You can [view the checksums in the user interface \(on page dclxii\)](#) or [retrieve them by using REST API \(on page dclxiii\)](#).

Viewing checksums in the user interface

9.2.3 Available from 9.2.3. To view checksums in the user interface, you must enable columns that show this data. It can be useful to verify checksums of single files. For bigger amounts of files, use REST API that can retrieve all collected checksums at once.

1. Log in to BigFix Inventory.
2. Go to **Reports > Scanned File Data**.
3. Hover over the **Manage Report View** icon , and click **Configure View**.
4. Select the **MD5** or **SHA-256** check box, and then click **OK**.

Checksums for each file are shown in the appropriate columns.

File Name	MD5	SHA-256
a04.sys	5ab92f5cff64054f2c5e64f45fc092f8	ce6b20ee7f7797e102f68d15099e...
a05.sys	10e5a5c00e2bed3ca40df5afd32fa...	f0af17449a83681de22db7ce1667...
a06.sys	b843ae5ec31fa7ac864b34985675...	86d5bc08c2eba828a8e3588e25a...
a07.sys	764f368a99e0c7c22c9ad735ec76...	364aba675151243a1e105cee426...

Retrieving checksums through REST API

9.2.3 Available from 9.2.3. To retrieve information about checksums, you can use REST API that is generally used to retrieve raw scanned file data. This API retrieves information about scanned files, such as name, size, version, path, and so on. It can also return the values of the MD5 and SHA-256 columns, which show the checksums.

Before you begin

This topic is only an excerpt from the REST API call, intended to show specific examples that can be used to retrieve the MD5 and SHA-256 checksums. For more information, see [REST API for retrieving raw scanned file data \(on page dccc/xxxiii\)](#).

Sample columns

The following response body shows default columns that can be retrieved by using this REST API. You can use the names of the shown columns while creating your REST API requests. If you do not specify any columns in the request, all of them are returned.

```
{
  "scan_file_id": 3,
  "computer_id": 5,
  "path": "C:\BES\BESAirgapTool",
  "name": "BESAirgapTool.exe",
  "full_path_sha1": "9eb46086883dcb6d2aefa2d6fd40bf8e736621b",
  "size": 92174,
  "version": null,
  "md5": "8cb2289800b34ef1a5f472f7177348e2",
  "sha256": "99e0c7c22c9ad735ec766914f0af17449a83681de22db7ce16672f16f37131bec0022371d4ace5d1854301e0",
  "valid_from": "2015-07-31T07:03:21Z"
  "valid_to": "9999-12-31T23:59:59Z"
}
```

Examples

Example 1: Retrieving information only about name, computer name, path, MD5, and SHA-256 columns.

Request:

```
https://hostname:port/api/sam/raw_file_facts?token=token
&columns[]=name&columns[]=computer.name&columns[]=path&columns[]=md5&columns[]=sha256
```

Output:

```
{
  "path": "C:\BES\BESAirgapTool",
  "name": "BESAirgapTool.exe",
  "md5": "8cb2289800b34ef1a5f472f7177348e2",
```

```

"sha256" :
"99e0c7c22c9ad735ec766914f0af17449a83681de22db7ce16672f16f37131bec0022371d4ace5d1854301e0",
"computer":{
  "name": "NC581058"}
}

```

Example 2: Retrieving information about files with a specific MD5 checksum.

This example can be used to quickly find a corrupted checksum.

Request:

```

https://hostname:port/api/sam/raw_file_facts?token=token
&criteria={"and":[{"md5","=", "8CB2289800B34EF1A5F472F7177348E2"}]}

```

Output:

```

{
  "scan_file_id": "6",
  "computer_id": "3",
  "path": "C:\Files",
  "name": "arch.exe",
  "full_path_shal": "9eb46086883dcbb6d2aefa2d6fd40bf8e736621b",
  "size": 2032,
  "version": null,
  "md5": "8cb2289800b34ef1a5f472f7177348e2",
  "sha256": null,
  "valid_from": "2015-12-10T11:37:14Z",
  "valid_to": "9999-12-31T23:59:59Z"
}

```

Example 3: Retrieving information about name, computer name, path, and MD5 columns for all **word.exe** files that have a different MD5 checksum than the specified one.

This example can be used if you know that a file should have a specific checksum. You can therefore retrieve all files with other, corrupted checksums.

Request:

```

https://hostname:port/api/sam/raw_file_facts?token=token
&columns[]=name&columns[]=computer.name&columns[]=path&columns[]=md5
&criteria={"and":[{"md5","!=", "8CB2289800B34EF1A5F472F7177348E2"}, {"name","=", "word.exe"}]}

```

Output:

```
{
  "path": "C:\\Programs",
  "name": "word.exe",
  "md5": "6cb1465800b34ef1a5f376f717743t6",
  "computer": {
    "name": "NC581058"
  }
}
```

Preview: Checking Common Vulnerabilities and Exposures (CVEs)

9.2.12 Available from 9.2.12. The software catalog contains information about Common Vulnerabilities and Exposures (CVEs). Browse the software catalog to check for any potential threats. Checking CVE in BigFix Inventory is a preview feature.

Common Vulnerabilities and Exposures (CVE) is a list of known security threats that are assigned identification numbers. BigFix Inventory uses CVE that is provided by the National Vulnerability Database at <https://nvd.nist.gov/> to help you identify potential threats in your environment.

Order of CVEs

Potential threats are displayed in the Vulnerability Risk (Preview) column. If matched, CVEs are sorted in descending order by the base score, and then the CVE identifier. They are not ordered by severity. The base score and severity are assigned according to the Common Vulnerability Scoring System (CVSS). When CVSS v3.0 is available, it takes precedence over CVSS v2.0.

CVE details

When you click the View Details icon  next to the CVE identification number, you are presented with the details of the relevant CVEs such as their names, severity and CVSS. If there are more CVEs matched with a particular component, you can view them on the detailed list. You can export the report view to CSV or PDF for additional processing. The exported report contains a full list of names of relevant CVEs.

Limitations

- CVEs that are listed in the National Vulnerability Database might impact software that is installed only on a particular operating system. BigFix Inventory does not take this fact into account while matching CVEs to components.
- If the name of a component or its publisher is different in BigFix Inventory and in the National Vulnerability Database, CVEs might not be matched in BigFix Inventory.
- If the detailed version of the component is significantly different from its version, CVEs might not be matched in BigFix Inventory.
- **9.2.14** The following aliases were added to improve the CPE generation and vulnerability matching:
 - RedHat alias for the Red Hat publisher.
 - Apache alias for the Apache Software Foundation publisher.

Log in to BigFix Inventory and open one of the following reports.

- **9.2.13** Go to **Reports > Software Classification**. To display the Vulnerability Risk (Preview) column, click the **Manage Report View** icon , click **Configure View**, and select the Vulnerability Risk (Preview) column to display it on the report.

CVEs on this report are matched with the particular software component through its detailed version.

<input type="checkbox"/>	Publisher...	Component Name	Component Ver...	Vulnerability Risk (Preview)	Product Name	Metric	Computer Name	Installation Path	Details																																												
<input checked="" type="checkbox"/>	Microsoft	Microsoft .Net Framework	1.0	CVE-2002-0369,...	Microsoft .NET Framework	Install Seats	NC9174039229	<See details>	DETAILS >																																												
				<table border="1"> <thead> <tr> <th>CVE Name</th> <th>CVSS</th> <th>CVSS Version</th> <th>Severity</th> </tr> </thead> <tbody> <tr><td>CVE-2002-0369</td><td>10.0</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2015-6108</td><td>9.3</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2015-2464</td><td>9.3</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2015-2463</td><td>9.3</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2015-2462</td><td>9.3</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2015-2460</td><td>9.3</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2015-2456</td><td>9.3</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2015-2455</td><td>9.3</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2014-0257</td><td>9.3</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2013-3132</td><td>9.3</td><td>2.0</td><td>High</td></tr> </tbody> </table>	CVE Name	CVSS	CVSS Version	Severity	CVE-2002-0369	10.0	2.0	High	CVE-2015-6108	9.3	2.0	High	CVE-2015-2464	9.3	2.0	High	CVE-2015-2463	9.3	2.0	High	CVE-2015-2462	9.3	2.0	High	CVE-2015-2460	9.3	2.0	High	CVE-2015-2456	9.3	2.0	High	CVE-2015-2455	9.3	2.0	High	CVE-2014-0257	9.3	2.0	High	CVE-2013-3132	9.3	2.0	High	Microsoft .NET Framework	Install Seats	NC9174039028	<See details>	DETAILS >
CVE Name	CVSS	CVSS Version	Severity																																																		
CVE-2002-0369	10.0	2.0	High																																																		
CVE-2015-6108	9.3	2.0	High																																																		
CVE-2015-2464	9.3	2.0	High																																																		
CVE-2015-2463	9.3	2.0	High																																																		
CVE-2015-2462	9.3	2.0	High																																																		
CVE-2015-2460	9.3	2.0	High																																																		
CVE-2015-2456	9.3	2.0	High																																																		
CVE-2015-2455	9.3	2.0	High																																																		
CVE-2014-0257	9.3	2.0	High																																																		
CVE-2013-3132	9.3	2.0	High																																																		
				CVE-2002-0369,...	Microsoft .NET Framework	Install Seats	NC9174039030	<See details>	DETAILS >																																												
				CVE-2002-0369,...	Microsoft .NET Framework	Install Seats	NC9174038048	<See details>	DETAILS >																																												
				CVE-2002-0369,...	Microsoft .NET Framework	Install Seats	NC9174039029	<See details>	DETAILS >																																												
				CVE-2015-6108,...	Microsoft .NET Framework	Install Seats	NC9174039028	<See details>	DETAILS >																																												
				CVE-2015-6108,...	Microsoft .NET Framework	Install Seats	NC9174039229	<See details>	DETAILS >																																												
				CVE-2015-6108,...	Microsoft .NET Framework	Install Seats	NC9174039030	<See details>	DETAILS >																																												
				CVE-2015-6108,...	Microsoft .NET Framework	Install Seats	NC9174038048	<See details>	DETAILS >																																												

- Go to **Reports > Software Components**. Common Vulnerabilities and Exposures are displayed in the Vulnerability Risk (Preview) column.

The report lists components in a particular version. However, CVEs that are matched relate to versions and their patches.

<input type="checkbox"/>	Publisher Name	Component Name	Component Version	End of Support	Vulnerability Risk (Preview)	Total Signatures																																												
<input checked="" type="checkbox"/>	Microsoft	Access 2000 for Micros...	2000.0	<not set>	CVE-2007-0065,...	Signatures: 1																																												
<input type="checkbox"/>	Microsoft	Access 2000 for Micros...	2000.0		<table border="1"> <thead> <tr> <th>CVE Name</th> <th>CVSS</th> <th>CVSS Version</th> <th>Severity</th> </tr> </thead> <tbody> <tr><td>CVE-2007-0065</td><td>10.0</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2003-0347</td><td>10.0</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2000-0854</td><td>10.0</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2000-0788</td><td>10.0</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2016-7182</td><td>9.8</td><td>3.0</td><td>Critical</td></tr> <tr><td>CVE-2015-6108</td><td>9.3</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2015-6107</td><td>9.3</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2015-6106</td><td>9.3</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2015-2510</td><td>9.3</td><td>2.0</td><td>High</td></tr> <tr><td>CVE-2015-2464</td><td>9.3</td><td>2.0</td><td>High</td></tr> </tbody> </table>	CVE Name	CVSS	CVSS Version	Severity	CVE-2007-0065	10.0	2.0	High	CVE-2003-0347	10.0	2.0	High	CVE-2000-0854	10.0	2.0	High	CVE-2000-0788	10.0	2.0	High	CVE-2016-7182	9.8	3.0	Critical	CVE-2015-6108	9.3	2.0	High	CVE-2015-6107	9.3	2.0	High	CVE-2015-6106	9.3	2.0	High	CVE-2015-2510	9.3	2.0	High	CVE-2015-2464	9.3	2.0	High	Signatures: 1
CVE Name	CVSS	CVSS Version	Severity																																															
CVE-2007-0065	10.0	2.0	High																																															
CVE-2003-0347	10.0	2.0	High																																															
CVE-2000-0854	10.0	2.0	High																																															
CVE-2000-0788	10.0	2.0	High																																															
CVE-2016-7182	9.8	3.0	Critical																																															
CVE-2015-6108	9.3	2.0	High																																															
CVE-2015-6107	9.3	2.0	High																																															
CVE-2015-6106	9.3	2.0	High																																															
CVE-2015-2510	9.3	2.0	High																																															
CVE-2015-2464	9.3	2.0	High																																															
<input type="checkbox"/>	Microsoft	Access 2000 for Micros...	2000.0		CVE-2007-0065,...	Signatures: 1																																												
<input type="checkbox"/>	Microsoft	AccessChk	6.1		CVE-2007-0065,...	Signatures: 1																																												
<input type="checkbox"/>	Microsoft	AccessChk	6.2		CVE-2007-0065,...	Signatures: 1																																												
<input type="checkbox"/>	Microsoft	AccessChk	6.0		CVE-2007-0065,...	Signatures: 1																																												
<input type="checkbox"/>	Microsoft	Account Lockout Status	1.0		CVE-2007-0065,...	Signatures: 1																																												
<input type="checkbox"/>	Microsoft	Active Accessibility Eve...	4.2		CVE-2007-0065,...	Signatures: 1																																												
<input type="checkbox"/>	Microsoft	Active Accessibility Obj...	7.2		CVE-2007-0065,...	Signatures: 1																																												



Note: **9.2.13** You can filter and sort both reports by CVE names.

- To show components for which any vulnerability was matched, specify the following filter:
Vulnerability Risk (Preview), is not empty.
- To show components for which a specific vulnerability was matched, specify the following filter:
Vulnerability: CVE Name, contains, and provide the name of the CVE.

For better performance, combine these two filters whenever you want to search for components for which a specific vulnerability was matched.

Updating information about Common Vulnerabilities and Exposures (CVE)

9.2.12

Available from 9.2.12. Common Vulnerabilities and Exposures (CVE) list is integrated with the software catalog and updated with each upgrade of BigFix Inventory. Use this procedure if you require more frequent updates.

- You must have access to the computer where the BigFix Inventory server is installed with permissions to paste files into the installation directory of BigFix Inventory.
- To perform the described actions, you must have the following tools installed:
 - GNU Wget
 - unzip tool
 - sha256sum

CVE contains a list of known security threats that are assigned identification numbers. Thanks to the import of CVE to BigFix Inventory, you can browse your software inventory and identify potential threats in your environment.

Information about CVE is automatically updated during every upgrade of BigFix Inventory. Use this procedure if you require to update CVE between the subsequent releases of BigFix Inventory.



Note: The procedure describes how to update CVE files with a semi-automated method. Alternatively, you can perform each step manually.

1. To download the **zip** CVE files, including **json** files and the relevant **meta** files from the [National Vulnerability Database](#), run the following command.

```
wget --secure-protocol=auto -r -A nvdCVE-1.0-20*.meta,nvdCVE-1.0-20*.json.zip
https://nvd.nist.gov/vuln/data-feeds
```

The files are downloaded to the following location: nvd.nist.gov/feeds/json/cve/1.0.



Important: Do not change names of the downloaded files.

2. **Optional:** Ensure that you downloaded the appropriate files.

- a. To unpack the packages that contain **json** files, run the following command.

```
unzip '*.zip'
```

- b. To generate SHA-256 for each **json** file, run the following command.

```
sha256sum file_name.json
```

- c. Compare SHA-256 from each **json** file with SHA-256 from the relevant **meta** file.
Both values should be the same.

3. Copy all **zip** CVE files to the following directory on the computer where the BigFix Inventory server is installed:

- **Linux** `installation_directory/cve_data`
- **Windows** `installation_directory\cve_data`

4. Wait for the scheduled import of data or run it manually.

The list of common vulnerabilities and exposures is populated to BigFix Inventory. To check information about CVE, go to **Reports > Software Components**. The vulnerabilities are listed in the Vulnerability Risk (Preview) column for each relevant component.

● (Base Report)		1 / 3650 rows (filtered)				
<input type="checkbox"/>	Publisher Name	Component Name	Component Version	End of Support	Vulnerability Risk (Preview)	Total Signatures
<input checked="" type="checkbox"/>	Microsoft	Access 2000 for Micros...	2000.0	<not set>	CVE-2007-0065,...	Signatures: 1
<input type="checkbox"/>	Microsoft	Access 2000 for Micros...	2000.0		CVE-2007-0065,...	Signatures: 1
<input type="checkbox"/>	Microsoft	Access 2000 for Micros...	2000.0		CVE-2007-0065,...	Signatures: 1
<input type="checkbox"/>	Microsoft	AccessChk	6.1			Signatures: 1
<input type="checkbox"/>	Microsoft	AccessChk	6.2			Signatures: 1
<input type="checkbox"/>	Microsoft	AccessChk	6.0			Signatures: 1
<input type="checkbox"/>	Microsoft	Account Lockout Status	1.0			Signatures: 1
<input type="checkbox"/>	Microsoft	Active Accessibility Eve...	4.2			Signatures: 1
<input type="checkbox"/>	Microsoft	Active Accessibility Obj...	7.2			Signatures: 1

CVE Name	CVSS	CVSS Version	Severity
CVE-2007-0065	10.0	2.0	High
CVE-2003-0347	10.0	2.0	High
CVE-2000-0854	10.0	2.0	High
CVE-2000-0788	10.0	2.0	High
CVE-2016-7182	9.8	3.0	Critical
CVE-2015-6108	9.3	2.0	High
CVE-2015-6107	9.3	2.0	High
CVE-2015-6106	9.3	2.0	High
CVE-2015-2510	9.3	2.0	High
CVE-2015-2464	9.3	2.0	High

Tutorials

Tutorials help you understand how to use BigFix Inventory. They consist of modules that focus on broad goals. Modules consist of tasks that show how to configure specific settings step-by-step.

Tutorial: Reporting subcapacity usage per computer group

9.2.2 Available from 9.2.2. In this tutorial, you will learn how to create computer groups that allow for managing software and generating audit snapshots for a subset of computers. You can use such groups to report PVU consumption for BigFix software installed in multiple BigFix subcapacity regions or to manage software per organizational business unit.

Managing software per computer group can be used in the following scenarios:

- You have BigFix software that is installed in multiple BigFix subcapacity regions. For more information, see: [BigFix subcapacity regions \(on page lxxxviii\)](#).
- You want to separately manage software that is installed in each business unit in your organization. For more information, see: [Organizational business units \(on page xciii\)](#).

Learning objectives

After you complete the lessons in this tutorial, you will know how to perform the following tasks:

- Divide the computers in your infrastructure into BigFix computer groups.
- Define the scope of audit snapshots by creating computer groups in BigFix Inventory.
- Import part numbers and manage software per computer group.
- Generate audit snapshots per computer group.

Time required

This module should take 45 minutes to complete. If you explore other concepts related to this tutorial, it can take longer to complete.

Audience

- Customers who have BigFix software installed in multiple BigFix subcapacity regions.
- Customers who want to manage software per organizational business unit.

Prerequisites

- Install the scanner on all computers that you want to monitor. For more information, see: [Installing the scanner \(on page ccii\)](#).
- Configure VM managers that control virtual machines in your environment. For more information, see: [Managing VM managers \(on page cccxviii\)](#).

Conventions used in this tutorial

This tutorial is based on the example of creating computer groups that reflect BigFix subcapacity regions.

Lessons in this tutorial

The tutorial contains the following lessons:

1. [Lesson 1 \(Optional\): Creating computer groups in the BigFix console \(on page dclxx\)](#)
2. [Lesson 2: Creating computer groups that define the scope of reports \(on page dclxxi\)](#)
3. [Lesson 3: Managing software and generating audit snapshots per computer group \(on page dclxxiii\)](#)

Lesson 1 (Optional): Creating computer groups in the BigFix console

9.2.2 Available from 9.2.2. You can divide the computers in your infrastructure into BigFix groups. They are used to manage the infrastructure, and run Fixlets against selected computers. They are not automatically copied to BigFix Inventory. However, they can be used as the basis for creating computer groups that define the scope of reports in BigFix Inventory.



Important: This lesson is optional. If you do not want to create separate scan schedules for each computer group, skip this lesson and proceed to the next one.

1. Log in to the BigFix console.
2. Create a computer group that contains all computers that are located in the first BigFix subcapacity region: North and South America.
 - a. In the top navigation bar, click **Tools > Create New Automatic Computer Group**.
 - b. Provide the name of the computer group.
For example, Americas.
 - c. Select the site and domain in which you want to create the computer group.
 - d. Specify criteria that must be met by the computers that belong to this group.
For example, add all computers whose DNS name contains `.us`.

- e. To save the group, click **Create**.
3. Create computer groups for the remaining BigFix subcapacity regions in which you have computers.
 4. For every computer group, schedule scans and uploads of scan results.
For more information, see: [Initiating software scans \(on page cciii\)](#) and [Uploading software scan results \(on page ccviii\)](#).

You created computer groups in the BigFix console. To generate audit snapshots for a subset of computers, create analogical computer groups in BigFix Inventory.

Lesson 2: Creating computer groups that define the scope of reports

9.2.2 Available from 9.2.2. To manage software and generate audit snapshots for subsets of computers, divide the computers in your infrastructure into groups in BigFix Inventory.

If you created computer groups in the BigFix console, run an import to ensure that information about these groups is available in BigFix Inventory.

1. Log in to BigFix Inventory.
2. Create a computer group that contains all computers that are in the first BigFix subcapacity region: North and South America.
 - a. In the top navigation bar, click **Management > Computer Groups**.
 - b. To create a computer group, click **New**.
 - c. Provide the name of the computer group.
For example, Americas.
 - d. **Optional:** Provide a description of the computer group.
 - e. Specify criteria that must be met by the computers that belong to this group.
To reuse the computer group that you created in the BigFix console, select **Data Source Groups, in set**, and select the proper BigFix computer group. In this case, select Americas.

! **Important:** Perform this step if you are creating computer groups to report PVU usage for BigFix software that is installed in multiple BigFix subcapacity regions. If you choose to calculate subcapacity data for the All Computers group, see: [BigFix subcapacity regions \(on page lxxxviii\)](#) to understand what is the difference between reports generated for each region separately and the report generated for the All Computers group.

Do not perform this step if you are creating computer groups to manage software that is installed in multiple business units. When you do not calculate subcapacity data for the All Computers group, you will not be able to generate audit snapshots for the entire environment.

5. To make the computer groups available in BigFix Inventory, wait for the scheduled import or run it manually.
6. **Optional:** If software in each computer group is managed by a separate person, create dedicated users, each with access to one computer group.

User Name	Roles	Authentication Method	Computer Group
admin	Administrators	Password	All Computers
americas_sam	Software Asset Managers, Infrastructure ...	Password	Americas
australia_sam	Software Asset Managers, Infrastructure ...	Password	Australia
europa_sam	Software Asset Managers, Infrastructure ...	Password	Europe and Africa

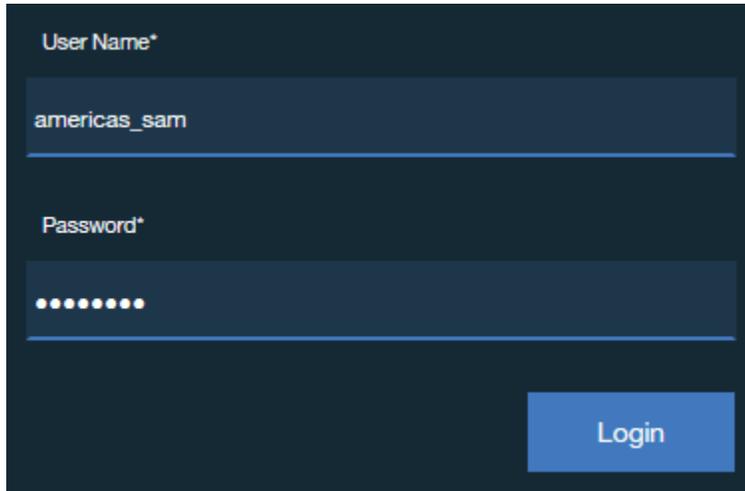
You created computer groups for which you can manage software and generate audit snapshots and dedicated users who can access these groups.

Lesson 3: Managing software and generating audit snapshots per computer group

9.2.2 Available from 9.2.2. After you create computer groups in BigFix Inventory, you can manage software and generate audit snapshot separately for each computer group.

1. Log in to BigFix Inventory.

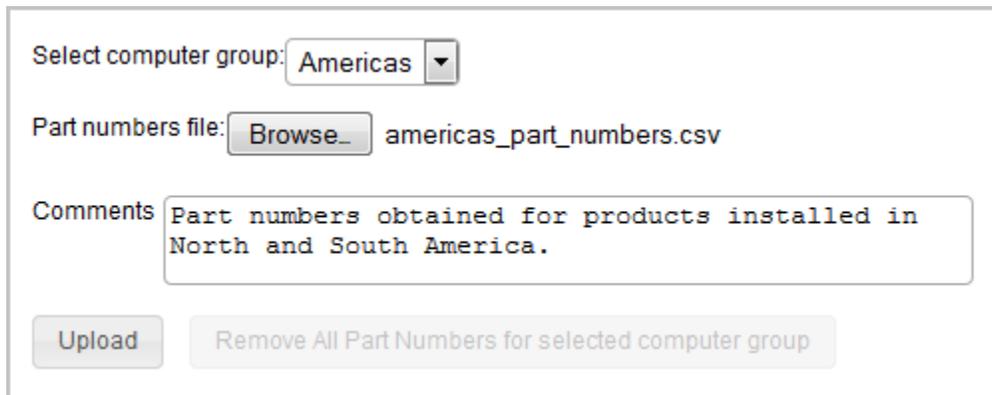
- If software in all computer groups is managed by one person, log in as the user that has access to all computer groups, for example the Administrator.
- If software in each group is managed by a different person, log in as the user with access only to the computer group that represents the first BigFix subcapacity region.



The screenshot shows a login form with a dark blue background. It has two input fields: "User Name*" with the text "americas_sam" and "Password*" with masked characters. A blue "Login" button is located at the bottom right.

2. If you have a file with part numbers related to the software that is installed in this region, upload the file.

- Click **Management > Part Numbers Upload**.
- Select the computer group for which you want to upload the part numbers. In this case, select Americas.
- Browse for the file with part numbers and provide a comment to the uploaded file. Then, click **Upload**.



The screenshot shows the "Part Numbers Upload" form. It includes a dropdown menu for "Select computer group:" set to "Americas", a "Part numbers file:" field with a "Browse..." button and the filename "americas_part_numbers.csv", and a "Comments" text area containing the text "Part numbers obtained for products installed in North and South America.". At the bottom, there are two buttons: "Upload" and "Remove All Part Numbers for selected computer group".

- To make the part numbers available in BigFix Inventory, wait for the scheduled import or run it manually.

3. To bundle the software, click **Reports > Software Classification**. The panel contains software from all computer groups to which the user that you used to log in has access. Assign each component to a product,

exclude, or suppress the instances that should not be included in the metric calculation, and confirm your actions. For more information, see: [Software classification \(on page dxxviii\)](#).

4. When you finish bundling the software, generate the audit snapshot for this region.
 - If you logged in as a user with access to multiple computer groups, open the Computer Groups report, and click the computer group that represents the first BigFix subcapacity region. Then, in the upper left corner, select All IBM Metrics or IBM PVU Subcapacity.



The report that opens contains information only about software that is installed on the computers in North and South America. Click **Audit Snapshot**. Provide a comment to the generated snapshot and upload any related files. Then, click **Generate**.

- If you logged in as a user with access only to the computer group that represents the first BigFix subcapacity region, open the All IBM Metrics or IBM PVU Subcapacity report, and click **Audit Snapshot**. Provide a comment to the generated snapshot and upload any related files. Then, click **Generate**

You generated the audit snapshot with information about PVU consumption in the first BigFix subcapacity region: North and South America.

Bundle software and generate audit snapshots for the remaining regions in which you have BigFix software installed. For every product, sum up the PVU values from all regions. The value that you obtain is the overall PVU usage for the particular product. This value might be different from the value that you obtain when you generate a single audit snapshot for the All Computers group. For more information, see: [BigFix subcapacity regions \(on page lxxviii\)](#).

Tutorial: Managing software in the service provider environment

9.2.2

Available from 9.2.2. In this tutorial, you will learn how to create computer groups that allow a service provider to manage software and generate audit snapshots per customer.

A service provider is a company that provides other organizations with access to software installed on machines that are owned and maintained by the service provider. There are three typical scenarios that are used in the service provider environment.

- Software that is installed on one virtual machine is used by one customer
- Some of the software that is installed on a virtual machine is used by the service provider
- Software that is installed on one virtual machine is used by many customers

For more information about the first two scenarios, see: [Service providers \(on page lxxvi\)](#). The last scenario is not supported by BigFix Inventory.

Learning objectives

After you complete the lessons in this tutorial, you will know how to perform the following tasks:

- Ensure that each customer can run scans and perform other actions in the BigFix console only on the computers on which the software that he uses is installed.
- Define the scope of audit snapshots by creating computer groups in BigFix Inventory.
- Import part numbers and manage software per computer group.
- Filter out software that is not used by the customer from reports.
- Generate audit snapshots per computer group.

Time required

This module should take 60 minutes to complete. If you explore other concepts related to this tutorial, it can take longer to complete.

Audience

Service providers who need to manage software and provide separate subcapacity reports for each customer.

Data visibility

Some of the items in BigFix Inventory are visible globally to all users. Other items are visible only for the user who has access to a particular computer group. If you grant your customers access to BigFix Inventory, review which items are visible globally and which are visible per computer group. For more information, see: [Service providers \(on page lxxxvi\)](#).

Prerequisites

- Install the scanner on all computers that you want to monitor. For more information, see: [Installing the scanner \(on page ccii\)](#).
- Configure VM managers that control virtual machines in your environment. For more information, see: [Managing VM managers \(on page cccxviii\)](#).

 **Tip:** If each customer should be allowed to view information about all VM managers in the infrastructure, use [basic VM management \(central\) \(on page cccxlvii\)](#). If each customer should be allowed to only view information about VM managers that control the virtual machines that he uses, use [advanced VM management \(distributed\) \(on page ccclxvii\)](#).

Conventions used in this tutorial

The tutorial is based on the example of a service provider whose customers are Bank ABC and IT Company.

Lessons in this tutorial

The tutorial contains the following lessons:

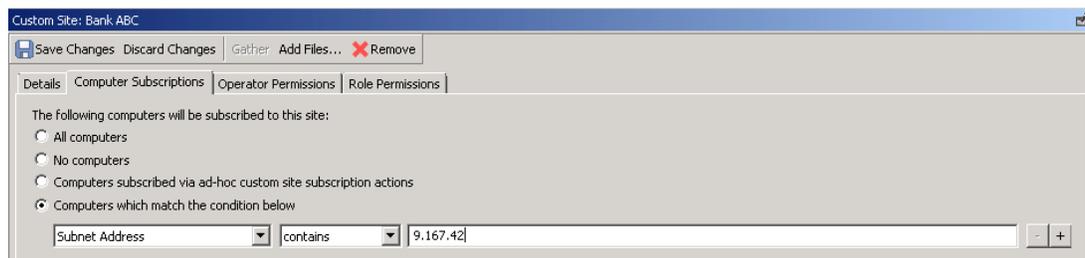
1. [Lesson 1 \(Optional\): Creating a site, a computer group, and an operator dedicated to each customer \(on page dclxxvii\)](#)
2. [Lesson 2: Creating computer groups to manage software per customer \(on page dclxxx\)](#)
3. [Lesson 3: Managing software and generating audit snapshots per customer \(on page dclxxxii\)](#)

Lesson 1 (Optional): Creating a site, a computer group, and an operator dedicated to each customer

9.2.2 Available from 9.2.2. If you grant your customers access to the BigFix console to allow them to run scans and perform other actions on computers, create a site, a computer group, and an operator that is dedicated to each customer. Then, provide the customer with credentials for the dedicated operator. This way, you ensure that each customer has access only to computers on which the software that he uses is installed.

! **Important:** This lesson is optional. If you do not grant you customers access to the BigFix console, you can skip this entire lesson and move to the next one. You can also create only computer groups to facilitate infrastructure and scan management as described in step 3. However, it is also optional.

1. Log in to the BigFix console.
2. Create a custom site that is dedicated to Bank ABC.
 - a. In the top navigation bar, click **Tools > Create Custom Site**.
 - b. Provide the name of the custom site. For example, Bank ABC. Then, click **OK**.
 - c. **Optional:** On the **Details** tab, provide the description of the site.
 - d. On the **Computer Subscriptions** tab, select **Computers which match the condition below**. Specify the condition that must be met by the computers that belong to this group. For example, choose all computers whose subnet address contains a particular string.



- e. Click **Save Changes**.
3. Create a computer group that contains the same computers as the site that you created in the previous step.
 - a. In the top navigation bar, click **Tools > Create New Automatic Computer Group**.
 - b. Provide the name of the computer group. For example, Bank ABC.

- c. From the **Create in site** drop-down list, select the site that you created in the previous step. In this case, select Bank ABC.
- d. Specify criteria that must be met by the computers that belong to this group. Choose the same criteria that you specified when you were creating the custom site.

Create Automatic Computer Group

Group name:

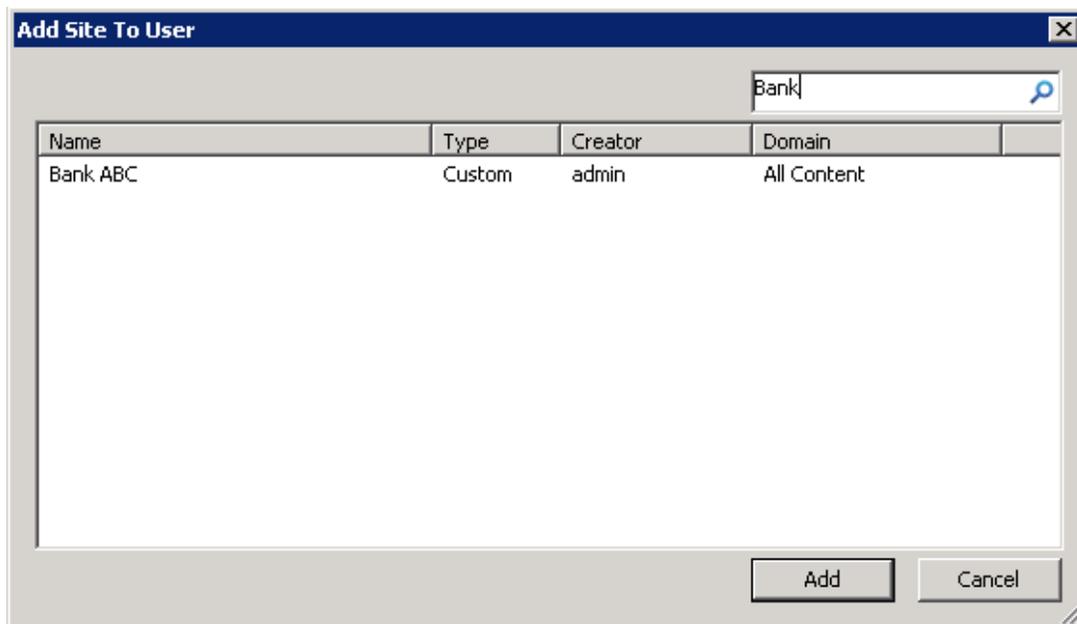
Create in site:

Create in domain:

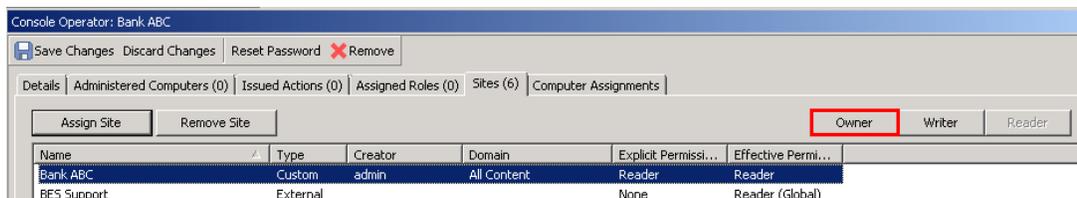
Include computers with the following property:

- +

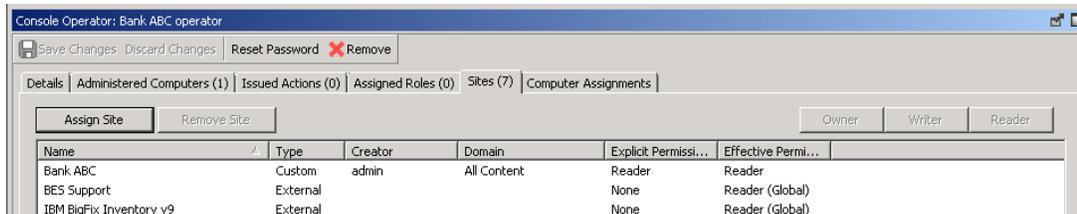
- e. Click **Create**.
4. Create a BigFix operator that is dedicated to Bank ABC. The operator has access only to computers on which the software that is used by this customer is installed.
 - a. In the top navigation bar, click **Tools > Create Operator**.
 - b. Provide the user name.
For example, Bank ABC.
 - c. Provide a password and confirm it. Then, click **OK**.
 - d. On the **Sites** tab, click **Assign Site**, and choose the custom site that is dedicated to this customer. In this case, choose Bank ABC. Then, click **Add**.



- e. To make operator the owner of the site, select the site and click **Owner**.

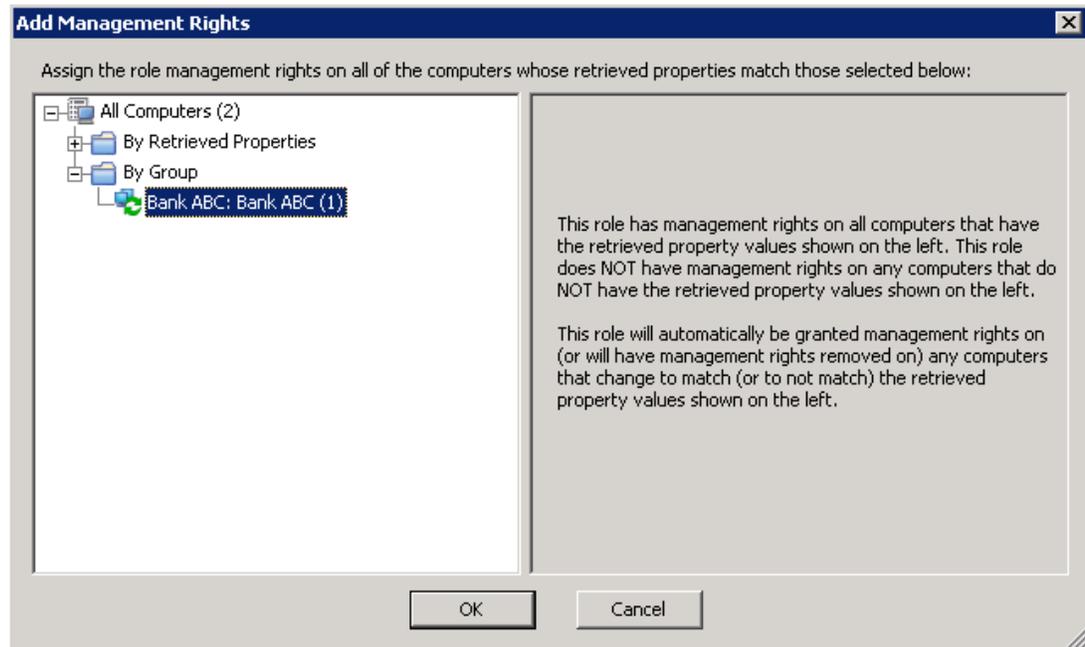


- f. Ensure that the customer does not have access to the sites dedicated to other customers. To remove access to a site, select the site and click **Remove Site**. Each customer should have access to his dedicated site, the BES Support site, and the BigFix Inventory site.



i **Tip:** If you cannot remove a site from the list of sites to which an operator has access, the site might have the read permission granted globally. To change the setting, open the site, go to the **Operator Permissions** tab, and clear the **Grant read permission globally** check box. Then, click **Save Changes**.

- g. On the **Computer Assignments** tab, click **Add**. Expand the By Group section, and select the computer group that is dedicated to Bank ABC. Then, click **OK**.



h. Click **Save Changes**.

You created a site, a computer group, and an operator dedicated to Bank ABC. You can now provide the customer with operator credentials. The customer can perform actions only against computers on which the software that he uses is installed.

Create sites, computer groups, and operators for the remaining customers.

Lesson 2: Creating computer groups to manage software per customer

9.2.2 Available from 9.2.2. To manage software and generate audit snapshots for each customer separately, create computer groups that represent the customers in BigFix Inventory.

If you created computer groups in the BigFix console, run an import to ensure that information about these groups is available in BigFix Inventory.

1. Log in to BigFix Inventory.
2. Create a computer group that contains all computers on which software that is used by Bank ABC is installed.
 - a. In the top navigation bar, click **Management > Computer Groups**.
 - b. To create a computer group, click **New**.
 - c. Provide the name of the computer group.
For example, Bank ABC.
 - d. **Optional:** Provide a description of the computer group.
 - e. Specify criteria that must be met by the computers that belong to this group.

If you created computer groups in the BigFix console, select **Data Source Groups, in set**, and select the proper BigFix computer group. In this case, select Bank ABC.

Create Computer Group

Parent

Name*

Description

Definition
 Specify the report filter which matches of the following conditions:

 **Restriction:** The size of data returned by the specified filter cannot exceed 1 MB.

f. **9.2.8** From the Type section, select **Reporting**.

g. Select license metrics for which you want to calculate utilization in this computer group.

Select license metrics for which you want to calculate utilization within this computer group. Disabling the calculations for metrics in which you are not interested can improve import performance, especially in larger environments.

PVU, RVU MAPC and VPC [?](#)

Install Seats, Install Instances, Microsoft Physical Core with SA, Microsoft Virtual Core with SA, Microsoft Single Processor, Microsoft Dual Processor, Oracle Processor Core and other metrics [?](#)

h. Specify the number of days for which the data is to be calculated.

This number defines the default period for which the data that is show on the All Metrics report is calculated and up-to-date.

Aggregate last days

For example, when you set the value to 30 days, and then display a report for the last 90 days, the report shows the following data:

- If you display the report for a newly created computer group, the report shows values for the last 30 days and is empty for the rest of the period.
- If you display the report for an existing group, the values might be outdated.

In both cases, recalculate the data.

- i. To save the group, click **Create**.
3. Create computer groups for the remaining customers.
 4. **Optional:** If you do not need a collective snapshot for all computers, but separate audit snapshots for each computer group, open the All Computers group, and clear **PVU and RVU MAPC**. Then, click **Save**.
 5. To make the computer groups available in BigFix Inventory, wait for the scheduled import or run it manually.
 6. **Optional:** If you give customers access to BigFix Inventory or when each customer account is managed by a separate person, create dedicated BigFix Inventory users, each with access to one computer group.

- a. To create a user, go to **Management > Users**, and click **New**.
- b. Select the role or roles that you want to assign to the user.

 **Restriction:** To ensure that each user has access to information only about computers that he uses, do not assign the user with the Administrator role. By default, this role has access to the All Computers group. Thus, a user who is assigned this role can view information about all computers in the infrastructure.

- c. Select the computer group to which you want the customer to have access.
- d. Choose the authentication method. If you choose Password, provide the password and confirm it.
- e. Click **Create**.

User Name	Roles	Authentication Method	Computer Group
admin	Administrators	Password	All Computers
Bank ABC	Software Asset Managers, Infrastructure ...	Password	Bank ABC
IT Company	Software Asset Managers, Infrastructure ...	Password	IT Company

You created computer groups, each dedicated to one customer, and users who can access these groups. You can manage software and generate audit snapshots for each customer separately. You can also provide your customers with credentials to their dedicated users and allow them to manage software and generate snapshots on their own.

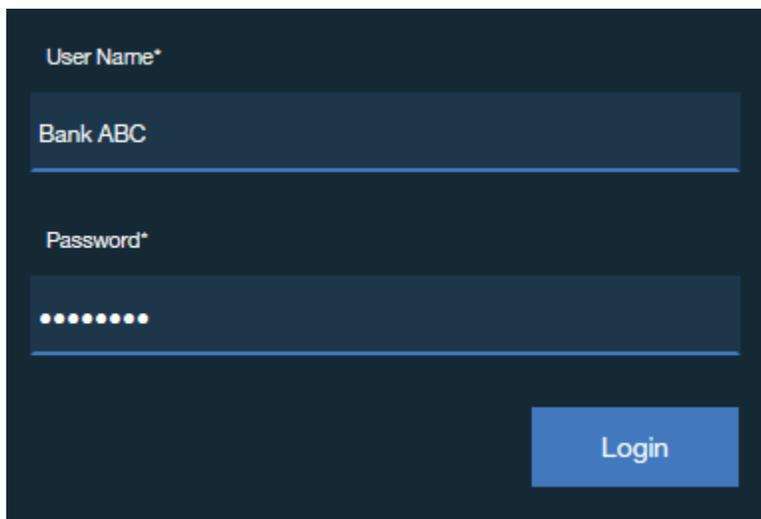
Lesson 3: Managing software and generating audit snapshots per customer

9.2.2 Available from 9.2.2. After you create computer groups in BigFix Inventory, you can manage software and generate audit snapshot separately for each customer.

Depending on your business model, you can perform these actions on your own or provide each customer with credentials to their dedicated user. In the latter case, customers can manage the software that they use and generate audit snapshots on their own.

1. Log in to BigFix Inventory.

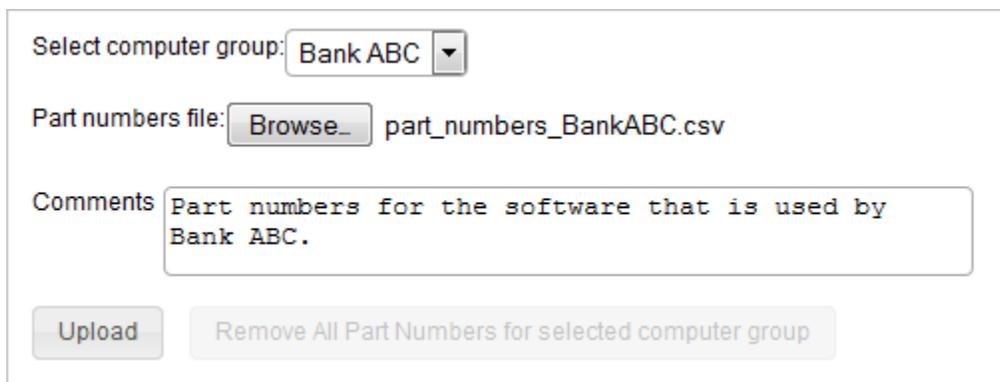
- If you manage customer software, log in as any user that has access to the computer group that represents Bank ABC, for example the Administrator.
- If you give customers access to BigFix Inventory or when each customer account is managed by a separate person, log in as the user with access only to the computer group that represents Bank ABC.



The screenshot shows a login form with a dark blue background. It features two input fields: 'User Name*' with the text 'Bank ABC' and 'Password*' with a masked password of seven dots. A blue 'Login' button is positioned at the bottom right of the form.

2. If you have a file with part numbers related to the software that is used by this customer, upload the file.

- Click **Management > Part Numbers Upload**.
- Select the computer group for which you want to upload the part numbers. In this case, select Bank ABC.
- Browse for the file with part numbers and provide a comment to the uploaded file. Then, click **Upload**.



The screenshot shows the 'Part Numbers Upload' form. It includes a dropdown menu for 'Select computer group' set to 'Bank ABC', a 'Part numbers file' field with a 'Browse...' button and the filename 'part_numbers_BankABC.csv', and a 'Comments' text area containing the text 'Part numbers for the software that is used by Bank ABC.'. At the bottom, there are two buttons: 'Upload' and 'Remove All Part Numbers for selected computer group'.

- To make the part numbers available in BigFix Inventory, wait for the scheduled import or run it manually.

3. To bundle the software, click **Reports > Software Classification**. The panel contains software from all computer groups to which the user that you used to log in has access. Assign each component to a product,

exclude, or suppress the instances that should not be included in the metric calculation, and confirm your actions. For more information, see: [Software classification \(on page dxxviii\)](#).

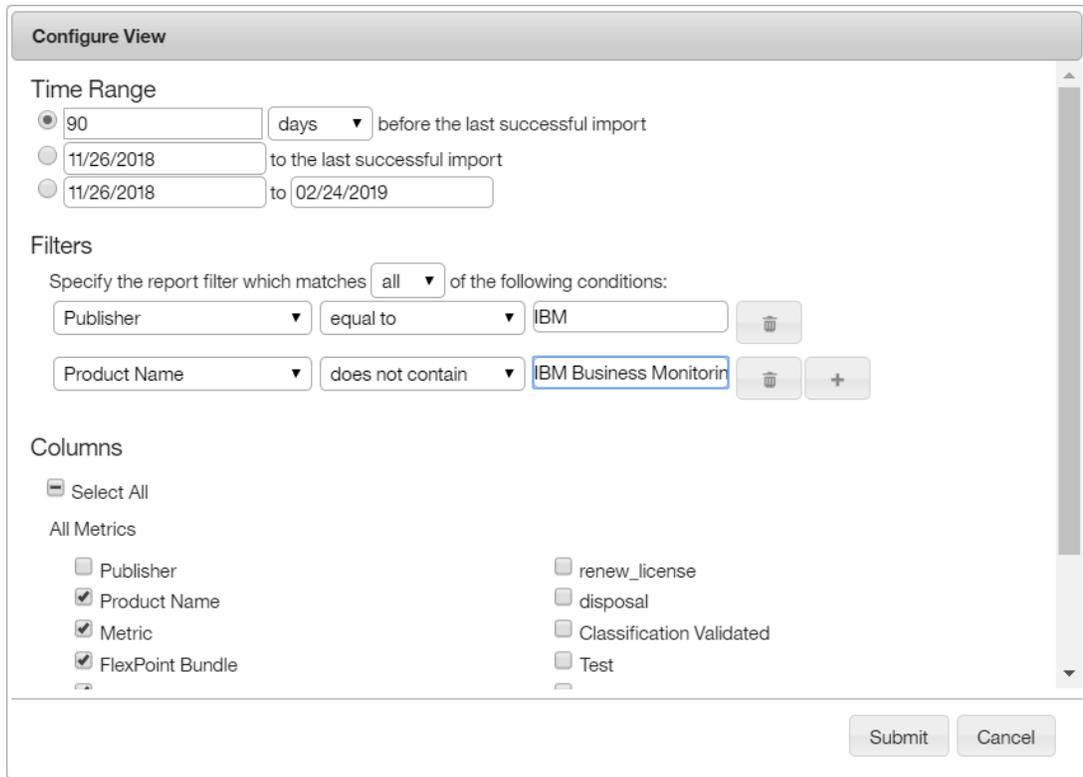
4. **Optional:** If some of the products are not used by the customer but by the service provider, filter out these products from the report to avoid charging the customer for products that he does not use.

 **Restriction:** All instances of the product that you filter out must be used by the service provider.

a. Open the All IBM Metrics or IBM PVU Subcapacity report.

b. To filter out the report, hover over the **Manage Report View** icon , and click **Configure View**. Then, add a filter to the filter that is already defined in this report.

For example, choose **Product, does not contain**, and provide the name of the product. Then, click **Submit**.



Configure View

Time Range

90 days before the last successful import

11/26/2018 to the last successful import

11/26/2018 to 02/24/2019

Filters

Specify the report filter which matches **all** of the following conditions:

Publisher equal to IBM

Product Name does not contain IBM Business Monitorin

Columns

Select All

All Metrics

Publisher

Product Name

Metric

FlexPoint Bundle

renew_license

disposal

Classification Validated

Test

Submit Cancel

c. **Optional:** You can set the filtered report as the default report view for this customer. To do this, hover over the **Manage Report View** icon , and click **Save As**. Then, provide the name of the report, and select **Set as default**. Then, click **Create**.

5. Generate the audit snapshot for the customer.

- If you logged in as a user with access to multiple computer groups, open the Computer Groups report, and click the computer group that represents Bank ABC. Then, in the upper left corner, select All IBM Metrics or IBM PVU Subcapacity.



The report that opens contains information only about software that is installed on the computers that are used by Bank ABC. Click **Audit Snapshot**. Provide a comment to the generated snapshot and upload any related files. Then, click **Generate**.

- If you logged in as a user with access only to the computer group that represents Bank ABC, open the All IBM Metrics or IBM PVU Subcapacity report, and click **Audit Snapshot**. Provide a comment to the generated snapshot and upload any related files. Then, click **Generate**.

You generated the audit snapshot with information about license consumption in Bank ABC.

Tutorial: Migrating software assignments between two BigFix servers

9.2.14 Available from 9.2.14 This tutorial teaches you how to migrate the software assignments for each computer that is moved to a new BigFix server using the REST API requests.

Learning objectives

In this tutorial you learn how to perform the following tasks:

- Retrieving information about the software assignments for the computer that is moved to the new BigFix server.
- Changing the Data Source Computer ID, performing test, and migrating the software assignments of the computer which is connected to the new BigFix server.

Time required

20 minutes. If you explore other concepts that are related to this tutorial, it might take longer to complete.

The order of lessons

This tutorial contains two lessons. Complete them in order.

Lessons in this tutorial

1. [Lesson 1: Retrieving information about software assignments for a specified computer \(on page dclxxxvi\)](#)
2. [Lesson 2: Performing tests and migrating the assignments \(on page dclxxxviii\)](#)

Lesson 1: Retrieving information about software assignments for a specified computer

9.2.14 Available from 9.2.14 This lesson shows you how to use the REST API request to retrieve software assignments of a computer that is moved to the new BigFix server.

 You must have the View Endpoints and View Hardware Inventory permissions to perform this task.

This tutorial is based on the old and new Data Source Computer ID of the computer that is migrated to the new BigFix server. This information is crucial for the REST API requests. To migrate software assignments for multiple computers, start with preparing a list of old and new Data Source Computer IDs for all affected computers. Then, for best results, migrate the software assignments for each computer in your environment separately. The number of software instances per single API request should not exceed 150.

1. Move the computer to the new BigFix server.
2. [Add a new data source to BigFix Inventory \(on page cdxxxviii\)](#).
3. [Upload software scan results \(on page ccviii\)](#) from the computer that is connected to the new BigFix server.
4. Collect a pair of old and new Data Source Computer IDs for the computer.
 - a. Log in to BigFix Inventory.
 - b. Go to **Reports > Computers**.
 - c. Hover over the **Manage Report View** icon , and click **Configure View**. Select **Data Source Name** and **Data Source Computer ID** from the list of columns, and click **Submit**.
 - d. Sort the report by **Computer Name** by clicking on the column header. By looking at the Computer Name and the Data Source Name prepare a pair of old and new Data Source Computer IDs of the computer. The following example is used in this tutorial.

- Data Source Computer ID that is assigned to the old data source on the report: **9596634**
- Data Source Computer ID that is assigned to the newly created data source on the report: **778906**

 **Tip:** If you plan to migrate software assignments for multiple computers, start with preparing a list of old and new Data Source Computer IDs for all affected computers to save you time and effort.

5. To retrieve information about software assignments related to this computer, run the following API query:

```
https://hostname:port/api/sam/v2/software_instances?limit=150&

columns[]=product_name&columns[]=component_name&columns[]=discovery_path&columns[]=is_charged&

columns[]=is_confirmed&columns[]=computer_bigfix_id&columns[]=discoverable_guid&columns[]=product_rele
ease_guid&
```

```
columns[]=metric_id&token=user_token&criteria={"and":[{"is_present","=","1"},
["computer_bigfix_id","=","9596634"]]}
```

The **computer_bigfix_id** is the same as the old Data Source Computer ID. In the example, the old Data Source Computer ID is **9596634**.

The following example shows the results of the query for this computer.

```
{
  "total": 2,
  "rows": [
    {
      "product_name": "BigFix Inventory",
      "component_name": "BigFix Inventory Server",
      "is_charged": 0,
      "is_confirmed": 1,
      "computer_bigfix_id": 9596634,
      "discoverable_guid": "cdf19da9-a9de-4ee9-ab41-ff09fa6eda92",
      "product_release_guid": "1ecbbb8d-1d1e-4beb-a40f-cb9b9b9462fb",
      "metric_id": -13103
    },
    {
      "product_name": "DB2 Enterprise Server Edition Unlicensed Product Base",
      "component_name": "DB2 Enterprise Server Edition Unlicensed Product Base",
      "is_charged": 1,
      "is_confirmed": 0,
      "computer_bigfix_id": 9596634,
      "discoverable_guid": "57628995-cf33-4335-984d-d2c8abde750e",
      "product_release_guid": "befa759d-0473-47bd-ae60-f1e16c13e8cc",
      "metric_id": 1
    }
  ]
}
```

- Copy the results of the query to the `Body` of the request, and change the value of the `computer_bigfix_id` to the new Data Source Computer ID.

The following example shows the modifications that are required to migrate software assignments to the computer on the new server. The new Data Source Computer ID is **778906**.

```
{
  "total": 2,
  "rows": [
    {
      "product_name": "BigFix Inventory",
```

```

    "component_name": "BigFix Inventory Server",
    "is_charged": 0,
    "is_confirmed": 1,
    "computer_bigfix_id": 778906,
    "discoverable_guid": "cdf19da9-a9de-4ee9-ab41-ff09fa6eda92",
    "product_release_guid": "1ecbbb8d-1dle-4beb-a40f-cb9b9b9462fb",
    "metric_id": -13103
  },
  {
    "product_name": "DB2 Enterprise Server Edition Unlicensed Product Base",
    "component_name": "DB2 Enterprise Server Edition Unlicensed Product Base",
    "is_charged": 1,
    "is_confirmed": 0,
    "computer_bigfix_id": 778906,
    "discoverable_guid": "57628995-cf33-4335-984d-d2c8abde750e",
    "product_release_guid": "befa759d-0473-47bd-ae60-f1e16c13e8cc",
    "metric_id": 1
  }
]
}

```

Lesson 2: Performing tests and migrating the assignments

9.2.14

Available from 9.2.14 This lesson shows you how to run a test and complete the migration of software assignments for a computer that is moved to the new BigFix server.



You must have the View Endpoints and View Hardware Inventory permissions to perform this task.

The **PUT** request needs to be made using a dedicated API tool, such as Postman or cURL.

1. Use the **Body** of the request that you prepared in [Step 6 at the end of Lesson 1 \(on page dclxxxvii\)](#). To run a test and simulate the migration, use the following **PUT** request.

```
https://hostname:port/api/sam/v2/software_instances?token=user_token&simulate=true&verbose=true
```

2. Check the results of the test.

The following sample shows the example of test results.

```

{
  "details": {
    "valid_instances": {
      "Bundled": {
        "1": "instance_ids: 314"
      }
    }
  }
}

```

```

    }
  },
  "invalid_instances": {},
  "unmodified_instances": {
    "1": "Bundling 'BigFix DB2 Enterprise Server Edition Unlicensed Product Base
10.5' to 'BigFix DB2 Enterprise Server Edition Unlicensed Product Base 10.5' on '9589066'"
  }
},
"summary": {
  "valid_instances": {
    "Bundled": 1
  },
  "invalid_instances": {},
  "unmodified_instances": 1
}
}

```

3. If the test is successful, migrate the software assignments with the following **PUT** request.

```
https://hostname:port/api/sam/v2/software_instances?token=user_token
```

The results of the API request should include a summary of changes that were made. The following example shows the sample of the results:

```

{
  "valid_instances": {
    "Bundled": 1
  },
  "invalid_instances": {},
  "unmodified_instances": 1
}

```

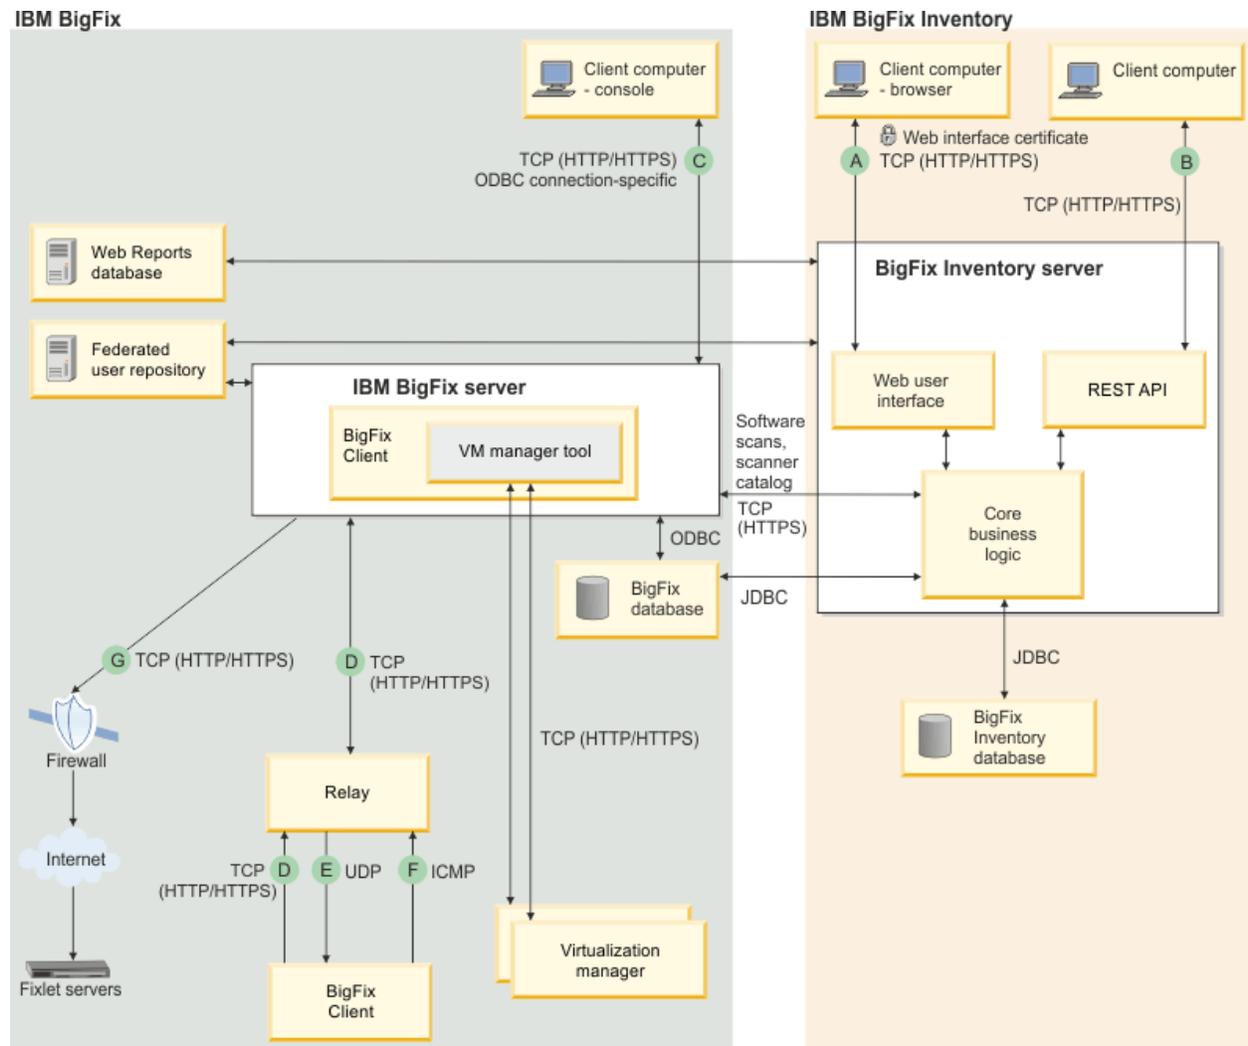
After migrating the software assignments with the API request, the data is updated in BigFix Inventory and reflected on Software Classification panel. After the next import of data, the license metric information is recalculated according to the changes.

Security

Configure different security features to adequately protect business assets and resources in the data model when using BigFix Inventory.

Flow of data

There are several different interactions that occur between the components of the BigFix Inventory infrastructure and between the user and tool.



BigFix Inventory domain

This table describes the flow of data for the BigFix Inventory domain. There are four columns and two rows. Each row is divided in to two rows from column three.

Inter- ac- tion	Type	Connection	Description
A	Web browser data traffic	Port	By default, the web browser connects to the BigFix Inventory server using port 9081 (HTTPS). You can disable the SSL/TLS connection tunnelling.
		Origination	The web browser connects to the BigFix Inventory server.
B	REST API data traffic	Port	By default, the web browser connects to the BigFix Inventory server using port 9081 (HTTPS). You can disable secure connection.
		Origination	A client that uses REST API connections.

BigFix domain

This table describes the flow of data for the BigFix domain. There are four columns and two rows. Each row is divided in to two rows from column three.

Inter- ac- tion	Type	Connection	Description
C	BigFix Console data traffic	Port	Consoles connect to root server using HTTPS 52311 for all interactions
		Origination	The BigFix console connects to the RootServer service.
		Network controls:	There is a "refresh rate" for each BigFix console user (default 15 seconds)
D	Gather, post, download	Port	Port 52311 is configurable by the BigFix administrator at installation time.
		Origination	The BigFix client initiates the request to the BigFix relay or server.
		Network controls:	<ul style="list-style-type: none"> • Configurable bandwidth throttling to BigFix relay or clients • Configurable gather interval. The default is 1 per day per fixlet site. • Configurable minimum time to wait between posts. The default is 15 seconds.

**Inter-
ac-
tion**

	Type	Connection	Description
			<ul style="list-style-type: none"> • Configurable temporal distribution (spread out downloads over time) per action • The ability to set "policy" to prevent computers from downloading files if they are not pointed at the proper BigFix relay
E	UDP "new information" message	Port	Port 52311 is configurable by the BigFix administrator at installation time.
		Origination	The UDP messages are sent from the BigFix clients' immediate "parent", which can be either a BigFix relay or server.
		Network controls:	<ul style="list-style-type: none"> • Configurable limit of the number of UDP messages sent at one time from a BigFix relay • Configurable limit of the amount of time to wait after sending UDP messages from a BigFix relay
F	Relay selection	Port	The ICMP protocol does not use a port.
		Origination	Each BigFix client sends progressive "rounds" of ICMP packets to each relay with increasing TTLs until a BigFix relay responds. For example, in a network of 2 relays, one 1 hop away and one 2 hops away, the BigFix client sends an ICMP message to both with TTL 1 and receives 2 "time exceeded" messages from the local router. The BigFix client then sends an ICMP message to both relays with TTL 2 and receives one "time exceeded" message and one reply message. The BigFix client then chooses the relay that is one hop away.
		Network controls	<ul style="list-style-type: none"> • Relay auto-selection can be disabled. • Configurable interval for when the BigFix clients perform auto-selection • Configurable limit on the maximum number of ICMP packets to send out in a time interval • Configurable limit on the maximum number of "rounds" to send out during relay auto-selection
G	New data download from external BigFix fixlet servers	Port	80 (typically); possibly 21, 443
		Origination	The BigFix server connects to the BigFix fixlet servers

Inter- ac- tion	Type	Connection	Description
		Network controls	There is a configurable interval that the BigFix server checks for new fixlet messages.

The following database protocols are used:

- ODBC
- JDBC

Security configuration scenarios

Check what security options need to be enabled on the BigFix server and the BigFix Inventory server to achieve each of the supported security scenarios.

TLS 1.2

To use TLS 1.2 protocol, perform the following configuration:

- [Enable enhanced security](#) on the BigFix server
- [Enable HTTPS \(on page dcxciii\)](#) in BigFix Inventory

FIPS 140-2

To achieve compliance with the FIPS 140-2 standard, perform the following configuration:

- [Enable enhanced security and SHA-256 downloads](#) on the BigFix server
- [Enable HTTPS \(on page dcxciii\)](#) in BigFix Inventory
- [Enable FIPS 140-2 \(on page dccii\)](#) in BigFix Inventory

For more information about the standard, see: [Federal Information Processing Standard 140-2 \(on page dccii\)](#).

SP800-131

To achieve compliance with SP800-131 standard, perform the following configuration:

- [Enable HTTPS \(on page dcxciii\)](#) in BigFix Inventory
- [Enable SP800-131 compliance \(on page dcciv\)](#) in BigFix Inventory

For more information about the standard, see: [SP800-131 compliance \(on page dcciii\)](#).

Configuring secure communication

To ensure secure communication, BigFix Inventory uses public key cryptography, which is based on algorithms that use two separate keys, a private key and a public key. This key pair is used to encrypt and decrypt communication.

The private key encrypts communication. The public key, which is contained in a certificate, decrypts communication. The use of encrypted communication requires that you create both a private key and a certificate that is associated with it. You can share the public key (certificate) with anyone because it is used only to read the communication. The safety of your communication depends mainly on your private key that proves your identity, and must therefore be securely stored. The keys are created in such a way that a message encrypted with the private key can be decrypted only with the public key that is associated with it. If someone receives your public key and can decrypt your communication with this key, they know for certain that you are the originator of the message and that it was not tampered with on the way. Otherwise, the public key would be unable to decrypt it.

BigFix Inventory provides self-signed certificates by default, but they are not intended for production environments. To improve security, create your own private key and a certificate signing request (CSR) that can be transformed into a certificate after it is signed by a certificate authority (CA). By signing your request, a CA approves your public key and certifies that it can be trusted. You can create your own private CA, use the CA of your organization, or an internationally trusted CA, such as Entrust, VeriSign, or other.

The private key and the associated certificate are uploaded to BigFix Inventory. After enabling the encrypted communication, anyone who connects to your server receives a certificate that contains your public key. All successive communication that originates from the server is encrypted with your private key. After a user receives the communication, it is decrypted with the certificate that they obtained from the server. If the certificate can decrypt the communication, it is known for certain that the server is the originator of the message and that it is valid.

Key pair requirements

Your key pair must meet the following requirements to be accepted by BigFix Inventory.

- Type: RSA or DSA.
- Key strength: maximum 2048 bits. This limitation is caused by IBM Java policy. You can use stronger keys if you substitute default policy files with the unlimited jurisdiction ones. For more information, see: [IBM SDK Policy files](#).
- Format: PEM-encoded. Such an encoding is ensured if you create the key pair by using openssl. You can also create your keys by using other methods, for example Makecert on Windows. Such keys are DER-encoded and therefore not supported by BigFix Inventory. However, you can convert other formats to PEM, for example by using openssl.
- Private key format: PKCS#8 (used by openssl). The pvk format is not supported.

Limitations

Key pair that is generated for BigFix Inventory can be used for Web Reports only if the private key is not password-protected.

Structure and format of the private key and certificate

- **Private key format**

- PEM-encoded without password protection. Ensure that the private key (`private.key`) is enclosed between the following statements:

```
-----BEGIN PRIVATE KEY-----
<<base64 stringfrom private.key>>
-----END PRIVATE KEY-----
```

- PEM-encoded with password protection. Ensure that the private key (`private.key`) is enclosed between the following statements:

```
-----BEGIN ENCRYPTED PRIVATE KEY-----
<<base64 stringfrom private.key>>
-----END ENCRYPTED PRIVATE KEY-----
```

- **X509 certificate format**

PEM-encoded. If you received the intermediate and root certificates as separate files, combine them into a single file. For example, if you have the primary certificate file (`certificate.crt`) and the intermediate certificate file (`ca_intermediate.crt`), combine them in the following order.

```
BEGIN CERTIFICATE-----
<<primary certificate: base64 stringfrom certificate.crt>>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<<intermediate certificate: base64 stringfrom ca_intermediate.crt>>
-----END CERTIFICATE-----
```

If you received the root certificate (`ca_root.crt`) in addition to the intermediate certificate, combine them in the following order.

```
BEGIN CERTIFICATE-----
<<primary certificate: base64 stringfrom certificate.crt>>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<<intermediate certificate: base64 stringfrom ca_intermediate.crt>>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<<root certificate: base64 stringfrom ca_root.crt>>
-----END CERTIFICATE-----
```

- **Single file (private key with certificates) format**

PEM-encoded. This file can contain the private key and the primary certificate, or the private key and the chain of certificates combined in the following order:

- Private key and primary certificate.

```
-----BEGIN CERTIFICATE-----
<<primary certificate: certificate.crt>>
-----END CERTIFICATE-----
-----BEGIN PRIVATE KEY-----
<<private key: base64 stringfrom private.key>>
-----END PRIVATE KEY-----
```

- Private key, primary certificate, and intermediate certificate.

```
BEGIN CERTIFICATE-----
<<primary certificate: base64 stringfrom certificate.crt>>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<<intermediate certificate: base64 stringfrom ca_intermediate.crt>>
-----END CERTIFICATE-----
-----BEGIN PRIVATE KEY-----
<<private key: base64 stringfrom private.key>>
-----END PRIVATE KEY-----
```

- Private key, primary certificate, intermediate certificate, and root certificate.

```
BEGIN CERTIFICATE-----
<<primary certificate: base64 stringfrom certificate.crt>>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<<intermediate certificate: base64 stringfrom ca_intermediate.crt>>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<<root certificate: base64 stringfrom ca_root.crt>>
-----END CERTIFICATE-----
-----BEGIN PRIVATE KEY-----
<<private key: base64 stringfrom private.key>>
-----END PRIVATE KEY-----
```

Procedure

Complete the following steps to create your key pair and to enable encrypted communication. If you already have a key pair or want to use the self-signed certificates, you can skip to enabling secure communication.

Step 1: Creating private keys and certificates

To improve security, create your own private key and a certificate instead of using the self-signed ones that are available in BigFix Inventory by default. You can use openssl to create a private key and a certificate signing request (CSR) that can be transformed into a certificate after it is signed by a certificate authority (CA).

This procedure is valid for all operating systems that support openssl.

If you are generating an encrypted private key in the pkcs8 format, add the following line to the `installation_dir/jre/lib/security/java.security` file:

```
security.provider.10=org.bouncycastle.jce.provider.BouncyCastleProvider
```

Then, restart the BigFix Inventory server.

1. Open the command line.
2. Create a new private key.

```
openssl genrsa -des3 -out key_name.key key_strength -sha256
```

For example, `openssl genrsa -des3 -out private_key.key 2048 -sha256`

Where:

-des3

Enables password for the private key. This is an optional parameter. You can also enable password for an existing private key by using the following command:

```
openssl rsa -des3 -in path_to_private_key.key -out key_name.key
```

key_name

File name for your new private key.

key_strength

Key strength, measured in bits. The maximum value that you can use for BigFix Inventory is 2048 bits.

3. Create a certificate signing request (CSR). The request is associated with your private key, and is later transformed into a certificate.

```
openssl req -new -key path_to_private_key.key -out csr_name.csr
```

For example, `openssl req -new -key private_key.key -out CSR.csr`

Where:

path_to_private_key

Path to your private key.

csr_name

File name for your certificate signing request (CSR).

After you run the command, you are asked to provide information that helps your users to identify your certificate and ensure that it can be trusted. The following excerpt from the command line is filled in with sample information:

```
Country Name (2 letter code) [XX]: US
State or Province Name (full name) []: New York
Locality Name (eg, city) [Default City]: New York
Organization Name (eg, company) [Default Company Ltd]: HCL (eg, section) []: Software
Common Name (eg, your name or your server's hostname) []: inventory.bigfix.com
Email Address []: inventory@bigfix.com
```

After completing these steps, two files are created, your private key (`.key`) and the certificate signing request (`.csr`).

Sign the request to transform it into the certificate. For information about how to create a private certificate authority (CA) to sign the request, see [Signing certificates \(on page dcxcviii\)](#).

Step 2: Signing certificates

Your certificate signing request (CSR) must be signed by a certificate authority (CA) to be transformed into a certificate that can be uploaded to BigFix Inventory. You can use the openssl cryptographic library to create a private CA and sign your request.

Using a private CA to sign your request is not the only way. You can also send the request to internationally trusted CAs, such as Entrust, VeriSign, and so on, or use the CA of your organization. The certificates of these CAs are often trusted by default and do not display any warnings in the browser. Warnings might be displayed if you use a private CA.

1. Create a private certificate authority (CA) and a certificate for it.
 - a. Create a private CA. This step creates a private key (`.key`) and a request (`.csr`) similar to those that you created in [Creating private keys and certificates \(on page dcxcvii\)](#).

```
openssl req -new -newkey rsa:key_strength -nodes
-out CA_csr_name.csr -keyout CA_key_name.key -sha256
```

For example, `openssl req -new -newkey rsa:2048 -nodes -out CA_CSR.csr -keyout CA_private_key.key -sha256`

Where:

key_strength

Key strength, measured in bits. The maximum value that you can use for BigFix Inventory is 2048 bits.

CA_csr_name

File name for the certificate signing request (CSR). The certificate authority (CA) requires a separate request.

CA_key_name

File name for the private key. The certificate authority (CA) requires a separate private key.

- b. Create a certificate for your private CA. This step creates a certificate (`.arm`) that you can use to sign your CSR.

```
openssl x509 -signkey path_to_CA_key.key -days
number_of_days -req -in path_to_CA_csr.csr
-out CA_certificate_name.arm -sha256
```

For example, `openssl x509 -signkey CA_private_key.key -days 90 -req -in CA_CSR.csr -out CA_certificate.arm -sha256`

Where:

key_strength

Key strength, measured in bits. The maximum value that you can use for BigFix Inventory is 2048 bits.

path_to_CA_csr

File name for the certificate signing request (CSR) that you created for the certificate authority (CA).

path_to_CA_key

File name for the private key that you created for the certificate authority (CA).

number_of_days

Number of days for the new certificate to be valid.

CA_certificate_name

File name for the certificate of your CA. This certificate is used to sign your CSR.

2. Use the CA certificate to sign the certificate signing request that you created in [Creating private keys and certificates \(on page dcxcvii\)](#).

```
openssl x509 -req -days number_of_days -in path_to_csr.csr -CA path_to_CA_certificate.arm
-CAkey path_to_CA_key.key -out new_certificate.arm -set_serial 01 -sha256
```

For example, `openssl x509 -req -days 90 -in CSR.csr -CA CA_certificate.arm -CAkey CA_private_key.key -out certificate.arm -set_serial 01 -sha256`

Where:

number_of_days

Number of days for the new certificate to be valid.

path_to_csr

Path to certificate signing request (CSR) that you want to sign.

path_to_CA_certificate

Path to certificate that you created for the certificate authority (CA).

path_to_CA_key

Path to the private key that you created for the certificate authority (CA).

new_certificate

File name for the new certificate that is created from your certificate signing request (CSR). You upload this certificate together with your private key to BigFix Inventory.

You signed your certificate signing request and obtained a new certificate.

Enable encrypted communication in BigFix Inventory and upload your private key and the certificate. These files replace the self-signed certificate that is already available in BigFix Inventory, and thus ensure secure communication. For more information, see: [Step 3: Enabling secure communication \(on page dcc\)](#).

Step 3: Enabling secure communication

You can enable encrypted communication (HTTPS) to ensure secure communication between your server and all users that access it. You can base your communication on self-signed certificates that are provided by default in BigFix Inventory, but these certificates are not intended for production environments. To improve security, create your own private key and certificate, and upload them to BigFix Inventory.

- The use of HTTPS is enabled by default, but this configuration is based on temporary self-signed certificates that are not intended for production environments.
- Enabling or disabling the use of HTTPS changes the web address of your BigFix Inventory server. Ensure that you run a data import afterward to update the address in the Fixlets that use it to download files from the server.

1. Log in to BigFix Inventory.
2. In the top navigation bar, click **Management > Server Settings**.
3. Select **Use HTTPS**. The Certificate subsection opens.
4. **Optional:** Select **Use TLSv1.2**.

**Important:**

- Enabling TLS 1.2 disables TLS 1.0.
- To use TLS 1.2, ensure that your browser supports TLS 1.2, and that it is enabled.
- To fulfill all the requirements for SP800-131 compliance, see: [Enabling SP800-131 compliance \(on page dcciv\)](#).

5. Provide information about the certificate.

- If you have a private key and a certificate:
 - a. Select **Import a PEM encoded private key and certificate**.
 - b. Click **Browse** to locate the files in the computer file system.
 - c. In the **Private key password** field, enter the password for the key. This field is required only if you set a password for your private key.
 - d. Click **Save**.



Note: The certificate and the key must be PEM-encoded.

- If you want to generate a new self-signed certificate:



Restriction: A self-signed certificate contains a public key, information about the owner of the certificate, and the owner's signature. Because such a certificate is signed by its own private key, it does not provide means to verify the origin of the certificate through a trusted certificate authority.

- a. Select **Generate a self-signed certificate**.
- b. Specify the certificate subject *common name*. The common name must correspond to the DNS name of the BigFix Inventory server.
- c. In the **Expiration Date** field, enter the date when the certificate expires.
- d. Click **Save**.



Note: Most browsers display a warning message when a self-signed certificate is used.

6. Restart the server.

You enabled secure communication on your server. All outgoing communication is now encrypted with the private key that you provided.

Related information

[SP800-131 compliance \(on page dcciii\)](#)

[Enabling SP800-131 compliance \(on page dcciv\)](#)

Enabling TLS 1.2 in IBM Java

By default, IBM Java has TLS 1.0 enabled. To enable TLS 1.2, change the value of the **Dcom.ibm.jsse2.overrideDefaultTLS** parameter to true in the `jvm.options` file.

9.2.13

Starting from application update 9.2.13, TLS 1.2 is enabled in IBM Java by default.

1. Go to the computer where the BigFix Inventory server is installed.
2. Go to one of the following directories and open the `jvm.options` file.
 - **Windows** `C:\Program Files\ibm\BFI\wlp\usr\servers\server1\`
 - **Linux** `/opt/ibm/BFI/wlp/usr/servers/server1/`
3. Set the value of the `-Dcom.ibm.jsse2.overrideDefaultTLS` to true.

```
-Dcom.ibm.jsse2.overrideDefaultTLS=true
```

4. Restart the BigFix Inventory server.
 - a. [Stop the server.](#) (on page *cd*)
 - b. [Start the server.](#) (on page *cccxcix*)

Assuring compliance with federal encryption standards

You can configure BigFix Inventory to be compliant with the Federal Information Processing Standard requirements that are related to encryption.

Federal Information Processing Standard 140-2

Federal Information Processing Standards (FIPS) are standards and guidelines that are issued by the National Institute of Standards and Technology (NIST) for federal government computer systems.

Government agencies and financial institutions use Federal Information Processing Standard (FIPS) to ensure that the products conform to specified security requirements. For more information about these standards, see the [NIST website](#).

FIPS 140-2 is the standard that defines the security requirements for cryptographic modules that are used within a system that handles sensitive but unclassified information. Compliance with the FIPS 140-2 standard has two aspects that affect BigFix Inventory: the algorithms that are used to manage sensitive data must be FIPS-approved and a FIPS-approved implementation must be used when data is transmitted with the SSL/TLS.

BigFix Inventory uses the FIPS 140-2 approved cryptographic providers for cryptography:

- IBMJCEFIPS (certificate 376)
- IBMJSSEFIPS (certificate 409)
- IBM Crypto for C (ICC) (certificate 384)

The certificates are listed on the [NIST web site](#).

Configuring the server to achieve FIPS compliance

You can assure compliance with the FIPS 140-2 standard by modifying the configuration properties for the underlying application server.

1. Edit your `java.security` file that is in the following directory: `installation_dir/jre/lib/security/`. Put the `com.ibm.crypto.fips.provider.IBMJCEFIPS` before the `IBMJCE` one in the provider list. Ensure that the list is correctly numbered.
2. Add the `-Dcom.ibm.jsse2.usefipsprovider=true` property to the `jvm.options` file. The property allows the Java™ Secure Socket Extension (JSSE2) provider to run in FIPS 140-2 mode.



Note: Your certificates must have a key that is at least 1024 bits long and can be signed with a DSA or RSA signature algorithm. You can use the IBM keytool utility to generate a compatible key pair.

3. To use the TLS protocol, [configure secure communication \(on page dcxciii\)](#).

A number of ciphers are supported by FIPS 140-2. The default HTTPS configuration automatically enables the FIPS 140-2 compliant ciphers when JSSE is running in FIPS mode. You can enable specific ciphers by listing them in the `enabledCiphers` attribute of the SSL service configuration element in the `server.xml` file.

SP800-131 compliance

SP800-131 requires longer key lengths and stronger cryptography. The specification also provides a transition configuration to enable users to move to a strict enforcement of SP800-131.

The transition configuration also enables users to run with a mixture of settings from both FIPS140-2 and SP800-131. SP800-131 can be run in two modes, transition and strict. The transition mode is offered to give you a setting to move your environment to SP800-131 strict mode. In transition mode, it is optional to use the SP800-131 required certificates and to set the protocol to SP800-131.

The following requirements must be fulfilled to allow for the strict enforcement of SP800-131:

- The use of the TLS version 1.2 protocol for the Secure Sockets Layer (SSL) context.
- Certificates must have a minimum length of 2048 bytes. An Elliptic Curve (EC) certificate requires a minimum size of 244-bit curves.
- Certificates must be signed with a signature algorithm of SHA256, SHA384, or SHA512. Valid signature algorithms include:
 - SHA256 with RSA
 - SHA384 with RSA
 - SHA512 with RSA
 - SHA256 with ECDSA
 - SHA384 with ECDSA
 - SHA512 with ECDSA
- SP800-131 approved cipher suites.

For more information about the SP800-131 standard, see the [web site](#) run by National Institute of Standards and Technology.

Related information

[Step 3: Enabling secure communication \(on page dcc\)](#)

[Enabling SP800-131 compliance \(on page dcciv\)](#)

Enabling SP800-131 compliance

You can set up a BigFix Inventory profile to meet the SP800-131 requirement that is originated by the National Institute of Standards and Technology (NIST).

You can configure BigFix Inventory to run in SP800-131 strict or transition mode.

- To configure the product to run in *strict mode*:

1. Ensure that your server certificates meet the criteria for SP800-131.

For more information about SP800-131, see the [National Institute of Standards and Technology Special Publication 800-131A](#).

2. [Modify your HTTPS configuration \(on page dcxciii\)](#) to use the TLS version 1.2 protocol.
3. Enable the Java Secure Socket Extension (JSSE) to run in SP800-131 strict mode: set the system property **com.ibm.jsse2.sp800-131** to *strict*. The property must be set in the `jvm.options` file, which is in the `installation_dir/wlp/usr/servers/server1` directory.

Example:

```
-Dcom.ibm.jsse2.sp800-131=strict
```



Note: If your server certificates do not meet the criteria for SP800-131 or if the TLS version 1.2 protocol is not used, then after you restart the server you are not able to connect to BigFix Inventory. In this event, you can remove the **com.ibm.jsse2.sp800-131** property from the `jvm.options` file, or set the property to *transition*.

- To configure the product to run in *transition mode*, enable JSSE to run in SP800-131 transition mode by setting the system property **com.ibm.jsse2.sp800-131** to *transition*. The property must be set in the `jvm.options` file, which is in the `installation_dir/wlp/usr/servers/server1` directory.

Example:

```
-Dcom.ibm.jsse2.sp800-131=transition
```

Related information

[SP800-131 compliance \(on page dccciii\)](#)

[Step 3: Enabling secure communication \(on page dcc\)](#)

Authenticating users with LDAP

BigFix Inventory supports authentication through a Lightweight Directory Access Protocol server. To use this feature, you must configure the BigFix Inventory server.

Configuring connection to a directory server

To use LDAP for authentication of BigFix Inventory users, you must first configure a connection to your directory server.

-  You must have the Manage Directory Servers permission to perform this task.
1. In the top navigation bar, click **Management > Directory Servers**.
 2. To create an LDAP connection, click **New**.
 3. Enter a name for the new directory service.
 4. In the LDAP server list, select the type of your LDAP server. If your LDAP server values are different from the defaults, select **Other** and enter the values of filters and attributes of your LDAP server. If you select Microsoft Active Directory **Global Catalog**, the Search Base field is disabled.

 **Important:** The default values might need to be modified in particular for openLDAP servers due to various implementations of openLDAP.

5. Type the name of Search Base. This parameter defines the location in the directory from which the LDAP search begins.
6. If your directory server uses Secure Socket Layer protocol, select the **SSL** check box.
7. If your server requires authentication, clear **Anonymous bind** and provide a name and a password for the user whose credentials are to be used for connecting to the directory server.

 **Tip:** If you selected Microsoft Active Directory, provide the user name as Active Directory logon name or User Principal Name, for example `username@domain.com`. Do not specify the user name in the following way: `DOMAIN/username`.

8. In the **Host** text field, provide the host name or IP address of your primary LDAP server.
9. Accept the default port value or provide a new one.
10. **Optional:** To add a backup server:
 - a. Click **add backup server**.
 - b. Provide its host name or IP address and the port number.

11. To verify whether all of the provided entries are valid, click **Test Connection**.
A confirmation pop-up window opens.
12. Click **Create**. A confirmation message is displayed in the middle of the page.

You configured a connection to your LDAP server.

Related information

[User provisioning \(on page dcccix\)](#)

Editing a directory server configuration

1. On the **Directory Servers** page, click the name of the directory server whose configuration you want to modify.
2. In the lower area of the window, enter the new parameters.
3. Click **Save**.

Deleting a directory server configuration

1. On the **Directory Servers** page, click the name of the directory server whose configuration you want to delete.
2. In the upper left area of the window, click **Delete**.

Configuring a directory server with load balancer or multiple domain controllers

If your LDAP server uses load balancer or multiple domain controllers that dynamically change the list of hosts, and the connection between LDAP and the BigFix Inventory server is secure, perform advanced configuration of the BigFix Inventory server. The LDAP server can be already configured in BigFix Inventory.

The `<install_dir>` is the directory where the BigFix Inventory server is installed. By default, it is the following path.

Linux	/opt/ibm/BFI
Windows	C:\Program Files\ibm\BFI

Substitute `<install_dir>` in the procedure with the directory where the BigFix Inventory server is installed in your environment.

- **9.2.9** For application update 9.2.9 or higher, perform the following steps to permanently configure the directory server.

1. Contact your LDAP server administrator to obtain a public certificate from the certificate authority (CA) that issues server certificates for LDAP. The certificate file should be in the PEM format, and should contain one or more certificates (full chain of trust).

For example:

```
-----BEGIN CERTIFICATE-----
MI IHZ jCCBk6gAwIBAgISKESJLWXAAAACtanBgkqhkiG9w0BAQUFADBNMRMwEQYK
```

```

CRWmyVBwPWQBUNd1lPKJRQwpeYKCZImiZPyLGQBGRYEQ354jTEgGG7GA1UEAiU5
.
.
.
MTAzMzQxWjBZMRMwEQYKCZImiZPJVGQBGRYDbmV0MRkwFwYKCZImiZPyLGQBGRYJ
bnNyY290ZGV2MScwJQYDVQQDEx5DaXRXAEludGVybmFsIERldmljZSBDQSawMyBM
-----END CERTIFICATE-----

```

- To create a custom truststore and import the public certificate for the LDAP server to that truststore, run the following command.

```

<install_dir>/jre/jre/bin/keytool -import -file <path_to_certificate> -alias ldapCA
-keystore <install_dir>/wlp/usr/servers/server1/resources/security/ldap_truststore
-storepass <password>

```

Where:

-file

Is the path to the public certificate of the LDAP server.

-storepass

Is the password that you want to set for the truststore.

- After you run the command, confirm it by typing `yes` in the following line.

```
Trust this certificate? [no]: yes
```

- To encode the truststore password, perform the following steps.

- Set the `JAVA_HOME` variable.

- **Linux** `export JAVA_HOME=<install_dir>/jre/jre`
- **Windows** `set JAVA_HOME=<install_dir>/jre/jre`

- Run the following command.

```
<install_dir>/wlp/bin/securityUtility encode --encoding=aes
```

- When prompted, enter and re-enter the password that you specified in step 2 (on page dcccvi).
- Save the encoded password. You will need it to complete this procedure.

- Back up the `server.xml` file before you make any changes. Then, provide information about the location of the truststore and its password in the `server.xml` file.

- Open the `server.xml` file that is in the `<install_dir>/wlp/usr/servers/server1` directory.
- Locate the `keyStore` entry that has the `defaultKeyStore` identifier.

```
<keyStore id='defaultKeyStore'>
```

- c. Paste the following lines after that entry. In the **password** parameter, provide the encoded password from step 3 (on page dccvii).

```
<keyStore id='ldapCustom'
location='<install_dir>/wlp/usr/servers/server1/resources/security/ldap_truststore'
password='{aes}xxxxXXXxxxxXXXxxxxXXXxxxx' />
```

- For versions older than application update 9.2.9, perform the following steps to configure the directory server. This solution must be applied after each application upgrade, as the certificates that you add according to this procedure are not preserved.



Note: Before you begin, contact the IBM Support for the password that is required during configuration.

1. Contact your LDAP server administrator, and obtain the public certificate from certificate authority (CA) that issues the server certificates for LDAP. The certificate file should be in the PEM format, and should contain one or more certificates.

For example:

```
-----BEGIN CERTIFICATE-----
      MI IHZ jCCBk6gAwIBAqISKESJLWXAAAACTANBgkqhkiG9w0BAQUFADBNMRMwEQYK
      CRWmyVBwPWQBUNDi1PKJRQwpeYKZImiZPyLGQBGRYEQ354jTEgGG7GA1UEAiU5
      .
      .
      .
      MTAzMzQxWjBZMRMwEQYKZImiZPJVQBGGRYDmV0MRkwFwYKZImiZPyLGQBGRYJ
      bnNyb290ZGV2MScwJQYDVQQDEx5DaXRXAEludGVybmFsIERldm1jZSBDQSAwMyBM
      -----END CERTIFICATE-----
```

2. Copy the certificate file to the following directory: `<install_dir>/jre/jre/lib/security/`.
3. Run the following command:

```
<install_dir>/jre/jre/bin/keytool -import -trustcacerts -file <certificate_file_name>
-alias certAliasName -keystore cacerts -storepass <password>
```

Where `<password>` is the password that was provided by the BigFix Support.

Linking users to directories

To complete an authentication process through LDAP, you must create a user that would link to the created directory.



You must have the Manage Users permission to perform this task.

1. In the top navigation bar, click **Management > Users**.
2. To create a user, click **New**.
3. In the **User Name** field, type the name of an existing user of an LDAP server.

 **Tip:** The login format depends on the configuration of your environment. It can be, for example, an e-mail address, or a user ID, and it might need to be followed by a domain name.

4. From the list, select a **Computer Group** to which the user would be assigned.
5. From the **Authentication Method** list, select the name of an LDAP directory.
6. Click **Create**.
7. To delete the created user, click its name. Then, in the upper left of the window click **Delete**.

 **Note:** The deleted user cannot be re-created.

To confirm authentication, log in to the BigFix Inventory server with the credentials of the LDAP user that you created in BigFix Inventory.

User provisioning

After configuring a Directory Server, you can integrate its users with BigFix Inventory using the LDAP protocol. With user provisioning, you can integrate whole groups of users instead of linking each of them separately.

[Configure a Directory Server \(on page dccv\)](#).

1. Log in to BigFix Inventory.
2. In the top navigation bar, click **Management > User Provisioning**.
3. Click **New** to create a new user provisioning rule.
4. In **Group Names**, start typing the name of a group of users from your Directory Server and then choose one of the autofilled results.
5. In **Roles**, select the roles that will be assigned to the new users.
6. In **Computer Group**, select the computer group to which the new users will be assigned.
7. Click **Create**.

You integrated your Directory Server users with BigFix Inventory. Each user is created in BigFix Inventory after you log in to the application with that user. The user is assigned a role that is defined in user provisioning. However, the role can be changed later on.

 **Important:** If you want to log in with Active Directory users, the user name must be followed by a domain name, for example `username@domain.com`.

Related information

[Configuring connection to a directory server \(on page dccv\)](#)

Integrating users with Web Reports

You can use the Web Reports component to allow your Lightweight Directory Access Protocol (LDAP) and Web Reports users to access BigFix Inventory.

Install the Web Reports component. The component is typically installed together with your BigFix server but you can also add it to your environment at any time. To do so, start the installation of BigFix and choose to install only Web Reports.

One of the Web Reports capabilities is integrating with an LDAP directory. This integration allows you to view information about your LDAP users through Web Reports and to grant them right privileges to access your BigFix environment. If you create an entry for each user in BigFix Inventory, they will be linked between the application and Web Reports. The linked users can then access BigFix Inventory with the same credentials that are specified in Web Reports. Whenever you change the credentials in Web Reports, they will also be valid in BigFix Inventory with no additional configuration.

1. Connect your BigFix Inventory server to the Web Reports database.
 - a. Log in to BigFix Inventory.
 - b. In the navigation bar, click **Management > Data Sources**.

- c. Click on your data source and fill in the connection parameters for the Web Reports database. The required information will differ depending on the type of the database that you use. For more information, see the following examples.

Web Reports Database

Database Type*

Host*

Database Name*

Authentication
 Windows Authentication
 SQL Server Authentication

User Name*

Password*

Web Reports Database

Database Type*

Host*

Port*

Database Name*

Authentication

User Name

Password*

2. Each of your Web Reports users must be manually added to BigFix Inventory. After you complete this action, the users will be linked with their equivalents in Web Reports:
 - a. In BigFix Inventory, click **Management > Users**.
 - b. Click **New** to create a new user.
 - c. Enter the user name that corresponds with a Web Reports user name.
 - d. Select the appropriate roles.
 The roles are not integrated between the applications and must be selected manually for each user.
 - e. In the Authentication Method, choose **Web Reports**.
 - f. Click **Create**.
 - g. Repeat this action for each of your Web Reports users.

The created user is linked with its equivalent in Web Reports. You can now use it to log in to BigFix Inventory by using the same password that is specified in Web Reports. Whenever you change this password in Web Reports, it will also be valid for logging in to BigFix Inventory.

Configuring and enabling single sign-on

Available from 9.2.1. You can now use the two-factor authentication and use single sign-on to log on to BigFix Inventory and maintain login consistency with other applications in the enterprise. You can configure BigFix Inventory to use two-factor authentication with single sign-on based either on the exchange of Security Assertion Markup Language (SAML 2.0) token and Microsoft™ Active Directory Federation Services as Identity Provider or you can use the IBM Lightweight Third-Party Authentication (LTPA) technology and IBM Security Access Manager for Web as the authentication service.

To enable debug logging for single sign-on in BigFix Inventory, edit the `web.xml` file and change the value of `config.sso.debug` to `true`.

```
<context-param>
  <param-name>config.sso.debug</param-name>
  <param-value>true</param-value>
</context-param>
```

The solution described in this section is based on the assumption that the connection with BigFix Inventory is established via the BigFix Inventory host name. For complex scenarios, you need to manually configure SAML provider in the `server.xml` file and perform additional configuration of the authentication service.

Option 1: Configuring single sign-on based on Security Assertion Markup Language token

You can configure single sign-on based on a Security Access Markup Language (SAML 2.0) token and an external Identity Provider server.

The following entities participate in the exchange of authentication data:

Application User

A person who uses several applications in one domain and wants to single sign-on to these applications.

Service Provider

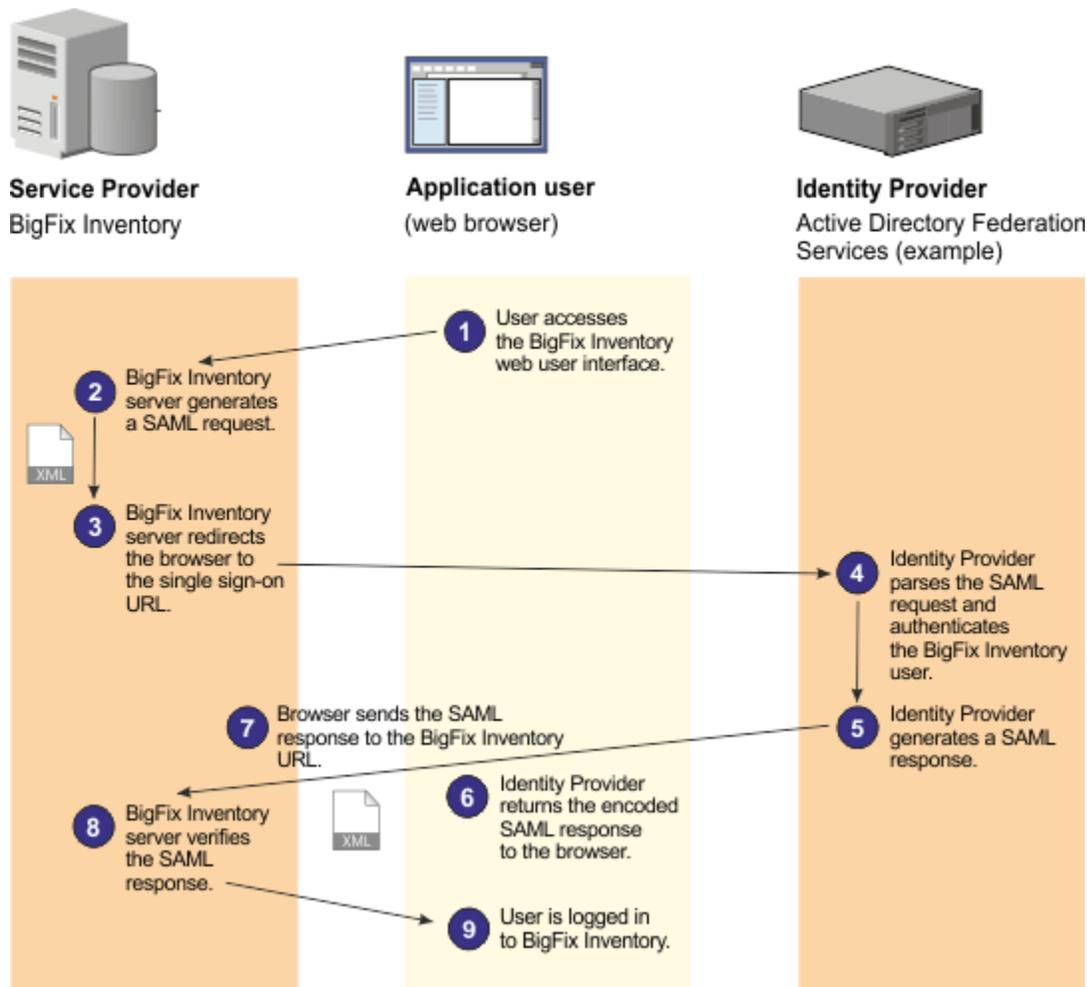
The application that requests the authentication service. In this case, it is BigFix Inventory.

Identity Provider

The service that authenticates the application users. BigFix Inventory supports only Active Directory Federation Services.

A web user authenticates to a SAML Identity Provider, which produces a SAML assertion. A SAML assertion is an XML-formatted token that is used to transfer user identity and attribute information from the Identity Provider of a user to a trusted Service Provider as part of a single sign-on request. The Service Provider consumes the SAML assertion to establish a security context for the web user.

The following diagram shows steps that are performed during a typical single sign-on based on the exchange of the SAML token.



Procedure

The following scenario presents a typical workflow of configuring BigFix Inventory to work with Active Directory Federation Services. However, you might want to use other software products for enabling single sign-on in your infrastructure.

Related information

[SAML 2.0 Web Browser Single-Sign-On](#)

[Configuring SAML Web Browser SSO in the Liberty profile](#)

Step 1: Configuring single sign-on settings in BigFix Inventory

As the first step, configure single sign-on settings in BigFix Inventory.

Gather necessary information

Before you start the configuration, gather the following information:

- URL to the login page of the Identity Provider. It is the URL to which an unauthenticated request is redirected. After the request is authenticated by the Identity Provider, the user is redirected to BigFix Inventory.

For example: `https://ADFS_host_name/adfs/ls/IdPInitiatedSignOn.aspx?LoginToRP=https://BFI_host_name:9081/ibm/saml20/defaultSP`.

- URL of the Trusted Issuer. It is the URL to the certificate issuer of the Identity Provider that is needed to establish a trust relationship.

For example, `http://ADFS_host_name/adfs/services/trust`.

- Public certificate of the Identity Provider in the `key_name.cer` format.

Enable SSL

Ensure that SSL is enabled in BigFix Inventory and in the Identity Provider.

Backup files

Before you start configuring single sign-on, back up the following files:

- `server.xml`

Linux `bfi_install_dir/wlp/usr/servers/server1`

Windows `bfi_install_dir\wlp\usr\servers\server1`

- `web.xml`

Linux `bfi_install_dir/wlp/usr/servers/server1/apps/tema.war/`

`WEB-INF`

Windows `bfi_install_dir\wlp\usr\servers\server1\apps\tema.war`

`\WEB-INF`



Note: If you set up the session timeout for Single Sign-On, remember that it should be longer than the session timeout that is set up for BigFix Inventory. Otherwise, change the settings in BigFix Inventory. For more information, see: [Setting session timeout \(on page cdi\)](#).

Create users

Create BigFix Inventory users (on page clxxxix) who will use the single sign-on. During the creation of the users, select **Single Sign-on** as the authentication method. Ensure that all user names are fully-qualified names that contain the full domain name, for example: `user@domain.example`. Also, ensure that at least one user is an Administrator.

Linux If the BigFix Inventory server is installed on Linux, and users in the Identity Provider use the camel-case naming convention, create users following the same convention in BigFix Inventory. Otherwise, the users are not be able to generate audit snapshots.

 **Note:** User token is not available after a single sign-on user is created. If you need the token, for example to run REST API calls, ask the BigFix Inventory administrator to provide it for you.

1. Log in to BigFix Inventory, and click **Management > Single Sign-On Settings**.
2. Select **SAML** as the single sign-on method.

The Instance ID field is automatically filled with the `defaultSP` value. It is the identifier of the BigFix Inventory service. Together with the BigFix Inventory URL, it forms the overall Service Provider ID:

`https://BFI_host_name:BFI_port/ibm/saml20/defaultSP`.

Based on this value, the SAML Assertion Consumer Service URL is built:

`https://BFI_host_name:BFI_port/ibm/saml20/defaultSP/acs`. The URL should be used for the configuration of the Identity Provider.

3. Specify the URL to the login page of the Identity Provider that you will use to single sign-on to BigFix Inventory. For example:

```
https://ADFS_host_name/adfs/ls/IdPInitiatedSignOn.aspx?LoginToRP=https://BFI_host_name:9081/ibm/saml20/defaultSP
```

 **Important:** Ensure that the URL that you specify is correct. The address is not validated. If you make a typo in the URL, you might need to manually [revert the SSO configuration \(on page dccxviii\)](#).

4. Provide the public certificate of the Identity Provider. Click **Browse** to locate the `key_name.cer` certificate that you created.
5. Provide the URL of the certificate issuer of the Identity Provider. It is the issuer name of the Identity Provider as it appears in the SAML assertion.

For example:

```
http://ADFS_host_name/adfs/services/trust
```

 **Important:** Ensure that the URL that you specify is correct. The address is not validated. If you make a typo in the URL, you might need to manually [revert the SSO configuration \(on page dccxviii\)](#).

6. Click **Save**.
7. **Optional:** To use a custom certificate for the SSO setup, see: [Using a CA-signed \(custom\) certificate for SSO based on SAML \(on page dccxvi\)](#). Otherwise, continue to the next step.

- Click the **Download Service Provider Metadata** link, and save the `spMetadata.xml` file.



Note: When the SAML single sign-on entry is created, only the **Delete** button, and the **Download SP Metadata** link are enabled. If the download link is not displayed, restart the BigFix Inventory server.

Based on the `spMetadata.xml` file, [configure Identity Provider for single sign-on \(on page dccxvii\)](#).

Using a CA-signed (custom) certificate for SSO based on SAML

By default, a self-signed certificate is used during the SSO configuration. However, you can use a custom certificate generated for the BigFix Inventory server to increase security of the configuration.

- Log in to the computer where Active Directory Federation Services are installed.
- Generate a certificate for the BigFix Inventory sever signed by a trusted CA.



Important: Ensure that you remember the certificate label that is used during certificate generation as it is needed in further steps.

- Export the certificate into a `.pfx` file. For example, `custom_cert.pfx`.
- Copy the `custom_cert.pfx` file to the computer where the BigFix Inventory server is installed and place it in the following location: `install_dir\wlp\usr\servers\server1\resources\security`.
- To delete the existing self-signed certificate and private key provided by HCL, run the following commands.

```
install_dir\jre\jre\bin\ikeycmd -cert -delete -label cert_label -db
install_dir\wlp\usr\servers\server1\resources\security\SPKeyStore.jceks -pw sso_password -type JCEKS
```

Where:

cert_label

Is the label of the custom certificate generated for the BigFix Inventory server in step 2 ([on page dccxvi](#)). If you do not know the certificate label, run the following command:

```
install_dir\jre\jre\bin\ikeycmd -cert -list -db custom_cert.pfx -pw custom_cert_password
-type pkcs12
```

sso_password

Is the password to the SSO keystore. For the default keystore password contact the HCL Support. Otherwise, provide the password that you configured.

- To import the custom certificate, run the following commands.

```
install_dir\jre\jre\bin\ikeycmd -cert -import -file custom_cert.pfx -pw custom_cert_password -type
pkcs12 -target
install_dir\wlp\usr\servers\server1\resources\security\SPKeyStore.jceks -target_pw sso_password
-target_type JCEKS -label cert_label -new_label samlsp
```

Where:

custom_cert_password

Is the password to the custom certificate generated in step 2 (on page dccxvi).

sso_password

Is the password to the SSO keystore.

cert_label

Is the label of the custom certificate generated in step 2 (on page dccxvi).

7. In BigFix Inventory go to **Management > Single Sign-On Settings**. Click **Download Service Provider Metadata**, and save the `spMetadata.xml` file.

Based on the `spMetadata.xml` file, [configure Identity Provider for single sign-on \(on page dccxvii\)](#).

Step 2: Configuring Identity Provider for single sign-on

As the second step, configure BigFix Inventory server as a relying party to consume claims from the Identity Provider. Perform the configuration based on the `spMetadata.xml` file that you downloaded from BigFix Inventory.

The following procedure is based on the example of Active Directory Federation Services (ADFS).

1. Log in to the computer where Active Directory Federation Services are installed.
2. Copy the `spMetadata.xml` file from your computer to a directory on the ADFS server.
3. Click the **Start** rectangle in the lower-left area of the screen in Windows 2012 and then click the **ADFS Management** tile.
4. In the left navigation tree of the ADFS application, expand **ADFS > Trust Relationships > Relying Party Trusts**.
5. In the **Relying Party Trusts** pane on the right, click **Add Relying Party Trust**. A wizard opens. Click **Start**.
6. Select **Import data about the relying party from a file**.
7. Click **Browse**, select the `spMetadata.xml` file and click **Open**. Click **Next**.
8. On the new pane, provide the **Display name** for your ADFS service. Click **Next**.
9. Leave the option **Permit all users to access the relying party** selected, and click **Next**.
10. On the **Ready to Add Trust** pane, click **Next**.
11. On the **Finish** pane, click **Close**. The Edit Claim rules window opens.
12. Click the **Add Rule** button in the lower left corner. The **Add Transform Claim Rule** wizard opens. Click **Next**.
13. In the **Claim Rule template**, type `Name ID rule`.
14. From the Attribute store drop-down list, select **Active Directory**.
15. In the **Mapping of LDAP Attributes to outgoing claim types** section, click the first drop-down list and select **User Principal Name**. From the second list, select **Name ID**.
16. Repeat the step to achieve the following configuration and click **Finish**.

Table 71. Mapping of LDAP Attributes to outgoing claim types

LDAP Attribute	Outgoing Claim Type
User-Principal-Name	Name ID

LDAP Attribute	Outgoing Claim Type
E-Mail-Addresses	E-Mail Address
Token-Groups - Qualified by Long Domain Name	Group
SAM-Account-Name	Windows™ account name

17. In the Edit Claim rules window, click **Apply** and **OK**.

[Enable single sign-on in BigFix Inventory \(on page dccxviii\)](#).

Step 3: Enabling the SAML single sign-on

As the final step, enable single sign-on in the BigFix Inventory web user interface.

 **Important:** Before enabling single sign-on, create a dedicated administrative user. For more information, see: [Creating single sign-on users \(on page dccxxxi\)](#).

1. Log in to BigFix Inventory, and click **Management > Single Sign-on Settings**. Then, click **Enable**.
2. [Stop the BigFix Inventory server \(on page cd\)](#).
3. [Start the BigFix Inventory server \(on page cccxcix\)](#). The BigFix Inventory web user interface opens and you are redirected to the login page of the Identity Provider.
4. Provide the user name and password and click sign-in.

 **Note:** All user names must be fully-qualified names that contain the full domain name. For example: *user@domain.example*.

Reverting disabled SSO configuration for SAML

You can revert to the default SAML SSO configuration with single sign-on disabled if there are problems with logging in to the application.

To revert to the disabled SAML SSO configuration, follow the automated procedure below.

1. [Stop the BigFix Inventory server \(on page cd\)](#).
2. Restore the `web.xml` file from the `web.xml.timestamp.backup` file. The backup file was created automatically in the following location:
 - `UNIX` `bfi_install_dir/wlp/usr/servers/server1/apps/tema.war/WEB-INF`
 - `Windows` `bfi_install_dir\wlp\usr\servers\server1\apps\tema.war\WEB-INF`
3. [Start the BigFix Inventory server \(on page cccxcix\)](#).

Reverting disabled SSO configuration for SAML manually

It is recommended to use the automated solution to revert to the disabled SAML SSO configuration. However, if the automated solution is not successful, follow the steps below to revert the default SAML SSO configuration manually.

1. Stop the BigFix Inventory server (on page cd).
2. Make changes in the `server.xml` file that is in the following directory.

- **UNIX** `bfi_install_dir/wlp/usr/servers/server1`
- **Windows** `bfi_install_dir\wlp\usr\servers\server1`

- a. Remove the `<application-bnd>` element that is inside of the `<application>` element.

```
<application autoStart='true' location="tema.war" context-root="/" name="tema" type="war">
  <classloader commonLibraryRef='tema,DatabaseLib' delegation='parentLast' />
  <application-bnd>
    <security-role id="TemaSSOAuthenticated" name="TemaSSOAuthenticated">
      <special-subject type="ALL_AUTHENTICATED_USERS" />
    </security-role>
  </application-bnd>
</application>
```

- b. Remove the `<feature>samlWeb-2.0</feature>` element that is inside the `<featureManager>` element.

- c. Remove the `<samlWebSso20>` element if it exists.

3. Make changes in the `web.xml` file that is in the following directory.

- **UNIX** `bfi_install_dir/wlp/usr/servers/server1/apps/tema.war/WEB-INF`
- **Windows** `bfi_install_dir\wlp\usr\servers\server1\apps\tema.war\WEB-INF`

- a. Set the value of the `<config.sso.enabled>` parameter to *false*.

```
<context-param>
  <param-name>config.sso.enabled</param-name>
  <param-value>false</param-value>
</context-param>
```

- b. Remove the `<security-constraint>` element.

```
<security-constraint>
  <display-name>TemaSSOAuthenticated</display-name>
  <web-resource-collection>
    <web-resource-name>index</web-resource-name>
  </web-resource-collection>
  <url-pattern>/</url-pattern>
  <url-pattern>/session/*</url-pattern>
  <url-pattern>/management/*</url-pattern>
  <url-pattern>/scm/*</url-pattern>
  <url-pattern>/sam/*</url-pattern>
  <url-pattern>/setup/*</url-pattern>
  <url-pattern>/internal/*</url-pattern>
  <url-pattern>/wait_for_import</url-pattern>
```

```

<url-pattern>/import_finalizing</url-pattern>
<url-pattern>/import_status</url-pattern>
<url-pattern>/missing_computer_group</url-pattern>
<url-pattern>/account/*</url-pattern>
<url-pattern>/autocomplete/*</url-pattern>
<url-pattern>/pagestates/*</url-pattern>
<url-pattern>/reports/*</url-pattern>
<url-pattern>/test/*</url-pattern>
<url-pattern>/help/*</url-pattern>
</web-resource-collection>
<auth-constraint>
  <role-name>TemaSSOAuthenticated</role-name>
</auth-constraint>
<user-data-constraint>
  <transport-guarantee>CONFIDENTIAL</transport-guarantee>
</user-data-constraint>
</security-constraint>

```

4. Start the BigFix Inventory server (on page cccxcix).

Option 2: Configuring single sign-on based on IBM Lightweight Third-Party Authentication

You can configure single sign-on based on IBM Lightweight Third-Party Authentication(LTPA)with © IBM Security Access Manager for Web.

Back up the following files before you start configuring single sign-on:

- `server.xml`
 - **UNIX** `installation_dir/wlp/usr/servers/server1`
 - **Windows** `installation_dir\wlp\usr\servers\server1`
- `web.xml`
 - **UNIX** `installation_dir/wlp/usr/servers/server1/apps/tema.war/WEB-INF`
 - **Windows** `installation_dir\wlp\usr\servers\server1\apps\tema.war\WEB-INF`

The following scenario presents a typical workflow for configuring BigFix Inventory to work with BigFix® Security Access Manager. However, you might want to use other software products for enabling single sign-on in your infrastructure.

1. Configure the connection to your directory server (on page dccv).
2. Create the users (on page clxxxix) that will be authenticated with the single sign-on server. You must create at least one user that has the Administrator role.



Important: Ensure that you select **Single Sign-on** from the **Authenticated method** drop-down list.

3. [Export the LDAP server SSL certificate \(on page dccxxi\)](#) embedded in BigFix® IBM Security Access Manager for Web.
4. [Configure LTPA single sign-on \(on page dccxxi\)](#) in BigFix Inventory web user interface.
5. [Import the LTPA keys \(on page dccxxii\)](#) into BigFix® Security Access Manager for Web.
6. [Import the BigFix Inventory server certificate \(on page dccxxiv\)](#) into BigFix® Security Access Manager for Web.
7. [Configure a Virtual Junction \(on page dccxxv\)](#) in BigFix® IBM Security Access Manager for Web.
8. [Enable single sign-on \(on page dccxxx\)](#) in BigFix Inventory.
9. [Optional: Update the WebUI shortcut \(Windows only\) \(on page dccxxxi\)](#)
10. [Optional: Reverting SSO configuration for LTPA \(on page dccxxxi\)](#).

You can revert to the default LTPA SSO configuration with single sign-on disabled if there are problems with logging in to the application.

Related information

[Configuring LTPA on the Liberty profile](#)

[Customizing SSO configuration using LTPA cookies for the Liberty profile](#)

Step 1: Exporting the LDAP server SSL certificate embedded in BigFix Security Access Manager

You must export the certificate from the LDAP server embedded in IBM Security Access Manager (ISAM) to be able to configure the single sign-on in the BigFix Inventory web user interface.

1. Log on to the ISAM web user interface.
2. **Manage System Settings > Secure Settings > SSL Certificates.**
3. Select **embedded_ldap_keys**, click the **Manage** drop-down list and then click **Edit SSL Certificate Database**.
4. Open the **Personal Certificates** tab. Select **server**, click the **Manage** drop-down list and then click **Export**. A pop-up window opens with the server certificate details.
5. Click **OK** and save the file in a directory on your computer.

You exported the LDAP server SSL certificate. You can use it in the next step while configuring LTPA single sign-on.

Step 2: Configuring LTPA single sign-on

You can use the **Single Sign-On Settings** pane to configure IBM Lightweight Third-Party Authentication (LTPA) settings in BigFix Inventory.



Note: If you set up the session timeout for Single Sign-On, remember that it should be longer than the session timeout that is set up for BigFix Inventory. Otherwise, change the settings in BigFix Inventory. For more information, see: [Setting session timeout \(on page cdi\)](#).

1. In the BigFix Inventory web UI, click **Management > Single Sign-On Settings**.
2. Select **LTPA** as the single sign-on method.
3. From the **Directory Server** drop-down list select the LDAP server that you want to be used for authenticating users.
4. Click **Browse** and select your certificate file.

Method* SAML LTPA

Instance ID* ISAMLdapRegistry

Directory Server* ISAM embedded LDAP ▼

Directory Server SSL Certificate server.cer

5. Click **Save**.
6. [Stop the BigFix Inventory server \(on page cd\)](#).
7. [Start the BigFix Inventory server \(on page cccxcix\)](#).

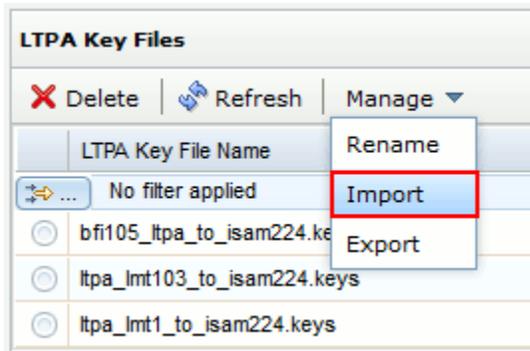
Step 3: Importing LTPA keys into BigFix® Security Access Manager

You must import the LTPA keys into IBM® Security Access Manager to add BigFix Inventory to the trusted list.

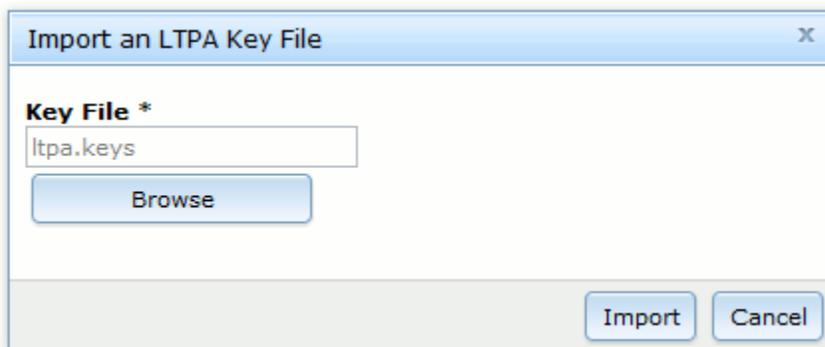
1. In the BigFix Inventory web user interface, click **Management > Single Sign-On Settings**.
2. Click **Download LTPA keys** and save the files in a directory on your computer.
3. Log on to the Security Access Manager web user interface.
4. Click **Secure Web Settings > Global Keys > LTPA Keys**.



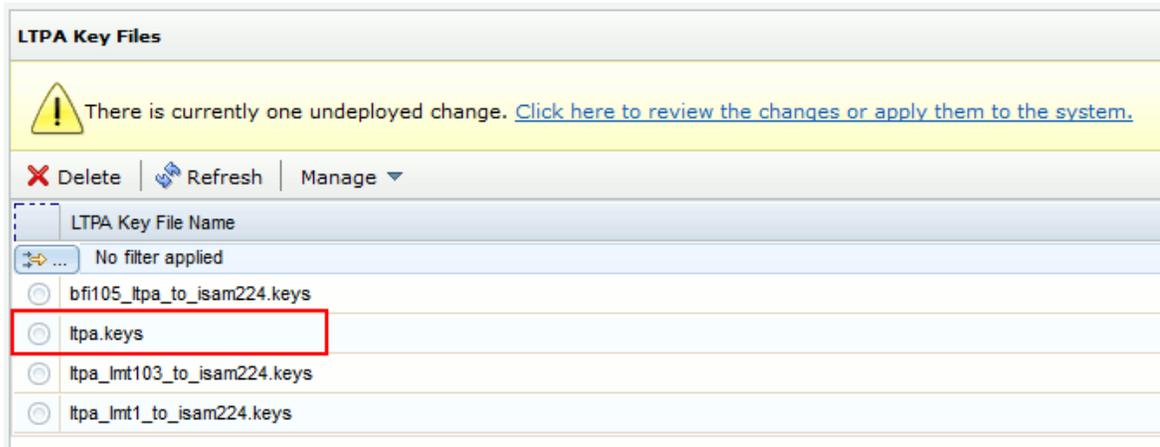
5. Click **Manage > Import**



6. Click **Browse** and select the file to be imported.
7. Click **Import**.



A message about one undeployed change is displayed. The LTPA keys are displayed together with other keys.



8. Click the link **Click here to review the changes or apply them to the system**.

You just imported the LTPA keys into IBM® Security Access Manager for Web.

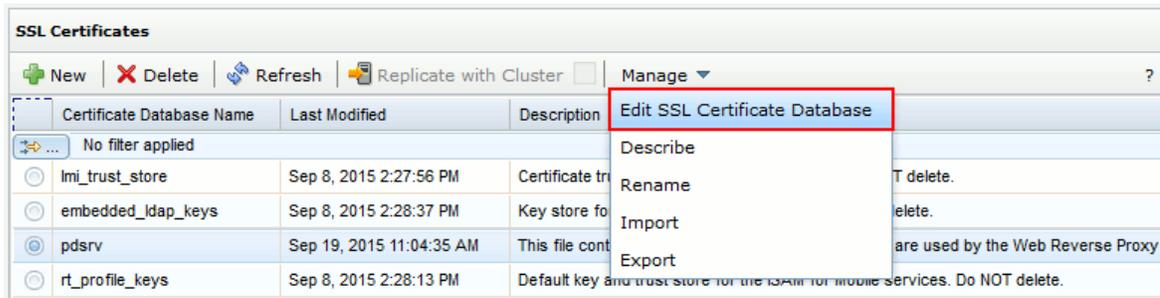
It is also advisable to rename the imported key file. To rename the file that you imported:

1. Click the **Manage** drop-down list and then click **Rename**.
2. In the pop-up window that opens, provide the new name for your LTPA keys file and click **Save**.
3. On the **LTPA Key Files** pane, click the link **Click here to review the changes or apply them to the system**.
4. In the **Deploy Pending Changes** pop-up window, click **Deploy**.

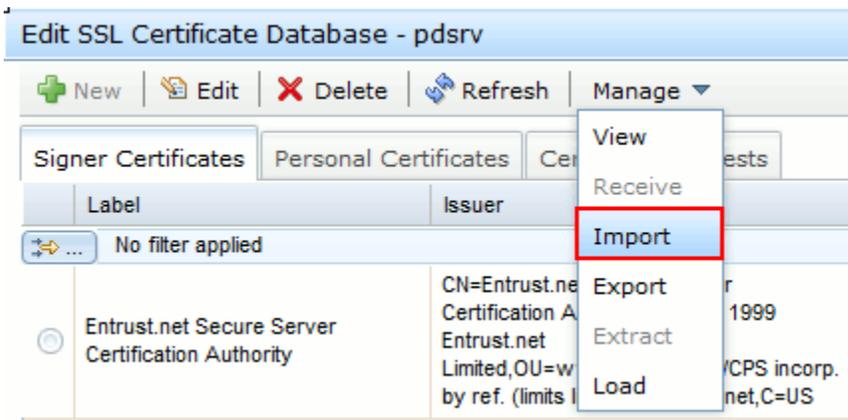
Step 4: Importing the BigFix Inventory server certificate into IBM® Security Access Manager

Import the BigFix Inventory server certificate into BigFix® Security Access Manager for Web to establish trusted server-to-server communication.

1. In the BigFix Inventory web UI, click **Management > Server Settings**.
2. Click **Download Certificate** and save the file in a directory on your computer.
3. Log on to the Security Access Manager web user interface.
4. Click **Manage System Settings > Secure Settings > SSL Certificates**.
5. Select **pdsrv** and then, from the **Manage** drop-down list, select **Edit SSL Certificate Database**.



6. From the **Manage** drop-down list, select **Import**. A new window opens.



7. Click **Browse** to select the certificate file and provide the certificate label.

8. Click **Import**.
9. Click **Close** to close the **Edit SSL certificate database - pdsrv** pop-up window.
A message about an undeployed change is displayed.
10. Click the link **Click here to review the changes or apply them to the system**.
11. In the **Deploy Pending Changes** pop-up window, click **Deploy**.
A message about the successful deployment is displayed.

Restart the default instance of the proxy server for the changes to take effect.

1. Click **Secure Web Settings** and then **Reverse Proxy**.
2. In the navigation bar, click **Restart**.

Step 5: Configuring a virtual junction

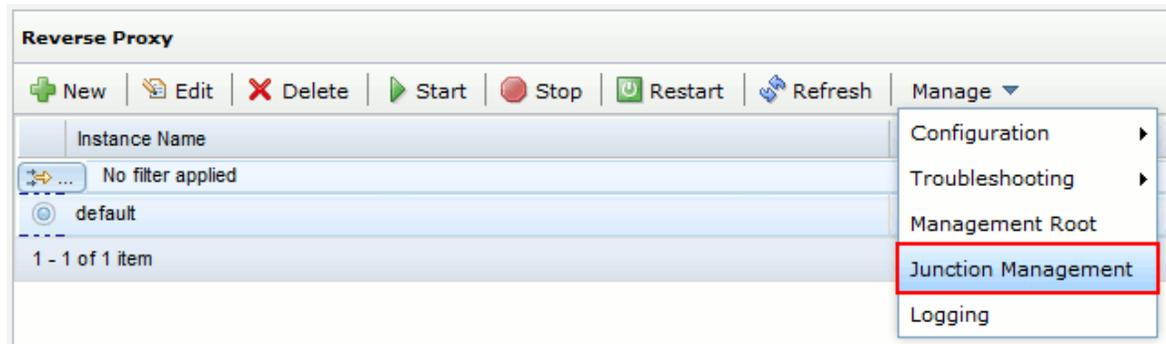
A virtual host junction is a mount point for specific content that is located on the WebSEAL server.

Ensure that the HTTPS port of the reverse proxy in BigFix Security Access Manager for web is set to 9081, which is the port number used by BigFix Inventory server. If the port is not set to this value, REST API might not work. To configure this port for the proxy server:

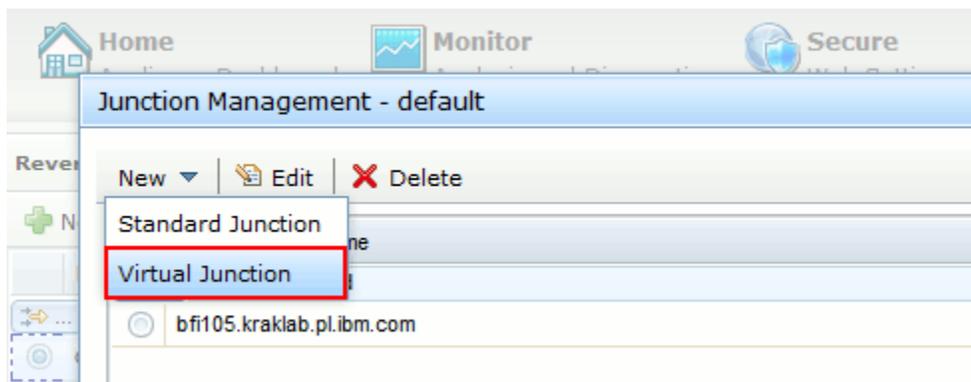
1. In the IBM Security Access Manager user interface, click **Secure Web Settings** and then, under **Manage** click **Reverse Proxy**.
 2. Select the **default** proxy and click **Edit**.
 3. In the **Reverse Proxy Basic Configuration - default** window, on the **Server** tab, provide 9081 as the value of the **HTTPS Port** and click **Save**.
1. Log on to BigFix® Security Access Manager.
 2. In the top navigation bar, click **Secure Web Settings > Manage > Reverse Proxy**.



3. Select the instance and then, from the drop-down list on the right of the Reverse Proxy bar, select **Manage > Junction Management**. A new pane opens.



4. From the drop-down list in the upper-left corner of the pane, click **New > Virtual Junction**.



5. On the new pane, specify the junction label, Virtual Host, Virtual Host Port, and SSL as the Junction Type, and click **Save**.

Create a Virtual Junction

Junction Servers Basic Authentication Identity SSO and LTPA

Creation of a junction for an initial server

Junction Label MyAppServer	Junction Type
<input type="checkbox"/> Stateful Junction	<input type="radio"/> TCP
DSC Environment 	<input checked="" type="radio"/> SSL
Virtual Host Label 	<input type="radio"/> TCP Proxy
Virtual Host NC142222.kraklab.pl.ibm.com	<input type="radio"/> SSL Proxy
Virtual Host Port 9081	

6. Configure a backend server for this junction: Click the **Servers** tab and then click **New**.

Create a Virtual Junction

Junction **Servers** Basic Authentication Identity SSO and LTPA

Target Backend Servers. At least one server is required to create a

New | Edit | Delete

Hostname

No filter applied

7. Specify the IP address or host name of the BigFix Inventory server and click **Save**.

Add TCP or SSL Servers

Hostname * 9.158.142.222	Query Contents <input type="text"/>
TCP or SSL Port * 9081	UUID of the Server <input type="text"/>
Local Address <input type="text"/>	Distinguished Name(DN) <input type="text"/>
	<input type="checkbox"/> Windows File System Support
	<input type="checkbox"/> Treat URL as case insensitive

8. Click the **Identity** tab and select all the items under **HTTP Header Identity Information**.

Create a Virtual Junction

Junction Servers Basic Authentication **Identity** SSO and LTPA General

Supply identity information in HTTP headers

HTTP Basic Authentication Header
Filter

GSO Resource or Group

HTTP Header Identity Information

- IV-USER
- IV-USER-L
- IV-GROUPS
- IV-CREDS

HTTP Header Encoding
None

Include session cookie

Insert client IP address

Enable TFIM SSO

9. Click the **SSO and LTPA** tab and select the following entries.

- **Enable LTPA cookie support**
- **Use Version 2 Cookies**

From the **LTPA Keyfile** drop-down list, select the file that you into IBM Security Access Manager and in the next field provide the LTPA keyfile password.

Create a Virtual Junction

Junction Servers Basic Authentication Identity **SSO and LTPA** General

WebSphere single signon (LTPA) junctions

Enable LTPA cookie Support

Use Version 2 Cookies

LTPA Keyfile

LTPA Keyfile Password

10. Click the **General** tab and leave all the default configuration unchanged.

Create a Virtual Junction

Junction Servers Basic Authentication Identity SSO and LTPA **General**

General Junction Options

FSSO Configuration File

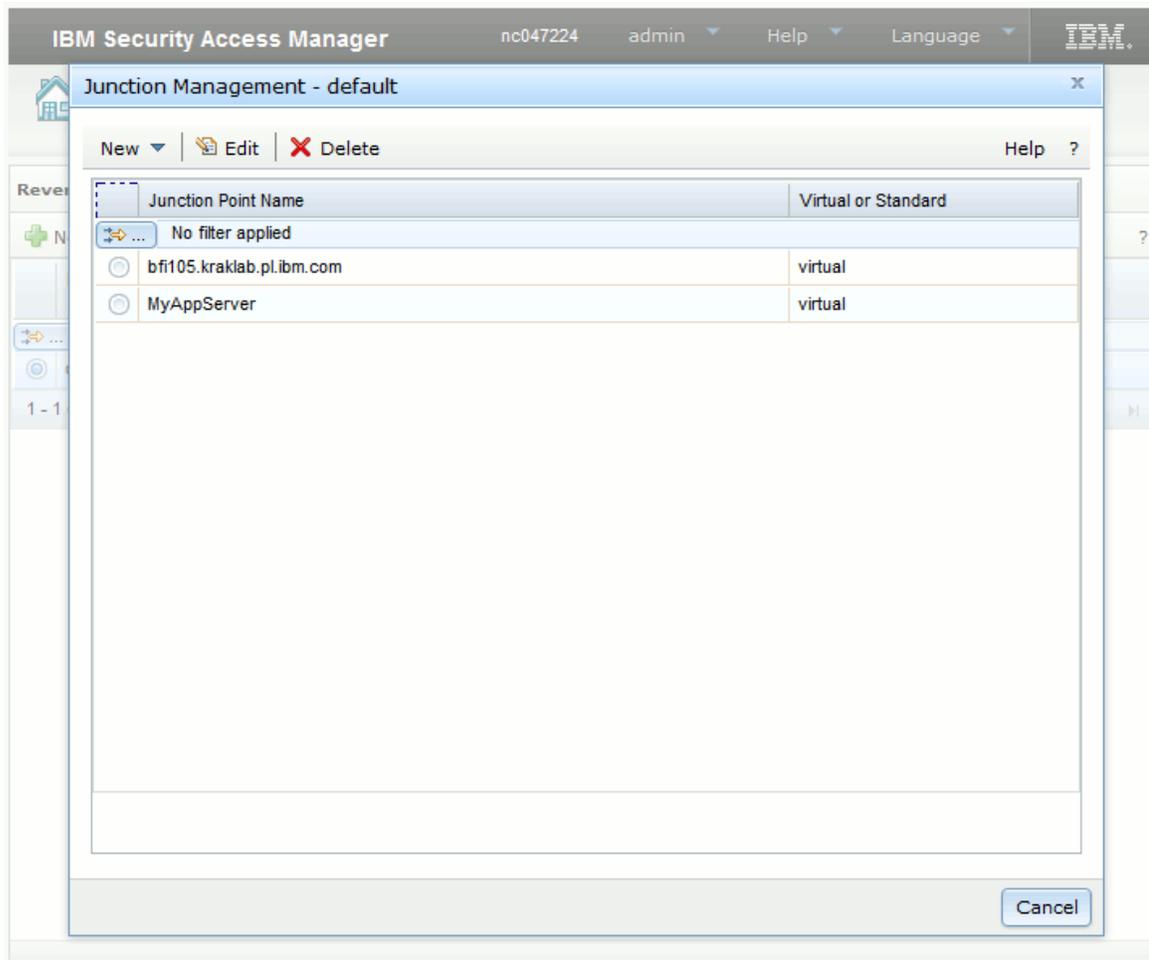
Percentage Value for Hard Limit of Worker Threads

Percentage Value for Soft Limit of Worker Threads

Include authorization rules decision information

11. Click **Save** to save the configuration and exit the wizard.

12. Click **Cancel** to exit the **Junction Management - default** pane.



Step 6: Enabling the LTPA single sign-on

As the final step, enable single sign-on in BigFix Inventory web user interface.

! **Important:** Before enabling single sign-on, create a dedicated administrative user. For more information, see: [Creating single sign-on users \(on page dccxxxi\)](#).

1. Open the BigFix Inventory web user interface.
2. Go to **Management > Single Sign-on Settings** and click **Enable**.
3. [Stop the BigFix Inventory server \(on page cd\)](#).
4. [Start the BigFix Inventory server \(on page cccxcix\)](#).
5. Access your BigFix Inventory server using the virtual host that you created in step [Step 5: Configuring a virtual junction \(on page dccxxv\)](#). Example: `https://virtual_host/sam`.



Note: Once LTPA single sign-on is enabled, you will not be able to access the BigFix Inventory web user interface using the URL `https://host_name:9081`.

6. Log in with the directory user that you created in the step [Setting up users \(on page clxxxix\)](#).

Optional: Updating the WebUI shortcut

You might need to edit the WebUI shortcut if you installed or upgraded to BigFix Inventory 9.2.1 and you also enabled single sign-on using LTPA. This URL is defined in the WebUI file that is in the following location: `C:\Program Files\BigFix Enterprise\BFI\admin\resources\WebUI`.

1. Open the WebUI file in a text editor, for example Notepad++, and change the URL from `https://localhost:port-number` to `https://virtual-junction-host-name:port-number`.
2. Clear the cache in your web browser.
3. End all your browser processes.
4. Restart the system on which your browser is installed.

Optional: Reverting disabled SSO configuration for LTPA

You can revert to the default LTPA SSO configuration with single sign-on disabled if there are problems with logging in to the application.

1. [Stop the BigFix Inventory server \(on page cd\)](#).
2. Restore the `web.xml` file from the `web.xml.timestamp.backup` file. The backup file was created automatically in the following location:
 - **UNIX** `bfi_installation_dir/wlp/usr/servers/server1/apps/tema.war/WEB-INF`
 - **Windows** `bfi_installation_dir\wlp\usr\servers\server1\apps\tema.war\WEB-INF`
3. [Start the BigFix Inventory server \(on page cccxcix\)](#).

Creating single sign-on users

To complete an authentication process through single sign-on, you must create users that can later log in to BigFix Inventory.



You must have the Manage Users permission to perform this task.

1. In the top navigation bar, click **Management > Users**.
2. To create a user, click **New**.
3. In the **User Name** field, type the name of an existing single sign-on user.



Tip: The login format depends on the configuration of your environment. It can be, for example, an e-mail address, or a user ID, and it might need to be followed by a domain name.

4. From the list, select a **Computer Group** to which the user would be assigned.
5. From the **Authentication Method** list, select **Single Sign-On**.
6. Click **Create**.
7. To delete the created user, click its name. Then, in the upper left of the window click **Delete**.

To confirm authentication, log in to the BigFix Inventory server with the credentials of the single sign-on user that you created in BigFix Inventory.

Logging out of the BigFix Inventory server

To log out of the BigFix Inventory user interface, simply clear your browser history and restart the browser.

1. Clear the browser history. The steps might vary depending on the web browser that you are using.
2. Restart the web browser.

Disabling the single sign-on configuration

You can disable the SSO configuration on the **Single Sign-on Settings** pane.

1. Log on to the BigFix Inventory web user interface.
2. In the top navigation bar, click **Management > Single Sign-On Settings**.
3. Click **Disable**.
4. Restart the BigFix Inventory server.
 - a. [Stop the server \(on page cd\)](#).
 - b. [Start the server \(on page cccxcix\)](#).

You can now log in to the BigFix Inventory web user interface with the credentials that you used before you enabled single sign-on.

To know more about reverting disabled SSO, refer to [Reverting disabled SSO configuration for SAML \(on page dccxviii\)](#) and [Optional: Reverting disabled SSO configuration for LTPA \(on page dccxxxi\)](#).

Deleting the single sign-on configuration

You can delete the SSO configuration on the **Single Sign-on Settings** pane.



Important: Do not delete the single sign-on configuration while it is enabled. You must first [disable the single sign-on configuration \(on page dccxxii\)](#) before you delete it.

1. Log on to the BigFix Inventory web user interface.
2. In the top navigation bar, click **Management > Single Sign-On Settings**.
3. Click **Delete**. A pop-up window appears.
4. Click **Delete** to confirm.

The `server.xml` file is backed up to `server.xml.timestamp.backup` in the following location. You can restore this file to undo these changes.

- **UNIX** `bfi_installation_dir/wlp/usr/servers/server1`
- **Windows** `bfi_installation_dir\wlp\usr\servers\server1`

Modifying port in BigFix Inventory that has single sign-on enabled

Modifying the port number on the **Server Settings** pane in BigFix Inventory while single sign-on is enabled invalidates the single sign-on configuration. To properly modify the port, you must repeat the single sign-on configuration steps.

1. [Disable the single sign-on configuration \(on page dccxviii\)](#).
2. Provide the new port value on the **Server Settings** page:
 - a. In the BigFix Inventory web user interface, click **Management > Server Settings**.
 - b. Provide the new port value in the **Port** field, and click **Save**.
3. Re-create the single sign-on configuration with the new port value. For more information, see either [Configuring SSO based on SAML token \(on page dccxii\)](#) or [Configuring SSO based on LTPA \(on page dccxx\)](#).

Configuring SSO keystore passwords and encryption

Configure unique passwords to the SSO keystores, and encrypt them with the AES encryption algorithm.

9.2.9 Starting from application update 9.2.9, for all fresh configurations of single sign-on in BigFix Inventory, SSO keystore passwords use AES as a default encryption method. The password encryption schema remains unchanged if the single sign-on was enabled before an upgrade to version 9.2.9.

Both SAML and LTPA SSO types require two separate keystores: Service Provider and Identity Provider. You can set up either identical or different passwords for these keystores. To configure new passwords for your SSO keystores, or change encryption method from XOR to AES, use the appropriate variables and complete the following procedure.

<SP_keystore_name>

Service Provider keystore name is either **SPKeyStore** for SAML SSO, or **LdapSSLTrustStore** for LTPA SSO.

<IP_keystore_name>

Identity Provider keystore name is either **SPKeyStore** for SAML SSO, or **LdapSSLTrustStore** for LTPA SSO.

<alias>

Is either **samlsp** for SAML SSO, or **default** for LTPA SSO.

1. [Stop the BigFix Inventory server. \(on page cd\)](#)
2. **Optional:** If you want to configure a custom SSO keystore password, follow the instructions below. If you only want to change password encryption method from XOR to AES, continue with step 3.

- a. To change the key password for the Service Provider keystore, run the following command.

```
Linux <Installation_directory>/jre/bin/keytool -keypasswd -keystore
<Installation_directory>/wlp/usr/servers/server1/resources/security/<SP_keystore_name>.jceks
-storetype JCEKS -alias <alias>
```

```
Windows <Installation_directory>\jre\bin\keytool -keypasswd -keystore
<Installation_directory>\wlp\usr\servers\server1\resources\security\<SP_keystore_name>.jceks
-storetype JCEKS -alias <alias>
```

When prompted, provide the SSO keystore password. For the default keystore password contact the BigFix Support.

```
Enter keystore password:
Enter key password for <alias>:
New key password for <alias>:
Re-enter new key password for <alias>:
Password change successful for alias <alias>
```

- b. To change the Service Provider keystore password, run the following command.

```
Linux <Installation_directory>/jre/jre/bin/keytool -storepasswd -keystore
<Installation_directory>/wlp/usr/servers/server1/resources/security/<SP_keystore_name>.jceks
-storetype JCEKS
```

```
Windows <Installation_directory>\jre\jre\bin\keytool.exe -storepasswd -keystore
<Installation_directory>\wlp\usr\servers\server1\resources\security\<SP_keystore_name>.jceks
-storetype JCEKS
```

Set up the new keystore password. The password must match the password that you set up for the key (step a).

```
Enter keystore password:
New keystore password:
Re-enter new keystore password:
```

- c. To change the Identity Provider keystore password, run the following command.

```
Linux <Installation_directory>/jre/jre/bin/keytool -storepasswd -keystore
<Installation_directory>/wlp/usr/servers/server1/resources/security/<IP_keystore_name>.jceks
-storetype JCEKS
```


If you upgraded from an earlier version and a user password does not fulfill these requirements, it remains valid until the user changes it. The new password must fulfill the default requirements. Similarly, if you introduce a password policy, existing passwords that do not meet the requirements remain valid until the next time they are changed.

1. Go to **Management > Advanced Server Settings**.

2. Define requirements for user passwords.

- To specify the minimal length of the password, set the **user_password_min_length** parameter. For example:

```
user_password_min_length = 8
```

- To specify the maximal length of the password, set the **user_password_max_length** parameter. For example:

```
user_password_max_length = 16
```

- To set a limit on the number of the same characters that can be used in a row, set the **user_password_max_of_identical_characters_in_row** parameter. For example:

```
user_password_max_of_identical_characters_in_row = 3
```

- To specify that at least one special character must be used, provide the list of acceptable characters in the **user_password_special_characters** parameter. For example:

```
user_password_special_characters = ?!%&
```

- To specify other requirements such as the need to use at least one uppercase character, create a regular expression that defines the requirement, and provide it in the **user_password_regular_expression** parameter. For example:

```
user_password_regular_expression = ^(?=.*[A-Z]).+$
```



Tip: To learn more about each parameter, see: [Advanced server settings \(on page cdii\)](#).

Configuring cryptographic keystore password and encryption

Configure a unique password to the cryptographic keystore, and encrypt it with the AES encryption algorithm.

9.2.7

Starting from application update 9.2.7, for all fresh installations of BigFix Inventory, keystore and database passwords use AES as a default encryption method. The password encryption schema remains unchanged for all the application instances upgraded to version 9.2.7.

To configure SSO keystore passwords and encryption, see: [Configuring SSO keystore passwords and encryption \(on page dccxxiii\)](#).

1. [Stop the BigFix Inventory server. \(on page cd\)](#)
2. **Optional:** If you want to change the keystore password into a custom password, follow the instructions below. If you only want to change password encryption method from XOR to AES, continue with step 3.

a. To list the contents of the application keystore, run the following command.

```
Linux Installation_directory/jre/jre/bin/keytool -list -keystore <inst_dir>/wlp/usr/servers/
server1/resources/security/key_server.jceks -storetype JCEKS
```

```
Windows Installation_directory\jre\jre\bin\keytool.exe -list -keystore <inst_dir>\wlp\usr
\servers\server1\resources\security\key_server.jceks -storetype JCEKS
```

When prompted, provide the keystore password. For the default keystore password contact the HCL Support.

```
Enter keystore password:

Keystore type: JCEKS
Keystore provider: IBMJCE

Your keystore contains 1 entry

default, Nov 15, 2013, keyEntry,
Certificate fingerprint (SHA1):
xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx
```



Note: The list can consist of more than one entry, if Single Sign On is configured for the product. The underlined element is an alias for the listed entry, and it is required to complete next steps.

b. Ensure that you change the password for all the listed entries. Run the following command, and specify the alias at the end of the command line.

```
Linux Installation_directory/jre/jre/bin/keytool -keypasswd -keystore <inst_dir>/wlp/usr/
servers/server1/resources/security/key_server.jceks -storetype JCEKS -alias default
```

```
Windows Installation_directory\jre\jre\bin\keytool.exe -keypasswd -keystore <inst_dir>\wlp\usr
\servers\server1\resources\security\key_server.jceks -storetype JCEKS -alias default
```

Set up a single password for all the listed entries. For more than one entry, ensure the passwords are identical.

```
Enter keystore password:
Enter key password for <default>:
New key password for <default>:
```


- To overwrite the default key that is used to encrypt passwords to VM managers, perform the following steps.

1. Go to the VM Manager Tool directory.

2. Stop VM Manager Tool by using the following command.

- **Linux** `./vmman.sh -stop`
- **Windows** `vmman.bat -stop`

3. Back up the `config` and `keydb` directories.

- **Linux** `/var/opt/BESClient/LMT/VMMAN/config//var/opt/BESClient/LMT/VMMAN/keydb/`
- **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\VMMAN\config\C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\VMMAN\keydb`

If an error occurs during the regeneration of the key, restore these directories to their current locations.

4. Run the following command.

- **Linux** `./vmman.sh -regenerateencryptionkey`
- **Windows** `vmman.bat -regenerateencryptionkey`

5. Start VM Manager Tool by using the following command.

- **Linux** `./vmman.sh -run`
- **Windows** `vmman.bat -run`

- To change the default password to the VM Manager Tool keystore, perform the following steps.

1. Go to the VM Manager Tool directory.

2. Stop VM Manager Tool by using the following command.

- **Linux** `./vmman.sh -stop`
- **Windows** `vmman.bat -stop`

3. Back up the VM Manager Tool keystore and configuration files.

- **Linux** `/var/opt/BESClient/LMT/VMMAN/keydb/keys.jceks/var/opt/BESClient/LMT/VMMAN/config/vmmmainconf.properties`
- **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\VMMAN\keydb\keys.jceksC:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\VMMAN\config\vmmmainconf.properties`

If an error occurs during changing the keystore password, restore these files to their current locations.

4. Create a `txt` file, for example `keystore_password.txt`. Provide the new keystore password in the `customPassword` parameter.

```
customPassword=<new_password>
```

5. To change the password, run the following command.

- **Linux** `./vmman.sh -changepassword -file /var/opt/BESClient/LMT/VMMAN/config/keystore_password.txt`
- **Windows** `vmman.bat -changepassword -file "C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\VMMAN\config\keystore_password.txt"`

Where `-file` is the path to the `txt` file in which you specified the new keystore password.

After you run the command, the password is encrypted and saved in the `vmmmainconf.properties` under the `vmm_keystore_password_do_not_change_it` parameter.

6. After the new password is set, remove the `txt` file in which you specified the password.

7. Start VM Manager Tool by using the following command.

- **Linux** `./vmman.sh -run`
- **Windows** `vmman.bat -run`

Changing the default secret key and password to the SAP Metric Data Collector keystore

9.2.9 Available from 9.2.9. To improve security of encrypting the password to the SAP Metric Data Collector, you can change the default secret key or the default password to the SAP Metric Data Collector keystore. These two procedures are independent. You can change the secret key, the keystore password or both depending on your needs.

Contact BigFix Support to obtain the current password to the SAP Metric Data Collector keystore.

- To substitute the default secret key with a custom one, perform the following steps.



Important: To run each of the following commands, provide the current password for the SAP Metric Data Collector keystore. You can obtain the default password from BigFix Support. If you already changed it, use the custom password that you set up.

1. To remove the default secret key, run the following command.

- **Linux** `"/var/opt/BESClient/LMT/SAPTOOL/java/jre/bin/keytool" -delete -keystore "/var/opt/BESClient/LMT/SAPTOOL/keydb/keys.jceks" -storetype JCEKS -alias bigfixsaptoolcustomkey`
- **Windows** `"C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\java\jre\bin\keytool.exe" -delete -keystore "C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\keydb\keys.jceks" -storetype JCEKS -alias bigfixsaptoolcustomkey`

2. To create a custom secret key, run the following command.



Important: During the creation of the new secret key, you are asked to provide its password. Specify the same password as the password for the SAP Metric Data Collector keystore. If you

- ! did not change the keystore password, provide the one that you obtained from BigFix Support. Otherwise, provide your custom password for the SAP Metric Data Collector keystore.

- **Linux** `"/var/opt/BESClient/LMT/SAPTOOL/java/jre/bin/keytool" -genseckey -keystore "/var/opt/BESClient/LMT/SAPTOOL/keydb/keys.jceks" -storetype JCEKS -keyalg AES -keysize 128 -alias bigfixsaptoolcustomkey`
- **Windows** `"C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\java\jre\bin\keytool.exe" -genseckey -keystore "C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\keydb\keys.jceks" -storetype JCEKS -keyalg AES -keysize 128 -alias bigfixsaptoolcustomkey`

Where:

-keyalg

Is the algorithm used to encrypt the secret key. Only AES is supported.

-keysize

Is the size of the secret key. The default size is 128 bits. To use a longer key, set up Java JCE unlimited jurisdiction policy. For more information, see: [IBM SDK Policy files](#).

- To change the default password to the SAP Metric Data Collector keystore, perform the following steps.

- ! **Important:** To run each of the following commands, provide the current password for the SAP Metric Data Collector keystore that you obtained from HCL Support.

1. To change the password to the keystore, run the following command.

Password to each key in the keystore must be the same as the password to the keystore. Thus, the **-all** parameter is used in the command.

- **Linux** `"/var/opt/BESClient/LMT/SAPTOOL/java/jre/bin/keytool" -storepasswd -all -keystore "/var/opt/BESClient/LMT/SAPTOOL/keydb/keys.jceks" -storetype JCEKS`
- **Windows** `"C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\java\jre\bin\keytool.exe" -storepasswd -all -keystore "C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\keydb\keys.jceks" -storetype JCEKS`

2. When prompted, provide the new keystore password.

3. To encode the password, run the following command.

- **Linux** `"/var/opt/BESClient/LMT/SAPTOOL/saptool.sh" -encode`
- **Windows** `"C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\saptool.bat" -encode`

The password is encoded and displayed in the command line.

4. Copy the encoded password to a file, and save it in a text format.

For example:

- **Linux** `"/var/opt/BESClient/LMT/SAPTOOL/config/custom_access_code.txt"`
- **Windows** `"C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\config\custom_access_code.txt"`



Note: It is recommended to save the `custom_access_code.txt` in the `config` directory. This directory remains unchanged after the SAP Metric Data Collector update.

5. Open the `saptool_config.properties` file that is in the following location.

- **Linux** `"/var/opt/BESClient/LMT/SAPTOOL/config/"`
- **Windows** `"C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\SAPTOOL\config\"`

Provide the path to the file with the encoded password in the `configuration_key_init_path` parameter.

For example:

- **Linux** `configuration_key_init_path=/var/opt/BESClient/LMT/SAPTOOL/config/custom_access_code.txt`
- **Windows** `configuration_key_init_path=C:\\Program Files (x86)\\BigFix Enterprise\\BES Client\\LMT\\SAPTOOL\\config\\custom_access_code.txt`



Note: Provide the path in one line and without quotation marks. To use a backslash, provide it twice (\\).

Related information

[Collecting utilization of SAP license metrics \(on page dxv\)](#)

[Measuring and collecting utilization of SAP license metrics \(on page dxx\)](#)

Configuring user account lockout

9.2.8

Available from 9.2.8. By default, user account is locked for 5 minutes after the user attempts to log in to BigFix Inventory more than 10 times within 5 minutes. You can change the default settings or disable the user account lockout.

The lockout mechanism is enabled by default. To disable it, set the value of the `user_lockout_enabled` parameter to `false`.

1. Go to **Management > Advanced Server Settings**.
2. Specify conditions under which the user account is locked.
 - To specify the maximal number of failed login attempts, set the value of the **user_max_login_retries** parameter. For example:

```
user_max_login_retries = 5
```

- To specifies the period during which the user can try to log in to the application, set the value of the **user_retry_time_period**. For example:

```
user_retry_time_period = 10
```

When the maximal number of failed login attempts is exceeded within this period, the account is locked.

- To specify the period for which a user cannot log in to the application after the account is locked, set the value of the **user_lockout_length**. For example:

```
user_lockout_length = 10
```

 **Tip:** To learn more about each parameter, see: [Advanced server settings \(on page cdii\)](#).

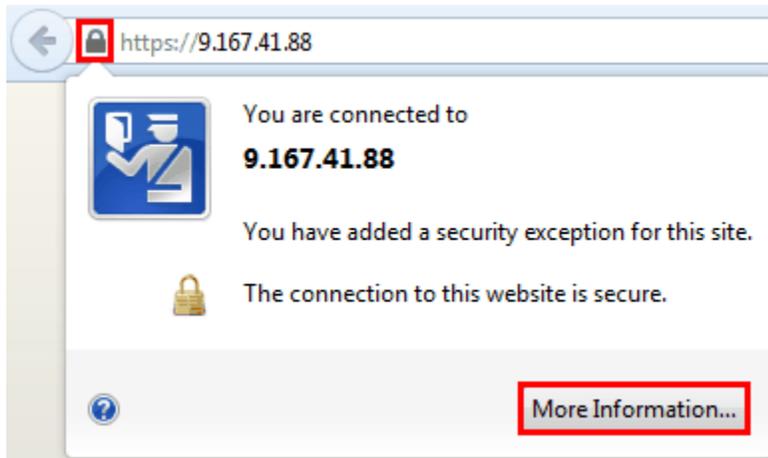
Configuring VM Manager Tool to accept trusted VM manager certificates

By default, the VM Manager Tool accepts all VM manager certificates regardless of whether they are trusted or not. You can change the default behavior to ensure that only trusted certificates are accepted by the VM Manager Tool.

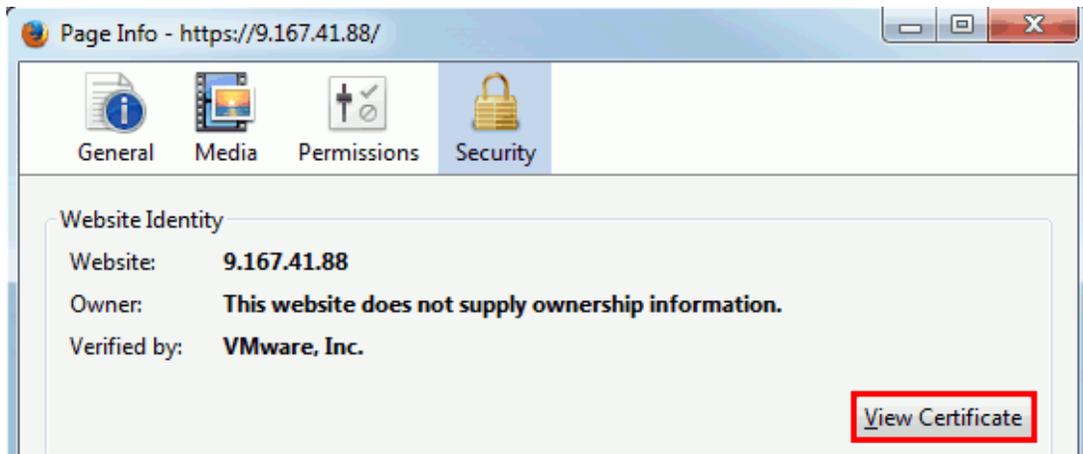
1. Extract the VM manager certificate to a file.

The following steps are provided as an example and show how to extract a VMware certificate in Firefox. The steps that you need to perform might differ depending on the virtualization type and the web browser that you are using. If you encounter problems with extracting VM manager certificates, refer to the documentation of the virtualization that you are using.

- a. Type the VM manager address in the web browser.
- b. Click the lock sign and click **More Information**.



c. Open the **Security** tab and click **View Certificate**.



d. Open the **Details** tab and click **Export**.

e. Save the file in the DER format.

2. Log in to the computer where the VM Manager Tool is installed and copy the VM manager certificate to that computer.

3. To define the certificate as trusted, open the command-line interface and run the following command.

```
vmman.bat -addcertificate -alias unique_alias -file vm_manager_certificate
```

Where:

-alias

Unique alias that will be associated with the VM manager certificate.

-file

Path to the VM manager certificate that you extracted.

 **Important:** Both switches are required.

4. Change the VM Manager Tool settings to accept only trusted certificates.
 - If you are using [basic VM management \(on page cccxlvi\)](#), log in to BigFix Inventory and go to **Management > Advanced Server Settings**. Then, set the value of the **vmmman_trust_all_vm_managers_certificates** parameter to false.
 - If you are using [advanced VM management \(on page ccclxvii\)](#), open the `<BES Client>\LMT\VMMAN\config\vmmainconf.properties` file on the computer where the VM Manager Tool is installed. Then, change the value of the **vmm_trust_all_vm_managers_certificates** parameter to false.

The VM Manager Tool accepts only trusted VM manager certificates. If a certificate of a VM manager that is defined in BigFix Inventory is not trusted, the status of the VM manager changes to *Connection failed*.

Relays

Relays lighten both upstream and downstream burdens on the server. Rather than communicating directly with a server, clients can instead be instructed to communicate with designated relays, considerably reducing both server load and client and server network traffic.

Relays work by:

- Relieving downstream traffic.
- Reducing upstream traffic.
- Reducing congestion on low-bandwidth connections.
- Reducing the load on the server.

Relays are an absolute requirement for any network with slow links or more than a few thousand clients. For more information about Relays, see [BigFix documentation](#).

Troubleshooting and support

Learn about solutions to common problems that might arise when you use BigFix Inventory and how to find logs and trace files that help you troubleshoot those problems.

Troubleshooting a problem

Troubleshooting is a systematic approach to solving a problem. The goal of troubleshooting is to determine why something does not work as expected and explain how to resolve the problem.

The first step in the troubleshooting process is to describe the problem completely. Problem descriptions help you and the BigFix Support person know where to start to find the cause of the problem. This step includes asking yourself basic questions:

- What are the symptoms of the problem?
- Where does the problem occur?
- When does the problem occur?
- Under which conditions does the problem occur?
- Can the problem be reproduced?

The answers to these questions typically lead to a good description of the problem, and that is the best way to start down the path of problem resolution.

What are the symptoms of the problem?

When starting to describe a problem, the most obvious question is "What is the problem?" This might seem like a straightforward question; however, you can break it down into several more-focused questions that create a more descriptive picture of the problem. These questions can include:

- Who, or what, is reporting the problem?
- What are the error codes and messages?
- How does the system fail? For example, is it a loop, hang, crash, performance degradation, or incorrect result?
- What is the business impact of the problem?

Where does the problem occur?

Determining where the problem originates is not always easy, but it is one of the most important steps in resolving a problem. Many layers of technology can exist between the reporting and failing components. Networks, disks, and drivers are only a few of the components to consider when you are investigating problems.

The following questions help you to focus on where the problem occurs to isolate the problem layer:

- Is the problem specific to one platform or operating system, or is it common across multiple platforms or operating systems?
- Is the current environment and configuration supported?

Remember that if one layer reports the problem, the problem does not necessarily originate in that layer. Part of identifying where a problem originates is understanding the environment in which it exists. Take some time to completely describe the problem environment, including the operating system and version, all corresponding software and versions, and hardware information. Confirm that you are running within an environment that is a supported configuration; many problems can be traced back to incompatible levels of software that are not intended to run together or have not been fully tested together.

When does the problem occur?

Develop a detailed timeline of events leading up to a failure, especially for those cases that are one-time occurrences. You can most easily do this by working backward: Start at the time an error was reported (as precisely as possible, even down to the millisecond), and work backward through the available logs and information. Typically, you need to look only as far as the first suspicious event that you find in a diagnostic log; however, this is not always easy to do and takes practice. Knowing when to stop looking is especially difficult when multiple layers of technology are involved, and when each has its own diagnostic information.

To develop a detailed timeline of events, answer these questions:

- Does the problem happen only at a certain time of day or night?
- How often does the problem happen?
- What sequence of events leads up to the time that the problem is reported?
- Does the problem happen after an environment change, such as upgrading or installing software or hardware?

Responding to questions like this helps to provide you with a frame of reference in which to investigate the problem.

Under which conditions does the problem occur?

Knowing which systems and applications are running at the time that a problem occurs is an important part of troubleshooting. These questions about your environment can help you to identify the root cause of the problem:

- Does the problem always occur when the same task is being performed?
- Does a certain sequence of events need to occur for the problem to surface?
- Do any other applications fail at the same time?

Answering these types of questions can help you explain the environment in which the problem occurs and correlate any dependencies. Remember that just because multiple problems might have occurred around the same time, the problems are not necessarily related.

Can the problem be reproduced?

From a troubleshooting standpoint, the *ideal* problem is one that can be reproduced. Typically, problems that can be reproduced have a larger set of tools or procedures at your disposal to help you investigate. Consequently, problems that you can reproduce are often easier to debug and solve. However, problems that you can reproduce can have a disadvantage: If the problem is of significant business impact, you do not want it to recur. If possible, re-create the

problem in a test or development environment, which typically offers you more flexibility and control during your investigation.

- Can the problem be re-created on a test system?
- Are multiple users or applications encountering the same type of problem?
- Can the problem be re-created by running a single command, a set of commands, a particular application, or a stand-alone application?

Troubleshooting software discovery

Problems with software discovery usually fall into one of three categories. The most common category is when the software is discovered but the components are assigned to an incorrect product. In this case, it is enough to reassign the component. The remaining categories include false-positive discovery, or no discovery due to the lack of the signature on the computer.

Incorrect software is displayed on reports due to incorrect bundling

Automated bundling assigns discovered software components to products based on bundling rules. The rules assess the probability with which a component is part of a particular product. Because automated bundling takes into account all products that are available in the software catalog, not only products that you purchased, it might happen that the component is assigned to a product that you do not have. As a result, the product that you have is not displayed on the reports, and some other product is displayed instead. To solve the problem, reassign the discovered component to the correct product.

1. Identify what components are used to discover the product that is not shown on the reports.
For example, check unconfirmed components that are installed on the computer where the product should be discovered. Such components are likely to be used to discover the product.
 - a. Log in to BigFix Inventory, and go to **Reports > Software Classification**.
 - b. Hover over the **Manage Report View** icon , and click **Configure View**. Then, filter the report to narrow down the list of instances to the unconfirmed components that are installed on a particular computer. Set up the following filters.
 - Confirm equal to No
 - Computer Name contains Computer Name
2. If any of the components was assigned to an incorrect product, reassign it to the correct one.
For example, if the component IBM WebSphere Application Server was assigned to the product IBM Rational Software Architect, but the product that you have is IBM WebSphere Application Server Network Deployment, reassign the component to IBM WebSphere Application Server Network Deployment.
 - a. Hover over the **Assign** icon , and click **Choose from Catalog**.
 - b. Select the product to which you want to reassign this component instance, and click **Assign**. For more information, see: [Assigning components to products \(on page dx/vj\)](#).

After you reassign the component, the correct product is displayed on the reports.

Software that you do not have is discovered due to false-positive discovery

It might happen that an older version of the software is discovered after the upgrade or the software signature is not removed from the computer during the uninstallation. False-positive discovery can also be caused by backups of software directories or installation media that are stored on the endpoints and are reported as separate software instances.

1. [Update the software catalog \(on page cdxliv\)](#) to the latest version.
2. Check what signature caused software detection.
 - a. In the top navigation bar, go to **Reports > Software Installations**.
 - b. Find the software that was discovered, and click the link in the **Details** column. Information about the signature or software ID tag that caused the detection and its location is displayed at the bottom of the page.



Check whether the signature or software ID tag that caused the detection was discovered in a backup directory or a directory that contains installation media. If yes, compress the directory, delete it, or exclude from software scans.

Otherwise, open a PMR and file it against the product that was discovered and provide the following information:

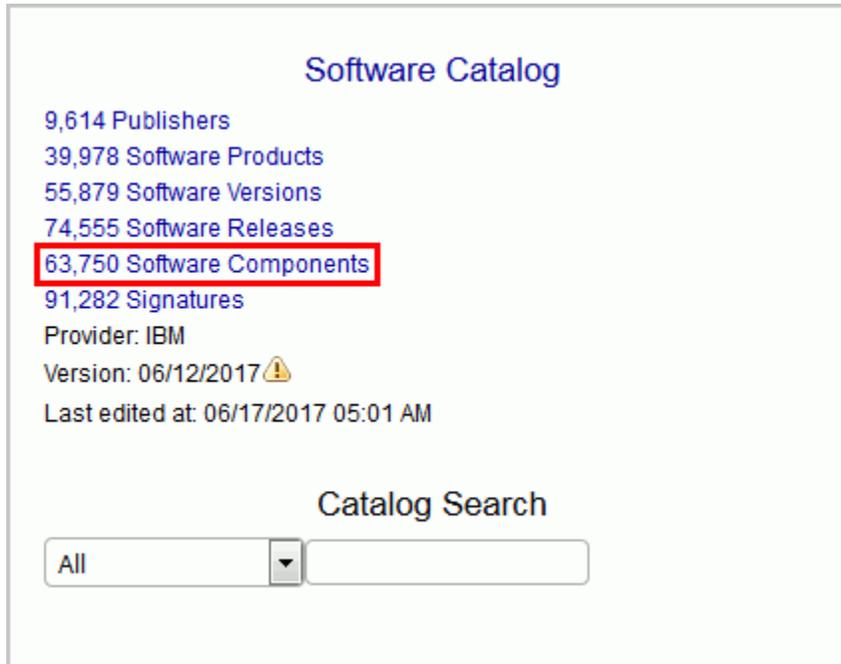
- Exact name of the product that was installed, including its release
- Signature or software ID tag that caused the discovery

Software that you have is not discovered

It might happen that software that is installed is not discovered because the signature or the software ID tag does not exist on the computer where the software is installed.

1. [Update the software catalog \(on page cdxliv\)](#) to the latest version.
2. Check what signature or software ID tag is used to discover the software.

- a. On the **Software Catalog** widget, click **Software Components**.



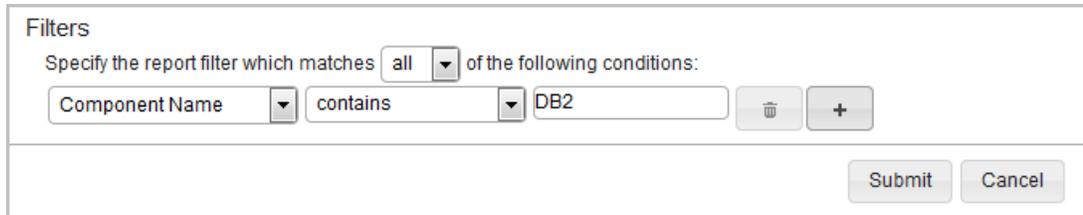
Software Catalog

9,614 Publishers
 39,978 Software Products
 55,879 Software Versions
 74,555 Software Releases
63,750 Software Components
 91,282 Signatures
 Provider: IBM
 Version: 06/12/2017 ⚠️
 Last edited at: 06/17/2017 05:01 AM

Catalog Search

All

- b. To filter out the report, hover over the **Manage Report View** icon , and click **Configure View**. Then, specify appropriate filters. For example, `Component Name, contains, DB2`.



Filters

Specify the report filter which matches of the following conditions:

- c. To view signatures that are used to discover this component, click the link in the **Component Name** column.

Component Name	Component Version	Total Signatures
DB2 Connect Enterprise Edition	8.1	3 Signatures
DB2 Connect Enterprise Edition	9.1	4 Signatures
DB2 Connect Enterprise Edition	7.1	4 Signatures

3. Check whether any of the signatures or the software ID tag exists on the computer where the software was not discovered.

- If the number of signatures is 0, the software is discovered by a software ID tag. Check whether `.swidtag` file is in the installation directory of the software that was not discovered.
- If the number of signatures is greater than 0, check whether any of discovery signatures exists on the computer.

▶ DB2 Connect Enterprise Edition 9.1

Signatures
4 Signatures

Software Installations
0 Software Installations on 0 Computers

▼ Component Signatures

Signature: 0f45c176-a724-43ab-85d9-7dbc56f66459

Definition
A computer running Windows has:
File Name: DB2SYSCS.EXE
Platform WINDOWS Discovery No Usage Yes

Signature: 46f267a6-0e62-41f6-9392-c57171647f0c

Definition
A computer running UNIX has:
A signature that uses: File Signature
File name: db2cons91_ee.sys
File size: 50

```
<MultipleInstance><Iterator name="INSTALL_PATH" export="true"><FindFilePathEx name="db2cons91_ee.sys"/></Iterator><Instance><Variable name="IS_INSTALLED" export="true"><FileInfoCompare absolutePath="$(INSTALL_PATH)/db2cons91_ee.sys" field="size" value="50" relation="eq"/></Variable><Condition withVariable="IS_INSTALLED"><Action do="SKIP" onValue="false"/></Condition></Instance></MultipleInstance>
```

Platform UNIX Discovery Yes Usage No

Signature: 6905d49b-8890-4210-a907-bb3cbc57dae4

Definition
A computer running Windows has:
File Name: DB2CONSV91_EE.SYS File Size: 50
Platform WINDOWS Discovery Yes Usage No

If none of the signatures nor the software ID tag exists on the computer, open a PMR and file it against the product that was not discovered. If a signature or the software ID tag exists on the computer but the software is not discovered, open a PMR and file it against BigFix Inventory. In both cases provide the following information:

- Exact name of the product that was not discovered, including its release
- Part number of the product that was not discovered
- Platform on which the product was not discovered

Troubleshooting problems with computers

When you encounter problems with scan results not being uploaded from a particular computer or you need to collect all logs from that computer, use the **Computer Support Data** panel. The panel is available in BigFix Inventory and is used to perform some troubleshooting actions without the need of using the BigFix console.

9.2.14 Checking whether the maximum archive size is exceeded

You can check whether the maximum archive size is exceeded through the Computer Support Data panel.



You must be an Administrator or have the Manage Support Data and View Endpoints permissions to perform this task.

1. Log in to BigFix Inventory, and go to **Reports > Computers**.
2. Select the name of the affected computer and then click **Computer Support Data**.



3. Check the information in the **Computer refresh** section to verify whether the maximum archive size is exceeded.
If it is, follow the documentation to resolve the problem. For more information, see: [Configuring VM manager for subcapacity reporting \(on page cccxcvii\)](#).

9.2.12 Refreshing scan data on a computer

If software scan, capacity scan, or VM Manager Tool scan finishes successfully on a particular computer but its results are not uploaded to the BigFix server, force the upload of the data. You can do it directly from BigFix Inventory without the need of accessing the BigFix console and running a separate fixlet for each type of scan.

 You must be an Administrator or have the Manage Support Data and View Endpoints permissions to perform this task.

1. Log in to BigFix Inventory, and go to **Reports > Computers**
2. Select the name of the computer from which you want to force the upload of scan results, and then click **Computer Support Data**.



3. Click **Request Data Refresh**.

Refresh of data forces upload of scan results from the BigFix client to the BigFix server.

Computer refresh

Refresh of data forces upload of scan results from the BigFix client to the BigFix server.
New data awaits importing - [please run an import](#).

Hardware scan results - file(s) updated at: 2018-05-18 09:04:02 UTC
Software scan results - file(s) updated at: 2018-05-30 08:46:18 UTC
VM Manager Tool results - data not refreshed yet.

Request Data Refresh

4. When the data is refreshed, run an import to make the data available in BigFix Inventory.

Collecting logs for troubleshooting purposes

If you are asked by BigFix Support to provide data for troubleshooting problems with computers in your infrastructure, go to the details of the computer. Open the Computer Support Data panel, download the log package, and provide it

to BigFix Support. The package contains log files that are needed for troubleshooting purposes such as BigFix client logs, scanner logs, scan configuration logs, VM manager configuration files, and files generated during scans.

 You must be an Administrator or have the Manage Support Data and View Endpoints permissions to perform this task.

If you want to collect the logs, follow the steps that are described below. The steps might be different for particular version. Choose the path suitable for your version.

- **9.2.11** For version 9.2.11 and higher, perform the following steps.

1. Log in to BigFix Inventory, and go to **Reports > Computers**.
2. Select the name of the computer from which you want to collect the logs, and then click **Computer Support Data**.



3. Download the logs.
 - If you are requested to provide only basic support data, click **Download**.
 - If you are requested to provide basic and supplementary data, click **Refresh Supplementary Data**. Then, refresh the page to see when the data is collected. When the supplementary data is available, click **Download**.

Log collector

Basic data is available. Supplementary data available from: 2018-03-08 09:35:14 UTC

Download

Refresh Supplementary Data

4. **Optional:** By default, the size of downloaded packages is limited to 20 MB. If you are not able to download the log package, change the value of the **_BESRelay_UploadManager_CompressedFileMaxSize** parameter. It is a one-time setup. To change the value, perform the following steps.
 - a. Log in to the BigFix console.
 - b. In the left navigation tree, click **Computers**.
 - c. Right-click the following computer, and click **Edit Computer Settings**.
 - BigFix relay with which the computer that cannot send the logs is communicating.
 - If there is no relay, computer on which the BigFix server is installed.
 - d. Set the value of the **_BESRelay_UploadManager_CompressedFileMaxSize** parameter to 100000000 (100 MB). If the parameter is not listed, click **Add** to add it.

```
_BESRelay_UploadManager_CompressedFileMaxSize=100000000
```

5. Provide the log package to BigFix Support.

- **9.2.8** For versions 9.2.8 to 9.2.10, perform the following steps.

1. Run a fixlet that creates a log package on the selected computer. The fixlet is available starting from site version 99.
 - a. Log in to the BigFix console.
 - b. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
 - c. Select **Collect Logs from Endpoints**, and click **Take Action**.
 - d. Choose the computer for which you need to provide the troubleshooting data, and click **OK**.
2. Download the log package.
 - a. Log in to BigFix Inventory.
 - b. In the navigation bar, click **Reports > Computers**.
 - c. Click the name of the computer on which the fixlet was run.
 - d. When the detailed report for that computer opens, add `/mustgather` to the report URL in the web browser to download the package to your computer. For example:

```
https://<host>:<port>/sam/computers/1/mustgather
```

3. Provide the log package to BigFix Support.

Disaster recovery

To perform disaster recovery for your BigFix environment, you need to regularly back up the BigFix Inventory database and the BigFix platform server.

Regular back up

Ensure you regularly back up:

- BigFix server.
- BigFix Inventory database.



Note: The BigFix Inventory data is stored in the database. Thus, you do not have to additionally back up the BigFix Inventory server.

Disaster recovery

To perform the disaster recovery of your BigFix Inventory environment:

- Restore the BigFix server. For more information, see: [Running server backup and restore](#)
- Restore your BigFix Inventory database. For more information, see: [Backing up and restoring the database \(on page cdxxiv\)](#)

Common problems

Learn how to solve some common problems with the server installation, configuration and administration.

The following list contains descriptions of common installation problems:

Agent problems

The **Deployment Health** widget provides the status of computers that do not meet one or more of the deployment health conditions. To resolve deployment health issues, perform the steps relevant for your particular problem. The BigFix console provides fixlets that can be run to resolve some deployment health issues. Issues that cannot be resolved by using fixlets must be resolved manually.

Disconnected: The computer has not connected to the server for the last 72 hours

These computers did not connect to the server for the last three days. You must check whether each computer listed in the report is up and running or why it is not connected to your network.

 **Important:** When you decommission a computer in your environment, the software installed on this computer is still listed on the metric reports. Remove the computer from BigFix database, to reflect the status change in the reports. For more information, see: [Removing inactive computers from BigFix \(on page cdxxxvii\)](#).

Low disk space: The drive on which the agent is installed has less than 100 MB of free disk space

Computers with low disk space have less than 100 MB of free space available. To identify the affected computers:

1. Log in to the BigFix console.
2. Click **Sites > External sites > BES Support > Fixlets and Tasks**.
3. In the upper right pane, click **BES Client Computer is Low on Free Disk Space**. The list of computers is displayed.
4. Free up disk space on each affected computer.

Out of sync: The time that is set on the computer differs by at least an hour from the time that is set on the server

Out of sync computers have incorrect system clock times. To update the clocks on affected computers:

1. Log in to the BigFix console.
2. Click **Sites > External sites > BES Support > Fixlets and Tasks**.
3. In the upper right pane, select **BES Clients Have Incorrect Clock Time** and click **Take Action**.
4. In the lower right pane of the window that opens, select the computers whose clocks you want to synchronize and click **OK**.

Missing prerequisites: Scanner prerequisites are not installed on the computer

Linux™ computers with missing prerequisites do not have the correct C++ runtime library versions installed.

1. Log in to the BigFix console.
2. Click **Sites > External sites > IBM BigFix Inventory v9 > Analyses**.
3. In the upper right pane, select the **Scanner Information** analysis and then click **Results** to see which computers are missing the correct C++ runtime library.
4. Install the C++ runtime libraries manually on each computer that is returned by this analysis. The required version of the `/usr/lib/libstdc++.*` or `/usr/lib64/libstdc++.*` library is 6.

If the scanner is still experiencing problems, ensure the endpoints have the correct version of the bzip2 archiver installed.

1. Log in to the BigFix console.
2. Click **Sites > External sites > IBM BigFix Inventory v9 > Analyses**.
3. In the upper right pane, select the **Software Scan Status** analysis and then click **Results** to see which computers return the Invalid archiver status.
4. Copy the correct version of the bzip2 archiver manually to each computer that is returned by this analysis, to the `/LMT/CIT/custom` folder. Name the file:

- **UNIX** bzip2
- **Windows** bzip2.exe

The correct version of the bzip2 archiver for your operating system can be obtained from your operating system vendor. It can also be compiled on the endpoint with resources from [bzip](#).

Catalog problems

Catalog problems are related to the inability to download a new catalog or upload it to BigFix Inventory.

The software catalog cannot be downloaded because the Software Catalog Update task is not applicable.



Note: The Software Catalog Update task is not relevant for BigFix Inventory versions before 9.2.11.

When you install BigFix Inventory, the target computer sends information about the new software to the BigFix server. Then, the BigFix server verifies that the computer is linked to this server. If it is, the server marks it with the custom setting `SUA_Server_Path_[user_ID]`. If the setting is set to the BigFix Inventory installation path, the task is applicable.

If the setting is not specified for a particular computer, the Software Catalog Update task might not be applicable. If you are sure that BigFix Inventory is installed on this computer, specify the setting manually.

1. Log in to the BigFix console that is linked to your BigFix server.
2. In the navigation bar, click **Computers**.

3. In the right-upper pane, right-click the computer on which you installed BigFix Inventory, and click **Edit Computer Settings**.
4. To add a setting, click **Add**.
5. In the **Setting Name**, type `SUA_Server_Path_{user_ID}`.



Note: `{user_ID}` is the ID of a Linux™ user who installed BigFix Inventory.

6. In the **Setting Value**, specify the BigFix Inventory installation path.
7. Click **OK**.

If the `SUA_Server_Path_{user_ID}` setting is set properly and the problem persists, ensure that the following conditions are fulfilled:

- The BigFix Inventory server is installed on a computer with a single static IP.
- The computer on which the BigFix Inventory server is installed can decompress the downloaded catalog file.
- A manually downloaded catalog is not in the `/opt/ibm/BFI/bfi_catalog` directory.
- The BigFix server can connect to the BigFix Inventory server by using port 9081.
- The BigFix Inventory server can connect to the BigFix server by using port 52311.
- The BigFix Inventory server has an entry for the BigFix server in the `etc/hosts` file.
- The BigFix server has an entry for the BigFix Inventory server in the `etc/hosts` file.
- The user that you provided when you configured the connection to the BigFix server has the right permissions to run the fixlet. For more information, see the topic [Operators permissions](#) in the BigFix documentation.

The Catalog Download task fails because the BigFix server cannot connect to the BigFix Inventory server.

The problem occurs when the BigFix and BigFix Inventory servers are installed in separate networks and the automatic lookup of their IP addresses does not work correctly. To ensure that IP addresses of the servers are correct and the communication between them works, change the default settings, and manually enter the IP addresses or domain names of the servers. For more information, see: [Configuring servers in separate networks \(on page clxxx\)](#).

When you upload a software catalog file, the following error is displayed: `Upload failed. The uploaded file does not appear to be a valid catalog. On inspection, file appears truncated.`

If there is a delay when initially downloading the catalog file from the server, then the downloaded file might be truncated. For example, if more than 30 seconds is taken to rename and save the file. This problem depends on how your browser handles the downloading of large files. To solve this problem, accept and save the file within a short time period or use a different browser.

Backups of software directories are reported as separate products.

If you store backups of software directories on your endpoints, they might be reported as separate software instances resulting in false discoveries and incorrect license consumption. To avoid this

problem, either [exclude the backups from scanning \(on page ccxliv\)](#) or compress them with a data compressor, such as zip or rar.

A duplicated software component is created during a second import of catalog customizations.

If you import a custom catalog definition that contains unique IDs to another BigFix Inventory instance, then modify a publisher or vendor in the signature xml file and import it again, the second import will create a new software component in the hierarchy. The signature definition will be removed from the original hierarchy.

This behavior is caused by using the same signature **uniqueid**. When the custom catalog definition XML file reuses a signature **uniqueid** that exists in the database, the signature is updated to reflect the new software component. To correct the data in the catalog, delete the first catalog entry and ensure that the new entry contains the needed data or reimport the new catalog with a different unique signature ID.

The number of discovered software instances is higher than the number of instances that are installed.

The problem might occur when the product is discovered both by a custom signature and a signature from the catalog provided by BigFix. It happens when you create a custom signature and later on a signature for the same product is delivered with the software catalog provided by BigFix. Because the catalogs are fully independent, the two signatures are not merged.

To verify whether it is the case, perform the following steps:

1. In the top navigation bar, click **Reports > Software Installations**.
2. Hover over the **Manage Report View** icon , and click **Manage Report View**. Then, select the **Component Definition Source** check box.
3. Filter the report to the product for which the number of discovered instances was too high.
4. Check whether the catalog that is listed in the **Component Definition Source** column is Custom for some instances and some other type for other instances of a particular product. If yes, delete the custom signature for this product.

The number of discovered instances should be normalized after the next software scan and import of data.

Part numbers are not matched with products after you upload the part number file.

The problem might be related to the outdated application data.

- Ensure that import of data was completed after the upload of the part number file. Log in to BigFix Inventory, and go to **Reports > Part Number Upload**. Look at the Upload and Import History and check the status of the latest action. After import completes, it should be marked as successful.
- Check if you have the latest version of the software catalog. If not, [update the software catalog to the latest version. \(on page cdxliv\)](#)

- Make sure that the list of discovered software is up-to-date by checking the date of the last scan in the Software Scan Status analysis in the BigFix console. To update the inventory data, run the Initiate Software Scan fixlet. Set up a single scan by clearing the Reapply this action check box on Execution tab.

After applying any changes, upload the part number file again and run the import of data.

The version of the software catalog is displayed in different formats in BigFix Inventory and in the BigFix console.

The catalog upload process is automatized and the task that triggers the catalog upload is changes. Because of these changes, the software catalog can be displayed in a date format or as a fixpack version. To ensure the version consistency, upgrade the BigFix Inventory server to the latest version.

The `<not set>` value appears under the Product Name on Software Classification panel

The product name is not set if the component does not have any product relation that is defined in the software catalog. The component is detected on its own.

To resolve this issue, make sure that you have the latest version of BigFix software catalog.

- To check your current version of software catalog, see: [Verifying the current version of BigFix Inventory and the software catalog \(on page clxxxii\)](#).
- To update the software catalog, see: [Updating the software catalog to ensure accuracy of software discovery and reporting \(on page cdxliv\)](#).



Note: If upgrading the software catalog does not help, contact the relevant product team and ask them to register the product. You can use custom bundling to create a new relation as a temporary workaround. For more information, see: [Assigning components to products \(on page dxlvi\)](#).

Updating the software catalog to ensure accuracy of software discovery and reporting

The software catalog is automatically updated during every upgrade of . However, as an exception, you can manually update the software catalog for troubleshooting purposes.

- **9.2.11** Starting from application update 9.2.11, the Software Catalog Update task automatically triggers the following actions:
 - The catalog is downloaded to the following directory:
 - **Linux** `/opt/ibm/BFI/wlp/usr/servers/server1/data/sam/catalog/BFI`
 - **Windows** `C:\Program Files\IBM\BFI\wlp\usr\servers\server1\data\sam\catalog\BFI`
 - The catalog is automatically uploaded to the BigFix Inventory server during the next import of data. No additional actions are required.

- For versions before 9.2.11, the software catalog is downloaded with the Software Catalog Update task but it still needs to be manually uploaded to the BigFix Inventory server according to this procedure.  You must have the Manage Uploads permission to perform this task.
- If the computer where the BigFix server is installed does not have the Internet access, ensure that the content of the BigFix Inventory site is up-to-date before you download the catalog. For more information, see: [Updating the fixlet site \(on page cccvii\)](#).

The software catalog is called `BFI_Catalog_version-timestamp.zip` and consists of the following files:

- Software catalog in canonical 2.0 format: `IBMSoftwareCatalog_canonical_2.0_form_date.xml`.
- **9.2.5** Software catalog in canonical 3.0 format: `IBMEExtendedCatalog_canonical_3.0_form_date.xml`. The file consists of low priority applications; for example, open source software, freeware, and games. The content of the file will be expanded to include high priority business applications with time. The file is included in each software catalog, however, it is ignored by BigFix Inventory versions lower than 9.2.5.
- Charge unit data file that contains information about charge unit definitions, their relations with products, and some additional parameters: `charge_unit_parameters_dataversion_version.csv`.
- Part numbers file that contains information about part numbers that are used for software licensing purposes: `part_numbers_dataversion_version.csv`.
- Catalog metadata file that describes the software catalog: `catalog_package.properties`.
- **9.2.13** FlexPoint bundles file that contains information about products that can be assigned to each of the available FlexPoint Offerings: `flexpoint_bundles_dataversion_version.csv`
- **9.2.11** For versions 9.2.11 and higher
 1. Log in to the BigFix console.
 2. In the navigation bar, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
 3. In the upper-right pane, select **Software Catalog Update**, and then click **Take Action**. The action is applicable only on the computer on which the BigFix Inventory server is installed. Select the computer, and click **OK**.
If the task is not applicable on that computer, see: [Catalog problems \(on page dcclviii\)](#).
 4. Wait for the scheduled import of data or run it manually.
- For versions before 9.2.11
 1. Download the software catalog
 - a. Log in to the BigFix console.
 - b. In the navigation bar, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.

- c. In the upper-right pane, select **Software Catalog Update**, and choose the option to download the software catalog for BigFix Inventory version before 9.2.11.0.
 - d. Optional: Copy the file to the computer from which you access the BigFix Inventory web user interface.
2. Upload the software catalog to BigFix Inventory.
 - a. Log in to BigFix Inventory.
 - b. In the navigation bar, click **Management > Catalog Upload**.
 - c. Click **Browse** and select the software catalog file.
 - d. To upload the file, click **Upload**. The software catalog file and the charge unit data are listed in the **Upload and Import History** table. Their status is **Pending**.
 - e. Wait for the scheduled import of data or run it manually.

During the import, scanner catalogs that are used for software discovery are created and automatically distributed to the computers in your infrastructure.

If the automatic distribution of scanner catalogs fails, computers on which the catalog was not updated have the **Outdated Catalog** status on the **Scan Health** widget. You must [manually update scanner catalogs \(on page cclxvii\)](#) on these computers. If you use the disconnected scanner, [distribute the downloaded software catalog \(on page cclv\)](#) to each computer.

Data import problems

Data import problems are related to situations in which the data import fails or the data that is displayed in reports is inconsistent as a result of a failed import.

During the initial import, the following error is written in the logs: `Error was getaddrinfo: name or service not known (SocketError)`.

During the initial import, the following error is written in the logs:

```
ERROR: Datasource file citsearch_0_4580013_cit.xml.bz2 raised an exception
while reading from {:port=>"52311", :path=>"/UploadReplication",
:query=>{:BaseDirectory=>1,
:Name=>"\\13\\4580013\\citsearch_0_4580013_cit.xml.bz2",
:shal=>"5B0FE15F7E097171566F0AC3B9BE93826FDC0D41", :offset=>0}}.
Error was getaddrinfo: name or service not known (SocketError)
```

The problem might be caused by incorrect DNS name settings. To solve the problem, ensure that BigFix Inventory can ping the BigFix server by using the DNS name that is specified in the fixlet site. To find the DNS name, on the computer where the BigFix server is installed, go to the `C:\Program Files (x86)\BigFix Enterprise\BES Installers\Server`, and find the `ActionSite.afxm` file. If BigFix Inventory cannot ping the BigFix server by using this DNS name, add the name to the `etc\hosts` file on the BigFix Inventory server.

The import fails. After you rerun the import, the software inventory is empty.

One of the scenarios in which the problem occurs is when you run an import and the BigFix server is not running. After you restart the server and rerun the import, the software inventory is empty. To solve the problem, manually initiate the scanner. Gather fresh scan data, and run the import.

As an alternative, you can run the Force Reupload of Software Scan Results task and then run the import of data. This task forces reupload of inventory data that was gathered by the software inventory and the file system scans to the BigFix server. The data is then imported to BigFix Inventory.

The import fails because the transaction log is full.

After a failed import, the import log contains the following error:

```
Batch failure. The batch was submitted, but at least one exception
occurred on an individual member of the batch. Use getNextException()
to retrieve the exceptions for specific batched elements.
ERRORCODE=-4229,
```

Also, the following error can be found in the `tema.log` file:

```
Batch execution error: Error for batch element #903:
The transaction log for the database is full.
```

To solve the problem, complete the following steps:

1. [Increase the size of the transaction log \(on page dcccxlvj\)](#) for the database.
2. Restart the DB2® and BigFix Inventory servers.

The import fails because the Java™ heap size is too low.

After a failed import, the import log contains the following error:

```
E SRVE0777E: Exception thrown by application class
'java.lang.StringBuilder.ensureCapacityImpl:342'
java.lang.OutOfMemoryError: Java heap space
```

To solve the problem, [increase the Java heap size \(on page dcccxliv\)](#).

The import fails and the following message is written in the logs: Overflow occurred during numeric data type conversion.

The problem occurs when you create a contract custom field that requires an integer value, and then enter a value that is greater than 32767. To solve the problem, enter a smaller value.

The import fails and the following message is written in the logs: The SQL statement or command failed because of a database system error. (Reason "optBitFilterSize less than min"). SQLCODE=-901, SQLSTATE=58004, DRIVER=4.14.111.

To solve the problem, create a script that reorganizes internal database tables and keeps statistics up-to-date, and run it against the BigFix Inventory database.

1. Create the `reorg.sh` script.

```
$ cat reorg.sh
#!/bin/ksh
db2 connect to TEMADB
db2 -x "select 'reorg table',substr(rtrim(tabschema)||'.'||rtrim(tabname),1,50),
'allow no access;'from syscat.tables where type = 'T' and tabschema not in
(NULLID,'SYSCAT','SYSFUN','SYSIBM','SYSPROC','SYSSTAT') order by tabschema,tabname
" > reorgs.sql
db2 -tvf reorgs.sql
db2 terminate

db2 connect to TEMADB
db2 -x "select 'runstats on table',substr(rtrim(tabschema)||'.'||rtrim(tabname),1,50),
' and indexes all;'from syscat.tables where type = 'T' and tabschema not in
(NULLID,'SYSCAT','SYSFUN','SYSIBM','SYSPROC','SYSSTAT') order by tabschema,tabname
" > runstats.sql
db2 -tvf runstats.sql
db2 terminate
```

2. Log in as a database instance owner to the computer where the DB2® is installed and run the script.

The import fails and the following message is written in the logs: INFO: ETL from Data Source `data_source_name`
- RawDataSourceFixletResult: Failed

The problem occurs because there is not enough disk space on the computer where the BigFix Inventory database is installed. To solve the problem, free some disk space.

The import fails and the following message is written in the logs: Snapshot isolation transaction failed accessing database 'TEMDB' because snapshot isolation is not allowed in this database. Use ALTER DATABASE to allow snapshot isolation.

To solve the problem, enable snapshot isolation in the MS SQL Server. For more information, see the [Microsoft™ SQL Server documentation](#).

After adding a large data source, the import fails and the following message is written in the logs: 500 Internal Server Error.

The problem might occur because there is not enough disk space on the computer where the BigFix database is installed. To calculate the amount of required disk space, perform the following steps:

1. Optimize the import of data from BigFix. Log in to BigFix Inventory and go to **Management > Advanced Server Settings**. Then, change the value of the `schema_next` parameter to true.
2. To calculate the required disk space, check how many objects exist in all fixlet sites that you have enabled in the BigFix console. An object is every computer group, analysis, fixlet, and task that exist in the console, including the ones that are not relevant. Every 1000 of objects requires

1GB of free disk space. For example, if you have 500 fixlets and tasks, 300 analyses, and 20 computer groups, you have 820 objects in total. The BigFix database requires 1 GB of disk space.

The import fails and the following message is written in the logs: INFO: ETL from Datasource -

RawDatasourceAnalysis (0x000000 - 0x00000035): Failed.

During the import, the following error is written in the logs:

```
INFO: ETL from Datasource - RawDatasourceAnalysis (0x000000 - 0x00000035): Failed
ERROR: Sequel::SerializationFailure: Java::ComMicrosoftSqlServerJdbc::SQLException:
Transaction (Process ID 1111) was deadlocked on lock resources with another process
and has been chosen as the deadlock victim. Rerun the transaction
```

To solve the problem, ensure that no actions such as a backup or recovery are taking place on the BigFix database. Then, rerun the import.

The import fails after you change the host name of the BigFix Inventory server.

Changing the host name of the BigFix Inventory server is not supported. When the application cannot recognize the original host name, the ETL step of a data import is failing and you cannot gather and process data from your endpoints. A new host name requires a new installation of BigFix Inventory.

The host name also cannot be changed for the BigFix server. In this case, the host name of the server is recorded into your license certificate during the installation. To change it, you must create a new license certificate that requires a new installation.

The import hangs after the database connection is lost and recovered

During the import of data, database connection is lost and information about connection problems is displayed on the user interface. After you restart the database, the user interface is refreshed but the import hangs. Additionally, the following or similar error is written in the logs.

```
ERROR: Sequel::DatabaseError: DBNAME: temadb25 -
Java::ComIbmDb2JccAm::SqlNonTransientConnectionException:
[jcc][t4][10335][10366][4.22.29] Invalid operation: Connection is closed. ERRORCODE=-4470,
SQLSTATE=08003
```

To solve the problem, restart the BigFix Inventory server.

A successful software scan is run but there are no changes to the data after the data import to the BigFix Inventory server.

Check the maximum archive file size to ensure that it is greater than the scan file size.

1. Log on to the BigFix console.
2. In the left navigation, click **Computers**, right-click the name of the appropriate computer, and then click **Edit Computer Settings**.
3. Check the **_BESClient_ArchiveManager_MaxArchiveSize** setting is greater than the size of the largest scan file. If needed, edit the value to increase the maximum archive size.

Verify that there are no software scan errors.

1. In the BigFix console navigation, click **Sites > External Sites > IBM BigFix Inventory v9 > Analyses**.
2. Select the **Software Scan Status** analysis.
3. In the lower pane, click the **Results** tab, and verify that the status of the software scan is OK for the computers.

Some of the products are not discovered.

If you optimized the volume of scanned file data, either during the post-upgrade configuration or manually, you must run an import for the changes to take effect. After the import, some software items might not be visible on the reports. It is an expected behavior. To ensure that the software inventory is properly reported, perform the following steps.

1. Ensure that the catalog that is uploaded to BigFix Inventory is in the canonical format. If the catalog is in the native format, upload a new catalog. If the catalog is in canonical format but a new version is available, upload the new catalog. To check the format of the uploaded catalog, click **Management > Catalog Upload** and check the **Catalog Format** column.
2. Stop the current scan.
 - a. Log in to the BigFix console and in the left navigation tree, click **Actions**.
 - b. In the upper-right pane, click **Initiate Software Scan** and then click **Stop**.
3. [Initiate a new software scan. \(on page cciii\)](#) Wait for the scan to finish.
4. Wait for the scheduled import or run it manually.

Files with a particular extension are not reported on the Scanned File Data report.

The problem might occur if you optimized the volume of scanned file data and removed the extension from the list of monitored extensions. To solve the problem, follow the steps that are described in the topic [Optimizing the volume of scanned file data \(on page cclxix\)](#). Add the file extension that you want to monitor to the following three files: `file_names_all.txt`, `file_names_unix.txt`, and `file_names_windows.txt`. They are in the `BFI_install_dir\wlp\usr\servers\server1\apps\tema.war\WEB-INF\domains\sam\config` directory.

After a failed import of scan data, the data that is displayed in reports is inconsistent

If there are inconsistencies in the data, and the last data import failed, the data inconsistencies might be the result of the import failure. To solve the problem, run the import again.

As an alternative, you can run the Force Reupload of Software Scan Results task and then run the import of data. This task forces reupload of inventory data that was gathered by the software inventory and the file system scans to the BigFix server. The data is then imported to BigFix Inventory.

The data import log cannot be viewed in the user interface.

The problem might occur if you enabled debug logging for the import process. Debug information significantly increases the size of the import log file and it might happen that the log file cannot be

displayed in the user interface. This problem impacts Internet Explorer and can be solved by accessing the user interface from a different browser. If that does not solve the problem, consider the following solutions:

- Instead of viewing the log file in the user interface, open it as a file. The log file is in one of the following directories:
 - **Linux** `installation_directory/wlp/usr/servers/server1/logs/imports`
 - **Windows** `installation_directory\wlp\usr\servers\server1\logs\imports`

To avoid problems in the user interface, move the log file to a different location so that BigFix Inventory does not load it into the user interface.

- [Server log file \(on page dcccxiv\)](#) to limit the amount of information saved in the log file.

Software is properly discovered by the scan but is not reported in BigFix Inventory after the import.

The problem occurs on BigFix 9.0 installed on Linux™. If the value of the sequence changed in the BigFix database and is higher than the sequence that was imported, scan results are not imported during a particular import. To solve the problem, wait for the next scheduled import or run it manually.

Data about a computer was not imported from the BigFix server.

Because BigFix Inventory imports live data that is constantly changing on the BigFix server, it might happen that some data, for example about computers, might not be imported. That happens because only the scope of data that is calculated at the beginning of an import is processed during the import.

Example: If a computer is saved in the BigFix database during a data import from BigFix to BigFix Inventory, it will not be imported.

 **Important:** It might also happen that after the second data import, the imported computers do not have important properties such as **Computer Name**, **IP Address**, or **Operating System**. To fix this problem:

1. In the navigation tree of the BigFix console, click **Computers** and then in the list pane, select the computer for which the properties are missing.
2. In the lower right pane, click **Send refresh** 3 times. All missing computers properties will be imported during the next data import.

Database problems

Database problems are related to issues with connecting to DB2 and MSSQL Server as well as retrieving information from these databases.

Windows SQL Server

During the import, SQL Server uses 99% of the physical memory.

To solve the problem, set the maximum amount of memory that can be used by SQL Server. For more information, see [Server Memory Options](#) and [How to: Set a Fixed Amount of Memory \(SQL Server Management Studio\)](#) in the SQL Server documentation.

The login process to SQL Server fails.

During a login process to SQL server, the following error is displayed:

```
Login failed for user 'username'. The user is not associated with a trusted SQL Server
connection.
(Microsoft SQL Server, Error: 18452).
```

The cause of this error is that SQL server is configured to use Windows Authentication mode and does not allow the use of SQL accounts. To solve this problem, enable the SQL Server Authentication in SQL Server. For more information, see [Enabling the SQL Server Authentication Mode](#).

Linux DB2

BigFix Inventory cannot connect to its database and the following error is written in the logs: `Connection refused. ERRORCODE=-4499, SQLSTATE=08001.`

The problem occurs because the DB2® database is not yet activated after a restart and is thus not accessible to BigFix Inventory. To solve the problem, run the following command after you restart the database:

```
db2 activate db database_name
```

BigFix Inventory cannot connect to the database. The default port range is modified.

The default local port range for Linux™ is 32768-61000. If the default range is modified, it is possible that the DB2® port number is allocated to another local process. If the port number is pre-allocated, BigFix Inventory cannot connect to DB2®. Complete the following steps:

- Ensure that IPv4 is being used.
- Use the netstat command to check whether DB2® is listening on the expected port, and no other socket is pre-allocated the DB2® port.
- If DB2® is not listening on the expected port, restart DB2® and check again.

It is impossible to establish a connection to a data source in the BigFix Inventory web user interface with single sign-on enabled.

An error message is displayed after you saved the connection parameters on the **Data Sources** pane:

```
Unexpected WebSEAL Response.
Code: 0x38cf04d3
```

```
Error: DPWWA1235E
```

```
Please contact your system administrator. This may indicate an insufficient proxy HTTPS timeout.
```

To be able to save the connection parameters successfully, increase the **https-timeout** parameter in the ISAM Reverse Proxy configuration. For example, to set the timeout to 5 minutes:

1. Log on to IBM Security Access Manager.
2. In the top navigation bar, click **Secure Web Settings > Manage > Reverse Proxy**.
3. Select the instance and then, from the drop-down list on the right of the Reverse Proxy bar, select **Manage > Configuration > Edit Configuration File**.
4. In the **Advanced Configuration File Editor**, locate the **https-timeout** parameter and specify the value of 300.

Example:

```
https-timeout = 300
```

5. Click **Save**.

Both databases

After you restore the BigFix database, new data is not displayed in BigFix Inventory.

To ensure that only newly created data is imported to BigFix Inventory, the data that is stored in the BigFix database is marked with the so-called sequence number. The number is incremented with every change. Every import contains data starting from the sequence number reported during the last successful import to the current sequence number.

When the BigFix database is restored, the sequence number is restored to the value from the time when the database backup was created. Thus, it might happen that the sequence number after the restoration is lower than the sequence number reported during the last successful import. In such case, the first import after the database restoration does not include any data. After that import, the sequence number is updated and the next import contains new data.

However, data from the period between the database restoration and the first successful import is not imported which creates a gap. To solve the problem, go to **Management > Advanced Server Settings**, and change the value of the **resynchronize_datasources_once** parameter to true. Then, run an import. This import will be time-consuming as it will import all software scan data from the restored BigFix database to cover the gap. The subsequent imports will run in the normal operation mode.

Checking the database name.

The default database name is `TEMADE`, however it changed since earlier versions, where the database was called `SUADB`. You can check the name of your database in the `database.yml` file in one of the following directories.

```
Linux /opt/ibm/BFI/wlp/usr/servers/server1/config/database.yml
```

Windows C:\Program Files\IBM\BFI\wlp\usr\servers\server1\config
 \database.yml

In earlier versions, the default installation path is C:\Program Files\IBM\SUA or /opt/ibm/SUA.

Installation and upgrade problems

Solve common problems that are related to installation, configuration, and upgrade of BigFix Inventory.

Installation of the server cannot continue and an error is displayed.

Installation of the BigFix Inventory server cannot continue and the following message is displayed:

```
./tools/getArch: line 108: print: command not found
./tools/getArch: line 109: print: command not found
./tools/getArch: line 116: print: command not found
setup-server-9.0-linux-x86_64.sh: line 52: print: command not found
```

The error might be caused by the fact the operating system is not fully configured. You might also need to reboot the operating system.

Installation of the BigFix server with a remote SQL Server database fails.

The problem might occur while installing the BigFix server on the Windows 2008 Server operating system. When the installer attempts to connect to the remote SQL Server database, the following message is displayed:

```
Computer Browser Error with Windows Authentication
```

To solve the problem, enable the file and printer sharing in Windows:

1. Go to **Control Panel > Network and Sharing Center**.
2. Click **Advanced sharing settings**.
3. Select the **Turn on file and printer sharing** check box.
4. Click **Save changes**.

Installation of the BigFix Inventory server does not start on Linux.

The problem occurs when the Symantec process **rtvscand** is running. To solve the problem, you can stop the **rtvscand** process or exclude the `/tmp` folder from Symantec scan. For more information about excluding the `/tmp` folder, see: [How to Configure Scan Exclusions in SEP for Linux](#).

Installation fails and a warning about the lack of free disk space in a given directory is displayed. However, there is enough free space in that directory.

To solve the problem, perform the following steps.

1. Open the command line.
2. Run the following command to change the variable that is responsible for checking disk space requirements.

- **Linux** `export tlm_debug_disable_disk_space_check=true`
- **Windows** `set tlm_debug_disable_disk_space_check=true`

3. Run the following command to start the installation.

- **Linux** `./setup-server-linux-x86_64.sh`
- **Windows** `setup-server-windows-x86_64.bat`

Installation cannot be started, because the `/tmp` directory has the `NOEXEC` flag.

In some environments, the `/tmp` directory might be mounted with the `NOEXEC` flag, which blocks the installation. The problem can be recognized by the following error in the `ia.log` file.

```
Launching installer...

./server/parts/CDROM_Installers/Disk1/InstData/Linux_64/VM/setupServer.b
in: line 3318: /tmp/install.dir.20400/Linux/resource/jre/jre/bin/java:
Permission denied
```

The `/tmp` directory is used to store temporary files during the installation. Since the BigFix Inventory installer is based on InstallAnywhere, you can change the `IATEMPDIR` environment variable that specifies the temporary directory of InstallAnywhere, and set it to a custom directory.

1. Create a custom temporary directory, for example:

```
mkdir /root/Install_tmp
```

2. Set the `IATEMPDIR` variable to the new directory:

```
export IATEMPDIR=/root/Install_tmp
```

3. Restart the installation.

During the configuration of the connection to a remote DB2 database, the following error is displayed: `The user specified for the communication with the database could not be found in the system.`

The problem occurs when the remote DB2 database is installed on an AIX computer. It means that the user that you want to use for connecting to the database does not exist on the LDAP server. To solve the problem, specify a different user or create a user with the specified credentials on the LDAP server.

During the installation of BigFix Inventory for non-English-language locales, some Java exceptions are displayed in English.

Even for non-English-language locales, some Java exceptions that might occur during the installation are displayed in English. However, the Details view that contains the exceptions also includes more information that can help you to understand and solve the issue. If you want to see the translated

description of a problem, you can switch to the Problems view where all the available details are provided.

During the configuration of the connection to the BigFix server, the following message is displayed: Could not determine character set of the BES database. Are you sure the BES server is running an agent?.

The problem occurs when the BigFix client was removed from the computer where the BigFix server is installed. To solve the problem, manually set the `_BESClient_DeploymentEncoding_IANAName` property to that of your deployment code page.

The list of non-English languages in the installation wizard is reduced.

To see your language as an option in the installation wizard, change your system locale to a chosen language:

1. On Linux, open the Terminal and run the following command:

```
export LC_ALL=language_code.UTF8
```

For example, export LC_ALL=en_US.UTF8.

2. Run the locale command to verify changes.
3. Restart the installation and choose the language.

Despite upgrading to the latest version, the service name still shows the old version.

The display name of the service on Windows does not change after the upgrade, which is a limitation. It does not, however, impact BigFix Inventory in any way. To update the service name, restart your computer.

The server cannot be upgraded because the task Upgrade to the latest version of BigFix Inventory is not relevant.

The problem occurs when the previous version of the server was installed manually on a computer without the BigFix client. To make the upgrade task relevant, perform the following steps:

1. Install the BigFix client on the computer with the BigFix Inventory server.
2. Set the value of the `SUA_Server_Path_{user_ID}` parameter to the installation path of the BigFix Inventory.
 - a. Log in to the BigFix console that is linked to your BigFix server.
 - b. In the navigation bar, click **Computers**.
 - c. In the upper-right pane, right-click the computer on which you installed BigFix Inventory, and click **Edit Computer Settings**.
 - d. To add a setting, click **Add**.
 - e. In the **Setting Name**, type `SUA_Server_Path_{user_ID}`.



Note: `{user_ID}` is the ID of the user who installed BigFix Inventory.

- f. In the **Setting Value**, specify the BigFix Inventory installation path.
- g. Click **OK**.

Upgrade of the server with a fixlet fails on Linux.

The problem occurs when the server was installed by a non-root user. To solve the problem, upgrade the server in [interactive \(on page cccxi\)](#) or [silent \(on page cccxii\)](#) mode.

After the upgrade, the Usage per Computer report is empty

To solve the problem, run the import of data.

Fields in the user interface are disarranged when using Internet Explorer.

During the initial configuration of BigFix Inventory, some fields in the user interface are doubled and disarranged. To solve this problem, add the BigFix Inventory server to the list of trusted sites in Internet Explorer.

1. Click the Settings icon in the top-right corner of your browser.
2. Click **Internet Options**, and switch to the **Security** tab.
3. Select **Trusted sites**, and then click **Sites**.
4. Enter the web address of your server, and click **Add**.
5. Refresh the page.

Manually changing the port number of BigFix Inventory.

If you cannot access the user interface (**Management > Server Settings**) to change the port number, you can also do it in the `server.xml` file.

1. Go to `install_dir/wlp/usr/servers/server1`, and edit the `server.xml` file.
2. In **httpsPort**, enter the new port number.

```
<httpEndpoint host="*" httpsPort="9081" id="tema">
```

The change takes effect immediately, the server does not need to be restarted.

Uninstalling BigFix Inventory fails because the BigFix Inventory service is unresponsive and cannot be stopped.

You can manually force an unresponsive service to stop.

1. Open the command line.
2. Run the following command:
 - **Linux** `service BFIServer status`
 - **Windows** `sc queryex BFIServer`
3. Identify the PID of the BigFix Inventory process.
4. In the same command line, run the following command:
 - **Linux** `kill -9 identified_PID`
 - **Windows** `taskkill /pid identified_PID /f`

Hardware inventory problems

Find the solution to the issues that you have encountered while working with virtual machine managers and hardware inventory.

Update of the VM Manager Tool fails

When you run the **Update VM Manager Tool to *version*** fixlet, the update fails. When you view the details of the action on a particular computer, the script can fail on various lines, for example:

```
waithidden cmd /c "rmdir "{parameter "homefolder"}" /s /q"
continue if {exit code of action = 0}
```

To solve the problem, perform the following steps:

1. Go to the computer on which the update of the VM Manager Tool failed and open the `<BES_Client>/LMT` folder.
2. If the folder does not contain the `VMMAN_copy` folder, create it.
3. Copy the `config` and `keydb` folders from the `<BES_Client>/LMT/VMMAN` folder to the `<BES_Client>/LMT/VMMAN_copy` folder.
4. Remove the `VMMAN` folder.
5. Run one of the following fixlets to complete the update. It might take some time before the fixlet becomes relevant.
 - Run the **Install VM Manager Tool *version number*** fixlet, if the update failed on the computer where the main instance of the VM Manager Tool is installed.
 - Run the **Install Additional VM Manager Tool (OPTIONAL) *version number*** fixlet if the update failed on the computer where an additional instance of the VM Manager Tool is installed.

Error message CODVM0005E is displayed when attempting to connect to the VM manager over SSL.

When trying to connect to the VM manager over SSL, an error message is displayed: CODVM0005E An error occurred when attempting to connect to the VM manager at the following address: hostName. To solve the problem, complete the following steps:

1. Go to <https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=jcesdk>.
2. Provide your BigFix ID and password and click **Sign in**. You might need to register with BigFix to download the files.
3. Select **Unrestricted SDK JCE Policy files for Java 5.0 SR16, Java 6 SR13, Java 7 SR4 and later versions** and then click **Continue**.
4. View the license agreement, select **I Agree**, and then click **I confirm**.
5. Click **Download Now**.
6. Extract the files and copy them to the following directory:

```
<BES_Client>/LMT/VMMAN/java/jre/security
```

7. Restart the server.

Server is running, but an exception in traces appears. The exception is related to the connection with the specified ESX with the following message: `javax.net.ssl.SSLException: Unrecognized SSL message, plaintext connection?`

1. Add the following lines to file `/etc/vmware/hostd/config.xml`:

```
..
  <ssl>
    <doVersionCheck> false </doVersionCheck>
    <handshakeTimeoutMs>30000</handshakeTimeoutMs>
  </ssl>
  <http>
    <readTimeoutMs>45000</readTimeoutMs>
    <writeTimeoutMs>45000</writeTimeoutMs>
    <blacklistPeriodMs>3000</blacklistPeriodMs>
  </http>
  <vmdb>
    ...
```

2. Restart the hostd service using `service mgmt-vmware restart`.
3. Verify that the exception does not occur.

Server is running, but an `org.xml.sax.SAXParseException` exception in traces appears. The exception is related to the connection with the specified ESX.

Make sure that you have the latest patches to your problematic ESX server installed.

Upgrade of vCenter server from v5.0 to v5.1 causing server connection failure.

The following message is displayed `CODVM0003E The VM manager denied access because of invalid credentials`. To solve this issue, perform the following steps:

1. In vCenter, stop all sessions under a user name that is defined as a user credential in VM manager for specific vCenter.
2. Remove this user name.
3. From vCenter, add the user name back with read-only or propagation authorities.
4. Redefine the VM manager entry for this specific vCenter using the same user name credential.

Specification of the domain for the user name for VM Managers is inconsistent.

Different definitions of users are used for each type of VM manager:

- For Microsoft Hyper-V, you must use the Administrator account. The user is defined as `user_name\domain` or `user_name@domain`. For example: `test\cluster.com` or `test@cluster.com`.
- For VMware, the user is defined as `domain\user_name`, for example: `cluster.com\test`.
- For RHEV-M, the user is defined as `user_name@domain`, for example: `test@cluster.com`.
- **9.2.12** For XenServer and Citrix XenServer, the user is defined as `user_name`, for example `root`.
- **9.2.14** For Oracle VM Server for x86, the user is defined as `user_name`, for example: `test`.
- **9.2.17** For Nutanix, the user is defined as `user_name`, for example: `test`.

Changes to the VM managers are not updated on the BigFix server.

Changes to the VM managers are not updated on the BigFix server and the following error message is displayed in the VM Managers panel: `The last modification of VM managers was not processed correctly on the BigFix server. The data is not synchronized with the VM Manager Tool.`

To solve the problem, make sure that the VM Manager Tool is installed. You can also [check the history \(on page dcclxxviii\)](#) of the Configure VM Manager Tool action in the BigFix console to investigate the failed step. Additional information can also be found in the BigFix console log files that are in one of the following directories:

- **Linux** `/var/opt/BES_Client/___BESData/___Global/Logs`
- **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES_Client___BESData___Global\Logs`

The connection test does not finish.

See the solution for: [Data from VM managers is not updated on the BigFix server. \(on page dcclxxvii\)](#)

Connection to Hyper-V fails with the following error: `The RPC server is unavailable.`

Ensure that port 135 as well as all Windows dynamic ports are open for communication.

Data from VM managers is not updated on the BigFix server.

If you run a connection test for a VM manager and it does not finish, it might indicate a problem with the BigFix client that is installed on the BigFix server. If the client is stopped, actions that you perform in BigFix Inventory are sent to the BigFix server but the status is marked as Not reported. To determine further actions, investigate the BigFix client that is installed on the BigFix server.

For the same reason, data from VM managers might not be uploaded to the BigFix server. Check the value in the **Last Successful Operation** column to verify if the data was recently sent to the server.

Actions that are performed in the VM Managers panel, such as testing the connection or adding a new VM manager, fail.

If actions that you perform in the VM Managers panel fail, you can log in to the BigFix console that is linked to your primary data source and check the history of recent actions, such as Configure VM Manager Tool or VM Manager Tool - connection test. By doing so, you can investigate details of the

failed step and determine the solution. If you are not sure which data source to connect to, log in to BigFix Inventory and click **Management > Data Sources**.

To check the history of recent actions, complete the following steps:

1. Log in to the BigFix console that is linked to your primary data source.
2. In the navigation tree, click **Computers**.
3. In the upper-right pane, select the computer that is defined as your primary data source.
4. In the lower-right pane, click the **Action History** tab.
5. Double-click on one of the recent actions that failed:
 - If the connection test was unsuccessful, check the VM Manager Tool - connection test action.
 - If modification of VM managers failed, check the Configure VM Manager Tool action.
6. In the new window, click the **Computers** tab.
7. Double-click on the action.
8. Locate the failed step and check the details.

The VM Managers panel is blocked or contains error messages.

The VM Managers panel is either blocked completely, with no possibility of performing any actions, or it contains one of the error messages that instruct you to install the VM Manager Tool or the BigFix services.

To solve the problem, complete the following steps:

- Ensure that the BigFix server and client are installed on the target endpoint.
- Install and start Web Reports on the BigFix server. For more information, see: [Installing the Web Reports](#).
- Subscribe the BigFix server to the **IBM BigFix Inventory v9** site.
- Make sure that the content of the **IBM BigFix Inventory v9** site is up-to-date. If your computer does not have access to the Internet, see: [Downloading files in an air-gapped environment](#).
- Subscribe the target endpoint to the **IBM BigFix Inventory v9** site.
- Make sure that the VM Manager Tool is installed. For more information, see: [Installing the VM Manager Tool \(on page ccclxxi\)](#).

CPU frequency that is displayed on the Hardware Inventory report equals zero

CPU frequency is an additional parameter. It is retrieved only from computers that are not managed by VM managers or from virtual machines that are not in the *OK* status. CPU frequency values do not influence PVU calculations.

Virtual machines that are managed by the VMware vCloud Director have duplicate BIOS UUIDs.

The problem occurs when virtual machines are deployed from catalog templates. To work around the problem, see: [BIOS UUIDs in vCloud Director are not unique when virtual machines are deployed from catalog templates \(2002506\)](#).

The Hardware Inventory report in BigFix Inventory displays *No Scan Data* after capacity scans were run on VMware guest operating systems.

The import log contains multiple occurrences of the following errors.

```
Some error occurred during importing the capacity scan from file :
capacity_scan_file_name.xml for endpoint :
endpoint_number.
```

```
getNodeInfo Error: VMWare VirtualMachine UUID (06XZXA0) do not match with the UUID pattern!
```

In the scan files uploaded by the clients, the UUID tag contains the hosts serial number instead of the virtual machine UUID. Example:

```
<VirtualMachineGuest version="1">
  <UUID>06CZFCV</UUID>
  <HypervisorType>VMware</HypervisorType>
</VirtualMachineGuest>
```

BigFix Inventory requires unique virtual machine UUID numbers to properly calculate the capacity for all virtual machines. This error might be caused by the host serial number appearing in the place where virtual machine UUID is normally recorded. If you install an operating system with the so-called Reseller Option Kit (ROK) media, some data might be obtained from the server BIOS, instead of from the virtual machine.

To enable the correct retrieval of UUID data from operating systems installed with the ROK media, edit the virtual machine's `.vmx` file and set the `reflectHost` parameter to `false`. Example:

```
SMBIOS.reflectHost = "false"
```

If the problem persists, force the upload of capacity scan data. To do this, perform the following steps.

1. Log in to the BigFix console.
2. In the navigation tree, go to **Fixlets and Tasks** and select the **Run Capacity Scan and Upload Results** fixlet.
3. In the lower pane, click **Take Action**, and choose **Click here to run a single capacity scan and force upload of results**.
4. Open the **Target** tab and select the computers that you want to scan. Then, click **OK**.
5. Wait for the scheduled import, or run it manually.

Data on the Hardware Inventory report is outdated.

The problem might occur when the import of capacity scan data fails. To solve the problem, perform the following steps:

1. Log in to the BigFix console.
2. In the navigation tree, go to **Fixlets and Tasks** and select the **Run Capacity Scan and Upload Results** fixlet.
3. In the lower pane, click **Take Action**, and choose **Click here to run a single capacity scan and force upload of results**.
4. Open the **Target** tab and select the computers that you want to scan. Then, click **OK**.
5. Wait for the scheduled import, or run it manually.

When the scan finishes, its results are uploaded to the BigFix server. After the data is imported to BigFix Inventory, the Hardware Inventory report should be up-to-date.

The Hardware Inventory report displays no Processor Brand String data.

The problem might occur after the migration, or upgrade. To see the Processor Brand String data on the report, force the upload of capacity data by completing the following steps.

1. Log in to the BigFix console.
2. In the navigation tree, go to **Fixlets and Tasks** and select the **Run Capacity Scan and Upload Results** fixlet.
3. In the lower pane, click **Take Action**, and choose **Click here to run a single capacity scan and force upload of results**.
4. Open the **Target** tab and select the computers that you want to scan. Then, click **OK**.
5. Wait for the scheduled import, or run it manually.

Computers have the *No Scan Data* status because results of the capacity scan were removed from the BigFix Inventory server

Results of the capacity scan are stored in a folder on the computer where the BigFix Inventory server is installed. If the folder is removed, the status of computers on the BigFix Capacity Completeness widget changes to *No Scan Data*. If the results of the next capacity scan differ from the results that currently exist on the scanned computers, they are uploaded to the BigFix Inventory server. Then, the status of computers changes from *No Scan Data* to *OK*. However, if the results of the next capacity scan are the same as the results that exist on the scanned computer, they are not uploaded and the computers remain in the *No Scan Data* status. To solve the problem, perform the following steps.

1. Log in to the BigFix console.
2. In the navigation tree, go to **Fixlets and Tasks** and select the **Run Capacity Scan and Upload Results** fixlet.
3. In the lower pane, click **Take Action**, and choose **Click here to run a single capacity scan and force upload of results**.
4. Open the **Target** tab and select the computers that you want to scan. Then, click **OK**.
5. Wait for the scheduled import, or run it manually.

Scan problems

The most common problems with software scans are reported on the Scan Health widget. Refer to this topic to learn how to solve problems that are reported on the widget as well as for solutions to other scan-related problems.

Scan Health widget

The Scan Health widget shows the health of scans that are running in your infrastructure. You can drill down to reports for specific computers with scan problems. By sorting the columns of a report, you can quickly understand which computers are failing which specific software scan types.

Failed Scan: At least one type of software scan (catalog-based, file system, package data, or software tags) did not complete successfully.

This problem might occur because the computer is out of space, is misconfigured, or the scan was stopped.

Missing Software Scan: The last attempt to initiate at least one type of software scan (catalog-based, file system, package data, or software tags) was more than 30 days ago.

This problem might occur because the computer is out of space, is misconfigured, or the scan was stopped.

Outdated Catalog: The version of the catalog on the agent is older than the version that is available on the server.

When you upload a software catalog to BigFix Inventory or edit your custom catalog, the status of the catalog is pending until you run an import. During the import, endpoint catalogs are created and the version of the catalog that is displayed on the Software Catalog widget is updated. Endpoint catalogs are then distributed to the computers in your infrastructure.

The time when the endpoint catalog reaches the computers depends, for example, on the computer connection status. When the endpoint catalog is updated on the computer, information about the new version of the catalog is gathered by the Software Scan Status analysis. After the next import, information about the version of the catalog that is available on the computer is updated in BigFix Inventory.

Because this entire flow might take some time, it can happen that a computer has the Outdated Catalog status even though there are no problems. To verify whether any action is needed, perform the following steps:

1. On the Software Catalog widget, check the date when the catalog was last edited.
2. On the Scan Health widget, click **Outdated Catalog** to view the list of affected computers and check the date in the **Latest Scan Import** column.

- If the last scan import was before the date when the catalog was last edited, wait for the next scan and the subsequent import. As described above, distribution of the endpoint catalog to the computers and reporting its version from the computers to BigFix Inventory might take some time.
- If the last scan import was after the date when the catalog was last edited, force the distribution of the software catalog to the computers. For more information, see: [Updating scanner catalogs \(on page dcccxxiv\)](#).

Scan Not Uploaded: Results of the catalog-based scan or file system scan were not uploaded to the BigFix server.

This problem might occur because the computer or a relay is offline, there is a network outage, or the last scan attempt was more than 30 days ago. Check whether each computer listed in the report is running and that the scan results are uploaded to the server on a regular basis. You can check the time of last scan attempt for all types of the software scan in the Software Scan Status analysis.

If the problem persists, it might indicate that the initial download of the catalog to the endpoints failed. To solve the problem, upload the latest catalog to the agent. For more information, see: [Updating scanner catalogs \(on page dcccxxiv\)](#).

IBM i

IBM i computer scans are excluded from the Scan Not Uploaded check.

Unscanned Shared Disks: A remote shared disk is not scanned by the agent.

License usage for BigFix software that is installed on an unscanned shared disk might be under-reported. To solve the problem, enable scans of remote shared disks.

9.2.12

The status is based on scanning shared disks by using the basic mode. Starting from application update 9.2.12, you can discover software that is installed on shared disks by using the optimized mode. This mode allows for reducing the load that scans produce on shared disks. For more information, see: [Discovering software on shared disks \(on page ccxxi\)](#).

Other scan problems

The software scans cannot be initiated because the Initiate Software Scans task is not applicable.

The software scans depend on the scanner catalogs that are used by the scanner to discover software. The scanner catalogs are created when you upload the software catalog to BigFix Inventory. If you cannot run the software scans, it might mean that the catalogs were not created. Before you update the scanner catalogs manually, complete the following steps to check the cause of this problem:

1. On the computer on which the task is not applicable, go to the `<BES_Client>\LMT\CIT` directory.
2. Check if the `catalog.bz2` file exists. The file contains the scanner catalogs.
3. Log in to the BigFix console.
4. In the navigation bar, click **Actions**.

5. In the upper-right pane, locate the **Catalog Download (Version version)** action and select it.
6. Check the details of the action. You can check on which computer the action failed and investigate the failed steps. Try to determine the cause of the problem and fix it.

If you cannot determine the cause, update the catalogs manually. For more information, see: [Updating scanner catalogs \(on page dcccxiv\)](#).

In the Software Scan Status analysis, the status of a software scan is Failed and the value in the Is archive file size exceeded column is True.

The problem occurs when the scan file is larger than the maximum archive file size that can be sent by the BigFix client. To solve the problem, perform the following steps:

1. Check the BigFix client log for information about the size of the file that exceeded the limit. By default, the log is in the following location.

Linux	/var/opt/BESClient/___BESData/___Global/Logs
Windows	C:\Program Files (x86)\BigFix Enterprise\BESClient _BESData___Global \Logs

2. Log in to the BigFix console.
3. In the left navigation, click **Computers**, right-click the computer on which the scan failed, and then click **Edit Computer Settings**.
4. Increase the value of the **_BESClient_ArchiveManager_MaxArchiveSize** parameter.

The software scan cannot be run on an LPAR because the Initiate Software Scan is not relevant on that computer

When a WPAR exists on an LPAR, all WPAR processes are visible on the level of the LPAR. When the software scan is running on the WPAR, the process is visible on the level of the LPAR and thus the Initiate Software Scan fixlet is not relevant on the LPAR. To run the software scan on the LPAR, wait for the scan to finish on the WPAR. Alternatively, perform the following steps:

1. Log in to the BigFix console.
2. Select the **Initiate Software Scan** fixlet and click **Take Action**.
3. On the **Target** tab, select **Dynamically target by property**.
4. Expand **By Retrieved Properties** and then expand **By Computer Name**.
5. Select the LPAR computer on which you want to run the software scan and click **OK**.

Upgrade of the scanner fails on a WPAR.

When a WPAR exists on an LPAR, you must first upgrade the scanner on the LPAR, and then on the WPAR.

Software scans are failing. Computers where the scans are failing have the Low Disk Space status on the Deployment Health widget.

The problem might be caused by insufficient disk space for the scanner cache. To solve the problem, move the scanner cache folder to a different location or optimize the cache. For more information, see: [Optimizing scanner cache configuration \(on page cclxvi\)](#).

Software components that are installed in the `/usr/lpp` directory on AIX are not discovered.

The problem occurs because the `/usr/lpp` directory is by default excluded from software scans. Starting from scanner version 2.8.0.5000, the directory is included in software scans, and components installed in this directory are discovered. Thus, to solve the problem, update the scanner to version 2.8.0.5000 or later, and wait for the next scheduled software scan. Alternatively, you can manually include the `/usr/lpp` directory in software scans. For more information, see: [Including the excluded directories back in scans \(on page ccxiv\)](#).

After the update of the scanner, the number of software components discovered on AIX increased and some duplicates appeared.

Previously, the `/usr/lpp` directory was excluded from software scans. However, components whose signatures existed only in this directory were not discovered. Thus, starting from scanner version 2.8.0.5000, the `/usr/lpp` directory was included in software scans. It caused that more components are discovered and displayed on the reports. However, components installed in the `/usr/lpp` might also be discovered in other directories. In such case, duplicate components appear on the reports. To solve the problem, suppress the duplicates. For more information, see: [Excluding and suppressing software instances \(on page dxlvii\)](#). You can also create a custom rule that suppresses components that are discovered in `/usr/lpp` and other directories. For more information, see: [Creating and managing custom rules \(on page dl\)](#).

New software titles were added to the catalog, a data import was run but the software is not visible in the BigFix Inventory web user interface.

One of the reasons for the missing software titles in the BigFix Inventory UI is that non-standard file types were used as software signatures. If you added a rule with non-standard file extensions, you must either wait until all the steps in the [catalog data flow \(on page cdxcix\)](#) complete or perform those steps yourself. To have accurate inventory data:

1. [Run a data import \(on page cclxii\)](#) to propagate the catalog to the endpoints.
2. Verify that the catalog was propagated successfully.
3. [Upload the scan data \(on page ccviii\)](#) to the BigFix Inventory server.
4. Run another data import.

The scanner cannot be updated on 32-bit Linux x86

The problem occurs on the 32-bit Linux x86 with the `libstdc++.so.5` library for which the last available version of the scanner is 2.8.0.3000. To solve the problem, update the library to `libstdc++.so.6` if available. Otherwise, the scanner cannot be updated.

The result of the file system scan is positive by default in the IBM i environment. The **File System Scan Successful** column is always set to yes.

UNIX After running a scan on a large system with high number of mount points or share drives, the log file contains multiple return codes related to insufficient memory

During file scan or software scan that is performed on large servers with high number of mount points or share drives, memory requirements for the scanner increase. If there is not enough memory available for the scanner to work, the log file might contain the following return codes:

- 125 which indicates memory allocation failure.
- 134 which translates into the operating system signal 6 - SIGABRT.
- 139 which translates into the operating system signal 11 - SIGSEGV.

For more information, see: [Software scan return codes \(on page dcccij\)](#).

To check whether the issue is related to problems with memory, monitor the memory usage of the scanner and determine if it reaches the limits set on the system. Ensure that your limit allows scanner to work properly. For example, for the scanner on a sever with 20 000 mount points or share drive the required memory is about 500 MB of RAM. Thus, the ulimit for `data seg size`, or `ulimit -d` must be set to at least 500000. For more information, see: [Ulimit command \(on AIX\)](#).

When you scan large systems with high number of mount points or share drives, consider the following aspects:

- The scan can be slow.
- The scanner logging must be set to minimum, as it has significant impact on performance in such environments.
- Do not set a CPU threshold, or if required, ensure that the limit is not too low.

Windows Optimized scan of shared disks does not complete on Windows

When you enable optimized scan of shared disk, the scan does not complete. In the BigFix console, when you go to **Actions**, the status of the Optimized Shared Disks Scan Update Resources List action is Pending Downloads and the following error is displayed in the Summary section.

```
HTTP Error 28: Timeout was reached: Connection timed out after 10000 milliseconds"
```

The problem occurs on Windows systems with multiple interfaces. To solve the problem, see: [Configuring servers in separate networks \(on page clxxx\)](#).

Security problems

Security problems in BigFix Inventory might include issues with logging in to the application or those related to the security of your credentials and your environment. However, you can easily recover from these problems.

Login credentials and the authenticity token are stored as plain text in the HTTP packet.

After logging in, the login form that contains the credentials is sent as plain text in the HTTP packet. You can solve this issue by configuring SSL.

When creating a new user, autocomplete is enabled for the password field.

When creating a new user, the password field might be filled by autocomplete based on the password that is stored in the browser.

The server is not working properly after certificates are modified.

If the server is not working properly after certificates are modified and the server is restarted, then delete the keystore file `key_server.jceks` and restart the server. The keystore file is regenerated with a self-signed certificate. You can investigate the problem in the `tema.log` file.

Difficulty establishing a connection with HTTPS.

If you have difficulty when establishing a connection with HTTPS and you are using SSL, check that your browser supports TLS 1.2 and that it is enabled.

The single sign-on configuration values are not updated automatically after you modify the server port.

Modifying the port number on the **Server Settings** pane in BigFix Inventory while single sign-on is enabled invalidates the single sign-on configuration. For information how to properly modify the port, see [Modifying port in BigFix Inventory that has single sign-on enabled \(on page dccxxiii\)](#).

If you already modified the BigFix Inventory server port and are experiencing issues signing onto BigFix Inventory, you need to:

1. [Revert the disabled SSO configuration for SAML \(on page dccxviii\)](#) or [Revert the disabled SSO configuration for LTPA \(on page dccxxi\)](#).
2. Provide the new port value on the **Server Settings** page. To access the page, click **Management > Server Settings**.
3. Re-create the single sign-on configuration with the new port value. For more information, see either [Configuring SSO based on SAML token \(on page dccxii\)](#) or [Configuring SSO based on LTPA \(on page dccxx\)](#).

After you log in to BigFix Inventory for the first time with single sign-on enabled, you are redirected to an BigFix icon instead of the overview page.

To recover from this error, follow the instructions in [Handling the favicon.ico file with Mozilla Firefox](#).

When you are importing a certificate in the PEM format and an encrypted private key in the pkcs8 format, an error about incorrect password for the private key is displayed.

When you are importing a certificate and a private key in the pkcs8 format, the following error is displayed:

```
Error when validating private key password: problem parsing ENCRYPTED
PRIVATE KEY: java.lang.SecurityException: JCE cannot authenticate the provider BC.
```

To solve the problem, add the following line to the `installation_dir/jre/lib/security/java.security` file:

```
security.provider.10=org.bouncycastle.jce.provider.BouncyCastleProvider
```

Then, restart the BigFix Inventory server.

When you log in to BigFix Inventory using the LDAP authentication, the following error message is displayed:

Error contacting the Directory Server for authentication.

The error might occur if the SSL LDAP certificate that is used to authenticate users in BigFix Inventory was recently updated. To refresh the certificate in the BigFix Inventory database, perform the following actions:

1. Log in to BigFix Inventory as a local administrator.
2. In the top navigation bar, click **Management > Directory Servers**.
3. Choose the LDAP server that is used to authenticate users.
4. Click **Test Connection**, and wait for connection test to finish.
5. Click **Save**.

Antivirus software detects the **LMT/CIT** directory as possible threat.

The **LMT/CIT** directory is one of the default scanner directories that is required by BigFix Inventory. It is not infected with any malicious software and does not pose any threat to your system. It is recommended to exclude this directory from antivirus scans.

Secure connection is not initialized and the **CWWKO0801E** error can be found in the `tema.log` file.

Secure connection is not initialized and the following error can be found in the `tema.log` file.

```
000000b7 com.ibm.ws.channel.ssl.internal.SSLHandshakeErrorTracker
CWWKO0801E: Unable to initialize SSL connection. Unauthorized access was denied
or security settings have expired. Exception is javax.net.ssl.SSLHandshakeException:
Client requested protocol TLSv1 not enabled or not supported.
```

To solve the problem, enable TLS 1.2 in IBM Java. For more information, see: [Enabling TLS 1.2 in IBM Java \(on page dcci\)](#).

Server operation problems

Server operation problems in BigFix Inventory might include general issues that you might encounter when you use the application and its functions. However, you can easily recover from these problems.

Scan configuration displays the status: **Server unavailable**

The scan schedule cannot be displayed and set because the BigFix server or the Web Reports service is not responding. To fix the problem:

1. Go to **Management > Data Sources**.
2. Check the connection parameters for the BigFix server.

3. Go to a browser on BigFix Inventory computer and test the following URL:

```
https://bigfix-server-address:52311/api/query?relevance=id%20of%20bes%20computers%20whose%20(root%20server%20flag%20of%20it=true)
```

4. When prompted, log in providing the connection parameters for the BigFix server, that you verified in [step 2 \(on page dcclxxxvii\)](#).

The expected result is:

```
<BESAPI xsi:noNamespaceSchemaLocation="BESAPI.xsd">
  <Query Resource="id of bes computers whose (root server flag of it=true)">
    <Result>
      <Answer type="integer">number value</Answer>
    </Result>
    <Evaluation>
      <Time>0.202ms</Time>
      <Plurality>Singular</Plurality>
    </Evaluation>
  </Query>
</BESAPI>
```

The following message indicates that Web Reports service is started but the required information is not collected yet. Wait a couple of minutes and retry the query.

```
<BESAPI xsi:noNamespaceSchemaLocation="BESAPI.xsd">
  <Query Resource="id of bes computers whose (root server flag of it=true)">
    <Result/>
    <Error>Singular expression refers to nonexistent object.</Error>
  </Query>
</BESAPI>
```

You can receive one of the following error messages:

- HTTP 401: Unauthorized

It indicates that you provided invalid credentials or the authorization step was canceled.

- Cannot perform relevance query evaluation at this time because there is no reachable Web Reports instance collecting data from this Server.

It indicates that Web Reports service is down or there is an issue with network connection, for example: proxy or firewall blocks access to Web Reports. Contact your network administrator for assistance to enable the communication between the BigFix server and Web Reports.

The RPM scanner fixlet fails on AIX® 6.1

The scanner fails and the RPM installation package returns `undefined` during the software scan. To remedy this problem, upgrade your BigFix server, console, and clients to version 8.2.1175.

Version 8.0 and 8.1 BigFix® Enterprise Server Clients cannot subscribe to the BigFix Inventory server.

The solution is to unsubscribe all computers from the current site or remove the site, and then subscribe the clients back.

A report was created correctly but it could not be sent as a PDF attachment to an email.

The PDF report could not be sent because the mail server in your company does not accept large email attachments. To fix the problem:

1. Determine the size of large reports that are generated by BigFix Inventory: generate a sample PDF by clicking the **PDF** icon in the BigFix Inventory web user interface.
2. Contact your mail server administrator and request a higher size limit for email attachments in both the outgoing and incoming mail server configuration.

Data that is gathered by an analysis property is incorrect.

The same analysis property can exist in multiple external sites. In such a situation, the data that is gathered by the analysis might come from any of the sites, not necessarily the intended one. If some unexpected data is gathered by an analysis property, ensure that the analysis that you defined comes from the correct site. Go to the following URLs:

- https://<bfi_host>:<port>/management/sam/unix_package_properties
- https://<bfi_host>:<port>/management/sam/package_properties
- https://<bfi_host>:<port>/management/sam/app_usage_properties

Click the name of the property, and in the **Data Source Property** list, check whether the site is correct. The name of the site is displayed below the name of the property. If the site is incorrect, change it.

BigFix Inventory does not start after rebooting the server.

The problem often occurs in environments where BigFix Inventory and DB2® are installed on the same server. After rebooting the server, the DB2® instance is not started which also blocks BigFix Inventory from starting.

To fix the problem, ensure that your DB2® instance is running and then start BigFix Inventory:

1. Log in as the DB2® instance owner and start the instance:

```
su db2inst1
db2start
```

2. Start the BigFix Inventory server:

```
/etc/init.d/SUAserver start
```

'We're sorry, but something went wrong.'

For more information about the error, check the server log file `tema.log` in the `installation_directory/wlp/usr/servers/server1/logs/` directory.

Software that is deployed under a `tmp` directory is not discovered.

Temporary directories (`tmp`) contain large amounts of irrelevant data. Thus, they are not scanned by default. However, some applications are installed under `tmp` directories. As a result, they are not discovered by software scans. To solve the problem, include all `tmp` directories except for the main temporary directory of the operating system into software scans. It ensures that applications that are installed under `tmp` directories are correctly discovered but the main temporary directory of the operating system is not unnecessarily scanned.

On the computer where the software is installed, go to the `BESClient/LMT/CIT/` directory and open the `exclude_path.txt` file. It lists all directories that are excluded from software scans.

- For UNIX operating systems
 1. Remove the `*/tmp/*` directory from the list.
 2. Add the following directories and save the file.

```
/tmp/*  
/var/tmp/*  
*/usr/tmp/*
```

- For Windows
 1. Remove the `*/tmp/*` directory from the list and save the file.

Return code 9 or 29 is displayed in the results of the Software Scan Status analysis.

The return code indicates that the scanner timed out. To solve the problem, perform the following tasks:

- Increase the scan timeout in the **Configure Scan Timeout** task.
- Increase the amount of CPU that can be consumed by the scanner by increasing CPU threshold in the **Initiate Software Scan** task.
- Exclude directories with backups from software scans.
- Decrease the scanner trace level in the **Edit Scanner Trace Settings** task.

If all other options fail, reinstall the scanner.

An error occurs while importing contracts through REST API.

When using a `curl` command to import contracts through REST API, the following error is displayed:

```
curl: (18) transfer closed with outstanding read data remaining
```

To solve the problem, upgrade `curl` to a later version:

Linux Curl must use the NSS library 3.14, or later. Run `curl --version`. The version of NSS is displayed in the output. Upgrade curl if the NSS version is earlier than 3.14.

```
curl 7.19.7 (x86_64-redhat-linux-gnu) libcurl/7.19.7 NSS/3.14.0.0
zlib/1.2.3 libidn/1.18 libssh2/1.4.2
```

Windows Upgrade to curl 7.27.0, or later.

When you make an API call, the call fails and the following message is written in the log: `The connection is closed.`

The problem occurs when you make an API during the import. To solve the problem, wait until the import finishes and ensure that the BigFix Inventory server is running. Then, repeat the API call.

The `tema.log` file contains the following error: `E CWWKF0002E: A bundle could not be found for com.ibm.ws.javaee.servlet.3.0/[1.0.0,1.0.100).`

The error occurs after you enable single sign-on. There are no other problems except for the errors in the log. To solve the problem, perform the following steps:

1. Go to the BigFix Inventory installation directory and open the `wlpstart.bat` file.
 - Linux: `/opt/ibm/BFI/wlp/bin/wlpstart.bat`
 - Windows: `C:\Program Files\ibm\BFI\wlp\bin\wlpstart.bat`
2. Add the `--clean` parameter to the last line in the file.

```
call "%WLP_PATH%\bin\server.bat" start server1 --clean
```

3. Stop the server (on page [cd](#)).
4. Start the server (on page [cccxcix](#)).
5. Verify that the errors no longer appear in the log and remove the `--clean` parameter from the `wlpstart.bat` file.

Emails with scheduled reports are not sent via email after the BigFix Inventory server was upgraded.

To fix this problem, restart the BigFix Inventory server.

User interface problems

User interface problems in BigFix Inventory are mostly related to information that is incorrectly displayed in the application or to errors that arise in the non-English-language versions of the software. However, you can easily recover from these problems.

After restarting the computer, the BigFix Inventory user interface cannot be launched.

When you restart the computer on which you have BigFix Inventory installed, you must start the DB2® database before starting the application. Otherwise, you will not be able to log in to BigFix Inventory.

If you restart your computer, complete the following steps before starting the application:

1. Switch to the DB2® instance owner:

```
su - db2inst1
```

2. Start the database:

```
db2start
```

3. Restart BigFix Inventory:

```
/etc/init.d/BFIserver restart
```

On a low-resolution monitor, the user interface is not displayed correctly.

The minimal supported screen resolution is 1360x768 pixels. Recommended screen resolution is 1920x1080 pixels.

On Linux, text columns in reports are not alphabetically sorted.

Text columns in BigFix Inventory are sorted differently on Linux, and Windows. If you click a column header, the values are sorted in a case-sensitive way on Linux (all the names with capital letters are displayed first) and case-insensitive on Windows (capitalized and non-capitalized words are mixed and sorted in alphabetical order).

On the Catalog Audit report, the values are displayed in an incorrect time zone.

Values that are related to dates and time are based on the time zone that is set in your operating system. However, the Daylight Saving Time is not considered. Thus, values are set to an incorrect time zone.

When you leave a page while modifying data, there is no warning that the data will be lost.

When you modify data, for example, in one of the panels from the **Catalog Customization** menu, the data will be lost if you change the view or close the page. No warning is displayed. Ensure that you save all data before navigating to another view.

Text in reports cannot be selected when using Internet Explorer.

When trying to copy the content of the tables in reports, the text cannot be selected and copied when using the Internet Explorer browser. To solve this issue, use a different browser.

The user interface of BigFix Inventory is not displayed correctly in Internet Explorer

The problem occurs on Internet Explorer version 8. To solve the problem, use Internet Explorer version 9 or higher.

The `Import is now running` message is not refreshed even after the import is complete.

The page with the `Import is now running` message is not refreshed and the message is still displayed even when the import is complete. To solve this issue, you can hover over the **Reports** menu and check the percentage status of the import on the bottom of the drop-down list. You can also refresh the page manually.

The PDF reports open in the same browser tab as BigFix Inventory.

The generated PDF reports open in the same tab as BigFix Inventory. Because of this issue, you might be forced to stop any activity in the application until the PDF opens. To solve the issue, change the settings in your browser.

The PDF reports that are downloaded by using Firefox are corrupted.

Although the PDF reports are correctly displayed in the browser, the downloaded file is corrupted. The problem occurs in Firefox. To solve the issue, use a different browser.

The Software Installations report cannot be exported to PDF when all columns are selected. The problem occurs in Internet Explorer.

The problem occurs because the maximum length of the URL in Internet Explorer is limited to 2083 characters. The URL of the Software Installations report with all columns selected exceeds the maximum length and the PDF cannot be generated. To solve the problem, use a supported version of Firefox or Chrome.

Autocomplete suggests irrelevant results.

In some cases, autocomplete might suggest irrelevant results. For example, when you type HCL® as a publisher and then start typing the software product, autocomplete will list all possible software products instead of only those that belong to HCL®.

When you open links to product details in new tabs, an error is displayed

When you open the All Metrics or the IBM PVU Subcapacity report and open links to products that are displayed on these reports in new tabs, an error is displayed. The problem occurs because information about the time range for which the metric quantity should be displayed on the report is lost while opening the links in new tabs. To solve the problem, refresh the page.

Paths for Korean and Japanese locales contain incorrect delimiter.

Paths for Korean and Japanese locales use a backslash (\) as a delimiter, instead of won (₩) or yen.

There is a missing space before the colon (:) for French locale.

When a colon is used in text, there is a missing space before it although it is required in French language.

The string on the Browse button on the Catalog Upload panel is not displayed in languages other than English.

This error is caused by the web browser internal configuration and does not depend on BigFix Inventory.

Some Chinese characters are not displayed correctly.

To solve the problem, install the Zhong Yi Song TrueType font for the GB18030 standard:

1. Obtain the `fonts-chinese-zysong` package.
2. Install the package by running the following command:

```
rpm -ivh fonts-chinese-zysong-0.1-5.e15.noarch.rpm --nodeps
```

The new font was installed in the `/usr/share/fonts/zh_CN/TrueType` directory. The characters should now be displayed correctly.

Messages are displayed in multiple languages.

The language in which some of some messages are displayed depends on the settings of the BigFix Inventory server. To change the language settings on the server:

1. Open the `jvm.options` file that is in the `installation_directory/wlp/usr/servers/server1/` directory.
2. Add the following lines to the file and specify the language to be used.

```
-Duser.country=country  
-Duser.language=language
```

For example:

```
-Duser.country=US  
-Duser.language=en
```

3. Restart the server.

Names of roles and computer properties are not translated.

Starting from version 9.2.3, the language in which names of roles and computer properties are displayed depends on the language that was chosen during the installation of the BigFix Inventory server. For earlier versions, the language depends on the language that was set in the web browser during the creation of database connections after the installation.

If you need to change the language in which these items are displayed, you must do it manually. To change the name of a role, go to **Management > Roles**, select the role that is displayed in a wrong language, and change its name. To change the name of a computer property, go to **Management > Computer Properties**, select the property that is displayed in a wrong language, and change its name. In both cases, wait for a scheduled import or run in manually for the changes to take effect.

The changes that you make apply to all users regardless of their language settings.

After exporting the catalog in Internet Explorer, the Catalog Upload page remains busy.

When you go to **Management > Catalog Upload** and click the **Export** button, the catalog is successfully exported but the cursor remains busy and the button itself inactive.

To solve this problem and make the button active again, refresh the page or go to a different section of BigFix Inventory and then return to **Catalog Upload**.

'We're sorry, but something went wrong.'

For more information about the error, check the server log file `tema.log` in the `installation_directory/wlp/usr/servers/server1/logs/` directory.

Migrated to SUA9 report column is blank.

When you migrate Software Use Analysis 2.2 endpoints to 9.2.16, the Migrated to SUA9 column in your software inventory and computer reports is blank. Check whether the **SUA 2.2 Endpoint Migration** analysis is activated.

1. In the BigFix console navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9** and click **Analyses**.
2. Select the **SUA 2.2 Endpoint Migration** analysis, right-click on it, and then select **Activate**.

Unable to log in with users from Microsoft™ Active Directory

If you are having problems with your Active Directory users, append the user name with a domain, for example `username@domain.com`. Adding a domain is required starting from Software Use Analysis 9.2.0.2, but it also concerns Active Directory servers that were configured in previous versions. See the following cases for more information about how this change might impact your existing users.

Active Directory with Global Catalog

LDAP Server* Global Catalog

- If you added users through User Provisioning in BigFix Inventory, new users are created whenever you log in with a user name and a domain. You can delete all Active Directory users that have only the user name specified because you will not be able to log in with those users.
- If you added users manually in BigFix Inventory, you must create them again, this time also specifying the domain. You can delete all Active Directory users that have only the user name specified because you will not be able to log in with those users.

Active Directory without Global Catalog

LDAP Server* Global Catalog
 Search Base*

If you did not use the Global Catalog, you must have specified a domain in the Search Base field. This domain is automatically added to all users in BigFix Inventory and no action is required from you.

URLs that point to the product web user interface do not work.

If you enabled single sign-on in BigFix Inventory, the link that is available in the **Start** menu on Windows™ might not work because an extra `https://` prefix is added to the url.

You can access this link on Windows™ 2008 by clicking **Start > All Programs > BigFix Inventory > Web Interface**. This URL is defined in the **WebUI** file that is in the following location: `C:\Program Files\ibm\BFI\admin\resources\WebUI`.

To fix this issue:

1. Open the **WebUI** file in a text editor, for example Notepad++, and remove the extra "https://" from the beginning of the URL.



Note: The default shortcut path is `https://localhost:9081`. If LTPA SSO is enabled, it needs to be updated to `https://virtual-junction-host-name:9081`.

2. Clear the cache in your web browser.
3. End all your browser processes.
4. Restart the system on which your browser is installed.

When opening an exported CSV report in Microsoft Excel, some data is corrupted.

Due to formatting issues in some versions of Microsoft Excel, the data might be displayed in incorrect columns. This problem mostly concerns components that have a comma in their name, which results in wrong formatting.

To fix this problem, complete the following steps.

1. Open Microsoft Excel.
2. In the navigation bar, click **Open**, and navigate to the folder that contains your report.
3. Change the valid file extensions to **All Files (*.*)** to see the report.
4. Open the file.
5. In the Text Import Wizard, select **Delimited**, and then choose **Comma** as a delimiter, and a quotation mark (") as a text qualifier.

After clicking **Finish**, the data will be divided among appropriate columns and displayed correctly.

Loading any report takes substantial amount of time in case of huge deployments.

In case of huge deployments, it can take a lot of time to generate a full report. Thus, if you do not need to view a full report as a default, you can use the work around and filter the report to limit the list of items.

1. Open a report and wait until it loads.
2. Hover over the **Manage Report View** icon , and click **Configure View**.
3. Set up adequate filters so that the report view is simple and limited to a few entries and click **Submit**.

4. Hover over the **Manage Report View** icon , and click **Save As**.
5. Name the report and save it as default.

This simple report that you created is now set as the default view. You no longer need to wait for the full report to load. You can further filter the data to view only the information that you require.

Fields in the user interface are disarranged when using Internet Explorer.

During the initial configuration of BigFix Inventory, some fields in the user interface are doubled and disarranged. To solve this problem, add the BigFix Inventory server to the list of trusted sites in Internet Explorer.

1. Click the Settings icon in the top-right corner of your browser.
2. Click **Internet Options**, and switch to the **Security** tab.
3. Select **Trusted sites**, and then click **Sites**.
4. Enter the web address of your server, and click **Add**.
5. Refresh the page.

Logs and return codes

Find logs and interpret return codes to troubleshoot issues.

Database creation logs and return codes

The log file is automatically generated when you create a database during the initial configuration of BigFix Inventory. The log file contains a return code that can help you check why database creation or validation failed. The file is in the `/tmp` directory and has a time stamp suffix, for example: `/tmp/createdb_20131018-131841`.

Linux

The log file is created only if you installed BigFix Inventory on Linux™ and use a DB2® server.

Return codes

Check the return code to find the reasons for database creation problems.

Table 72. Return codes for the create database script

The table consists of two columns and 9 rows.

Return code	Description and possible solutions
0	Database successfully created.
1	Help message is displayed, or incorrect syntax.
100	One or more errors occurred during the database creation. Database creation error. Check the database creation log: creation log <code>/tmp/createdb_latest_timestamp</code> .
101	Database already created.

Table 72. Return codes for the create database script

The table consists of two columns and 9 rows.

(continued)

Return code	Description and possible solutions
102	One or more errors occurred during the database validation. There is something wrong with the DB2® commands run from the script. The script requires a correctly initialized DB2® environment.
103	The DB2® version could not be correctly read from the db2level command. The db2level command returned the DB2® version string in an unexpected format.
104	Unsupported database version. ®You can check the version by running the db2level and db2licm -l commands as the SYSADM user.
105	An unsupported edition of the database was found in the system. You can check the version by running the db2level and db2licm -l commands as the SYSADM user.
127	Cannot run the db2 command. Incorrectly configured DB2® instance. The user that runs the script must have SYSADM authority. Try to run the script as the DB2® instance owner.

Server installation and upgrade logs

Installation and upgrade log files are in the same directories, because those processes are completed by using the same installer. If you encounter any problems while installing or upgrading the server, refer to those log files for more information about any occurring errors.

Log files

There are several log files that are created either during the installation or upgrade, or only when the process failed or completed successfully. Each path starts with a variable that is specific to the operating system. Those variables are explained below the table.

Table 73. Installation and upgrade log files

Installation Status	Windows	Linux
In progress	%USERPROFILE%\BFI_9.2.16	\$HOME/BFI_9.2.16
	%USERPROFILE%\ia.log	\$HOME/ia.log
	BESclient_installation_directory\LMT\BFI_upgrade.log	BESclient_installation_directory\LMT\BFI_upgrade.log

Table 73. Installation and upgrade log files

(continued)

Installation Status	Windows	Linux
Failed	<code>%USERPROFILE%\BFI_9.2.16_timestamp_uniqueID_logs.zip</code>	<code>\$HOME/BFI_9.2.16_timestamp_uniqueID_logs.tar.gz</code>
Successful	<code>installation_directory\BFI_9.2.16_timestamp_uniqueID_logs.zip</code>	<code>installation_directory/BFI_9.2.16_timestamp_uniqueID_logs.tar.gz</code>

%USERPROFILE%

Home directory of the user who installed the server, for example `C:\Users\Administrator\`.



Important: When you upgrade by using a Fixlet, the upgrade runs under the `NT AUTHORITY\SYSTEM` user whose default home directory is `C:\Windows\System32\config\systemprofile\`.

\$HOME

Home directory of the user who installed the BigFix client, for example `/root/`.

Log file contents

To understand the root cause of installation and upgrade problems, review logs for messages with a severity level of WARN or ERROR. The following example shows a warning message.

```
<Message Id="CODIN0215W" Severity="WARN">
  <Time Millis="1381929782997"> 2013-10-16 15:23:02.997+02:00</Time>
  <Server Format="IP">NC040226.kraklab.pl.ibm.com</Server>
  <ProductI>dCOD</ProductI>d
  <Component>Install</Component>
  <ProductInstance></ProductInstance>
  <LogText><![CDATA[CODIN0215W The following ports are in use: 9081, .
  The installation process can be continued but the server will not be started.
  The server has to be started manually after resolving the ports conflict.]]></LogText>
  <Source FileName="com.ibm.license.mgmt.install.ia.common.CommunicationCommon"
  Method="okToContinue"/>
  <TranslationInfo Type="JAVA" Catalog="userLocales.InstallMessageEWI" MsgKey="checkPorts">
  <Param><![CDATA[9081, ]]></Param></TranslationInfo>
  <Principal></Principal>
</Message>
```

To learn more about errors and return codes that can appear in the logs, go to [Server installation and uninstallation return codes \(on page dccc\)](#).

Server installation and uninstallation return codes

If the server installation or uninstallation fails, check the return code to learn about the reason of the problem and possible solutions.

Table 74. Server installation and uninstallation return codes

The table consists of two columns and 32 rows.

Return code	Possible cause and solutions
0	The server was installed successfully.
5	An unexpected error occurred.
6	An unexpected exception occurred.
7	An internal error occurred. The installer failed to save the file with information that was collected or generated during the preinstallation stage.
8	The installation was canceled.
9	A post-installation step was terminated before it was finished. Problems with resuming the installation might occur.
11	Validation of the communication ports failed. Either the same port is specified for more than one parameter or the specified port is in use. If you want to specify a port that is temporarily used but will be available later, set the RSP_DISABLE_COMMUNICATION_WARNINGS parameter to true. For more information, see: Server installation response file (on page cxlv) .
13	Validation of the license agreement, or the file path failed. Either the license agreement was not accepted, the path to the installation response file is not absolute, or the command is too long. To accept the license agreement, set the RSP_LICENSE_ACCEPTED parameter to true. If the problem persists, ensure that you provided an absolute path to the response file, or move the files to a different directory to shorten the command, for example <code>/root/BFI</code> or <code>C:\BFI</code> .
14	There is not enough space for the installation. To check how much free disk space is required to proceed with the installation, see the following installation log: <code>installation_dir/BFI_9.2_-timestamp_logs.tar.gz</code> .
18	Validation of the installation path failed. The specified path is incorrect or the installation directory is read-only.
20	An unknown response file parameter was specified. Remove the parameter from the installation response file.

Table 74. Server installation and uninstallation return codes

The table consists of two columns and 32 rows.

(continued)

Return code	Possible cause and solutions
21	The response file was not found. It is either empty or contains Windows™ line endings instead of UNIX™ ones.
23	The command-line interface or another application from the BigFix Inventory installation path is still running. End the process manually or set the RSP_AUTO_CLOSE_PROCESSES parameter to true.
26	An internal error occurred. Creation of the log directory failed.
27	It was impossible to recognize the environment, for example, installed products.
28	The upgrading scenario is not supported.
29	A part of BigFix Inventory that is already installed, is corrupted.
30	The uninstallation wizard could not find product information in registries. BigFix Inventory was already uninstalled.
31	The host name was not obtained. To verify the host name, in the command-line interface, enter the following command: <pre>nslookup host_name</pre>
32	An exception was detected while reading the <code>setup.ini</code> file.
33	An attempt of creating a log directory in the installation path failed because a file called <code>BFI9.2</code> already exists. To proceed with the installation, delete the file.
34	The log directory is read-only.
35	The system <code>TEMP</code> environment variable does not point to a valid directory.
36	Installation in console mode is not supported. Use interactive or silent mode.
37	The required resources could not be extracted from the installation image.
38	The required resources could not be found inside the installation image.
41	The post-installation failed.
42	Another instance of the installer is already running.
46	The post-installation was interrupted.
50	Resuming a failed installation in silent mode is not supported.
55	All elements of the infrastructure are already installed.
59	An internal error occurred. Contact IBM® support.

Table 74. Server installation and uninstallation return codes

The table consists of two columns and 32 rows.

(continued)

Return code	Possible cause and solutions
208	The installation could not finish because of the <code>NOEXEC</code> on the <code>tmp</code> directory. If all installation steps are completed, no further action is required. Otherwise, remove <code>NOEXEC</code> from the <code>tmp</code> directory or specify a new <code>tmp</code> directory in the following environment variable: <code>IATEMPDIR</code> . For example, <code>mkdir /root/Install_tmp; export IATEMPDIR=/root/Install_tmp</code> . Next, repeat the installation.
214	The uninstallation process could not connect to the X server. Verify that the <code>DISPLAY</code> variable is properly set and points to a working X server.

TASK

[Uninstalling the server on Windows in interactive mode \(on page cli\)](#)

[Uninstalling the server on Windows in silent mode \(on page clii\)](#)

TASK

[Uninstalling the server on Linux in silent mode \(on page clxxvii\)](#)

[Uninstalling the server on Linux in interactive mode \(on page clxxvii\)](#)

Software scan return codes

If the software scan fails, an error code that indicates why the scan failed is returned. Check what is the possible cause of scan failure that is indicated by each code and how to solve the problem.

Locating the return codes

The error codes for the software scan are returned in the results of the Software Scan Status analysis.

Software scan return codes

Table 75. Software scan return codes

Return code	Possible cause and solutions
0	No errors.
3	A signature file could not be parsed. Update the scanner catalog (on page dcccxiv) on the endpoints where the problem occurred.

Table 75. Software scan return codes

(continued)

Return code	Possible cause and solutions
4	<p>Scan output file could not be saved to the output directory. Ensure that the scan output directory is writable. By default, the directory is:</p> <pre data-bbox="391 541 1349 625"> Linux /var/opt/BESClient/LMT/CIT Windows C:\Program Files (x86)\BigFix Enterprise\BESClient\LMT\CIT </pre>
6	<p>Scanner catalog does not exist on the endpoint. Update the scanner catalog (on page dcccxiv) on the endpoints where the problem occurred.</p>
8	<p>Scanner internal error occurred. Gather scanner logs and contact BigFix Support. By default, the log files are in the following directory:</p> <pre data-bbox="391 884 1019 961"> Linux Scanner_installation_directory/logs Windows Scanner_installation_directory\logs </pre>
9	<p>The scanner timed out. Perform the following tasks:</p> <ul style="list-style-type: none"> • Increase the scan timeout in the Configure Scanner Query Timeout task. • Increase the amount of CPU that can be consumed by the scanner by increasing CPU threshold in the Initiate Software Scan task. • Exclude directories with backups from software scans. For more information, see: Excluding directories from being scanned (on page ccxli). • Decrease the scanner trace level in the Edit Scanner Trace Settings task. <p>If all other options fail, reinstall the scanner.</p>
10	<p>The scanner is being upgraded. Wait until the upgrade finishes.</p>
11	<p>Scan output file or directory is read only. Ensure that the scan output file or directory is writable for the user that is running the BigFix client. By default, the directory is:</p> <pre data-bbox="391 1585 1349 1663"> Linux /var/opt/BESClient/LMT/CIT Windows C:\Program Files (x86)\BigFix Enterprise\BESClient\LMT\CIT </pre>
12	<p>The <code>cit.ini</code> file was not found. Install the scanner. By default, the <code>cit.ini</code> file is located in the following directory:</p>

Table 75. Software scan return codes

(continued)

Return code	Possible cause and solutions
	<pre>Linux /etc/cit Windows C:\Windows\cit</pre>
13	The scanner configuration file was not found. Reinstall or update the scanner.
14	The scanner trace file was not found. Reinstall or update the scanner.
18	<p>One of the scanner files could not be open. Ensure that the administrator or root of the computer on which the problem occurred has the read and write permissions to the following files:</p> <ul style="list-style-type: none"> Scanner configuration file <code>CitHWConfig.xml</code>, scanner properties file <code>Cit.properties</code>, and scanner trace file <code>CitTrace.properties</code>, which by default are in the following directory: <pre>Linux /opt/tivoli/cit/config/ Windows C:\Program Files\tivoli\cit\config</pre>
19	<p>One of the scanner files could not be renamed. Ensure that the administrator or root of the computer on which the problem occurred has the read and write permissions to the following files:</p> <ul style="list-style-type: none"> Scanner configuration file <code>CitHWConfig.xml</code>, scanner properties file <code>Cit.properties</code>, and scanner trace file <code>CitTrace.properties</code>, which by default are in the following directory: <pre>Linux /opt/tivoli/cit/config/ Windows C:\Program Files\tivoli\cit\config</pre>
20	<p>One of the scanner files could not be deleted. Ensure that the administrator or root of the computer on which the problem occurred has the read and write permissions to the following files:</p> <ul style="list-style-type: none"> Scanner configuration file <code>CitHWConfig.xml</code>, scanner properties file <code>Cit.properties</code>, and scanner trace file <code>CitTrace.properties</code>, which by default are in the following directory: <pre>Linux /opt/tivoli/cit/config/ Windows C:\Program Files\tivoli\cit\config</pre>
21	The scanner configuration file is corrupted. Reinstall or upgrade the scanner.
22	The scanner trace file is corrupted. Reinstall or upgrade the scanner.
28	A required shared library file is not available. Reinstall or upgrade the scanner.
29	<p>The scanner timed out. Perform the following tasks:</p> <ul style="list-style-type: none"> Increase the scan timeout in the Configure Scanner Query Timeout task. Increase the amount of CPU that can be consumed by the scanner by increasing CPU threshold in the Initiate Software Scan task.

Table 75. Software scan return codes

(continued)

Return code	Possible cause and solutions
	<ul style="list-style-type: none"> Exclude directories with backups from software scans. For more information, see: Excluding directories from being scanned (on page ccxli). Decrease the scanner trace level in the Edit Scanner Trace Settings task.
	If all other options fail, reinstall the scanner.
30	The scanner query failed. Gather scanner logs and contact BigFix Support.
31	The scan process was interrupted. Stop the scan action that is running on the endpoints where the problem occurred and start the scans again.
37	Scanner internal error occurred. Gather scanner logs and contact BigFix Support.
40	An error occurred while creating the warning file during the software scan. Ensure that the administrator or root of the computer on which the problem occurred has the read and write permissions to the following directory: <pre data-bbox="386 987 1352 1060"> Linux /var/opt/BESClient/LMT/CIT Windows C:\Program Files (x86)\BigFix Enterprise\BESClient\LMT\CIT </pre>
41	Scanner process initialization failed. Gather scanner logs and contact BigFix Support.
42	The signature catalog schema cannot be found. Reinstall or update the scanner. If it does not help, gather scanner logs and contact BigFix Support.
49	The user that runs the scan does not have the read permission to the scanner properties file <code>cit.properties</code> . Ensure that the administrator or root of the computer on which the problem occurred has the read permissions to this file. By default, the file is in the following directory: <pre data-bbox="386 1428 1006 1501"> Linux /opt/tivoli/cit/config/ Windows C:\Program Files\tivoli\cit\config </pre>
51	The scanner configuration file was not found. Reinstall or upgrade the scanner.
52	Scanner internal error occurred. Gather scanner logs and contact BigFix Support.
54	The scan file cannot be compressed. Ensure that the administrator or root of the computer on which the problem occurred has the read and write permissions to the following directory: <pre data-bbox="386 1774 1352 1848"> Linux /var/opt/BESClient/LMT/CIT Windows C:\Program Files (x86)\BigFix Enterprise\BESClient\LMT\CIT </pre>

Table 75. Software scan return codes

(continued)

Return code	Possible cause and solutions
56	The scan output signature file cannot be created. Ensure that the administrator or root of the computer on which the problem occurred has the read and write permissions to the following directory:

Linux /var/opt/BESClient/LMT/CIT

Windows C:\Program Files (x86)\BigFix Enterprise\BESClient\LMT\CIT

125 Memory allocation failed. Gather scanner logs and contact BigFix Support.

Return codes greater than 128

UNIX Some return codes are triggered by the problems that are signaled by the operating system. These are the return codes greater than 128. To check what is the system signal, use the following calculation:

```
return code - 128 = operating system signal
```

To identify the issue, check the meaning of the signal in the operating system documentation.

The following table presents the list of the most common return codes, their possible cause and solution.

Table 76. Scanner return codes greater than 128

Return code	Signal	Possible cause and solution
134	6 - SIGABRT	There might not be enough memory available for the process to work. Increase the memory limit for this process by changing ulimit for <code>data</code> <code>seg size: ulimit -d</code> . If the problem persists, contact BigFix Support.
138	10 - SIGBUS	There is not enough disk space on the computer. Free some disk space. If the problem persists, ensure that you use the latest version of the scanner.
139	11 - SIGSEGV	There might not be enough memory available for the process to work. Increase the memory limit for this process by changing ulimit for <code>data</code>

Table 76. Scanner return codes greater than 128 (continued)

Return code	Signal	Possible cause and solution
		<code>seg size: ulimit -d</code> . If the problem persists, contact BigFix Support.

Capacity scan return codes

If the capacity scan fails, an error code that indicates why the scan failed is returned. Check what is the possible cause of scan failure that is indicated by each code and how to solve the problem.

Locating the return codes

If the capacity scan fails, first check whether the return code is related to the scanner, or to running the `checkCapacityScanOutput.vbs` script. If the log refers to the script, the return code is specific for the script. It indicates that the capacity scan results are incomplete. The problem is usually solved after the next scan or after the upgrade of the scanner to the latest version.

Check for the error codes:

- The scanner error codes for capacity scan are written in the BigFix client logs. The logs are stored in the following locations:

Linux `/var/opt/BESClient/___BESData/___Global/Logs`

Windows `C:\Program Files (x86)\BigFix Enterprise\BESClient___BESData
___Global\Logs`

- To check the error codes that are related to running the `checkCapacityScanOutput.vbs` script, check the Run Capacity Scan and Upload Results action in the BigFix console. On the Computers tab, identify the computer on which the action failed by checking its status. Double-click this computer. The error code is written in execution details.

Capacity scan return codes

Table 77. Capacity scan return codes

Return code	Possible cause and solutions
0	No errors.
1	Wrong arguments were used during the capacity scan. Ensure that the capacity scan is triggered with an official fixlet.
2	The scan configuration file is corrupted. Update the fixlet site (on page cccvii) .
5	The user that runs the scan does not have read permission to the scanner configuration file. Ensure that the BigFix client is running as SYSTEM user on Windows, or root for other operating systems.

Table 77. Capacity scan return codes (continued)

Return code	Possible cause and solutions
8	Scanner internal error occurred. Gather scanner logs and contact HCL Support. By default, the log files are in the following directory: <div style="margin-left: 20px;"> <pre data-bbox="391 472 959 499">Linux /usr/ibm/tivoli/common/CIT/logs</pre> <pre data-bbox="391 516 1133 543">Windows C:\Program Files\ibm\tivoli\common\CIT\logs</pre> </div> <p data-bbox="310 590 1409 655">If the logs are not in the default directory, check the value of the <code>common_trace_path</code> parameter in the following file for information about the location of the logs:</p> <div style="margin-left: 20px;"> <pre data-bbox="391 703 1045 730">Linux /opt/tivoli/cit/config/Cit.properties</pre> <pre data-bbox="391 747 1219 774">Windows C:\Program Files\tivoli\cit\config\Cit.properties</pre> </div>
11	Scan output file or directory is read only. Ensure that the scan output file or directory is writable for the user that is running the BigFix client. By default, the directory is: <div style="margin-left: 20px;"> <pre data-bbox="391 949 889 976">Linux /var/opt/BESClient/LMT/CIT</pre> <pre data-bbox="391 993 1349 1020">Windows C:\Program Files (x86)\BigFix Enterprise\BESClient\LMT\CIT</pre> </div>
13	The scanner configuration file was not found. Reinstall or update the scanner.
18	One of the scanner files could not be open. Ensure that the administrator or root of the computer on which the problem occurred has the read and write permissions to the following files: <ul style="list-style-type: none"> <li data-bbox="375 1247 1386 1312">• Scanner configuration file <code>CitHWConfig.xml</code>, scanner properties file <code>Cit.properties</code>, and scanner trace file <code>CitTrace.properties</code>, which by default are in the following directory: <div style="margin-left: 20px;"> <pre data-bbox="467 1327 924 1354">Linux /opt/tivoli/cit/config/</pre> <pre data-bbox="467 1371 1084 1398">Windows C:\Program Files\tivoli\cit\config</pre> </div>
21	The scanner configuration file is corrupted. Reinstall or upgrade the scanner.
49	The user that runs the scan does not have the read permission to the scanner properties file <code>Cit.properties</code> . Ensure that the administrator or root of the computer on which the problem occurred has the read permissions to this file. By default, the file is in the following directory: <div style="margin-left: 20px;"> <pre data-bbox="391 1640 846 1667">Linux /opt/tivoli/cit/config/</pre> <pre data-bbox="391 1684 1003 1711">Windows C:\Program Files\tivoli\cit\config</pre> </div>
54	The scan file cannot be compressed. Ensure that the administrator or root of the computer on which the problem occurred has the read and write permissions to the following directory:

Table 77. Capacity scan return codes (continued)

Return code	Possible cause and solutions
	<p>Linux /var/opt/BESClient/LMT/CIT</p> <p>Windows C:\Program Files (x86)\BigFix Enterprise\BESClient\LMT\CIT</p>

Return codes greater than 128

UNIX Some return codes are triggered by the problems that are signaled by the operating system. These are the return codes greater than 128. To check what is the system signal, use the following calculation:

```
return code - 128 = operating system signal
```

To identify the issue, check the meaning of the signal in the operating system documentation.

The following table presents the list of the most common return codes, their possible cause and solution.

Table 78. Scanner return codes greater than 128

Return code	Signal	Possible cause and solution
134	6 - SIGABRT	There might not be enough memory available for the process to work. Increase the memory limit for this process by changing ulimit for <code>data seg size: ulimit -d</code> . If the problem persists, contact BigFix Support.
138	10 - SIGBUS	There is not enough disk space on the computer. Free some disk space. If the problem persists, ensure that you use the latest version of the scanner.
139	11 - SIGSEGV	There might not be enough memory available for the process to work. Increase the memory limit for this process by changing ulimit for <code>data seg size: ulimit -d</code> . If the problem persists, contact BigFix Support.

Return codes related to the `checkCapacityScanOutput.vbs` script

Table 79. Return codes linked to the `checkCapacityScanOutput.vbs` script

p

Return code	Possible cause	Systems other than Windows	Solution	
	Windows systems			
1	A general VBS script error. The Windows Script Host could be disabled.	n/a	Enable Windows Script Host.	
2	<code>ComponentID</code> is missing the <code>SerialNumber</code> attribute in capacity scan output file.	Update the scanner to the latest version. For more information, see: <code>#unique_198</code> .	Capacity scan output file is missing the <code>PhysicalProcessor</code> section.	
3	<code>SerialNumber</code> attribute in the capacity scan output file does not contain any value.		Capacity scan output file is missing the <code>OperatingSystem</code> section.	
4	Capacity scan output file is missing the <code>PhysicalProcessor</code> section.		Capacity scan output file is missing the <code>Lpar</code> section.	
5	Capacity scan output file is missing the <code>OperatingSystem</code> section.		n/a	
6	Capacity scan output file is missing the <code>Lpar</code> section.			
7	Capacity scan output file is missing the <code>ComponentID</code> section.		Capacity scan output file is missing the <code>ComponentID</code> section.	n/a

Return codes of capacity scans on virtualization hosts

If the capacity scan that is run by using the Run Capacity Scan on Virtualization Hosts task fails, an error code that indicates why the scan failed is returned. The code is presented in the results of the Status of Capacity Scan on Virtualization Hosts analysis. Check what is the possible cause of scan failure that is indicated by each code and how to solve the problem.

Table 80. Return codes of capacity scans on virtualization hosts

Return code	Possible cause and solution
0	No errors.
20	Unsupported virtualization type. To solve the problem, perform the following steps:

Table 80. Return codes of capacity scans on virtualization hosts (continued)

Return code	Possible cause and solution
	<ul style="list-style-type: none"> • Verify that the computer on which you ran the task is a virtualization host. If it is not, you do not need to run the task on this computer. • Verify that the computer runs on a supported virtualization (on page c). If the virtualization is a supported version of PowerKVM, KVM x86, Citrix XenServer, or Xen Hypervisor, contact BigFix Support. If it is a different type of a supported virtualization, configure VM managers to collect capacity data from these hosts. For more information, see: Adding VM managers in central mode (on page cccl).
30	KVM virtualization architecture could not be recognized. KVM virtualization was detected, but testing the virsh command to determine whether the architecture is x86 or Power failed. Contact BigFix Support.
40	Determining the number of processor sockets or cores failed. The returned result is not a number. Contact BigFix Support.
50	No valid host UUID could be retrieved by using the xe, dmidecode, and xl command. Verify that the UUID of the computer is correctly returned by SMBIOS. If it is not correctly returned, contact the hardware vendor to solve the problem.
Other codes	If other codes are returned, check the <code>run_vtech_scan.log</code> log file to learn running which command returned the code. The log file is in the installation directory of the BigFix client, by default: <code>var/opt/BESClient/LMT/vtech/run_vtech_scan.log</code> .

Related information

[Collecting capacity data from virtualization hosts for Xen and KVM \(on page ccxvi\)](#)

Import logs

Import logs contain information about the import of data from the BigFix server to BigFix Inventory. They list steps related to the import process and show their status. Each action in the log file is reported according to the UTC time zone, but the file also shows information about the local time zone of your server.

The logs are stored in one of the following directories:

- **Linux** `installation_directory/wlp/usr/servers/server1/logs/imports`
- **Windows** `installation_directory\wlp\usr\servers\server1\logs\imports`

The log from the latest import is also accessible from the BigFix Inventory user interface. To view the log, click **Management > Data Imports**. The last megabyte of the log data is displayed in the web user interface.

To download an earlier version of the import log from the user interface, identify the table row with the import that you are interested in, and click **Download**. The log file is saved in the download directory of your web browser.

Import Settings

Enabled

Imports per day: (times specified in UTC +01:00)

Time:

Import History

Start Time	User Name	Duration	Download
04/21/2015 10:...	Scheduled	0:01:27	Download

Start Time: 04/21/2015 08:09 AM
 Status: Successful
 Duration: 0:01:27

Import Log

The last megabyte of data from the import log:

```
duler for System x", :version=>"2.2.0", :tag_version=>0, :version_scheme=>"multipartnumeric", :delta=>nil,
:file_name=>"1988-03.com.ibm.ibm.Platform.Resource.Scheduler_for_System_x-2.2.0.swidtag", :license
2015-04-21 08:11:00 (+0:00:00.000) DEBUG: SAM::IsotagFact Finished parsing scan file. Lines: 2
2015-04-21 08:11:00 (+0:00:00.001) INFO: [sequel/lib/jdbc.rb:249] set identity_insert [sam].[isotags] off
2015-04-21 08:11:00 (+0:00:00.002) INFO: [sequel/lib/jdbc.rb:249] set identity_insert [sam].[isotags] on
```

Log level

By default, the data import logging level is set to INFO. You can change the setting to collect debug information about broken relationships, for example deleted component relationships. To enable debug logging, go to **Management > Advanced Server Settings**, and change the value of the **debug_logging_for_imports** parameter to true.

Enabling the debug logging level increases the size of the data import log, which might lead to errors while displaying the log in the user interface. This problem often impacts Internet Explorer and can be solved by accessing the user interface from a different browser. For more information, see [Troubleshooting \(on page dcclxvii\)](#).

Import statuses

Successful

The import finished successfully.

Incomplete

The import could not finish. It might happen when the BigFix Inventory server is stopped during the import. Ensure that the server is running, check the import logs, and rerun the import.

Failed

The import failed. Ensure that the BigFix server is running, check the import logs, and rerun the import.

Time zone

The time of all actions saved in the import log file is specified according to the UTC time zone. You can compare these actions to the local time zone of your server to learn when exactly a particular action occurred. Information about your server time zone is written in the beginning of the log file and is repeated every 1000 entries. The following examples are the excerpts from the log file:

Linux

```
[INFO]: Local Server Time Zone is +0200 Europe/Warsaw
```

Windows

```
[INFO]: Local Server Time Zone is CurrentTimeZone=-300 (UTC-06:00) DaylightBias=-60
```

Reviewing the import log

Review the following INFO messages in the import log to check how much data was transferred during the import.

Summary for complex table

Information about	Items specified in the import log	Description
Infrastructure	Computer items	The total number of computers in your environment. A computer is a system with a BigFix client that provides data to BigFix Inventory.
Software and hardware	SAM::ScanFile items	The number of files that have input data for the following items: <ul style="list-style-type: none"> • File system scan information (SAM::FileFact items) • Catalog-based scan information (SAM::CitFact items) • Software identification tag scan information (SAM::IsotagFact items)
	SAM::FileFactDelta items	The total count of information pieces about files that changed between the last two full file system scans.
	SAM::FileFact items	The total count of information pieces about files from all computers in your environment (contained in the processed scan files).
	SAM::CitFact items	The total count of information pieces from catalog-based scans (contained in the processed scan files).
	SAM::IsotagFact items	The total count of information pieces from software identification tag scans (contained in the processed scan files).
Installed packages	SAM::PackageFact items	The total count of information pieces about Windows™ packages that were gathered by the package data scan.
	SAM::UnixPackageFact items	The total count of information pieces about UNIX™ packages that were gathered by the package data scan.
Software usage	SAM::AppUsagePropertyValue items	The total number of processes that were captured during scans on the systems in your infrastructure.

Example:

```
INFO: Computer items: 15000
INFO: SAM::AppUsagePropertyValue items: 4250
INFO: SAM::ScanFile items: 30000
```

```

INFO: Delta changes applied on model SAM::FileFact: 0 rows
INFO: Number of computers processing delta file scan data: 0
INFO: SAM::FileFactDelta items: 0
INFO: Number of computers processing full file scan data: 16
INFO: Inserting new 28423 rows into SAM::FileFact
INFO: SAM::FileFact items: 15735838
INFO: SAM::IsotagFact items: 0
INFO: SAM::CitFact items: 149496
INFO: SAM::PackageFact items: 406687
INFO: SAM::UnixPackageFact items: 1922564

```

Server log file

The server log file `tema.log` is in the `installation_directory/wlp/usr/servers/server1/logs/` directory. The log file saves all actions related to the server and is the primary source of information for troubleshooting purposes. Each action in the log file is reported according to the UTC time zone, but the file also shows information about the local time zone of your server.



Note: This directory also contains the `production.log` file. It is an additional log that contains messages related to the database server and can be used to investigate database-related problems. Most of the information that is written in this log is also available in the `tema.log`.

Settings that you can modify

Logging properties are set in the `server.xml` file that is in one of the following directories:

- **Linux** `installation_directory/wlp/usr/servers/server1`
- **Windows** `installation_directory\wlp\usr\servers\server1`

You can change the following settings:

messageFileName

The server log file name. The default name is `tema.log`.

logDirectory

The server log file location. The default location is:

- **Linux** `installation_directory/wlp/usr/servers/server1`
- **Windows** `installation_directory\wlp\usr\servers\server1`

maxFileSize

The maximum size (in MB) that a log file can reach before it is rolled over. To disable this attribute, set the value to 0. The default value is 10.

maxFiles

If an enforced maximum file size exists, this setting is used to determine the maximum number of iterations of the server log file. If the log file reaches the maximum size that is defined in the **maxFileSize** parameter, the remaining messages roll over to another iteration of the file.

Log level

The log level controls which events are recorded in the server log file. The log level is set in the `log4j.properties` file that is in one of the following directories:

- **Linux** `installation_directory/wlp/usr/servers/server1/config`
- **Windows** `installation_directory\wlp\usr\servers\server1\config`

You can use the `log4j.rootCategory` property to change the logging level. The default level is INFO. The following is a list of the possible log levels:

FATAL

Task cannot continue and component, application, and server cannot function.

ERROR

Error events that might still allow the application to continue running.

WARN

Potential error or impending error. This level can also indicate a progressive failure.

INFO

General information that outlines the overall task progress.

DEBUG

Fine-grained informational events that are most useful to debug an application.

You need to change the log level to debug the BigFix Inventory server. To change the log level, you must modify the `log4j.properties` file in one of the following directories:

- **Linux** `installation_directory/wlp/usr/servers/server1/config`
- **Windows** `installation_directory\wlp\usr\servers\server1\config`

and `jvm.options` in one of the following directories:

- **Linux** `installation_directory/wlp/usr/servers/server1`
- **Windows** `installation_directory\wlp\usr\servers\server1`

To modify the file, perform these steps:

1. Change `#-DTEMA_LOG_DEBUG=true` to `-DTEMA_LOG_DEBUG=true`.
2. Restart BigFix Inventory.

Time zone

The time of all actions saved in the log file is specified according to the UTC time zone. You can compare these actions to the local time zone of your server to learn when exactly a particular action occurred. Information about your server time zone is written in the beginning of the log file and is repeated every 1000 entries. The following examples are the excerpts from the log file:

Linux

```
[INFO]: Local Server Time Zone is +0200 Europe/Warsaw
```

Windows

```
[INFO]: Local Server Time Zone is CurrentTimeZone=-300 (UTC-06:00) DaylightBias=-60
```

Enabling the logging of memory consumption

You can log information about the memory that is used by BigFix Inventory.

1. On the computer where the BigFix Inventory server is installed, go to one of the following directories, and create or modify the `server.env` file.

- **Linux** `install_dir/wlp/usr/servers/server1`
- **Windows** `install_dir\wlp\usr\servers\server1`

2. Add the following variable to the file.

```
LOG_MEMORY=true
```

3. Restart the BigFix Inventory server.

Each log returns memory consumption information.

```
INFO: Memory Used,458MB,Memory Diff,0kB,Memory Committed,743MB,
DEBUG: SF = # <SAM::ScanFile @values={:scan_file_id=>1, :file_size=>18681, :computer_id=>4,
:file_name=>*itsitsearch_0_3205710.xml.bz2*}>
INFO: Memory Used,460MB,Memory Diff,1835kB,Memory Committed,743MB,
DEBUG: Checking cache for /itsitsearch_0_3205710.xml.bz2
```

VM Manager Tool return codes

If you encounter any problems while configuring or managing the VM Manager Tool, refer to the log files for more information about any occurring errors. The log files for the VM Manager Tool are in the following directories:

- Trace log files: `BES Client\LMT\VMMAN\logs`
- Installation log files: `BES Client\LMT\VMMAN\logs\install`

For more information, see: [Log files \(on page cccxci\)](#).

The log file contains a return code that can help you check why the issue occurs and solve the problem.

Table 81. VM Manager Tool return codes

The table consists of two columns and 17 rows.

Return code	Description and possible solutions
0	No errors.
1	It failed to initialize the logs and the log files are not created. Ensure that the log4j.properties file is located in its default directory: <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;">Linux /var/opt/BESClient/LMT/VMMAN/config/</div> <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;">Windows C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\VMMAN\config\</div>
2	The security module did not start successfully. Ensure that <code>keys.jackes</code> keystore is located in its default directory: <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;">Linux /var/opt/BESClient/LMT/VMMAN/keydb/</div> <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;">Windows C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\VMMAN\keydb\</div>
3	It failed to lock the <code>.lock</code> file. When the file is locked, it means that the application is working properly. Locking the file might fail because the application is already running or because some other process is blocking the file. To solve the issue, find the process that is blocking the <code>.lock</code> file and stop it.
4	The used command is incorrect. Follow the displayed instructions and type the correct command.
5	It is not possible to read the configuration files. Ensure that the following folder contains all configuration files and that they are correct: <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;">Linux /var/opt/BESClient/LMT/VMMAN/config/</div> <div style="background-color: #f0f0f0; padding: 5px; margin: 5px 0;">Windows C:\Program Files (x86)\BigFix Enterprise\BES Client\LMT\VMMAN\config\</div>
6	The command did not complete successfully. For more information, check command output and the log file.
7	The application did not start in running mode because the service is not installed. Run the command <code>-install</code> and install the service.

Table 81. VM Manager Tool return codes

The table consists of two columns and 17 rows.

(continued)

Return code	Description and possible solutions
8	The service did not start in running mode. For more information, check the log file.
10	It failed to connect to VM Manager using the command <code>test connection</code> . For more information, check the command output.
12	It was not possible to add the VM Manager certificate to the keystore. The alias that was used to create the certificate already exists. Use another alias.
13	Adding the certificate did not complete successfully. For more information, see the log files.
14	It failed to change the password for keystore. For more information, check the command output and log files.
15	It is not possible to regenerate the encryption key for VM Manager. For more information, check the command output and log files.
16	The command <code>run once</code> did not perform any action because there are no VM Managers defined.

Changing and analyzing scanner trace settings

You can change the scanner trace settings to collect diagnostic data to help in investigating problems.

Changing scanner trace settings

If requested by HCL® support or to limit the amount of traces generated by the scanner tool, you can change the scanner trace settings.

1. Log on to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
3. Select **Edit Scanner Trace Settings**, specify the following values:
 - Trace level
 - Number of trace files
 - Trace file size

In the lower pane, click **Take Action**.

4. To select a subset of computers on which you want to create the capacity configuration, open the **Target** tab, and then click the computers.

Analyzing scanner trace settings

You can check scanner trace settings for computers, and see the number of trace files and trace file size.

1. Log on to the BigFix console.
2. In the navigation tree, click **Sites > External Sites > IBM BigFix Inventory v9 > Analyses**.
3. In the upper-right pane, select an **Scanner Trace Settings** analysis, and in the lower-right pane, click the **Result** tab.

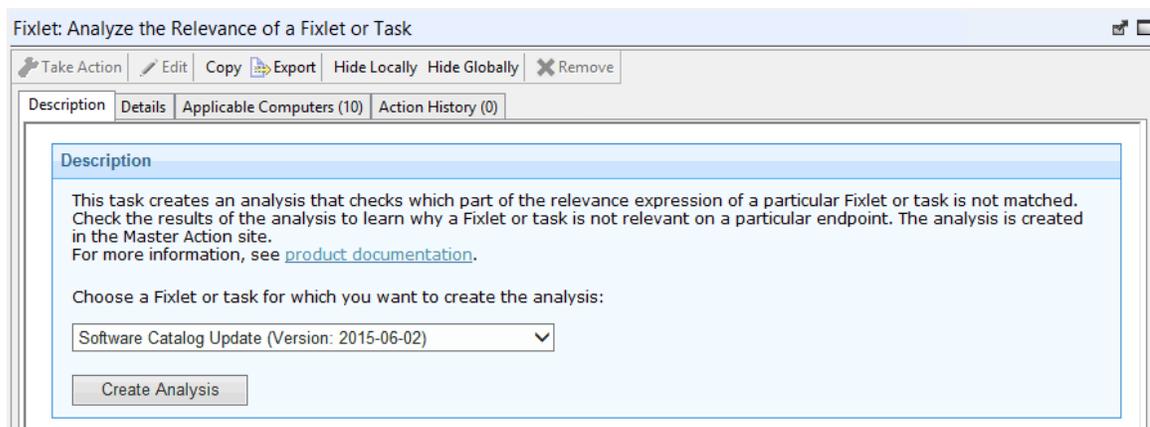
For each endpoint, you can view the following information:

- Computer host name
- Computer ID
- Trace level
- Number of trace files
- Trace file size

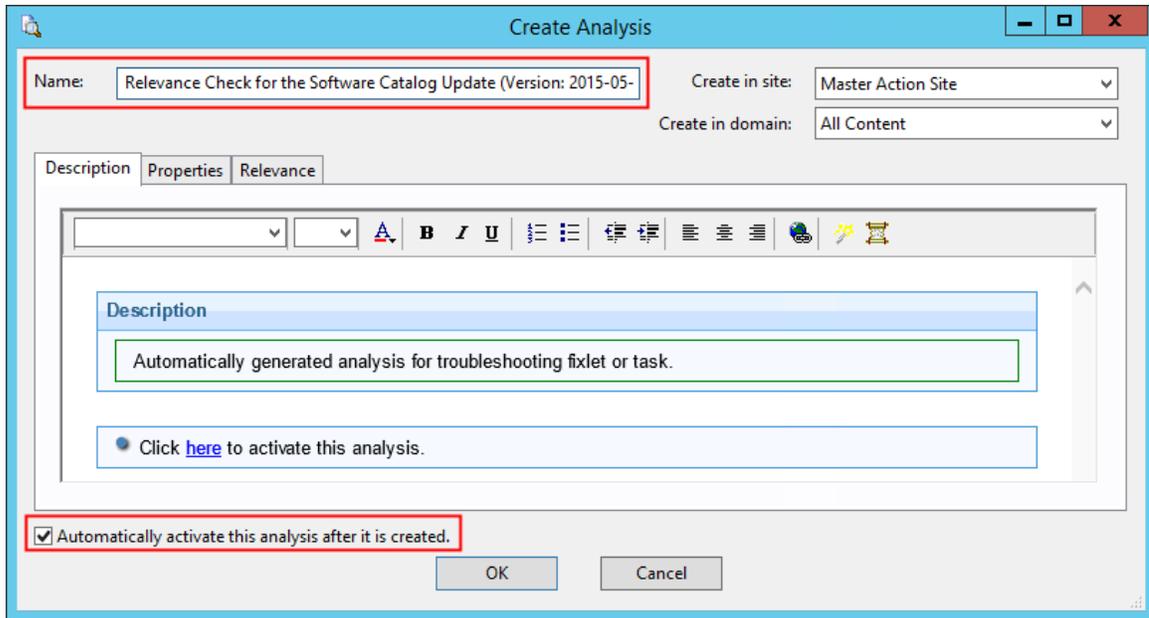
9.2.1 Checking why a fixlet or task is not relevant

9.2.1 Available from 9.2.1. When a fixlet or task is not relevant on a particular endpoint, you can create an analysis that checks which part of the relevance expression is not matched. This way, you can identify the exact reason why the fixlet or task cannot be run on an endpoint and troubleshoot the problem.

1. In the navigation tree of BigFix, click **Sites > External Sites > IBM BigFix Inventory v9 > Fixlets and Tasks**.
2. In the upper right pane, select **Analyze the Relevance of a Fixlet or Task**.
3. Choose the fixlet or task that you want to analyze. Then, click **Create Analysis**.



4. **Optional:** In the window that opens, you can change the display name of the analysis or decide whether the analysis is activated automatically.



5. To create the analysis, click **OK**.

By default, the analysis is created in the Master Action site and is activated globally.

6. Open the Master Action site and click **Analyses**.

7. In the upper right pane, select **Relevance Check for the name Task/Fixlet from the BigFix Inventory Site** and open the **Results** tab. The part of the relevance expression that causes that the fixlet or task is not relevant on a particular endpoint has the value `False`.

In this example, the second and third part of the relevance is not matched and thus causes that the entire task is not relevant.

Computer Name	Software Catalog Update (Version: 2015-05-13)_2	Software Catalog Update (Version: 2015-05-13)_3	Software Catalog Update (Version: 2015-05-13)_4
NC040208	True	False	False
NC91431261...	True	True	False
NC91431270...	True	True	False

8. When you know which part of the relevance expression is not matched, open the **Details** tab to see the content of the relevance.

In this example, the task is not relevant for two reasons. First, the targeted endpoint does not meet the operating system requirements (relevance 2). Second, the directory `SUA_Server_Path/properties/`

`version` or `SUA_Server_Path/iso-swid` does not exist on the endpoint or the latest software catalog is already available on that endpoint (relevance 3).

▼ Properties

Software Catalog Update (Version: 2015-05-13)_4
Period 30 days

```
( (if (name of operating system as lowercase starts with "win") then (true) else (name of operating system as lowercase starts with "linux" and (exists package "zip" of rpm) and (exists package "unzip" of rpm))) )
```

Software Catalog Update (Version: 2015-05-13)_3
Period 30 days

```
exists settings whose (name of it starts with "SUA_Server_Path" AND ((exists folder (value of it & "/properties/version") AND exists files of folder (value of it & "/properties/version"))) OR (exists folder (value of it & "/iso-swid") AND exists files of folder (value of it & "/iso-swid"))) AND not exists file (value of it & "/sua_catalog/SUA2_June_2015.zip") of client
```

Software Catalog Update (Version: 2015-05-13)_2
Period 30 days

```
( (name of operating system contains "Linux Red Hat Enterprise Server" and version of operating system >= "6.3" and version of operating system < "8.0") or (name of operating system as lowercase starts with "win2008r2") or (name of operating system as lowercase starts with "win2012") )
```

Software Catalog Update (Version: 2015-05-13)_1
Period 30 days

```
( (if (name of it as lowercase starts with "win") then (true) else ((name of it as lowercase starts with "linux") OR (name of it as lowercase starts with "aix") OR (name of it as lowercase starts with "hp-ux" AND ((architecture of it as lowercase contains "ia64") OR (family name of main processor as lowercase contains "pa-risc") OR (exists match (regex "^PA8[0-9]{3}(\s+)?$") of (family name of main processor)))) OR (exists match (regex "sunos 5\.(8|9|10|11)") of (name of it as lowercase) of operating system)) of operating system AND (if exists property "in proxy agent context" then ( not in proxy agent context ) else true) )
```

When you know the exact reason why a fixlet or task is not relevant on an endpoint, you can troubleshoot the problem.

The analysis is created in the Master Action site and thus is activated on all endpoints that report to BigFix, not only the endpoints that are subscribed to the BigFix Inventory site. To avoid performance issues, deactivate the analysis when you identify and troubleshoot the problem with the relevance of a fixlet or task. To deactivate an analysis, right-click it and click **Deactivate**.

Removing the server manually

If you encounter any problems with uninstalling the BigFix Inventory server, you can remove it manually.

1. Run the following command to verify that the process associated with the server is running:

```
ps -ef | grep <installation dir>/cli | grep server1
```

 **Tip:** The ID of the process is the number next to the user name.

2. End the process by running the following command:

```
kill -9 <process ID>
```

3. Remove the directory in which you installed BigFix Inventory:

```
rm -rf <installation dir>
```

For example:

```
rm -rf /opt/IBM/BFI
```

4. Edit the `.com.zerog.registry.xml` file that is in the `/var` directory. If you installed the server as a non-root user, the registry file is in `$HOME/`.
5. In the registry file, locate the entries that are related to BigFix Inventory. Remove the entries that represent the product and all the related components. You can recognize them by the product name and a common file path. If the registry file contains entries only for BigFix Inventory, you can delete the whole file.

TASK

[Uninstalling the server on Windows in interactive mode \(on page cli\)](#)

[Uninstalling the server on Windows in silent mode \(on page clii\)](#)

TASK

[Uninstalling the server on Linux in silent mode \(on page clxxvii\)](#)

[Uninstalling the server on Linux in interactive mode \(on page clxxvii\)](#)

Checking the client installation date

You can create a property on your BigFix server that enables the date when your endpoints first connected to the server.

This procedure might be useful if you are using an automated tool to silently install many BigFix clients. To verify that the clients were installed and check their installation dates, you can add a column that shows the dates when your clients first connected to the server.

1. Create a property on your BigFix server that enables the new column:

- a. Log in to the BigFix console.
- b. Click **Tools > Manage Properties**, and then click **Add New**.
- c. In the **Name** field, enter **First Report Time**.
- d. In the **Relevance** field, enter the following relevance expression:

```
(month of it as two digits & "/" & day_of_month of it as two digits & "/" & year of it as string)
of date (local time zone) of (minimum of subscribe times of sites)
```



Note: The preceding expression changes the date format to `DD/MM/YYYY`, which makes it suitable for sorting in a CSV format. If you want to keep the default format (Mon, 1 Jan 2014), use the following relevance expression: `minimum of subscribe times of sites`.

- e. In the **Evaluate** field, choose **12 hours**.
 - f. Click **OK**. The new column can be viewed in the **Computers** page.
2. **Optional:** Add the created column to Web Reports where your data can be viewed as reports, or exported to a CSV or PDF format:
 - a. Log in to Web Reports.



Tip: You can access Web Reports under the following web address:

`http://BigFix_server_hostname/webreports`.

- b. Click **Explore Data**.
- c. In the **Computer** section, click **Edit Columns**, and then select the **First Report Time** column.
- d. Click **Save Report**.
- e. **Optional:** To export your report, click **Export to CSV** or **Export to PDF**.

9.2.3 Changing the PVU per core value

9.2.3 Available from 9.2.3. The number of processor value units is assigned per processor core based on the information that is provided in the PVU table. It might happen that the value that is assigned to a processor in your environment is incorrect. In such case, manually change the PVU per core value to ensure that subcapacity licenses are properly calculated. When the problem that causes incorrect detection of the PVU value is solved, reset the PVU per core to the value that is specified in the PVU table.



You must have the Manage Hardware Inventory permission to perform this task.

The Manually Adjust the PVU per Core Value fixlet that was used in earlier versions of BigFix Inventory is deprecated. The PVU per core value can now be changed on the **Hardware Inventory** panel. If you previously changed the PVU per core value on some computers by using the fixlet, run the deprecated fixlet on these computers to remove the adjustment. Do it to avoid a situation in which the PVU values changed by the fixlet overwrite the values that you specify on the **Hardware Inventory** panel.

To properly calculate PVU values for Linux on IBM Power Systems, set the PVU per core value to 70.

1. In the top navigation bar, click **Reports > Hardware Inventory**.
2. Click the pencil in the **PVU per Core** column, and click **Change PVU per Core**.



Note: The pencil is not available for computers with the *No Scan Data* status. Capacity data is not available for such computers, and thus it is not possible to determine whether the PVU per core value was properly assigned or requires changing.

Status	Computer N...	IP Address...	Partition ...	Server ID	Server C...	Vendor	Brand	Type	Model	PVU per ...	Changed ...	Default P...
OK	NC042208	9.167.42.2...	2	Microsoft A...	2	Microsoft(R)	Microsoft(...	One core	All Existing	70	False	False
OK	NC042207	169.254.1....	2	IBM Corp. ...	8	Intel(R)	Xeon(R)	Multi-core	3400-36		Change PVU per Core	False
No VM ...	NC040191	9.167.40.1...	2	TLM_VM_...	2	Intel(R)	Xeon(R), ...	Multi-core	6500-659...	120	False	False

3. Provide the PVU value that you want to assign to the processor.

It will be changed for the server that is listed in the **Server ID** column and all virtual machines that are hosted on that server.



Important: Ensure that you provide a correct PVU per core value as it affects your PVU reports. In case of doubts, contact BigFix Support for guidance.

4. Provide a justification for changing the PVU value, for example a support case number. The justification is listed on the Audit Trail report. Then, click **Change**.

The changed PVU per core value is displayed on the Hardware Inventory report. To reflect the change on other reports, recalculate the data. Open the All Metrics report, and click **Recalculate**.

When the problem that causes incorrect detection of the PVU value is solved, reset the changed PVU per core value to the value that is provided in the PVU table. To do this, click the pencil in the **PVU per Core** column, and click **Reset PVU per Core**. Then, click **Reset**. The PVU per core value is reset on the server that is listed in the **Server ID** column and all virtual machines that are hosted on that server.

Status	Computer...	IP Address...	Partition ...	Server ID	Server Co...	Vendor	Brand	Type	Model	PVU per ...	Changed ...	Default P...
OK	NC042208	9.167.42.2...	2	Microsoft A...	2	Microsoft(R)	Microsoft(...	One core	All Existing	120	True	False
OK	NC042207	169.254.1....	2	IBM Corp. ...	8	Intel(R)	Xeon(R)	Multi-core	3400-36		Change PVU per Core	False
No VM ...	NC040191	9.167.40.1...	2	TLM_VM_4...	2	Intel(R)	Xeon(R), M...	Multi-core	6500-65		Reset PVU per Core	False

Updating scanner catalogs

Scanner catalogs are used by the scanner to discover software on the endpoints. The catalogs are automatically updated after each import of the BigFix software catalog. Use this procedure only if the automatic update of the scanner catalogs fails.

Before you force the update of scanner catalogs, ensure that the following requirements are met.

- The BigFix Inventory server is visible to the BigFix server.
- If Secure Socket Layer (SSL) is enabled in BigFix Inventory, the BigFix server recognizes SSL certificates of BigFix Inventory as valid.

1. Check whether the action that automatically updates scanner catalogs was created.
 - a. Log in to the BigFix console.
 - b. In the navigation bar, click **Actions**.
 - c. In the upper-right pane, locate the **Catalog Download (Version version)** action.
The source of the action should be Master Operator Site.
2. If the action exists, check whether it failed due to prefetch problems. If it does not exist, go to step 3 ([on page dcccxxv](#)).
 - a. Select the **Catalog Download (Version version)** action.
 - b. Open the **Computers** tab, and double-click a computer on which the status of the action is *Failed*.
 - c. Check whether the reason of the failure is a problem with prefetching the catalog.

```
Failed prefetch catalog.xml.bz2
```

If the status is *Failed*, communication between the BigFix server and the BigFix Inventory server is blocked. Ensure that the servers can communicate. Then, download the Catalog Download Fixlet to force the catalog update. If you cannot change the configuration to allow communication between servers, edit the fixlet so that the BigFix Inventory server can download catalogs from the BigFix server, and then run the fixlet.

3. Download the fixlet for forcing the update of scanner catalogs.
 - a. Log in to BigFix Inventory.
 - b. In the top navigation bar, click **Management > Catalog Upload**.
 - c. Click the question mark sign . Then, click **Catalog Download Fixlet**. Choose the location where you want to save the `catalog_download.bes` file, and click **Save**.
4. **Optional:** If the problem is caused by the lack of communication between servers, edit the `catalog_download.bes` file and substitute the `host_name` and `port` of the BigFix Inventory server with values that allow the server for downloading scanner catalogs from the BigFix server.

```
prefetch catalog.xml.bz2 sha1:24dcb13c743f2f92b0c5e9887e9df1d4491c4a66
size:398083 http://host_name:port/sam/catalogs/CIT_catalog_WINDOWS.xml.bz2
sha256:1e81c865d7fc96468649dbd5c334a2d77b12c5dd252671e22a7e5df0bd7ccbbe
```



Important: Change the `host_name` and `port` for all catalogs.

5. Upload the `catalog_download.bes` file to the BigFix console and run the Catalog Download (Version: `version`) fixlet.
 - a. Copy the file to the computer where the BigFix console is installed.
 - b. Log in to the BigFix console.
 - c. To import the file to the console, click **File > Import**.

- d. Open the directory where you store the `catalog_download.bes` file, select the file, and click **Open**. The file is imported.
 - e. In the left pane, click **Sites > Master Action Site > Fixlets and Tasks**. The list of available fixlets opens in the upper right pane.
 - f. Select **Catalog Download (Version: *version*)**, and click **Take Action**.
 - g. Select computers on which you want to run the fixlet, and click **OK**.
6. Wait for the next scheduled software scan and import of data or run these actions manually. For more information, see: [Initiating software scans \(on page cccii\)](#) and [Scheduling imports of data \(on page cclxii\)](#).
 7. **Optional:** If the problem is not solved, manually copy the catalog files from the BigFix Inventory server to the BigFix server.

- a. Log in to the computer where the BigFix Inventory server is installed and go to the following directory.

- **Linux** `/opt/ibm/BFI/wlp/usr/servers/server1/data/sam/public/catalogs`
- **Windows** `C:\Program Files\IBM\BFI\wlp\usr\servers\server1\data\sam\public\catalogs`

- b. Copy the following files to a temporary folder on the computer where the BigFix server is installed.

- `CIT_catalog_AIX.xml.bz2`
- `CIT_catalog_HPUX.xml.bz2`
- `CIT_catalog_I5OS.xml.bz2`
- `CIT_catalog_LINUX.xml.bz2`
- `CIT_catalog_SUN.xml.bz2`
- `CIT_catalog_WINDOWS.xml.bz2`

- c. Open the `catalog_download.bes` file that you downloaded in step 3 (on page dcccxxv) in a text editor, and check the `sha1` value for each catalog file.

For example, the `sha1` value for the catalog on Windows is `24dcb13c743f2f92b0c5e9887e9df1d4491c4a66`.

```
prefetch catalog.xml.bz2 sha1:24dcb13c743f2f92b0c5e9887e9df1d4491c4a66
size:398083 http://host_name:port/sam/catalogs/CIT_catalog_WINDOWS.xml.bz2
sha256:1e81c865d7fc96468649dbd5c334a2d77b12c5dd252671e22a7e5df0bd7ccbbe
```

Change the names of all catalog files to their `sha1` values.

- d. Copy the renamed catalog files to the following directory.

- **Linux** `/var/opt/BESserver/wwwrootbes/bfmirror/downloads/sha1`
- **Windows** `C:\Program Files (x86)\BigFix Enterprise\BES Server\wwwrootbes\bfmirror\downloads\sha1`

- e. Wait for the next scheduled software scan and import of data or run these actions manually.

Scanner catalogs are imported to the computers in your infrastructure and are used to discover the installed software.

Preserving bundling when the BigFix client is reinstalled or reverted from a snapshot

When the BigFix client is reinstalled or reverted from a snapshot, the next time it registers at the BigFix server, it receives a new Data Source Computer ID. The BigFix client with the old Data Source Computer ID becomes inactive because it does not report with the same Computer ID. As a result, the BigFix console shows duplicated entries for the same computer. Additionally, the affected computer loses bundling information for software that was discovered on this computer. To avoid this situation, store some registry keys data aside of the computer whose identity you want to preserve.

To preserve the identity of a computer on which the BigFix client is already installed, you need to prepare a gold image of this computer and store some registry keys data aside. The gold image is a simple VM template that can be deployed multiple times. This procedure is applicable in the following cases:

- You revert the computer from a VM snapshot with registry keys data stored aside. In this case, the computer keeps its previous Data Source Computer ID.
 - You revert the computer from a VM snapshot without registry keys data stored aside. In this case, the BigFix client needs to register on the BigFix server to receive its unique Data Source Computer ID.
1. The BigFix server can match the data that is stored aside to the BigFix client that is reinstalled or reverted from a snapshot when the **ClientIdentityMatch** parameter of the BigFix server is set to 100. By default, the parameter is set to 0. To change the value of the parameter, go to the computer on which the BigFix server is installed and perform the following steps.

- **Windows** Go to **Start > BigFix Administrative Tool > Advanced Options** and set the value of the **ClientIdentityMatch** parameter to 100.
- **Linux** Run the following command.

```
./BESAdmin.sh -setadvancedoptions -sitePvkLocation=/root/backup/license.pvk
-sitePvkPassword=pippo000 -update clientIdentityMatch=100
```

Ensure that you complete this action before you install the BigFix client on the computer whose identity you want to preserve. For more information about the **ClientIdentityMatch** parameter, see: [List of advanced options](#).

2. Install the BigFix client on the computer whose identity you want to preserve. For available methods, see: [Installing the BigFix clients \(on page cxxxvi\)](#).

After the installation of the BigFix client completes, the computer automatically registers at the BigFix server and receives a unique Data Source Computer ID.

3. Include the BigFix client in a computer image that can be deployed multiple times.
 - a. Stop the BigFix client.
 - **Windows** Open the Windows services dialog box and stop the BigFix client service.
 - **Linux** Run the following command.

```
BESClient stop
```

b. Go to the following location.

- **Windows** `HKEY_LOCAL_MACHINE\SOFTWARE\BigFix\EnterpriseClient\GlobalOptions`
- **Linux** `/var/opt/BESClient/besclient.config` under the section `[Software\BigFix\EnterpriseClient\GlobalOptions]`.

Remove values for the following parameters.

- `RegCount`
- `ComputerID`
- `ReportSequenceNumber`

c. Delete `__BESData` and `KeyStorage` folders from the default installation directory of the BigFix client.

- **Windows** `C:\Program Files\BigFix Enterprise\BES Client`
- **Linux** `/var/opt/BESClient`

Your computer is now ready to be imaged. It is important to save the gold image when the BigFix client service is stopped. To prepare a VM template, follow documentation for the virtualization technology that you use.

4. Reload the image without any computer identity that is stored aside and start the BigFix client.

- **Windows** Open the Windows services dialog box and start the BigFix client service.
- **Linux** Run the following command.

```
BESClient start
```

After you start the BigFix client service, the computer that is reloaded from this image registers automatically at the BigFix server and receives a new Data Source Computer ID.

5. To preserve the new identity of this computer, store some registry keys data aside before you power off the computer.

a. Stop the BigFix client service.

b. Store the following data.

- The registry key value for `ComputerID` from the following location.
 - **Windows** `HKEY_LOCAL_MACHINE\SOFTWARE\BigFix\EnterpriseClient\GlobalOptions`
 - **Linux** `/var/opt/BESClient/besclient.config` under the section `Software\BigFix\EnterpriseClient\GlobalOptions`.
- The `BES Client\KeyStorage` folder from the default installation directory of the BigFix client.

6. Reload the computer from the gold image with the BigFix client service stopped.

7. Import the data that you stored in step 5 (on page dcccxxviii).

8. Start the BigFix client.

The gold image of your computer is ready and you can reload it at any time. Thanks to the registry keys data stored aside, the computer keeps its identity and bundling information of software that was discovered on this computer.

You can write a script to automate some of the steps. For example, the script can:

- Start and stop the BigFix client service.
- Store and restore registry keys data that is needed to identify the computer.

If you want to deploy the image multiple times in your environment, it is beneficial to include the scripts in the gold image to help you make the process more effective.

Setting the code page for double-byte languages

Double-byte character corruption can occur in some environments because BigFix Inventory uses the code page to transcode data. For double-byte languages such as Korean, Japanese, and Chinese, set the code page to avoid character corruption.

The **_BESClient_DeploymentEncoding_IANAName** setting of the BESClient running on the BigFix server must use the correct character set for double-byte languages. For example, for simplified Chinese the setting is CP936. One double-byte code page translation is allowed per BigFix server that is configured as a data source for BigFix Inventory.

1. Log in to the BigFix console.
2. Go to your **Subscribed Computers**, and select the computer on which you have both the BigFix server and the BESClient installed.
3. Go to **Edit Settings**, and change the **_BESClient_DeploymentEncoding_IANAName** setting to the correct encoding.
4. Stop the BigFix Inventory server.
5. Log in to the database server as a user with DB2 authority, open a DB2 command line and run the following commands:

```
db2 "connect to database database_name"
db2 "update dbo.datasources_sequences set last_sequence=cast( x'0000000000000000' as char(8) for bit
data)'"
db2 commit
```

6. Start the BigFix Inventory server and run a data import.

Contacting Support

HCL Support provides assistance with product defects.

Before contacting Support, your company must have an active software maintenance contract, and you must be authorized to submit problems to HCL. For information about the types of maintenance contracts available, see "Enhanced Technical Support" in the *Software Support Handbook* at: <https://support.hcltechsw.com/csm>.

Complete the following steps to contact Support with a problem:

1. Define the problem, gather background information, and determine the severity of the problem. For help, see "Contacting Software Support" in the *Software Support Handbook*.
2. Gather diagnostic information.
3. Submit your problem to Support in one of the following ways:

dcccxxx

- Online: Visit the Support site: <https://support.hcltechsw.com/csm>.
- By phone: For the phone number to call in your country, go to the “*Contact Information*” section of the *Software Support Handbook*.

If the problem you submit is for a software defect or for missing or inaccurate documentation, HCL Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, HCL Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. HCL publishes resolved APARs on the HCL Support website daily, so that other users who experience the same problem can benefit from the same resolution.

Tuning performance

Learn how to plan the infrastructure of BigFix Inventory and to configure the application server to achieve optimal performance. The following guidelines are applicable in big data environments as well as in smaller environments that are running on low-performance hardware.

Your deployment architecture depends on the number of endpoints that you want to have in your audit reports. The BigFix Inventory server stores its data in a dedicated database, either DB2 or MS SQL Server. When you plan and install the BigFix Inventory infrastructure, consider the hardware requirements for all the components. For more information, see: [Hardware requirements \(on page cvii\)](#).

To benefit from all the good practices, it is recommended to use the BigFix server version 9.5 and the latest version of BigFix Inventory.

Infrastructure

The main factor that affects the performance of the BigFix Inventory server is the performance of the underlying infrastructure on which the application server and the DB2 database are running. You must ensure that the recommended requirements are met.

 **Important:** The official hardware recommendations were prepared with assumption that the BigFix Inventory server is running on a dedicated physical server or servers. If virtual machines are used, the specified hardware (processor, RAM, or disk space) must be exclusively dedicated to those virtual machines – not shared with any others.

Even if you meet the recommended requirements, the application might work slowly because of issues with the underlying hardware. The issues might be caused by, for example, sharing resources with other machines within the virtual environment. You can run some diagnostic methods to ensure that the underlying infrastructure is in good shape.

Make sure that you configured the VM Manager Tool efficiently. For more information, see: [Best practices for configuring VM managers \(on page cccxcvi\)](#).

Memory performance requirements

RAM performance metrics such as seek time or memory bandwidth, are strictly dependent on the memory type, which is linked with the processor type.

 **Restriction:** Do not run your server on a virtual machine for which memory compression was enabled, for example in x86 virtualization that is provided by VMware. It is especially important in case of the computer on which the DB2® database is running.

Pay particular attention to the amount of memory that is available for BigFix Inventory and the DB2® application. It is important to monitor the RAM usage on a daily basis because memory usage differs greatly depending on how the application is being used. The main factors besides the environment size include:

- The number of concurrent web UI users
- The frequency of data imports

Memory monitoring is especially important on the database computer as insufficient amount of RAM can significantly slow down background tasks, such as PVU data aggregation because of memory swapping. Swapping, also known as paging, is the use of secondary disk storage to store and retrieve application data for use in RAM. It is automatically performed by the operating system and typically occurs when the available RAM is depleted. Swapping can have a significant impact on DB2® performance and should be avoided. To avoid swapping, ensure that sufficient RAM is available on the computer so processes do not consume all the available RAM. Use the following Linux™ tools to monitor RAM usage and memory swapping:

- top
- free
- vmstat
- or sar

If you notice that there is almost no free memory available to BigFix Inventory and the DB2® processes are being swapped to disk, the best solution is to increase the RAM memory size.

Storage performance requirements

To ensure good performance of BigFix Inventory, it is essential to have good storage performance on the database server computer. The storage performance of the BigFix Inventory server alone is not that important.

The recommended average disk speed for the database server computer is specified in the following table.

Table 82. Recommended average disk speed for the database server computer in specific environments

Recommended average disk speed for the database server computer in specific environments

Deployment Size	Disk Speed
Small (up to 5000 endpoints)	100 MB/sec
Medium (5000 - 30000 endpoints)	250 MB/sec
Large (more than 30000 endpoints)	400 MB/sec

Running hard disk tests

Learn some diagnostic methods and use them to ensure that the hardware infrastructure is in good shape. On Linux™, you can use the `hdparm` and `dd` commands or the Bonnie++ application to determine the hard disk write rate. You can also use the `iostat` command to determine average disk usage.

Determining the write rate with the `hdparm` command

In the Linux operating system, there is a built-in `hdparm` command that can be used to determine the hard disk write rate.

In the Linux console, enter `hdparm -t path_to_the_test_file`

```
hdparm -t /dev/sdal
Timing buffered disk reads: 200 MB in 1.19 seconds = 167.47 MB/sec
```

Running the disk test with the DD application

The DD test measures the most basic single-threaded disk access, which is a large sequential write, followed by a large sequential read. It is relevant for the database performance testing effort because it gives you the maximum speed for sequential scans of large tables.



Note: Check the RAM size on the testing computer to correctly run a DD test and avoid file system caching that can distort the test results.

1. Create a file that is twice the size of the RAM capacity on the test computer and copy it to the target computer hard disk.

In your tests use a 4 KB block size because the default page size of the DB2® table space that BigFix Inventory is using is 4 KB. For example, if you have 16 GB of RAM, run the following command so that a 32 GB file named `ddfile` is created in the current directory:

```
time sh -c "dd if=/dev/zero of=ddfile bs=4k count=8000000 && sync".
```



Note: Some implementations of DD report the write rate, but the `time` command always returns the execution time that is larger than the DD time. The time and rate reported by DD represent the rate without any lag or synchronization time. Divide the data size by the time reported by the `time` command to get the real synchronous file writing rate.

Example:

```
8000000+0 records in
8000000+0 records out
32768000000 bytes (32 GB) copied, 38.924 s, 842 MB/s
real 3m49.192s
```

```
user 0m0.450s
sys 0m38.282s
```

In this example, the rate for writes including caching is 842 MB/s, but including the synchronization time it is much poorer – only about 136.46 MB/s (32 GB / 3 minutes 49 seconds).

2. To flush out the file system cache so that you read directly from the disk later, write to the disk another large file: `dd if=/dev/zero of=ddfile2 bs=4k count=4000000`
3. Read the first large file. Since the file system cache is filled with the second file, this test returns a valid read rate result:

```
time dd if=ddfile of=/dev/null bs=4k
8000000+0 records in
8000000+0 records out
32768000000 bytes (32 GB) copied, 186.456 seconds, 167.6 MB/s
real 3m6.496s
user 0m1.652s
sys 0m10.753s
```

4. Compare the results with those in the table [Recommended average disk write and read rates for the database server computer \(on page dcccxxxii\)](#) to determine whether your storage can handle the DB2® database for your environment. The system should easily be able to deal with a database of up to 10000 clients.

Determining disk usage

If response times are slow, check the average disk usage with the `iostat` command. Disk usage times of greater than 80% can reduce I/O performance, it is desirable to have average disk usage below 40%.

In the Linux console, enter `iostat -x interval duration`

Where:

- `interval` is the time interval in seconds. If no interval is specified, the output reflects the values over the entire period since the system was rebooted
- `duration` is the number of times to run the command.

Example:

```
iostat -x 30 5
```

Using relays to increase the performance of BigFix

To take advantage of the speed and scalability that is offered by BigFix, it is often necessary to tune the settings of the BigFix deployment.

A relay is a client that is enhanced with a relay service. It performs all client actions to protect the host computer, and in addition, delivers content and software downloads to child clients and relays. Instead of requiring every networked computer to directly access the server, relays can be used to offload much of the burden. Hundreds of clients can

point to a relay for downloads, which in turn makes only a single request to the server. Relays can connect to other relays as well, further increasing efficiency.

Reducing the BigFix server load

For all but the smallest BigFix deployments (< 500 BigFix clients), a primary BigFix relay should be set for each BigFix client even if they are not in a remote location.

The reason for this is that the BigFix server performs many tasks including:

- Gathering new Fixlet content from the BigFix server
- Distributing new Fixlet content to the clients
- Accepting and processing reports from the BigFix clients
- Providing data for the BigFix consoles
- Sending downloaded files (which can be large) to the BigFix client, and much more.

By using BigFix relays, the burden of communicating directly with every client is effectively moved to a different computer (the BigFix relay computer), which frees the BigFix server to do other tasks. If the relays are not used, you might observe that performance degrades significantly when an action with a download is sent to the BigFix server.

Setting up BigFix relays in appropriate places and correctly configuring clients to use them is the most important change that has highest impact on performance. To configure a relay, you can:

- Allow the clients to auto-select their closest BigFix relay.
- Manually configure the BigFix clients to use a specific relay.

For more information, see [Managing relays](#).

Setting up scans and import

Learn how to improve the scanning and importing activities to optimize the performance in your infrastructure.

Good practices for running scans and imports

Apply the good practices when running scans and imports to ensure that your infrastructure works efficiently.

Initial configuration task

- Run the initial import

It is a good practice to run the initial import before you schedule scans and activate analysis. This import uploads the software catalog from the installation directory to BigFix Inventory and extracts basic data about endpoints from the BigFix server. Make sure that the scan data from the first scan group is available in BigFix server and then run the second import. When the scan results from the second scan group are available in the BigFix server, you can run the next import of data.

Scanning related activities

- Plan scan frequency

After you find the optimal size of scan groups, set the frequency of software scans. The most common frequency is weekly so that every endpoint is scanned once a week. If your environment has more than 100 000 endpoints, consider performing scans less frequently, for example monthly. If scans are running daily, take into account system updates. When many files are modified, the next data import runs longer.

- Avoid scanning when it is not needed

Scan frequency depends on how often software changes on the endpoints and on your reporting needs. Group endpoints on which software changes dynamically together and scan them more frequently, for example once a week. Group endpoints with a more stable set of software together and scan them less frequently, for example once a month.

- Limit the number of computer properties gathered during scans

By default, BigFix Inventory retrieves four computer properties from the BigFix server: computer name, DNS name, IP address, and operating system. Imports can be much longer if you extract more properties from BigFix during each import. It is a good practice to limit the number of computer properties to 10 or fewer.

- Limit the number of BigFix Inventory computer groups

Create only as many computer groups as are needed. Import of data gets longer with a growing number of computer groups. If the size of your environment requires that you create many computer groups, consider skipping calculation of extended software aggregates. By skipping these calculations, you can reduce the length of data imports in very large environments. For more information, see: [Disabling the calculation of extended software aggregates \(on page dcccxlj\)](#).

- Enable the collection of checksums for a small group of endpoints

Each change to the configuration of the cryptographic hash collection (enabling, disabling, adding new types) significantly lengthens the first data import that follows the change. Because of multiple modifications on the file system, the new configuration triggers a complete data import instead of the delta one, in which only modifications are imported. This first data import might take up to three times longer, and the subsequent ones about 10% longer than data imports without file hashes. The impact of subsequent data imports is considered as moderate. Before you enable the collection of file hashes, divide your environment into scan groups to distribute the load of the imported data. For more information, see: [Enabling the collection of checksums \(on page dclx\)](#).

Import related activities

- Maintain frequent imports

After the installation, imports are scheduled to run once a day. Do not change this configuration. However, you might want to change the hour when the import starts. If your import is longer than 24 hours, you can

improve the scan groups configuration. Alternatively, you can preserve the current configuration because BigFix Inventory handles overlapping imports. If an import is running, no other import is started.

- Consider web user interface

Data import is a computation-intensive task so be prepared to experience slower user interface response times while you are using BigFix Inventory. Thus, it is better to schedule the imports to take place at other times when you are not likely to use the application web UI.

If, in case of huge deployments, loading any report takes substantial amount of time, see: [User interface problems \(on page dccxci\)](#).

Dividing the infrastructure into scan groups

To avoid running into performance issues, you should divide the computers in your infrastructure into scan groups and properly set the scan schedule. You should start by creating a benchmark scan group on which you can try different configurations to achieve optimal import time. After the import time is satisfactory for the benchmark group, you can divide the rest of your infrastructure into analogous scan groups.

1. Create a single scan group that will be your benchmark.

The size of the scan group might vary depending on the size of your infrastructure. However, the recommendation is to avoid creating a group larger than 20 000 endpoints. For more information about creating computer groups, see: [Lesson 1 \(Optional\): Creating computer groups in the BigFix console \(on page dclxx\)](#)

2. Run **Initiate Software Scan** fixlet and scan the computers in this scan group. For more information, see: [Initiating software scans \(on page cciii\)](#).
3. When the scan finishes, use **Upload Software Scan Results** fixlet to upload the results of the scan to the BigFix server. For more information, see: [Uploading software scan results \(on page ccviii\)](#).
4. Run an import of data. For more information about running imports, see: [Scheduling imports of data \(on page cclxii\)](#). Check the import time and decide whether its duration is satisfactory.
5. If you are not satisfied with the import time, go to the following directory and check the import log.
 - **Linux** `installation_directory\wlp\usr\servers\server1/logs/imports`
 - **Windows** `installation_directory\wlp\usr\servers\server1\logs\imports`
6. Undertake one of the following actions:
 - If you see that the duration of the import of raw file system scan data or package data takes longer than one third of the ETL time and the volume of the data is large (a few millions of entries), create a smaller group.
 - If you see that the duration of the import of raw file system scan data or package data takes longer than one third of the ETL time but the volume of the data is low, fine-tune hardware. For information about processor and RAM requirements as well as network latency and storage throughput, see: [Hardware requirements \(on page cvii\)](#).

7. If you see that processing of usage data takes an excessive amount of time and you are not interested in collecting usage data, disable gathering of usage data. For more information, see: [Disabling the collection of software usage \(on page cclxviii\)](#).
8. After you adjust the first scan group, run the software scan again, upload its results to the BigFix server and run an import of data.
9. When you achieve an import time that is satisfactory, decide whether you want to have a shorter scan cycle.

If you have an environment that consists of 42 000 endpoints and you created seven scan groups of 6000 endpoints each, your scan cycle will last seven days. To shorten the scan cycle, you can try increasing the number of computers in a scan group, for example, to 7000. It will allow you for shortening the scan cycle to six days.



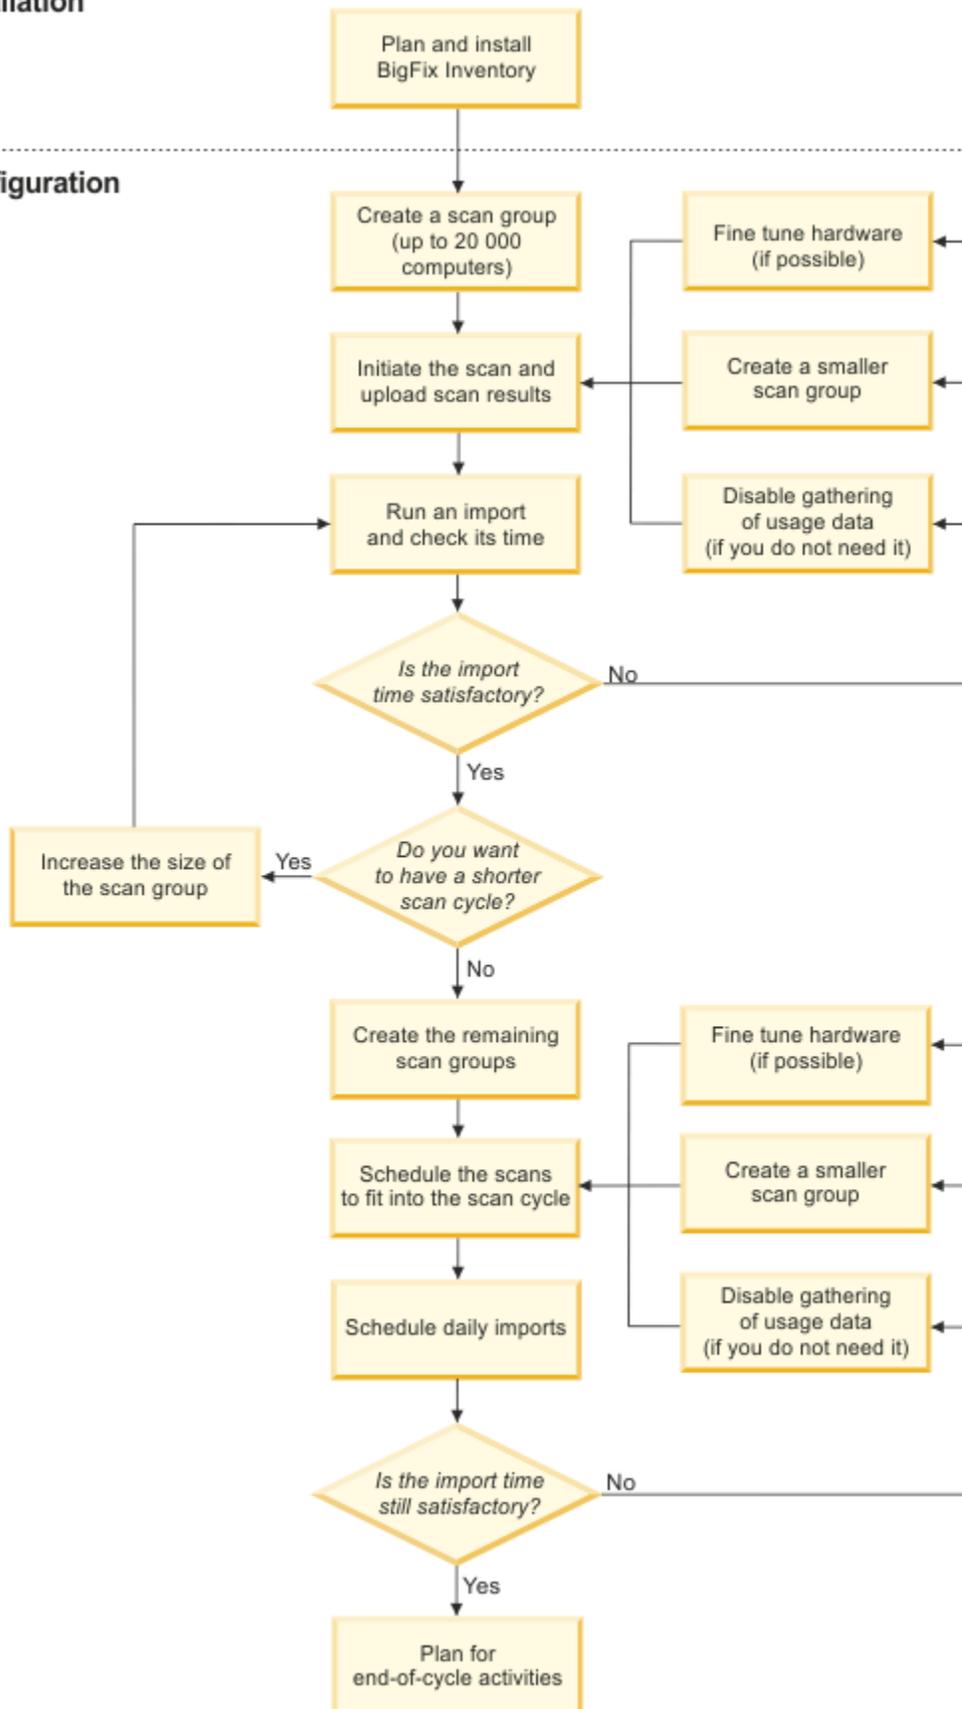
Note: After you increase the scan group size, observe the import time to ensure that its performance remains on an acceptable level.

When you are satisfied with the performance of the benchmark scan group, create the remaining groups. Schedule scans so that they fit into your preferred scan cycle. Then, schedule import of data from BigFix server. Observe the import time. If it is not satisfactory, adjust the configuration as you did in the benchmark scan group. When you achieve suitable performance, make room for end-of-scan-cycle activities and plan to have a data export to other integrated solutions (i.e. SmartCloud Control Desk through BigFix Tivoli Integration Composer) at the end of a 1- or 2-week cycle. For more information, see: [Integrating with IBM Control Desk \(on page mcxxxv\)](#).

Use the following diagram to get an overview of actions and decisions that you will have to undertake to achieve optimal performance of BigFix Inventory.

Installation

Configuration



Scan schedule

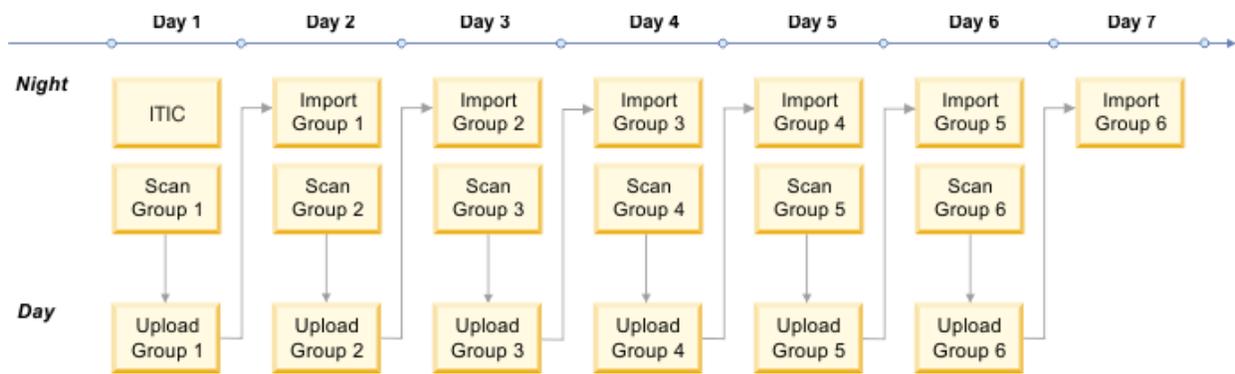
It is critical for BigFix Inventory performance that you accurately schedule scans in the scan groups in your environment. If the configuration is not well-balanced, you might experience long import times.

For environments larger than 35 000 endpoints, divide your endpoints into separate scan groups. The system administrator can then set a different scanning schedule for every scan group in your environment.

Example

If you have 60 000 endpoints, you can create six scan groups (every group containing 10 000 endpoints). The first scan group has the scanning schedule set to Monday, the second to Tuesday, and so on. Using this configuration, every endpoint is scanned once a week. At the same time, the BigFix server receives data only from 1/6 of your environment daily and for every daily import the BigFix Inventory server needs to process data only from 10 000 endpoints (instead of 60 000 endpoints). This environment configuration shortens the BigFix Inventory import time.

The image below presents a scan schedule for an infrastructure that is divided into six scan groups. You might achieve such a schedule after you implement recommendations that are contained in this guide. The assumption is that both software scans and imports of scan data to BigFix Inventory are scheduled to take place at night, while uploads of scan data from the endpoints to the BigFix server occur during the day.



If you have a powerful server computer and longer import time is acceptable, you can create fewer scan groups with greater number of endpoints in the BigFix console. Remember to monitor the import log to analyze the amount of data that is processed and the time it takes to process it.

For information how to create scan groups, see the topic [Computer groups](#) that is available in the BigFix documentation.

Recommendations for the service provider environments

The service provider functionality is used to create separate subcapacity reports for different sets of computers. You can create such a report for each computer group. By default, the calculation of PVU and RVU MAPC consumption is disabled for new computer groups. After you enable it, a computer group becomes a subcapacity computer group, and can have its own subcapacity report. This, however, impacts performance due to increased number of subcapacity calculations.

If you decide to create more than 10 subcapacity computer groups, adjust the use of CPU resources and tune the default configuration of BigFix Inventory. The most important processes for subcapacity computer groups are aggregation and reaggregation. Performance of these processes depends on the maximum number of threads that can be run for the calculations, which by default is 2. You improve performance by increasing the number of available threads.

To increase the number of available threads, specify the values of the **maxAggregationThreads** and **maxReaggregationThreads** parameters. To specify the values of these parameters, log in to BigFix Inventory and go to **Management > Advanced Server Settings**. In general, you should provide two processor cores for each thread on your database server. For example, if you can provide 12 processor cores for the aggregation process itself, increase the number of threads to 6 by specifying **maxAggregationThreads=6**. For more information, see: [Advanced server settings \(on page cdii\)](#).

Additional tasks for improving performance

Learn about additional tasks that you might want to perform to improve the performance in your infrastructure.

Changing the mode of sorting bundling options on the **Software Classification** panel

If the Software Classification panel runs slowly or the BigFix Inventory server is under a heavy load, you might want to change the mode of sorting bundling options.

After you invoke the function **Assign Components to Product** on the Software Classification panel, you might experience a performance issue while loading the data. By default, the bundling options that are displayed when you reassign a software component are sorted by confidence. To improve the performance, sort the displayed bundling options alphabetically. You can use a REST add-on in your web browser or the curl tool to perform this task.

- You can use the REST add-on in your web browser to sort the displayed bundling options alphabetically.
 - Open the REST add-on in your web browser, for example Advanced REST client.
 - To set the value of the **blockUiBundlingComputations** parameter to *true*, run the following REST API query:

```
PUT http://bfi_server_host_name:port_number/rest/configs?
token=token&name=blockUiBundlingComputations&value=true
```

Example:

```
PUT
http://localhost:9981/api/sam/configs?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623&name=blockUiBundlingComputations&value=true
```

- You can use the curl tool to sort the displayed bundling options alphabetically.
 - Start the curl tool.
 - Enter the following command.

```
curl -v PUT
  http://bfi_server_host_name:port_number/rest/configs?token=token&name=blockUiBundlingComputati
ons&value=true
```

Example:

```
curl -v PUT
  http://localhost:9981/api/sam/configs?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623&name=blockUiBundlingComputations&value=true
```

The bundling options are sorted alphabetically and are displayed more quickly because no additional computation is added to the server workload.

Related information

[Advanced server settings \(on page cdii\)](#)

Disabling the calculation of extended software aggregates

You can noticeably reduce the length of data imports in very large environments if you disable the calculation of extended software aggregates. Software aggregates are pieces of data that are gathered during software scans on your endpoints and which are later processed during successive data imports. They are displayed on the **Inventory Exploration** report.

You can disable the gathering of the following software aggregates:

- Total Runs
- Total Run Time
- Average Runs per Day
- Average Run Time
- First Used

1. Enable the new feature that allows for skipping the processing of extended software aggregates.
 - a. Open the following URL in your web browser: `https://bfi_server_name:port/management/feature`.
 - b. Select **Skip Extended Software Aggregates** and click **Save**.
2. To disable the calculation of extended software aggregates for a specific computer group:
 - a. In the BigFix Inventory web user interface, click **Management > Computer Groups**.
 - b. Expand the **All Computers** computer group and click the group for which you want to disable the calculation of software aggregates.
 - c. Select **Skip Extended Software Aggregates** and click **Save**.

After you run a data import, the individual entries for extended software aggregates are displayed on the **Inventory Exploration** report as **Not calculated**.

Run a data import to recalculate the aggregates.

Recovering from accumulated scans

If scan results are not imported to BigFix Inventory for a longer period of time, they accumulate on the BigFix server. They are all processed during the first import after the break which might lead to long import times. To avoid the problem, skip importing the results of the software scan during the first import after the break. Then, apply good practices for distributing scans in your infrastructure, including dividing the computers into scan groups, and planning the scan schedule. It will allow for importing scan results from the consecutive scan groups without overloading the import.

Recovering from accumulated scans is necessary in the following situations:

Import of data to BigFix Inventory was failing for a longer period of time

If the import of data fails for a longer period of time, scan results accumulate on the BigFix server. They are all processed during the first import after the break which might lead to long import times.

BigFix Inventory was reinstalled and scans were run in the past

With each reinstallation of BigFix Inventory, a new instance of BigFix Inventory database is created. A fresh installation triggers the import of all historically collected scan files.

BigFix Inventory V9.x coexists with Software Use Analysis V1.3, V2.2, License Metric Tool V7.x or Tivoli Asset Discovery for Distributed V7.x

Different versions of License Metric Tool 9.x and BigFix Inventory 9.x can coexist with each other and the earlier versions as well as Tivoli Asset Discovery for Distributed. For each additionally installed or reinstalled instance of the License Metric Tool 9.x or BigFix Inventory 9.x server, a new database is created. It triggers the import of all historically collected scan files.

For more information about coexistence scenarios, see [Coexistence scenarios](#).

1. Log in to BigFix Inventory and go to the following URL: <https://host:port/management/feature>.
2. Check the **Disable the import of accumulated software scan results (next import only)** check box.
3. Run the import of data (ETL).

If the accumulated scan results are successfully skipped, the following information is written in the import log:

```
2015-08-10 22:11:35 (+0:00:00.003) INFO: Accumulated results of software scans
are SKIPPED
2015-08-10 22:11:35 (+0:00:00.119) INFO: Accumulated results of software scans
were successfully skipped during the import. The setting applied only to this import.
The subsequent imports will contain the results of the software scan.
```

If the import is successful, the function will be disabled.

If the computers in your environment are divided into scan groups, run the scans according to the specified schedule and wait for the scan results. Otherwise, [divide the computers into scan groups \(on page dcccxxvii\)](#) and distribute scans

over the days of the week according to the good practices for running scans. Then, [schedule incremental imports of data \(on page dcccxxv\)](#) to BigFix Inventory.

Shortening the retention period gradually to avoid problems with growing database size

Since BigFix Inventory 9.2.1, retention period is set to 7 days in a fresh installation, so this procedure does not apply in case of fresh 9.2.1 installations. If you did not configure the retention period after installation in earlier versions of BigFix Inventory and the data accumulation continued for several months, you must enable and repeatedly change the data retention configuration to make the retention period as short as possible. With each successive import, the historical data is deleted gradually and thus problems with growing database size are avoided.

For information about how to configure the retention period in the product web user interface, see [Configuring data retention period for raw utilization data \(on page ccxcviii\)](#).



Note:

- After this configuration was applied, the transaction log is used heavily because all the data is deleted in one transaction.
 - The database size is not decreased automatically. If you want to reclaim disk space that is occupied by database tables, run the command that is appropriate for your database management system.
 - If the data imports during which historical data was deleted failed, increase the retention period and run the import again.
1. Determine the date of the first successful import.
 2. Calculate the number of days since the first successful import. For example: If the import occurred 6 months ago: 6 months * 30 days = 180 days.
 3. Set the retention period to begin 30 days after the first successful data import.
Example: 5 months * 30 days = 150 days.
 4. Run the data import.
 5. Set the retention period to begin 60 days after the first successful data import, that is, set the retention period to 120 days.
Example: 4 months * 30 days = 120 days.
 6. Repeat the step as many times as necessary, each time decreasing the retention period by 30 days.
 7. When the retention period is short enough, enter the target configuration of 7 days.

Tuning performance in medium and large environments

You can learn the possible causes of the most common performance issues in medium and large environments. Follow the tips and the known solutions to these issues to keep the system in good shape and avoid performance problems.

Scalability of multiple concurrent UI users

The number of users that are simultaneously logged in to BigFix Inventory and perform actions on the user interface influences the application performance. When BigFix Inventory is running on the hardware that meets the minimal hardware requirements, the optimal number of concurrent users is five.

Tuning DB2 application database

To avoid performance issues in medium and large environments, configure the location of the transaction log and adjust the log size. If you are using DB2 as BigFix Inventory database, you might want also to change the swappiness parameter in Linux or configure the **DB2_COMPATIBILITY_VECTOR**.

Make sure that the statement concentrator is set to OFF on the DB2 instance where the BigFix Inventory is installed. For more information, see: [stmt_conc - Statement concentrator configuration parameter](#).

Maintaining the DB2 database

To ensure that the DB2® database works properly, you need to regularly back up your data and perform maintenance actions.

The configuration of DB2® has a significant impact on performance. You should perform some standard actions to administer the database correctly. If you have DB2® administration professionals, they can monitor the DB2® health and performance. In other cases, you should follow the procedures that are described in this topic to keep your database healthy.

1. Perform regular backups of the data that is stored in the database. It is advisable to back up the database before updating the software catalog or upgrading the server to facilitate recovery in case of failure.
 - a. Stop the server.
 - b. In the DB2® command-line interface, run the following command: `db2 backup database TEMADB.`



Note: TEMADB is the default database name. If you are unsure whether it applies to your database, see: [Checking the database name \(on page dcclxx\)](#).

- c. Start the server

The backup is created in the current working directory, and the file name contains the instance name and time stamp of the backup procedure (for example: `TLMA.0.db2inst1.NODE0000.CATN0000.20101105000715.001`).

For more information about database backup strategies, see: [Developing a backup and recovery strategy](#).

2. Reorganize the table to match the index and to reclaim space:
 - a. Stop the BigFix Inventory server.
 - b. In the DB2® command-line interface, run the following commands:

- `db2 connect to TEMADB`
- `db2 -x "select 'reorg table',substr(rtrim(tabschema)||'.'||rtrim(tabname),1,50),' allow no access;'from syscat.tables where type = 'T' and tabschema in ('ADM','SAM','DBO') order by tabschema,tabname " > reorgs.sql`
- `db2 -tvf reorgs.sql`
- `db2 terminate`

c. Start the BigFix Inventory server.

3. Keep the statistics up-to-date. By default, DB2® statistics are run automatically. If this option is disabled, you must manually run the following commands:

a. Stop the BigFix Inventory server.

b. In the DB2® command-line interface, run the following commands:

- `db2 connect to TEMADB`
- `db2 -x "select 'runstats on table',substr(rtrim(tabschema)||'.'||rtrim(tabname),1,50),' and indexes all;'from syscat.tables where type = 'T' and tabschema in ('DBO','SAM','ADM') order by tabschema,tabname " > runstats.sql`
- `db2 -tvf runstats.sql`
- `db2 terminate`

c. Start the server.

For a more advanced investigation of queries that are used by the BigFix Inventory server, you should use the DB2® design advisor command (`db2advvis`). It helps in finding new indexes that can improve the database performance. The Design Advisor uses the output from the DB2® monitors to suggest the creation of new indexes. The suggestions are based on the queries that are found in monitors. For more information about Design Advisor, consult the DB2® documentation:

- DB2 10.1
 - [db2advvis - DB2 Design Advisor command](#)
 - [DB2® topic: Design Advisor](#)
- DB2 10.5
 - [db2advvis - DB2 Design Advisor command](#)
 - [DB2® topic: Design Advisor](#)
- DB2 11.1
 - [db2advvis - DB2 Design Advisor command](#)
 - [DB2® topic: Design Advisor](#)

Configuring the transaction logs size

If your environment consists of many endpoints, increase the transaction logs size to improve performance.

The transaction logs size can be configured through the **LOGFILSIZ** DB2 parameter that defines the size of a single log file. To calculate the value that can be used for this parameter, you must first calculate the total disk space that is required for transaction logs in your specific environment and then multiply it, thus obtaining the size of one transaction log. The required amount of disk space depends on the number of endpoints in your environment and the number of endpoints in the biggest scan group for which data is processed during the import.

 **Important:** Use the provided formula to calculate the size of transaction logs that are generated during the import of data. More space might be required for transaction logs that are generated when you remove the data source.

1. Use the following formula to calculate the disk space that is needed for transaction logs:

```
<The number of computers> × 0.0012 GB + <the number of computers in the biggest scan group> × 0.0012
GB + 17 GB
```

2. To obtain the size of a single transaction log file that can be specified in the **LOGFILSIZ** DB2 parameter, multiply the result by 1852.

 **Note:** The number 1852 expresses the relation between the primary and secondary log files and is necessary to calculate the size of a single transaction log file (**LOGFILSIZ**). The factor was calculated assuming the default number of log files (**LOGPRIMARY** = 25 and **LOGSECOND** = 110).

3. Run the following command to update the transaction log size in your database. Substitute *value* with the size of a single transaction log.

 **Note:** **TEMADB** is the default database name. If you are unsure whether it applies to your database, see: [Checking the database name \(on page dclclxx\)](#).

```
db2 update database configuration for TEMADB using logfilsiz value
```

4. For the changes to take effect, restart the database. Run the following commands:

```
db2 deactivate db TEMADB
db2stop
db2start
db2 activate db TEMADB
```

5. Restart the BigFix Inventory server.

- a. To stop the server, run the following command:

```
/etc/init.d/BFIserver stop
```

b. To start the server, run the following command:

```
/etc/init.d/BFIserver start
```

Calculating the single transaction log size for 100 000 endpoints and 15 000 scan results:

```
100 000 × 0.0012 GB + 15 000 × 0.0012 GB + 17 GB = 155 GB
155 × 1852 = 287060
```

287060 is the value to be specified in the **LOGFILSIZ** parameter.

Configuring the transaction log location for DB2

To increase database performance, move the DB2® transaction log to a file system that is separate from the DB2® file system.

Medium environments: **Strongly advised**

Large environments: **Required**

Very large environments: **Required**

1. To move the DB2® transaction log to a file system that is separate from the DB2® file system, update the DB2® **NEWLOGPATH** parameter for your BigFix Inventory database:

```
UPDATE DATABASE CONFIGURATION FOR TEMADB USING NEWLOGPATH value
```

Where *value* is a directory on a separate disk (different from the disk where the DB2® database is installed) where you want to keep the transaction logs. This configuration is strongly advised.

2. For the changes to take effect, restart the database. Run the following commands:

```
DEACTIVATE DB TEMADB
DB2STOP
DB2START
ACTIVATE DB TEMADB
```

3. Restart the BigFix Inventory server.

a. To stop the server, run the following command:

```
/etc/init.d/BFIserver stop
```

b. To start the server, run the following command:

```
/etc/init.d/BFIserver start
```

Configuring swappiness in Linux hosting DB2 database server

Swappiness determines how quickly processes are moved from RAM to hard disk to free memory. It can assume the value 0 - 100. A low value means that your Linux™ system swaps out processes rarely while a high value means that processes are written to disk immediately. Swapping out runtime processes should be avoided on the DB2® server on Linux™, so it is advisable to set the swappiness kernel parameter to a low value or zero.

1. Log in to the Linux™ system as root.
2. Set the swappiness parameter to a low value or 0.
 - Option A:
 - a. Open the file `/etc/sysctl.conf` in a text editor and enter the **vm.swappiness** parameter of your choice.

Example:

```
vm.swappiness = 0
```

- b. Restart the operating system to load the changes.
 - Option B: To change the value without restarting the operating system, run the following command:
`sysctl -w vm.swappiness=0.`

Configuring the DB2_COMPATIBILITY_VECTOR variable for improved UI performance

For environments with 5000 or more clients in your infrastructure, it is advisable to set the value of the **DB2_COMPATIBILITY_VECTOR** variable to **MYS**. This change might result in a UI response time that is significantly faster on some BigFix Inventory installations.

For information about how to modify this registry variable, see [DB2_COMPATIBILITY_VECTOR registry variable](#) in HCL® Knowledge Center.

Tuning MS SQL application database

To avoid performance issues in medium and large environments, configure the location of the transaction log. If you are using MS SQL as BigFix Inventory database, you might want to shrink the transaction log or update query optimization statistics.

Optimizing the tempdb database in Microsoft SQL Server

tempdb is a system database in SQL Server whose main functions are to store temporary tables, cursors, stored procedures, and other internal objects that are created by the database engine.

By default, the database size is set to 8 MB and it can grow by 10% automatically. In large environments, its size can be as large as 15 GB. It is therefore important to optimize the tempdb database because the location and size of this database can negatively affect the performance of the BigFix Inventory server.

For information about how to set the database size and how to determine the optimal number of files, see the TechNet article [Optimizing tempdb Performance](#).

MS SQL Server transaction log related activities

Perform the following activities to improve the performance and the reliability of your environment.

- Shrink the transaction log once a month. If the environment is large, it is advisable to shrink the log after every data import.

For more information, see: [How to: Shrink a File \(SQL Server Management Studio\)](#).

- The transaction log in MS SQL Server records all transactions and the database changes caused by each transaction.

For information how to configure the database transaction log, see: [Move the Database Transaction Log to Another Drive](#).

Rebuilding database indexes

Rebuilding database indexes and updating the statistics is mandatory to allow BigFix Inventory operate smoothly. In case you experience delay in UI response or importing data, make sure the defragmentation is done and statistics are up to date.



Note: It is recommended to run the operation once in a week to rebuild database.

The below example shows the simple approach to rebuild indexes and update statistics. It rebuilds all indexes independent of their fragmentation. The script assumes standard database name as TEMADB.

```
USE TEMADB
GO
IF EXISTS (SELECT
*
FROM dbo.imports
WHERE success IS NULL)
BEGIN
PRINT N'CANNOT RUN index rebuild. BFI import is running!'
PRINT N'Wait until BFI import finishes'
END
ELSE
BEGIN
DECLARE table_cursor CURSOR FOR
SELECT
table_schema,
table_name
FROM INFORMATION_SCHEMA.TABLES
WHERE table_type = 'BASE TABLE'
OPEN table_cursor
```

```

DECLARE @tableName sysname

DECLARE @tableSchema sysname

FETCH NEXT FROM table_cursor INTO @tableSchema, @tableName

WHILE @@fetch_status != -1

BEGIN

PRINT N'START alter index all on ' + @tableSchema
+ N'.' + @tableName + N' rebuild';

EXECUTE (N'alter index all on ' + @tableSchema + N'.'
+ @tableName
+
N' rebuild')

PRINT N'END alter index all on ' + @tableSchema
+ N'.' + @tableName + N' rebuild';

FETCH NEXT FROM table_cursor INTO @tableSchema,
@tableName

END

CLOSE table_cursor

DEALLOCATE table_cursor

PRINT N'START sp_updatestats';

EXECUTE sp_updatestats

PRINT N'END sp_updatestats';

END

GO

```

REST API considerations

You can use BigFix Inventory REST API to retrieve large amounts of data that is related to computer systems, software instances, and license usage in your environment. Such information can then be passed to other applications for further processing and analysis.

Although using single API requests to retrieve data only from a selected subset of computers does not greatly impact the performance of BigFix Inventory, this is not true when retrieving data in bulk for all your computer systems at the same time. Such an action requires the processing of large amounts of data and it always influences the application performance and when possible should not be executed during data import.

In general, the API requests should not be used together with other performance intensive tasks, like software scans or data imports. Each user that is logged in to the application, as well as the number of actions that are performed in the web user interface during the REST API calls also decrease the performance.



Important: Each time you want to retrieve data through REST API, ensure that the use of BigFix Inventory at a moderate level, so that the extra workload resulting from REST API does not overload the application and create performance problems.

Special considerations

1. The application that calls the API controls the size of individual response for API calls. BigFix Inventory server should have enough memory to receive and save the responses.
2. Check only for changes in data and do not fetch all the data every time. Use *valid_from* / *valid_to* or *discovery_start* / *discovery_end* fields to retrieve changes.
3. When you complete your data import, BigFix Inventory no more calls APIs. In case, you attempt to perform this process, the request may fail with a misleading error. It is recommended to implement an option for the application to retry the request after some time.

Optimizing response time

Avoiding calculation of the number of records returned

When executing an API call, use `countSwitch=1` attribute to avoid calculating number of reported rows. To return total number of the rows for an API query, use `countSwitch=2` and then retrieve pages of the records.

1. Retrieve total number of available records.

```
https://hostname:port/api/sam/raw_file_facts?token=token&countSwitch=2
```

2. Retrieve first 10000 records by using the **limit** parameter.

```
https://hostname:port/api/sam/raw_file_facts?token=token&countSwitch=1&limit=10000&offset=0
```

3. Retrieve next 10000 records. You can ignore the already retrieved records by using the **offset** parameter.

```
https://hostname:port/api/sam/raw_file_facts?token=token&countSwitch=1&limit=10000&offset=10000
```

Options to limit size of the individual response

Using limit and offset

The limit and offset parameters works as same as their representation in SQL/RDBMS. This allows to page over the result set. The first page, Tradeoff is available in the shortest time and other subsequent pages load slowly. It needs more data to load the subsequent pages and it cannot be optimized.

When you retrieve data in bulk, you can also make several API requests and use the limit and offset parameters to paginate your results instead of retrieving all the data at the same time:

- Use the limit parameter to specify the number of retrieved results:

```
https://hostname:port/URL?token=token&countSwitch=1&limit=10000&offset=0
```

For example:

```
https://192.0.2.2:9081/api/sam/v2/computers?token=token&countSwitch=1&limit=10000&offset=0
```

- If you limit the first request to 100000 results, append the next request with the `offset=100000` parameter to omit the records that you already retrieved:

```
https://hostname:port/URL?token=token&countSwitch=1&limit=10000&offset=10000
```

For example:

```
https://192.0.2.2:9081/api/sam/v2/computers?token=token&countSwitch=1&limit=10000&offset=10000
```



Note: The limit and offset parameters can be omitted if you are retrieving data for 50 endpoints. For environments with approximately 200000 endpoints, you are advised to retrieve data in pages of 100000 rows for computer systems, 200000 rows for software instances, and 300000 rows for license usage.

Using criteria

The alternative method depends on the API instance and is based on using criteria. It allows to optimize response time for subsequent pages as the application can extract information directly about the page, without preparing the whole data.

Limit and offset parameters require to have the results ready according to some order and then extract it partially. With this approach, the API, when requested, yields the entire result. The pages are reported at different time. However, size of response varies due to the data structure.

Software Instance Paging example using criteria

There `instance_id` field provides an unique value of the entry. It does not guarantee that there will be no gaps in `instance_id` number. First `instance_id` is 1.

To get maximum value of `instance_id`, use the below example:

```
https://hostname:port/api/sam/v2/software_instances?
columns[]=instance_id&order[]=instance_id%20desc&offset=0&limit=1&countSwitch=1
```

The return value is:

```
{"total":-1,"rows":[{"instance_id":162}]}
```

Specifying order (`order[]=instance_id desc`) assures that as first will be reported with maximum value, so we can get just top row (`limit=1`). Returned value is maximum value of `instance_id`.

You can page over the `software_instances` using `instance_id` and criteria.

Below is an example request: `https://hostname:port/api/sam/v2/software_instances?`

```
columns[]=computer_id&columns[]=instance_id&columns[]=component_name&offset=0&limit=10000&countSwitch=1&criteria={"and":
[["instance_id",">=", "1"],["instance_id ","<","10000"]]}
```

Limit and offset values are additional parameters but it is recommended to use them.

dcccliv

Control over paging is moved to the criteria part then:

```
criteria={"and":[  
  
  ["instance_id", ">=", "1"],  
  
  ["instance_id ", "<", "10000"]  
  
]}
```

Next page is returned by passing:

```
criteria={"and":[  
  
  ["instance_id", ">=", "10000"],  
  
  ["instance_id ", "<", "20000"]  
  
]}
```

This method leverage data indexes to extract immediately desired subset of data and should provide more stable response times over subsequent queries.

Synchronizing with Data Imports

The below available APIs provide status information of the data import:

1. https://hostname:port/api/import_status.json?token=token
2. https://help.hcltechsw.com/bigfix/9.5/inventory/Inventory/integration/r_get_import_status_bfi.html

Consider synchronizing API bulk calls with data import by:

- checking if there is no import in progress (mode).
- checking if the time of last successful data import (*last_success_time*) is later than in previous synchronization.

Understanding data changes

By default, raw data remains in the database for 7 days. To change this period, go to *Management > Advanced Server Settings* and change the value of the `raw_data_api_history_keep_days` parameter.

This applies to:

- `/api/sam/raw_package_facts.json`
- `/api/sam/raw_unix_package_facts.json`
- `/api/sam/raw_file_facts`

The APIs enables the below fields to be displayed:

Table 83.

valid_from	The date when you imported information about the file for the first time.	✓	String
valid_to	The date when you attempted to import information about the file for the first time. The value 9999-12-31T23:59:59Z indicates that the file is still being discovered.	✓	String

See below an example for filtering criteria for a request:

```
{ "or": [
  { "and": [
    ["valid_from", ">", "${import_time_previous}"],
    ["valid_from", "<=", "${import_time_last}"]
  ]}
,
  { "and": [
    ["valid_to", ">", "${import_time_previous}"],
    ["valid_to", "<=", "${import_time_last}"]
  ]}
]}
```

Where:

- `import_time_previous` – Previous successful data import time.
- `import_time_last` – Last successful data import time during initial load.

This API provides information about the hardware associated to the application.

Table 84.

first_seen	Date and time when the computer were entered in the application. The time is shown in GMT.		String
is_deleted	Information whether the computer was removed.	✓	Boolean

Table 84. (continued)

<code>deletion_date</code>	Date and time when the computer was removed. The time is shown in GMT.		String
----------------------------	--	--	--------

The API does not provide details of the computer.

This API shows information when the hardware details were updated. It shows details of the new or latest updated hardwares.

Table 85.

<code>updated_at</code>	Date and time when the current information was imported to BigFix Inventory.		String
-------------------------	--	--	--------

This API shows information related to software associated to the application. The following information are applicable for the software.

Table 86.

<code>is_present</code>	Information whether the component is still installed in the environment.		Boolean
<code>discovery_start</code>	Date and time when the component was reported for the first time. The time is shown in GMT.	✓	String
<code>discovery_end</code>	Date and time when the component was reported for the last time. The time is shown in GMT.	✓	String
<code>is_suppressed</code>	Information whether the component is suppressed on the computer on which it is installed.	✓	Boolean

Integrating

External systems integration is one of the key features of BigFix Inventory. Business logic is enabled for integration and interfaces are provided for common integration points.

 **Restriction:** The database cannot be modified. The data can be retrieved only through REST API, not directly from the database.

REST API

Managing large amounts of information by using the application user interface can be time-consuming. You can use the REST APIs as an alternative to the graphical user interface to reduce the time that is needed to manage your software inventory and the content of your software catalog.

 **Important:** If you want to use REST API when you are not logged in to BigFix Inventory, you must provide the **token** parameter for each query. To check the token, log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.

The **token** parameter is not required when you are logged in to BigFix Inventory.

Although using single API requests to retrieve data only from a selected subset of computers does not greatly impact the performance of BigFix Inventory, this is not true when retrieving data in bulk for all your computer systems at the same time. To improve the performance of a REST API, retrieve the data in chunks. For more information, see: [REST API considerations \(on page dccccli\)](#).

 **Restriction:** Do not use REST API during the import as you might encounter a problem with the database connection. If the problem occurs, wait until the import finishes and ensure that the BigFix Inventory server is up and running. Then, repeat the API call.

Abbreviations

The following abbreviations are used in the documentation that is related to REST API:

HTTP

Hyper Text Transfer Protocol. HTTP version 1.1 is defined in RFC 2616. Unless otherwise noted, the use of the term HTTP indicates both HTTP and HTTPS.

HTTPS

Hyper Text Transfer Protocol Secure, as defined in RFC 2818

JSON

JavaScript™ Object Notation, as defined in ECMA-262

REST

Representational State Transfer, as originally and informally described in *Architectural Styles and the Design of Network-based Software Architectures*

URI

Uniform Resource Identifier, as defined in RFC 3986, Draft 5

Software Products and Software Components

BigFix Inventory detects software components which are bundled with the product.

The list shows the available software component. You can add a software component to Inventory as required. Below given is an example to BigFix Platform Agent which is installed at `/opt/BESClient/bin/`.

<input type="checkbox"/>	Publisher Name	Component Name	Component Ver...	Product Name	Metric	Installation Path	Process Na...	Total Runs	Total Run T...	First Used	Details
<input type="checkbox"/>	IBM	IBM BigFix Platform Agent	9.5	IBM BigFix Inventory	Client Device	/opt/BESClient/bin/swidtag	besclient	3	20 days, 21...	09/02/2019 ...	DETAILS >

Use the Software Component listing to see the number of active and inactive components.

Each row shows details, such as usage, product and metric that it is added to, the installation path about the selected component. You can select multiple components.

Publishe...	Component Name	Comp...	Product Name	Metric	Installation Path	Process ...	Total Runs	Total Ru...	First Used	Last Used	Average ...	Average ...	Has Usage	D
IBM	IBM DB2 Workgroup Server E...	10.5	IBM BigFix Inventory	Client De...	/opt/ibm/db2/V10.5/properties...	<no data>	<no data>	<no data>	<no data>	<no data>	<no data>	<no data>	No	D
IBM	IBM BigFix Platform Agent	9.5	IBM BigFix Inventory	Client De...	/opt/BESClient/bin/swidtag	besclient	3	20 days, ...	09/02/20...	09/23/20...	0.141	6 days, 2...	Yes	D
IBM	IBM BigFix Inventory Server	9.2	IBM BigFix Inventory	Client De...	/opt/ibm/BFI/iso-swid	<no data>	<no data>	<no data>	<no data>	<no data>	<no data>	<no data>	No	D

You can access information about computer properties through *Computer Properties (Management > Computers)*.

You can access the computer properties through the user interface or the API.



Note: The API shows the computer properties by its number and ID and not by its name.

Table 87. API

<code>computer_property_number</code>	Custom computer property that was added on the Computer Properties panel. To view the list of all custom properties, view the computer_details schema . The name of the property as displayed on the user interface is shown in the <code>title</code> parameter.
---------------------------------------	--

You can use this feature to access user information, such as username from the system. It provides the basic user information which is already available at the Platform level. The user information may vary for each application.

Create Computer Property

Name*

Link to Data Source

Data Source Property*

User Name

ActionSite

User Name

USER NAME

REST API resources and HTTP methods

The operations of the BigFix Inventory REST API protocol are defined as HTTP methods on certain REST resources.

Table 88. Overview of REST operations

The table consists of three columns and 26 rows.

Target REST operation URI	HTTP methods	Purpose of the operation
api/get_token	POST (on page dccclxxi)	Returns a unique token that is required to authenticate REST API requests.
api/import_status.json	GET (on page dccclxxviii)	Returns the current status of the data import.
api/reports	POST (on page mciii)	Imports a saved report view.
api/reports/report_ID	GET (on page mc)	Exports a saved report view.
api/sam/about	GET (on page dccclxxv)	Returns the version of BigFix Inventory.
api/sam/clusters	GET (on page cmlxxiv)	Returns information about clusters that group your host computer systems.
api/sam/computer_systems	GET (on page cmlxvii)	Returns a list of computer systems in your infrastructure.
api/sam/configs	GET (on page cmlvi)	Returns information about the current settings of the administration server.
	PUT (on page cmlix)	Changes the current settings of the administration server.
api/sam/contracts/import	POST (on page mcvi)	Imports contracts from a CSV file.

Table 88. Overview of REST operations

The table consists of three columns and 26 rows.

(continued)

Target REST operation URI	HTTP methods	Purpose of the operation
<code>api/sam/license_usages</code>	GET (on page cmlxxxvi)	Returns information about license usage reported by your computer systems.
<code>api/sam/raw_app_usage_property_values</code>	GET (on page dcccxcix)	Retrieves raw metering data.
<code>api/sam/raw_file_facts</code>	GET (on page dcccxxxiii)	Retrieves raw scanned file data.
<code>api/sam/raw_package_facts</code>	GET (on page dcccxxxix)	Retrieves raw package data.
<code>api/sam/raw_unix_package_facts</code>	GET (on page dcccxciv)	Retrieves raw UNIX package data.
<code>api/sam/software_instances</code>	GET (on page cmlxxix)	Returns a list of installed software instances.
<code>api/sam/swinventory/confirm</code>	POST (on page cmxliv)	Confirms instances to bundle or assign.
<code>api/sam/swinventory/confirmRelease</code>	POST (on page cmxlv)	Confirms the assignment of instances for a release.
<code>api/sam/swinventory/exclude</code>	POST (on page cmxxxv)	Excludes instances from pricing calculations.
<code>api/sam/swinventory/include</code>	POST (on page cmxxxiii)	Includes instances in pricing calculations.
<code>api/sam/swinventory/instanceToShare</code>	GET (on page cmxxvii)	Returns a list of software releases that can share a particular instance.
<code>api/sam/swinventory/product/{product_id}/releases</code>	GET (on page cmix)	Returns a list of releases of a particular software product that is identified by its identifier.
<code>api/sam/swinventory/products</code>	GET (on page cmiv)	Returns a list of software products.
<code>api/sam/swinventory/reassign</code>	POST (on page cmxxxviii)	Reassigns instances to a product.
<code>api/sam/swinventory/reassignRelease</code>	POST (on page cmxli)	Reassigns instances to a product for a list of releases.

Table 88. Overview of REST operations

The table consists of three columns and 26 rows.

(continued)

Target REST operation URI	HTTP methods	Purpose of the operation
<code>api/sam/swinventory/release/{release_id}/instances</code>	GET (on page cmxiii)	Returns a list of instances for the requested software product release.
<code>api/sam/swinventory/share</code>	POST (on page cmxxx)	Shares an instance with a list of software products.
<code>api/sam/swinventory/targetBundlesOfInstances</code>	GET (on page cmxvii)	Returns a list of possible releases to which the requested software instance can be reassigned.
<code>api/sam/swinventory/targetBundlesOfReleases</code>	GET (on page cmxx)	Returns a list of possible releases to which the instances of the requested software releases can be reassigned.
<code>api/sam/swinventory/targetInstances</code>	GET (on page cmxxiv)	Returns a list of target instances that will be reassigned to a particular release.
9.2.8 <code>api/sam/v2/computers</code>	GET (on page mxxx)	Returns information about hardware inventory in your infrastructure.
9.2.10 <code>api/sam/v2/license_usage</code>	GET (on page mix)	Returns information about utilization of license metrics by products that are installed in your infrastructure.
	PUT (on page mxix)	Sets up license metric thresholds as well as values of custom fields that were added on the All Metrics report.
9.2.8 <code>api/sam/v2/software_instances</code>	GET (on page cmxcii)	Returns information about software inventory in your infrastructure.
<code>api/sam/vmmanagers</code>	GET (on page cmxlix)	Returns information about VM managers that are defined in your infrastructure.
	PUT (on page cmlii)	Adds or updates VM managers.
9.2.8 <code>api/v1/servers</code>	PUT (on page mcxii)	Changes Oracle core factor that is assigned to a server.
9.2.11 <code>sam/component_lifecycles</code>	POST (on page mxc)	Sets the end of support date for software components.

Table 88. Overview of REST operations

The table consists of three columns and 26 rows.

(continued)

Target REST operation URI	HTTP methods	Purpose of the operation
9.2.11 <code>sam/software_components/software_ component_id/component_lifecycles</code>	DELETE (on page mxciii)	Removes the end of support date for a software component.
9.2.12 <code>api/sam/v2/detailed_hw_ip_address- es</code>	GET (on page mxliv)	Retrieves information related to IP addresses of computers in your infrastructure.
9.2.12 <code>api/sam/v2/detailed_hw_lpars</code>	GET (on page mxlix)	Retrieves information related to the logical partition and its CPU resources.
9.2.12 <code>api/sam/v2/detailed_hw_memories</code>	GET (on page mliv)	Retrieves information related to information related to the system memory of the scanned system.
9.2.12 <code>api/sam/v2/detailed_hw_network_ adapters</code>	GET (on page mlix)	Retrieves information related to information related to the type, model and MAC addresses of the network adapters.
9.2.12 <code>api/sam/v2/detailed_hw_operating_ systems</code>	GET (on page mlix)	Retrieves information related to information related to the operating system on the scanned computer.
9.2.12 <code>api/sam/v2/detailed_hw_partitions</code>	GET (on page mlxix)	Retrieves information related to information related to the file system that is mounted on a scanned system: Unix mount points and Windows logical drives.
9.2.12 <code>api/sam/v2/detailed_hw_physical_ processors</code>	GET (on page mlxxiv)	Retrieves information related to information related to all active physical processors.
9.2.12 <code>api/sam/v2/detailed_hw_smbios</code>	GET (on page mlxxx)	Retrieves information related to information related to SMBIOS of the scanned system.
9.2.12 <code>api/sam/v2/detailed_hw_storages</code>	GET (on page mlxxv)	Retrieves information related to information related to storage devices on the scanned system.

Common REST API elements

You use connectors and operators to build query segments that filter on matching conditions and values.

Connectors

Table 89. Connectors

Connector	Description
and	Add all clauses with an and concatenator

Table 89. Connectors (continued)

Connector	Description
<code>or</code>	Add all clauses with an <code>or</code> concatenator

Parameters

Table 90. Parameters

Parameter	Description
<code>countSwitch</code>	<p>Specifies whether the query returns the requested data, the number of retrieved rows, or both. The operator can have the following values:</p> <ul style="list-style-type: none"> • 1 - only requested data is returned • 2 - only the number of rows is returned • No value or any other value - both the requested data and the number of rows are returned

Operators

Table 91. Operators

Operator	Description	Applicable to
<code>=</code>	Equal to.	
<code>!=</code>	Not equal to.	
<code><</code>	Less than.	Numeric and date fields only.
<code><=</code>	Less than or equal to.	Numeric and date fields only.
<code>></code>	Greater than.	Numeric and date fields only.
<code>>=</code>	Greater than or equal to.	Numeric and date fields only.
<code>last</code>	Relative period of time in the past (for example, last 7 days).	Date and time fields only.
<code>next</code>	Relative period of time in the future (for example, next 7 days).	Date and time fields only.
<code>starting</code>	Relative period of time starting in the past or in the future (for example, 7 days ago, or 7 days ahead) and lasting until infinity. For syntax, see Table 92: Relative date and time operators (on page dccclxv) .	Date and time fields only.

Table 91. Operators (continued)

Operator	Description	Applicable to
ending	Relative period of time ending in the past or in the future (for example, 7 days ago, or 7 days ahead) and including all data until this point. For syntax, see Table 92: Relative date and time operators (on page dccclxv) .	Date and time fields only.
begins_with	String begins with the specified value.	String fields only.
not_begins_with	String does not begin with the specified value.	String fields only.
contains	String contains the specified value.	String fields only.
not_contains	String does not contain the specified value.	String fields only.
ends_with	String ends with the specified value.	String fields only.
not_ends_with	String does not end with the specified value.	String fields only.
9.2.13 is_empty	Column does not contain any value. Example:	Selected columns only.
	<ul style="list-style-type: none"> <pre>https://hostname:port/api/sam/v2/software_instances?token=token&criteria={"and":[{"exclusion_or_suppress_comment","is_empty"]}}</pre> 	
9.2.13 is_not_empty	Column contains a value. Example:	Selected columns only.
	<ul style="list-style-type: none"> <pre>https://hostname:port/api/sam/v2/software_instances?token=token&criteria={"and":[{"exclusion_or_suppress_comment","is_not_empty"]}}</pre> 	

Syntax of relative date and time operators

These operators can be used for all columns that use the date and time format.

Table 92. Relative date and time operators

Operator	Syntax
last	Px[D W M Y]
next	<p>where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively. For example:</p> <ul style="list-style-type: none"> Retrieve data from computers that reported within last 7 days: <pre>https://hostname:port/api/sam/computer_systems?token=token &criteria={"and":[{"last_seen","last","P7D"}]}</pre> Retrieve contracts whose entitlement ends within next 7 days: <pre>https://hostname:port/api/sam/contracts?token=token &criteria={"and":[{"entitlement_end","next","P7D"}]}</pre>
starting	[- +]Px[D W M Y]
ending	<p>where x is a number in the 1-999 range, -/+ stands for ago or ahead, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively. For example:</p> <ul style="list-style-type: none"> Retrieve contracts whose maintenance starts in a period of time starting 1 week ahead. This API retrieves only future contracts. <pre>https://hostname:port/api/sam/contracts?token=token &criteria={"and":[{"maintenance_start","starting","+P1W"}]}</pre> Retrieve contracts whose entitlement ends in a period of time ending 1 day ago. This API retrieves all possible contracts from the past until a day ago: <pre>https://hostname:port/api/sam/contracts?token=token &criteria={"and":[{"entitlement_end","ending","-P1D"}]}</pre>

Response elements for GET requests

Table 93. Response elements for GET requests

Element	Description
total	Number of all records that meet query parameter. The number of returned records might be smaller, for example if you use the limit parameter.
rows	Records that are returned by the query.

Mapping REST API versions

9.2.8 Available from 9.2.8. For some REST APIs, version 2 is created to provide improved functions or experience. In such cases, the original API becomes deprecated with time and only version 2 is further developed. If you have custom tools or internal processes that are based on such REST APIs, it is best to adjust the tools to use the new version of the APIs.

Mapping columns between **software_instances** and **v2/software_instances** REST API

9.2.8 Available from 9.2.8. If you have custom tools or internal processes that use the **software_instances** REST API, learn how columns in this API map to the columns in the **v2/software_instances** API. It is best to adjust the tools to use REST API in version 2.

The amount of data returned by the **v2/software_instances** API might be greater than the amount of data returned by the **software_instances** API. It happens because the **v2/software_instances** API returns historical data about uninstalled software and removed computers. It also returns multiple records if software is installed in multiple locations on the same computer. The **software_instances** API does not return this data.



Note: If a column from the **computer_systems** API that you use does not map to any column in the **v2/computers** API or the mapping provides insufficient granularity of information, [submit an idea](#).

Table 94. Mapping software_instances to v2/software_instances REST API

software_instances API	v2/software_instances API	Comment
id		
software_fact_id		
computer_system_id		Computer systems are substituted with computers.
computer_id	computer_id	
discoverable_guid	discoverable_guid	
default_product_guid		The default_product_guid column provides information about the default bundling. REST API version 2, provides information about the product to which the component is in fact assigned. The

Table 94. Mapping software_instances to v2/software_instances REST API

(continued)

software_instances API	v2/software_instances API	Comment
		information is provided in the product_name column.
first_used	usage_data.first_used	
last_used	usage_data.last_used	
valid_from	discovery_start	
valid_to	discovery_end	The valid_to column always shows <code>9999-12-31T23:59:59Z</code> because the software_instances REST API does not provide historical data. The discovery_end column provides the actual date when the component was reported for the last time.
updated_at		
signature_count		
total_time	usage_data.total_time	
total_runs	usage_data.total_runs	
avg_run_time	usage_data.avg_run_time	
avg_runs_per_day	usage_data.avg_runs_per_day	
process	usage_data.process	
deleted	is_present	The deleted column always shows <code>false</code> . The is_present column provides the actual information.
catalog_dimension.software_title_name	product_name	The catalog_dimension columns provide information about the default product.
catalog_dimension.publisher_name	product_publisher_name	
catalog_dimension.software_title_release_name	product_release_name	REST API version 2, provides information about the product to which the component is in fact assigned.
catalog_dimension.version	product_release	

Table 94. Mapping software_instances to v2/software_instances REST API

(continued)

software_instances API	v2/software_instances API	Comment
catalog_dimension.software_title_- version_name		

Mapping columns between **computer_systems** and **v2/computers** REST API

9.2.8 Available from 9.2.8. If you have custom tools or internal processes that use the **computer_systems** REST API, learn how columns in this API map to the columns in the **v2/computers** API. It is best to adjust the tools to use REST API in version 2.



Note: If a column from the **computer_systems** API that you use does not map to any column in the **v2/computers** API or the mapping provides insufficient granularity of information, [submit an idea](#).

Table 95. Mapping computer_systems to v2/computers REST API

computer_systems API	computers API	Comments
id		Computer systems are substituted with computers.
parent_id		Computer systems are substituted with computers. Additionally, REST API version 2 is not hierarchical.
computer_id	id	
computer_remote_id	bigfix_id	
server_id	computer_hardware.server_id	
datasource_name		
datasource_id		
type	computer_hardware.computer_type	REST API version 2 has three possible values: physical, virtual, and computer running on a public cloud.
os	os	
host_name	name	
dns_name	dns_name	
ip_address	ip_address	
last_seen	last_seen	

Table 95. Mapping computer_systems to v2/computers REST API (continued)

computer_systems API	computers API	Comments
hardware_manufacturer		The information is contained in the computer_hardware.server_name column.
hardware_model		The information is contained in the computer_hardware.server_name column.
hardware_serial_number		The information is contained in the computer_hardware.server_name column.
hardware_type		The information is contained in the computer_hardware.server_name column.
hardware_name	computer_hardware.server_name	
processor_brand_string	computer_hardware.processor_brand_string	
processor_type	computer_hardware.processor_type	
processor_brand	computer_hardware.processor_brand	
processor_vendor	computer_hardware.processor_vendor	
partition_cores	computer_hardware.partition_cores	
server_processors	computer_hardware.node_total_processors	
server_cores	computer_hardware.server_cores	
pvu_per_core	computer_hardware.pvu_per_core	
uuid		
cluster_id		Cluster ID is substituted with cluster name that is available in the computer_hardware.cluster_name column.

Mapping functions of the **software_instances** REST API

9.2.14 Available from 9.2.14. If you have custom tools or internal processes that use the `software_instances` REST API, learn how to achieve your goals by using the `v2/license_usage` and `v2/software_instances` REST API. It is best to adjust the tools to use REST API in version 2.

Mapping `software_instances` REST API functions

The `software_instances` REST API allows you to retrieve information about the software instance that contributed to the peak in the utilization of a license metric by a product during the reporting period. By default, the reporting period is set to 90 days.

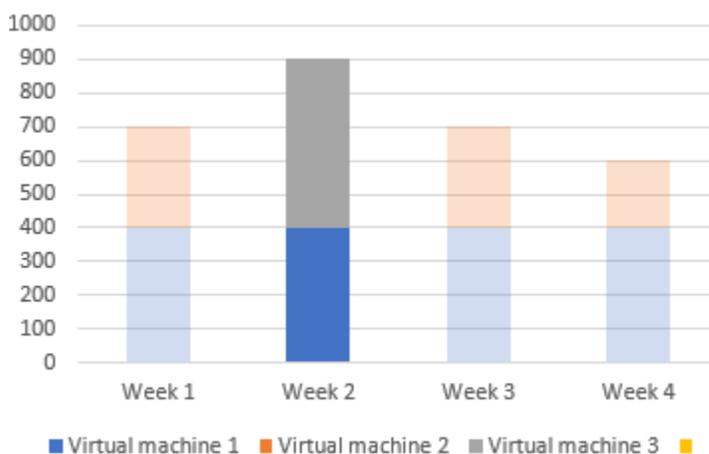
The `v2/software_instances` REST API allows you to retrieve information about all software instances in your environment. Additionally, you can retrieve data from any relevant period, including the details of existing components as well as historical data about components that were removed. By filtering the data you can narrow down the results to the instances that contributed to the peak in the utilization of a license metric by a product. Combining the capabilities of `v2/license_usage` and `v2/software_instances` APIs can not only substitute all functions of the old version of `software_instances` API, but also guarantees flexibility that allows you to alter the queries and retrieve the required data sample.

Examples

Scenario: Retrieving information about the software instances that contributed to the peak in the utilization of a license metric by a product

- `software_instances` **API** allows you to only retrieve information about the software instances that contributed to the peak of a license metric of a product over a reporting period.

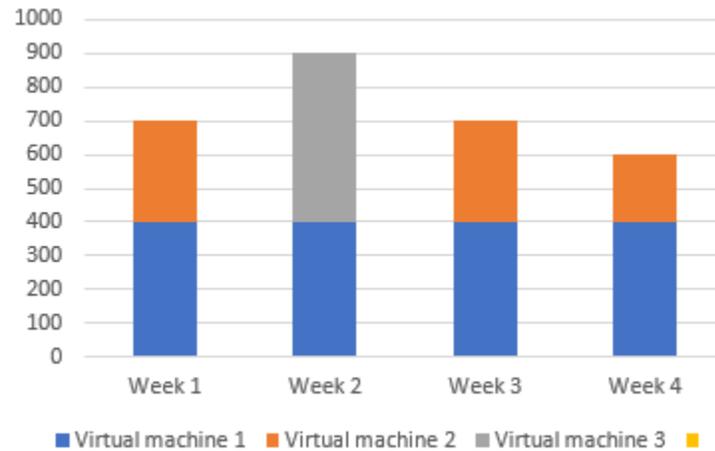
Figure 8. Software instances that can be retrieved with the `software_instances` REST API



- `v2/license_usage` and `v2/software_instances` **APIs** allow you to retrieve the following information:

- The list of all software instances that were used within a specified period

Figure 9. Software instances that can be retrieved with the `v2/software_instances` REST API



- The time of the peak in the utilization of a license metric by a product with the `hwm_peak_time` parameter value
- The list of software instances that contributed to the peak
- The overall utilization of a license metric of a product during the peak

To retrieve information about the software instances that contributed to the peak in the utilization of a license metric by a product proceed with the following steps.

1. Run the `v2/license_usage` REST API for a product and the license metric that is assigned to this product.

Example

```
GET
api/sam/v2/license_usage?columns[]=product_name&columns[]=metric_code_name&columns[]=hwm_quantity&
columns[]=hwm_peak_time&criteria={"and":[{"product_name","=","IBM DB2 Advanced Enterprise Server Edition PVU Option"]}]
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

2. See the value of the `hwm_peak_time` column for the product-license pair in the response.

Example

```
[{
"product_name": "IBM DB2 Advanced Enterprise Server Edition PVU Option",
```

```
"metric_code_name": "PVU_SUB_CAP",
"hwm_quantity": 960,
"hwm_peak_time": "2018-10-11T01:12:26Z"
}]
```

3. Run the `v2/software_instances` REST API with the filter on the following columns: `Discovery Start` and `Discovery End`. Base the dates on the value of `hwm_peak_time`.

Example

```
GET
api/sam/v2/software_instances?columns[]=discovery_start&columns[]=discovery_end&columns[]=product_name&columns[]=metric_id&columns[]=computer_name&criteria={"and":[{"product_name","=", "IBM DB2 Advanced Enterprise Server Edition PVU Option"}, {"metric_id","=", "5"}, {"discovery_start", "<=", "2018-10-11T01:12:26Z"}, {"discovery_end", ">", "2018-10-11T01:12:26Z"}]}&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```



Note: The example uses the timestamp format used for PVU, RVP MPAC and VPC metrics. For different metrics, see: [Retrieval of software inventory \(v2\) \(on page cmxcii\)](#).

4. See the results for a list of software instances that contributed to the peak.

REST API for retrieving authentication token

You use the `POST` operation on the `api/get_token` element to request your unique token that is required to authenticate the REST API requests.

To request your unique token, use the following URL:

```
https://hostname:port/api/get_token
```



Important: This REST API call does not work when single sign-on authentication is enabled. To view the token, log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.

Table 96. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	POST /api/get_token
Purpose	Returns the authentication token
HTTP method	POST
Resource URI	<code>https://server_host_name:port_number/api/get_token</code>
URL link relation	n/a
URI query parameters	n/a
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	<pre>{ "user" : "username", "password" : "password" }</pre>
Request Content-Type	application/json
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p>

Table 96. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
	Specifies the language of the response content. If this header is not specified, the content is returned in the server language.
Response payload	Token element
Response Content-Type	application/json
Normal HTTP response codes	200 - OK
Error HTTP response codes	500 - "Bad Request" if a query parameter contains errors or is missing
	Message body includes an error message with details.

Example HTTP conversation

Request

```
POST api/get_token
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Request header

```
Content-Type: application/json
```

Request payload

```
{
  "user" : "admin",
  "password" : "password"
}
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

Response body (JSON)

```
{
  "token": "44072fb20fbe38322b5e67a7e780978e20abbc80"
}
```

REST API for retrieving the BigFix Inventory version

You use the `GET` operation on the `api/sam/about` element to request information about the version of BigFix Inventory.

To retrieve information about the current version of BigFix Inventory, use the following URL:

```
https://hostname:port/api/sam/about?token=token
```

**Important:**

- To use this API, you must be assigned to the All Computers group and have the View Hardware Inventory permission.
- Each API request must be authenticated with the **token** parameter. You can retrieve it by using [REST API for retrieving authentication token \(on page dcccclxxii\)](#). You can also log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.
- By default, the retrieved data is sorted by `id`.

Table 97. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/sam/about</code>
Purpose	Returns the version of BigFix Inventory
HTTP method	<code>GET</code>
Resource URI	<code>https://server_host_name:port_number/api/sam/about</code>
URL link relation	n/a
URI query parameters	n/a
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p>

Table 97. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Request payload	Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.
Request <code>Content-Type</code>	n/a
Response headers	<code>application/json</code>
Response headers	Header
Response headers	<code>Content-Type</code>
Response headers	Values
Response headers	<code>application/json</code>
Response headers	Specifies the content type of the response.
Response headers	Header
Response headers	<code>Content-Language</code>
Response headers	Values
Response headers	en-US, ...
Response headers	Specifies the language of the response content. If this header is not specified, the content is returned in the server language.
Response payload	<code>About</code> element
Response <code>Content-Type</code>	<code>application/json</code>
Normal HTTP response codes	200 - OK
Error HTTP response codes	500 - "Bad Request" if a query parameter contains errors or is missing
Error HTTP response codes	Message body includes an error message with details.

Example HTTP conversation

Request

```
GET /api/sam/about
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

Response body (JSON)

```
{
  "product": "BigFix Inventory",
  "version": "9.2.20140331-2015"
}
```

REST API for retrieving raw scan results

9.2.2 Available from 9.2.2. Retrieving raw scan results through REST API is intended to improve the monitoring of your environment by allowing you to quickly discover software that was installed or modified recently. By using shortened data imports, data about your software reaches the server much quicker, and at the same time allows you to increase the number of imports that you can run each day.

A data import consists of several steps, out of which the most important ones are the retrieval of raw scan results, aggregation of this data, and finally calculation of your license consumption. In big environments, a complete data import can take up to several hours, which increases the time that is required for BigFix Inventory to show information about recent changes.

When using REST API to retrieve raw scan results, you can run a Raw Data Only type of data import that retrieves raw scan results, but omits later steps, such as aggregation and license calculations. This data import is much shorter, but the imported raw data can only be retrieved through REST API. Since the data is not aggregated, it is not displayed in the BigFix Inventory UI. The UI still shows the state of your environment from the last successful complete data import.

Before you start retrieving raw scan results, ensure that you enable the storing of raw scan results, because not all such data is stored by default. You can then run a Raw Data Only import and retrieve the results through REST API.

Running data imports

9.2.2 Available from 9.2.2. You can use this REST API to check the status of a data import, or to run a *Raw Data Only*, or a *Complete* data import. All these actions are completed through a common API request. You distinguish between them by appending the request with specific parameters. To work with this API, you use the `GET` operation on the `api/import_status.json` element, or its alias `api/management/imports.json`.

Before you begin



You must have the Manage Imports permission to perform this task.

Checking the import status

Before you start running data imports, retrieve the status of a data import to ensure that no imports are currently in progress. The statuses that can be retrieved are explained in [Example HTTP conversation \(on page dcccclxxxii\)](#). To check the import status, use the following request:

```
https://hostname:port/api/import_status.json?token=token
```

Running a Raw Data Only import

Running a *Raw Data Only* import is much quicker, because it only retrieves raw scan results. This type of data import is used for retrieving raw scan results through REST API. A *Raw Data Only* data import includes the following steps of the import process:

- Initialization of models:

```
Calling Model.before_snapshot
```

- Initialization of each data source:

```
Initialize datasource #{datasource.name}
```

- Computer data:

```
RawComputerId
Computer
ComputerPropertyValue
ComputerDimension
```

- Computer-related files:

```
RawDatasourceFile
DatasourceFile
SAM::ScanFile
```

- Scanned file, package, UNIX package, and application usage data:

```
SAM::FileFactDelta
SAM::FileFact
```

```
SAM::RawPackageFact
SAM::PackageFact
SAM::RawUnixPackageFact
SAM::UnixPackageFact
SAM::RawAppUsagePropertyValue
SAM::AppUsagePropertyValue
```

To start importing your raw data, append the API request with the **run** parameter, and specify the value as *raw_only*:

```
https://hostname:port/api/import_status.json?token=token&run=raw_only
```



Note: After running a *Raw Data Only* data import, the raw data can be retrieved through REST API, but is not visible in the UI. The UI still shows the state of your environment from the last successful data import. The only exception is the Metering Data report that shows the most recent data even after a raw only data import. For more details, see [Importing raw scan data](#)

Running a complete data import

Running a *Complete* import through REST API has the same effect as running the data import from the UI. To start importing your data, append the API request with the **run** parameter, and specify the value as *Complete*:

```
https://hostname:port/api/import_status.json?token=token&run=complete
```

Operation descriptions



Note: Each API request must be authenticated with the token parameter. You can retrieve it by using [REST API for retrieving authentication token \(on page dccclxxii\)](#). You can also log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.

Table 98. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	GET /api/import_status.json, OR GET /api/management/imports.json
Purpose	Returns the status of the data import
HTTP method	GET
Resource URI	https://server_host_name:port_number/api/import_status.json , or https://server_host_name:port_number/api/management/imports.json

Table 98. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
URL link relation	n/a
URI query parameters	n/a
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p>
Response payload	<p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...;</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response <code>Content-Type</code>	<code>Import status</code> element
Response <code>Content-Type</code>	<code>application/json</code>

Table 98. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Normal HTTP response codes	
	200 - OK
Error HTTP response codes	
	500 - "Bad Request" if a query parameter contains errors or is missing
	Message body includes an error message with details.

Query parameters

You can use query parameters to narrow down the results of your search or perform extra actions. The following table presents query parameters that you can use for the `api/import_status.json` and `api/management/imports.json` elements.

Table 99. Query parameters for running a data import and retrieving its status

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
format	Specifies the format of the retrieved information. The possible values are <code>json</code> or <code>xml</code> . Retrieve information in the <code>json</code> format: <pre>URL?format=json</pre>	No	String
token	A unique user authentication identifier. You can view your token in the Profile preferences of BigFix Inventory.	Yes	Alphanumeric
run	Specifies the type of a data import to run. The possible values are: <ul style="list-style-type: none"> <code>complete</code> - runs a Complete data import <code>raw_only</code> - runs a Raw Data Only import that retrieves raw scan results. The data is not 	No	String

Table 99. Query parameters for running a data import and retrieving its status

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	aggregated and therefore not shown in the UI. You can only retrieve it through REST API.		

Example HTTP conversation

Request

```
GET api/import_status.json
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

Response body (JSON)

If a data import has never been initiated:

```
{
  "mode": "none",
  "can_run_imports": true,
  "import_status_url": "/import_status",
  "import_create_url": "/management/imports",
}
```

If a data import is in progress:

```
{
  "mode": "running",
  "can_run_imports": true,
  "import_status_url": "/import_status",
  "import_create_url": "/management/imports",
  "progress": 59
}
```

If a data import is not running:

```
{
  "mode": "idle"/"pending",
  "can_run_imports": true,
  "import_status_url": "/import_status",
  "import_create_url": "/management/imports",
  "last_status": successful,
  "last_success_time": "2015-06-18T04:00:29Z"
  "last_import_time": "2015-06-18T04:00:29Z",
  "last_import_type": "raw_only"
  "next_import_time": "2015-06-20T06:00:00Z",
  "next_import_time_in": 75057,
  "next_import_type": "complete"
}
```

Where:

- `mode` - status of the data import, it can assume the following values:
 - `none` - a data import has never been initiated
 - `idle` - no data imports are currently running
 - `running` - a data import is in progress
 - `pending` - an action performed in the user interface requires a data import to be started for the change to take effect
- `can_run_imports` - specifies whether the user retrieving this API can run imports,
- `import_status_url` - web address of this API,
- `import_create_url` - web address of the data imports panel in BigFix Inventory,
- `progress` - percentage status of the running import,
- `last_status` - status of the last import,
- `last_success_time` - time of the last successful import,
- `last_import_time` - time of the last import,
- `last_import_type` - type of the last import,
- `next_import_time` - time of the next scheduled data import,
- `next_import_time_in` - remaining time (in seconds) to the next scheduled data import,
- `next_import_type` - type of the next scheduled data import.

Retrieving raw scanned file data

9.2.2 Available from 9.2.2. You use the `GET` operation on the `/api/sam/raw_file_facts` element to retrieve information about all files discovered on the computers in your infrastructure. The API returns details of existing files as well as historical data about files that were removed. By default, raw data is preserved for 7 days.

To change this period, go to [Management > Advanced Server Settings](#) and change the value of the `raw_data_api_history_keep_days` parameter.

Note



To obtain information about detected software use `api/sam/v2/software_instances` API.

Prerequisites

- To improve performance, retrieve raw scanned file data in chunks of 10 000 records.

- Retrieve the total number of available records.

```
https://hostname:port/api/sam/raw_file_facts?token=token&countSwitch=2
```

- Retrieve the first 10 000 records by using the **limit** parameter.

```
https://hostname:port/api/sam/raw_file_facts?token=token&countSwitch=1&limit=10000&offset=0
```

- Retrieve the next 10 000 records. You omit the already retrieved records by using the **offset** parameter.

```
https://hostname:port/api/sam/raw_file_facts?token=token&countSwitch=1&limit=10000&offset=10000
```

Permissions



You must have the View Raw Data permission to perform this task.

Resource URL

```
https://hostname:port/api/sam/raw_file_facts?token=token
```

Resource information

Table 100. Resource information

The table consists of two columns and 6 rows.

Operation details	Description
HTTP method	GET
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>

Table 100. Resource information

The table consists of two columns and 6 rows.

(continued)

Operation details	Description
Request format	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	<code>File Fact</code> element
Response format	<code>application/json</code>
Response codes	<p><code>200 - OK</code></p> <p><code>500 - "Bad Request"</code> if a query parameter contains errors or is missing</p>

9.2.10 Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/schemas/raw_file_fact.json?token=token
```

Available columns

Table 101. Columns with information about scanned files

Column	Description	Displayed by default	Type
computer_id	Identifier of the computer on which the file was discovered.	✓	String
path	Path where the file was discovered.	✓	String
name	Name of the discovered file.	✓	String
size	Size of the discovered file.	✓	Integer
version	Version of the discovered file if available.	✓	String
md5	MD5 checksum of the file.	✓	String
sha256	SHA256 checksum of the file.	✓	String
valid_from	Date of the first import that contained information about the file.	✓	String
valid_to	Date of the first import that did not include information about the file. The value <code>9999-12-31T23:59:59Z</code> indicates that the file is still being discovered.	✓	String

Applicable associations

You can additionally retrieve data from the following associations:

- [computer](#) (on page [mcxvi](#))

Query parameters

Table 102. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve the name and size of a file: <pre>URL?columns[]=name&columns[]=size</pre>		String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name.		String

Table 102. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<p>Example: Order by name</p> <pre>URL?order[]=name desc</pre>		
limit	<p>Specify the number of rows to retrieve. If you omit this parameter, 100 000 rows are retrieved. The number is defined by the raw_data_api_default_limit parameter on the Advanced Server Settings panel.</p> <p>Example: Retrieve 100 records</p> <pre>URL?limit=100</pre>		Numeric
offset	<p>Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results.</p> <p>Example: Retrieve 50 records starting after record 150:</p> <pre>URL?limit=50&offset=150</pre>		Numeric
token	<p>A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dccclxxii). You can also log in to BigFix Inventory, hover over the User icon , and click Profile. Then, click Show token.</p>	✓	Alphanumeric
criteria	<p>Retrieve records which match specific conditions. The parameter should have the following structure, written in one line:</p> <pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criterion> [{ <comma> <criterion> }...] <right-bracket> <right-brace></pre> <pre><boolean-operator> ::= "and" "or"</pre> <pre><criterion> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket></pre> <pre><column> ::= <json-string></pre> <pre><operator> ::= <json-string></pre> <pre><value> ::= <json-array> <json-string> <json-numver> <json-null></pre> <p>For more information about operators, see Common connectors and operators (on page dccclxii).</p> <p>Example: Retrieve scanned file data from computer systems with ID greater than 10:</p>		

Table 102. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<p><code>URL?criteria={ "and": [["computer_id", ">", "10"]] }</code></p> <p>For columns that use the date and time values, such as Last Seen, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <operator>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.</p> <p>Example: Retrieve computers that were first reported in BigFix Inventory within last 7 days</p>		
	<p><code>URL?criteria={ "and": [["valid_from", "last", "P7D"]] }</code></p>		

Example conversation - default columns

Request

```
GET api/sam/raw_file_facts
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
{
  "computer_id": 5,
  "path": "C:\BES\BESAirgapTool",
  "name": "BESAirgapTool.exe",
  "size": 92174,
  "version": null,
  "valid_from": "2015-07-31T07:03:21Z"
  "valid_to": "9999-12-31T23:59:59Z"
}
{
  "computer_id": 5,
  "path": "C:/Program Files/ibm/SQLLIB/BIN",
  "name": "db2set.exe",
```

```
"size": 81768,
"version": "10.1",
"valid_from": "2015-07-31T07:03:21Z",
"valid_to": "9999-12-31T23:59:59Z"
}
```

Retrieving raw package data

9.2.2 Available from 9.2.2. This API retrieves raw package data, which contains information about all packages discovered on your computers. This information is retrieved from the Windows registry. Unlike in the UI, through this API you can retrieve all historical data that is stored in the database.

To obtain information about detected software use `api/sam/v2/software_instances` API.

Before you begin

-  You must have the View Raw Data permission to perform this task.

To improve performance, retrieve raw package data in chunks of 10 000 records.

1. Retrieve the total number of available records.

```
https://hostname:port/api/sam/raw_package_facts?token=token&countSwitch=2
```

```
{"total":49999,"rows":[]}
```

2. Retrieve only the first 10 000 records by using the **limit** parameter.

```
https://hostname:port/api/sam/raw_package_facts?token=token&countSwitch=1&limit=10000&offset=0
```

3. Retrieve the next 10 000 records. You omit the already retrieved records by using the **offset** parameter.

```
https://hostname:port/api/sam/raw_package_facts?token=token&countSwitch=1&limit=10000&offset=10000
```

 **Note:** Each API request must be authenticated with the token parameter. You can retrieve it by using [REST API for retrieving authentication token \(on page dccclxxii\)](#). You can also log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.

Table 103. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	GET /api/sam/raw_package_facts

Table 103. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Purpose	Returns raw package data.
HTTP method	GET
Resource URI	<code>https://server_host_name:port_number/api/sam/raw_package_facts</code>
URL link relation	n/a
URI query parameters	n/a
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	<code>Package Fact</code> element

Table 103. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Response Content-Type	application/json
Normal HTTP response codes	200 - OK
Error HTTP response codes	500 - "Bad Request" if a query parameter contains errors or is missing Message body includes an error message with details.

Available columns

Table 104. Available columns

Column	Description	Displayed by default	Type
computer_id	Identifier of the computer as defined in BigFix Inventory.	✓	Nu-mer-ic
description	Description of the package.	✓	String
name	Name of the package.	✓	String
version	Version of the package.	✓	String
vendor	Vendor of the package.	✓	String
valid_from	Date and time when the package was reported for the first time. The time is specified in the GMT time zone.	✓	String
valid_to	Date and time when the package was reported for the last time. The time is specified in the GMT time zone.	✓	String

Query parameters

You can use query parameters to narrow down the results of your search.

Table 105. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, a set of default columns is retrieved. Example: Retrieve the name and computer_id columns: <pre>URL?columns[]=name&columns[]=computer_id</pre>	No	String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by name descending: <pre>URL?order[]=name desc</pre>	No	String
limit	Specify the number of rows to retrieve. If you omit this parameter, the limit is set according to the value of the <code>raw_data_api_default_limit</code> server setting, which by default is 100 000.	No	Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150: <pre>URL?limit=50&offset=150</pre>	No	Numeric
token	A unique user authentication identifier. You can view your token in the Profile preferences of BigFix Inventory.	Yes	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written on one line: <pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criterion> [{ <comma> <criterion> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criterion> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket></pre>		

Table 105. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<pre><column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null></pre>		
	<p>For more information about operators, see Common connectors and operators (on page dcccxcxii).</p> <p>Example: Retrieve package data from computer systems with ID greater than 10:</p> <pre>URL?criteria={ "and": [["computer_id", ">", "10"]] }</pre>		
	<p>For columns that use the date and time values, such as Last Seen, you can retrieve data also for a instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <i><operator></i>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively. For example, to retrieve computer systems that reported within last 7 days, use the following API request:</p> <pre>URL?criteria={ "and": [["last_seen", "last", "P7D"]] }</pre>		

Example HTTP conversation

Request

```
GET /api/sam/raw_package_facts
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

Response body (JSON)

```
{
  "package_property_id": 1,
  "computer_id": 3,
```

```

"description": "{FA7394B8-CE65-4F9E-AC99-F372AD365424}",
"name": "SQL Server 2008 R2 Database Engine Services",
"version": "10.50.1600.1",
"vendor": "Microsoft Corporation",
"valid_from": "2015-08-05T14:41:19Z",
"valid_to": "9999-12-31T23:59:59Z"
}
{
"package_property_id": 1,
"computer_id": 3,
"description": "{0AECF03B-E9FC-4A17-999D-6641A2851B26}",
"name": "BigFix Client",
"version": "9.2.3.68",
"vendor": "IBM Corp.",
"valid_from": "2015-08-05T14:41:19Z",
"valid_to": "9999-12-31T23:59:59Z"
}

```

Retrieving raw UNIX package data

9.2.2

Available from 9.2.2. This API retrieves raw UNIX package data, which contains information about all packages discovered on the UNIX operating systems. Unlike in the UI, through this API you can retrieve all historical data that is stored in the database.

To obtain information about detected software use `api/sam/v2/software_instances` API.

Before you begin

-  You must have the View Raw Data permission to perform this task.

To improve performance, retrieve raw UNIX package data in chunks of 10 000 records.

1. Retrieve the total number of available records.

```
https://hostname:port/api/sam/raw_unix_package_facts?token=token&countSwitch=2
```

```
{"total":49999,"rows":[]}
```

2. Retrieve only the first 10 000 records by using the **limit** parameter.

```
https://hostname:port/api/sam/raw_unix_package_facts?token=token&countSwitch=1&limit=10000&offset=0
```

3. Retrieve the next 10 000 records. You omit the already retrieved records by using the **offset** parameter.

```
https://hostname:port/api/sam/raw_unix_package_facts?token=token&countSwitch=1&limit=10000&offset=10000
```



Important: Each API requests must be authenticated with the **token** parameter. You can retrieve it by using [REST API for retrieving authentication token \(on page dcccclxxii\)](#). You can also log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.

Table 106. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/sam/raw_unix_package_facts</code>
Purpose	Returns raw scanned file data.
HTTP method	<code>GET</code>
Resource URI	<code>https://server_host_name:port_number/api/sam/raw_unix_package_facts</code>
URL link relation	n/a
URI query parameters	n/a
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request Content-Type	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p>

Table 106. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Values	
en-US, ...	
	Specifies the language of the response content. If this header is not specified, the content is returned in the server language.
Response payload	
	Unix Package Fact element
Response Content-Type	
	application/json
Normal HTTP response codes	
	200 - OK
Error HTTP response codes	
	500 - "Bad Request" if a query parameter contains errors or is missing
	Message body includes an error message with details.

Available columns

Table 107. Available columns

Column	Description	Displayed by default	Type
computer_id	Identifier of the computer as defined in BigFix Inventory.	✓	Nu-mer-ic
description	Description of the package.	✓	String
name	Name of the package.	✓	String
version	Version of the package.	✓	String
vendor	Vendor of the package.	✓	String
valid_from	Date and time when the package was reported for the first time. The time is specified in the GMT time zone.	✓	String

Table 107. Available columns (continued)

Column	Description	Displayed by default	Type
valid_to	Date and time when the package was reported for the last time. The time is specified in the GMT time zone.	✓	String

Query parameters

You can use query parameters to narrow down the results of your search.

Table 108. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, a set of default columns is retrieved. Example: Retrieve the name and computer_id columns: <pre>URL?columns[]=name&columns[]=computer_id</pre>	No	String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by name descending: <pre>URL?order[]=name desc</pre>	No	String
limit	Specify the number of rows to retrieve. If you omit this parameter, the limit is set according to the value of the <code>raw_data_api_default_limit</code> server setting, which by default is 100 000.	No	Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150: <pre>URL?limit=50&offset=150</pre>	No	Numeric

Table 108. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
token	A unique user authentication identifier. You can view your token in the Profile preferences of BigFix Inventory.	Yes	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written on one line:		
	<pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criteria> [{ <comma> <criteria> }...] <right-bracket> <right-brace></pre>		
	<pre><boolean-operator> ::= "and" "or" <criteria> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null></pre>		

For more information about operators, see [Common connectors and operators \(on page dcccxi\)](#).

Example: Retrieve UNIX package data from computer systems with ID greater than 10:

```
URL?criteria={ "and": [ ["computer_id", ">", "10"]] }
```

For columns that use the date and time values, such as Last Seen, you can retrieve data also for a period instead of a specific date. To do so, use *last* or *next* as **<operator>**, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively. For example, to retrieve computer systems that reported within last 7 days, use the following API request:

```
URL?criteria={ "and": [ ["last_seen", "last", "P7D"]] }
```

Example HTTP conversation

Request

```
GET api/sam/raw_unix_package_facts
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
```

```
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

Response body (JSON)

```
{
  "unix_package_property_id": -1,
  "computer_id": 14,
  "name": "lohit-oriya-fonts",
  "version": "2.4.3-6.el6",
  "vendor": "Red Hat, Inc.",
  "description": null,
  "type": "Rpm",
  "index_sha1": "176d82d8994b5c7b27f5ba8446cb40a802b2e8f5",
  "valid_from": "2016-02-29T15:33:10Z",
  "valid_to": "9999-12-31T23:59:59Z"
}
{
  "unix_package_property_id": -1,
  "computer_id": 14,
  "name": "libcollection",
  "version": "0.6.0-9.el6",
  "vendor": "Red Hat, Inc.",
  "description": null,
  "type": "Rpm",
  "index_sha1": "69ac3e6fc5674c4feb1c90aae88975d9d0ebe615",
  "valid_from": "2016-02-29T15:33:10Z",
  "valid_to": "9999-12-31T23:59:59Z"
}
```

Retrieving raw metering data

9.2.2 Available from 9.2.2. This API retrieves raw metering data, which contains information about the use of software items that are installed on your computers. Unlike in the UI, through this API you can retrieve all historical data that is stored in the database.

To obtain information about detected software use [api/sam/v2/software_instances](#) API.

Before you begin

-  You must have the View Raw Data permission to perform this task.

To improve performance, retrieve raw metering data in chunks of 10 000 records.

1. Retrieve the total number of available records.

```
https://hostname:port/api/sam/raw_app_usage_property_values?token=token&countSwitch=2
{"total":49999,"rows":[]}
```

2. Retrieve only the first 10 000 records by using the **limit** parameter.

```
https://hostname:port/api/sam/raw_app_usage_property_values?token=token&countSwitch=1
&limit=10000&offset=0
```

3. Retrieve the next 10 000 records. You omit the already retrieved records by using the **offset** parameter.

```
https://hostname:port/api/sam/raw_app_usage_property_values?token=token&countSwitch=1
&limit=10000&offset=10000
```



Note: Each API request must be authenticated with the token parameter. You can retrieve it by using [REST API for retrieving authentication token \(on page dccccxxii\)](#). You can also log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.

Table 109. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	GET /api/sam/raw_app_usage_property_values
Purpose	Returns raw package data.
HTTP method	GET
Resource URI	<code>https://server_host_name:port_number/api/sam/raw_app_usage_property_values</code>
URL link relation	n/a
URI query parameters	n/a
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p>

Table 109. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
	Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.
Request payload	n/a
Request <code>Content-Type</code>	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	<code>App Usage Property Value</code> element
Response <code>Content-Type</code>	<code>application/json</code>
Normal HTTP response codes	200 - OK
Error HTTP response codes	<p>500 - "Bad Request" if a query parameter contains errors or is missing</p> <p>Message body includes an error message with details.</p>

Query parameters

You can use query parameters to narrow down the results of your search.

Table 110. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, a set of default columns is retrieved. Example: Retrieve the name and computer_id columns: <pre>URL?columns[]=name&columns[]=computer_id</pre>	No	String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by name descending: <pre>URL?order[]=name desc</pre>	No	String
limit	Specify the number of rows to retrieve. If you omit this parameter, the limit is set according to the value of the <code>raw_data_api_default_limit</code> server setting, which by default is 100 000.	No	Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150: <pre>URL?limit=50&offset=150</pre>	No	Numeric
token	A unique user authentication identifier. You can view your token in the Profile preferences of BigFix Inventory.	Yes	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written on one line: <pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criterion> [{ <comma> <criterion> }...] <right-bracket> <right-brace></pre>		

Table 110. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<pre><boolean-operator> ::= "and" "or" <criterion> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null></pre>		
	<p>For more information about operators, see Common connectors and operators (on page dcccclxii).</p> <p>Example: Retrieve metering data from computer systems with ID greater than 10:</p> <pre>URL?criteria={ "and": [["computer_id", ">", "10"]] }</pre>		
	<p>For columns that use the date and time values, such as Last Seen, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <operator>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively. For example, to retrieve computer systems that reported within last 7 days, use the following API request:</p> <pre>URL?criteria={"and": [{"last_seen", "last", "P7D"}]}</pre>		

Example HTTP conversation

Request

```
GET api/sam/raw_app_usage_property_values
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

Response body (JSON)

```

{
  "app_usage_property_id": -1,
  "computer_id": 2,
  "process": "BESClient.exe",
  "first_used": "2015-03-27T16:28:47Z",
  "last_used": "2015-08-03T11:46:16Z",
  "total_runs": 5,
  "total_time": 10888494,
  "avg_runs_per_day": 0.03943099252280677,
  "avg_run_time": 2177698.8,
  "valid_from": "2015-08-05T14:41:19Z",
  "valid_to": "9999-12-31T23:59:59Z"
}
{
  "app_usage_property_id": -1,
  "computer_id": 2,
  "process": "BESConsole.exe",
  "first_used": "2015-07-09T11:27:50Z",
  "last_used": "2015-07-13T13:33:42Z",
  "total_runs": 2,
  "total_time": 16767,
  "avg_runs_per_day": 0.4893077201884741,
  "avg_run_time": 8383.5,
  "valid_from": "2015-08-05T14:41:19Z",
  "valid_to": "9999-12-31T23:59:59Z"
}

```

REST API for software inventory management

You can reduce the time needed to manage your software inventory in a large environment by using REST API instead of the application user interface. You can use REST API to get a list of software products, releases, and release instances. You can also find out how your software can be bundled.



Note: REST API for software inventory management work only for IBM products. The REST APIs return only instances that are assigned to default license metrics.

Retrieval of software products

You use the `GET` operation on the `api/sam/swinventory/products` element to request information about the software products that are installed in your infrastructure.



You must have the Manage Software Classification permission to perform this task.



Important: Use this REST API only for IBM products. This REST API returns only instances that are assigned to default license metrics.

Table 111. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	GET /api/sam/swinventory/products
Purpose	Returns a list of software products.
HTTP method	GET
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/products</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmviii) .
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request Content-Type	Application/json
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>Application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p>

Table 111. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Values	
en-US	
	Specifies the language of the response content. If this header is not specified, the content is returned in the server language.
Response payload	
Response Content-Type	Products element
	Application/json
Normal HTTP response codes	
	200 - OK
Error HTTP response codes	
	400 - "Bad Request" if a query parameter contains errors or is missing
	401 - "Unauthorized" if the user is not authorized to perform the operation
	416 - "Requested Range Not Satisfiable" if the user provided a start or count range that cannot be satisfied
	Message body includes an error message with details.

Available columns

Table 112. Available columns

The table consists of four columns and 9 rows.

Column	Description	Displayed by default	Type
nmbOfAllRows	Number of products that are returned by the REST API request.	✓	Integer

Table 112. Available columns

The table consists of four columns and 9 rows.

(continued)

Column	Description	Displayed by default	Type
licenseType	<p>Indicates the type of the license that is assigned to the product. Possible values are:</p> <ul style="list-style-type: none"> • 1 - Other • 5 - PVU • 6 - RVU MAPC <p> Tip: To retrieve more detailed information about the type of the license, use the api/sam/v2/software_instances (on page cmxcii) REST API.</p>	✓	Integer
level	Indicates whether the discovered software is a component or a product. In this REST API, the returned value is always <code>product</code> .	✓	String
children	Indicates whether any instances of this product release are discovered.	✓	Boolean
confidence	Confidence of the default assignment of a software component to the software product that is expressed in per cents.	✓	Integer
productReleaseComponent	Name of the software product.	✓	String
isConfirmed	<p>Indicates whether the software classification was manually confirmed. Possible values are:</p> <ul style="list-style-type: none"> • 0 - Not confirmed • 1 - Confirmed 	✓	Integer
id	Identifier of the software product.	✓	Integer
type	For internal use only. In this REST API, the returned value is always <code>root</code> .	✓	String

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `swinventory/products` element.

Table 113. Query parameters for retrieving software products

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
<code>confirmed_filter</code>	The state of bundling confirmation. If the parameter is not specified, both confirmed and unconfirmed software instances are displayed.	No	truefalse
<code>count</code>	The number of rows to be returned.	No	Numeric Default: 80
<code>endDate</code>	Discovery end date. If the parameter is not specified, the current date is used.	No	Date in the YYYY-MM-DD format
<code>part_number_filter</code>	The part number to which the software item was assigned according to the imported part number list.	No	String
<code>product_release_component_filter</code>	The name of the software product, release, or component.	No	String
<code>start</code>	The number of the first row of data to be returned.	No	Numeric Default: 0
<code>startDate</code>	Discovery start date. If the parameter is not specified, the start date is set either to 90 days back, or to the server installation date, depending on which one of these dates is closer to the current date.	No	Date in the YYYY-MM-DD format
<code>token</code>	A unique user authentication identifier.	Yes	Alphanumeric



Note: The token includes information about the computer group to which the user has access. Data that is retrieved is limited to data from computers that belong to this group.

Example HTTP conversation

Request

```
GET /api/sam/swinventory/products
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623&startDate=2012-10-19
```

```
&endDate=2012-10-19 HTTP/1.1
Host: localhost:9080
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

Response body (JSON)

```
[{
  "nmbOfAllRows" : "1967",
  "licenseType" : "1",
  "level" : "product",
  "children" : "true",
  "confidence" : "99",
  "productReleaseComponent" : "AIX 5.2 Workload Partitions for AIX 7",
  "isConfirmed" : "false",
  "id" : "30369",
  "type" : "root"
```

Retrieval of releases of a product

You use the `GET` operation on the `api/sam/swinventory/product/{product_id}/releases` element to request information about the releases of a particular software product that is installed in your infrastructure.

 You must have the Manage Software Classification permission to perform this task.

 **Important:** Use this REST API only for IBM products. This REST API returns only instances that are assigned to default license metrics.

Table 114. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/sam/swinventory/product/{product_id}/releases</code>
Purpose	Returns a list of releases of a particular software product that is identified by its identifier.
HTTP method	<code>GET</code>

Table 114. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/product/{product_id}/releases</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmxi) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	<code>Application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>Application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	<code>Releases</code> element

Table 114. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Response Content-Type	Application/json
Normal HTTP response codes	200 - OK
Error HTTP response codes	<p>400 - "Bad Request" if a query parameter contains errors or is missing</p> <p>401 - "Unauthorized" if the user is not authorized to perform the operation</p> <p>404 - "Not Found" if a release or instance of the product or version that is requested does not exist</p> <p>416 - "Requested Range Not Satisfiable" if the user provided a start or count range that cannot be satisfied</p>
	Message body includes an error message with details.

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `swinventory/product/{product_id}/releases` element.

Table 115. Query parameters for retrieving releases of a software product

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
confirmed_filter	The state of bundling confirmation. If the parameter is not specified, both confirmed and unconfirmed software instances are displayed.	No	truefalse
count	The number of rows to be returned.	No	Numeric Default: 80
endDate	Discovery end date. If the parameter is not specified, the current date is used.	No	Date in the YYYY-MM-DD format

Table 115. Query parameters for retrieving releases of a software product

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
part_number_filter	The part number to which the software item was assigned according to the imported part number list.	No	String
product_release_-component_filter	The name of the software product, release, or component.	No	String
start	The number of the first row of data to be returned.	No	Numeric Default: 0
startDate	Discovery start date. If the parameter is not specified, the start date is set either to 90 days back, or to the server installation date, depending on which one of these dates is closer to the current date.	No	Date in the YYYY-MM-DD format
token	A unique user authentication identifier.	Yes	Alphanumeric

 **Note:** The token includes information about the computer group to which the user has access. Data that is retrieved is limited to data from computers that belong to this group.

Example HTTP conversation

Request

```
GET /api/sam/swinventory/product/38818/releases
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623&startDate=2012-10-19
&endDate=2012-10-19 HTTP/1.1
Host: localhost:9080
Accept: application/json Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

Response body (JSON)

```
[{
  "nmbOfAllRows": "1",
  "level": "release",
  "productReleaseComponent": "DB2 UDB Query Patroller 6.1",
  "id": "61922",
  "numberOfAllInstances": "1",
  "children": "true"
}]
```

Retrieval of release instances

You use the `GET` operation on the `api/samswinventory/release/{release_id}/instances` element to request information about the instances of a software product release for a particular release that is installed in your infrastructure.

 You must have the Manage Software Classification permission to perform this task.

 **Important:** Use this REST API only for IBM products. This REST API returns only instances that are assigned to default license metrics.

Table 116. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/sam/swinventory/release/{release_id}/instances</code>
Purpose	Returns a list of instances for the requested software product release.
HTTP method	<code>GET</code>
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/release/{release_id}/instances</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmxv) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>

Table 116. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Request payload	n/a
Request Content-Type	Application/json
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>Application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	Instances element. Each instance is a single component on a single agent.
Response Content-Type	Application/json
Normal HTTP response codes	200 - OK
Error HTTP response codes	<p>400 - "Bad Request" if a query parameter contains errors or is missing</p> <p>401 - "Unauthorized" if the user is not authorized to perform the operation</p> <p>404 - "Not Found" if a release or instance of the product or version that is requested does not exist</p> <p>416 - "Requested Range Not Satisfiable" if the user provided a start or count range that cannot be satisfied</p>

Table 116. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
	Message body includes an error message with details.

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `swinventory/release/{release_id}/instances` element.

Table 117. Query parameters for retrieving release instances

The table consists of four columns and 12 rows.

Parameter	Description	Required	Value
confirmed_filter	The state of bundling confirmation. If the parameter is not specified, both confirmed and unconfirmed software instances are displayed.	No	truefalse
count	The number of rows to be returned.	No	Numeric Default: 80
endDate	Discovery end date. If the parameter is not specified, the current date is used.	No	Date in the YYYY-MM-DD format
name_host_filter	The name of the computer on which the software item is installed.	No	String
name_server_filter	The name of the server on which the software item is installed.	No	String
operating_system_filter	Operating system of the computer on which the software item is installed.	No	String
part_number_filter	The part number to which the software item was assigned according to the imported part number list.	No	String
product_release_component_filter	The name of the software product, release, or component.	No	String
start	The number of the first row of data to be returned.	No	Numeric Default: 0
startDate	Discovery start date. If the parameter is not specified, the start date is set either to 90 days back, or	No	Date in the YYYY-MM-DD format

Table 117. Query parameters for retrieving release instances

The table consists of four columns and 12 rows.

(continued)

Parameter	Description	Required	Value
	to the server installation date, depending on which one of these dates is closer to the current date.		
token	A unique user authentication identifier.	Yes	Alphanumeric
	 Note: The token includes information about the computer group to which the user has access. Data that is retrieved is limited to data from computers that belong to this group.		

Example HTTP conversation

Request

```
GET /api/sam/swinventory/release/61922/instances
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623&startDate=2012-10-19
&endDate=2012-10-19 HTTP/1.1
Host: localhost:9080
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

Response body (JSON)

```
[{
  "isCharged": "true",
  "operatingSystem": "Linux Red Hat Enterprise Server 5.4 (2.6.18-164.el5)",
  "isConfirmed": "false",
  "currentServerId": "TLM_VM_VMware-42 3b 3b 73 2a 12 4a c5-e0 56 d1 30 74 6b 53 2a",
  "updateTime": "1374082814268",
  "children": "false",
  "hostname": "NC042189",
  "productInventoryId": "347",
```

```

"id": "347",
"level": "instance",
"bundleRules": "the relation in the software catalog, the stand-alone product discovery",
"processorType": "Intel(R) Xeon(R) Multi-core 3400-3699 or 5500-5699",
"isAgentDeleted": "false",
"productReleaseComponent": "IBM CICS Transaction Gateway 6.0",
"pvuPerCore": "70",
"installationPaths": "/nfs/bak/SLES10.2/usr/catalog_test/linux01/1",
"nmbOfAllRows": "4",
"isSimple": "false"
}]

```

Retrieval of releases to which a software instance can be reassigned

You use the `GET` operation on the `api/sam/swinventory/targetBundlesOfInstances` element to request information about the possible releases to which the requested software instance can be reassigned.



You must have the Manage Software Classification permission to perform this task.



Important: Use this REST API only for IBM products. This REST API returns only instances that are assigned to default license metrics.

Table 118. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/sam/swinventory/targetBundlesOfInstances</code>
Purpose	Returns a list of possible releases to which the requested software instance can be reassigned.
HTTP method	GET
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/targetBundlesOfInstances</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmxix) .
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p>

Table 118. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
	<p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <small>Content-Type</small>	<p><code>Application/json</code></p>
Response headers	<p>Header</p> <p><small>Content-Type</small></p> <p>Values</p> <p>Application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><small>Content-Language</small></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	<code>targetReleases</code> element
Response <small>Content-Type</small>	<p><code>Application/json</code></p>
Normal HTTP response codes	200 - OK

Table 118. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Error HTTP response codes	
	400 - "Bad Request" if a query parameter contains errors or is missing
	401 - "Unauthorized" if the user is not authorized to perform the operation
	404 - "Not Found" if a release or instance of the product or version that is requested does not exist
	Message body includes an error message with details.

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `swinventory/targetBundlesOfInstances` element.

Table 119. Query parameters for retrieving releases to which a software instance can be reassigned

The table consists of four columns and three rows.

Parameter	Description	Required	Value
productInventoryId	A list of unique product identifiers separated with a comma.	Yes	Numeric
token	A unique user authentication identifier.	Yes	Alphanumeric



Note: The token includes information about the computer group to which the user has access. Data that is retrieved is limited to data from computers that belong to this group.

Example HTTP conversation

Request

```
GET /api/sam/swinventory/targetBundlesOfInstances
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623&productInventoryId=200032 HTTP/1.1
Host: localhost:9080
```

```
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

Response body (JSON)

```
[{
  "productId": "69558",
  "isSelected": "false",
  "productInventoryId": "200032",
  "appliedRules": "the relation in the software catalog,
the stand-alone product discovery",
  "productName": "BigFix Smart Analytics System 2050 Departmental Base Remote Standby Server 1.0",
  "id": "69558",
  "branchType": "0",
  "isShared": "false"
}]
```

Retrieval of releases to which a release instance can be reassigned

You use the `GET` operation on the `api/sam/swinventory/targetBundlesOfReleases` element to request information about the possible releases to which the instances of the requested software releases can be reassigned.

 You must have the Manage Software Classification permission to perform this task.

 **Important:** Use this REST API only for IBM products. This REST API returns only instances that are assigned to default license metrics.

Table 120. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/sam/swinventory/targetBundlesOfReleases</code>
Purpose	Returns a list of possible releases to which the instances of the requested software releases can be reassigned.
HTTP method	<code>GET</code>

Table 120. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/targetBundlesOfReleases</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmxxii) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	<code>Application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p>Application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	<code>targetReleases</code> element

Table 120. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Response Content-Type	Application/json
Normal HTTP response codes	200 - OK
Error HTTP response codes	<p>400 - "Bad Request" if a query parameter contains errors or is missing</p> <p>401 - "Unauthorized" if the user is not authorized to perform the operation</p> <p>404 - "Not Found" if a release or instance of the product or version that is requested does not exist</p>
	Message body includes an error message with details.

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `swinventory/targetBundlesOfReleases` element.

Table 121. Query parameters for retrieving releases to which a release instance can be reassigned

The table consists of four columns and 11 rows.

Parameter	Description	Required	Value
confirmed_filter	The state of bundling confirmation. If the parameter is not specified, both confirmed and unconfirmed software instances are displayed.	No	true/false
endDate	Discovery end date. If the parameter is not specified, the current date is used.	No	Date in the YYYY-MM-DD format
name_host_filter	The name of the computer on which the software item is installed.	No	String
name_server_filter	The name of the server on which the software item is installed.	No	String

Table 121. Query parameters for retrieving releases to which a release instance can be reassigned

The table consists of four columns and 11 rows.

(continued)

Parameter	Description	Required	Value
operating_system_filter	Operating system of the computer on which the software item is installed.	No	String
part_number_filter	The part number to which the software item was assigned according to the imported part number list.	No	String
product_release_component_filter	The name of the software product, release, or component.	No	String
releases	A list of unique release identifiers separated with a comma.	Yes	Numeric
startDate	Discovery start date. If the parameter is not specified, the start date is set either to 90 days back, or to the server installation date, depending on which one of these dates is closer to the current date.	No	Date in the YYYY-MM-DD format
token	A unique user authentication identifier.	Yes	Alphanumeric

 **Note:** The token includes information about the computer group to which the user has access. Data that is retrieved is limited to data from computers that belong to this group.

Example HTTP conversation

Request

```
GET /api/sam/swinventory/targetBundlesOfReleases
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
&releases=64050&startDate=2012-09-29&endDate=2012-10-02 HTTP/1.1
Host: localhost:9080
Accept: application/json
Accept-Language: en-US
```

Response header

```

HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US

```

Response body (JSON)

```

[ {
  "productId": "72563",
  "isSelected": "false",
  "productInventoryId": "200223",
  "appliedRules": "the relation in the software catalog",
  "productName": "BigFix",
  "id": "72563",
  "branchType": "0",
  "isShared": "false"
} ]

```

Retrieval of instances to reassign to a release

You use the `GET` operation on the `api/sam/swinventory/targetInstances` element to request information about the target instances that can be reassigned to a particular release.

 You must have the Manage Software Classification permission to perform this task.

 **Important:** Use this REST API only for IBM products. This REST API returns only instances that are assigned to default license metrics.

Table 122. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/sam/swinventory/targetInstances</code>
Purpose	Returns a list of target instances that can be reassigned to a particular release.
HTTP method	<code>GET</code>
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/targetInstances</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmxxvi) .

Table 122. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Request headers	Header
	<code>Accept-Language</code> (optional)
	Values
	en-US (only English is supported)
	Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.
Request payload	n/a
Request <code>Content-Type</code>	
	<code>Application/json</code>
Response headers	Header
	<code>Content-Type</code>
	Values
	<code>Application/json</code>
	Specifies the content type of the response.
	Header
	<code>Content-Language</code>
	Values
	en-US, ...
	Specifies the language of the response content. If this header is not specified, the content is returned in the server language.
Response payload	<code>targetInstances</code> element
Response <code>Content-Type</code>	
	<code>Application/json</code>
Normal HTTP response codes	
	200 - OK

Table 122. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Error HTTP response codes	
	400 - "Bad Request" if a query parameter contains errors or is missing
	401 - "Unauthorized" if a user is not authorized to perform the operation
	404 - "Not Found" if a release or instance of the product or version that you requested for does not exist
	Message body includes an error message with details.

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `GET/swinventory/targetInstances` element.

Table 123. Query parameters for retrieving instances to reassign to a release

The table consists of four columns and three rows.

Parameter	Description	Required	Value
releaseId	A unique identifier of a release.	Yes	Numeric
token	A unique user authentication identifier.	Yes	Alphanumeric



Note: The token includes information about the computer group to which the user has access. Data that is retrieved is limited to data from computers that belong to this group.

Example HTTP conversation

Request

```
GET /api/sam/swinventory/targetInstances
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623&releaseId=72717 HTTP/1.1
Host: localhost:9080
```

```
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

Response body (JSON)

```
[{
  "componentNameVersion": "BigFix Tivoli Directory Server (SP) 6.0",
  "hostname": "Agent20",
  "path": "C:\\Documents and Settings\\Administrator\\Signatures",
  "updateTime": "1349270690593",
  "productInventoryId": "200027",
  "appliedRules": "the relation in the software catalog, the infrastructure co-location",
  "productId": "72717",
  "id": "200027",
  "productNameVersion": "IBM WebSphere Process Server Hypervisor Edition for Novell SLES for x86 6.2",
  "isShared": "false"
}]
```

Retrieval of instances shared by releases

You use the `GET` operation on the `api/sam/swinventory/instanceToShare` element to request information about the releases that can share a particular instance.



You must have the Manage Software Classification permission to perform this task.



Important: Use this REST API only for IBM products. This REST API returns only instances that are assigned to default license metrics.

Table 124. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/sam/swinventory/instanceToShare</code>
Purpose	Returns a list of software releases that can share a particular instance.
HTTP method	<code>GET</code>

Table 124. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/instanceToShare</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmxxix) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	<code>Application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p>Application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	<code>targetReleases</code> element

Table 124. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Response <code>Content-Type</code>	<code>Application/json</code>
Normal HTTP response codes	200 - OK
Error HTTP response codes	<p>400 - "Bad Request" if a parameter contains errors or is missing</p> <p>401 - "Unauthorized" if you are not authorized for the operation</p> <p>404 - "Not Found" if a release or instance of the product or version that is requested does not exist</p>
	Message body includes an error message with details.

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `swinventory/instanceToShare` element.

Table 125. Query parameters for retrieving an instance shared by releases

The table consists of four columns and three rows.

Parameter	Description	Required	Value
<code>productInventoryId</code>	A unique identifier of a product.	Yes	Numeric
<code>token</code>	A unique user authentication identifier.	Yes	Alphanumeric



Note: The token includes information about the computer group to which the user has access. Data that is retrieved is limited to data from computers that belong to this group.

Example HTTP conversation

Request

```
GET /api/sam/swinventory/instanceToShare
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623&productInventoryId=200032 HTTP/1.1
Host: localhost:9080
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

Response body (JSON)

```
[{
  "productId": "64050",
  "isSelected": "true",
  "productInventoryId": "200032",
  "appliedRules": "the relation in the software catalog, the stand-alone product discovery",
  "productName": "IBM Alphablox for Linux, UNIX and Windows 9.5",
  "id": "64050",
  "branchType": "0",
  "isShared": "false"
}]
```

REST API for software classification

You can reduce the time that is needed to manage your software instances in a large environment by using REST API instead of the application user interface. You can use REST API to share instances, include or exclude instances from a pricing calculation, confirm and assign instances, and reassign instances.



Important: REST APIs for software classification work only for BigFix products and return only instances that are assigned to default license metrics.

Share an instance with more than one product

You use the `POST` operation on the `api/sam/swinventory/share` element to share an instance of a component with more than one product.



You must have the Manage Software Classification permission to perform this task.



Important: Use this REST API only for BigFix products. The REST API returns only instances that are assigned to default license metrics.

Table 126. Operation descriptions

Operation details	Description
Operation	<code>/api/sam/swinventory/share</code>
Purpose	Shares an instance with a list of software products.
HTTP method	POST
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/share</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmxxxii) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	n/a
Response headers	<p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	n/a
Response <code>Content-Type</code>	n/a
Normal HTTP response codes	<code>204 - No content</code>

Table 126. Operation descriptions (continued)

Operation details	Description
Error HTTP response codes	
	400 - "Bad Request" if a query parameter contains errors or is missing
	401 - "Unauthorized" if the user is not authorized to perform the operation
	404 - "Not Found" if a release or instance of the product or version that the user requested does not exist
	422 - "Unprocessable Entity" if the request was well-formed but was unable to be followed due to semantic errors
Message body includes an error message with details.	

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you use for the `POST/swinventory/share` element.

Table 127. Query parameters for sharing an instance with more than one product

The table consists of four columns and four rows.

Parameter	Description	Required	Value
productIds	A list of unique product identifiers separated with a comma. At least one identifier must be specified. Use the <code>GET</code> operation on the <code>swinventory/instance-ToShare</code> element to get product identifiers.	Yes	Numeric
productInventoryId	A unique identifier of an instance that is shared by products.	Yes	Numeric
token	A unique user authentication identifier.	Yes	Alphanumeric
updateTime	Timestamp of the last modification time of the instance expressed in milliseconds. This parameter is used to handle concurrent actions. Use the <code>GET</code> operation on the <code>swinventory/release/{release_id}/instances</code> element to get the update time.	Yes	Numeric

Example HTTP conversation

Request

```
POST /api/sam/swinventory/share?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
&productInventoryId=200032&updateTime=1349237658578&productIds=70665,70656 HTTP/1.1
Host: localhost:9080
Accept-Language: en-US
```

Response header

```
HTTP/1.1 204 OK
Content-Type: application/json
Content-Language: en-US
```

Related information

[Retrieval of release instances \(on page cmxiii\)](#)

[Retrieval of instances shared by releases \(on page cmxxvii\)](#)

Inclusion of instances in a pricing calculation

You use the `POST` operation on the `api/sam/swinventory/include` element to include instances in the pricing calculation.

 You must have the Manage Software Classification permission to perform this task.



Important: Use this REST API only for BigFix products. The REST API returns only instances that are assigned to default license metrics.

Table 128. Operation descriptions

Operation details	Description
Operation	<code>/api/sam/swinventory/include</code>
Purpose	Includes instances in a pricing calculation.
HTTP method	<code>POST</code>
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/include</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmxxxiv) .
Request headers	Header

Table 128. Operation descriptions (continued)

Operation details	Description
	<p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	n/a
Response headers	<p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	n/a
Response <code>Content-Type</code>	n/a
Normal HTTP response codes	<p><code>204 - No content</code></p>
Error HTTP response codes	<p><code>400 - "Bad Request"</code> if a query parameter contains errors or is missing</p> <p><code>401 - "Unauthorized"</code> if the user is not authorized to perform the operation</p> <p><code>404 - "Not Found"</code> if a release or instance of the product or version that the user requested does not exist</p> <p><code>422 - "Unprocessable Entity"</code> if the request was well-formed but was unable to be followed due to semantic errors</p> <p>Message body includes an error message with details.</p>

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `POST/swinventory/include` element.

Table 129. Query parameters for including instances in pricing calculations

The table consists of four columns and three rows.

Parameter	Description	Required	Value
productInventoryId	A list of unique identifiers of instances that are to be included in the calculation separated with a comma. At least one identifier must be specified.	Yes	Numeric
token	A unique user authentication identifier.	Yes	Alphanumeric
updateTime	A list of timestamps of the last modification time of instances expressed in milliseconds. The timestamps are separated with a comma. The first timestamp in the list corresponds to the first instance, and so on. This parameter is used to handle concurrent actions. Use the <code>GET</code> operation on the <code>swinventory/release/{release_id}/instances</code> element to get the update times.	Yes	Numeric

Example HTTP conversation

Request

```
POST /api/sam/swinventory/include?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
&productInventoryId=198,300201&updateTime=1349237658578,1349237658588 HTTP/1.1
Host: localhost:9080
Accept-Language: en-US
```

Response header

```
HTTP/1.1 204 OK
Content-Type: application/json
Content-Language: en-US
```

Related information

[Retrieval of release instances \(on page cmxiii\)](#)

Exclusion of instances from pricing calculations

You use the `POST` operation on the `api/sam/swinventory/exclude` element to exclude instances from pricing calculations.

 You must have the Manage Software Classification permission to perform this task.

 **Important:** Use this REST API only for BigFix products. The REST API returns only instances that are assigned to default license metrics.

Table 130. Operation descriptions

Operation details	Description
Operation	<code>/api/sam/swinventory/exclude</code>
Purpose	Excludes instances from a pricing calculation.
HTTP method	POST
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/exclude</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmxxxvii) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	n/a
Response headers	<p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	n/a
Response <code>Content-Type</code>	n/a

Table 130. Operation descriptions (continued)

Operation details	Description
Normal HTTP response codes	
	204 - No content
Error HTTP response codes	
	400 - "Bad Request" if a query parameter contains errors or is missing
	401 - "Unauthorized" if the user is not authorized to perform the operation
	404 - "Not Found" if a release or instance of the product or version that the user requested does not exist
	422 - "Unprocessable Entity" if the request was well-formed but was unable to be followed due to semantic errors
Message body excludes an error message with details.	

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `POST/swinventory/exclude` element.

Table 131. Query parameters for excluding instances from pricing calculations

Parameter	Description	Required	Value
productInventoryId	A list of unique identifiers of instances to be excluded from the calculation separated with a comma. At least one identifier must be specified.	Yes	Numeric
reason	The reason for excluding an instance from the pricing calculations.	Yes	<code>backupbetacomponentevaluation-</code> <code>no_licensingnot_compatibleother</code>
comment	Additional comments for the reason of exclusion.	Only when reason is other	String
token	A unique user authentication identifier.	Yes	Alphanumeric
updateTime	A list of timestamps of the last modification time of the instances expressed in milliseconds. Each timestamp is separated with a comma. The first timestamp in the list corresponds to the first instance, and so on. This	Yes	Numeric

Table 131. Query parameters for excluding instances from pricing calculations (continued)

Parameter	Description	Required	Value
	parameter is used to handle concurrent actions.		
	Use the <code>GET</code> operation on the <code>swinventory/release/{release_id}/instances</code> element to get the update times.		

Example HTTP conversation

Request

```
POST /api/sam/swinventory/exclude
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
&productInventoryId=300001,500001&updateTime=1349237658578,1349237658588
&reason=no_licensing HTTP/1.1
Host: localhost:9080
Accept-Language: en-US
```

Response header

```
HTTP/1.1 204 OK
Content-Type: application/json
Content-Language: en-US
```

Related information

[Retrieval of release instances \(on page cmxiii\)](#)

Reassignment of instances to a product

You use the `POST` operation on the `api/sam/swinventory/reassign` element to reassign instances to a different product.

 You must have the Manage Software Classification permission to perform this task.



Important: Use this REST API only for BigFix products. The REST API returns only instances that are assigned to default license metrics.

Table 132. Operation descriptions

Operation details	Description
Operation	<code>/api/sam/swinventory/reassign</code>

Table 132. Operation descriptions (continued)

Operation details	Description
Purpose	Reassigns instances to a different product.
HTTP method	POST
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/reassign</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmxl) .
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request Content-Type	n/a
Response headers	<p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	n/a
Response Content-Type	n/a
Normal HTTP response codes	204 - No content
Error HTTP response codes	<p>400 - "Bad Request" if a query parameter contains errors or is missing</p> <p>401 - "Unauthorized" if the user is not authorized to perform the operation</p>

Table 132. Operation descriptions (continued)

Operation details	Description
	404 - "Not Found" if a release or instance of the product or version that the user requested does not exist
	422 - "Unprocessable Entity" if the request was well-formed but was unable to be followed due to semantic errors
Message body excludes an error message with details.	

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `POST/swinventory/reassign` element.

Table 133. Query parameters for reassigning instances to a different product

The table consists of four columns and four rows.

Parameter	Description	Required	Value
productId	A unique identifier of a product to which the instance is to be reassigned. Use the <code>GET</code> operation on the <code>swinventory/target-BundlesOfInstances</code> element to get product identifiers.	Yes	Numeric
productInventoryId	A list of unique identifiers of instances to be reassigned to a product separated with a comma. At least one identifier must be specified.	Yes	Numeric
updateTime	A list of timestamps of the last modification time of the instances expressed in milliseconds. The timestamps are separated with a comma. The first timestamp in the list corresponds to the first instance, and so on. This parameter is used to handle concurrent actions. Use the <code>GET</code> operation on the <code>swinventory/release/{release_id}/instances</code> element to get the update times.	Yes	Numeric
token	A unique user authentication identifier.	Yes	Alphanumeric

Example HTTP conversation

Request

```
POST /api/sam/swinventory/reassign?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
&productInventoryId=200027,300001&productId=72786&updateTime=
1349237658578,1349237658588 HTTP/1.1
Host: localhost:9080
Accept-Language: en-US
```

Response header

```
HTTP/1.1 204 OK
Content-Type: application/json
Content-Language: en-US
```

Related information

[Retrieval of release instances \(on page cmxiii\)](#)

[Retrieval of releases to which a software instance can be reassigned \(on page cmxvii\)](#)

Reassignment of instances to a product for a list of releases

You use the `POST` operation on the `api/sam/swinventory/reassignRelease` element to reassign instances for a list of releases to a product.



You must have the Manage Software Classification permission to perform this task.



Important: Use this REST API only for BigFix products. The REST API returns only instances that are assigned to default license metrics.

Table 134. Operation descriptions

Operation details	Description
Operation	<code>/api/sam/swinventory/reassignRelease</code>
Purpose	Reassigns instances to a product for a list of releases.
HTTP method	<code>POST</code>
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/reassignRelease</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmxliv) .

Table 134. Operation descriptions (continued)

Operation details	Description
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	n/a
Response headers	<p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	n/a
Response <code>Content-Type</code>	n/a
Normal HTTP response codes	<p><code>204 - No content</code></p>
Error HTTP response codes	<p><code>400 - "Bad Request"</code> if a query parameter contains errors or is missing</p> <p><code>401 - "Unauthorized"</code> if the user is not authorized to perform the operation</p> <p><code>404 - "Not Found"</code> if a release or instance of the product or version that the user requested does not exist</p> <p><code>422 - "Unprocessable Entity"</code> if the request was well-formed but was unable to be followed due to semantic errors</p> <p>Message body excludes an error message with details.</p>

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `POST/swinventory/reassignRelease` element.

Table 135. Query parameters for retrieving instances to reassign to a release

The table consists of four columns and ten rows.

Parameter	Description	Required	Value
endDate	Discovery end date. If the parameter is not specified, the current date is used.	No	Date in the YYYY-MM-DD format
name_host_filter	The name of the computer on which the software item is installed.	No	String
name_server_filter	The name of the server on which the software item is installed.	No	String
operating_system_filter	Operating system of the computer on which the software item is installed.	No	String
part_number_filter	The part number to which the software item was assigned according to the imported part number list.	No	String
product_release_component_filter	The name of the software product, release, or component.	No	String
productId	A unique identifier of a product to which the instance is to be assigned. Use the <code>GET</code> operation on the <code>swinventory/target-BundlesOfReleases</code> element to get product identifiers.	Yes	Numeric
releases	A list of unique release identifiers separated with a comma. At least one identifier must be specified.	Yes	Numeric
startDate	Discovery start date. If the parameter is not specified, the first date that is not covered by a signed audit report is used. If there are no signed reports, the server installation date is used.	No	Date in the YYYY-MM-DD format
token	A unique user authentication identifier.	Yes	Alphanumeric

Example HTTP conversation

Request

```
POST /api/sam/swinventory/reassignRelease
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
&releases=67757,95947&productId=67756&startDate=2012-09-29&endDate=2012-10-02
Host: localhost:9080
Accept-Language: en-US
```

Response header

```
HTTP/1.1 204 OK
Content-Type: application/json
Content-Language: en-US
```

Related information

[Retrieval of releases to which a release instance can be reassigned \(on page cmxx\)](#)

Confirmation of instances to bundle or assign

You use the `POST` operation on the `api/sam/swinventory/confirm` element to confirm the bundling or assignment of instances.



You must have the Manage Software Classification permission to perform this task.



Important: Use this REST API only for BigFix products. The REST API returns only instances that are assigned to default license metrics.

Table 136. Operation descriptions

Operation details	Description
Operation	<code>/api/sam/swinventory/confirm</code>
Purpose	Confirms the bundling or assignment of instances.
HTTP method	<code>POST</code>
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/confirm</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmxlv) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p>

Table 136. Operation descriptions (continued)

Operation details	Description
	en-US (only English is supported)
	Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.
Request payload	n/a
Request <code>Content-Type</code>	n/a
Response headers	<p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	n/a
Response <code>Content-Type</code>	n/a
Normal HTTP response codes	<p><code>204 - No content</code></p>
Error HTTP response codes	<p><code>400 - "Bad Request"</code> if a query parameter contains errors or is missing</p> <p><code>401 - "Unauthorized"</code> if the user is not authorized to perform the operation</p> <p><code>404 - "Not Found"</code> if a release or instance of the product or version that the user requested does not exist</p> <p><code>422 - "Unprocessable Entity"</code> if the request was well-formed but was unable to be followed due to semantic errors</p> <p>Message body excludes an error message with details.</p>

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `POST/swinventory/confirm` element.

Table 137. Query parameters for confirming instances

The table consists of four columns and three rows.

Parameter	Description	Required	Value
productInventoryId	A list of unique identifiers of instances to be confirmed separated with a comma. At least one identifier must be specified.	Yes	Numeric
updateTime	A list of timestamps of the last modification time of instances expressed in milliseconds. The timestamps are separated with a comma. The first timestamp in the list corresponds to the first instance, and so on. This parameter is used to handle concurrent actions. Use the <code>GET</code> operation on the <code>swinventory/release/{release_id}/instances</code> element to get the update times.	Yes	Numeric
token	A unique user authentication identifier.	Yes	Alphanumeric

Example HTTP conversation

Request

```
POST /api/sam/swinventory/confirm?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
&productInventoryId=200027,300001&updateTime=1349237658578,1349237658588 HTTP/1.1
Host: localhost:9080
Accept-Language: en-US
```

Response header

```
HTTP/1.1 204 OK
Content-Type: application/json
Content-Language: en-US
```

Related information

[Retrieval of release instances \(on page cmxiii\)](#)

Confirmation of the assignment of instances for a release

You use the `POST` operation on the `api/sam/swinventory/confirmRelease` element to confirm the assignment of instances for a release.



You must have the Manage Software Classification permission to perform this task.



Important: Use this REST API only for BigFix products. The REST API returns only instances that are assigned to default license metrics.

Operation details	Description
Operation	<code>api/sam/swinventory/confirmRelease</code>
Purpose	Confirms assignment of instances for a release.
HTTP method	POST
Resource URI	<code>https://server_host_name:port_number/api/sam/swinventory/confirmRelease</code>
URL link relation	n/a
URI query parameters	No
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (English language only supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	n/a
Response headers	<p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	n/a
Response <code>Content-Type</code>	n/a
Normal HTTP response codes	204 - No content

Operation details	Description
Error HTTP response codes	
	400 - "Bad Request" if a query parameter contains errors or is missing
	401 - "Unauthorized" if the user is not authorized to perform the operation
	404 - "Not Found" if a release or instance of the product or version that the user requested does not exist
	422 - "Unprocessable Entity" if the request was well-formed but was unable to be followed due to semantic errors
Message body excludes an error message with details.	

Query parameters

The following table presents query parameters that you can use for the `POST/swinventory/confirmRelease` element.

Parameter	Description	Required	Value
endDate	Discovery end date. If the parameter is not specified, the current date is used.	No	Date in the YYYY-MM-DD format
name_host_filter	The name of the computer on which the software item is installed	No	String
name_server_filter	The name of the server on which the software item is installed	No	String
operating_system_filter	Operating system of the computer on which the software item is installed	No	String
part_number_filter	The part number to which the software item was assigned according to the imported part number list.	No	String
product_release_component_filter	The name of a software product, software release, or a component	No	String
releases	A list of unique release identifiers. Each identifier is separated by a comma. At least one ID must be entered	Yes	Numeric
startDate	Discovery start date. If the parameter is not specified, the first date that is not covered by a signed audit report is used. If there are no signed reports, the server installation date is used.	No	Date in the YYYY-MM-DD format

Parameter	Description	Required	Value
token	A unique user authentication identifier.	Yes	Alphanumeric

Example HTTP conversation

Request

```
POST /api/sam/swinventory/confirmRelease
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623&releases=95185
&startDate=2012-10-18&endDate=2013-01-26 HTTP/1.1
Host: localhost:9080
Accept-Language: en-US
```

Response header

```
HTTP/1.1 204 OK
Content-Type: application/json
Content-Language: en-US
```

REST API for managing VM managers

You can use REST API requests to add and change VM managers in your infrastructure.

Retrieval of VM managers

You use the `GET` operation on the `api/sam/vmmanagers` element to request information about VM managers that are defined in your infrastructure. The REST API returns more information than is visible on the user interface. It returns information about VM managers that are managed both in central and distributed mode, VM managers that were inactive for more than 90 days as well as those that were removed. It also returns information from virtualization technologies that do not require defining connection to the VM manager such as Xen or KVM without RHEV-M. The scope of VM managers from which data is collected depends on which VM Manager Tools are installed on computers that belong to the computer group to which you are assigned.

To retrieve information about VM managers, you must have the Manage VM Managers and Servers permission.

Table 138. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/sam/vmmanagers</code>
Purpose	Returns a list of VM managers.
HTTP method	<code>GET</code>

Table 138. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Resource URI	<code>https://server_host_name:port_number/api/sam/vmmanagers?token=token</code>
Request Content-Type	<code>application/json</code>
Response Content-Type	<code>application/json</code>
Normal HTTP response codes	<code>200 - OK</code>
Error HTTP response codes	<p><code>400 - "Bad Request"</code> if a query parameter contains errors or is missing</p> <p><code>401 - "Unauthorized"</code> if the user does not have the Manage VM Managers and Servers permission</p> <p><code>403 - "Forbidden"</code> if the computer where the central VM Manager Tool is installed does not belong to the computer group that the user is assigned</p> <p><code>500 - "Internal Server Error"</code> if no VM Manager Tool is defined in the BigFix Inventory database.</p> <p>Message body includes an error message with details.</p>

Response attributes

Table 139. Response attributes and their possible values

Attribute	Possible values
<code>deleted</code>	Indicates whether the VM manager was removed from BigFix Inventory.
<code>lastmodified</code>	The date and time when the VM manager was last modified.
<code>login</code>	User name that is used to access the VM manager.
<code>protocol</code>	<p>Communication protocol used by the VM manager. The attribute is applicable only for Hyper-V.</p> <ul style="list-style-type: none"> • 0 - WinRM • 1 - PowerShell • null - the VM manager is not Hyper-V

Table 139. Response attributes and their possible values (continued)

Attribute	Possible values
<code>sharedcredentials</code>	Indicates whether the VM manager shares credentials with other hosts in the same cluster. The attribute is applicable only for Hyper-V.
<code>status</code>	<p>Status of the VM manager.</p> <ul style="list-style-type: none"> • -1 - Invalid credentials - suspended • 0 - Pending • 1 - OK • 2 - Connection failed • 3 - Invalid credentials - attempting • 6 - Duplicated address • 7 - Hard timeout - suspended • 8 - OK - duplicated UUIDs discarded • 9 - No VM Manager Data • 98 - Inactive • 99 - Unknown problem <p>For more information about each status, see: VM manager statuses (on page cccli).</p>
<code>type</code>	<p>Type of the VM manager.</p> <ul style="list-style-type: none"> • 1 - Hyper-V • 2 - VMware vSphere or VMware ESX • 3 - RHEV-M • 9.2.12 4 - XenServer or Citrix XenServer • 99 - Other. The value is used for VM managers for which data is collected by using the Run Capacity Scan on Virtualization Hosts task, for example Xen and KVM without RHEV-M.
<code>url</code>	Web address of the VM manager.
<code>vmmanagerid</code>	ID of the VM manager.
<code>vmmanagertoolid</code>	ID of the VM Manager Tool that collects information from the particular VM manager.

Example HTTP conversation**Request**

```
GET http://localhost:9081/api/sam/vmmanagers
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Response body (JSON)

```
[{
  "vmmanagerid":1,
  "protocol":null,
  "deleted":false,
  "vmmanagertoolid":1,
  "sharecredentials":false,
  "lastmodified":"2016-06-12 09:06:51.723",
  "login":"admin@internal",
  "type":3,
  "url":"https://9.156.44.146:8443/api",
  "status":98
}]
```

Adding and updating VM managers

You use the `PUT` operation on the `api/sam/vmmanagers` element to add or update VM managers. You can perform these actions only for VM managers that are managed in central mode.

To add or update VM managers, the following requirements must be met:

- You must have the Manage VM Managers and Servers permission
- The computer group to which you are assigned must contain the computer on which the central VM Manager Tool is installed

Table 140. Operation descriptions

The table consists of two columns and 15 rows.

Operation details	Description
Operation	<code>PUT api/sam/vmmanagers</code>
Purpose	Adds a new VM manager or changes an existing one.
HTTP method	<code>PUT</code>
Resource URI	<code>https://server_host_name:port_number/api/sam/vmmanagers?token=token</code>
Request Content-Type	<code>application/json</code>

Table 140. Operation descriptions

The table consists of two columns and 15 rows.

(continued)

Operation details	Description
Response Content-Type	application/json
Normal HTTP response codes	200 - OK
Error HTTP response codes	<p>400 - "Bad Request" if a query parameter contains errors or is missing</p> <p>401 - "Unauthorized" if the user does not have the Manage VM Managers and Servers permission</p> <p>403 - "Forbidden" if the computer where the central VM Manager Tool is installed does not belong to the computer group that the user is assigned</p> <p>500 - "Internal Server Error" if no VM Manager Tool is defined in the BigFix Inventory database. The code might also indicate an unknown problem that requires contacting BigFix support.</p>

Request attributes

Table 141. Request attributes

Attribute	Description
login	User name that is used to access the VM manager.
password	Password that is used to access the VM manager. The password is encrypted after it is sent.
protocol	<p>Communication protocol used by the VM manager. The attribute is applicable only for Hyper-V.</p> <ul style="list-style-type: none"> • 0 - WinRM • 1 - PowerShell
sharecredentials	Indicates whether the VM manager shares credentials with other hosts in the same cluster. The attribute is applicable only for Hyper-V.

Table 141. Request attributes (continued)

Attribute	Description
	<ul style="list-style-type: none"> • <code>true</code> - the VM manager shares the credentials and it is not necessary to specify them • <code>false</code> - the VM manager does not share the credentials
<code>type</code>	<p>Type of the VM manager.</p> <ul style="list-style-type: none"> • 1 - Hyper-V • 2 - vCenter • 3 - RHEV-M • 9.2.12 4 - XenServer or Citrix XenServer
<code>url</code>	<p>Web address of the VM manager. You can specify a full URL, a host name, or an IP address. The default URL differs depending on the virtualization type.</p> <ul style="list-style-type: none"> • vCenter <ul style="list-style-type: none"> ◦ For versions up to 5.1: <code>https://<vCenter_IP_address>/sdk</code> ◦ For version 5.5 and higher: <code>https://<vCenter_IP_address>/sdk/vimService.wsdl</code> • RHEV-M <ul style="list-style-type: none"> ◦ For version 3.0: <code>https://<RHEV-M_IP_address>:8443/api</code> ◦ For version 3.1 and higher: <code>https://<RHEV-M_IP_address>/api</code> ◦ For version 4.0: <code>https://<RHEV-M_IP_address>/ovirt-engine/api</code> • Hyper-V <ul style="list-style-type: none"> ◦ <code>https://<Hyper-V_IP_address>/wsman</code> • 9.2.12 XenServer or Citrix XenServer <ul style="list-style-type: none"> ◦ <code>https://<xen-server></code> • 9.2.14 Oracle VM Server for x86 <ul style="list-style-type: none"> ◦ <code>https://<OVM-Manager_IP_address>:7002/ovm/core/wsapi/rest</code> • 9.2.17 Nutanix <ul style="list-style-type: none"> ◦ <code>https://<NUTANIX_PRISM_IP_ADDRESS>:9440/PrismGateway/services/rest/v2.0</code>

Table 141. Request attributes (continued)

Attribute	Description
<code>vmmanagerid</code>	ID of the VM manager that you want to change. To obtain the ID, use REST API for the retrieval of VM managers (on page cmxlx) .
<code>vmmanagertoolid</code>	ID of the VM Manager Tool that collects data from the VM manager that you want to change. To obtain the ID, use REST API for the retrieval of VM managers (on page cmxlx) .

Example HTTP conversation - adding a VM manager

To add a new VM manager, provide the following attributes:

- `login`
- `password`
- `url`
- `type`

Request

```
PUT http://localhost:9081/api/sam/vmmanagers
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
[ {
  "login": "user1",
  "password": "vZaW1Ve6",
  "url": "192.0.2.24",
  "type": 2
}]
```

Response

```
200 OK
```

Example HTTP conversation - changing a VM manager

To change an existing VM manager, provide the following attributes:

- `login`
- `password`
- `url`
- `type`
- `vmmanagerid`
- `vmmanagertoolid`

! **Important:** All attributes must be provided in the request regardless of how many of the attributes you want to change.

Request

```
PUT http://localhost:9081/api/sam/vmmanagers
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
[ {
  "vmmanagerid":1,
  "vmmanagertoolid":1,
  "login":"user1",
  "password":"vZaW1Ve6",
  "url":"192.0.2.24",
  "type":3
}]
```

Response

```
200 OK
```

REST API for changing advanced server settings

9.2.3 Available from 9.2.3. You can change the parameters of the BigFix Inventory server by using REST API.

Retrieval of the current values of advanced server settings

9.2.3 Available from 9.2.3. You use the `GET` operation on the `api/sam/configs` element to request information about the current values of the BigFix Inventory server settings. The output is returned in the JSON format.

Table 142. Operation descriptions

The table consists of 2 columns and 15 rows.

Operation details	Description
Operation	<code>GET /api/sam/configs</code>
Purpose	Returns information about the current values of the BigFix Inventory server settings.
HTTP method	<code>GET</code>
Resource URI	<code>https://server_host_name:port_number/api/sam/configs</code>
URL link relation	n/a
URL query parameters	For a list of query parameters, see: Query parameters (on page cmlviii) .

Table 142. Operation descriptions

The table consists of 2 columns and 15 rows.

(continued)

Operation details	Description
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	<code>Application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p>Application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	<code>configs</code> element
Response <code>Content-Type</code>	<code>Application/json</code>
Normal HTTP response codes	n/a
Error HTTP response codes	n/a

Query parameters

The following table presents query parameters that you can use for the `configs` element.

Table 143. Query parameters for retrieving current values of server settings

The table consists of four columns and two rows.

Parameter	Description	Required	Value
name	Returns the configuration of a single parameter. For a complete list of possible values and their explanation, see: Advanced server settings (on page cdii) .	No	Name of the parameter
token	A unique user authentication identifier.	Yes	Alphanumeric

Example HTTP conversation

Checking the current value of all parameters.

Request

```
GET https://localhost:9081/api/sam/configs?token=
7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Excerpt for the response body (JSON)

```
[
  {
    "valueMax":10080,
    "valueMin":0,
    "valueType":"minutes",
    "name":"vmman_transfer_period","value":"720"},
  {
    "valueType":"boolean",
    "name":"vmman_check_uniqueness_enabled",
    "value":"true"},
  {
    "valueType":"boolean",
    "name":"blockUiBundlingComputations",
    "value":"false"},
  {
    "valueMax":10080,
    "valueMin":0,
    "valueType":"minutes",
    "name":"catalog_download_task_stagger_interval",
    "value":"0"},
```

```
{ "valueType": "boolean",
  "name": "storeHwDataForAllVMMManagerNodes",
  "value": "false" },
]
```

Checking the current value of a single parameter.

Request

```
GET https://localhost:9081/api/sam/configs?token=
7adc3efb175e2bc0f4484bdd2efca54a8fa04623&name=maxVMMManagerInactivity
```

Response body (JSON)

```
[
  {
    "valueMax": 90,
    "valueMin": 1,
    "valueType": "days",
    "name": "maxVMMManagerInactivity",
    "value": "3"
  }
]
```

Configuration of advanced server settings

9.2.3 Available from 9.2.3. You use the `PUT` operation on the `api/sam/configs` element to change the values of the BigFix Inventory server settings.

 You must be an Administrator to perform this task.

Table 144. Operation descriptions

The table consists of two columns and 15 rows.

Operation details	Description
Operation	<code>PUT /api/sam/configs</code>
Purpose	Changes the current value of a server setting.
HTTP method	<code>PUT</code>
Resource URI	<code>https://server_host_name:port_number/api/sam/configs</code>
URL link relation	n/a
URL query parameters	For a list of query parameters, see: Query parameters (on page cm1x) .
Request headers	Header

Table 144. Operation descriptions

The table consists of two columns and 15 rows.

(continued)

Operation details	Description
	<code>Accept-Language</code> (optional) Values en-US (only English is supported) Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.
Request payload	n/a
Request <code>Content-Type</code>	<code>application/x-www-form-urlencoded</code>
Response headers	Header <code>Content-Language</code> Values en-US, ... Specifies the language of the response content. If this header is not specified, the content is returned in the server language.
Response payload	n/a
Response <code>Content-Type</code>	n/a
Normal HTTP response codes	204
Error HTTP response codes	<code>400 - "Bad Request"</code> if the provided value is invalid. For example, it is out of range or has a wrong type. <code>404 - "Not Found"</code>

Query parameters

The following table presents query parameters that you can use for the `configs` element.

Table 145. Query parameters for changing current values of server settings

The table consists of four columns and three rows.

Parameter	Description	Required	Value
name	Specifies the parameter whose value is to be changed. For a complete list of possible values and their explanation, see: Advanced server settings (on page cdi) .	Yes	Name of the parameter
token	A unique user authentication identifier.	Yes	Alphanumeric
value	Specifies the new value of the parameter.	Yes	Value that is within the range specific for the parameter.

Example HTTP conversation

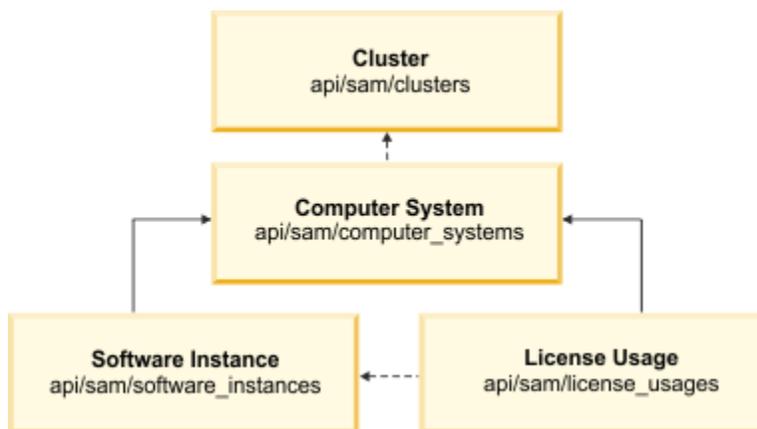
Request

```
PUT http://localhost:9981/api/sam/configs
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623&name=
maxVMManagerInactivity&value=30
```

REST API for retrieving computer systems, clusters, software instances, and license usage

You can use this REST API to quickly retrieve large amounts of data related to your computer systems, clusters, software instances, and license usage. Once the data is retrieved, it can be passed to other applications for further processing and analysis.

The following image represents the relations between all resources included in this REST API:



Scenarios

Choose one of the scenarios to learn how to determine the license usage for all of your computer systems or only for the selected ones. The first approach is recommended if you want to retrieve your data in bulk and calculate the license usage for the whole environment. The second one, however, allows you to quickly target specific information by retrieving data for a chosen subset of computer systems.

Determining license usage for all computer systems

You can extract your data in bulk to determine total license usage for software on all your computer systems. Instead of using reports in the application user interface, you can make API requests to quickly retrieve large amounts of data.

- Each API request must be authenticated with the **token** parameter. You can retrieve it by using [REST API for retrieving authentication token \(on page dccclxxii\)](#). You can also log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.
- To obtain the list of query parameters to narrow down your results and the list of available columns, see the information related to a particular API type.

Performance considerations

- Retrieving large amounts of data might impact the performance of your environment, therefore API should not be used together with other performance intensive tasks, like data imports. For more information, see [REST API considerations \(on page dcccli\)](#).
- Retrieve your data in pages rather than query for each computer ID separately. You can make several API requests and use the **limit** (on page cmlxxi) and **offset** (on page cmlxxi) parameters to paginate your results.



Note:

- For environments with approximately 200 000 endpoints, it is recommended to retrieve your data in pages of 100 000 rows for computer systems, 200 000 rows for software instances, and 300 000 rows for license usage. If you limit the first request to 100 000 results, append the next request with the **offset=100000** parameter to omit the already retrieved records. Adjust the values according to the size of your deployment.
- If you have a small number of endpoints, you can omit the **limit** and **offset** parameters, and retrieve your data by using only one API request.

1. Retrieve the list of your computer systems:

```
https://hostname:port/api/sam/computer_systems?
token=token&limit=100000&offset=100000
```

Result: Each computer system has a different **id**. You can later use this **id** to create a match between your software instances and computer systems.

```

{
  "id": 182,
  "parent_id": 175,
  "type": "virtual",
  "os": "Win2008R2 6.1.7601",
  "host_name": "NC9128109187",
  "dns_name": "NC9128109187",
  "ip_address": [
    "9.128.109.187"],
  "last_seen": "2014-06-06T03:56:39Z",
  "hardware_manufacturer": "-",
  "hardware_model": "-",
  "hardware_serial_number": "TLM_VM_4236ac43",
  "processor_type": "Multi-core",
  "processor_brand": "Xeon(R), 3 or 4 Socket",
  "processor_vendor": "Intel(R)",
  "processor_model": "E3-12xx, E7-28xx, E7-48xx",
  "partition_cores": "1.0",
  "server_processors": 1,
  "server_cores": 1
}

```

2. Retrieve the list of your software instances. The `software_title_name` column that allows you to recognize the name of your software is hidden by default, which means that you have to append the URL with the `columns[]` parameter followed by the name of a hidden column.

The following example retrieves the `computer_system_id` and `software_title_name` columns so that you can recognize which software is installed on which computer system. If you want to retrieve complete information, append the URL with the names of all columns. You can find the column names in [response body \(on page cmlxxxv\)](#):

```

https://hostname:port/api/sam/software_instances?
token=token&limit=100000&offset=100000&
columns[]=computer_system_id&columns[]=catalog_dimension.software_title_name

```

Result: Each software instance contains a `computer_system_id` column that represents an ID of a computer that a particular instance is installed on. Use this column to match your software instances with computer systems. For example, if rows 152-155 contain a `computer_system_id=182` column, it means that all those software instances are installed on a computer system with ID 182.

```

{
  "computer_system_id": 182,

```

```
"catalog_dimension": {
  "software_title_name": "BigFix Inventory"
}
```

- Retrieve the license usage information. Append the URL with the `software_title_dimension.name` column to be able to recognize the name of the software that the license usage is presented for.

The following example retrieves the `computer_system_id`, `metric_name`, `peak_value`, and `software_title_dimension.name` columns. If you want to retrieve complete information, append the URL with the names of all columns. You can find the column names in [response body \(on page cmxci\)](#):

```
https://hostname:port/api/sam/license_usages?
token=token&limit=100000&offset=100000&
columns[]=computer_system_id&columns[]=metric_name&columns[]=peak_value&
columns[]=software_title_dimension.name
```

Result: Each record contains a `computer_system_id` column that represents an ID of a computer for which the license usage is calculated. Use this column to match the license usage with computer systems. For example, if rows 152-155 contain a `computer_system_id=182` column, it means that all those license usage records are presented for a computer system with ID 182. The `peak_value` column represents the peak license usage (over last 90 days) for a particular software title (which is described by `software_title_dimension.name`). The `metric_name` column allows you to recognize whether the license type is PVU or RVU, full or subcapacity.

```
{
  "computer_system_id": 182,
  "metric_name": "RVU_FULL_CAP",
  "peak_value": 2,
  "software_title_dimension": {
    "name": "BigFix Inventory"
  }
},
{
  "computer_system_id": 182,
  "metric_name": "RVU_SUB_CAP",
  "peak_value": 2,
  "software_title_dimension": {
    "name": "BigFix Inventory"
  }
}}
```

- Determine the total license usage for a software title by summing up the values of all `peak_value` columns retrieved for this software title from all your computer systems. For example, sum up all `peak_values` for BigFix Inventory on all computer systems that contain entries for this particular software. Do not combine the metric types, but calculate the `PVU_FULL_CAP`, `RVU_FULL_CAP`, `PVU_SUB_CAP`, and `RVU_SUB_CAP` separately.

Determining license usage for selected computer systems

You can narrow down the results of your API requests to retrieve data only from selected computer systems. This approach is recommended if you want to quickly target specific information.

- Each API request must be authenticated with the **token** parameter. You can retrieve it by using [REST API for retrieving authentication token \(on page dccclxxii\)](#). You can also log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.
- To obtain the list of query parameters to narrow down your results and the list of available columns, see the information related to a particular API type.

1. Retrieve selected computer systems by querying for their ID:

```
https://hostname:port/api/sam/computer_systems?
token=token&criteria={"or":[{"id","=","1"}, {"id","=","2"}]}
```

Result: Each computer system has a different `id`. You can later use this `id` to create a match between your software instances and computer systems. In API for software instances and license usage, this `id` is represented by the `computer_system_id` column.

```
{
  "id": 2,
  "parent_id": 175,
  "type": "virtual",
  "os": "Win2008R2 6.1.7601",
  "host_name": "NC9128109187",
  "dns_name": "NC9128109187",
  "ip_address": [
    "9.128.109.187"],
  "last_seen": "2014-06-06T03:56:39Z",
  "hardware_manufacturer": "-",
  "hardware_model": "-",
  "hardware_serial_number": "TLM_VM_4236ac43",
  "processor_type": "Multi-core",
  "processor_brand": "Xeon(R), 3 or 4 Socket",
  "processor_vendor": "Intel(R)",
  "processor_model": "E3-12xx, E7-28xx, E7-48xx",
  "partition_cores": "1.0",
  "server_processors": 1,
  "server_cores": 1
}
```

2. Retrieve software instances for selected computer systems by querying for their ID:

The following example retrieves the `computer_system_id` and `software_title_name` columns so that you can recognize which software is installed on which computer system. If you want to retrieve complete information, append the URL with the names of all columns. You can find the column names in [response body \(on page cmlxxiii\)](#):

```
https://hostname:port/api/sam/software_instances?
token=token&
criteria={"or":[{"computer_system_id","=","1"},{"computer_system_id","=","2"}]}
&columns[]=computer_system_id&columns[]=catalog_dimension.software_title_name
```

Result: Each software instance contains a `computer_system_id` column that represents an ID of a computer that a particular instance is installed on. Use this column to match your software instances with computer systems. For example, if rows 152-155 contain a `computer_system_id=2` column, it means that all those software instances are installed on a computer system with ID 2.

```
{
  "computer_system_id": 2,
  "catalog_dimension": {
    "software_title_name": "BigFix Inventory"
  }
}
```

- Retrieve the license usage for selected computer systems by querying for their ID. Append the URL with the `software_title_dimension.name` column to be able to recognize the name of the software that the license usage is presented for.

The following example retrieves the `computer_system_id`, `metric_name`, `peak_value`, and `software_title_dimension.name` columns. If you want to retrieve complete information, append the URL with the names of all columns. You can find the column names in [response body \(on page cmlxxiii\)](#):

```
https://hostname:port/api/sam/license_usages?
token=token&
criteria={"or":[{"computer_system_id","=","1"},{"computer_system_id","=","2"}]}
&columns[]=computer_system_id&columns[]=metric_name&columns[]=peak_value&
columns[]=software_title_dimension.name
```

Result: Each record contains a `computer_system_id` column that represents an ID of a computer for which the license usage is calculated. Use this column to match the license usage with computer systems. For example, if rows 152-155 contain a `computer_system_id=2` column, it means that all those license usage records are presented for a computer system with ID 2. The `peak_value` column represents the peak license usage (over last 90 days) for a particular software title (which is described by `software_title_dimension.name`). The `metric_name` column allows you to recognize whether the license type is PVU or RVU, full or subcapacity.

```
{
  "computer_system_id": 2,
```

```

"metric_name": "RVU_SUB_CAP",
"peak_value": 2,
"software_title_dimension": {
  "name": "BigFix Inventory"
}}

```

You retrieved the list of software instances that are installed on a particular computer system, as well as the license usage that they generate. This information is only a part of the total license usage generated on all your computer systems. To determine the total license usage for your environment, see [Determining license usage for all computer systems \(on page cmlxii\)](#).

Retrieval of computer systems

You use the `GET` operation on the `api/sam/computer_systems` element to request information about computer systems in your infrastructure.



Important: This REST API is planned to be entirely substituted with [api/sam/v2/computers \(on page mxxx\)](#) REST API. If you have custom tools or internal processes that are based on `api/sam/computer_systems` REST API, it is best to adjust the tools to use the new version of the API. For information, see: [Mapping columns between computer_systems and v2/computers REST API \(on page dccclxviii\)](#).

This API retrieves information about physical and virtual computer systems. Apart from computers that have an BigFix client installed, it also includes host systems on which the client is not deployed because some virtualization types do not allow for software to be installed on the host level (VMware ESXi, IBM PowerVM, and so on). Data related to such host computer systems must however be retrieved to build a proper hierarchy between the server and its virtual machines and to correctly report the PVU and RVU utilization.

The `type` property determines whether a computer system is a host computer or a virtual one. A host can be, for example, a virtualization server, ESX hypervisor, or a laptop with agent installed. A virtual machine refers only to virtual machines that can be deployed on a host. Because of this, you should not compare the outcome of the REST API with single reports, especially with the Computers, and Hardware Inventory reports.

To retrieve information about your computer systems, use the following URL:

```
https://hostname:port/api/sam/computer_systems?token=token
```



Important:



- To use this API, you must be assigned to the All Computers group and have the View Hardware Inventory permission.
- Each API request must be authenticated with the **token** parameter. You can retrieve it by using [REST API for retrieving authentication token \(on page dccclxxii\)](#). You can also log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.
- By default, the retrieved data is sorted by `id`.

Table 146. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/sam/computer_systems</code>
Purpose	Returns a list of computer systems.
HTTP method	<code>GET</code>
Resource URI	<code>https://server_host_name:port_number/api/sam/computer_systems</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmlxx) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p>

Table 146. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
	Specifies the content type of the response.
	Header
	<code>Content-Language</code>
	Values
	en-US, ...
	Specifies the language of the response content. If this header is not specified, the content is returned in the server language.
	Header
	<code>Import-Mode</code>
	Values
	<code>none, idle, running, pending</code>
	Specifies the data import status.
	Header
	<code>Import-Progress</code>
	Values
	A percentage value, for example 59.
	Specifies the percentage progress of a data import.
	Header
	<code>Import-Last-Status</code>
	Values
	<code>successful, failed</code>
	Specifies the status of the last data import.
	Header
	<code>Import-Last-Success-Time</code>
	Values

Table 146. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
	Date, for example <code>2014-06-18T04:00:29Z</code> .
	Specifies the time of the last successful data import.
Response payload	<code>Computer Systems</code> element
Response Content-Type	<code>application/json</code>
Normal HTTP response codes	<code>200 - OK</code>
Error HTTP response codes	<code>500 - "Bad Request"</code> if a query parameter contains errors or is missing
	Message body includes an error message with details.

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `api/sam/computer_systems` element.

Table 147. Query parameters for retrieving software products

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, a set of default columns is retrieved. For available columns, see Response body (on page cmlxxiii) . Example: Retrieve the name and version columns: <pre>URL?columns[]=name&columns[]=version</pre>	No	String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you	No	String

Table 147. Query parameters for retrieving software products

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<p>want to specify a descending sort, append <code>desc</code> to the column name. Example:</p> <p>Order by type descending:</p> <pre>URL?order[]=type desc</pre>		
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved.	No	Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example:	No	Numeric
	<p>Retrieve 50 records starting after record 150:</p> <pre>URL?limit=50&offset=150</pre>		
token	A unique user authentication identifier. You can view your token in the Profile preferences of BigFix Inventory.	Yes	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written on one line:		
	<pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criteria> [{ <comma> <criteria> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criteria> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null></pre>		

For more information about operators, see [Common connectors and operators \(on page dcccixii\)](#).

Example: Retrieve computer systems whose operating system contains "AIX" OR the last seen date within a specific date range:

Table 147. Query parameters for retrieving software products

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<pre>URL?criteria={ "or": [["os", "contains", "aix"], { "and": [["last_seen", ">", "1970-01-01T00:00:00+00:00Z"], ["last_seen", "<", "1970-01-02T00:00:00+00:00Z"]] }] }</pre>		
	<p>For columns that use the date and time values, such as Last Seen, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <operator>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively. For example, to retrieve computer systems that reported within last 7 days, use the following API request:</p>		
	<pre>URL?criteria={"and":[["last_seen", "last", "P7D"]]}</pre>		

Example HTTP conversation

Request

```
GET api/sam/computer_systems
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

The above response header may be followed by different entries, depending on the current status of the data import which is represented by `Import-Mode`. To understand the returned values, see the following definitions:

- `none` - a data import has never been initiated
- `idle` - no data imports are currently running
- `running` - a data import is in progress
- `pending` - an action performed in the user interface requires a data import to be started for the change to take effect

If a data import has never been initiated:

```
Import-Mode: none
```

If a data import is in progress:

```
Import-Mode: running
Import-Progress: 41
```

If a data import is not running:

```
Import-Mode: idle/pending
Import-Last-Status: successful
Import-Last-Success-Time: Mon, 23 Jun 2014 12:18:29 GMT
```

Response body (JSON)

If a particular entry is hidden by default, it is not retrieved by using the general URL. To retrieve such data, you must use query parameters to specify the name of the hidden column. For example, you can retrieve the `server_id` and `datasource_id` columns by using the `columns[]` parameter:

```
URL?columns[]=server_id&columns[]=datasource_id
```

```
{
  "id": 25,
  "parent_id": 9, //ID of a host for a VM
  "computer_id": 2, //hidden by default
  "computer_remote_id": 123, //hidden by default
  "server_id": 24, //hidden by default
  "datasource_name": "Data source", //hidden by default
  "datasource_id": 1, //hidden by default
  "type": "virtual", //type: virtual or host
  "os": "Linux Red Hat Enterprise Server 6.2)",
  "host_name": "NC040221",
  "dns_name": "NC040221.kraklab.pl.ibm.com",
  "ip_address": [
    "9.167.40.221",
    "192.168.122.1" ],
  "last_seen": "2014-04-08T14:33:41Z",
  "hardware_manufacturer": "IBM",
  "hardware_model": "System x3550 M2 -[794662G]-",
  "hardware_serial_number": "99B7166",
  "hardware_type": "7946", //hidden by default
  "hardware_name": "IBM Corp. 7946", //hidden by default
  "processor_brand_string": "Intel(R) Xeon(R) CPU E7-3400 @ 2.40GHz",
```

```

"processor_type": "Multi-core",
"processor_brand": "Xeon(R)",
"processor_vendor": "Intel(R)",
"processor_model": "3400-3699 or 5500-5699",
"partition_cores" : 1,           //null for host serves
"server_processors": 1,
"server_cores": 8,
"pvu_per_core": 70,             //hidden by default
"uuid": "50305bd3-1f09-7294-7033-b903767d4605" //hidden by default
"cluster_id": "1"               // hidden by default
} //ID of the cluster to which the computer system
belongs. Available only for type: host.

```

Retrieval of clusters

You use the `GET` operation on the `api/sam/clusters` element to request information about clusters that group your host computer systems. This information represents the current state of your environment and does not include any historical data.

After you retrieve the IDs of your clusters, compare them with the `cluster_id` column retrieved by [API for computer systems \(on page cmlxvii\)](#) to recognize which of these systems belong to which clusters. This information can also be viewed on the Hardware Inventory report in the user interface.

To retrieve information about clusters, use the following URL:

```
https://hostname:port/api/sam/clusters?token=token
```

Important:

- To use this API, you must be assigned to the All Computers group and have the View Hardware Inventory permission.
- Each API request must be authenticated with the **token** parameter. You can retrieve it by using [REST API for retrieving authentication token \(on page dcclxxii\)](#). You can also log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.
- By default, the retrieved data is sorted by `id`.

Table 148. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/sam/clusters</code>
Purpose	Returns a list of clusters.
HTTP method	GET
Resource URI	<code>https://server_host_name:port_number/api/sam/clusters</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmlxx) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p>

Table 148. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
	<p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p> <p>Header</p> <p><code>Import-Mode</code></p> <p>Values</p> <p><code>none, idle, running, pending</code></p>
	<p>Specifies the data import status.</p> <p>Header</p> <p><code>Import-Progress</code></p> <p>Values</p> <p>A percentage value, for example <code>59</code>.</p>
	<p>Specifies the percentage progress of a data import.</p> <p>Header</p> <p><code>Import-Last-Status</code></p> <p>Values</p> <p><code>successful, failed</code></p>
	<p>Specifies the status of the last data import.</p> <p>Header</p> <p><code>Import-Last-Success-Time</code></p> <p>Values</p> <p>Date, for example <code>2014-06-18T04:00:29Z</code>.</p>
Response payload	<code>Clusters</code> element
Response Content-Type	<code>application/json</code>

Table 148. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Normal HTTP response codes	
	200 - OK
Error HTTP response codes	
	500 - "Bad Request" if a query parameter contains errors or is missing
	Message body includes an error message with details.

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `api/sam/clusters` element.

Table 149. Query parameters for retrieving clusters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, a set of default columns is retrieved. For available columns, see Response body (on page cmlxxix) . Example: Retrieve the name and id columns: <pre>URL?columns[]=name&columns[]=id</pre>	No	String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by type descending: <pre>URL?order[]=type desc</pre>	No	String
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved.	No	Numeric

Table 149. Query parameters for retrieving clusters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150: <code>URL?limit=50&offset=150</code>	No	Numeric
token	A unique user authentication identifier. You can view your token in the Profile preferences of BigFix Inventory.	Yes	Alphanumeric

Example HTTP conversation

Request

```
GET api/sam/clusters
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

The above response header may be followed by different entries, depending on the current status of the data import which is represented by `Import-Mode`. To understand the returned values, see the following definitions:

- `none` - a data import has never been initiated
- `idle` - no data imports are currently running
- `running` - a data import is in progress
- `pending` - an action performed in the user interface requires a data import to be started for the change to take effect

If a data import has never been initiated:

```
Import-Mode: none
```

If a data import is in progress:

```
Import-Mode: running
Import-Progress: 41
```

If a data import is not running:

```
Import-Mode: idle/pending
Import-Last-Status: successful
Import-Last-Success-Time: Mon, 23 Jun 2014 12:18:29 GMT
```

Response body (JSON)

```
{
  "id": 1,
  "name": "Production Cluster 1",
  "cores_count": 24,
  "vm_manager_url": "https://198.51.100/sdk",
}
```

Retrieval of software instances

You use the `GET` operation on the `api/sam/software_instances` element to request information about software installed in your infrastructure.

! **Important:** This REST API is planned to be entirely substituted with `api/sam/v2/software_instances` (on page [cmxcii](#)) REST API. If you have custom tools or internal processes that are based on `api/sam/software_instances` REST API, it is best to adjust the tools to use the new version of the API. For information, see: [Mapping columns between software_instances and v2/software_instances REST API \(on page dccclxvi\)](#).

To retrieve information about software instances installed on your computer systems, use the following URL:

```
https://hostname:port/api/sam/software_instances?token=token
```

! **Important:**

- To use this API, you must be assigned to the All Computers group and have the View Endpoints permission.
- Each API request must be authenticated with the **token** parameter. You can retrieve it by using [REST API for retrieving authentication token \(on page dccclxxii\)](#). You can also log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.
- By default, the retrieved data is sorted by `id`.

Table 150. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	GET /api/sam/software_instances
Purpose	Returns a list of software.
HTTP method	GET
Resource URI	<code>https://server_host_name:port_number/api/sam/software_instances</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmlxxxii) .
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request Content-Type	application/json
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p>

Table 150. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
	<p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p> <p>Header</p> <p><code>Import-Mode</code></p> <p>Values</p> <p><code>none, idle, running, pending</code></p>
	<p>Specifies the data import status.</p> <p>Header</p> <p><code>Import-Progress</code></p> <p>Values</p> <p>A percentage value, for example <code>59</code>.</p>
	<p>Specifies the percentage progress of a data import.</p> <p>Header</p> <p><code>Import-Last-Status</code></p> <p>Values</p> <p><code>successful, failed</code></p>
	<p>Specifies the status of the last data import.</p> <p>Header</p> <p><code>Import-Last-Success-Time</code></p> <p>Values</p> <p>Date, for example <code>2014-06-18T04:00:29Z</code>.</p>
Response payload	<p><code>Software_Instances</code> element</p>
Response Content-Type	<p><code>application/json</code></p>

Table 150. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Normal HTTP response codes	
	200 - OK
Error HTTP response codes	
	500 - "Bad Request" if a query parameter contains errors or is missing
	Message body includes an error message with details.

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `api/sam/software_instances` element.

Table 151. Query parameters for retrieving software products

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, a set of default columns is retrieved. Example: Retrieve the name and version columns: <pre>URL?columns[]=name&columns[]=version</pre>	No	String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by computer system ID descending: <pre>URL?order[]=computer_system_id desc</pre>	No	String
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved.	No	Numeric

Table 151. Query parameters for retrieving software products

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150: <pre>URL?limit=50&offset=150</pre>	No	Numeric
token	A unique user authentication identifier. You can view your token in the Profile preferences of BigFix Inventory.	Yes	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written on one line: <pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criterion> [{ <comma> <criterion> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criterion> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null></pre>		

For more information about operators, see [Common connectors and operators \(on page dcccxi\)](#).

Example: Retrieve software instances whose publisher name contains "IBM" OR the last seen date within a specific date range:

```
URL?criteria={ "or": [ ["publisher_name", "contains", "IBM"],
{ "and": [ ["last_seen", ">", "1970-01-01T00:00:00+00:00Z"],
["last_seen", "<", "1970-01-02T00:00:00+00:00Z"] ] } ] }
```

For columns that use the date and time values, such as Last Seen, you can retrieve data also for a period instead of a specific date. To do so, use *last* or *next* as `<operator>`, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively. For example, to retrieve software instances that reported within last 7 days, use the following API request:

Table 151. Query parameters for retrieving software products

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<code>URL?criteria={"and":[{"last_seen","last","P7D"}]}</code>		

Example HTTP conversation

Request

```
GET api/sam/software_instances
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

The above response header may be followed by different entries, depending on the current status of the data import which is represented by `Import-Mode`. To understand the returned values, see the following definitions:

- `none` - a data import has never been initiated
- `idle` - no data imports are currently running
- `running` - a data import is in progress
- `pending` - an action performed in the user interface requires a data import to be started for the change to take effect

If a data import has never been initiated:

```
Import-Mode: none
```

If a data import is in progress:

```
Import-Mode: running
Import-Progress: 41
```

If a data import is not running:

```

Import-Mode: idle/pending
Import-Last-Status: successful
Import-Last-Success-Time: Mon, 23 Jun 2014 12:18:29 GMT

```

Response body (JSON)

If a particular entry is hidden by default, it is not retrieved by using the general URL. To retrieve such data, you must use query parameters to specify the name of the hidden column. For example, you can retrieve the `computer_id` and `software_title_name` columns by using the `columns[]` parameter:

```
URL?columns[]=computer_id&columns[]=catalog_dimension.software_title_name
```

```

{
  "id": 123,
  "software_fact_id": 123, //hidden by default
  "computer_system_id": 3,
  "computer_id": 3, //hidden by default
  "discoverable_guid": "0768fb15-383c-4124-a7e2-0d76dda06874",
  "default_product_guid": "78d380e0-9fb9-11e3-a151-005056872dc7", //hidden by default
  "first_used": null,
  "last_used": null,
  "valid_from": "2014-04-02T14:24:04Z",
  "valid_to": "9999-12-31T23:59:59Z",
  "updated_at": "2014-04-02T14:24:04Z",
  "signature_count": 1,
  "total_time": 0,
  "total_runs": 0,
  "avg_run_time": null,
  "avg_runs_per_day": null,
  "process": null,
  "deleted": false, //hidden by default
  "catalog_dimension": //hidden by default
  {
    "software_title_name": "BigFix Platform Agent",
    "publisher_name": "IBM",
    "software_title_version_name": "BigFix Platform Agent",
    "software_title_release_name": "BigFix Platform Agent",
    "version": "9.0"
  }
}

```

Retrieval of license usage

You use the `GET` operation on the `api/sam/license_usages` element to request information about license usage reported by your computer systems.

This API retrieves peak license usage over the last 90 days. PVU and RVU full capacity and subcapacity license usage is retrieved only on the host computer system level. To retrieve information about license usage reported by your computer systems, use the following URL:

```
https://hostname:port/api/sam/license_usages?token=token
```



Important:

- To use this API, you must be assigned to the All Computers group and have the View License Metrics permission.
- Each API request must be authenticated with the **token** parameter. You can retrieve it by using [REST API for retrieving authentication token \(on page dcccclxxii\)](#). You can also log in to BigFix Inventory, hover over the **User** icon , and click **Profile**. Then, click **Show token**.
- This API retrieves data for the last 90 days, until the last successful import.
- By default, the data retrieves is sorted by `computer_system_id`, `software_title_id`, and `metric_name` columns.

Table 152. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/sam/license_usages</code>
Purpose	Returns information about license usage
HTTP method	<code>GET</code>
Resource URI	<code>https://server_host_name:port_number/api/sam/license_usages</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmlxxxviii) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p>

Table 152. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
	<p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	n/a
Request <code>Content-Type</code>	<p><code>application/json</code></p>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p> <p>Header</p> <p><code>Import-Mode</code></p> <p>Values</p> <p>none, idle, running, pending</p> <p>Specifies the import status.</p> <p>Header</p> <p><code>Import-Progress</code></p> <p>Values</p> <p>A percentage value, for example 59.</p>

Table 152. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
	<p>Specifies the percentage progress of an import.</p> <p>Header</p> <p><code>Import-Last-Status</code></p> <p>Values</p> <p>successful, failed</p>
	<p>Specifies the status of the last import.</p> <p>Header</p> <p><code>Import-Last-Success-Time</code></p> <p>Values</p> <p>Date, for example <code>2014-06-18T04:00:29Z</code>.</p>
Response payload	<p><code>License Usages</code> element</p>
Response Content-Type	<p><code>application/json</code></p>
Normal HTTP response codes	<p><code>200 - OK</code></p>
Error HTTP response codes	<p><code>500 - "Bad Request"</code> if a query parameter contains errors or is missing</p> <p>Message body includes an error message with details.</p>

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `api/sam/license_usages` element.

Table 153. Query parameters for retrieving license usage

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, a set of default columns is retrieved. Example: Retrieve the name and version columns: <pre>URL?columns[]=name&columns[]=version</pre>	No	String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by peak license usage descending: <pre>URL?order[]=peak_value desc</pre>	No	String
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved.	No	Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150: <pre>URL?limit=50&offset=150</pre>	No	Numeric
token	A unique user authentication identifier. You can view your token in the Profile preferences of BigFix Inventory.	Yes	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written on one line: <pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criterion> [{ <comma> <criterion> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criterion> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string></pre>		

Table 153. Query parameters for retrieving license usage

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<pre><operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null></pre>		
	<p>For more information about operators, see Common connectors and operators (on page dcccclxii).</p> <p>Example: Retrieve license usage for software whose publisher name contains "BigFix":</p> <pre>URL?criteria={ "and": [["publisher_name", "contains", "BigFix"]] }</pre>		

Example HTTP conversation

Request

```
GET api/sam/license_usages
?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Language: en-US
```

The above response header may be followed by different entries, depending on the current status of the data import which is represented by `Import-Mode`. To understand the returned values, see the following definitions:

- `none` - a data import has never been initiated
- `idle` - no data imports are currently running
- `running` - a data import is in progress
- `pending` - an action performed in the user interface requires a data import to be started for the change to take effect

If a data import has never been initiated:

```
Import-Mode: none
```

If a data import is in progress:

```
Import-Mode: running
Import-Progress: 41
```

If a data import is not running:

```
Import-Mode: idle/pending
Import-Last-Status: successful
Import-Last-Success-Time: Mon, 23 Jun 2014 12:18:29 GMT
```

Response body (JSON)

If a particular entry is hidden by default, it is not retrieved by using the general URL. To retrieve such data, you must use query parameters to specify the name of the hidden column. For example, you can retrieve the `software_title_dimension.name` column by using the `columns[]` parameter:

```
URL?columns[]=software_title_dimension.name
```

```
{
  "software_title_id": 36181,
  "metric_name": "RVU_FULL_CAP",
  "computer_system_id": 4,
  "peak_value": 1,
  "peak_date": "2014-04-14",
  "peak_time": "2014-04-14T13:05:10Z",
  "reporting_period_start_date": "2014-04-01",
  "reporting_period_end_date": "2014-04-30",
  "software_title_dimension": //hidden by default
  {
    "name": "BigFix Protection",
    "guid": "d78048cf-842b-44e8-8036-e7e2bf8afb31",
    "publisher_name": "BigFix",
    "publisher_guid": "8a759f0c-b91a-4d7d-8c4a-a9d85e06c13d"
  }
}
```

REST API for retrieving software inventory and metric utilization and classifying software (v2)

9.2.8 Available from 9.2.8. You can use REST API requests to quickly retrieve large amounts of data related to license metric utilization as well as your software inventory, including historical data about uninstalled software. You can then pass this data to other applications for further processing and analysis.

Retrieval of software inventory (v2)

9.2.8 Available from 9.2.8. You use the `GET` operation on the `api/sam/v2/software_instances` element to request information about software installed in your infrastructure. The API returns details of existing components as well as historical data about components that were removed.

To obtain information about detected software use `api/sam/v2/software_instances` API.

Permissions



You must have the View Endpoints and View Hardware Inventory permissions to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/software_instances?token=token
```

Resource information

Table 154. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	<code>GET</code>
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p>

Table 154. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Values	Description
	en-US, ...	Specifies the language of the response content. If the header is not specified, the content is returned in the server language.
Response payload	Software Instances	element
Response format	application/json	
Response codes	200 - OK	
	500 - "Bad Request"	if a query parameter contains errors or is missing

9.2.10 Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/software_instance.json?token=<token>
```

Available columns

Table 155. Available columns

The table consists of three columns and 26 rows.

Column	Description	Displayed by default	Type
instance_id	Identifier of the component instance.		Nu-mer-ic
computer_id	Identifier of the computer as defined in BigFix Inventory.		Nu-mer-ic
computer_bigfix_id	Identifier of the computer as defined in BigFix.		Nu-mer-ic

Table 155. Available columns

The table consists of three columns and 26 rows.

(continued)

Column	Description	Displayed by default	Type
computer_dns_name	DNS of the computer.		String
computer_group_id	Identifier of the BigFix Inventory computer group to which the computer belongs. The value 0 represents the All Computers group.		Nu-mer-ic
computer_name	Name of the computer.	✓	String
computer_ip_address	IP address of the computer.	✓	String
computer_os	Operating system of the computer.		String
computer_os_type	Type of the operating system of the computer.		String
component_name	Component name.	✓	String
component_release	Component release number.	✓	String
9.2.11 component_detailed_version	Version of the component based on the highest version of SWID tag or file that caused its detection.		String
discoverable_guid	Component GUID.		String
discoverable_family_guid	Component family GUID. For example, if the component is IBM DB2 Enterprise Server Edition, component family is IBM DB2.		String
component_publisher_name	Component publisher.		String
component_has_usage	Information whether usage data is available for the component.		Boolean
is_present	Information whether the component is still installed in the environment.		Boolean
discovery_start	Date and time when the component was reported for the first time. The time is specified in the GMT time zone.	✓	String
discovery_end	Date and time when the component was reported for the last time. The time is specified in the GMT time zone.	✓	String
discovery_path	Path under which the component is installed. Available for IBM products only. For non-IBM products, the returned value is <code>null</code> .	✓	String
product_name	Name of the product to which the component is assigned.	✓	String
product_family_guid	Product family GUID. For example, if the product is IBM DB2 Enterprise Server Edition, product family is IBM DB2.		String

Table 155. Available columns

The table consists of three columns and 26 rows.

(continued)

Column	Description	Displayed by default	Type
product_publisher_name	Product publisher.		String
product_release_name	Name of the product release to which the component is assigned.		String
product_release	Product release number.	✓	String
product_release_guid	Product release GUID.		String
metric_id	Identifier of the license metric that is used by the product. For information about the meaning of each <code>metric_id</code> , see: Metric IDs and code names (on page mxxiv) .		String
metric_code_name	Unique code name of the license metric that is used by the product. For information about the meaning of each <code>metric_code_name</code> , see: Metric IDs and code names (on page mxxiv) .		String
is_confirmed	Information whether the assignment of the component to the product is confirmed.		Boolean
is_excluded	Information whether the product is excluded from pricing calculations.		Boolean
is_suppressed	Information whether the component is suppressed on the computer on which it is installed.	✓	Boolean
exclusion_or_suppress_comment	Comment that was provided during the exclusion or suppression.		String
9.2.13 bundle_id	Identifier of the FlexPoint offering to which the software instance is bundled. Possible values include: <ul style="list-style-type: none"> - 1 - The software instance cannot be assigned to any FlexPoint offering because the component is not assigned to any product - 0 - The software instance is not bundled to any FlexPoint offering - Other values - Identifier of the FlexPoint offering to which the software instance is assigned 		String
9.2.13 bundle_name	Name of the FlexPoint offering to which the software instance is assigned. Possible values include:		String

Table 155. Available columns

The table consists of three columns and 26 rows.

(continued)

Column	Description	Displayed by default	Type
	<ul style="list-style-type: none"> • Empty column - The software instance cannot be assigned to any FlexPoint offering because the component is not assigned to any product • <code>None</code> - The software instance is not bundled to any FlexPoint offering • Other values - Name of the FlexPoint offering to which the software instance is assigned 		
9.2.13 <code>is_vulnerable</code>	Information whether any Common Vulnerabilities and Exposures (CVE) are matched to the component. Possible values include: <ul style="list-style-type: none"> • 0 - No CVEs are matched • 1 - At least one CVE is matched 		Boolean
9.2.14 <code>is_charged</code>	Information whether the component to product assignment is confirmed.		Boolean

Applicable associations

You can additionally retrieve data from the following associations:

- **9.2.12** `component_cpe` (on page [mcxv](#))
- **9.2.14** `component_tags` (on page [mcxxxiii](#))
- `computer_hardware` (on page [mcxxi](#))
- `computer_health` (on page [mcxx](#))
- **9.2.13** `cve` (on page [mcxxv](#))
- **9.2.11** `discovery_details` (on page [mcxxvii](#))
- **9.2.14** `instance_tags` (on page [mcxxxiii](#))
- **9.2.13** `release_component_lifecycle` (on page [mcxxxii](#))
- `usage_data` (on page [mcxxxiv](#))



Important: All associations provide data about items that currently exist in the infrastructure. For removed computers and uninstalled software, these associations return `null`.

Query parameters

Table 156. Query parameters

The table consists of four columns and six rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve product name and release <pre>URL?columns[]=product_name&columns[]=product_release</pre>		String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by computer ID descending <pre>URL?order[]=computer_id desc</pre>		Alphanumeric
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved. Example: Retrieve 100 records <pre>URL?limit=100</pre>		Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150 <pre>URL?limit=50&offset=150</pre>		Numeric
token	A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dccclxxii) . You can also log in to BigFix Inventory, hover over the User icon  , and click Profile . Then, click Show token .	✓	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written in one line: <pre><criteria> ::= <left-brace> <boolean-operator> <colon> <left-bracket> <criteria> [{ <comma> <criteria> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or"</pre>		String

Table 156. Query parameters

The table consists of four columns and six rows.

(continued)

Parameter	Description	Required	Value
	<pre><criteria> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-number> <json-null></pre>		

For more information about operators, see: [Common connectors and operators \(on page dccclxii\)](#).

Example: Retrieve software instances whose product name contains "BigFix" OR the discovery start is within a specific date range

```
URL?criteria={ "or": [ ["product_name", "contains",
"BigFix"],
{ "and": [ ["discovery_start", ">",
"2018-10-01T00:00:00+00:00Z"],
["discovery_start", "<",
"2018-10-02T00:00:00+00:00Z"] ] } ] }
```

For columns that use the date and time values, you can retrieve data also for a period instead of a specific date. To do so, use *last* or *next* as **<operator>**, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.

Example: Retrieve software instances which took part in the license utilization peak. The date and time of the peak is: **2018-10-10T00:00:00+00:00Z**. Query searches for the instances that meet the following criteria: product name contains "DB2" AND use the given metric AND discovery start within the peak.

For PVU, RVU MAPC and VPC

```
URL?criteria={ "and": [ ["discovery_start", "<=",
"2018-10-10T00:00:00+00:00Z"],
```

Table 156. Query parameters

The table consists of four columns and six rows.

(continued)

Parameter	Description	Required	Value
	<pre>["discovery_end", ">", "2018-10-10T00:00:00+00:00Z"], ["product_name", "contains", "DB2"], ["metric_id", "=", "5"]] }</pre>		
	<p>For other metrics</p> <pre>URL?criteria={ "and": [["discovery_start", "<", "2018-10-11T00:00:00+00:00Z"], ["discovery_end", ">=", "2018-10-10T00:00:00+00:00Z"], ["product_name", "contains", "DB2"], ["metric_id", "=", "9"]] }</pre>		
	<p>Example: Retrieve software instances that were first reported within the last 7 days</p> <pre>URL?criteria={"and":[{"discovery_start","last","P7D"}]}</pre>		

Example conversation - default columns

Request

```
GET api/sam/v2/software_instances?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "computer_name": "NC107069",
  "computer_ip_address": ["10.0.107.69"],
  "component_name": "BigFix Platform Agent",
  "component_release": "9.0",
  "is_present": 1,
  "discovery_start": "2016-04-28T15:37:26Z",
  "discovery_end": "2016-11-17T15:11:10Z",
  "discovery_path": "/opt/BESClient/bin/properties/version",
  "product_name": "BigFix for Power Management",
```

m

```
"product_release": "9.0"  
}]
```

Example conversation - all columns

Request

```
GET api/sam/v2/software_instances?columns[]=discovery_path  
&columns[]=instance_id&columns[]=computer_id&columns[]=computer_bigfix_id  
&columns[]=computer_dns_name&columns[]=computer_group_id&columns[]=computer_name  
&columns[]=computer_ip_address&columns[]=computer_os&columns[]=computer_os_type  
&columns[]=component_name&columns[]=component_release&columns[]=discoverable_guid  
&columns[]=component_publisher_name&columns[]=component_has_usage&columns[]=is_present  
&columns[]=discovery_start&columns[]=discovery_end&columns[]=product_name  
&columns[]=product_publisher_name&columns[]=product_release_name&columns[]=product_release  
&columns[]=metric_id&columns[]=metric_code_name&columns[]=is_confirmed  
&columns[]=is_excluded&columns[]=is_suppressed&columns[]=exclusion_or_suppress_comment  
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623  
  
Host: localhost:9081  
Accept: application/json  
Accept-Language: en-US
```

Response body

```
[{  
  "instance_id": 1,  
  "computer_id": 3,  
  "computer_bigfix_id": null,  
  "computer_dns_name": "NC107069",  
  "computer_group_id": 0,  
  "computer_name": "NC107069",  
  "computer_ip_address": "10.0.107.69",  
  "computer_os": "AIX 5.3",  
  "computer_os_type": "AIX",  
  "component_name": "BigFix Platform Agent",  
  "component_release": "9.0",  
  "discoverable_guid": "0768FB15-383C-4124-A7E2-0D76DDA06874",  
  "component_publisher_name": "IBM",  
  
  "component_has_usage": 1,  
  "is_present": 0,  
  "discovery_start": "2016-04-28T15:37:26Z",  
  "discovery_end": "2016-11-17T15:11:10Z",
```

```

"discovery_path":"/opt/BESClient/bin/properties/version",
"product_name":"BigFix for Power Management",
"product_publisher_name":"IBM",
"product_release_name":"BigFix for Power Management",
"product_release":"9.0",
"metric_id":"6",
"metric_code_name":"INTERNAL_RVU_MAPC",
"is_confirmed":0,
"is_excluded":0,
"is_suppressed":0,
"exclusion_or_suppress_comment":null
}]

```

Example conversation - additional column

Request

```

GET api/sam/v2/software_instances?columns[]=component_name
&columns[]=product_name&columns[]=is_confirmed
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US

```

Response body

```

[ {
  "component_name":"BigFix Platform Agent",
  "product_name":"BigFix for Power Management",
  "is_confirmed":0
} ]

```

Example conversation - association

Request

```

GET api/sam/v2/software_instances?columns[]=computer_id
&columns[]=computer_health.catalog_version
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US

```

Response body

```
[{
  "computer_id":1,
  "catalog_version":1304630
}]
```

Classifying software (v2)

9.2.14 Available from 9.2.14. You can use the `PUT` operation on the `api/sam/v2/software_instances` element to define, change, or migrate software classification.

Permissions

 You must have the View Endpoints and View Hardware Inventory permissions to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/software_instances?token=token
```

 **Important:** You can use this API to update or migrate software bundlings. However, you can only update the bundlings that are defined in the software catalog. You cannot create custom bundlings through the API.

Resource information

Table 157. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	<code>PUT</code>
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p>

Table 157. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
	<p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response payload	n/a element
Response format	<code>application/json</code>
Response codes	<p>207 - OK</p> <p>500 - "Bad Request" if a query parameter contains errors or is missing</p>

Parameters



Important:

The following parameters are required to successfully complete the query:

- **token**
- **is_charged**
- **metric_id**
- **product_release_guid**
- **instance_id**, or, alternatively, a combination of parameters including: a parameter that identifies a computer: **computer_dns_name** or **computer_bigfix_id** and a parameter that identifies software: **discoverable_guid** that can be followed by the **discovery_path** parameter for precision.

Other parameters that are listed below are optional. You can include them in your query according to your needs. The out-of-range values for all the optional parameters are ignored.

Table 158. Parameters

The table consists of three columns and a few rows.

Parameter	Description	Type
instance_id	Identifier of the component instance.	Nu-mer-ic
computer_bigfix_id	Identifier of the computer as defined in BigFix.	Nu-mer-ic
computer_dns_name	DNS of the computer.	String
discoverable_guid	Component GUID.	String
discovery_path	Path under which the component is installed. Available for BigFix products only. For non-BigFix products, the returned value is <code>null</code> .	String
product_release_guid	Product release GUID.	String
metric_id	Identifier of the license metric that is used by the product. For information about the meaning of each <code>metric_id</code> , see: Metric IDs and code names (on page mxxiv) .	String
is_confirmed	Information whether the assignment of the component to the product is confirmed.	Boolean
is_excluded	Information whether the product is excluded from pricing calculations.	Boolean
is_suppressed	Information whether the component is suppressed on the computer on which it is installed.	Boolean
exclusion_or_suppress_comment	Comment that was provided during the exclusion or suppression.	String
9.2.14 is_charged	Information whether the component to product assignment is charged.	Boolean
token	A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dccclxxii) . You can also log in to BigFix Inventory, hover over the User icon  , and click Profile . Then, click Show token .	Al-phanu-mer-ic
verbose	Descriptive information about the result of an API request. By default, this parameter is set to false.	Boolean
simulate	By adding this parameter to API request you can test the call to view its results and status. By default, this parameter is set to false.	Boolean



Note: You can change multiple values during a single request. For example, you can use a single request to reassign the component, and change its status to confirmed.

Example conversation

Request

```
PUT api/sam/v2/software_instances?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
200 - OK
```

Example conversation - migrating software assignments between BigFix servers

For more information, see: [Tutorial: Migrating software assignments between two BigFix servers \(on page dclxxxv\)](#).

Retrieval of software components (v2)

9.2.14

Available from 9.2.14. You use the `GET` operation on the `api/sam/v2/software_components` element to request information about software components in your catalog. The API returns details of existing components as well as historical data about components that were removed.

Permissions



You must have the View Software Catalog and Signatures permission to perform this task.

Resource URL

```
https://hostname:port/api/sam/v2/software_components?token=token
```

Resource information

Table 159. Resource information

Operation details	Description
HTTP method	GET
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p>

Table 159. Resource information

(continued)

Operation details	Description
	en-US (only English is supported)
	Negotiates the language of the response. If the header is not specified, the content is returned in the server language.
Request format	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response payload	<code>Software Components</code> element
Response format	<code>application/json</code>
Response codes	<p>200 - OK</p> <p>500 - "Bad Request" if a query parameter contains errors or is missing</p>

9.2.10 Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/software_component.json?token=<token>
```

Available columns

Table 160. Available columns

The table consists of three columns and 26 rows.

Column	Description	Displayed by default	Type
id	Identifier of the component instance.	✓	Nu-mer-ic
name	Name of the component.	✓	String
release	Release of the component.	✓	String
guid	Component GUID.	✓	String
is_BigFix	Information whether the component is provided by BigFix.		Boolean

Applicable associations

You can additionally retrieve data from the following associations:

- [tags](#) (on page mcxxxiii)

Query parameters

Table 161. Query parameters

The table consists of four columns and six rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve component name and id <pre>URL?columns[]=name&columns[]=id</pre>		String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by component ID descending <pre>URL?order[]=id desc</pre>		Alphanumeric

Table 161. Query parameters

The table consists of four columns and six rows.

(continued)

Parameter	Description	Required	Value
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved. Example: Retrieve 100 records <code>URL?limit=100</code>		Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150 <code>URL?limit=50&offset=150</code>		Numeric
token	A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dccclxxii) . You can also log in to BigFix Inventory, hover over the User icon  , and click Profile . Then, click Show token .	✓	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written in one line: <pre><criteria> ::= <left-brace> <boolean-operator> <colon> <left-bracket> <criterion> [{ <comma> <criterion> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criterion> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-number> <json-null></pre> For more information about operators, see: Common connectors and operators (on page dccclxii) .		String

Example conversation - default columns

Request

```
GET api/sam/v2/software_components?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":1,
  "name":"DB2 8.1",
  "release":8,
  "guid":"XYZ",
  "is_ibm":1,
}]
```

Example conversation - association

Request

```
GET api/sam/v2/software_components?columns[ ]=id&columns[ ]=name
&columns[ ]=tags.name
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":1,
  "name":"DB2 8.1",
  "tags":
  [{"name":"ibm"}]
}]
```

Retrieval of license metric utilization (v2)

9.2.10 Available from 9.2.10. You use the `GET` operation on the `api/sam/v2/license_usage` element to request information about utilization of license metrics by products that are installed in your infrastructure. By default, the results are returned for the computer group of the user whose token is used for authentication and cover the period for which data is aggregated in this group. They also include custom fields that were added to the All Metrics report.

Permissions

 You must have the View License Metrics permission to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/license_usage?token=token
```

Resource information

Table 162. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	GET
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	application/json
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p> <p>Header</p> <p>startdate</p> <p>Values</p> <p>YYYY-MM-DD</p>

Table 162. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
	Specifies the date starting from which the data is retrieved.
	Header
	<code>enddate</code>
	Values
	<code>YYYY-MM-DD</code>
	Specifies the date until which the data is retrieved.
	Header
	<code>computerGroupId</code>
	Values
	Integer
	Specifies ID of the computer group for which data is retrieved.
Response payload	<code>License Usage</code> element
Response format	<code>application/json</code>
Response codes	200 - OK
	400 - "Bad Request" if a query parameter contains errors or is missing

Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/license_usage.json?token=token
```

Available columns

Table 163. Available columns

The table consists of three columns and 26 rows.

Column	Description	Displayed by default	Type
product_publisher_name	Name of the product publisher.	✓	String
product_id	Identifier of the software product.		Integer
product_name	Name of the software product.	✓	String
product_family_guid	GUID of the software product.		String
metric_id	Identifier of the license metric. For explanation of the returned values, see: Metric IDs and code names (on page mxxiv) .		String
metric_name	Name of the license metric. The value that is returned by the <code>metric_name</code> parameter is provided as reference and can slightly differ from the value that is provided as a metric description in Table 167: Metric ID, code name, and description (on page mxxv) . It is best to retrieve metrics by using <code>metric_id</code> or <code>metric_code_name</code> parameters and then check the exact metric description by referring to Table 167: Metric ID, code name, and description (on page mxxv) .		String
metric_code_name	Code name of the license metric. For explanation of the returned values, see: Metric IDs and code names (on page mxxiv) .	✓	String
hwm_quantity	The highest number of metric units that the product used within the period for which the data is retrieved. When metric quantity is not measured for a particular license metric, the value returned by the hwm_quantity parameter is -1. The value -1 cannot be used for sorting nor filtering.	✓	Integer
threshold	The maximum number of metric units that the product is entitled to use within a computer group. The value is set manually and is used to calculate the metric threshold delta.		Integer
threshold_delta	It is calculated by subtracting metric quantity from the threshold. When you specify a threshold for a license metric that is not calculated, the value returned by the threshold_delta parameter is 2147483647. The value 2147483647 cannot be used for sorting nor filtering.		Integer

Table 163. Available columns

The table consists of three columns and 26 rows.

(continued)

Column	Description	Displayed by default	Type
imported_part_numbers	Part number that was imported to BigFix Inventory. It represents the product that is listed in the product_name column and its license metric.		String
is_reaggregation_needed	Specifies whether recalculation is needed for the product. The parameter cannot be used to filter or sort the results.		Boolean
custom_field_number	Custom field that was added to the All Metrics report. To view the list of all custom fields, view the license_usage.json schema (on page mxi) .		Various
9.2.14 hwm_peak_time	<p>The date and time when the highest number of metric units was used by a product in the selected period of time. If the value of hwm_quantity parameter is -1, the value of hwm_peak_time is meaningless.</p> <p> Note: Retrieving hwm_peak_time might noticeably increase the time of retrieving the data.</p>		String

Query parameters

Table 164. Query parameters

The table consists of four columns and six rows.

Parameter	Description	Required	Value
columns[]	<p>Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved.</p> <p>Example: Retrieve product name and threshold delta</p> <pre>URL?columns[]=product_name&columns[]=threshold_delta</pre>		String
computer-GroupId	Specify ID of the computer group for which you want to retrieve the data. If you do not specify this parameter, the data is retrieved for the computer group of the user whose token is used		Integer

Table 164. Query parameters

The table consists of four columns and six rows.

(continued)

Parameter	Description	Required	Value
	<p>for authentication. If you specify the parameter, you can retrieve the data for a subgroup of this computer group.</p> <p>To view IDs of computer groups, log in to BigFix Inventory and go to Reports > Computer Groups. Then, hover of the Manage Report View icon , click Configure View, and select the ID column to display it on the report.</p> <p>Example: Retrieve data for computer group 5</p> <pre>URL?computerGroupId=5</pre>		
order	<p>Specify how to sort the retrieved data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name.</p> <p>Example: Order by threshold delta</p> <pre>URL?order[]=threshold_delta desc</pre>		Alphanumeric
limit	<p>Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved.</p> <p>Example: Retrieve 100 records</p> <pre>URL?limit=100</pre>		Numeric
offset	<p>Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results.</p> <p>Example: Retrieve 50 records starting after record 150</p> <pre>URL?limit=50&offset=150</pre>		Numeric
startdate	<p>Specify the date starting from which you want to retrieve the data. Specify it in the <code>YYYY-MM-DD</code> format. If you do not specify the filter, its default value is the date of the last successful import of data to BigFix Inventory minus the number of days for which data is calculated in the computer group (90 days by default).</p> <p>Example: Retrieve data starting from 14th July 2017</p>		Date

Table 164. Query parameters

The table consists of four columns and six rows.

(continued)

Parameter	Description	Required	Value
enddate	<p><code>URL?startdate=2017-07-14</code></p> <p>Specify the date until which you want to retrieve the data. Specify it in the <code>YYYY-MM-DD</code> format. If you do not specify the filter, its default value is the date of the last successful import of data to BigFix Inventory.</p> <p>Example: Retrieve data from 1st October 2017 to 31st October 2017</p> <p><code>URL?startdate=2017-10-01&enddate=2017-10-31</code></p>		Date
token	<p>A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dccclxxii). You can also log in to BigFix Inventory, hover over the User icon  and click Profile. Then, click Show token.</p>	✓	Alphanumeric
criteria	<p>Retrieve records which match specific conditions. The parameter should have the following structure, written in one line:</p> <pre><criteria> ::= <left-brace> <boolean-operator> <colon> <left-bracket> <criterion> [{ <comma> <criterion> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criterion> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-number> <json-null></pre>		String



Note: The `license_usage` REST API does not support nested filtering criteria.

Example 1: Retrieve software instances whose product name contains "BigFix" AND the threshold delta is below 0

Table 164. Query parameters

The table consists of four columns and six rows.

(continued)

Parameter	Description	Required	Value
<pre>URL? criteria={"and":[{"product_name","contains","BigFix"}, ["threshold_delta","<","0"]] }</pre>	<p>In case of the <code>threshold</code>, and <code>custom_field_number</code> fields, you can retrieve all entries for which the value is specified or for which it is not specified.</p>		
<p>Example 2: Retrieve software products for which the threshold is specified</p>	<pre>URL?criteria={"and":[{"threshold","!=" ,[]]}</pre>		
<p>Example 3: Retrieve software products for which the threshold is not specified</p>	<pre>URL?criteria={"and":[{"threshold","=" ,[]]}</pre>		
<p>If you created custom fields that use the date values, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <code><operator></code>, and then specify the time value in the following convention: <code>PxD/PxW/PxM/PxY</code>, where <i>x</i> is a number in the 1-999 range, and <i>D</i>, <i>W</i>, <i>M</i>, or <i>Y</i> is a designator that represents days, weeks, months, or years respectively.</p>	<p>Example 4: Retrieve software products for which entitlement ends within next month</p>		
<pre>URL?criteria={"and":[{"custom_field_1","next","P1M"]}]}</pre>	<p>For more information about operators, see: Common connectors and operators (on page dccclxii).</p>		

Example conversation - default columns

Request

```
GET api/sam/v2/license_usage?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
```

```
Accept: application/json
Accept-Language: en-US
```

Response header

```
Status Code: 200 OK
Content-Type: application/json
computerGroupId: 0
enddate: 2017-10-31
startdate: 2017-10-01
```

Response body

```
[{
  "product_publisher_name": "IBM",
  "product_name": "WebSphere Service Registry and Repository",
  "metric_code_name": "PVU_FULL_CAP",
  "hwm_quantity": 480
}]
```

Example conversation - all columns

Request

```
GET api/sam/v2/license_usage?columns[]=product_publisher_name&columns[]=product_id
&columns[]=product_name&columns[]=product_family_guid&columns[]=metric_id
&columns[]=metric_name&columns[]=metric_code_name&columns[]=hwm_quantity
&columns[]=threshold&columns[]=threshold_delta&columns[]=imported_part_numbers
&columns[]=is_reaggregation_needed&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "product_publisher_name": "IBM",
  "product_id": 29258,
  "product_name": "WebSphere Service Registry and Repository",
  "product_family_guid": "3b31a72e-468d-47bb-825a-ea26c8e85199",
  "metric_id": 3,
  "metric_code_name": "PVU_FULL_CAP",
  "metric_name": "PVU Full Capacity",
```

```
"hwm_quantity": 480,
"threshold": null,
"threshold_delta": null,
"imported_part_numbers": null,
"is_reaggregation_needed": 0
}]
```

Example conversation - additional column

Request

```
GET api/sam/v2/license_usages?columns[]=product_name
&columns[]=metric_name&columns[]=threshold_delta
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "product_name": "WebSphere Service Registry and Repository",
  "metric_name": "PVU Full Capacity",
  "threshold_delta": 100
}]
```

Example conversation - custom field

To retrieve data from custom fields that were added to the All Metrics report, start by viewing the [license_usage.json](#) schema. The schema lists all columns, including custom fields. Identify the custom field from which you want to retrieve the data.

Request - check list of created custom fields

```
GET api/sam/v2/schemas/license_usage.json?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Response - list of all columns, including custom fields

```
[{
  "product_name":
    {
      "type": "string",
      "description": "Name of the software product."
    },
  ...
}]
```

```
"custom_field_1":
  {
    "type": "date",
    "title": "Entitlement End"
  }
}]
```

After you identify the name of the custom field, you can use it in the REST API request.

Request

```
GET api/sam/v2/license_usages?columns[]=product_name
&columns[]=custom_field_1&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "product_name": "WebSphere Service Registry and Repository",
  "custom_field_1": "2017-10-01"
}]
```

Setting license metric thresholds and custom field values (v2)

9.2.10 Available from 9.2.10. You use the `PUT` operation on the `api/sam/v2/license_usage` element to set up license metric thresholds as well as values of custom fields that were added on the All Metrics report.

Permissions

 You must have the Manage Contracts permission to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/license_usage?token=token
```

Resource information

Table 165. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	<code>PUT</code>

Table 165. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	<p><code>application/json</code></p>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p> <p>Header</p> <p><code>computerGroupId</code></p> <p>Values</p> <p>Integer</p> <p>Specifies ID of the computer group for which values are changed.</p>
Response payload	<p>n/a</p>
Response format	<p><code>application/json</code></p>
Response codes	<p>200 - OK</p>

Table 165. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
	400 - "Bad Request" if a parameter contains errors
	401 - "Unauthorized user" if the user whose token is used for authentication is not authorized to access the data
	404 - "Not Found" if a parameter does not exist
	406 - "Not Acceptable" if the import is in progress and the values cannot be updated. For information about checking the status of the import, see: Running data imports (on page dccclxxviii) .

Schema description

To retrieve the list of all columns that are used by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/license_usage.json?token=token
```

Query parameters

Table 166. Query parameters

The table consists of three columns and 26 rows.

Column	Description	Required	Type
computerGroupId	Identifier of the computer group for which you want to update the values. If you do not specify this parameter, the values are updated for the computer group of the user whose token is used for authentication. To view IDs of computer groups, log in to BigFix Inventory and go to Reports > Computer Groups . Then, hover of the Manage Report View icon  , click Configure View , and select the ID column to display it on the report.		
product_id	Identifier of the software product.	✓	Integer

Table 166. Query parameters

The table consists of three columns and 26 rows.

(continued)

Column	Description	Required	Type
<code>metric_id</code>	Identifier of the license metric. For the list of metric identifiers, see: Metric IDs and code names (on page mxxiv) .	✓	String
<code>threshold</code>	The maximum number of metric units that the product is entitled to use within a computer group.		Integer
<code>custom_field_number</code>	Custom field that was added to the All Metrics report. To view the list of all custom fields, view the license_usage.json schema (on page mxxi) .		Various



Note: You can change multiple values during a single request. For example, you can use a single request to set the metric threshold, set a value of one custom field, and clear the value of another custom field.

Example conversation - setting a license metric threshold

- To retrieve information about the identifier of the product for which you want to set the license metric threshold, use the following `GET` request. The request returns product ID and name as well as metric ID and code name.

Request

```
GET api/sam/v2/license_usage?columns[]=product_id&columns[]=product_name
&columns[]=metric_id&columns[]=metric_code_name&token=7adc3efb175e2bc0f4484bdd2efca54a8fa
04623
```

Response

```
[{
  "product_id": 29258,
  "product_name": "WebSphere Service Registry and Repository",
  "metric_id": 3,
  "metric_code_name": "PVU_FULL_CAP"
}]
```

- To change the value of the license metric threshold, use the following `PUT` request.

A relation between the product specified by the `product_id` parameter and the license metric specified by the `metric_id` parameter must exist in the software catalog. Otherwise, the request returns an error. If the

relation does not exist in the software catalog, you can create it on the Products & Metrics panel. For more information, see: [Assigning an additional metric to a product \(on page cdxvii\)](#).

Request

```
PUT api/sam/v2/license_usage?product_id=29258&metric_id=3
&threshold=300&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Response

```
200 - OK
```



Tip: To automate the process, you can write a script that parses the results returned by the `GET` request, iterates through these results, and updates them one by one by using the `PUT` request.

Example conversation - setting the value of a custom field

1. To check the list of custom fields that were created on the All Metrics report, view the `license_usage.json` schema. The schema lists all columns, including custom fields. Identify the custom field for which you want to set the value.

Request - check the list of existing custom fields

```
GET api/sam/v2/schemas/license_usage.json?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Response - list of all columns, including custom fields

```
[{
  "product_name":
    {
      "type": "string",
      "description": "Name of the software product."
    },
  ...
  "custom_field_1":
    {
      "type": "boolean",
      "title": "Requires Extension"
    }
}]
```

2. After you identify the name of the custom field, use the following `PUT` request.

Request

```
PUT api/sam/v2/license_usage?product_id=29258&metric_id=3
&custom_field_1=1&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Response

```
200 - OK
```

Example conversation - clearing a license metric threshold

1. To retrieve information about the identifier of the product for which you want to clear the license metric threshold, use the following `GET` request. The request returns product ID and name, metric ID and name, and the threshold value.

Request

```
GET api/sam/v2/license_usage?columns[]=product_id&columns[]=product_name
&columns[]=metric_id&columns[]=metric_name&columns[]=threshold
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Response

```
[{
  "product_id": 29258,
  "product_name": "WebSphere Service Registry and Repository",
  "metric_id": 3,
  "metric_name": "PVU Full Capacity"
  "threshold": 300
}]
```

2. To clear the license metric threshold, specify an empty value in the `threshold` parameter. The same method is used for clearing values of custom fields.

Request

```
PUT api/sam/v2/license_usage?product_id=29258&metric_id=3
&threshold=&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Response

```
200 - OK
```

9.2.8 Metric IDs and code names

9.2.8 Available from 9.2.8. The `api/sam/v2/software_instances` and `api/sam/v2/license_usage` REST APIs return information about metric that is used by a product in the form of a unique metric ID or metric code name. Learn what is the relation between the metric ID and code name as well as what is the description of every metric.

Table 167. Metric ID, code name, and description

Metric ID	Metric code name	Metric description
-32603	CONCURRENT_USER	Concurrent User
-32470	VU_VALUE_UNIT	VU Value Unit
-32337	APPSERVER_CONNECTED	AppServer Connected
-30184	LIMITED_USE_VIRTUAL_SERVER	Limited Use Virtual Server
-29478	POPULATED_SOCKET	Populated Socket
-27870	MANAGED_SWITCH	Managed Switch
-27189	BED	Bed
-26157	SWITCH	Switch
-25214	TOKEN	Token
-24771	AUTHORIZED_USER_SINGLE_INSTALL	Authorized User Single Install
9.2.16 -24466	5000_FLOATING_USERS	5000 Floating Users
-24282	PRINTER	Printer
-23594	ESTABLISHMENT	Establishment
-21525	CURRENCY_VALUE_UNIT	Currency Value Unit
-21171	MILLION_SPEND_CONVERSION_UNITS	Million Spend Conversion Units
-18783	AUTHORIZED_USER	Authorized User
-18750	PROCESSOR	Processor
-18667	FEED	Feed
-18577	STORAGE_CAPACITY_UNIT_(101-250)	Storage Capacity Unit (101-250)
-17749	STORAGE_CAPACITY_UNIT_(751-1250)	Storage Capacity Unit (751-1250)
-16893	TEN_MONTHLY_JOBS	10 Monthly Jobs
-16798	TERABYTE	Terabyte
-16762	INSTALL_WITH_PAGES	Install with up to <number> Pages
-15896	RESOURCE_VALUE_UNIT	Resource Value Unit
-15529	INSTANCE	Instance
-15520	CLIENT_USER	Client User

Table 167. Metric ID, code name, and description (continued)

Metric ID	Metric code name	Metric description
-13938	STORE	Stores
-13494	AUTHORIZED_USER_SINGLE_SESSION	Authorized User Single Session
-13103	CLIENT_DEVICE	Client Device
-12983	TRANSACTION	Transaction
-11982	SIMULTANEOUS_SESSION	Simultaneous Session
-11743	FILE	File
-10658	TERMINAL	Terminal
-10377	HOST_SERVER	Host Server
-10210	PAGE	Page
-9208	APPLIANCE_INSTALL	Appliance Install
-8670	RACK	Rack
-8617	MILLION_OF_SERVICE_UNIT_PER_HOURS	Millions of Service Units per Hour
-8333	STORAGE_CAPACITY_UNIT_(251-500)	Storage Capacity Unit (251-500)
-8152	STORAGE_DEVICE	Storage Device
-7054	VIRTUAL_CORE_BASED	Microsoft Virtual Core with SA
-6957	TIERED_TERABYTES	Tiered Terabytes
-6399	PROCESSOR_DAY	Processor Day
-5373	TIVOLI_MANAGEMENT_POINT	Tivoli Management Point
-5368	STORAGE_CAPACITY_UNIT_(1-100)	Storage Capacity Unit (1-100)
-5256	STORAGE_CAPACITY_UNIT_(501-750)	Storage Capacity Unit (501-750)
-5048	ONE_M_ORDER_LINES	1 Million Order Lines
- 4563	ORACLE_PROCESSORS	Oracle Processor Core
-4608	MANAGED_DEVICE	Managed Device Oracle Core Based
9.2.16 -4050	ADDRESSABLE_DEVICE	Addressable Device
-2984	LIMITED_USE_AUTHORIZED_USER	Limited Use Authorized User
-2809	STORAGE_CAPACITY_UNIT_(2001+)	Storage Capacity Unit (2001+)

Table 167. Metric ID, code name, and description (continued)

Metric ID	Metric code name	Metric description
-2758	FLOATING_USER_SINGLE_SESSION_SINGLE_INSTALL	Floating User Single Session Single Install
-2583	SERVER_DEVICE	Server
-2368	CLIENT	Client
-1827	LIMITED_USE_MANAGED_SERVER	Limited Managed Server Unit
1	INTERNAL_OTHER	Unknown
3	PVU_FULL_CAP Use this metric ID or code name to retrieve or set up thresholds of license metric utilization by using the <code>api/sam/v2/license_usage</code> REST API.	PVU Full Capacity
4	PVU_SUB_CAP Use this metric ID or code name to retrieve or set up thresholds of license metric utilization by using the <code>api/sam/v2/license_usage</code> REST API.	PVU Subcapacity
5	INTERNAL_PVU Use this metric ID or code name to retrieve information about the license metric that is assigned to a product by using the <code>api/sam/v2/software_instances</code> REST API.	PVU
6	INTERNAL_RVU_MAPC Use this metric ID or code name to retrieve information about the license metric that is assigned to a product by using the <code>api/sam/v2/software_instances</code> REST API.	RVU MAPC
7	RVU_FULL_CAP Use this metric ID or code name to retrieve or set up thresholds of license metric utilization by using the <code>api/sam/v2/license_usage</code> REST API.	RVU MAPC Full Capacity
8	RVU_SUB_CAP	RVU MAPC Subcapacity

Table 167. Metric ID, code name, and description (continued)

Metric ID	Metric code name	Metric description
	Use this metric ID or code name to retrieve or set up thresholds of license metric utilization by using the <code>api/sam/v2/license_usage</code> REST API.	
9	INTERNAL_INSTALL	Install Seats
10	FLEXPOINT	Flexpoint
488	TEN_AUTHORIZED_USER	10 Authorized User
526	MAILBOX	Mailbox
865	VIRTUAL_SERVER	Virtual Server
9.2.16 2110	500_FLOATING_USERS	500 Floating Users
2692	CONNECTOR_FOR_DEMAND_SIDE	Connector for Demand-side
3058	CONNECTION	Connection
3729	FLOATING_USER_SINGLE_INSTALL	Floating User Single Install
3774	USER_VALUE_UNIT	User Value Unit
9.2.16 3827	MANAGED_CONTAINER	Managed Container
5264	REGISTERED_USER	Registered User
5844	VIRTUAL_PROCESSOR_CORE	Virtual Processor Core
5856	ELIGIBLE_PARTICIPANT	Eligible Participant
6083	10_MANAGED_VIRTUAL_SERVERS	10 Managed Virtual Servers
6089	PER_TOKEN	Per Token
6420	MANAGED_VIRTUAL_NETWORK_DEVICE	Managed Virtual Network Device
9.2.16 6554	10000_FLOATING_USERS	10000 Floating Users
6967	APPLICATION	Application
7057	MONTHLY_MILLION_RULES_DECISIONS	Monthly Million Rules Decisions
7884	MILLION_SQUARE_FEET	Million Square Feet
8321	LIMITED_USE_SOCKET	Limited Use Socket
8604	INSTALL_WITH_3_AUTHORIZED_USERS	Install with 3 Authorized Users
8843	CONCURRENT_SESSION	Concurrent Session

Table 167. Metric ID, code name, and description (continued)

Metric ID	Metric code name	Metric description
9033	THOUSAND_ADDRESSABLE_DEVICES	Thousand Addressable Devices
10157	CONNECTOR	Connector
10310	PORT	Port
10989	APPLICATION_INSTANCE	Application Instance
11165	LINEAR_ASSET	Linear Asset
11328	USER	Users
12599	HUNDRED_THOUSAND_SQUARE_METERS	Hundred Thousand Square Meters
13557	FLOATING_USER	Floating User
14325	MONTHLY_THOUSAND_MANAGED_DECISION_ARTIFACTS	Monthly Thousand Managed Decision Artifacts
14692	SINGLE_PHYSICAL_PROCESSOR_BASED	Microsoft Single Processor
15006	FLOATING_USER_SINGLE_SESSION	Floating User Single Session
15282	DECIMAL_TERABYTE	Decimal Terabyte
16958	BASE_PROCESSOR	Base Processors
17437	SERVER_LESS_THAN_1000	Server Less Than 1000 Nodes
17502	MANAGED_CHASSIS	Managed Chassis
19131	MONTHLY_JOBS	Monthly Order Request
20170	SOCKET	Socket
20249	TICKETS	Tickets
21058	SERVER_GREATER_OR_EQUAL_TO_1000	Server Greater Or Equal to 1000 Nodes
21426	NETWORK_NODE	Network Node
9.2.16 21610	1000_FLOATING_USERS	1000 Floating Users
21704	CONNECTOR_FOR_SUPPLY_SIDE	Connector for Supply-side
21936	MANAGED_TERABYTE	Managed Terabyte
22225	STORAGE_CAPACITY_UNIT_(1251-2000)	Storage Capacity Unit (1251-2000)
22989	IDLE_STANDBY_SERVER	Idle Standby Server

Table 167. Metric ID, code name, and description (continued)

Metric ID	Metric code name	Metric description
23720	PHYSICAL_CORE_BASED	Microsoft Physical Core with SA
24179	MANAGED_CLIENT_DEVICES	Managed Client Devices
9.2.16 26571	MONTHLY_HOUR	Monthly Hour
27777	ASSET	Asset
28226	GENERAL_PURPOSE_GRAPHICS_PROCESSING_UNIT	General Purpose Graphics Proc Unit
28591	MANAGED_VIRTUAL_SERVER	Managed Virtual Server
28953	INSTALL	Install Instances
28967	MANAGED_SERVER	Managed Server
29313	ADDRESSABLE_DEVICES	Addressable Devices
31619	STG_TIER	STG Tier
31705	SERVER_WITH_ONE_PROCESSOR	Server with One Processor
31997	ADDITIONAL_PROCESSOR	Additional Processor
32455	DUAL_PHYSICAL_PROCESSOR_BASED	Microsoft Dual Processor

REST API for hardware (v2)

9.2.8 Available from 9.2.8. You can use REST API requests to quickly retrieve large amounts of data related to hardware inventory, including historical data about removed computers. You can then pass this data to other applications for further processing and analysis.

Retrieval of hardware inventory (v2)

9.2.8 Available from 9.2.8. You use the `GET` operation on the `api/sam/v2/computers` element to request information about computers in your infrastructure. The API returns details of existing computers as well as historical data about computers that were removed.

Permissions

 You must have the View Endpoints and View Hardware Inventory permission to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/computers?token=token
```

Resource information

Table 168. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	GET
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	application/json
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response payload	Computer Systems element
Response format	application/json
Response codes	<p>200 - OK</p> <p>500 - "Bad Request" if a query parameter contains errors or is missing</p>

9.2.10 Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/computer.json?token=token
```

Available columns

Table 169. Available columns

Column	Description	Displayed by default	Type
id	Identifier of the computer as defined in BigFix Inventory.		Nu-mer-ic
bigfix_id	Identifier of the computer as defined in BigFix.		Nu-mer-ic
computer_group_id	Identifier of the BigFix Inventory computer group to which the computer belongs. The value 0 represents the All Computers group.		Nu-mer-ic
name	Name of the computer.	✓	String
dns_name	DNS of the computer.		String
ip_address	IP address of the computer.	✓	String
os	Operating system of the computer.	✓	String
os_type	Type of the operating system of the computer.		String
first_seen	Date and time when the computer reported to BigFix for the first time. The time is specified in the GMT time zone.		String
last_seen	Date and time when the computer reported to BigFix for the last time. The time is specified in the GMT time zone.	✓	String
is_deleted	Information whether the computer was removed.	✓	Boolean
deletion_date	Date and time when the computer was removed. The time is specified in the GMT time zone.		String
is_managed_by_vm_manager	Information whether the computer is managed by a VM manager.		Boolean

Applicable associations

You can additionally retrieve data from the following associations:

- **9.2.10** `computer_details` (on page [mcxviii](#))
- `computer_hardware` (on page [mcxxi](#))
- `computer_health` (on page [mcxx](#))
- **9.2.12** `mounted_shared_disks` (on page [mcxxx](#))
- **9.2.12** `detailed_hw_ip_addresses` (on page [mxliv](#))
- **9.2.12** `detailed_hw_lpars` (on page [mxlix](#))
- **9.2.12** `detailed_hw_memories` (on page [mliv](#))
- **9.2.12** `detailed_hw_network_adapters` (on page [mlix](#))
- **9.2.12** `detailed_hw_operating_systems` (on page [mlxiv](#))
- **9.2.12** `detailed_hw_partitions` (on page [mlxix](#))
- **9.2.12** `detailed_hw_physical_processors` (on page [mlxxiv](#))
- **9.2.12** `detailed_hw_smbios` (on page [mlxxx](#))
- **9.2.12** `detailed_hw_storages` (on page [mlxxxv](#))



Note: Retrieving data from multiple API associations in a single request might increase the response time.

For more information about how to use API associations, see: [REST API associations \(on page mcxv\)](#).



Important: All associations provide data about items that currently exist in the infrastructure. For removed computers and uninstalled software, these associations return `null`.

Query parameters

Table 170. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
<code>columns[]</code>	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve product name and release <pre>URL?columns[]=computer_group_id&columns[]=first_seen</pre>		String
<code>order</code>	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by computer ID descending: <pre>URL?order[]=id desc</pre>		String

Table 170. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved. Example: Retrieve 100 records <code>URL?limit=100</code>		Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150 <code>URL?limit=50&offset=150</code>		Numeric
token	A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dccclxxii) . You can also log in to BigFix Inventory, hover over the User icon  , and click Profile . Then, click Show token .	✓	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written in one line: <pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criterion> [{ <comma> <criterion> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criterion> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null></pre>		String

For more information about operators, see: [Common connectors and operators \(on page dccclxii\)](#).

Example: retrieve computers that reported for the first time within a specific date range:

Table 170. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<pre>URL?criteria={"and":[{"first_seen", ">", "1970-01-01T00:00:00+00:00Z"}, {"first_seen", "<", "1970-01-02T00:00:00+00:00Z"}]}</pre>		
	<p>For columns that use the date and time values, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <i><operator></i>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.</p> <p>Example: Retrieve computers that were first reported within the last 7 days</p>		
	<pre>URL?criteria={"and":[{"first_seen", "last", "P7D"}]}</pre>		

Example conversation - default columns

Request

```
GET api/sam/v2/computers?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "name": "NC9143126194",
  "ip_address": ["9.143.126.194"],
  "os": "Win2012R2 6.3.9600",
  "last_seen": "2016-09-22T10:32:12Z",
  "is_deleted": 1
}]
```

Example conversation - all columns

Request

```
GET api/sam/v2/computers?columns[]=id&columns[]=bigfix_id
&columns[]=computer_group_id&columns[]=name&columns[]=dns_name
&columns[]=ip_address&columns[]=os&columns[]=os_type
&columns[]=first_seen&columns[]=last_seen&columns[]=is_deleted
&columns[]=deletion_date&columns[]=is_managed_by_vm_manager
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":2,
  "bigfix_id":11368943,
  "computer_group_id":0,
  "name":"NC9143126194",
  "dns_name":"NC9143126194",
  "ip_address":["9.143.126.194"],
  "os":"Win2012R2 6.3.9600",
  "os_type":"Windows",
  "first_seen":"2016-10-04T09:42:55Z",
  "last_seen":"2016-09-22T10:32:12Z",
  "is_deleted":1,
  "deletion_date":"2016-12-19T09:40:14Z",
  "is_managed_by_vm_manager":0
}]
```

Example conversation - additional column

Request

```
GET api/sam/v2/computers?columns[]=id
&columns[]=is_deleted&columns[]=deletion_date
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":21,
  "is_deleted":1,
```

```
"deletion_date": "2016-10-18T14:51:52Z"
}]
```

Example conversation - association

Request

```
GET api/sam/v2/computers?columns[ ]=id
&columns[ ]=computer_health.catalog_version
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":1,
  "catalog_version":1304630
}]
```

Retrieval of shared disks (v2)

9.2.12 Available from 9.2.12. You use the `GET` operation on the `api/sam/v2/shared_disks` element to request information about shared disks that are used in your infrastructure.

Prerequisites

Before you can retrieve information about shared disks by using the REST API, discover the shared disks in your infrastructure. For more information, see: [Step 1: Discovering remote shared disks \(on page ccxxiii\)](#).

Permissions



You must have the View Shared Disks permissions to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/shared_disks?token=token
```

Resource information

Table 171. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	GET
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	application/json
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response payload	Shared Disks element
Response format	application/json
Response codes	<p>200 - OK</p> <p>403 - "Forbidden" if the user is not authorized to perform the action</p> <p>500- "Bad Request" if a query parameter contains errors or is missing</p>

Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/shared_disk.json?token=<token>
```

Available columns

Table 172. Available columns

The table consists of four columns and 8 rows.

Column	Description	Displayed by default	Type
id	Identifier of the shared disk instance.	✓	Integer
status	Indicates the status of the computer that is designated to scan the shared disk instance. Possible values are: <ul style="list-style-type: none"> • 0 - Shared disk detected • 1 - Child shared disk • 2 - Designating a computer • 3 - Waiting for software scan • 4 - OK • 5 - Designating new top-level directory • 6 - Designating a new computer • 7 - Waiting for software scan on the newly designated computer 	✓	Integer
parent_id	Identifier of the parent of the shared disk instance.	✓	Integer
top_level	Indicates whether the shared disk is a top-level directory in the shared disk structure. Possible values are: <ul style="list-style-type: none"> • 0 - The directory is not top-level • 1 - The directory is top-level 		Boolean
exported_directory	Exported directory of the shared disk.	✓	String

Table 172. Available columns

The table consists of four columns and 8 rows.

(continued)

Column	Description	Displayed by default	Type
type	Type of the shared disk. Possible values are: <ul style="list-style-type: none"> • 0 - Network File System • 1- Samba File System 	✓	Integer
computers_count	Number of computers on which the shared disk is mounted.	✓	Integer

Applicable associations

You can additionally retrieve data from the following associations:

- [designated_computer](#) (on page [mcxxviii](#))
- [shared_disk_members](#) (on page [mcxxxii](#))

Query parameters

Table 173. Query parameters

The table consists of four columns and six rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve product name and release <code>URL?columns[]=product_name&columns[]=product_release</code>		String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by computer ID descending <code>URL?order[]=computer_id desc</code>		Alphanumeric

Table 173. Query parameters

The table consists of four columns and six rows.

(continued)

Parameter	Description	Required	Value
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved. Example: Retrieve 100 records <code>URL?limit=100</code>		Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150 <code>URL?limit=50&offset=150</code>		Numeric
token	A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dccclxxii) . You can also log in to BigFix Inventory, hover over the User icon  , and click Profile . Then, click Show token .	✓	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written in one line: <pre><criteria> ::= <left-brace> <boolean-operator> <colon> <left-bracket> <criteria> [{ <comma> <criteria> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criteria> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-number> <json-null></pre>		String

For more information about operators, see: [Common connectors and operators \(on page dccclxii\)](#).

Example: Retrieve software instances whose product name contains "BigFix" OR the discovery start is within a specific date range

Table 173. Query parameters

The table consists of four columns and six rows.

(continued)

Parameter	Description	Required	Value
	<pre>URL?criteria={ "or": [["product_name", "contains", "BigFix"], { "and": [["discovery_start", ">", "1970-01-01T00:00:00+00:00Z"], ["discovery_start", "<", "1970-01-02T00:00:00+00:00Z"]] }] }</pre>		
	<p>For columns that use the date and time values, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <operator>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.</p> <p>Example: Retrieve software instances that were first reported within the last 7 days</p>		
	<pre>URL?criteria={"and":[["discovery_start", "last", "P7D"]]}</pre>		

Example conversation - default columns

Request

```
GET api/sam/v2/shared_disks?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id": "5",
  "status": "0"
  "type": "0"
  "designated_computer_id": "29"
  "parent_id": "3"
  "computers_count": "10"
  "exported_directory": "192.0.2.24:/file_server/tlcm/ISO/RHEL"
```

Example conversation - all columns

Request

```
GET api/sam/v2/shared_disks?columns[]=id&columns[]=status
&columns[]=parent_id&columns[]=top_level&columns[]=exported_directory
&columns[]=type&columns[]=computers_count&columns[]=designated_copmuter_id
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Response body

```
[{
  "id": "1",
  "status": "0",
  "parent_id": "5",
  "top_level": "0",
  "exported_directory": "192.0.2.24:/file_server/tlcm/HOME"
  "type": "0",
  "computers_count": "10",
  "designated_computer_id": "29"
```

Example conversation - association

Request

```
GET api/sam/v2/shared_disks?columns[]=exported_directory
&columns[]=computers_count&columns[]=shared_disk_members.computer_id
&columns[]=shared_disk_members.mount_point&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "exported_directory": "192.0.2.24:/file_server/tlcm/HOME",
  "computers_count": "1",
  "members":
    {
      "computer_id": "5",
      "mount_point": "/mnt/fs"
    }
}]
```

REST API for retrieving detailed hardware information (v2)

You can use REST API requests to retrieve data related to detailed hardware information, such as memory, operating systems, storage, processors, partitions, network adapters, SMBIOS data, IP addresses and logical partition capacity data. You can then pass this data to other applications for further processing and analysis.



Note: The APIs from this section can be used as associations to the `api/sam/v2/computers` API.

Retrieval of IP addresses (v2)

9.2.12

Available from 9.2.12. You use the `GET` operation on the `api/sam/v2/detailed_hw_ip_addresses` element to request information related to IP addresses of computers in your infrastructure.

Permissions



You must have the View Endpoints and View Hardware Inventory permission to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/detailed_hw_ip_addresses?token=token
```

Resource information

Table 174. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	<code>GET</code>
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p>

Table 174. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
	Specifies the content type of the response.
	<p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p>
	Specifies the language of the response content. If the header is not specified, the content is returned in the server language.
Response payload	IP Addresses element
Response format	application/json
Response codes	<p>200 - OK</p> <p>500 - "Bad Request" if a query parameter contains errors or is missing</p>

Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/detailed_hw_ip_address.json?token=token
```

Available columns

The list of attributes related to IP addresses on the scanned system. The retrieved addresses are in the IPv4 format, except for the **ipv6_address** attribute.

The output might contain multiple instances of each attribute from this group.

Table 175. Columns with information related to IP addresses

Property	Description	Type
id	Identifier of the record.	Integer

Table 175. Columns with information related to IP addresses (continued)

Property	Description	Type
computer_id	Identifier of the computer as specified in the BigFix Inventory database.	Integer
updated_at	Date and time when the current information was imported to BigFix Inventory.	String
address	IP address of the computer.	String
hostname	Host name of the computer with the given IP.	String
primary_dns	Primary domain name system (DNS).	String
permanent_mac_address	Permanent media access control (MAC) address associated with the given IP address.	String
ipv6_address	IP address in version 6.	String

Related APIs

This API can be associated with the following API:

- [Retrieval of hardware inventory \(v2\) \(on page mxxx\)](#)

Query parameters

Table 176. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve the computer ID: <pre>URL?columns[]=computer_id</pre>		String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by computer ID descending: <pre>URL?order[]=computer_id desc</pre>		String

Table 176. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
limit	<p>Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved.</p> <p>Example: Retrieve 100 records</p> <pre>URL?limit=100</pre> <p> Note: By default the limit parameter for this API is set to 100000.</p>		Numeric
offset	<p>Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results.</p> <p>Example: Retrieve 50 records starting after record 150</p> <pre>URL?limit=50&offset=150</pre>		Numeric
token	<p>A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dcccclxxii). You can also log in to BigFix Inventory, hover over the User icon  and click Profile. Then, click Show token.</p>	✓	Alphanumeric
criteria	<p>Retrieve records which match specific conditions. The parameter should have the following structure, written in one line:</p> <pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criterion> [{ <comma> <criterion> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criterion> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null></pre> <p>For more information about operators, see: Common connectors and operators (on page dcccclxii).</p>		String

Table 176. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<p>Example: retrieve computers with computer ID greater than 1000.</p> <pre>URL?criteria={"and":[{"computer_id", ">", "1000"}]}</pre> <p>For columns that use the date and time values, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <i><operator></i>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.</p>		

Example conversation - default columns

Request

```
GET /api/sam/v2/detailed_hw_ip_addresses?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":1
  "computer_id":1,
  "updated_at":"2018-05-16T14:37:08Z",
  "address":"9.174.39.183",
  "hostname":"NC9174039183",
  "primary_dns":"9.20.136.11",
  "permanent_mac_address":"00:50:56:94:7C:95",
  "ipv6_address":"fe80:0000:0000:0000:44a3:f150:0968:869e%2"
}]
```

Example conversation - selected columns

Request

```
GET api/sam/v2/detailed_hw_ip_addresses?columns[]=computer_id
&columns[]=primary_dns
```

```
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "computer_id":1,
  "primary_dns":"9.20.136.11"}
]
```

Retrieval of information related to logical partition capacity (v2)

9.2.12 Available from 9.2.12. You use the `GET` operation on the `api/sam/v2/detailed_hw_lpars` element to request information related to the logical partition and its CPU resources. The data retrieved by this API is not limited to LPARs on POWER. It includes data related to every supported virtualization technology.

Permissions

 You must have the View Endpoints and View Hardware Inventory permission to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/detailed_hw_lpars?token=token
```

Resource information

Table 177. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	GET
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	application/json

Table 177. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p><code>en-US, ...</code></p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response payload	<code>LPARS</code> element
Response format	<code>application/json</code>
Response codes	<p><code>200 - OK</code></p> <p><code>500 - "Bad Request"</code> if a query parameter contains errors or is missing</p>

Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/detailed_hw_lpar.json?token=token
```

Available columns

The list of attributes related to the logical partition and its CPU resources.

The scan output contains a single instance of each attribute from this group.

Table 178. Columns with information about the logical partition and its CPU resources

Property	Description	Type
id	Identifier of the record.	Integer
computer_id	Identifier of the computer as specified in the BigFix Inventory database.	Integer
updated_at	Date and time when the current information was imported to BigFix Inventory.	String
lpar_capacity_in_cores	The number of processor cores assigned to the partition.	String
lpar_online_processor_count	The number of active logical processors in the partition. Logical processors are the smallest CPU entities visible in the operating system, which are dividable and schedulable. Depending on the hardware platform, a logical processor might correspond to a chip, the core of a multi-core chip, or a hardware CPU thread on a multithreaded chip or core (provided that the feature is enabled).	String
lpar_online_core_count	The number of active cores in a logical partition. The core is active if any of the logical processors in the system is running on this core.	String
lpar_online_package_count	The number of active processor packages in the partition. The processor package is active if any of the logical processors in the system is running on this package.	String

Related APIs

This API can be associated with the following API:

- [Retrieval of hardware inventory \(v2\) \(on page mxxx\)](#)

Query parameters

Table 179. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved.		String

Example: Retrieve the computer id:

```
URL?columns[]=computer_id
```

Table 179. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by computer ID descending: <pre>URL?order[]=computer_id desc</pre>		String
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved. Example: Retrieve 100 records <pre>URL?limit=100</pre>  Note: By default the limit parameter for this API is set to 100000.		Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150 <pre>URL?limit=50&offset=150</pre>		Numeric
token	A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dcccxxii) . You can also log in to BigFix Inventory, hover over the User icon  , and click Profile . Then, click Show token .	✓	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written in one line: <pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criterion> [{ <comma> <criterion> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criterion> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string></pre>		String

Table 179. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<pre><value> ::= <json-array> <json-string> <json-numver> <json-null></pre>		
	<p>For more information about operators, see: Common connectors and operators (on page dccclxii).</p> <p>Example: retrieve computers with computer ID greater than 1000.</p> <pre>URL?criteria={"and":[{"computer_id", ">", "1000"}]}</pre>		
	<p>For columns that use the date and time values, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <operator>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.</p>		

Example conversation - default columns

Request

```
GET api/sam/v2/detailed_hw_lpars?token=c77a88cfdde276de605c6fdee689b7961ddb93ef
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":1,
  "computer_id":1,
  "updated_at":"2018-05-28T01:22:16Z",
  "lpar_capacity_in_cores":"-1.000000",
  "lpar_online_processor_count":"2.000000",
  "lpar_online_core_count":"2.000000",
  "lpar_online_package_count":"2.000000"
}]
```

Example conversation - selected columns

Request

```
GET api/sam/v2/detailed_hw_lpars?columns[]=id
&columns[]=lpar_capacity_in_cores
&token=c77a88cfdde276de605c6fdee689b7961ddb93ef
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":1,
  "lpar_capacity_in_cores":"-1.000000"
}]
```

Retrieval of information related to system memory (v2)

9.2.12 Available from 9.2.12. You use the `GET` operation on the `api/sam/v2/detailed_hw_memories` element to request information related to the system memory of the scanned system.

Permissions

 You must have the View Endpoints and View Hardware Inventory permission to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/detailed_hw_memories?token=token
```

Resource information

Table 180. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	<code>GET</code>
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p>

Table 180. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
	Negotiates the language of the response. If the header is not specified, the content is returned in the server language.
Request format	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response payload	System Memory element
Response format	<code>application/json</code>
Response codes	<p><code>200 - OK</code></p> <p><code>500 - "Bad Request"</code> if a query parameter contains errors or is missing</p>

Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/detailed_hw_memory.json?token=token
```

Available columns

The list of attributes related to the memory on a scanned system.

The scan output contains a single instance of each attribute from this group.

Table 181. Columns with information about the system memory

Property	Description	Type
id	Identifier of the record.	Integer
computer_id	Identifier of the computer as specified in the BigFix Inventory database.	Integer
updated_at	Date and time when the current information was imported to BigFix Inventory.	String
total_physical	Total amount of installed memory in kilobytes (KB).	String
free_physical	Total amount of free memory on the system in kilobytes (KB).	String

Related APIs

This API can be associated with the following API:

- [Retrieval of hardware inventory \(v2\) \(on page mxxx\)](#)

Query parameters

Table 182. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve the computer id: <pre>URL?columns[]=computer_id</pre>		String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by computer ID descending: <pre>URL?order[]=computer_id desc</pre>		String
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved.		Numeric

Table 182. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<p>Example: Retrieve 100 records</p> <pre>URL?limit=100</pre> <p> Note: By default the limit parameter for this API is set to 100000.</p>		
offset	<p>Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results.</p> <p>Example: Retrieve 50 records starting after record 150</p> <pre>URL?limit=50&offset=150</pre>		Numeric
token	<p>A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dccclxxii). You can also log in to BigFix Inventory, hover over the User icon , and click Profile. Then, click Show token.</p>	✓	Alphanumeric
criteria	<p>Retrieve records which match specific conditions. The parameter should have the following structure, written in one line:</p> <pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criterion> [{ <comma> <criterion> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criterion> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null></pre> <p>For more information about operators, see: Common connectors and operators (on page dccclxii).</p> <p>Example: retrieve computers with computer ID greater than 1000.</p>		String

Table 182. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<pre>URL?criteria={"and":[{"computer_id", ">", "1000"}]}</pre>		
	<p>For columns that use the date and time values, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <i><operator></i>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.</p>		

Example conversation - default columns

Request

```
GET api/sam/v2/detailed_hw_memories?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":1,
  "computer_id":1,
  "updated_at":"2018-05-28T01:22:16Z",
  "total_physical":"8010840",
  "free_physical":"221176"
}]
```

Example conversation - selected columns

Request

```
GET api/sam/v2/detailed_hw_memories?columns[]=id
&columns[]=total_physical
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":1,
  "total_physical":"8010840"
}]
```

Retrieval of information related to network adapters (v2)

9.2.12 Available from 9.2.12. You use the `GET` operation on the `api/sam/v2/detailed_hw_network_adapters` element to request information related to the type, model and MAC addresses of the network adapters.

Permissions



You must have the View Endpoints and View Hardware Inventory permission to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/detailed_hw_network_adapters?token=token
```

Resource information

Table 183. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	<code>GET</code>
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p>

Table 183. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
Header	
	<code>Content-Language</code>
Values	
	en-US, ...
	Specifies the language of the response content. If the header is not specified, the content is returned in the server language.
Response payload	<code>Network Adapter</code> element
Response format	<code>application/json</code>
Response codes	<p>200 - OK</p> <p>500 - "Bad Request" if a query parameter contains errors or is missing</p>

Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/detailed_hw_network_adapter.json?token=token
```

Available columns

The list of attributes related to the network adapters.

The scan output might contain multiple instances of each attribute from this group.

Table 184. Columns with information about the network adapters

Property	Description	Type
id	Identifier of the record.	Integer
computer_id	Identifier of the computer as specified in the BigFix Inventory database.	Integer
updated_at	Date and time when the current information was imported to BigFix Inventory.	String

Table 184. Columns with information about the network adapters (continued)

Property	Description	Type
permanent_address	The permanent media access control (MAC) address associated with the adapter.	String
current_address	The current network address of the adapter.	String
type	The type of network adapter.	String
model	The model of the network adapter.	String

Related APIs

This API can be associated with the following API:

- [Retrieval of hardware inventory \(v2\) \(on page mxxx\)](#)

Query parameters

Table 185. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve the computer id: <pre>URL?columns[]=computer_id</pre>		String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by computer ID descending: <pre>URL?order[]=computer_id desc</pre>		String
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved. Example: Retrieve 100 records <pre>URL?limit=100</pre>		Numeric

Table 185. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	 Note: By default the limit parameter for this API is set to 100000.		
offset	<p>Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results.</p> <p>Example: Retrieve 50 records starting after record 150</p> <pre>URL?limit=50&offset=150</pre>		Numeric
token	<p>A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dccclxxii). You can also log in to BigFix Inventory, hover over the User icon , and click Profile. Then, click Show token.</p>	✓	Alphanumeric
criteria	<p>Retrieve records which match specific conditions. The parameter should have the following structure, written in one line:</p> <pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criteria> [{ <comma> <criteria> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criteria> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null></pre> <p>For more information about operators, see: Common connectors and operators (on page dccclxxii).</p> <p>Example: retrieve computers with computer ID greater than 1000.</p> <pre>URL?criteria={"and":[{"computer_id", ">", "1000"}]}</pre> <p>For columns that use the date and time values, you can retrieve data also for a period instead of a specific date. To do so, use</p>		String

Table 185. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<i>last</i> or <i>next</i> as <code><operator></code> , and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.		

Example conversation - default columns

Request

```
GET api/sam/v2/detailed_hw_network_adapters?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":13,
  "computer_id":4,
  "updated_at":"2018-05-28T01:22:16Z",
  "permanent_address":"00:50:56:94:A1:39",
  "current_address":"00:50:56:94:A1:39",
  "type":"Ethernet Adapter",
  "model":"vmxnet3 Ethernet Adapter"
}]
```

Example conversation - selected columns

Request

```
GET api/sam/v2/detailed_hw_network_adapters?columns[]=id
&columns[]=id&columns[]=type&columns[]=model
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":13,
  "type":"Ethernet Adapter",
  "model":"vmxnet3 Ethernet Adapter"
}]
```

Retrieval of information related to the operating system (v2)

9.2.12

Available from 9.2.12. You use the `GET` operation on the `api/sam/v2/detailed_hw_operating_systems` element to request information related to the operating system on the scanned computer.

Permissions



You must have the View Endpoints and View Hardware Inventory permission to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/detailed_hw_operating_systems?token=token
```

Resource information

Table 186. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	<code>GET</code>
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p>

Table 186. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
	Specifies the content type of the response.
	Header
	<code>Content-Language</code>
	Values
	en-US, ...
	Specifies the language of the response content. If the header is not specified, the content is returned in the server language.
Response payload	Operating System element
Response format	<code>application/json</code>
Response codes	200 - OK
	500 - "Bad Request" if a query parameter contains errors or is missing

Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/detailed_hw_operating_system.json?token=token
```

Available columns

The list of attributes related to the operating system that is installed on the scanned system.

The scan output contains a single instance of each attribute from this group.

Table 187. Columns with information about the operating system

Property	Description	Type
id	Identifier of the record.	Integer

Table 187. Columns with information about the operating system (continued)

Property	Description	Type
computer_id	Identifier of the computer as specified in the BigFix Inventory database.	Integer
updated_at	Date and time when the current information was imported to BigFix Inventory.	String
name	The name of an operating system; for example, Windows 10.	String
major_version	The major version of an operating system.	String
minor_version	The minor version of an operating system.	String
Windows instal- l_date	The date when the operating system was installed.	String
os_arch	The version of the operating system architecture.	String
os_kernel_mode	The kernel mode of the operating system (32-bit or 64-bit).	String
Linux de- scription_string	The description of the operating system.	String

Related APIs

This API can be associated with the following API:

- [Retrieval of hardware inventory \(v2\) \(on page mxxx\)](#)

Query parameters

Table 188. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve the computer id: <pre>URL?columns[]=computer_id</pre>		String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name.		String

Table 188. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<p>Example: Order by computer ID descending:</p> <pre>URL?order[]=computer_id desc</pre>		
limit	<p>Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved.</p> <p>Example: Retrieve 100 records</p> <pre>URL?limit=100</pre> <p> Note: By default the limit parameter for this API is set to 100000.</p>		Numeric
offset	<p>Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results.</p> <p>Example: Retrieve 50 records starting after record 150</p> <pre>URL?limit=50&offset=150</pre>		Numeric
token	<p>A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dcccclxxii). You can also log in to BigFix Inventory, hover over the User icon , and click Profile. Then, click Show token.</p>	✓	Alphanumeric
criteria	<p>Retrieve records which match specific conditions. The parameter should have the following structure, written in one line:</p> <pre><criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criterion> [{ <comma> <criterion> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criterion> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null></pre>		String

Table 188. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<p>For more information about operators, see: Common connectors and operators (on page dccclxii).</p> <p>Example: retrieve computers with computer ID greater than 1000.</p> <pre>URL?criteria={"and":[{"computer_id", ">", "1000"}]}</pre>		
	<p>For columns that use the date and time values, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <operator>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.</p>		

Example conversation - default columns

Request

```
GET api/sam/v2/detailed_hw_operating_systems?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":13,
  "computer_id":22,
  "updated_at":"2018-05-28T01:22:16Z",
  "name":"Red Hat Enterprise Linux Server release 6.8 (Santiago)",
  "major_version":"2",
  "minor_version":"6",
  "install_date":"2017-10-31-08.37.06.000000",
  "os_arch":"x86_64",
  "os_kernel_mode":"64",
  "description_string":"Red Hat Enterprise Linux Server release 6.8 (Santiago)"
}]
```

Example conversation - selected columns

Request

```
GET /api/sam/v2/detailed_hw_operating_systems?columns[]=id
&columns[]=name&columns[]=major_version&columns[]=minor_version
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":13,
  "name":"Red Hat Enterprise Linux Server release 6.8 (Santiago)",
  "major_version":"2",
  "minor_version":"6"
}]
```

Retrieval of information related to partitions (v2)

9.2.12 Available from 9.2.12. You use the `GET` operation on the `api/sam/v2/detailed_hw_partitions` element to request information related to the file system that is mounted on a scanned system: Unix mount points and Windows logical drives.

Permissions

 You must have the View Endpoints and View Hardware Inventory permission to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/detailed_hw_partitions?token=token
```

Resource information

Table 189. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	GET
Request headers	Header Accept-Language (optional)

Table 189. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
	<p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response payload	<code>Partitions</code> element
Response format	<code>application/json</code>
Response codes	<p>200 - OK</p> <p>500 - "Bad Request" if a query parameter contains errors or is missing</p>

Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/detailed_hw_partition.json?token=token
```

Available columns

The list of attributes related to the file system that is mounted on a scanned system: Unix mount points and Windows logical drives.

The scan output might contain multiple instances of each attribute from this group.

Table 190. Columns with information about partitions

Property	Description	Type
id	Identifier of the record.	Integer
computer_id	Identifier of the computer as specified in the BigFix Inventory database.	Integer
updated_at	Date and time when the current information was imported to BigFix Inventory.	String
type	The type of the partition; for example, Logical Drive or Remote Drive.	String
media_type	The media type that contains the partition; for example, Local Disc or CD-ROM.  Note: The media type is Unknown when the scanner cannot recognize it.	String
device_name	The name of the device or a label of a Windows drive; for example, <code>/dev/dsk/c1t3d0s0</code> or Storage.	String
physical_size	The size of a drive that contains the partition.	String
fs_total_size	Total size of the partition in kilobytes (KB).	String
fs_free_size	The amount of free space on the partition in kilobytes (KB).	String

Related APIs

This API can be associated with the following API:

- [Retrieval of hardware inventory \(v2\) \(on page mxxx\)](#)

Query parameters

Table 191. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve the computer id: <pre>URL?columns[]=computer_id</pre>		String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by computer ID descending: <pre>URL?order[]=computer_id desc</pre>		String
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved. Example: Retrieve 100 records <pre>URL?limit=100</pre>  Note: By default the limit parameter for this API is set to 100000.		Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150 <pre>URL?limit=50&offset=150</pre>		Numeric
token	A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dccc1xxii) . You can also log in to BigFix Inventory, hover over the User icon  , and click Profile . Then, click Show token .	✓	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written in one line:		String

Table 191. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<pre> <criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criteria> [{ <comma> <criteria> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criteria> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null> </pre>		
	<p>For more information about operators, see: Common connectors and operators (on page dcccclxii).</p> <p>Example: retrieve computers with computer ID greater than 1000.</p> <pre> URL?criteria={"and":[{"computer_id", ">", "1000"}]} </pre>		
	<p>For columns that use the date and time values, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <operator>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.</p>		

Example conversation - default columns

Request

```

GET api/sam/v2/detailed_hw_partitions??token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US

```

Response body

```

[ {
  "id": 4,

```

```

"computer_id":1,
"updated_at":"2018-05-28T01:22:16Z",
"type":"Logical Drive",
"media_type":"Local Disk",
"device_name":"tmpfs",
"physical_size":"4005420",
"fs_total_size":"4005420",
"fs_free_size":"4005408"
}]

```

Example conversation - selected columns

Request

```

GET api/sam/v2/detailed_hw_partitions?columns[]=id
&columns[]=media_type&columns[]=physical_size
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US

```

Response body

```

[ {
  "id":4,
  "media_type":"Local Disk",
  "physical_size":"4005420"
}]

```

Retrieval of information related to physical processors (v2)

9.2.12 Available from 9.2.12. You use the `GET` operation on the `api/sam/v2/detailed_hw_physical_processors` element to request information related to all active physical processors.

Permissions

 You must have the View Endpoints and View Hardware Inventory permission to use this API.

Resource URL

```

https://hostname:port/api/sam/v2/detailed_hw_physical_processors?token=token

```

Resource information

Table 192. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	GET
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	application/json
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response payload	Physical Processors element
Response format	application/json
Response codes	<p>200 - OK</p> <p>500 - "Bad Request" if a query parameter contains errors or is missing</p>

Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/detailed_hw_physical_processor.json?token=token
```

Available columns

The list of attributes related to all active physical processors including single chips and processor packages. You can retrieve values for each platform working as a host which translates into one operating system per physical processor.



Note: To retrieve information about physical processors for guest operating systems, such as VMware or BigFix LPAR, the complete information for the physical machine must be available.

The scan output might contain multiple instances of each attribute from this group.

Table 193. Columns with information about physical processors

Property	Description	Type
id	Identifier of the record.	Integer
computer_id	Identifier of the computer as specified in the BigFix Inventory database	Integer
updated_at	Date and time when the current information was imported to BigFix Inventory.	String
core_per_package_count	The number of cores on a physical processor; for example, 8 cores for Intel Xeon E5-4620, and 12 for POWER9.	String
logical_proc_per_core	The number of logical processors assigned for each core; for example, 2 for Intel Xeon E5-4620, and 8 for POWER9.	String
manufacturer	Processor manufacturer; for example Intel, AMD, or BigFix.  Note: You can only retrieve information about the manufacturers that are supported by the scanner. Otherwise, the displayed value is Unknown.	String
family	Processor family; for example, Xeon, or POWER9.	String

Table 193. Columns with information about physical processors (continued)

Property	Description	Type
	 Note: You can only retrieve information about the processor families that are supported by the scanner. Otherwise, the displayed value is Unknown.	
type	<p>Processor type; for example, E5-4620 for Intel Xeon E5-4620.</p>  Note: If the scanner does not support type identification of the processor family the value is empty. If the scanner supports type identification for the processor family, however, the given types is not recognize, the displayed value is Unknown.	String
cpu_freq	<p>Processor frequency in megahertz (MHz). This information is optional, and vendor-specific. Depending on the processor family the value might represent either nominal or current CPU frequency.</p>  Note: CPU frequency might be interpreted differently across processor architectures and families. If this information is not applicable the value is 0.	String
brandname	<p>Processor name. This information is optional, and specific to platform and vendor. Thus, it can be retrieved from a different source in case of each platform, or even processor family. If available, it is a copy of a value reported by CPUID instruction, BIOS or OS.</p>	String
active_processor_count	<p>The number of active processors in the physical processor. A processor is active if a logical processor is running on this physical processor.</p>	String
active_core_count	<p>The number of active cores in the physical processor. The core is active if a logical processor is running on this core.</p>	String

Related APIs

This API can be associated with the following API:

- [Retrieval of hardware inventory \(v2\) \(on page mxxx\)](#)

Query parameters

Table 194. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve the computer id: <pre>URL?columns[]=computer_id</pre>		String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by computer ID descending: <pre>URL?order[]=computer_id desc</pre>		String
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved. Example: Retrieve 100 records <pre>URL?limit=100</pre>  Note: By default the limit parameter for this API is set to 100000.		Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150 <pre>URL?limit=50&offset=150</pre>		Numeric
token	A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dccc1xxii) . You can also log in to BigFix Inventory, hover over the User icon  , and click Profile . Then, click Show token .	✓	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written in one line:		String

Table 194. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<pre> <criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criteria> [{ <comma> <criteria> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criteria> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null> </pre>		
	<p>For more information about operators, see: Common connectors and operators (on page dcccclxii).</p> <p>Example: retrieve computers with computer ID greater than 1000.</p> <pre> URL?criteria={"and":[{"computer_id", ">", "1000"}]} </pre>		
	<p>For columns that use the date and time values, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <i><operator></i>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.</p>		

Example conversation - default columns

Request

```

GET api/sam/v2/detailed_hw_physical_processors?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US

```

Response body

```

[ {
  "id": 7,

```

mlxxx

```
"computer_id":4,
"updated_at":"2018-05-30T01:15:05Z",
"core_per_package_count":"1",
"logical_proc_per_core":"1",
"manufacturer":"Intel",
"family":"Xeon",
"type":"E7-8880",
"cpu_freq":"2200",
"brandname":"Intel(R) Xeon(R) CPU E7-8880 v4 @ 2.20GHz",
"active_processor_count":"1",
"active_core_count":"1"
}]
```

Example conversation - selected columns

Request

```
GET api/sam/v2/detailed_hw_physical_processors?columns[]=computer_id
&columns[]=manufacturer&columns[]=family&columns[]=type
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
"computer_id":4,
"manufacturer":"Intel",
"family":"Xeon",
"type":"E7-8880"
}]
```

Retrieval of information related to SMBIOS (v2)

9.2.12 Available from 9.2.12. You use the `GET` operation on the `api/sam/v2/detailed_hw_smbios` element to request information related to SMBIOS of the scanned system.

Permissions

 You must have the View Endpoints and View Hardware Inventory permission to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/detailed_hw_smbios?token=token
```

Resource information

Table 195. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	GET
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	application/json
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response payload	SMBIOS element
Response format	application/json
Response codes	<p>200 - OK</p> <p>500 - "Bad Request" if a query parameter contains errors or is missing</p>

Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/detailed_hw_smbios.json?token=token
```

Available columns

The list of attributes related to SMBIOS of the scanned system.

The scan output contains a single instance of each attribute from this group.

Table 196. Columns with information about SMBIOS

Property	Description	Type
id	Identifier of the record.	Integer
computer_id	Identifier of the computer as specified in the BigFix Inventory database	Integer
updated_at	Date and time when the current information was imported to BigFix Inventory.	String
<div style="display: flex; align-items: center;"> <div style="background-color: #800040; color: white; padding: 2px 5px; margin-right: 5px;">Linux</div> <div style="background-color: #800040; color: white; padding: 2px 5px; margin-right: 5px;">Windows</div> <div style="margin-right: 5px;">bios_</div> <div>vendor</div> </div>	The manufacturer of the BIOS system.	String
<div style="display: flex; align-items: center;"> <div style="background-color: #800040; color: white; padding: 2px 5px; margin-right: 5px;">Linux</div> <div style="background-color: #800040; color: white; padding: 2px 5px; margin-right: 5px;">Windows</div> <div style="margin-right: 5px;">bios_</div> <div>version</div> </div>	The version of the BIOS system.	String
<div style="display: flex; align-items: center;"> <div style="background-color: #800040; color: white; padding: 2px 5px; margin-right: 5px;">Linux</div> <div style="margin-right: 5px;">sys-</div> <div>tem_version</div> </div>	The version of the system.	String
<div style="display: flex; align-items: center;"> <div style="background-color: #800040; color: white; padding: 2px 5px; margin-right: 5px;">Linux</div> <div style="background-color: #800040; color: white; padding: 2px 5px; margin-right: 5px;">Windows</div> <div style="margin-right: 5px;">sys-</div> <div>tem_serial_number</div> </div>	The serial number of the system.	String

Related APIs

This API can be associated with the following API:

- [Retrieval of hardware inventory \(v2\) \(on page mxxx\)](#)

Query parameters

Table 197. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve the computer id: <pre>URL?columns[]=computer_id</pre>		String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by computer ID descending: <pre>URL?order[]=computer_id desc</pre>		String
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved. Example: Retrieve 100 records <pre>URL?limit=100</pre>  Note: By default the limit parameter for this API is set to 100000.		Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150 <pre>URL?limit=50&offset=150</pre>		Numeric
token	A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dccc1xxii) . You can also log in to BigFix Inventory, hover over the User icon  , and click Profile . Then, click Show token .	✓	Alphanumeric
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written in one line:		String

Table 197. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
	<pre> <criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket> <criteria> [{ <comma> <criteria> }...] <right-bracket> <right-brace> <boolean-operator> ::= "and" "or" <criteria> ::= <criteria> <left-bracket> <column> <comma> <operator> <comma> <value> <right-bracket> <column> ::= <json-string> <operator> ::= <json-string> <value> ::= <json-array> <json-string> <json-numver> <json-null> </pre>		
	<p>For more information about operators, see: Common connectors and operators (on page dccclxii).</p> <p>Example: retrieve computers with computer ID greater than 1000.</p> <pre> URL?criteria={"and":[{"computer_id", ">", "1000"}]} </pre>		
	<p>For columns that use the date and time values, you can retrieve data also for a period instead of a specific date. To do so, use <i>last</i> or <i>next</i> as <operator>, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.</p>		

Example conversation - default columns

Request

```

GET api/sam/v2/detailed_hw_smbios?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US

```

Response body

```

[ {
  "id":12,

```

```

"computer_id":19,
"updated_at":"2018-05-28T01:22:16Z",
"bios_vendor":"Phoenix Technologies LTD",
"bios_version":"6.00",
"system_version":"None",
"system_serial_number":"VMware-42 14 ea ea 03 c0 d5 03-fd f6 43 d6 44 a2 f8 52"
}]

```

Example conversation - selected columns

Request

```

GET api/sam/v2/detailed_hw_smbios?columns[]=computer_id
&columns[]=bios_vendor&columns[]=bios_version
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US

```

Response body

```

[ {
  "computer_id":19,
  "bios_vendor":"Phoenix Technologies LTD",
  "bios_version":"6.00"
} ]

```

Retrieval of information related to storage (v2)

9.2.12 Available from 9.2.12. You use the `GET` operation on the `api/sam/v2/detailed_hw_storages` element to request information related to storage devices on the scanned system.

Permissions



You must have the View Endpoints and View Hardware Inventory permission to use this API.

Resource URL

```

https://hostname:port/api/sam/v2/detailed_hw_storages?token=token

```

Resource information

Table 198. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	GET
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	application/json
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response payload	Storage Device element
Response format	application/json
Response codes	<p>200 - OK</p> <p>500 - "Bad Request" if a query parameter contains errors or is missing</p>

Schema description

To retrieve the list of all columns that are returned by this REST API together with their descriptions, use the following request.

```
GET api/sam/v2/schemas/detailed_hw_storage.json?token=token
```

Available columns

The list of attributes related to storage devices on the scanned system.



Note:

AIX The detailed hardware scan cannot retrieve information about storage for AIX WPARs because of the specification of these virtual machines.

Windows When you set up a quota on a non-administrative user account, the free disk space might be greater than the space that is available on a disk volume. Thus, the storage information might not be accurate. According to the Microsoft policy, you must be an administrator to retrieve information about storage quota that is allocated to other users.

The scan output might contain multiple instances of each attribute from this group.

Table 199. Columns with information about storage devices

Property	Description	Type
id	Identifier of the record.	Integer
computer_id	Identifier of the computer as specified in the BigFix Inventory database	Integer
updated_at	Date and time when the current information was imported to BigFix Inventory.	String
type	The type of the storage device.	String
model	The model of the storage device.	String
manufacturer	The manufacturer of the storage device.	String
total_size	The size of the hard drive in megabytes (MB).	String
device	The unique system identifier of the storage device. This information is optional.	String

Related APIs

This API can be associated with the following API:

- [Retrieval of hardware inventory \(v2\) \(on page mxxx\)](#)

Query parameters

Table 200. Query parameters

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
columns[]	Specify which columns to retrieve. If you do not specify this parameter, only default columns are retrieved. Example: Retrieve the computer id: <pre>URL?columns[]=computer_id</pre>		String
order	Specify how to sort the returned data. The default direction for sorting columns is ascending. If you want to specify a descending sort, append <code>desc</code> to the column name. Example: Order by computer ID descending: <pre>URL?order[]=computer_id desc</pre>  Note: By default the limit parameter for this API is set to 100000.		String
limit	Specify the number of rows to retrieve. If you omit this parameter, all rows are retrieved. Example: Retrieve 100 records <pre>URL?limit=100</pre>		Numeric
offset	Specify the number of rows to skip for retrieving results. You can use it together with the limit parameter to paginate results. Example: Retrieve 50 records starting after record 150 <pre>URL?limit=50&offset=150</pre>		Numeric
token	A unique user authentication identifier. You can retrieve it by using REST API for retrieving authentication token (on page dcccclxxii) . You can also log in to BigFix Inventory, hover over the User icon  , and click Profile . Then, click Show token .	✓	Alphanumeric

Table 200. Query parameters

The table consists of four columns and nine rows.

(continued)

Parameter	Description	Required	Value
criteria	Retrieve records which match specific conditions. The parameter should have the following structure, written in one line:		String

```
<criteria> ::= <left-brace> <boolean-operator><colon> <left-bracket>
<criteria> [{ <comma> <criteria> }...] <right-bracket>
<right-brace>
<boolean-operator> ::= "and" | "or"
<criteria> ::= <criteria> | <left-bracket> <column> <comma>
<operator> <comma> <value> <right-bracket>
<column> ::= <json-string>
<operator> ::= <json-string>
<value> ::= <json-array> | <json-string> | <json-numver> |
<json-null>
```

For more information about operators, see: [Common connectors and operators \(on page dccclxii\)](#).

Example: retrieve computers with computer ID greater than 1000.

```
URL?criteria={"and":[[ "computer_id", ">", "1000"]]}
```

For columns that use the date and time values, you can retrieve data also for a period instead of a specific date. To do so, use *last* or *next* as **<operator>**, and then specify the time value in the following convention: PxD/PxW/PxM/PxY, where x is a number in the 1-999 range, and D, W, M, or Y is a designator that represents days, weeks, months, or years respectively.

Example conversation - default columns

Request

```
GET api/sam/v2/detailed_hw_storages?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "id":1,
  "computer_id":1,
  "updated_at":"2018-05-28T01:22:16Z",
  "type":"Hard Disk",
  "model":"Virtual disk",
  "manufacturer":"VMware",
  "total_size":"40960",
  "device" "/dev/sda"
}]
```

Example conversation - selected columns

Request

```
GET api/sam/v2/detailed_hw_storages?columns[]=computer_id
&columns[]=model&columns[]=manufacturer
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response body

```
[{
  "computer_id":1,
  "model":"Virtual disk",
  "manufacturer":"VMware"
}]
```

REST API for setting and deleting the end of support date

9.2.11 Available from 9.2.11. You can set and delete the end of support date for any software component by using REST API.

Setting the end of support date for software components

9.2.11 Available from 9.2.11. You use the `POST` operation on the `sam/component_lifecycles` element to set the end of support date for software components. The end of support date is added on the Products & Metrics report. If you want to change the end of support date, you must first delete the current date and then set a new one by using this API.

Permissions



You must have the Manage Catalogs permission to use this API.

Resource URL

```
https://hostname:port/sam/component_lifecycles?token=token
```



Note: The URL is going to change in the future. If you are going to use it for integration, check for updates in future releases.

Resource information

Table 201. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	POST
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	application/json
Request payload	<pre>{ "component_lifecycle": { "eos_date": "YYYY-MM-DD", "software_component_id": "software_component_id" } }</pre>
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p>

Table 201. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
	Specifies the language of the response content. If the header is not specified, the content is returned in the server language.
Response format	application/json
Response codes	200 - OK
	422 - "Unprocessable entity" if a parameter contains errors or is missing

Example conversation - setting end of support date for a software component

- To retrieve information about the software component ID for which you want to set the end of support date, perform the following steps:
 - Log in to BigFix Inventory.
 - Go to **Reports > Software Components**.
 - Hover over the component name and check the URL address at the bottom of the page. The software component ID is a number at the end of the URL. For example: `sam/catalog/software_components/10379`
- To set the end of support date, use the following `POST` request.

Request header

```
POST sam/component_lifecycles?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Request Body

```
{ "component_lifecycle":
  {
    "eos_date": "2020-12-22",
    "software_component_id": "10379"
  }
}
```

Response Header

```
200 - OK
```

Response Body

```
{
  "created_at": "2018-01-26T10:30:06Z",
  "detailed_version": null,
  "eos_date": "2020-12-23",
  "id": 2,
  "is_release": true,
  "software_component_id": 10379,
  "source_id": 1,
  "source_type": 0,
  "updated_at": "2018-01-26T10:30:06Z"
}
```

Deleting the end of support date for software components

9.2.11 Available from 9.2.11. You use the `DELETE` operation on the `sam/software_components/software_component_id/component_lifecycles/` element to remove the end of support date for a software component.

Permissions

 You must have the Manage Catalogs permission to use this API.

Resource URL

```
https://hostname:port/sam/software_components/software_component_id/component_lifecycles?token=token
```

Resource information

Table 202. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	<code>DELETE</code>
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>

Table 202. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
Request format	<code>application/json</code>
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p><code>en-US, ...</code></p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response format	<code>application/json</code>
Response codes	<code>200 - OK</code>

Example conversation - clearing the end of support date for a software component

- To retrieve information about the software component ID for which you want to clear the end of support date, perform the following steps:
 - Log in to BigFix Inventory.
 - Go to **Reports > Software Components**.
 - Hover over the component name and check the URL address at the bottom of the page. The software component ID is a number at the end of the URL. For example: `sam/catalog/software_components/10379`
- To remove the end of support date, use the following `DELETE` request.

Request header

```
DELETE
sam/software_components/10379/component_lifecycles?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
Host: localhost:9081
```

```
Accept: application/json
Accept-Language: en-US
```

Response

```
200 - OK
```

REST API for adding and removing tags

9.2.14 Available from 9.2.14. You can add and remove tags for any software component by using REST API.

Adding tags for software components

9.2.14 Available from 9.2.14. You use the `POST` operation on the `api/sam/v2/tags/software_components` element to add tags for software components.

Permissions



You must have the Manage Catalogs permission to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/tags/software_components?token=token
```

Resource information

Table 203. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	<code>POST</code>
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	<code>application/x-www-form-urlencoded</code>
Request payload	<code>ids=id_1, id_2, id_3&tags=tag_1, tag_2, tag_3</code>
Response headers	Header

Table 203. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
	<p><code>Content-Type</code></p> <p>Values</p> <p><code>application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p> <p>Values</p> <p><code>en-US, ...</code></p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response format	<code>application/json</code>
Response codes	<p><code>200 - OK</code></p> <p><code>400 - Bad request</code></p>

Example conversation - adding tags for the selected software components

Note: The length of a single tag is limited to 254 characters.

Request

```
POST
  api/sam/v2/tags/software_components?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623&ids=1,2,3&tags=tag1,tag2
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response

```
{
  "updated_components": 3
}
```

Removing tags for software components

9.2.14 Available from 9.2.14. You use the `DELETE` operation on the `api/sam/v2/tags/software_components` element to remove the previously added tags.

Permissions

 You must have the Manage Catalogs permission to use this API.

Resource URL

```
https://hostname:port/api/sam/v2/tags/software_components?token=token
```

Resource information

Table 204. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	DELETE
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	application/x-www-form-urlencoded
Request payload	ids=id_1, id_2, id_3&tags=tag_1, tag_2, tag_3
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>application/json</p> <p>Specifies the content type of the response.</p>

Table 204. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
	<p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response format	application/json
Response codes	<p>200 - OK</p> <p>400 - "Bad request"</p>

Example conversation - removing tags for the selected software components



Note: The length of a single tag is limited to 254 characters.

Request

```
DELETE
api/sam/v2/tags/software_components?token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623&ids=1,2,3&tags=tag1,tag2
Host: localhost:9081
Accept: application/json
Accept-Language: en-US
```

Response

```
{
  "updated_components": 3
}
```

REST API for export and import of saved report views

You can use REST API requests to export a saved report view from one instance of BigFix Inventory and import it to another instance of the application.

Exporting and importing a saved report view

To export a saved report view to another instance of BigFix Inventory, check the ID of the report view in the source instance of BigFix Inventory. Then, use the GET operation to retrieve the definition of the saved report view. Next, use the POST operation to import it to another instance of BigFix Inventory.

The following procedure uses cURL command-line tool for negotiating API requests.

1. Check the identifier of the saved report view.
 - a. Log in to the instance of BigFix Inventory from which you want to export the saved report view.
 - b. In the top navigation bar, click **Reports > Saved Reports**. Open the saved report view and check the last number in the report URL.
In the following example, the report ID is 2.

```
http://server_host_name:port_number/sam/pvuonlysubcapreports#32fe0f54dc719893faacc1d0f38a0c9045863729/2
```

2. Obtain API tokens from both instances of BigFix Inventory.
 - a. Hover over the **User** icon , and click **Profile**.
 - b. In the API Token line, click **Show token**.
3. Open the command line interface and change to the location where cURL is installed.
4. To export the report view, use the following GET request:

```
curl -o C:\saved_reports\report_definitions\pvu_subcapacity.txt
-X GET server_host_name:port_number/api/reports/report_ID?
token=5edd5aey7cd91467h08450bc258c31f0ce706543
```

Where:

-o

Specifies the path to the file where the definition of the saved report view is to be saved.

-X

Specifies the type of HTTP request.

report_ID

Specifies the identifier of the saved report view that is to be exported.

If the request is successful, the following message is displayed:

```
HTTP/1.1 200 OK
```

5. To import the saved report view to another instance of BigFix Inventory, use the following POST request:

```
curl -H "Content-Type: application/json"
-X POST -d "@C:\saved_reports\report_definitions\pvu_subcapacity.txt"
```

```
server_host_name:port_number/api/reports?
token=5cd3gh78499496e89a3246ab343474e85d8bc8fc
```

Where:

- H**
Specifies the header of the request.
- X**
Specifies the type of HTTP request.
- d**
Specifies the path to the file where the definition of the saved report view was saved.

The saved report view was imported to the target instance of BigFix Inventory.

Export of saved report views

You use the `GET` operation on the `api/reports/report_ID` element to export a saved report view which can be then imported to another instance of BigFix Inventory.

Table 205. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>GET /api/reports/report_ID</code>
Purpose	Returns the saved report view with the specified ID.
HTTP method	<code>GET</code>
Resource URI	<code>https://server_host_name:port_number/api/reports/report_ID</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmviii) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>

Table 205. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
Request payload	n/a
Request Content-Type	Application/json
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>Application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	Saved reports element
Response Content-Type	Application/json
Normal HTTP response codes	200 - OK
Error HTTP response codes	<p>401- There is no match for the provided user name and password</p> <p>401 - You are not assigned a Computer Group. You will not be able to access the system until you are assigned a valid Computer Group. Contact your administrator for assistance</p> <p>404 - Sequel::RecordNotFound</p>
	<p>For more details about each error code, check the <code>tema.log</code> log file in the <code>install_dir/wlp/usr/servers/server1/logs/</code> directory.</p>

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `api/reports` element.

Table 206. Query parameters for retrieving saved report views

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
token	A unique user authentication identifier.	Yes	Alphanumeric
id	An identifier of the saved report. To find the report identifier, open the saved report and check the last number in the report URL. In the following example, the report ID is 3.	No	Numeric

```
server_host_name:port/sam/app_usage_property_val
ues#cafcac39cd0b242b82729377b0b2b872d3af8d8a/3
```

Example HTTP conversation

Request

```
GET/api/reports/1
&token=a070afeca411d2fs25f5s4c962de2d002b14352f
```

Response header

```
HTTP/1.1 200 OK
Content-Type: Application/json
Content-Language: en-US
```

Response body (JSON)

```
{
  "id":1,
  "user_id":1,
  "pagestate_id":"eb5743b7cd23316f0bc8a4dd3f63b90d73549042",
  "name":"pvu_test",
  "path":"/sam/pvuonlysubcapreports",
  "private":true,
  "state":{
    "columns":["product","quantity","hwm_history"],
    "criteria":{"and":[{"quantity",">","200"}]},
    "grid_options":["autosize_columns"],
```

```

"order":{"asc":true,"col":null},
"time_range":{"
  "max":"2014-07-21T10:18:04Z",
  "min":"2014-04-21T22:00:00Z",
  "type":"absolute_to_now",
  "units":"days",
  "value":"90"
},
"column_order":{"
  "hwm_history":1,
  "product":0,
  "quantity":2
},
"column_widths":{"
  "hwm_history":180,
  "product":524,
  "quantity":719
}}

```

Import of saved report views

You use the `POST` operation on the `api/reports` element to import a saved report view from another instance of BigFix Inventory.

Table 207. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

Operation details	Description
Operation	<code>POST /api/reports</code>
Purpose	Imports a saved report view from another instance of BigFix Inventory.
HTTP method	<code>POST</code>
Resource URI	<code>https://server_host_name:port_number/api/reports</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page cmviii) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p>

Table 207. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
	<p>en-US (only English is supported)</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p>
Request payload	Saved reports element
Request <small>Content-Type</small>	<p><code>Application/json</code></p>
Response headers	<p>Header</p> <p><small>Content-Type</small></p> <p>Values</p> <p><code>Application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><small>Content-Language</small></p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If this header is not specified, the content is returned in the server language.</p>
Response payload	n/a
Response <small>Content-Type</small>	<p><code>Application/json</code></p>
Normal HTTP response codes	<p>200 - OK</p>
Error HTTP response codes	<p>401- There is no match for the provided user name and password</p> <p>401 - You are not assigned a Computer Group. You will not be able to access the system until you are assigned a valid Computer Group. Contact your administrator for assistance.</p> <p>404 - Not Found</p>

Table 207. Operation descriptions

The table consists of two columns and 15 rows. For the eleventh row there are two levels.

(continued)

Operation details	Description
	404 - Sequel::RecordNotFound
	500 - Name is already taken
	500 - There was a problem with your request
	500 - We're sorry, but something went wrong. Please contact your BigFix Analytics administrator if the problem persists
	500 if the JSON format is invalid, a detailed message with the problem explanation is also returned. For example, if a name field is missing, the following message is displayed: The property '#/' did not contain a required property of 'name' in schema 69e40f9b-9a3e-53d1-a7d8-61d63eb191e8#
	For more details about each error code, check the <code>tema.log</code> log file in the <code>install_dir/wlp/usr/servers/server1/logs/</code> directory.

Query parameters

You can use query parameters to narrow down the results of your search. The following table presents query parameters that you can use for the `api/reports` element.

Table 208. Query parameters for retrieving saved report views

The table consists of four columns and nine rows.

Parameter	Description	Required	Value
<code>token</code>	A unique user authentication identifier.	Yes	Alphanumeric

Example HTTP conversation

Request header

```
Content-Type: application/json
Content-Language: en-US
```

Request

```
POST /api/reports?token=465c33848de3db7ch5699023ea22deb5b1a476d1
{
  "id":4,
  "user_id":2,
```

```
"pagestate_id": "7ebb0b0a018ab55fab8e6e40d5eb62529ea38fad",
"name": "pvu_report",
"path": "/sam/pvuonlysubcapreports",
"private": true,
"state": {
  "columns": ["product", "quantity", "hwm_history"],
  "criteria": {"and": [{"quantity", ">", "200"}]},
  "grid_options": ["autosize_columns"],
  "order": {
    "asc": true,
    "col": "quantity"
  },
  "time_range": {
    "max": "2014-07-21T10:18:04Z",
    "min": "2014-04-21T22:00:00Z",
    "type": "absolute_to_now",
    "units": "days",
    "value": "90"
  },
  "column_order": {
    "hwm_history": 1,
    "product": 0,
    "quantity": 2
  },
  "column_widths": {
    "hwm_history": 180,
    "product": 524,
    "quantity": 719
  }
}}
```

Response

```
HTTP/1.1 200 OK
```

REST API for import of contracts

You can use REST API requests to import contracts from a csv file.

Import of contracts

You use the `POST` operation on the `api/sam/contracts/import` element to import contracts from a csv file



You must have the Manage Software Classification permission to perform this task.

Table 209. Operation descriptions

The table consists of two columns and 14 body rows.

Operation details	Description
Operation	<code>/api/sam/contracts/import</code>
Purpose	Uploads the contents of the csv file with contract definitions
HTTP method	POST
Resource URI	<code>https://server_host_name:port_number/api/sam/contracts/import</code>
URL link relation	n/a
URI query parameters	For a list of applicable query parameters, see: Query parameters (on page mcviii) .
Request headers	<p>Header</p> <p><code>Accept-Language</code> (optional)</p> <p>Values</p> <p>en-US, ...</p> <p>Used to negotiate the language of the response. If this header is not specified, the content is returned in the server language.</p> <p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p>multipart/form-data</p>
Request payload	Contract definitions as form-encoded data specified as a file attachment.
Request <code>Content-Type</code>	multipart/form-data
Response headers	<p>Header</p> <p><code>Content-Type</code></p> <p>Values</p> <p><code>Application/json</code></p> <p>Specifies the content type of the response.</p> <p>Header</p> <p><code>Content-Language</code></p>

Table 209. Operation descriptions

The table consists of two columns and 14 body rows.

(continued)

Operation details	Description
Values	
en-US, ...	
	Specifies the language of the response content. If this header is not specified, the content is returned in the server language.
Response payload	Status of imported elements
Response Content-Type	Application/json
Normal HTTP response codes	
	200 - No content
Error HTTP response codes	
	400 - "Bad Request" if a query parameter contains errors or is missing
	401 - "Unauthorized" if the user is not authorized to perform the operation
	Message body excludes an error message with details.

Query parameters

Use the **upload[contracts]** query parameter to specify the path to the csv file with contracts.

Table 210. Query parameters for specifying the path to the csv file with contracts

The table consists of four columns and two rows.

Parameter	Description	Required	Value
upload[contracts]	Path to the csv file with contracts	Yes	String

Example HTTP conversation

Request header

```
Content-Type: multipart/form-data;
```

Request

```
POST /api/sam/contracts/import?token=a2279dc218ecd1f12a5e2fc1b5d77a4e6fec6598
Contract Name,Software Names,Computer Group Name,Seats,Acquisition Cost,
Maintenance Cost,Entitlement Start,Entitlement End,Maintenance Start,
Maintenance End,custStrcon1,"IBM Software Assembly Toolkit,
License Metric Tool 7,License Metric Tool 9.0",All Computers,
10,10.0,5.0,2015-04-23,2015-05-30,2015-04-23,2015-05-30,con2,
BigFix Inventory, All Computers,0,1.0,2.0,2015-04-23,,,
```

Response header

```
HTTP/1.1 200 OK
Content-Type: application/json
```

Response

```
{
  "status":[{
    "type":"CONTRACTS",
    "description":{
      "contract_name":"con1",
      "software_names":"IBM Software Assembly Toolkit,
      License Metric Tool 7,
      License Metric Tool 9.0",
      "user_name":"admin",
      "access_computer_group_name":"All Computers",
      "computer_group_name":"All Computers",
      "seats":"10",
      "acquisition_cost":"10.0",
      "annual_maintenance_cost":"5.0",
      "entitlement_start":"2015-04-23",
      "entitlement_end":"2015-05-30",
      "maintenance_start":"2015-04-23",
      "maintenance_end":"2015-05-30",
      "custom_fields":{
        "custStr":null
      }
    },
    "actions":[{
      "error":"","
      "status":"SUCCESS",
      "type":"ADD"
    }
  ]
},{
```

```

"type": "CONTRACTS",
"description": {
  "contract_name": "con2",
  "software_names": "BigFix Inventory",
  "user_name": "admin",
  "access_computer_group_name": "All Computers",
  "computer_group_name": "All Computers",
  "seats": null,
  "acquisition_cost": "1.0",
  "annual_maintenance_cost": "2.0",
  "entitlement_start": "2015-04-23",
  "entitlement_end": null,
  "maintenance_start": null,
  "maintenance_end": null,
  "custom_fields": {
    "custStr": null
  }
},
"actions": [
  {
    "error": "",
    "status": "SKIPPED",
    "type": "ADD"
  }
]
}

```

Importing contracts from a file in CSV format

You can use the cURL command-line tool to issue a POST request to import a csv file with contracts.

The following procedure uses cURL a command-line tool for running REST API requests.

1. Create a text file with the following CSV format.

```

Contract Name,Software Names,Computer Group Name,Seats,Acquisition Cost,
Maintenance Cost,Entitlement Start,Entitlement End,Maintenance Start,
Maintenance End,Contract Custom Fields...

```

where

Contract Name

The name of the contract. It must be unique. If the contract with a given name exists in BigFix Inventory, the importing of that row with contract definition is skipped.

Software Names

Names of software products, versions or/and releases that are separated by a comma.

Computer Group Name

Name of the computer group. The computer group with given name must exist in BigFix Inventory. Otherwise, the import of that row fails.

Seats

Number of users who have access to the software product under a particular contract. In case of an ELA license, the column must be empty. (optional)

Acquisition Cost

The cost of purchasing a software product.

Maintenance Cost

The cost of purchasing maintenance plan for a software product.

Entitlement Start

The start date for a particular contract.

Entitlement End

The end date for a particular contract. (optional)

Maintenance Start

The maintenance plan start date for the software products. (optional)

Maintenance End

The maintenance plan end date for the software products. (optional)

Contract Custom Fields

The list of contract custom fields values separated by a comma. All custom fields must exist in BigFix Inventory. Otherwise, the import of that row fails. Moreover, all custom fields that are required in BigFix Inventory must be listed in the file.

Example

```
Contract Name,Software Names,Computer Group Name,Seats,Acquisition Cost,
Maintenance Cost,Entitlement Start,Entitlement End,Maintenance Start,
Maintenance End,custom integer,custom boolean
Name1,"BigFix Software Assembly Toolkit,License Metric Tool 7,
License Metric Tool 9.0",All Computers,10,10.0,5.0,2015-04-24,
2015-05-31,2015-04-24,2015-05-31,1,true
Name2,BigFix Inventory,All Computers,0,1.0,2.0,2015-04-24,,,,,false
```

2. Obtain API tokens from both instances of BigFix Inventory.

- a. Hover over the **User** icon  and click **Profile**.
- b. In the API Token line, click **Show token**.

- Open the command-line interface and change to the location where cURL is installed.
- To import a file with contracts into BigFix Inventory, use the following POST request.

```
curl -H "Content-Type: multipart/form-data" -X POST -F
"upload[contracts]=@path_to_file" http://server_host_name:port_number/
api/sam/contracts/import?token=token
```

Where:

@path_to_file

Specifies the path to the file where the definition of contracts was saved.

server_host_name

Address of the BigFix Inventory server.

port_number

The communications port number.

token

API token from profile view.

Example:

```
curl -H "Content-Type: multipart/form-data" -X POST -F
"upload[contracts]=@C:\temp\contracts.csv" http://localhost:9081/
api/sam/contracts/import?token=a2279dc218ecd1f12a5e2fc1b5d77a4e6fec6593
```

If the request is successful, a response in JSON format is displayed.

You successfully imported contracts from a report in CSV format.

REST API for updating the Oracle core factor

9.2.8

Available from 9.2.8. You use the `PUT` operation on the `api/v1/servers` element to change the processor core factor that is applied to a group of servers for the purpose of calculating utilization of the Oracle Processor Core metric.

Permissions



You must have the Manage Hardware Inventory permissions to use this API.

Resource URL

```
https://hostname:port/api/v1/servers?token=token
```

For more information about using the REST API, see: [Changing the Oracle core factor through REST API \(on page dlxxiii\)](#).

Resource information

Table 211. Resource information

The table consists of two columns and seven rows.

Operation details	Description
HTTP method	PUT
Request headers	<p>Header</p> <p>Accept-Language (optional)</p> <p>Values</p> <p>en-US (only English is supported)</p> <p>Negotiates the language of the response. If the header is not specified, the content is returned in the server language.</p>
Request format	application/json
Request payload	Oracle Core Factors element
Response headers	<p>Header</p> <p>Content-Type</p> <p>Values</p> <p>application/json</p> <p>Specifies the content type of the response.</p> <p>Header</p> <p>Content-Language</p> <p>Values</p> <p>en-US, ...</p> <p>Specifies the language of the response content. If the header is not specified, the content is returned in the server language.</p>
Response format	application/json
Response codes	<p>200 - OK</p> <p>400 - "Bad Request" if a query parameter contains errors or is missing</p> <p>401 - "Unauthorized" if the user is not authorized to perform the operation</p>

Table 211. Resource information

The table consists of two columns and seven rows.

(continued)

Operation details	Description
	422 - "Unprocessable Entity" if the request is well-formed but cannot be processed due to semantic errors

Available columns

Table 212. Available columns

The table consists of three columns and 26 rows.

Column	Description	Required	Type
server_id	ID of the server for which you want to change the core factor.	✓	Nu-mer-ic
oracle_core_factor	Oracle core factor that you want to assign to server.	✓	Float

Example conversation

Request header

```
Content-Type: application/json
Content-Language: en-US
```

Response body

```
POST api/v1/servers?token=465c33848de3db7ch5699023ea22deb5b1a476d1
{
  "oracle_core_factors":
  [
    {
      "server_id": 1,
      "oracle_core_factor": 1.0
    },
    {
      "server_id": 2,
      "oracle_core_factor": 1.5
    }
  ]
}
```

REST API associations

9.2.8 Available from 9.2.8. For REST APIs in version 2, you can use associations to retrieve additional data as part of API requests. For example, you can retrieve detailed information about computer hardware as part of the `software_instances` API request.

Syntax

To retrieve the additional data, use the following syntax of the query:

```
<URL>?columns[ ]=<association>.<column>&token=<token>
```

Where:

- `<URL>` is the REST API request
- `<association>` is the name of the association from which you want to retrieve the column
- `<column>` is the name of the column that you want to retrieve
- `<token>` is the unique user authentication identifier

For example:

```
GET api/sam/v2/software_instances?columns[]=computer_health.catalog_version
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```



Note: Use API associations to retrieve information that cannot be obtained directly through an existing API.

component_cpe association

9.2.12 Available from 9.2.12. Use the `component_cpe` association to retrieve information about the discovered component as specified in the Common Platform Enumeration (CPE) dictionary as part of other REST API requests.

Association

`component_cpe`

Applicable REST APIs

You can use the `component_cpe` association with the following REST APIs:

- [/api/sam/v2/software_instances](#) (on page [cmxcii](#))

Syntax

```
<URL>?columns[ ]=component_cpe.<column>&token=<token>
```

Schema description

To retrieve the list of all columns that are returned by this association together with their descriptions, use the following request.

```
GET api/sam/schemas/associations/component_cpe.json?token=<token>
```

Available columns

Table 213. Columns with information about components as specified in the CPE standard

Property	Description	Type
cpe22	Information about the discovered component in the CPE v2.2 format.	String
cpe23	Information about the discovered component in the CPE v2.3 format.	String

Limitations

- BigFix Inventory does not support the following fields that are specified in the CPE standard: `update`, `edition`, `lang`.
- Marking of software as `o` (operating system) is automatic and might not always be accurate.
- Dots (.) and dashes (-) are not escaped as defined in the CPE standard. Otherwise, CPE information that is provided by BigFix Inventory could not be matched with the CVE data.

computer association

9.2.10 Available from 9.2.10. Use the `computer` association to retrieve information about details of the computer and custom computer properties as part of REST API requests.

Association

```
computer
```

Applicable REST APIs

You can use the `computer` association with the following REST APIs:

- [/api/sam/raw_file_facts](#) (on page [dcccclxxxiii](#))

Syntax

```
<URL>?columns[]=computer.<column>&token=<token>
```

For example:

```
GET api/sam/raw_file_facts?columns[]=computer.ip_address
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Schema description

To retrieve the list of all columns that are returned by this association together with their descriptions, use the following request. The list contains custom computer properties that were added on the **Computer Properties** panel.

```
GET api/sam/schemas/associations/computer.json?token=<token>
```

Available columns

Table 214. Columns with information about computer details

Column	Description	Type
id	Identifier of the computer as specified in the BigFix Inventory database.	Integer
remote_id	Identifier of the computer as specified in the BigFix database.	Integer
last_seen	Date and time when the BigFix client connected to the BigFix server the last time.	String
datasource_id	Identifier of the BigFix server to which the computer reports.	Integer
datasource_name	Name of the BigFix server to which the computer reports.	String
created_at	Date and time when the information about the computer was imported from BigFix for the first time.	String
valid_to	The date and time when the computer stopped being discovered, for example because it was removed from the infrastructure, the data source was removed, or the BigFix client was uninstalled.	String
os	Operating system of the computer.	String
dns_name	Name of the computer that is specified in the Domain Name System.	String
name	Host name of the computer.	String

Table 214. Columns with information about computer details (continued)

Column	Description	Type
<code>ip_address</code>	IP address of the computer.	String
<code>computer_property_number</code>	Custom computer property that was added on the Computer Properties panel. To view the list of all custom properties, view the computer schema (on page mcxvii) . The name of the property as displayed on the user interface is shown in the <code>title</code> parameter.	String

computer_details association

9.2.10

Available from 9.2.10. Use the `computer_details` association to retrieve information about details of the computer and custom computer properties as part of REST API requests.

Association

```
computer_details
```

Applicable REST APIs

You can use the `computer_details` association with the following REST APIs:

- [api/sam/v2/computers \(on page mxxx\)](#)

Syntax

```
<URL>?columns[]=computer_details.<column>&token=<token>
```

For example:

```
GET api/sam/v2/software_instances?columns[]=computer_details.ip_address
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Schema description

To retrieve the list of all columns that are returned by this association together with their descriptions, use the following request. The list contains custom computer properties that were added on the **Computer Properties** panel.

```
GET api/sam/schemas/associations/computer_details.json?token=<token>
```

Available columns

Table 215. Columns with information about computer details

Column	Description	Type
id	Identifier of the computer as specified in the BigFix Inventory database.	Integer
remote_id	Identifier of the computer as specified in the BigFix database.	Integer
last_seen	Date and time when the BigFix client connected to the BigFix server the last time.	String
datasource_id	Identifier of the BigFix server to which the computer reports.	Integer
datasource_name	Name of the BigFix server to which the computer reports.	String
created_at	Date and time when the information about the computer was imported from BigFix for the first time.	String
valid_to	The date and time when the computer stopped being discovered, for example because it was removed from the infrastructure, the data source was removed, or the BigFix client was uninstalled.	String
os	Operating system of the computer.	String
dns_name	Name of the computer that is specified in the Domain Name System.	String
name	Host name of the computer.	String
ip_address	IP address of the computer.	String
computer_property_number	Custom computer property that was added on the Computer Properties panel. To view the list of all custom properties, view the computer_details schema (on page mcxviii) . The name of the property as displayed on the	String

Table 215. Columns with information about computer details (continued)

Column	Description	Type
	user interface is shown in the <code>title</code> parameter.	

computer_health association

9.2.8

Available from 9.2.8. Use the `computer_health` association to retrieve information about health of the computer, such as the status of software scans, or version of the installed BigFix client, as part of REST API requests.

Association

```
computer_health
```

Applicable REST APIs

You can use the `computer_health` association with the following REST APIs:

- [api/sam/v2/computers](#) (on page *mxxx*)
- [api/sam/v2/software_instances](#) (on page *cmxcii*)

Syntax

```
<URL>?columns[]=computer_health.<column>&token=<token>
```

For example:

```
GET api/sam/v2/software_instances?columns[]=computer_health.catalog_version
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

9.2.10**Schema description**

To retrieve the list of all columns that are returned by this association together with their descriptions, use the following request.

```
GET api/sam/schemas/associations/computer_health.json?token=<token>
```

Available columns**Table 216. Columns with information about computer health**

Column	Description	Type
<code>agent_version</code>	Version of the BigFix client that is deployed on the computer.	Numeric

Table 216. Columns with information about computer health (continued)

Column	Description	Type
catalog_version	Version of the BigFix-provided software catalog that is available on the computer.	Numeric
catalog_custom_content_version	Version of the custom software catalog that is available on the computer.	Numeric
is_catalog_scan_successful	Indicates whether the catalog scan was successful.	Boolean
is_filesys_scan_successful	Indicates whether the file system scan was successful.	Boolean
is_idtag_scan_successful	Indicates whether the last software tags scan was successful.	Boolean
is_low_on_disk_space	Indicates whether the computer has less than 100 MB of free disk space.	Boolean
is_missing_prereqs	Indicates whether any scanner prerequisites are missing on the computer.	Boolean
is_out_of_date	Indicates whether the latest version of the BigFix client is installed on the computer. For some operating systems, no newer version of the client is available.	Boolean
is_out_of_sync	Indicates whether the time that is set on the computer is at least an hour from the time that is set on the BigFix Inventory server.	Boolean
is_package_scan_successful	Indicates whether the latest package scan was successful.	Boolean
last_scan_attempt	Date and time when the last scan was initiated.	String

computer_hardware association

9.2.8 Available from 9.2.8. Use the `computer_hardware` association to retrieve detailed hardware information as part of other REST API requests.

Association

```
computer_hardware
```

Applicable REST APIs

You can use the `computer_hardware` association with the following REST APIs:

- [api/sam/v2/computers](#) (on page mxxx)
- [api/sam/v2/software_instances](#) (on page cmxcii)

Syntax

```
<URL>?columns[]=computer_hardware.<column>&token=<token>
```

For example:

```
GET api/sam/v2/software_instances?columns[]=computer_hardware.computer_type
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

9.2.10 Schema description

To retrieve the list of all columns that are returned by this association together with their descriptions, use the following request.

```
GET api/sam/schemas/associations/computer_hardware.json?token=<token>
```

Available columns

Table 217. Columns with information about computer hardware

Name on the Hardware Inventory report	Description	Type
<code>computer_type</code>	Computer Type. Possible values are: <ul style="list-style-type: none"> • 0 - physical • 1- virtual • 2- computer running on a public cloud 	Numeric
<code>shared_pool_id</code>	Identifier of IBM POWER shared processor pool.	String

Table 217. Columns with information about computer hardware (continued)

Name on the Hardware Inventory report	Description	Type
system_model	Name of the hardware system or virtualization technology that is discovered by the BigFix client.	String
cluster_name	Name of the cluster to which the computer belongs.	String
cluster_cores_count	Number of processor cores available within the cluster to which the computer belongs.	Numeric
Partition Cores	Number of processor cores on the partition on which the BigFix client is installed. The number does not take into account hyper-threading on the x86 architecture.	Numeric
status	<p>Status of the computer. Possible values are:</p> <ul style="list-style-type: none"> • 1 - OK • 2 - No Scan Data • 3 - No Host Scan Data • 4 - No VM Manager Data • 9.2.10 5 - Outdated VM Manager Data 	String
server_id	Internal identifier of the physical server that hosts a computer. In case of a virtual machine that is not fully defined, the ID is a placeholder for a physical server and it changes when you complete defining the virtual machine.	Numeric
server_name	A unique system in the infrastructure. For a physical machine, it is the hardware manufacturer, type, and a machine serial number. For a virtual ma-	String

Table 217. Columns with information about computer hardware (continued)

Name on the Hardware Inventory report	Description	Type
	<p>chine, it is the manufacturer and host name. For a virtual machine with an incomplete definition, it is a UUID that is prefixed with <code>TLM_VM_</code>.</p> <p>You can specify the data that is returned by this property by configuring the following advanced server settings: managedServerTagTemplate and standaloneServerTagTemplate. For more information, see: Configuring advanced server settings (on page cdi).</p>	
9.2.16 <code>server_serial_number</code>	A serial number of a physical server. In case of a virtual machine that is not fully defined, the Server Serial Number is a UUID that is prefixed with <code>TLM_VM_</code> .	Numeric
9.2.16 <code>server_type</code>	A type of a physical server.	String
9.2.16 <code>server_vendor</code>	A type of a physical server.	String
9.2.16 <code>server_model</code>	A type of a physical server.	String
<code>node_total_processors</code>	Number of active sockets on the physical server for placing CPU units.	Numeric
<code>server_cores</code>	Number of active physical processor cores on the physical server.	Numeric
<code>server_id</code>	Identifier of the server.	
<code>pvu_per_core</code>	Number of PVUs that are assigned to a processor core on the computer. By default, the value is taken from the PVU table.	Numeric
<code>default_pvu_value</code>	Indicates whether a default number of processor value units is assigned to the computer.	Boolean

The following properties identify the Current Physical Server Processor. In case of a virtual machine that is not fully defined, the information is gathered on VM level.

Table 218. Columns with properties that identify the Current Physical Server Processor

Property	Name on the Hardware Inventory report	Description	Type
processor_brand_string	Processor Brand String	Full specification of the processor including its brand, model and speed as read from the computer operating system.	String
processor_brand	Brand	Brand of the processor that was matched in the PVU table. If the brand is "Other", the processor is not listed in the PVU table.	String
processor_model	Model	Model of the processor.	String
processor_type	Type	Type of the processor according to the number of cores.	String
processor_vendor	Vendor	Vendor of the processor.	String

cve association

9.2.13 Available from 9.2.13. Use the `cve` association to retrieve information about Common Vulnerabilities and Exposures (CVEs) that were matched with a particular component through its detailed version as part of other REST API requests.

Association

`cve`

Applicable REST APIs

You can use the `cve` association with the following REST API:

- [/api/sam/v2/software_instances](#) (on page [cmxcii](#))

Syntax

```
<URL>?columns[]=cve.<column>&token=<token>
```

Schema description

To retrieve the list of all columns that are returned by this association together with their descriptions, use the following request.

```
GET api/sam/schemas/associations/cve.json?token=<token>
```

Available columns

Table 219. Columns with information about CVEs

Property	Description	Type
name	A unique identifier of the vulnerability as assigned in the National Vulnerability Database.	String
base_severity	Severity of the CVE. Possible values: <ul style="list-style-type: none"> • 0 - None • 1 - Low • 2 - Medium • 3 - High • 4 - Critical 	Integer
base_score	Base score of the CVE that represents innate characteristics of the vulnerability. It maps to the severity of the CVE. For example, base score in range between 0.1 and 3.9 is an equivalent of low severity in CVSS v3.0. For more information, see: Vulnerability Metrics .	Float
vector_string	Vector string of the CVE in the CVSS system. It is a textual representation of metric values that are used to score the vulnerability.	String
exploitability_score	Exploitability sub-score of the CVE that represent the ease and technical means by which the vulnerability can be exploited.	Float
impact_score	Impact sub-score of the CVE that indicates how important will be the	Float

Table 219. Columns with information about CVEs (continued)

Property	Description	Type
	consequences if the vulnerability is exploited.	
source_version	Version of CVSS that is used as the source of information about the severity and metrics of the CVE. Possible values: <ul style="list-style-type: none"> • 2 - CVSS v2.0 • 3 - CVSS v3.0 	Integer
publish_date	The date and time when the CVE was published in the National Vulnerability Database.	String
modified_date	The date and time when the CVE was last modified in the National Vulnerability Database.	String
description	Description of the CVE details.	String

Sorting and filtering



Note: It is not recommended to filter information by more than two CVE properties at a time because of the high consumption of database resources.

- When you filter REST API responses by any CVE property, the filter shows all components that meet the specified criteria by searching through the full list of CVEs of each reported component.
- CVEs that are returned for a single component are sorted by the **base_score** value.
- REST API responses cannot be sorted by columns from the `cve` association because the association returns multiple CVEs for each component.

discovery_details association

9.2.11 Available from 9.2.11. Use the `discovery_details` association to retrieve details of a file that caused component discovery as part of a REST API request. The details include the name and size of the file as well as the path under which it was found.

Association

```
discovery_details
```

Applicable REST APIs

You can use the `discovery_details` association with the [api/sam/v2/software_instances](#) (on page [cmxcii](#)) REST API.

Syntax

```
<URL>?columns[]=discovery_details.<column>&token=<token>
```

For example:

```
GET api/sam/v2/software_instances?columns[]=discovery_details.file_path
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Available columns

Table 220. Columns with information about file that caused component discovery

Property	Description	Type
<code>discovery_path</code>	Path to the file that caused component discovery.	String
<code>file_name</code>	Name of the file that caused component discovery.	String
<code>file_size</code>	Size of the file that caused component discovery.	String

designated_computer association

9.2.12

Available from 9.2.12. Use the `designated_computer` association to retrieve information about computers that were designated to scan shared disks as part of other REST API requests.

Association

`designated_computer`

Applicable REST APIs

You can use the `designated_computer` association with the following REST APIs:

- [api/sam/v2/shared_disks](#) (on page [mxxxvii](#))

Syntax

```
<URL>?columns[]=designated_computer.<column>&token=<token>
```

For example:

```
GET api/sam/v2/shared_disks?columns[]=designated_computer.ip_address
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Schema description

To retrieve the list of all columns that are returned by this association together with their descriptions, use the following request.

```
GET api/sam/schemas/associations/designated_computer.json?token=<token>
```

Available columns

Table 221. Columns with information about computer hardware

Property	Description	Type
id	Identifier of the computer that was designated by BigFix Inventory to scan the shared disk as specified in the BigFix Inventory database.	Integer
remote_id	Identifier of the computer that was designated by BigFix Inventory to scan the shared disk as specified in the BigFix database.	Integer
last_seen	Date and time when the computer reported to BigFix for the last time. The time is specified in the GMT time zone.	String
datasource_id	Identifier of the BigFix server to which the computer reports.	Integer
datasource_name	Name of the BigFix server to which the computer reports.	String
created_at	Date and time when the information about the computer was imported from BigFix for the first time.	String
valid_to	The date and time when the computer stopped being discovered, for example because it was removed from the infrastructure, the data source was removed, or the BigFix client was uninstalled.	String

Table 221. Columns with information about computer hardware (continued)

Property	Description	Type
os	Operating system of the computer.	String
dns_name	Name of the computer that is specified in the Domain Name System.	String
name	Host name of the computer.	String
ip_address	IP address of the computer.	String

mounted_shared_disks association

9.2.12 Available from 9.2.12. Use the `mounted_shared_disks` association to retrieve information about shared disks that are mounted on each computer as part of other REST API requests.

Association

`mounted_shared_disks`

Applicable REST APIs

You can use the `mounted_shared_disks` association with the following REST APIs:

- [api/sam/v2/computers](#) (on page *mxxx*)

Syntax

```
<URL>?columns[]=mounted_shared_disks.<column>&token=<token>
```

For example:

```
GET api/sam/v2/computers?columns[]=mounted_shared_disks.mount_point
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Schema description

To retrieve the list of all columns that are returned by this association together with their descriptions, use the following request.

```
GET api/sam/schemas/associations/mounted_shared_disks.json?token=<token>
```

Available columns

Table 222. Columns with information about computer hardware

Property	Description	Type
<code>shared_disk_id</code>	Identifier of the shared disk that is mounted on the computer.	Integer
<code>mount_point</code>	Directory in which the shared disk is mounted.	String
<code>exported_directory</code>	Exported directory of the shared disk.	String

`release_component_lifecycle` association

9.2.13

Available from 9.2.13. Use the `release_component_lifecycle` association to retrieve information about when a component reaches its end of support as part of other REST API requests.

Association

`release_component_lifecycle`

Applicable REST APIs

You can use the `release_component_lifecycle` association with the following REST API:

- `/api/sam/v2/software_instances` (on page [cmxcii](#))

Syntax

```
<URL>?columns[]=release_component_lifecycle.<column>&token=<token>
```

Schema description

To retrieve the list of all columns that are returned by this association together with their descriptions, use the following request.

```
GET api/sam/schemas/associations/release_component_lifecycle.json?token=<token>
```

Available columns

Table 223. Columns with information about when a component reaches its end of support

Property	Description	Type
<code>source_type</code>	Indicates whether the end of support date is predefined or was set by a user. Possible values:	Integer

Table 223. Columns with information about when a component reaches its end of support (continued)

Property	Description	Type
	<ul style="list-style-type: none"> • 0 - Set by a user • 1 - Predefined 	
eos_date	Date when the component reaches its end of support.	String
updated_at	Date and time when the end of support date was last modified.	String

shared_disk_members association

9.2.12 Available from 9.2.12. Use the `shared_disk_members` association to retrieve information about computers on which shared disks are mounted as part of other REST API requests.

Association

`shared_disk_members`

Applicable REST APIs

You can use the `shared_disk_members` association with the following REST APIs:

- [api/sam/v2/shared_disks](#) (on page *mxxxvii*)

Syntax

```
<URL>?columns[]=shared_disk_members.<column>&token=<token>
```

For example:

```
GET api/sam/v2/shared_disks?columns[]=shared_disk_members.mount_point
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

Schema description

To retrieve the list of all columns that are returned by this association together with their descriptions, use the following request.

```
GET api/sam/schemas/associations/shared_disk_members.json?token=<token>
```

Available columns

Table 224. Columns with information about computer hardware

Property	Description	Type
computer_id	Identifier of the computer on which the shared disk is mounted.	Integer
mount_point	Directory in which the shared disk is mounted.	String

tags association

9.2.14 Available from 9.2.14. Use the `tags` association to retrieve information about tags as part of other REST API requests.

Association

`tags`

Applicable REST APIs

You can use the `tags` association with the following REST APIs:

- [/api/sam/v2/software_instances](#) (on page [cmxcii](#))

Syntax:

```
<URL>?columns[]=component_tags.<column>&token=<token>
```

```
<URL>?columns[]=instance_tags.<column>&token=<token>
```

- [/api/sam/v2/software_components](#) (on page [mv](#))

Syntax:

```
<URL>?columns[]=tags.<column>&token=<token>
```

Schema description

To retrieve the list of all columns that are returned by this association together with their descriptions, use the following request.

```
GET api/sam/schemas/associations/tag.json?token=<token>
```

Available columns

Table 225. Columns with information about components as specified in the CPE standard

Property	Description	Type
ID	Information about the unique identifier of a tag.	Integer
name	Name of a tag.	String

usage_data association

9.2.8 Available from 9.2.8. Use the `usage_data` association to retrieve information about application usage as part of other REST API requests.

Association

```
usage_data
```

Applicable REST APIs

You can use the `usage_data` association with the [api/sam/v2/software_instances](#) (*on page cmxcii*) REST API.

Syntax

```
<URL>?columns[]=usage_data.<column>&token=<token>
```

For example:

```
GET api/sam/v2/software_instances?columns[]=usage_data.has_usage
&token=7adc3efb175e2bc0f4484bdd2efca54a8fa04623
```

9.2.10 Schema description

To retrieve the list of all columns that are returned by this association together with their descriptions, use the following request.

```
GET api/sam/schemas/associations/usage_data.json?token=<token>
```

Available columns

Table 226. Columns with information about usage

Column	Description	Type
avg_run_time	Average amount of time during which a process or software was used. It	String

Table 226. Columns with information about usage (continued)

Column	Description	Type
	is calculated by dividing the total run time by the number of total runs.	
avg_runs_per_day	Average number of times a process or software item is started per day.	Numeric
first_used	Date and time when the information about the process was first collected by the BigFix client.	String
has_usage	Indicates whether a usage signature exists for the particular software.	Boolean
last_used	Date and time when the software item was last used or a process was last run.	String
process	Name of the process on which the calculation of software usage is based.	String
total_runs	Number of times a software item or process was started.	Numeric
total_time	Indicates how long the software item or process was running.	String

Integrating with IBM Control Desk

IBM Control Desk is an integrated service management solution that helps you manage a comprehensive range of IT processes, services, and assets. You can use one of its components, Integration Composer, to import and then process data from external applications, like BigFix Inventory.

Integration Composer

Integration Composer is an integration tool that imports hardware and software inventory data from external databases into the Maximo database that is used by IBM Control Desk. The tool can be used to retrieve your hardware inventory from BigFix and software inventory from BigFix Inventory. The retrieved data can be then passed to IBM Control Desk to create reports and license calculations.

Integration Composer uses integration adapters to facilitate data imports. Those adapters specify how to transform and import data for a specific discovery tool. Such integration adapters are therefore required to describe the data that is being imported and to provide instructions on how to transform it.

Integration adapters

Integration adapters consist of a data schema and a mapping file. They provide instructions to Integration Composer on how to transform and import your data. Your complete inventory, consisting of information about your hardware and software, is imported in two batches which requires two adapters to be used. The adapter for BigFix must be used first and is responsible for retrieving your hardware inventory directly from the BigFix platform. After that data is imported, you use the dedicated BigFix Inventory adapter to retrieve your software inventory from BigFix Inventory. The names of the data schemas as well as the mapping files listed in the following table are needed to define the data source connections and to create proper mappings that are prerequisites for starting the import.

Table 227. Integration adapters for BigFix Inventory

This table consists of five columns, one header row and two body rows.

BigFix Inventory version	Supported IBM Control Desk version	Required adapters		Retrieved data
		Data schema	Mapping file	
BigFix Inventory	IBM Control Desk 7.5.1.2, including Integration Composer 7.5.1.2	BigFix 9.0	IEM90ToD-PA75.fsn	Hardware inventory
		BigFix SUA 9.1	EMSUA91ToD-PA75.fsn	Software inventory, PVU and RVU data

Limitations

Importing your data to IBM Control Desk might result in some discrepancies between what you are importing and what is displayed in the application, mostly caused by different approaches to classifying software products between BigFix Inventory and IBM Control Desk. Also, because of the fact that your hardware and software inventory are imported separately, the right order of running the adapters is significant, not to mention cases in which you might want to import the data from both BigFix Inventory and other discovery tools, like Tivoli® Asset Discovery for Distributed.

Integration with BigFix Inventory and Tivoli Asset Discovery for Distributed

The BigFix Inventory adapter often imports more extensive data about software inventory than the Tivoli Asset Discovery for Distributed adapter. Because each software inventory import overwrites the data from another software inventory import, run the BigFix Inventory adapter as the last one to avoid loss of more extensive data. The following order is recommended when you run multiple adapters:

1. Tivoli Asset Discovery for Distributed - to import software data.
2. BigFix - to import hardware data.
3. BigFix Inventory - to import software data.

Importing your data

Data related to your hardware and software inventory must be imported separately. It means that first you must use the adapter for BigFix to import your hardware inventory directly from the platform, and then the adapter for BigFix Inventory to import your software inventory.

Hardware inventory

To import complete hardware inventory from BigFix, Integration Composer retrieves a set of computer properties that describe details of assets that are deployed in your environment, such as the number of processor cores, memory, host names, and so on. These properties are included in a number of analyses that collect this data from your endpoints and upload it to BigFix. All of these properties along with analyses and action sites in which they are contained are described in the mapping table for BigFix. Before you use the BigFix adapter to import you hardware inventory to IBM Control Desk, ensure that all of the action sites listed in the mapping table are enabled and that the corresponding analyses are activated. For more information, see [BigFix mapping table](#).

When all of the analyses are activated, you can start importing your hardware inventory to IBM Control Desk by using Integration Composer. To complete this task, you must define the data source connections so that both the source and the target database is recognized by Integration Composer, create a mapping that describes the way in which your data should be transformed and imported, and then run this mapping to start the import:

1. [Define your data source connections.](#)

This procedure must be completed twice to define the connection for a source and for a target. To define the source, choose the BigFix data schema that is called **BigFix 9.0**. The target is the Maximo database that is used by IBM Control Desk. This data schema is called **Deployed Assets 7.5**.

2. [Create a mapping.](#)

While creating a mapping, you import the mapping file that provides instructions to Integration Composer on how to transform and import your data. The mapping file for BigFix is `IEM90ToDPA75.fsn`.

3. [Run a mapping.](#)

Run the mapping to start importing your data. You can do it from a command line which is the quickest option, or by using a properties file or a prefilled script. In case of frequent imports, you might want to choose the prefilled script because you can save the required parameters and then reuse them with each import.

Software inventory

Your software inventory, together with PVU and RVU data, is imported directly from BigFix Inventory. No additional configuration is required because all the necessary computer properties are already available in BigFix Inventory. The procedure for importing your data is similar to the one for BigFix. In the same way, you define the data source connections and create a mapping, however you must use the dedicated data schema and mapping file:

1. [Define your data source connections.](#)

This procedure must be completed twice to define the connection for a source and for a target. To define the source, choose the BigFix Inventory data schema that is called **BigFix SUA 9.1**. The target is the Maximo database that is used by IBM Control Desk. This data schema is called **Deployed Assets 7.5**.

2. [Create a mapping.](#)

The mapping file for BigFix Inventory is `IEMSUA91ToDPA75.fsn`.

3. [Run a mapping.](#)

Run the mapping to start importing your data.

Results

After you run the mapping, your data is imported from BigFix Inventory. You can now log in to IBM Control Desk to view the imported data.

Related information

[Integration best practices](#)

Viewing your data

You can view the imported data in the IBM Control Desk interface. Check the data about computers in your environment and software applications that are installed on them. You can also view the imported PVU and RVU consumption.

Viewing computers and installed software

View the imported data about computers and software applications that are installed on them.

1. Log in to IBM Control Desk.
2. In the navigation bar, click **Assets > Deployed Assets > Computers**.
3. Click the **Refresh** icon to update the list from the database.

The list contains data for your computers that was imported from BigFix Inventory.

4. To view software installed on your computers, click on one of them.
5. Click the **Software** tab and then click **Applications**.

The list contains software applications that are installed on a particular computer.

Viewing PVU and RVU consumption

View the consumed PVU and RVU data for your software applications.

1. Log in to IBM Control Desk.
2. In the navigation bar, click **Assets > Licenses**.
3. In the **More Actions** section, click **Add/View Software Consumption Data**.

The list contains your software applications and the related PVU and RVU consumption.



Note: The VALUNIT capacity unit represents the PVU consumption.

Related information

[Managing licenses](#)

Glossary

This glossary provides terms and definitions for the BigFix Inventory software and products.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

[A \(on page mcxl\)](#) [B \(on page mcxl\)](#) [C \(on page mcxli\)](#) [D \(on page mcxlii\)](#) [E \(on page mcxliii\)](#) [F \(on page mcxliii\)](#) [H \(on page mcxliii\)](#) [I \(on page mcxliii\)](#) [K \(on page mcxliv\)](#) [L \(on page mcxliv\)](#) [M \(on page mcxlv\)](#) [O \(on page mcxlv\)](#) [P \(on page mcxlv\)](#) [R \(on page mcxlv\)](#) [S \(on page mcxlv\)](#) [T \(on page mcxlvii\)](#) [V \(on page mcxlvii\)](#)

A

activated processor core

A processor core that is managed or used by a product, regardless of whether the capacity of the processor core can be limited through virtualization technologies.

administration server

The database and web interface that software asset managers use for maintaining information about license entitlements and instances of installed products.

agent

A process that performs an action on behalf of a user or other program without user intervention or on a regular schedule, and reports the results back to the user or program.

API

See [application programming interface \(on page mcxl\)](#).

application programming interface (API)

An interface that allows an application program that is written in a high-level language to use specific data or functions of the operating system or another program.

audit

In ITIL, the act of confirming compliance with a standard or set of guidelines, or comparing actual measurements to targets, or verifying the accuracy of recorded information.

authentication

In computer security, a process that ensures that the identities of both the sender and the receiver of a network transaction are true.

B

bundle

To package a collection of individually orderable components or products into a single offering, often for promotional purposes. Software manufacturers typically offer a single license to cover all components of a bundled offering.

bundling

A process during which the user assigns a component to a product.

C**CA**

See [certificate authority \(on page mcxli\)](#).

certificate

In computer security, a digital document that binds a public key to the identity of the certificate owner, thereby enabling the certificate owner to be authenticated. A certificate is issued by a certificate authority and is digitally signed by that authority. See also [certificate authority \(on page mcxli\)](#), [certificate signing request \(on page mcxli\)](#).

certificate authority (CA)

A trusted third-party organization or company that issues the digital certificates. The certificate authority typically verifies the identity of the individuals who are granted the unique certificate. See also [certificate \(on page mcxli\)](#).

certificate signing request (CSR)

An electronic message that an organization sends to a certificate authority (CA) to obtain a certificate. The request includes a public key and is signed with a private key; the CA returns the certificate after signing with its own private key. See also [certificate \(on page mcxli\)](#), [keystore \(on page mcxliv\)](#).

client

A software program or computer that requests services from a server. See also [host \(on page mcxliv\)](#).

cluster

A set of independent systems or logical partitions (called nodes) that are organized into a network for the purpose of sharing resources and communicating with each other.

compliance

A state of being in accordance with established software and security specifications on target computers, or the process of becoming so.

component

A software item that is part of a software product, and might be separately identified, but is not individually licensed.

confirmed instance

An instance of an installed software product that has been explicitly assigned to one or more bundles.

constant special item ID list (CSIDL)

In Windows environments, a list that identifies frequently used special folders whose location might vary on different systems.

contract

A set of information about a software license for a product or products, its cost and entitlement period. When a contract is assigned to a computer group, it indicates which computers are entitled to the licenses described by that contract.

core

A single chip that houses a central processing unit (CPU) and is a component in the larger circuit design of a computer. A single chip can only contain one CPU, but a processor can contain multiple cores. See also [dual-core \(on page mcxlii\)](#), [multi-core \(on page mcxlv\)](#), [processor value unit \(on page mcxlv\)](#).

CSIDL

See [constant special item ID list \(on page mcxlii\)](#).

CSR

See [certificate signing request \(on page mcxli\)](#).

cURL

A command line tool for retrieving and posting files using URL syntax.

D

data source

The source of data itself, such as a database or XML file, and the connection information necessary for accessing the data.

discover

To identify resources within a network environment.

DNS

See [domain name server \(on page mcxlii\)](#).

domain name server (DNS)

A server program that supplies name-to-address conversion by mapping domain names to IP addresses.

dual-core

Referring to a system that integrates two processors into one virtual processor. See also [core \(on page mcxlii\)](#), [multi-core \(on page mcxlv\)](#).

E

endpoint

A server, computer, machine or device that is monitored.

entitlement

In software licensing, the maximum allowed allocation of capacity as determined by a license agreement.

F

fixlet

A message that provides instructions to an agent to perform a management or reporting action.

full-capacity

Pertains to a software licensing scheme that bases charges on the capacity of the entire machine or cluster of machines that is available to the licensed program, rather than on just one or more partitions. See also [processor value unit \(on page mcxlv\)](#), [subcapacity \(on page mcxlvii\)](#).

H

host

A computer that is connected to a network and that provides an access point to that network. The host can be a client, a server, or both a client and server simultaneously. See also [client \(on page mcxli\)](#).

host name

In Internet communication, the name given to a computer. The host name might be a fully qualified domain name such as mycomputer.city.company.com, or it might be a specific subname such as mycomputer.

HTTP

See [Hypertext Transfer Protocol \(on page mcxliii\)](#).

Hypertext Transfer Protocol (HTTP)

An Internet protocol that is used to transfer and display hypertext and XML documents on the web.

hypervisor

Software or a physical device that enables multiple instances of operating systems to run simultaneously on the same hardware.

I

instance

A software product or component that is installed on a server or logical partition (LPAR).

IT infrastructure

All of the hardware, software, networks, and other facilities (but not people or processes) required to develop, test, deliver, monitor, control, or support IT services.

K**keystore**

In security, a file or a hardware cryptographic card where identities and private keys are stored, for authentication and encryption purposes. Some keystores also contain trusted or public keys. See also [certificate signing request \(on page mcxli\)](#).

knowledge base

See [software knowledge base \(on page mcxlvii\)](#).

L**LDAP**

See [Lightweight Directory Access Protocol \(on page mcxliv\)](#).

license

A legal agreement that authorizes the use of proprietary information including, but not limited to, copyrighted or patented information.

Lightweight Directory Access Protocol (LDAP)

An open protocol that uses TCP/IP to provide access to directories that support an X.500 model and that does not incur the resource requirements of the more complex X.500 Directory Access Protocol (DAP). For example, LDAP can be used to locate people, organizations, and other resources in an Internet or intranet directory.

locale

A setting that identifies language or geography and determines formatting conventions such as collation, case conversion, character classification, the language of messages, date and time representation, and numeric representation.

logical partition (LP, LPAR)

One or more virtualized images of a hardware computing system that can include shared and dedicated resources assigned from the pool of resources available on a physical server. Each image appears to the operating system running within it to be a unique instance of a physical server. See also [virtualization \(on page mcxlvii\)](#).

LP

See [logical partition \(on page mcxliv\)](#).

LPAR

See [logical partition \(on page mcxliv\)](#).

M

multi-core

Referring to a system that integrates multiple processors into one virtual processor. See also [core \(on page mcxlii\)](#), [dual-core \(on page mcxlii\)](#).

O

offering

The element or integrated set of elements (hardware, software, services) designed to satisfy the wants and needs of current and/or prospective customers. A solution is the application of the offering in a specific customer environment.

P

partition

A logical division of storage on a fixed disk. See also [subcapacity \(on page mcxlvii\)](#), [virtualization \(on page mcxlvii\)](#).

PID

See [product identifier \(on page mcxlv\)](#).

pricing metric

A measurement that defines how capacity is counted and applied against the customer's entitlement. See also [entitlement \(on page mcxliii\)](#).

processor value unit (PVU)

A unit of measure that is assigned to each processor core for software licensing purposes. PVUs vary according to chip architecture. See also [core \(on page mcxlii\)](#), [full-capacity \(on page mcxliii\)](#), [subcapacity \(on page mcxlvii\)](#).

product ID

See [product identifier \(on page mcxlv\)](#).

product identifier (PID, product ID)

A unique value that identifies an BigFix software product. Every mainframe and distributed BigFix software product has a PID.

PVU

See [processor value unit \(on page mcxlv\)](#).

R

raw data

Unprocessed data that contains information about software items and manufacturers, obtained through software scans.

Representational State Transfer (REST)

A software architectural style for distributed hypermedia systems like the World Wide Web. The term is also often used to describe any simple interface that uses XML (or YAML, JSON, plain text) over HTTP without an additional messaging layer such as SOAP.

resource value unit (RVU)

A unit of measure by which a program can be licensed that is based on the number of units of a specific resource used or managed by the program.

response file

A file that can be customized with the setup and configuration data that automates an installation. During an interactive installation, the setup and configuration data must be entered, but with a response file, the installation can proceed without any intervention.

REST

See [Representational State Transfer \(on page mcxlvii\)](#).

RVU

See [resource value unit \(on page mcxlvii\)](#).

S

scan

To systematically search a computer or a network for information about hardware, software, or configuration. See also [software scan \(on page mcxlviii\)](#).

signature

The set of unique information that identifies a software application, such as the name, version, and file size of an application.

silent mode

A method for installing or uninstalling a product component from the command line with no GUI display. When using silent mode, you specify the data required by the installation or uninstallation program directly on the command line or in a file (called an option file or response file).

site

A subdivision of an organization that can track inventory and other data separately from other sites.

software catalog

A portable representation of the contents of the software knowledge base that serves IT management applications including asset management, license management, and software provisioning. Catalog format and content might vary, depending on the target application. See also [software knowledge base \(on page mcxlvii\)](#).

software hierarchy

The combination of software product, version, release, and variation that represents a software item in the software knowledge base. The product is the root of the hierarchy. See also [software knowledge base \(on page mcxlvii\)](#).

software knowledge base

A collection of information about distributed software products and components, dependencies between them, the means to discover them, and their basic licensing properties. The knowledge base is used to generate software catalogs that asset management tools can use for software inventory identification, license compliance, and software provisioning. See also [software catalog \(on page mcxlvii\)](#), [software hierarchy \(on page mcxlvii\)](#).

software product

A software item that is licensed independently of other software items. For licensing purposes, software products are sometimes bundled into a single sales offering.

software scan

An automated process that discovers instances of software installed on the computers in a network. The software scan results include details about discovered software, such as version numbers. See also [scan \(on page mcxlvii\)](#).

subcapacity

Pertains to a software licensing scheme that bases charges on the capacity of the partition where the licensed program is used, rather than on the total capacity in the server. See also [full-capacity \(on page mcxlvii\)](#), [partition \(on page mcxlvii\)](#), [processor value unit \(on page mcxlvii\)](#), [virtualization \(on page mcxlvii\)](#), [virtualization \(on page mcxlvii\)](#).

T**TLS**

See [Transport Layer Security \(on page mcxlvii\)](#).

Transport Layer Security (TLS)

A set of encryption rules that uses verified certificates and encryption keys to secure communications over the Internet. TLS is an update to the SSL protocol.

V**virtualization**

The creation of a virtual computing resource such as an operating system, server, storage device, or network resources in a one-to-many or many-to-one association. For example, a single operating system might be divided into multiple partitions that operate as independent systems, or multiple disk devices might appear as a single logical drive. See also [logical partition \(on page mcxliv\)](#), [partition \(on page mcxlv\)](#), [subcapacity \(on page mcxlvii\)](#).

virtual machine (VM)

An instance of a data-processing system that appears to be at the exclusive disposal of a single user, but whose functions are accomplished by sharing the resources of a physical data-processing system.

VM

See [virtual machine \(on page mcxlviii\)](#).

Notices

This information was developed for products and services offered in the US.

HCL may not offer the products, services, or features discussed in this document in other countries. Consult your local HCL representative for information on the products and services currently available in your area. Any reference to an HCL product, program, or service is not intended to state or imply that only that HCL product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any HCL intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-HCL product, program, or service.

HCL may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

HCL

330 Potrero Ave.

Sunnyvale, CA 94085

USA

Attention: Office of the General Counsel

For license inquiries regarding double-byte character set (DBCS) information, contact the HCL Intellectual Property Department in your country or send inquiries, in writing, to:

HCL

330 Potrero Ave.

Sunnyvale, CA 94085

USA

Attention: Office of the General Counsel

HCL TECHNOLOGIES LTD. PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. HCL may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-HCL websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this HCL product and use of those websites is at your own risk.

HCL may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

mcl

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

HCL

330 Potrero Ave.

Sunnyvale, CA 94085

USA

Attention: Office of the General Counsel

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by HCL under terms of the HCL Customer Agreement, HCL International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-HCL products was obtained from the suppliers of those products, their published announcements or other publicly available sources. HCL has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-HCL products. Questions on the capabilities of non-HCL products should be addressed to the suppliers of those products.

Statements regarding HCL's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to HCL, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. HCL, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS," without warranty of any kind. HCL shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from HCL Ltd. Sample Programs.

Trademarks

HCL Technologies Ltd. and HCL Technologies Ltd. logo, and hcl.com are trademarks or registered trademarks of HCL Technologies Ltd., registered in many jurisdictions worldwide.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of HCL or other companies.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the HCL website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of HCL.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of HCL.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

HCL reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by HCL, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

HCL MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.