

**HCL OneDB 2.0.1**

**OneDB Guide to SQL: Reference**



# Contents

<b>Chapter 1. Guide to SQL: Reference.....</b>	<b>4</b>		
System catalog tables.....	4		
Objects that the system catalog tables track.....	4		
Using the system catalog.....	5		
Structure of the System Catalog.....	11		
SYSAGGREGATES.....	14		
SYSAMS.....	15		
SYSATTRTYPES.....	18		
SYSAUTOLOCATE.....	19		
SYSBLOBS.....	20		
SYSCASTS.....	20		
SYSCHECKS.....	21		
SYSCHECKUDRDEP.....	21		
SYSCOLATTRIBS.....	22		
SYSCOLAUTH.....	23		
SYSCOLDEPEND.....	23		
SYSCOLUMNS.....	24		
SYSCONSTRAINTS.....	29		
SYSDEFAULTS.....	30		
SYSDEPEND.....	31		
SYSDIRECTIVES.....	31		
SYSDISTRIB.....	32		
SYSDOMAINS.....	34		
SYSERRORS.....	34		
SYSEXTCOLS.....	35		
SYSEXTDFILES.....	35		
SYSEXTERNAL.....	36		
SYSFRAGAUTH.....	36		
SYSFRAGDIST.....	37		
SYSFRAGMENTS.....	39		
SYSINDEXES.....	41		
SYSINDICES.....	43		
SYSINHERITS.....	46		
SYSLANGAUTH.....	46		
SYSLOGMAP.....	47		
SYSOBJSTATE.....	47		
SYSOPCLASSES.....	48		
SYSOPCLSTR.....	49		
SYSPROCAUTH.....	51		
SYSPROCBODY.....	51		
SYSPROCCOLUMNS.....	52		
SYSPROCEDURES.....	53		
SYSPROCPPLAN.....	56		
SYSREFERENCES.....	56		
SYSROLEAUTH.....	57		
SYSROUTINELANGS.....	58		
SYSSECLABELAUTH.....	58		
SYSSECLABELCOMPONENTS.....	58		
SYSSECLABELCOMPONENTELEMENTS.....	59		
SYSSECLABELNAMES.....	59		
SYSSECLABELS.....	60		
SYSSECPOLICIES.....	60		
SYSSECPOLICYCOMPONENTS.....	61		
SYSSECPOLICYEXEMPTIONS.....	61		
SYSSEQUENCES.....	62		
SYSSURROGATEAUTH.....	62		
SYSSYNONYMS.....	63		
SYSSYNTABLE.....	63		
SYSTABAMDATA.....	64		
SYSTABAUTH.....	65		
SYSTABLES.....	65		
SYSTRACECLASSES.....	69		
SYSTRACEMSGS.....	69		
SYSTRIGBODY.....	70		
SYSTRIGGERS.....	71		
SYSUSERS.....	72		
SYSVIEWS.....	72		
SYSVIOLATIONS.....	73		
SYSXADATASOURCES.....	73		
SYSXASOURCETYPES.....	74		
SYSXTDDESC.....	74		
SYSXTDTYPEAUTH.....	75		
SYSXTDTYPES.....	75		
Information Schema.....	77		
Data types.....	81		
Summary of data types.....	81		
Description of Data Types.....	87		
Built-In Data Types.....	118		
Extended Data Types.....	128		
Data Type Casting and Conversion.....	132		
Operator Precedence.....	137		
Environment variables.....	138		
Types of environment variables.....	138		
Limitations on environment variables.....	139		
Using environment variables on UNIX™.....	139		
Using environment variables on Windows™.....	144		
Environment variables in HCL OneDB™ products.....	147		
Appendixes.....	220		
The stores_demo Database.....	220		
The superstores_demo database.....	222		

**Index..... 226**

# Chapter 1. Guide to SQL: Reference

The *HCL OneDB™ Guide to SQL: Reference* contains the reference information for the system catalog tables, data types, and environment variables of the HCL OneDB™ dialect of the SQL language, as implemented in HCL OneDB™. These topics also include information about the `stores_demo`, `sales_demo`, and `superstore_demo` databases that are included with HCL OneDB™.

This information is intended for the following users:

- Database users
- Database administrators
- Database security administrators
- Database application programmers.

This information assumes that you have the following background:

- A working knowledge of your computer, your operating system, and the utilities that your operating system provides.
- Some experience working with relational databases or exposure to database concepts.
- Some experience with computer programming.

## System catalog tables

The *system catalog* consists of tables and views that describe the structure of the database. Sometimes called the *data dictionary*, these table objects contain everything that the database knows about itself. Each system catalog table contains information about specific elements in the database. Each database has its own system catalog.

These topics provide information about the structure, content, and use of the system catalog tables. It also contains information about the Information Schema, which provides information about the tables, views, and columns in all the databases of the HCL OneDB™ instance to which your user session is currently connected.

## Objects That the System Catalog Tables Track

The system catalog tables maintain information about the database, including the following categories of database objects:

- Tables, views, synonyms, and table fragments
- Columns, constraints, indexes, and index fragments
- Distribution statistics for tables, indexes, and fragments
- Triggers on tables, and INSTEAD OF triggers on views
- Procedures, functions, routines, and associated messages
- Authorized users, roles, and privileges to access database objects
- LBAC security policies, components, labels, and exemptions
- Data types and casts
- User-defined aggregate functions
- Access methods and operator classes

- Sequence objects
- Storage spaces for BLOB and CLOB objects
- External optimizer directives
- Inheritance relationships
- XA data sources and XA data source types
- Trusted user and surrogate user information

## Using the system catalog

HCL OneDB™ automatically generate the system catalog tables when you create a database. You can query the system catalog tables as you would query any other table in the database. The system catalog tables for a newly created database are located in a common area of the disk called a *dbspace*. Every database has its own system catalog tables. All tables and views in the system catalog have the prefix **sys** (for example, the **sys** system catalog table).

Not all tables with the prefix **sys** are true system catalog tables. For example, the **syscdr** database supports the Enterprise Replication feature. Non-catalog tables, however, have a **tabid**  $\geq$  100. System catalog tables all have a **tabid**  $<$  100. See later in this section and [SYSTABLES on page 65](#) for more information about **tabid** numbers that the database server assigns to tables, views, synonyms, and (in HCL OneDB™) sequence objects.



**Tip:** Do not confuse the system catalog tables of a database with the tables in the **sysmaster**, **sysutils**, **syscdr**, or (for HCL OneDB™) the **sysadmin** and **sysuser** databases. The names of tables in those databases also have the **sys** prefix, but they contain information about an entire database server, which might manage multiple databases. Information in the **sysadmin**, **sysmaster**, **sysutils**, **syscdr**, and **sysuser** tables is primarily useful for database server administrators (DBSAs). See also the *HCL OneDB™ Administrator's Guide* and *HCL OneDB™ Administrator's Reference*.

The database server accesses the system catalog constantly. Each time an SQL statement is processed, the database server accesses the system catalog to determine system privileges, add or verify table or column names, and so on.

For example, the following CREATE SCHEMA block adds the **customer** table, with its indexes and privileges, to the **stores\_demo** database. This block also adds a view, **california**, which restricts the data of the **customer** table to only the first and last names of the customer, the company name, and the telephone number for all customers who reside in California.

```
CREATE SCHEMA AUTHORIZATION maryl
CREATE TABLE customer (customer_num SERIAL(101), fname CHAR(15),
    lname CHAR(15), company CHAR(20), address1 CHAR(20), address2 CHAR(20),
    city CHAR(15), state CHAR(2), zipcode CHAR(5), phone CHAR(18))
GRANT ALTER, ALL ON customer TO cathl WITH GRANT OPTION AS maryl
GRANT SELECT ON customer TO public
GRANT UPDATE (fname, lname, phone) ON customer TO nhowe
CREATE VIEW california AS
    SELECT fname, lname, company, phone FROM customer WHERE state = 'CA'
CREATE UNIQUE INDEX c_num_ix ON customer (customer_num)
CREATE INDEX state_ix ON customer (state)
```

To process this CREATE SCHEMA block, the database server first accesses the system catalog to verify the following information:

- The new table and view names do not already exist in the database. (If the database is ANSI-compliant, the database server verifies that the new names do not already exist for the specified owners.)
- The user has permission to create tables and grant user privileges.
- The column names in the CREATE VIEW and CREATE INDEX statements exist in the **customer** table.

In addition to verifying this information and creating two new tables, the database server adds new rows to the following system catalog tables:

- **systables**
- **syscolumns**
- **sysviews**
- **systabauth**
- **syscolauth**
- **sysindexes**
- **sysindices**

### Rows added to the systables system catalog table

The following two new rows of information are added to the **systables** system catalog table after the CREATE SCHEMA block is run.

Column name	First row	Second row
<b>tablename</b>	customer	california
<b>owner</b>	maryl	maryl
<b>partnum</b>	16778361	0
<b>tabid</b>	101	102
<b>rowsize</b>	134	134
<b>ncols</b>	10	4
<b>nindexes</b>	2	0
<b>nrows</b>	0	0
<b>created</b>	01/26/2007	01/26/2007
<b>version</b>	1	0
<b>tabtype</b>	T	V
<b>locklevel</b>	P	B
<b>npused</b>	0	0
<b>fextsize</b>	16	0
<b>nextsize</b>	16	0

Column name	First row	Second row
flags	0	0
site		
dbname		

Each table recorded in the **sytables** system catalog table is assigned a **tabid**, a system-assigned sequential number that uniquely identifies each table in the database. The system catalog tables receive 2-digit **tabid** numbers, and the user-created tables receive sequential **tabid** numbers that begin with 100.

### Rows added to the syscolumns system catalog table

The CREATE SCHEMA block adds 14 rows to the **syscolumns** system catalog table. These rows correspond to the columns in the table **customer** and the view **california**, as the following example shows.

colname	tabid	colno	coltype	collength	colmin	colmax
customer_num	101	1	262	4		
fname	101	2	0	15		
lname	101	3	0	15		
company	101	4	0	20		
address1	101	5	0	20		
address2	101	6	0	20		
city	101	7	0	15		
state	101	8	0	2		
zipcode	101	9	0	5		
phone	101	10	0	18		
fname	102	1	0	15		
lname	102	2	0	15		
company	102	3	0	20		
phone	102	4	0	18		

In the **syscolumns** table, each column within a table is assigned a sequential column number, **colno**, that uniquely identifies the column within its table. In the **colno** column, the **fname** column of the **customer** table is assigned the value 2 and the **fname** column of the view **california** is assigned the value 1.

The **colmin** and **colmax** columns are empty. These columns contain values when a column is the first key (or the only key) in an index, has no NULL or duplicate values, and the UPDATE STATISTICS statement has been run.

## Rows added to the sysviews system catalog table

The database server also adds rows to the **sysviews** system catalog table, whose **viewtext** column contains each line of the CREATE VIEW statement that defines the view. In that column, the **x0** that precedes the column names in the statement (for example, **x0.fname**) operates as an alias that distinguishes among the same columns that are used in a self-join.

## Rows added to the systabauth system catalog table

The CREATE SCHEMA block also adds rows to the **systabauth** system catalog table. These rows correspond to the user privileges granted on **customer** and **california** tables, as the following example shows.

grantor	grantee	tabid	tabauth
maryl	public	101	su-idx--
maryl	cathl	101	SU-IDXAR
maryl	nhowe	101	--*-----
	maryl	102	SU-ID---

The **tabauth** column specifies the table-level privileges granted to users on the **customer** and **california** tables. This column uses an 8-byte pattern, such as **s** (Select), **u** (Update), **\*** (column-level privilege), **i** (Insert), **d** (Delete), **x** (Index), **a** (Alter), and **r** (References), to identify the type of privilege. In this example, the user **nhowe** has column-level privileges on the **customer** table. A hyphen (-) means the user has not been granted the privilege whose position the hyphen occupies within the **tabauth** value.

If the **tabauth** privilege code is in uppercase (for example, **s** for Select), the user has this privilege and can also grant it to others; but if the privilege code is lowercase (for example, **s** for Select), the user cannot grant it to others.

## Rows added to the syscolauth system catalog table

In addition, three rows are added to the **syscolauth** system catalog table. These rows correspond to the user privileges that are granted on specific columns in the **customer** table as the following example shows.

grantor	grantee	tabid	colno	colauth
maryl	nhowe	101	2	-u-
maryl	nhowe	101	3	-u-
maryl	nhowe	101	10	-u-

The **colauth** column specifies the column-level privileges that are granted on the **customer** table. This column uses a 3-byte pattern such as **s** (Select), **u** (Update), and **r** (References), to identify the type of privilege. For example, the user **nhowe** has Update privileges on the second column (because the **colno** value is 2) of the **customer** table (indicated by **tabid** value of 101).



## Rows added to the sysindexes or the sysindices table

The CREATE SCHEMA block adds two rows to the **sysindexes** system catalog table (the **sysindices** table for HCL OneDB™). These rows correspond to the indexes created on the **customer** table, as the following example shows.

idxname	c_num_ix	state_ix
owner	maryl	maryl
tabid	101	101
idxtype	U	D
clustered		
part1	1	8
part2	0	0
part3	0	0
part4	0	0
part5	0	0
part6	0	0
part7	0	0
part8	0	0
part9	0	0
part10	0	0
part11	0	0
part12	0	0
part13	0	0
part14	0	0
part15	0	0
part16	0	0
levels		
leaves		
nunique		
clust		
idxflags		

In this table, the **idxtype** column identifies whether the created index requires unique values (U) or accepts duplicate values (D). For example, the **c\_num\_ix** index on the **customer.customer\_num** column is unique.

## Accessing the system catalog

Normal user access to the system catalog is read-only. Users with Connect or Resource privileges cannot alter the catalog, but they can access data in the system catalog tables on a read-only basis using standard SELECT statements.

For example, the following SELECT statement displays all the table names and corresponding **tabid** codes of user-created tables in the database:

```
SELECT tabname, tabid FROM systables WHERE tabid > 99
```

When you use DB-Access, only the tables that you created are displayed. To display the system catalog tables, enter the following statement:

```
SELECT tabname, tabid FROM systables WHERE tabid < 100
```

You can use the **SUBSTR** or the **SUBSTRING** function to select only part of a source string. To display the list of tables in columns, enter the following statement:

```
SELECT SUBSTR(tabname, 1, 18), tabid FROM systables
```

Although user **informix** can modify most system catalog tables, you should not update, delete, or insert any rows in them. Modifying the content of system catalog tables can affect the integrity of the database. However, you can safely use the ALTER TABLE statement to modify the size of the next extent of system catalog tables. Changing the next extent size does not affect extents that already exist.

For certain catalog tables of HCL OneDB™, however, it is valid to add entries to the system catalog tables. For instance, in the case of the **syserrors** system catalog table and the **sysracemsgs** system catalog table, a DataBlade® module developer can directly insert entries that are in these system catalog tables.

## Update system catalog data

If you use the UPDATE STATISTICS statement to update the system catalog before executing a query or other data manipulation language (DML) statement, you can ensure that the information available to the query execution optimizer is current.

In HCL OneDB™, the optimizer determines the most efficient strategy for executing SQL queries and other DML operations. The optimizer allows you to query the database without requiring you to consider fully which tables to search first in a join or which indexes to use. The optimizer uses information from the system catalog to determine the best query strategy.

When you delete or modify a table, the database server does not automatically update the related statistical data in the system catalog. For example, if you delete one or more rows in a table with the DELETE statement, the **nrows** column in the **systables** system catalog table, which holds the number of rows for that table, is not updated automatically.

The UPDATE STATISTICS statement causes the database server to recalculate data in the **systables**, **sysdistrib**, **syscolumns**, and **sysindices** system catalog tables, and in the **sysindexes** view. (For operations on fragmented tables where the STATLEVEL attribute is set to FRAGMENT, it also updates the **sysfragdist** and **sysfragments** system catalog tables.) After

you run UPDATE STATISTICS, the **systables** system catalog table holds the correct value in the **nrows** column. If you specify MEDIUM or HIGH mode when you run UPDATE STATISTICS, the **sysdistrib** and (for fragment-level statistics) the **sysfragdist** system catalog tables hold the updated column-distribution data.

Whenever you modify a data table extensively, use the UPDATE STATISTICS statement to update data in the system catalog. For more information about the UPDATE STATISTICS statement, see the *HCL OneDB™ Guide to SQL: Syntax*.

## Structure of the System Catalog

The following system catalog tables describe the database objects in a database.

---

### System Catalog Tables

---

[SYSAGGREGATES on page 14](#)

---

[SYSAMS on page 15](#)

---

[SYSATTRTYPES on page 18](#)

---

[SYSAUTOLOCATE on page 19](#)

---

[SYSBLOBS on page 20](#)

---

[SYSCASTS on page 20](#)

---

[SYSCHECKS on page 21](#)

---

[SYSCHECKUDRDEP on page 21](#)

---

[SYSCOLATTRIBS on page 22](#)

---

[SYSCOLAUTH on page 23](#)

---

[SYSCOLDEPEND on page 23](#)

---

[SYSCOLUMNS on page 24](#)

---

[SYSCONSTRAINTS on page 29](#)

---

[SYSDEFAULTS on page 30](#)

---

[SYSDEPEND on page 31](#)

---

[SYSDIRECTIVES on page 31](#)

---

[SYSDISTRIB on page 32](#)

---

[SYSDOMAINS on page 34](#)

---

[SYSERRORS on page 34](#)

---

[SYSEXTCOLS on page 35](#)

---

[SYSEXTDFILES on page 35](#)

---

[SYSEXTERNAL on page 36](#)

---

---

**System Catalog Tables**

---

[SYSECLABELAUTH on page 58](#)

---

[SYSECLABELCOMPONENTS on page 58](#)

---

[SYSECLABELCOMPONENTELEMENTS on page 59](#)

---

[SYSECLABELNAMES on page 59](#)

---

[SYSECLABELS on page 60](#)

---

[SYSECPOLICIES on page 60](#)

---

[SYSECPOLICYCOMPONENTS on page 61](#)

---

[SYSECPOLICYEXEMPTIONS on page 61](#)

---

[SYSEQUENCES on page 62](#)

---

[SYSECLABELAUTH on page 58](#)

---

[SYSECLABELCOMPONENTS on page 58](#)

---

[SYSECLABELCOMPONENTELEMENTS on page 59](#)

---

[SYSECLABELNAMES on page 59](#)

---

[SYSECLABELS on page 60](#)

---

[SYSECPOLICIES on page 60](#)

---

[SYSECPOLICYCOMPONENTS on page 61](#)

---

[SYSECPOLICYEXEMPTIONS on page 61](#)

---

[SYSECLABELAUTH on page 58](#)

---

[SYSECLABELCOMPONENTS on page 58](#)

---

[SYSECLABELCOMPONENTELEMENTS on page 59](#)

---

[SYSECLABELNAMES on page 59](#)

---

[SYSECLABELS on page 60](#)

---

[SYSECPOLICIES on page 60](#)

---

[SYSECPOLICYCOMPONENTS on page 61](#)

---

[SYSECPOLICYEXEMPTIONS on page 61](#)

---

[SYSECLABELAUTH on page 58](#)

---

[SYSECLABELCOMPONENTS on page 58](#)

---

[SYSECLABELCOMPONENTELEMENTS on page 59](#)

---

[SYSECLABELNAMES on page 59](#)

---

[SYSECLABELS on page 60](#)

---

[SYSECPOLICIES on page 60](#)

---

[SYSECPOLICYCOMPONENTS on page 61](#)

---

[SYSECPOLICYEXEMPTIONS on page 61](#)

---

---

**System Catalog Tables**


---

[SYSSURROGATEAUTH](#) on page 62

---

[SYSSYNONYMS](#) on page 63

---

[SYSSYNTABLE](#) on page 63

---

[SYSTABAMDATA](#) on page 64

---

[SYSTABAUTH](#) on page 65

---

[SYSTABLES](#) on page 65

---

[SYSTRACECLASSES](#) on page 69

---

[SYSTRACEMSGS](#) on page 69

---

[SYSTRIGBODY](#) on page 70

---

[SYSTRIGGERS](#) on page 71

---

[SYSUSERS](#) on page 72

---

[SYSVIEWS](#) on page 72

---

[SYSVIOLATIONS](#) on page 73

---

[SYSXADATASOURCES](#) on page 73

---

[SYSXASOURCETYPES](#) on page 74

---

[SYSXTDDESC](#) on page 74

---

[SYSXTDTYPEAUTH](#) on page 75

---

[SYSXTDTYPES](#) on page 75

---

In case-sensitive databases that use the default database locale (U. S. English, ISO **8859-1** code set), character columns in these tables are CHAR and VARCHAR data types. For all other locales, character columns are the NLS data types, NCHAR and NVARCHAR. For information about differences in the collation order of character data types, see the *HCL OneDB™ GLS User's Guide*. See also the [Data types on page 81](#) chapter of this publication.

**Character columns in databases that are not case-sensitive**

In databases that are created with the NLSCASE INSENSITIVE keywords and that use the default database locale (U. S. English, ISO **8859-1** code set), character columns in system catalog tables are CHAR and VARCHAR data types, which support case-sensitive queries. For all other database locales, character column data types in the system catalog tables are the NLS data types, NCHAR and NVARCHAR, but with the following specific exceptions:

<i>Table_name.Column_name</i>	<b>Data type</b>
<b>sysams.am_sptype</b>	CHAR(3)

<i>Table_name.Column_name</i>	Data type
<b>syscolauth.colauth</b>	CHAR(3)
<b>sysdefaults.class</b>	CHAR(1)
<b>sysfragauth.fragauth</b>	CHAR(6)
<b>sysinherits.class</b>	CHAR(1)
<b>syslangauth.langauth</b>	CHAR(1)
<b>sysprocauth.procauth</b>	CHAR(1)
<b>sysprocedures.mode</b>	CHAR(1)
<b>systabauth.tabauth</b>	CHAR(9)
<b>systriggers.event</b>	CHAR(1)
<b>sysxdtypeauth.auth</b>	CHAR(2)

In each of these columns, case-sensitive encoding can record information that utilities of the database server require in queries on those system catalog tables. In a database that is case-insensitive, queries might return incorrect results from data stored in NCHAR or NVARCHAR columns, if different attributes of database objects are encoded as different cases of the same letter. To avoid the loss of information, CHAR data types are used for the system catalog columns listed above.

## SYSAGGREGATES

The **sysaggregates** system catalog table records user-defined aggregates (UDAs). The **sysaggregates** table has the following columns.

**Table 1. SYSAGGREGATES table column descriptions**

Column	Type	Explanation
<b>name</b>	VARCHAR(128)	Name of the aggregate
<b>owner</b>	CHAR(32)	Name of the owner of the aggregate
<b>aggid</b>	SERIAL	Unique code identifying the aggregate
<b>init_func</b>	VARCHAR(128)	Name of initialization UDR
<b>iter_func</b>	VARCHAR(128)	Name of iterator UDR

**Table 1. SYSAGGREGATES table column descriptions**

(continued)

Column	Type	Explanation
<b>combine_func</b>	VARCHAR(128)	Name of combine UDR
<b>final_func</b>	VARCHAR(128)	Name of finalization UDR
<b>handlesnulls</b>	BOOLEAN	NULL-handling indicator: <ul style="list-style-type: none"> <li>• t = handles NULLs</li> <li>• f = does not handle NULLs</li> </ul>

Each user-defined aggregate has one entry in **sysaggregates** that is uniquely identified by its identifying code (the **aggid** value). Only user-defined aggregates (aggregates that are not built in) have entries in **sysaggregates**.

Both a simple index on the **aggid** column and a composite index on the **name** and **owner** columns require unique values.

## SYSAMS

The **sysams** system catalog table contains information that is required for using built-in access methods and those created by the CREATE ACCESS\_METHOD statement of SQL.

The **sysams** table has the following columns.

**Table 2. SYSAMS table column descriptions**

Column	Type	Explanation
<b>am_name</b>	VARCHAR(128, 0)	Name of the access method
<b>am_owner</b>	CHAR(32)	Name of the owner of the access method
<b>am_id</b>	INTEGER	Unique identifying code for an access method  This corresponds to the <b>am_id</b> columns in the <b>sysables</b> , <b>sysindices</b> , and <b>sysopclasses</b> tables.
<b>am_type</b>	CHAR(1)	Type of access method: P = Primary; S = Secondary
<b>am_sptype</b>	CHAR(3)	Types of spaces where the access method can exist: <ul style="list-style-type: none"> <li>• <b>A</b> means the access method supports extspaces and sbspaces. If the access method is built in, such as a B-tree, it also supports dbspaces.</li> <li>• <b>D</b> or <b>a</b> means the access method supports dbspaces only.</li> </ul>

Table 2. SYSAMS table column descriptions (continued)

Column	Type	Explanation
		<ul style="list-style-type: none"> <li>• <b>DS</b> means the access method supports dbspaces and sbspaces.</li> <li>• <b>S</b> or <b>s</b> means the access method supports sbspaces only.</li> <li>• <b>X</b> or <b>x</b> means the access method supports extspaces only.</li> <li>• <b>SX</b> means the access method supports sbspaces and extspaces.</li> </ul>
<b>am_defopclass</b>	INTEGER	<p>Unique identifying code for default-operator class</p> <p>Value is the <b>opclassid</b> from the entry for this operator class in the <b>sysopclasses</b> table.</p>
<b>am_keyscan</b>	INTEGER	<p>Whether a secondary access method supports a key scan</p> <p>(An access method supports a key scan if it can return a key and a rowid from a call to the <b>am_getnext</b> function.) (0 = FALSE; Non-zero = TRUE)</p>
<b>am_unique</b>	INTEGER	<p>Whether a secondary access method can support unique keys (0 = FALSE; Non-zero = TRUE)</p>
<b>am_cluster</b>	INTEGER	<p>Whether a primary access method supports clustering (0 = FALSE; Non-zero = TRUE)</p>
<b>am_rowids</b>	INTEGER	<p>Whether a primary access method supports rowids (0 = FALSE; Non-zero = TRUE)</p>
<b>am_readwrite</b>	INTEGER	<p>Whether a primary access method can both read and write ( 0 = access method is read-only; Non-zero = access method is read/write )</p>
<b>am_parallel</b>	INTEGER	<p>Whether an access method supports parallel execution (0 = FALSE; Non-zero = TRUE)</p>
<b>am_costfactor</b>	SMALLFLOAT	<p>The value to be multiplied by the cost of a scan to normalize it to costing done for built-in access methods</p> <p>The scan cost is the output of the <b>am_scancost</b> function.</p>
<b>am_create</b>	INTEGER	<p>The routine specified for the AM_CREATE purpose for this access method</p> <p>Value = <b>procid</b> for the routine in the <b>sysprocedures</b> table.</p>



Table 2. SYSAMS table column descriptions (continued)

Column	Type	Explanation
<b>am_drop</b>	INTEGER	The routine specified for the AM_DROP purpose function for this access method
<b>am_open</b>	INTEGER	The routine specified for the AM_OPEN purpose function for this access method
<b>am_close</b>	INTEGER	The routine specified for the AM_CLOSE purpose function for this access method
<b>am_insert</b>	INTEGER	The routine specified for the AM_INSERT purpose function for this access method
<b>am_delete</b>	INTEGER	The routine specified for the AM_DELETE purpose function for this access method
<b>am_update</b>	INTEGER	The routine specified for the AM_UPDATE purpose function for this access method
<b>am_stats</b>	INTEGER	The routine specified for the AM_STATS purpose function for this access method
<b>am_scancost</b>	INTEGER	The routine specified for the AM_SCANCOST purpose function for this access method
<b>am_check</b>	INTEGER	The routine specified for the AM_CHECK purpose function for this access method
<b>am_beginscan</b>	INTEGER	The routine specified for the AM_BEGINSCAN purpose function for this access method
<b>am_endscan</b>	INTEGER	The routine specified for the AM_ENDSCAN purpose function for this access method
<b>am_rescan</b>	INTEGER	The routine specified for the AM_RESCAN purpose function for this access method
<b>am_getnext</b>	INTEGER	The routine specified for the AM_GETNEXT purpose function for this access method
<b>am_getbyid</b>	INTEGER	The routine specified for the AM_GETBYID purpose function for this access method
<b>am_build</b>	INTEGER	The routine specified for the AM_BUILD purpose function for this access method
<b>am_init</b>	INTEGER	The routine specified for the AM_INIT purpose function for this access method

**Table 2. SYSAMS table column descriptions (continued)**

Column	Type	Explanation
<b>am_truncate</b>	INTEGER	The routine specified for the AM_TRUNCATE purpose function for this access method
<b>am_expr_pushdown</b>	INTEGER	Reserved for future use Whether parameter descriptors are supported (0 = FALSE; Non-zero = TRUE)

For each of the columns that contain a routine for a purpose function, the value is the **sysprocedures.procid** value for the corresponding routine.

A composite index on the **am\_name** and **am\_owner** columns in this table allows only unique values. The **am\_id** column has a unique index.

For information about access method functions, see the documentation of your access method.

## SYSATTRTYPES

The **sysattrtypes** system catalog table contains information about members of a complex data type. Each row of **sysattrtypes** contains information about elements of a collection data type or fields of a row data type.

The **sysattrtypes** table has the following columns.

**Table 3. SYSATTRTYPES table column descriptions**

Column	Type	Explanation
<b>extended_id</b>	INTEGER	Identifying code of an extended data type  Value is the same as in the <b>sysxdtypes</b> table ( <a href="#">SYSXDTYPES on page 75</a> ).
<b>seqno</b>	SMALLINT	Identifying code of an entry having <b>extended_id</b> type
<b>levelno</b>	SMALLINT	Position of member in collection hierarchy
<b>parent_no</b>	SMALLINT	Value in the <b>seqno</b> column of the complex data type that contains this member
<b>fieldname</b>	VARCHAR(128)	Name of the field in a row type  Null for other complex data types
<b>fieldno</b>	SMALLINT	Field number sequentially assigned by system (from left to right within each row type)
<b>type</b>	SMALLINT	Code for the data type

**Table 3. SYSATTRTYPES table column descriptions (continued)**

Column	Type	Explanation
		See the description of <b>syscolumns.coltype</b> (page <a href="#">SYSCOLUMNS on page 24</a> ).
<b>length</b>	SMALLINT	Length (in bytes) of the member
<b>xtd_type_id</b>	INTEGER	Code identifying this data type  See the description of <b>sysxdtypes.extended_id</b> ( <a href="#">SYSXDTYPES on page 75</a> ).

Two indexes on the **extended\_id** column and the **xtd\_type\_id** column allow duplicate values. A composite index on the **extended\_id** and **seqno** columns allows only unique values.

## SYSAUTOLOCATE

The **sysautolocate** system catalog table indicates which dbspaces are available for automatic table fragmentation. The **sysautolocate** system catalog table is reserved for future use.

**Table 4. SYSAUTOLOCATE table column descriptions**

Column	Type	Explanation
<b>dbsnum</b>	INTEGER	The ID number of the dbspace. 0 indicates multiple dbspaces. Reserved for future use.
<b>dbsname</b>	VARCHAR(128,0)	The name of the dbspace. An asterisk (*) indicates multiple dbspaces. Reserved for future use.
<b>pagesize</b>	SMALLINT	The page size of the dbspace. 0 indicates multiple page sizes. Reserved for future use.
<b>flags</b>	INTEGER	<ul style="list-style-type: none"> <li>• 1 = On. The dbspace is available for automatic table fragmentation.</li> <li>• 2 = Off. The dbspace is not available for automatic table fragmentation.</li> </ul>

Reserved for future use.

You add or remove dbspace from the list of available dbspace by running the `task()` or `admin()` SQL administration API function with one of the `autolocate` database arguments.

The **sysautolocate** system catalog table does not necessarily list every dbspace. For example, if all dbspaces are available for automatic table fragmentation, the table contains one row:

dbsnum	dbsname	pagesize	flags
0	*	0	1

If all but one dbspace is available, the table contains two rows, for example:

dbsnum	dbsname	pagesize	flags
0	*	0	1
12	dbs12	8	2

If all but two dbspaces are unavailable, the table contains three rows, for example:

dbsnum	dbsname	pagesize	flags
0	*	0	2
12	dbs12	8	1
13	dbs13	4	1

## SYSBLOBS

The **sysblobs** system catalog table specifies the storage location of BYTE and TEXT column values. Its name is based on a legacy term for BYTE and TEXT columns, blobs (also known as *simple large objects*), and does not refer to the BLOB data type of HCL OneDB™. The **sysblobs** table contains one row for each BYTE or TEXT column, and has the following columns.

**Table 5. SYSBLOBS table column descriptions**

Column	Type	Explanation
<b>spacename</b>	VARCHAR(128)	Name of partition, dbspace, or family
<b>type</b>	CHAR(1)	Code identifying the type of storage media: M = Magnetic Code identifying the type of storage media: M = Magnetic O = Optical
<b>tabid</b>	INTEGER	Code identifying the table
<b>colno</b>	SMALLINT	Column number within its table

A composite index on **tabid** and **colno** allows only unique values.

For information about the location and size of chunks of blobspaces, dbspaces, and sbspaces for TEXT, BYTE, BLOB, and CLOB columns, see the *HCL OneDB™ Administrator's Guide* and the *HCL OneDB™ Administrator's Reference*.

## SYSCASTS

The **syscasts** system catalog table describes the casts in the database. It contains one row for each built-in cast, each implicit cast, and each explicit cast that a user defines. The **syscasts** table has the following columns.

**Table 6. SYSCASTS table column descriptions**

Column	Type	Explanation
<b>owner</b>	CHAR(32)	Owner of cast (user <b>informix</b> for built-in casts and <i>user</i> name for implicit and explicit casts)
<b>argument_type</b>	SMALLINT	Source data type on which the cast operates
<b>argument_xid</b>	INTEGER	Code for the source data type specified in the <b>argument_type</b> column

**Table 6. SYSCASTS table column descriptions (continued)**

Column	Type	Explanation
<b>result_type</b>	SMALLINT	Code for the data type returned by the cast
<b>result_xid</b>	INTEGER	Data type code of the data type named in the <b>result_type</b> column
<b>routine_name</b>	VARCHAR(128)	Function or procedure implementing the cast
<b>routine_owner</b>	CHAR(32)	Name of owner of the function or procedure specified in the <b>routine_name</b> column
<b>class</b>	CHAR(1)	Type of cast: E = Explicit cast I = Implicit cast S = Built-in cast

If **routine\_name** and **routine\_owner** have NULL values, this indicates that the cast is defined without a routine. This can occur if both of the data types specified in the **argument\_type** and **result\_type** columns have the same length and alignment, and are passed by reference, or passed by value.

A composite index on columns **argument\_type**, **argument\_xid**, **result\_type**, and **result\_xid** allows only unique values. A composite index on columns **result\_type** and **result\_xid** allows duplicate values.

## SYSCHECKS

The **syschecks** system catalog table describes each check constraint defined in the database. Because the **syschecks** table stores both the ASCII text and a binary encoded form of the check constraint, it contains multiple rows for each check constraint. The **syschecks** table has the following columns.

**Table 7. SYSCHECKS table column descriptions**

Column	Type	Explanation
<b>constrid</b>	INTEGER	Unique code identifying the constraint
<b>type</b>	CHAR(1)	Form in which the check constraint is stored: B = Binary encoded s = Select T = Text
<b>seqno</b>	SMALLINT	Line number of the check constraint
<b>checktext</b>	CHAR(32)	Text of the check constraint

The text in the **checktext** column associated with **B** type in the type column is in computer-readable format. To view the text associated with a particular check constraint, use the following query with the appropriate **constrid** code:

```
SELECT * FROM syschecks WHERE constrid=10 AND type='T'
```

Each check constraint described in the **syschecks** table also has its own row in the **sysconstraints** table.

A composite index on the **constrid**, **type**, and **seqno** columns allows only unique values.

## SYSCHECKUDRDEP

The **syscheckudrdep** system catalog table describes each check constraint that is referenced by a user-defined routine (UDR) in the database. The **syscheckudrdep** table has the following columns.

**Table 8. SYSCHECKUDRDEP table column descriptions**

Column	Type	Explanation
<b>udr_id</b>	INTEGER	Unique code identifying the UDR
<b>constraint_id</b>	INTEGER	Unique code identifying the check constraint

Each check constraint described in the **syscheckudrdep** table also has its own row in the **sysconstraints** system catalog table, where the **constrid** column has the same value as the **constraint\_id** column of **syscheckudrdep**.

A composite index on the **udr\_id** and **constraint\_id** columns requires that combinations of these values be unique.

## SYSCOLATTRIBS

The **syscolattrs** system catalog table describes the characteristics of smart large objects, namely CLOB and BLOB data types.

It contains one row for each sbspace referenced in the PUT clause of the CREATE TABLE statement or of the ALTER TABLE statement.

**Table 9. SYSCOLATTRIBS table column descriptions**

Column	Type	Explanation
<b>tabid</b>	INTEGER	Code uniquely identifying the table
<b>colno</b>	SMALLINT	Number of the column that contains the smart large object
<b>extentsize</b>	INTEGER	Pages in smart-large-object extent, expressed in KB
<b>flags</b>	INTEGER	Integer representation of the combination (by addition) of hexadecimal values of the following parameters: <ul style="list-style-type: none"> <li>• LO_NOLOG (0x00000001 = 1) = The smart large object is not logged.</li> <li>• LO_LOG (0x00000010 = 2) = Logging of smart large objects conforms to current log mode of the database.</li> <li>• LO_KEEP_LASTACCESS_TIME (0x00000100 = 4) = Keeps a record of when this column was most recently accessed by a user.</li> <li>• LO_NOKEEP_LASTACCESS_TIME (0x00001000 = 8) = No record is kept of when this column was most recently accessed by a user.</li> </ul>

Table 9. SYSCOLATTRIBS table column descriptions

(continued)

Column	Type	Explanation
		<ul style="list-style-type: none"> <li>• HI_INTEG (0x00010000= 16) = Sbspace data pages have headers and footers to detect incomplete writes and data corruption.</li> <li>• MODERATE_INTEG (0x00100000= 32) = Data pages have headers but no footers.</li> </ul>
<b>flags1</b>	INTEGER	Reserved for future use
<b>sbspace</b>	VARCHAR(128)	Name of the sbspace

A composite index on the **tabid**, **colno**, and **sbspace** columns allows only unique combinations of these values.

## SYSCOLAUTH

The **syscolauth** system catalog table describes each set of discretionary access privileges granted on a column. It contains one row for each set of column-level privileges that are currently granted to a user, to a role, or to the PUBLIC group on a column in the database. The **syscolauth** table has the following columns.

Column	Type	Explanation
<b>grantor</b>	VARCHAR(32)	Authorization identifier of the grantor
<b>grantee</b>	VARCHAR(32)	Authorization identifier of the grantee
<b>tabid</b>	INTEGER	Code uniquely identifying the table
<b>colno</b>	SMALLINT	Column number within the table
<b>colauth</b>	CHAR(3)	3-byte pattern specifying column privileges: s or S = Select, u or U = Update, r or R = References

If the **colauth** privilege code is uppercase (for example, **S** for Select), a user who has this privilege can also grant it to others. If the **colauth** privilege code is lowercase (for example, **s** for Select), the user who has this privilege cannot grant it to others. A hyphen ( - ) indicates the absence of the privilege corresponding to that position within the **colauth** pattern.

A composite index on the **tabid**, **grantor**, **grantee**, and **colno** columns allows only unique values. A composite index on the **tabid** and **grantee** columns allows duplicate values.

## SYSCOLDEPEND

The **syscoldepend** system catalog table tracks the table columns specified in check constraints and in NOT NULL constraints. Because a check constraint can involve more than one column in a table, the **syscoldepend** table can contain multiple rows for each check constraint; one row is created for each column involved in the constraint. The **syscoldepend** table has the following columns.

Column	Type	Explanation
<b>constrid</b>	INTEGER	Code uniquely identifying the constraint
<b>tabid</b>	INTEGER	Code uniquely identifying the table
<b>colno</b>	SMALLINT	Column number within the table

A composite index on the **constrid**, **tabid**, and **colno** columns allows only unique values. A composite index on the **tabid** and **colno** columns allows duplicate values.

See also the **syscheckudrdep** system catalog table in [SYSCHECKUDRDEP on page 21](#), which lists every check constraint that is referenced by a user-defined routine.

See also the **sysreferences** table in [SYSREFERENCES on page 56](#), which describes dependencies of referential constraints.

## SYSCOLUMNS

The **syscolumns** system catalog table describes each column in the database.

One row exists for each column that is defined in a table or view.

**Table 10. The SYSCOLUMNS table**

Column	Type	Explanation
<b>colname</b>	VARCHAR(128)	Column name
<b>tabid</b>	INTEGER	Identifying code of table containing the column
<b>colno</b>	SMALLINT	Column number  The system sequentially assigns this (from left to right within each table).
<b>coltype</b>	SMALLINT	Code indicating the data type of the column:  <ul style="list-style-type: none"> <li>0 = CHAR</li> <li>1 = SMALLINT</li> <li>2 = INTEGER</li> <li>3 = FLOAT</li> <li>4 = SMALLFLOAT</li> <li>5 = DECIMAL</li> <li>6 = SERIAL <sup>1</sup></li> <li>7 = DATE</li> <li>8 = MONEY</li> <li>9 = NULL</li> <li>10 = DATETIME</li> </ul>



Table 10. The SYSCOLUMNS table (continued)

Column	Type	Explanation
		11 = BYTE
		12 = TEXT
		13 = VARCHAR
		14 = INTERVAL
		15 = NCHAR
		16 = NVARCHAR
		17 = INT8
		18 = SERIAL8 <sup>1</sup>
		19 = SET
		20 = MULTISSET
		21 = LIST
		22 = ROW (unnamed)
		23 = COLLECTION
		40 = LVARCHAR fixed-length opaque types <sup>2</sup>
		41 = BLOB, BOOLEAN, CLOB variable-length opaque types <sup>2</sup>
		43 = LVARCHAR (client-side only)
		45 = BOOLEAN
		52 = BIGINT
		53 = BIGSERIAL <sup>1</sup>
		2061 = IDSSECURITYLABEL <sup>2,3</sup>
		4118 = ROW (named)
<b>collength</b>	Any of the following data types: <ul style="list-style-type: none"> <li>• Integer-based</li> <li>• Varying-length character</li> <li>• Time</li> <li>• Fixed-point</li> <li>• Simple-large-object</li> <li>• <a href="#">IDSSECURITYLABEL on page 99</a></li> </ul>	The value depends on the data type of the column. For some data types, the value is the column length (in bytes). See <a href="#">Storing Column Length on page 27</a> for more information.
<b>colmin</b>	INTEGER	Minimum column length (in bytes)
<b>colmax</b>	INTEGER	Maximum column length (in bytes)
<b>extended_id</b>	INTEGER	Data type code, from the <b>sysxdtypes</b> table, of the data type specified in the <b>coltype</b> column

**Table 10. The SYSCOLUMNS table (continued)**

Column	Type	Explanation
<b>seclabelid</b>	INTEGER	The label ID of the security label associated with the column if it is a protected column. NULL otherwise.
<b>colattr</b>	SMALLINT	<p><b>HIDDEN</b></p> <p>1 - Hidden column</p> <p><b>ROWVER</b></p> <p>2 - Row version column</p> <p><b>ROW_CHKSUM</b></p> <p>4 - Row key column</p> <p><b>ER_CHECKVER</b></p> <p>8 - ER row version column</p> <p><b>UPGRD1_COL</b></p> <p>16 - ER auto primary key column</p> <p><b>UPGRD2_COL</b></p> <p>32 - ER auto primary key column</p> <p><b>UPGRD3_COL</b></p> <p>64 - ER auto primary key column</p> <p><b>PK_NOTNULL</b></p> <p>128 - NOT NULL by PRIMARY KEY</p>



**Note:**

<sup>1</sup> In DB-Access, an offset value of 256 is always added to these **coltype** codes because DB-Access sets SERIAL, SERIAL8, and BIGSERIAL columns to NOT NULL.

<sup>2</sup> The built-in opaque data types do not have a unique **coltype** value. They are distinguished by the **extended\_id** column in the [SYSXTDTYPES on page 75](#) system catalog table.

<sup>3</sup> DISTINCT OF VARCHAR(128).

A composite index on **tabid** and **colno** allows only unique values.

The **coltype** codes can be incremented by bitmaps showing the following features of the column.

Bit Value	Significance When Bit Is Set
0x0100	NULL values are not allowed

Bit Value	Significance When Bit Is Set
0x0200	Value is from a host variable
0x0400	Float-to-decimal for networked database server
0x0800	DISTINCT data type
0x1000	Named ROW type
0x2000	DISTINCT type from LVARCHAR base type
0x4000	DISTINCT type from BOOLEAN base type
0x8000	Collection is processed on client system

For example, the **coltype** value 4118 for named row types is the decimal representation of the hexadecimal value 0x1016, which is the same as the hexadecimal **coltype** value for an unnamed row type (0x016), with the named-row-type bit set. The file `$ONEDB_HOME/incl/esql/sqltypes.h` contains additional information about **syscolumns.coltype** codes.

The following table lists the **coltype** values for the built-in opaque data types:

### NOT NULL constraints

Similarly, the **coltype** value is incremented by 256 if the column does not allow NULL values. To determine the data type for such columns, subtract 256 from the value and evaluate the remainder, based on the possible **coltype** values. For example, if the **coltype** value is 262, subtracting 256 leaves a remainder of 6, indicating that the column has a SERIAL data type.

### Storing the column data type

The database server stores the **coltype** value as bitmap, as listed in [SYSCOLUMNS on page 24](#).

### Storing column length

The **collength** column value depends on the data type of the column.

### Integer-based data types

A **collength** value for a BIGINT, BIGSERIAL, DATE, INTEGER, INT8, SERIAL, SERIAL8, or SMALLINT column is machine-independent. The database server uses the following lengths for these integer-based data types of the SQL language.

Integer-based data types	Length (in bytes)
SMALLINT	2
DATE, INTEGER, and SERIAL	4
INT8 and SERIAL8	10
BIGINT and BIGSERIAL	8

## Varying-length character data types

For HCL OneDB™ columns of the LVARCHAR type, **collength** has the value of *max* from the data type declaration, or 2048 if no maximum was specified.

For VARCHAR or NVARCHAR columns, the *max\_size* and *min\_space* values are encoded in the **collength** column using one of these formulas:

- If the **collength** value is positive:

$$\text{collength} = (\text{min\_space} * 256) + \text{max\_size}$$

- If the **collength** value is negative:

$$\text{collength} + 65536 = (\text{min\_space} * 256) + \text{max\_size}$$

## Time data types

As noted previously, DATE columns have a value of 4 in the **collength** column.

For columns of type DATETIME or INTERVAL, **collength** is determined using the following formula:

$$(\text{length} * 256) + (\text{first\_qualifier} * 16) + \text{last\_qualifier}$$

The length is the physical length of the DATETIME or INTERVAL field, and *first\_qualifier* and *last\_qualifier* have values that the following table shows.

Field qualifier	Value	Field qualifier	Value
YEAR	0	FRACTION(1)	11
MONTH	2	FRACTION(2)	12
DAY	4	FRACTION(3)	13
HOUR	6	FRACTION(4)	14
MINUTE	8	FRACTION(5)	15
SECOND	10		

For example, if a DATETIME YEAR TO MINUTE column has a length of 12 (such as YYYY:DD:MO:HH:MI), a *first\_qualifier* value of 0 (for YEAR), and a *last\_qualifier* value of 8 (for MINUTE), then the **collength** value is 3080 (from  $(256 * 12) + (0 * 16) + 8$ ).

## Fixed-point data types

The **collength** value for a MONEY or DECIMAL (*p*, *s*) column can be calculated using the following formula:

$$(\text{precision} * 256) + \text{scale}$$

## Simple-large-object data types

If the data type of the column is BYTE or TEXT, **collength** holds the length of the descriptor.

## Storing Maximum and Minimum Values

The **colmin** and **colmax** values hold the second-smallest and second-largest data values in the column, respectively. For example, if the values in an indexed column are 1, 2, 3, 4, and 5, the **colmin** value is 2 and the **colmax** value is 4. Storing the second-smallest and second-largest data values lets the query optimizer make assumptions about the range of values in the column and, in turn, further refine search strategies.

The **colmin** and **colmax** columns contain values only if the column is indexed and the UPDATE STATISTICS statement has explicitly or implicitly calculated the column distribution. If you store BYTE or TEXT data in the tblspace, the **colmin** value is encoded as -1.

The **colmin** and **colmax** columns are valid only for data types that fit into four bytes: SMALLFLOAT, SMALLINT, INTEGER, and the first four bytes of CHAR. The values for all other noninteger column types are the initial four bytes of the maximum or minimum value, which are treated as integers.

It is better to use UPDATE STATISTICS MEDIUM than to depend on **colmin** and **colmax** values. UPDATE STATISTICS MEDIUM gives better information and is valid for all data types.

HCL OneDB™ does not calculate **colmin** and **colmax** values for user-defined data types. These columns, however, have values for user-defined data types if a user-defined secondary access method supplies them.

## SYSCONSTRAINTS

The **sysconstraints** system catalog table lists the constraints placed on the columns in each database table. An entry is also placed in the **sysindexes** system catalog table (or **sysindices** view for HCL OneDB™) for each unique, primary key, or referential constraint that does not already have a corresponding entry in **sysindexes** or **sysindices**. Because indexes can be shared, more than one constraint can be associated with an index. The **sysconstraints** table has the following columns.

**Table 11. SYSCONSTRAINTS table column descriptions**

Column	Type	Explanation
<b>constrid</b>	SERIAL	Code uniquely identifying the constraint
<b>constrname</b>	VARCHAR(128)	Name of the constraint
<b>owner</b>	VARCHAR(32)	Name of the owner of the constraint
<b>tabid</b>	INTEGER	Code uniquely identifying the table
<b>constrtype</b>	CHAR(1)	Code identifying the constraint type: <ul style="list-style-type: none"> <li>• C = Check constraint</li> <li>• N = Not NULL</li> <li>• P = Primary key</li> <li>• R = Referential</li> </ul>

**Table 11. SYSCONSTRAINTS table column descriptions**

(continued)

Column	Type	Explanation
		<ul style="list-style-type: none"> <li>• T = Table</li> <li>• U = Unique</li> </ul>
<b>idxname</b>	VARCHAR(128)	Name of index corresponding to constraint
<b>collation</b>	CHAR(32)	Collating order at the time when the constraint was created.

A composite index on the **constrname** and **owner** columns allows only unique values. An index on the **tabid** column allows duplicate values, and an index on the **constrid** column allows only unique values.

For check constraints (where **constrtype** = c), the **idxname** is always NULL. Additional information about each check constraint is contained in the **syschecks** and **syscoldepend** system catalog tables.

## SYSDEFAULTS

The **sysdefaults** system catalog table lists the user-defined defaults that are placed on each column in the database. One row exists for each user-defined default value.

The **sysdefaults** table has the following columns:

**Table 12. SYSDEFAULTS table column descriptions**

Column	Type	Explanation
<b>tabid</b>	INTEGER	Code uniquely identifying a table. When the <b>class</b> column contains the code P, then the <b>tabid</b> column references a procedure ID not a table ID.
<b>colno</b>	SMALLINT	Code uniquely identifying a column.
<b>type</b>	CHAR(1)	Code identifying the type of default value: <ul style="list-style-type: none"> <li>C = Current@</li> <li>L = Literal value</li> <li>N = NULL</li> <li>S = Dbservername or Sitename</li> <li>T = Today</li> <li>U = User</li> </ul>
<b>default</b>	CHAR(256)	If <b>sysdefaults.type</b> = L, a literal default value.
<b>class</b>	CHAR(1)	Code identifying what kind of column:

**Table 12. SYSDEFAULTS table column descriptions (continued)**

Column	Type	Explanation
		T = table t = ROW type P = procedure

If no default is specified explicitly in the CREATE TABLE or the ALTER TABLE statement, then no entry exists for that column in the **sysdefaults** table.

If you specify a literal for the default value, it is stored in the **default** column as ASCII text. If the literal value is not of one of the data types listed in the next paragraph, the **default** column consists of two parts. The first part is the 6-bit representation of the binary value of the default value structure. The second part is the default value in ASCII text. A blank space separates the two parts.

If the data type of the column is not CHAR, NCHAR, NVARCHAR, or VARCHAR, or (for HCL OneDB™) BOOLEAN or LVARCHAR, a binary representation of the default value is encoded in the **default** column.

A composite index on the **tabid**, **colno**, and **class** columns allows only unique values.

## SYSDEPEND

The **sysdepend** system catalog table describes how each view or table depends on other views or tables. One row exists in this table for each dependency, so a view based on three tables has three rows. The **sysdepend** table has the following columns.

**Table 13. SYSDEPEND table column descriptions**

Column	Type	Explanation
<b>btid</b>	INTEGER	Code uniquely identifying the base table or view
<b>btype</b>	CHAR(1)	Base object type: T = Table V = View
<b>dtid</b>	INTEGER	Code uniquely identifying a dependent table or view
<b>dtype</b>	CHAR(1)	Code for the type of dependent object; currently, only view (V = View) is implemented

The **btid** and **dtid** columns are indexed and allow duplicate values.

## SYSDIRECTIVES

The **sysdirectives** table stores external optimizer directives that can be applied to queries. Whether queries in client applications can use these optimizer directives depends on the setting of the **IFX\_EXTDIRECTIVES** environment variable

on the client system, as described in Chapter 3, and on the EXT\_DIRECTIVES setting in the configuration file of the database server.

The **sysdirectives** table has the following columns:

**Table 14. SYSDIRECTIVES table column descriptions**

Column	Type	Explanation
<b>id</b>	SERIAL	Unique code identifying the optimizer directive
<b>query</b>	TEXT	Text of the query as it exists in the application
<b>directives</b>	TEXT	Text of the optimizer directive, without comments
<b>directive_code</b>	BYTE	Encoded directive
<b>active</b>	SMALLINT	Integer code that identifies whether this entry is active ( = 1 ) or test only ( = 2 )
<b>hash_code</b>	SMALLINT	For internal use only

NULL values are not valid in the **query** column. There is a unique index on the **id** column.

## SYSDISTRIB

The **sysdistrib** system catalog table stores data-distribution information for the query optimizer to use. Data distributions provide detailed table and column information to the optimizer to improve the choice of execution paths of SELECT statements.

The **sysdistrib** table has the following columns.

**Table 15. SYSDISTRIB table column descriptions**

Column	Type	Explanation
<b>tabid</b>	INTEGER	Code identifying the table from which data values were gathered
<b>colno</b>	SMALLINT	Column number in the source table
<b>seqno</b>	INTEGER	Ordinal number for multiple entries
<b>constructed</b>	DATETIME YEAR TO FRACTION(5)	Date when the data distribution was created
<b>mode</b>	CHAR(1)	Optimization level: M = Medium H = High
<b>resolution</b>	SMALLFLOAT	Specified in the UPDATE STATISTICS statement
<b>confidence</b>	SMALLFLOAT	Specified in the UPDATE STATISTICS statement
<b>encdat</b>	STAT	Statistics information



Table 15. SYSDISTRIB table column descriptions (continued)

Column	Type	Explanation
<b>type</b>	CHAR(1)	Type of statistics: A = <b>enccdat</b> has ASCII-encoded histogram in fixed-length character field S = <b>enccdat</b> has user-defined statistics
<b>smplsize</b>	SMALLFLOAT	A value greater than zero up to 1.0 indicating a proportion of the total rows in the table that UPDATE STATISTICS samples. Values greater than 1.0 indicate the actual number of rows used that UPDATE STATISTICS samples. A value of zero indicates that no sample size is specified. UPDATE STATISTICS HIGH always updates statistics for all rows.
<b>rowssmpld</b>	FLOAT	Number of rows in the sample
<b>constr_time</b>	DATETIME YEAR TO FRACTION(5)	Time when the distribution was recorded
<b>ustnrows</b>	FLOAT	Rows in fragment when distribution was calculated.
<b>ustbuildduration</b>	INTERVAL HOUR TO FRACTION(5)	Time spent calculating the distribution statistics for this column
<b>nupdates</b>	FLOAT	Number of updates to the table
<b>ndeletes</b>	FLOAT	Number of deletes to the table
<b>ninserts</b>	FLOAT	Number of inserts to the table

Information is stored in the **sysdistrib** table when an UPDATE STATISTICS statement with mode MEDIUM or HIGH is executed for a table. (UPDATE STATISTICS LOW does not insert a value into the **mode** column.)

Only user **informix** can select the **enccdat** column.

Each row in the **sysdistrib** system catalog table is keyed by the **tabid** and **colno** for which the statistics are collected.

For built-in data type columns, the **type** field is set to A. The **enccdat** column stores an ASCII-encoded histogram that is broken down into multiple rows, each of which contains 256 bytes.

In HCL OneDB™, for columns of user-defined data types, the **type** field is set to S. The **enccdat** column stores the statistics collected by the **statcollect** user-defined routine in multirepresentational form. Only one row is stored for each **tabid** and **colno** pair. A composite index on the **tabid**, **colno**, and **seqno** columns requires unique combinations of values.

The following three DML counter columns record counts of how many DML operations modifying data rows were performed on the table at the time of generation of column distribution statistics:

- UPDATE operations in **nupdates**
- DELETE operations in **ndeletes**
- and INSERT operations in **ninserts**

These counts can also include rows modified by MERGE statements.

These DML counter columns store the values of the counters from the server partition that exists when distribution statistics are generated. If the AUTO\_STAT\_MODE configuration parameter, or the AUTO\_STAT\_MODE session environment setting, or the AUTO keyword of the UPDATE STATISTICS statement has enabled selective updating of data distribution statistics, the **ninserts**, **ndeletes**, and **nupdates** values can affect whether UPDATE STATISTICS operations refresh existing data distribution statistics. When the UPDATE STATISTICS statement runs in MEDIUM or HIGH mode against the table, the database server compares the stored values in these columns with the current values in the partition. Column distribution statistics for the table are not updated if the sum of the stored values differs from the sum of these current **sysdistrib** DML counter values from the partition page by less than the threshold specified by the setting of the STATCHANGE table attribute or of the STATCHANGE configuration parameter.

## SYSDOMAINS

The **sysdomains** view is not used. It displays columns of other system catalog tables. It has the following columns.

**Table 16. SYSDOMAINS table column descriptions**

Column	Type	Explanation
<b>id</b>	SERIAL	Unique code identifying the domain
<b>owner</b>	CHAR(32)	Name of the owner of the domain
<b>name</b>	VARCHAR(128)	Name of the domain
<b>type</b>	SMALLINT	Code identifying the type of domain

There is no index on this view.

## SYSERRORS

The **syserrors** system catalog table stores information about error, warning, and informational messages returned by DataBlade® modules and user-defined routines using the **mi\_db\_error\_raise()** DataBlade® API function.

For a description of an error message, use the finderr utility.

The **syserrors** table has the following columns.

Column	Type	Explanation
<b>sqlstate</b>	CHAR(5)	SQLSTATE value associated with the error.

Column	Type	Explanation
<b>locale</b>	CHAR(36)	The locale with which this version of the message is associated (for example, <b>en_us.8859-1</b> )
<b>level</b>	SMALLINT	Reserved for future use
<b>seqno</b>	SMALLINT	Reserved for future use
<b>message</b>	VARCHAR(255)	Message text

To create a new message, insert a row directly into the **syserrors** table. By default, all users can view this table, but only users with the DBA privilege can modify it.

A composite index on the **sqlstate**, **locale**, **level**, and **seqno** columns allows only unique values.

## SYSEXTCOLS

The **sysextcols** system catalog table contains a row that describes each of the internal columns in external table **tabid** of format type (**fmttype**) FIXED.

The **sysextcols** table has the following columns.

Column	Type	Explanation
<b>tabid</b>	INTEGER	Unique identifying code of a table
<b>colno</b>	SMALLINT	Code identifying the column
<b>exttype</b>	SMALLINT	Code identifying an external column type
<b>extstart</b>	SMALLINT	Starting position of column in the external data file
<b>extlength</b>	SMALLINT	External column length (in bytes)
<b>nullstr</b>	CHAR(256)	Represents NULL in external data
<b>decprec</b>	SMALLINT	Precision for external decimals
<b>extstype</b>	VARCHAR(128,0)	External type name

No entries are stored in **sysextcols** for DELIMITED or HCL OneDB™ format external files.

You can use the DBSCHEMA utility to write out the description of the external tables. To query these system catalog tables about an external table, use the **tabid** as stored in **systables** with **tabtype** = 'E'.

An index on the **tabid** column allows duplicate values.

## SYSEXTDFILES

The **sysextdfiles** system catalog table contains identifying codes and the paths of external tables.

For each external table, at least one row exists in the **sysextdfiles** system catalog table, which has the following columns.

Column	Type	Explanation
<b>tabid</b>	INTEGER	Unique identifying code of an external table
<b>dfentry</b>	CHAR(469)	Absolute source or target file path
<b>blobdir</b>	CHAR(344)	Absolute or relative directory name
<b>clobdir</b>	CHAR(344)	Absolute or relative directory name

You can use DBSCHEMA to write out the description of the external tables. To query these system catalog tables about an external table, use the **tabid** as stored in **systables** with **tabtype** = 'E'.

An index on the **tabid** column allows duplicate values.

## SYSEXTERNAL

For each external table, a single row exists in the **sysexternal** system catalog table.

The **tabid** column associates the external table record in this system catalog table with an entry in **systables**.

Column	Type	Explanation
<b>tabid</b>	INTEGER	Unique identifying code of an external table
<b>fmttype</b>	CHAR(1)	Type of format: D = (delimited) F = (fixed) I = (HCL OneDB™)
<b>codeset</b>	VARCHAR(128)	Reserved for future use
<b>recdelim</b>	VARCHAR(128)	The record delimiter
<b>flddelim</b>	CHAR(4)	The field delimiter
<b>datefmt</b>	CHAR(8)	Reserved for future use
<b>moneyfmt</b>	CHAR(20)	Reserved for future use
<b>maxerrors</b>	INTEGER	Number of errors to allow
<b>rejectfile</b>	CHAR(464)	Name of the reject file
<b>flags</b>	INTEGER	Optional <b>load</b> flags
<b>ndfiles</b>	INTEGER	Number of data files in <b>sysextdfiles</b>

You can use the **dbschema** utility to write out the description of the external tables. To query these system catalog tables about an external table, use the **tabid** as stored in **systables** with **tabtype** = 'E'.

An index on the **tabid** column allows only unique values.

## SYSFRAGAUTH

The **sysfragauth** system catalog table stores information about the privileges that are granted on table fragments. This table has the following columns.

Table 17. SYSPRAGAUTH table column descriptions

Column	Type	Explanation
<b>grantor</b>	CHAR(32)	Name of the grantor of privilege
<b>grantee</b>	CHAR(32)	Name of the grantee of privilege
<b>tabid</b>	INTEGER	Identifying code of the fragmented table
<b>fragment</b>	VARCHAR(128)	Name of dbspace where fragment is stored
<b>fragauth</b>	CHAR(6)	A 6-byte pattern specifying fragment privileges (including 3 bytes reserved for future use): <ul style="list-style-type: none"> <li>• u or U = Update</li> <li>• i or I = Insert</li> <li>• d or D = Delete</li> </ul>

In the **fragauth** column, an uppercase code (such as **U** for Update) means that the grantee can grant the privilege to other users; a lowercase (for example, **u** for Update) means the user cannot grant the privilege to others. Hyphen (-) indicates the absence of the privilege for that position within the pattern.

A composite index on the **tabid**, **grantor**, **grantee**, and **fragment** columns allows only unique values. A composite index on the **tabid** and **grantee** columns allows duplicate values.

The following example displays the fragment-level privileges for one base table, as they exist in the **sysfragauth** table. In this example, the grantee **rajesh** can grant the Update, Delete, and Insert privileges to other users.

grantor	grantee	tabid	fragment	fragauth
dba	omar	101	dbsp1	-ui--
dba	jane	101	dbsp3	--i--
dba	maria	101	dbsp4	--id--
dba	rajesh	101	dbsp2	-UID--

## SYSPRAGDIST

The **sysfragdist** system catalog table stores fragment-level column statistics for fragmented tables and indexes. One row exists for each table fragment or index fragment.

Only columns in fragmented tables are described here. (For table-level column statistics, see the **sysdistrib** system catalog table.)

The **sysfragdist** table has the following columns.

Column	Type	Explanation
<b>tabid</b>	INTEGER	Unique identifying code of table (= <b>systables.tabid</b> )
<b>fragid</b>	INTEGER	Unique identifying code of fragment (= <b>sysfragments.partnum</b> )
<b>colno</b>	SMALLINT	Unique identifying code of column (= <b>syscolumns.colno</b> )
<b>seqno</b>	SMALLINT	Sequence number (for distributions that span multiple rows)
<b>mode</b>	CHAR(1)	UPDATE STATISTICS mode (H = high, or M = medium)
<b>resolution</b>	SMALLFLOAT	Average percentage of the sample in each bin
<b>confidence</b>	SMALLFLOAT	Estimated likelihood that a MEDIUM mode sample value is equivalent to an exact HIGH mode result
<b>rowssampled</b>	FLOAT	Number of rows in the sample
<b>ustbuildduration</b>	INTERVAL HOUR TO FRACTION(5)	Time spent to calculate the distribution for this column
<b>constr_time</b>	DATETIME YEAR TO FRACTION(5)	Time when the distribution was recorded
<b>ustnrows</b>	FLOAT	Rows in fragment when distribution was calculated.
<b>minibinsize</b>	FLOAT	For internal use only
<b>nupdates</b>	FLOAT	Number of updates to the table
<b>ndeletes</b>	FLOAT	Number of deletes to the table
<b>ninserts</b>	FLOAT	Number of inserts to the table
<b>version</b>	INTEGER	Reserved for future use
<b>dbsnum</b>	INTEGER	Unique identifying code of sbspace where <b>encdist</b> is stored
<b>encdist</b>	STAT	Encrypted fragment distributionReserved for future use

The set of rows with a given combination of **tabid**, **fragid**, and **colno** values identifies the column statistics for that fragment of a table. These statistics can span multiple rows by using the **seqno** column for sequence numbering.

The *mode*, *resolution* and *confidence* values that are specified in the UPDATE STATISTICS MEDIUM or HIGH statement that calculate the column statistics for the fragment are recorded in the **sysfragdist** columns of the same names. To use existing fragment statistics to build table statistics, these three parameters should not change between UPDATE STATISTICS statements that reference the fragments of the same table. The only exception to this is that “H” mode fragmented statistics can be used to build “M” mode table statistics.

Column distribution statistics for the fragment are stored in the column **enclist**. The **dbnum** column stores the identifying code of the smart blob space where the **enclist** object describing this fragment is stored. By default, the SBSPACENAME configuration parameter setting is the identifier of the sbpace whose identifying code is in the **dbnum** column.

The following three columns record counts of how many DML operations modifying data rows were performed on the fragment at the time of generation of column distribution statistics:

- UPDATE operations in **nupdates**
- DELETE operations in **ndeletes**
- and INSERT operations in **ninserts**

These counts can also include rows modified by MERGE statements.

These DML counter columns store the values of the counters from the server partition that existed when distribution statistics were generated. When UPDATE STATISTICS runs in MEDIUM or HIGH mode against the fragmented table with fragment level statistics, the database server compares the stored values in these columns with the current values in the partition.

When the AUTO\_STAT\_MODE configuration parameter, or the AUTO\_STAT\_MODE session environment setting, or the AUTO keyword of the UPDATE STATISTICS statement has enabled selective updating of data distribution statistics, the **ninserts**, **ndeletes**, and **ninserts** values can affect whether UPDATE STATISTICS operations refresh existing data distribution statistics for the fragment. Column statistics for the fragment corresponding to the row in the **sysfragdist** table are not updated if the sum of the stored values differs from the sum of these current DML counter values for the partition page by less than the threshold specified by the setting of the STATCHANGE table attribute or of the STATCHANGE configuration parameter.

## SYSFRAGMENTS

The **sysfragments** system catalog table stores fragmentation information and LOW mode statistical distributions for individual fragments of tables and indexes. One row exists for each table fragment or index fragment.

The **sysfragments** table has the following columns.

Column	Type	Explanation
<b>fragtype</b>	CHAR(1)	Code indicating the type of fragmented object: <ul style="list-style-type: none"> <li>• I = Original index fragment</li> <li>• T = Original table fragment</li> </ul>
<b>tabid</b>	INTEGER	Unique identifying code of table
<b>indexname</b>	VARCHAR(128)	Name of index
<b>colno</b>	INTEGER	Identifying code of TEXT or BYTE column, or the upper limit on the number of rolling window fragments
<b>partn</b>	INTEGER	Identifying code of physical storage location

Column	Type	Explanation
<b>strategy</b>	CHAR(1)	Code for type of fragment distribution strategy: <ul style="list-style-type: none"> <li>• R = Round-robin distribution strategy</li> <li>• E = Expression-based distribution strategy</li> <li>• I = IN DBSPACE clause specifies a storage location as part of distribution strategy</li> <li>• N = raNge-iNterval (or rolliNg wiNdown) distribution strategy</li> <li>• N = raNge-iNterval distribution strategy</li> <li>• L = List distribution strategy</li> <li>• T = Table-based distribution strategy</li> <li>• H = table is a subtable within a table Hierarchy</li> </ul>
<b>location</b>	CHAR(1)	Reserved for future use; shows L for local
<b>servername</b>	VARCHAR(128)	Reserved for future use
<b>evalpos</b>	INTEGER	Position of fragment in the fragmentation list.  For fragmentation by INTERVAL, one of the following values that indicates the type of information in the <b>exprtext</b> field: <ul style="list-style-type: none"> <li>• -1 = List of dbspaces for interval fragments</li> <li>• -2 = Interval value</li> <li>• -3 = Fragmentation key</li> <li>• -4 = Rolling window fragment</li> </ul> Fragmentation by LIST also uses the -3 value.
<b>exprtext</b>	TEXT	Expression for fragmentation strategy  For fragmentation by INTERVAL, LIST, or rolling window, provides the information corresponding to the value of the <b>evalpos</b> field.  For fragmentation by INTERVAL or LIST, provides the information corresponding to the value of the <b>evalpos</b> field.
<b>exprbin</b>	BYTE	Binary version of expression
<b>exprarr</b>	BYTE	Range-partitioning data to optimize expression in range-expression fragmentation strategy
<b>flags</b>	INTEGER	Used internally
<b>dbspace</b>	VARCHAR(128)	Name of dbspace storing this fragment
<b>levels</b>	SMALLINT	Number of B-tree index levels



Column	Type	Explanation
<b>nused</b>	FLOAT	For table-fragmentation strategies: the number of data pages  For index-fragmentation strategies: the number of leaf pages  For rolling window tables: the units for the storage size limit in nrows
<b>nrows</b>	FLOAT	For tables: the number of rows in the fragment.  For indexes: the number of unique keys.  For rolling window tables: the upper limit on storage size in the purge policy.
<b>clust</b>	FLOAT	Degree of index clustering; smaller numbers correspond to greater clustering.
<b>partition</b>	VARCHAR(128)	Fragment name. This can match the name of the dbspace that stores the fragment, or can be an arbitrary name.
<b>version</b>	SMALLINT	Number that increments when fragment statistics is updated
<b>nupdates</b>	FLOAT	Number of updates to the fragment
<b>ndeletes</b>	FLOAT	Number of deletes to the fragment
<b>ninserts</b>	FLOAT	Number of inserts to the fragment

Every fragment has a row in this table. The **evalpos** and **evaltext** fields contain information about individual fragments.

Tables and indexes created with fragmentation by INTERVAL or LIST have additional rows containing information about the fragmentation strategy.

The **strategy** type `T` is used for attached indexes. (This is a fragmented index whose fragmentation strategy is the same as for the table fragmentation.)

For information about the **nupdates**, **ndeletes**, and **ninserts** columns, which in **sysfragments** tabulate DML operations on a table since the most recent recalculation of its distribution statistics, see the description of the three columns that have the same names in the [SYSDISTRIB on page 32](#) system catalog table.

In HCL OneDB™, a composite index on the **fragtype**, **tabid**, **indexname**, and **evalpos** columns allows duplicate values.

## SYSINDEXES

The **sysindexes** table is a view on the **sysindices** table. It contains one row for each index in the database.

The **sysindexes** table has the following columns.

**Table 18. SYSINDEXES table column descriptions**

Column	Type	Explanation
<b>idxname</b>	VARCHAR(1 28)	Index name
<b>owner</b>	VARCHAR(32)	Owner of index (user <b>informix</b> for system catalog tables and <i>username</i> for database tables)
<b>tabid</b>	INTEGER	Unique identifying code of table
<b>idxtype</b>	CHAR(1)	Index type:  U = Unique D = Duplicates allowed G = Nonbitmap generalized-key index g = Bitmap generalized-key index u = unique, bitmap d = nonunique, bitmap
<b>clustered</b>	CHAR(1)	Clustered or nonclustered index (C = Clustered)
<b>part1</b>	SMALLINT	Column number ( <b>colno</b> ) of a single index or the 1st component of a composite index
<b>part2</b>	SMALLINT	2nd component of a composite index
<b>part3</b>	SMALLINT	3rd component of a composite index
<b>part4</b>	SMALLINT	4th component of a composite index
<b>part5</b>	SMALLINT	5th component of a composite index
<b>part6</b>	SMALLINT	6th component of a composite index
<b>part7</b>	SMALLINT	7th component of a composite index
<b>part8</b>	SMALLINT	8th component of a composite index
<b>part9</b>	SMALLINT	9th component of a composite index
<b>part10</b>	SMALLINT	10th component of a composite index
<b>part11</b>	SMALLINT	11th component of a composite index
<b>part12</b>	SMALLINT	12th component of a composite index
<b>part13</b>	SMALLINT	13th component of a composite index
<b>part14</b>	SMALLINT	14th component of a composite index
<b>part15</b>	SMALLINT	15th component of a composite index

**Table 18. SYSINDEXES table column descriptions**

(continued)

Column	Type	Explanation
<b>part16</b>	SMALLINT	16th component of a composite index
<b>levels</b>	SMALLINT	Number of B-tree levels
<b>leaves</b>	INTEGER	Number of leaves
<b>nunique</b>	INTEGER	Number of unique keys in the first column
<b>clust</b>	INTEGER	Degree of clustering; smaller numbers correspond to greater clustering
<b>idxflags</b>	INTEGER	Bitmap storing the current locking mode of the index

As with most system catalog tables, changes that affect existing indexes are reflected in this table only after you run the UPDATE STATISTICS statement.

Each **part1** through **part16** column in this table holds the column number (**colno**) of one of the 16 possible parts of a composite index. If the component is ordered in descending order, the **colno** is entered as a negative value. The columns are filled in for B-tree indexes that do not use user-defined data types or functional indexes. For generic B-trees and all other access methods, the **part1** through **part16** columns all contain zeros.

The **clust** column is blank until the UPDATE STATISTICS statement is run on the table. The maximum value is the number of rows in the table, and the minimum value is the number of data pages in the table.

## SYSINDICES

The **sysindices** system catalog table describes the indexes in the database. It stores LOW mode statistics for all indexes, and contains one row for each index that is defined in the database.

**Table 19. sysindices system catalog table columns**

Column	Type	Explanation
<b>idxname</b>	VARCHAR(28)	Name of index
<b>owner</b>	VARCHAR(32)	Name of owner of index (user <b>informix</b> for system catalog tables and <i>username</i> for database tables)
<b>tabid</b>	INTEGER	Unique identifying code of table
<b>idxtype</b>	CHAR(1)	Uniqueness status

Table 19. sysindices system catalog table columns

(continued)

Column	Type	Explanation
		U = Unique values required D = Duplicates allowed
<b>clustered</b>	CHAR(1)	Clustered or nonclustered status (C = Clustered)
<b>levels</b>	SMALLINT	Number of tree levels
<b>leaves</b>	FLOAT	Number of leaves
<b>unique</b>	FLOAT	Number of unique keys in the first column
<b>cluster</b>	FLOAT	Degree of clustering; smaller numbers correspond to greater clustering. The maximum value is the number of rows in the table, and the minimum value is the number of data pages in the table. This column is blank until UPDATE STATISTICS is run on the table.
<b>rows</b>	FLOAT	Estimated number of rows in the table (zero until UPDATE STATISTICS is run on the table)
<b>index keys</b>	INDEXKEYARRAY	Internal representation of the index keys. Column can have up to three fields, in the format: <b>procid</b> , ( <i>col1,col2,...,coln</i> ), <b>opclassid</b> where $1 < n < 341$
<b>amid</b>	INTEGER	Unique identifying code of the access method that implements this index. (Value = <b>am_id</b> for that access method in the <b>sysams</b> table.)
<b>amparam</b>	LVARCHAR(2048)	List of parameters used to customize the <b>amid</b> access method behavior
<b>collation</b>	CHAR(32)	Database locale whose collating order was in effect at the time of index creation
<b>page size</b>	INTEGER	Size of the page, in bytes, where this index is stored
<b>nhashcols</b>	SMALLINT	Number of hashed columns in a FOT index
<b>nbuckets</b>	SMALLINT	Number of subtrees (buckets) in a forest of trees (FOT) index

Table 19. sysindices system catalog table columns

(continued)

Col umn	Type	Explanation
<b>ustlo</b>	DATETIME	Date and time when index statistics were last recorded
<b>wts</b>	YEAR TO FRACTION	
<b>ustbu</b>	INTERVAL	Time required to calculate index statistics
<b>ilddu</b>	HOUR TO FRACTION	
<b>ion</b>	(5)	
<b>nupdat</b>	FLOAT	Number of updates to the table
<b>ndele</b>	FLOAT	Number of deletes to the table
<b>ninserts</b>	FLOAT	Number of inserts to the table
<b>firstext</b>	INT	Size (in KB) of the first extent of the index
<b>nextext</b>	INT	Size (in KB) of the next extent of the index
<b>indexattr</b>	INT	<ul style="list-style-type: none"> <li>• 0x00000001 = The index has a partial column key</li> <li>• 0x00000002 = The index is compressed</li> <li>• 0x00000004 = The index is on a BSON column</li> </ul>
<b>bson</b>	LVARCHAR(2048)	BSON index information



**Tip:** This system catalog table is changed from Version 7.2 of HCL OneDB™. The earlier schema of this system catalog table is still available as a view that can be accessed under its original name: **sysindexes**. See [SYSINDEXES on page 41](#).

Changes that affect existing indexes are reflected in this system catalog table only after you run the UPDATE STATISTICS statement.

The fields within the **indexkeys** columns have the following significance:

- The **procid** (as in **sysprocedures**) exists only for a functional index on return values of a function defined on columns of the table.
- The list of columns (*col1, col2, ... , coln*) in the second field identifies the columns on which the index is defined. The maximum is language-dependent: up to 341 for an SPL or Java™ UDR; up to 102 for a C UDR.
- The **opclassid** identifies the secondary access method that the database server used to build and to search the index. This is the same as the **sysopclasses.opclassid** value for the access method.

For information about the **nupdates**, **ndeletes**, and **ninserts** columns, which in **sysindices** tabulate DML operations on an index since the most recent recalculation of its distribution statistics, see the description of the three columns that have the same names in the [SYSDISTRIB on page 32](#) system catalog table.

The **fectsize** column shows the user-defined first extent size (in kilobytes) that the optional EXTENT SIZE clause specified in the CREATE INDEX statement that defined the index. Similarly, the **nextsize** column shows the user-defined next extent size (in kilobytes) that the optional NEXT SIZE clause specified in the CREATE INDEX statement. Each of these columns displays a value of zero ( 0 ) if the corresponding EXTENT SIZE or NEXT SIZE clause was omitted when the index was created.

If the CREATE INDEX statement that defines a new index includes no explicit extent size specifications, the database server automatically calculates the first and next extent sizes, but the **fectsize** and **nextsize** column values are set to 0. When the database server is converted from a release earlier than Version 11.70, the **fectsize** and **nextsize** values for every migrated index are 0.

The **tabid** column is indexed and allows duplicate values. A composite index on the **idxname**, **owner**, and **tabid** columns allows only unique values.

## SYSINHERITS

The **sysinherits** system catalog table stores information about table hierarchies and named ROW type inheritance. Every supertype, subtype, supertable, and subtable in the database has a corresponding row in the **sysinherits** table.

Column	Type	Explanation
<b>child</b>	INTEGER	Identifying code of the subtable or subtype
<b>parent</b>	INTEGER	Identifying code of the supertable or supertype
<b>class</b>	CHAR(1)	Inheritance class: <b>T</b> = named ROW type <b>T</b> = table

The **child** and **parent** values are from **sysxdtypes.extended\_id** for named ROW types, or from **sysstables.tabid** for tables. Simple indexes on the **child** and **parent** columns allow duplicate values.

## SYSLANGAUTH

The **syslangauth** system catalog table contains the authorization information about computer languages that are used to write user-defined routines (UDRs).

**Table 20. SYSLANGAUTH table column descriptions**

Column	Type	Explanation
<b>grantor</b>	VARCHAR(32)	Name of the grantor of the language authorization
<b>grantee</b>	VARCHAR(32)	Name of the grantee of the language authorization
<b>langid</b>	INTEGER	Identifying code of language in <b>sysroutinelangs</b> table
<b>langauth</b>	CHAR(1)	The language authorization:  u = Usage privilege granted U = Usage privilege granted WITH GRANT OPTION

A composite index on the **langid**, **grantor**, and **grantee** columns allows only unique values. A composite index on the **langid** and **grantee** columns allows duplicate values.

## SYSLOGMAP

The **syslogmap** system catalog table contains fragmentation information.

**Table 21. SYSLOGMAP table column descriptions**

Column	Type	Explanation
<b>tabloc</b>	INTEGER	Code for the location of a table in another database
<b>tabid</b>	INTEGER	Unique identifying code of the table
<b>fragid</b>	INTEGER	Identifying code of the fragment
<b>flags</b>	INTEGER	Bitmap of modifiers from declaration of fragment

A simple index on the **tabloc** column and a composite index on the **tabid** and **fragid** columns do not allow duplicate values.

## SYSOBJSTATE

The **sysobjstate** system catalog table stores information about the state (object mode) of database objects. The types of database objects that are listed in this table are indexes, triggers, and constraints.

Every index, trigger, and constraint in the database has a corresponding row in the **sysobjstate** table if a user creates the object. Indexes that the database server creates on the system catalog tables are not listed in the **sysobjstate** table because their object mode cannot be changed.

The **sysobjstate** table has the following columns.

**Table 22. SYSOBJSTATE table column descriptions**

Col umn	Type	Explanation
<b>objtype</b>	CHAR(1)	Code for the type of database object: <ul style="list-style-type: none"> <li>• C = Constraint</li> <li>• I = Index</li> <li>• T = Trigger</li> </ul>
<b>owner</b>	VARCHAR(32)	Authorization identifier of the owner of the database object
<b>name</b>	VARCHAR(128)	Name of the database object
<b>tabid</b>	INTEGER	Identifying code of table on which the object is defined
<b>state</b>	CHAR(1)	The current state (object mode) of the database object. This value can be one of the following codes: <ul style="list-style-type: none"> <li>• D = Disabled</li> <li>• E = Enabled</li> <li>• F = Filtering with no integrity-violation errors</li> <li>• G = Filtering with integrity-violation error</li> </ul>

A composite index on the **objtype**, **name**, **owner**, and **tabid** columns allows only unique combinations of values. A simple index on the **tabid** column allows duplicate values.

## SY SOPCLASSES

The **sysopclasses** system catalog table contains information about operator classes associated with secondary access methods. It contains one row for each operator class that has been defined in the database. The **sysopclasses** table has the following columns.

Column	Type	Explanation
<b>opclassname</b>	VARCHAR(128)	Name of the operator class



Column	Type	Explanation
<b>owner</b>	VARCHAR(32)	Name of the owner of the operator class
<b>amid</b>	INTEGER	Identifying code of the secondary access method associated with this operator class
<b>opclassid</b>	SERIAL	Identifying code of the operator class
<b>ops</b>	LVARCHAR(2048)	List of names of the operators that belong to this operator class
<b>support</b>	LVARCHAR(2048)	List of names of support functions defined for this operator class

The **opclassid** value corresponds to the **sysams.am\_defopclass** value that specifies the default operator class for the secondary access method that the **amid** column specifies.

The **sysopclasses** table has a composite index on the **opclassname** and **owner** columns and an index on **opclassid** column. Both indexes allow only unique values.

## SYSOPCLSTR

The **sysopclstr** system catalog table defines each optical cluster in the database. The table contains one row for each optical cluster.

The **sysopclstr** table has the following columns.

Column	Type	Explanation
<b>owner</b>	VARCHAR(32)	Name of the owner of the optical cluster
<b>clstrname</b>	VARCHAR(128)	Name of the optical cluster
<b>clstrsize</b>	INTEGER	Size of the optical cluster
<b>tabid</b>	INTEGER	Unique identifying code for the table
<b>blobcol1</b>	SMALLINT	BYTE or TEXT column number 1
<b>blobcol2</b>	SMALLINT	BYTE or TEXT column number 2
<b>blobcol3</b>	SMALLINT	BYTE or TEXT column number 3
<b>blobcol4</b>	SMALLINT	BYTE or TEXT column number 4
<b>blobcol5</b>	SMALLINT	BYTE or TEXT column number 5
<b>blobcol6</b>	SMALLINT	BYTE or TEXT column number 6
<b>blobcol7</b>	SMALLINT	BYTE or TEXT column number 7
<b>blobcol8</b>	SMALLINT	BYTE or TEXT column number 8
<b>blobcol9</b>	SMALLINT	BYTE or TEXT column number 9
<b>blobcol10</b>	SMALLINT	BYTE or TEXT column number 10

Column	Type	Explanation
<b>blobcol11</b>	SMALLINT	BYTE or TEXT column number 11
<b>blobcol12</b>	SMALLINT	BYTE or TEXT column number 12
<b>blobcol13</b>	SMALLINT	BYTE or TEXT column number 13
<b>blobcol14</b>	SMALLINT	BYTE or TEXT column number 14
<b>blobcol15</b>	SMALLINT	BYTE or TEXT column number 15
<b>blobcol16</b>	SMALLINT	BYTE or TEXT column number 16
<b>clstrkey1</b>	SMALLINT	Cluster key number 1
<b>clstrkey2</b>	SMALLINT	Cluster key number 2
<b>clstrkey3</b>	SMALLINT	Cluster key number 3
<b>clstrkey4</b>	SMALLINT	Cluster key number 4
<b>clstrkey5</b>	SMALLINT	Cluster key number 5
<b>clstrkey6</b>	SMALLINT	Cluster key number 6
<b>clstrkey7</b>	SMALLINT	Cluster key number 7
<b>clstrkey8</b>	SMALLINT	Cluster key number 8
<b>clstrkey9</b>	SMALLINT	Cluster key number 9
<b>clstrkey10</b>	SMALLINT	Cluster key number 10
<b>clstrkey11</b>	SMALLINT	Cluster key number 11
<b>clstrkey12</b>	SMALLINT	Cluster key number 12
<b>clstrkey13</b>	SMALLINT	Cluster key number 13
<b>clstrkey14</b>	SMALLINT	Cluster key number 14
<b>clstrkey15</b>	SMALLINT	Cluster key number 15
<b>clstrkey16</b>	SMALLINT	Cluster key number 16

The contents of this table are sensitive to CREATE OPTICAL CLUSTER, ALTER OPTICAL CLUSTER, and DROP OPTICAL CLUSTER statements that have been executed on databases that support optical cluster subsystems. Changes that affect existing optical clusters are reflected in this table only after you run the UPDATE STATISTICS statement.

A composite index on the **clstrname** and **owner** columns allows only unique values. A simple index on the **tabid** column allows duplicate values.

## SYSPROCAUTH

The **sysprocauth** system catalog table describes the privileges granted on a procedure or function. It contains one row for each set of privileges that is granted. The **sysprocauth** table has the following columns.

**Table 23. SYSPROCAUTH table column descriptions**

Column	Type	Explanation
<b>grantor</b>	VARCHAR(32)	Name of grantor of privileges to access the routine
<b>grantee</b>	VARCHAR(32)	Name of grantee of privileges to access the routine
<b>procid</b>	INTEGER	Unique identifying code of the routine
<b>procauth</b>	CHAR(1)	Type of privilege granted on the routine: e = Execute privilege on routine E = Execute privilege WITH GRANT OPTION

A composite index on the **procid**, **grantor**, and **grantee** columns allows only unique values. A composite index on the **procid** and **grantee** columns allows duplicate values.

## SYSPROCBODY

The **sysprocbody** system catalog table describes the compiled version of each procedure or function in the database. Because the **sysprocbody** table stores the text of the routine, each routine can have multiple rows. The **sysprocbody** table has the following columns.

**Table 24. SYSPROCBODY table column descriptions**

Column	Type	Explanation
<b>procid</b>	INTEGER	Unique identifying code for the routine
<b>datakey</b>	CHAR(1)	Type of information in the <b>data</b> column: A = Routine alter SQL (will not change this value after update statistics) D = Routine user documentation text E = Time of creation information L = Literal value (that is, literal number or quoted string) P = Interpreter instruction code (p-code) R = Routine return value type list

**Table 24. SYSPROCBODY table column descriptions (continued)**

Column	Type	Explanation
		S = Routine symbol table T = Routine text creation SQL
<b>seqno</b>	INTEGER	Line number within the routine
<b>data</b>	CHAR(256)	Actual text of the routine

The A flag indicates the procedure modifiers are altered. ALTER ROUTINE statement updates only modifiers and not the routine body. UPDATE STATISTICS updates the query plan and not the routine modifiers, and the value of datakey will not be changed from A. The A flag marks all the procedures and functions that have altered modifiers, including overloaded procedures and functions. The T flag is used for routine creation text.

The **data** column contains actual data, which can be in one of these formats:

- Encoded return values list
- Encoded symbol table
- Literal data
- P-code for the routine
- Compiled code for the routine
- Text of the routine and its documentation

A composite index on the **procid**, **datakey**, and **seqno** columns allows only unique values.

## SYSPROCCOLUMNS

The **sysproccolumns** system catalog table stores information about return types and parameter names of all UDRs in SYSPROCEDURES.

A composite index on the **procid** and **paramid** columns in this table allows only unique values.

**Table 25. SYSPROCCOLUMNS table column descriptions**

Column	Type	Explanation
<b>procid</b>	INTEGER	Unique identifying code of the routine
<b>paramid</b>	INTEGER	Unique identifying code of the parameter
<b>paramname</b>	VARCHAR (IDENTSIZE)	Name of the parameter
<b>paramtype</b>	SMALLINT	Identifies the type of parameter
<b>paramlen</b>	SMALLINT	Specifies the length of the parameter
<b>paramxid</b>	INTEGER	Specifies the extended type ID for the parameter

**Table 25. SYSPROCCOLUMNS table column descriptions (continued)**

Column	Type	Explanation
<b>paramattr</b>	INTEGER	0 = Parameter is of unknown type 1 = Parameter is INPUT mode 2 = Parameter is INOUT mode 3 = Parameter is multiple return value 4 = Parameter is OUT mode 5 = Parameter is a return value

## SYSPROCEDURES

The **sysprocedures** system catalog table lists the characteristics for each function and procedure that is registered in the database. It contains one row for each routine.

Each function in **sysprocedures** has a unique value, **procid**, called a *routine identifier*. Throughout the system catalog, a function is identified by its routine identifier, not by its name.

The **sysprocedures** table has the following columns.

**Table 26. SYSPROCEDURES table column descriptions**

Column	Type	Explanation
<b>procname</b>	VARCHAR(128)	Name of routine
<b>owner</b>	VARCHAR(32)	Name of owner
<b>procid</b>	SERIAL	Unique identifying code for the routine
<b>mode</b>	CHAR(1)	Mode type:  D or d = DBA O or o = Owner P or p = Protected R or r = Restricted T or t = Trigger
<b>retsize</b>	INTEGER	Compiled size (in bytes) of returned values
<b>symsize</b>	INTEGER	Compiled size (in bytes) of symbol table
<b>datasize</b>	INTEGER	Compiled size (in bytes) of constant data
<b>codesize</b>	INTEGER	Compiled size (in bytes) of routine code
<b>numargs</b>	INTEGER	Number of arguments to routine
<b>isproc</b>	CHAR(1)	Specifies if the routine is a procedure or a function:  p = procedure f = function

**Table 26. SYSPROCEDURES table column descriptions (continued)**

Column	Type	Explanation
<b>specificname</b>	VARCHAR(128)	Specific name for the routine
<b>externalname</b>	VARCHAR(255)	Location of the external routine. This item is language-specific in content and format.
<b>paramstyle</b>	CHAR(1)	Parameter style: <b>T</b> = HCL OneDB™
<b>langid</b>	INTEGER	Language code (in <b>sysroutinelangs</b> table)
<b>paramtypes</b>	RTNPARAMTYPES	Information describing the parameters of the routine
<b>variant</b>	BOOLEAN	Whether the routine is VARIANT or not:  t = is VARIANT f = is not VARIANT
<b>client</b>	BOOLEAN	Reserved for future use
<b>handlesnulls</b>	BOOLEAN	NULL handling indicator:  t = handles NULLs f = does not handle NULLs
<b>percallcost</b>	INTEGER	Amount of CPU per call  Integer cost to execute UDR: cost/call - 0 -(2^31-1)
<b>commutator</b>	VARCHAR(128)	Name of commutator function
<b>negator</b>	VARCHAR(128)	Name of the negator function
<b>selfunc</b>	VARCHAR(128)	Name of function to estimate selectivity of the UDR
<b>internal</b>	BOOLEAN	Specifies if the routine can be called from SQL:  t = routine is internal, not callable from SQL f = routine is external, callable from SQL
<b>class</b>	CHAR(18)	CPU class by which the routine should be executed
<b>stack</b>	INTEGER	Stack size in bytes required per invocation
<b>parallelizable</b>	BOOLEAN	Parallelization indicator for UDR:

**Table 26. SYSPROCEDURES table column descriptions (continued)**

Column	Type	Explanation
		t = parallelizable f = not parallelizable
<b>costfunc</b>	VARCHAR(128)	Name of the cost function for the UDR
<b>selconst</b>	SMALLFLOAT	Selectivity constant for UDR
<b>procflags</b>	INTEGER	For internal use only
<b>collation</b>	CHAR(32)	Collating order at the time when the routine was created

In the **mode** column, the R mode is a special case of the O mode. A routine is in restricted (R) mode if it was created with a specified owner who is different from the routine creator. If routine statements involving a remote database are executed, the database server uses the access privileges of the user who executes the routine instead of the privileges of the routine owner. In all other scenarios, R-mode routines behave the same as O-mode routines.

The database server can create protected routines for internal use. The **sysprocedures** table identifies these protected routines with the letter **p** or **P** in the **mode** column, where **p** indicates an SPL routine. Protected routines have the following restrictions:

- You cannot use the ALTER FUNCTION, ALTER PROCEDURE, or ALTER ROUTINE statements to modify protected routines.
- You cannot use the DROP FUNCTION, DROP PROCEDURE, or DROP ROUTINE statements to unregister protected routines.
- You cannot use the dbschema utility to display protected routines.

In earlier versions, protected SPL routines were indicated by a lowercase **p**. Starting with version 9.0, protected SPL routines are treated as DBA routines and cannot be Owner routines. Thus **D** and **O** indicate DBA routines and Owner routines, while **d** and **o** indicate protected DBA routines and protected Owner routines.

The trigger mode designates user-defined SPL routines that can be invoked only from the FOR EACH ROW section of a triggered action.



**Important:** After you issue the SET SESSION AUTHORIZATION statement, the database server assigns a restricted mode to all Owner routines that you created while using the new identity.

A unique index is defined on the **procid** column. A composite index on the **procname**, **isproc**, **numargs**, and **owner** columns allows duplicate values, as does a composite index on the **specificname** and **owner** columns.

## SYSPROCPLAN

The **sysprocplan** system catalog table describes the query-execution plans and dependency lists for data-manipulation statements within each routine. Because different parts of a routine plan can be created on different dates, this table can contain multiple rows for each routine.

**Table 27. SYSPROCPLAN table column descriptions**

Column	Type	Explanation
<b>procid</b>	INTEGER	Identifying code for the routine
<b>planid</b>	INTEGER	Identifying code for the plan
<b>datakey</b>	CHAR(1)	Type of information stored in <b>data</b> column:  D = Dependency list I = Information record Q = Execution plan
<b>seqno</b>	INTEGER	Line number within the plan
<b>created</b>	DATE	Date when plan was created
<b>datasize</b>	INTEGER	Size (in bytes) of the list or plan
<b>data</b>	CHAR(256)	Encoded (compiled) list or plan

Before a routine is run, its dependency list in the **data** column is examined. If the major version number of a table accessed by the plan has changed, or if any object that the routine uses has been modified since the plan was optimized (for example, if an index has been dropped), then the plan is optimized again. When **datakey** is I, the **data** column stores information about UPDATE STATISTICS and PDQPRIORITY.

It is possible to delete all the plans for a given routine by using the DELETE statement on **sysprocplan**. When the routine is subsequently executed, new plans are automatically generated and recorded in **sysprocplan**. The UPDATE STATISTICS FOR PROCEDURE statement also updates this table.

A composite index on the **procid**, **planid**, **datakey**, and **seqno** columns allows only unique values.

## SYSREFERENCES

The **sysreferences** system catalog table lists all referential constraints on columns. It contains a row for each referential constraint in the database.



**Table 28. SYSREFERENCES table column descriptions**

Column	Type	Explanation
<b>constrid</b>	INTEGER	Code uniquely identifying the constraint
<b>primary</b>	INTEGER	Identifying code of the corresponding primary key
<b>ptabid</b>	INTEGER	Identifying code of the table that is the primary key
<b>updrule</b>	CHAR(1)	Reserved for future use; displays an R
<b>delrule</b>	CHAR(1)	Whether constraint uses cascading delete or restrict rule:  C = Cascading delete R = Restrict (default)
<b>matchtype</b>	CHAR(1)	Reserved for future use; displays an N
<b>pendant</b>	CHAR(1)	Reserved for future use; displays an N

The **constrid** column is indexed and allows only unique values. The **primary** column is indexed and allows duplicate values.

## SYSROLEAUTH

The **sysroleauth** system catalog table describes the roles that are granted to users. It contains one row for each role that is granted to a user in the database. The **sysroleauth** table has the following columns.

**Table 29. SYSROLEAUTH table column descriptions**

Column	Type	Explanation
<b>rolename</b>	VARCHAR(32)	Name of the role
<b>grantee</b>	VARCHAR(32)	Name of the grantee of the role
<b>is_grantable</b>	CHAR(1)	Specifies whether the role is grantable:  Y = Grantable N = Not grantable

The **is\_grantable** column indicates whether the role was granted with the WITH GRANT OPTION of the GRANT statement.

A composite index on the **rolename** and **grantee** columns allows only unique values.

## SYSROUTINELANGS

The **sysroutinelangs** system catalog table lists the supported programming languages for user-defined routines (UDRs). It has these columns.

Column	Type	Explanation
<b>langid</b>	SERIAL	Code uniquely identifying a supported language
<b>langname</b>	CHAR(30)	Name of the language, such as C or SPL
<b>langinitfunc</b>	VARCHAR(128)	Name of initialization function for the language
<b>langpath</b>	CHAR(255)	Directory path for the UDR language
<b>langclass</b>	CHAR(18)	Name of the class of the UDR language

An index on the **langname** column allows duplicate values.

## SYSSECLABELAUTH

The **sysseclabelauth** system catalog table records the LBAC labels that have been granted to users. It has these columns.

Column	Type	Explanation
<b>GRANTEE</b>	CHAR(32)	The name of the label grantee
<b>secpolicyid</b>	INTEGER	The ID of the security policy to which the security label belongs.
<b>readseclabelid</b>	INTEGER	The security label ID of the security label granted for read access
<b>writeseclabelid</b>	INTEGER	The security label ID of the security label granted for write access

## SYSSECLABELCOMPONENTS

The **sysseclabelcomponents** system catalog table records security label components. It has these columns.

Column	Type	Explanation
<b>compname</b>	VARCHAR(128)	Component name
<b>compid</b>	SERIAL	Component ID
<b>comptype</b>	CHAR(1)	The component type:

Column	Type	Explanation
		A = array S = set T = tree
<b>numelements</b>	INTEGER	Number of elements in the component
<b>coveringinfo</b>	VARCHAR(128)	Internal encoding information
<b>numalters</b>	SMALLINT	Numbers of alter operations that have been performed on the component

## SYSSECLABELCOMPONENTELEMENTS

The **sysseclabelcomponentelements** system catalog table records the values of component elements of security labels. It has these columns.

Column	Type	Explanation
<b>compid</b>	INTEGER	Component ID
<b>element</b>	VARCHAR(32)	Element name
<b>elementencoding</b>	CHAR(8)	Encoded form of the element
<b>parentelement</b>	VARCHAR(32)	The name of the parent elements for tree components. The value is NULL for the following items:  Set components Array components Root nodes of a tree component
<b>alterversion</b>	SMALLINT	The number of the alter operation when the element is added. This value is used by the <b>dbexport</b> and <b>dbimport</b> commands.

## SYSSECLABELNAMES

The **sysseclabelnames** system catalog table records the security label names. It has these columns.

Column	Type	Explanation
<b>secpolicyid</b>	INTEGER	The ID of the security policy to which the security label belongs.
<b>seclabelname</b>	VARCHAR(128)	The name of the security label

Column	Type	Explanation
<b>seclabelid</b>	INTEGER	The ID of the security label

## SYSSECLABELS

The **sysseclabels** system catalog table records the security label encoding. It has these columns.

Column	Type	Explanation
<b>secpolicyid</b>	INTEGER	ID of the security policy to which the security label belongs
<b>seclabelid</b>	INTEGER	Security label ID
<b>sysseclabelnames</b>	VARCHAR(128)	Security label encoding

## SYSSECPOLICIES

The **syssecpolicies** system catalog table records security policies It has these columns.

Column	Type	Explanation
<b>secpolicyname</b>	VARCHAR(128)	Security policy name
<b>secpolicyid</b>	SERIAL	Security policy ID
<b>numcomps</b>	SMALLINT	Number of security label components in the security policy
<b>comptypelist</b>	CHAR(16)	An ordered list of the type of each component in the policy.  A = array S = set T = tree - = Beyond NUMCOMPS
<b>overrideseclabel</b>	CHAR(1)	Indicates the behavior when a user's security label and exemption credentials do not allow them to insert or update a data row with the security that is label provided on the INSERT or UPDATE SQL statement.

Column	Type	Explanation
		<ul style="list-style-type: none"> <li>• Y: The security label provided is ignored and replaced by the user's security label for write access.</li> <li>• N: Return an error when not authorized to write a security label.</li> </ul>

## SYSSECPOLICYCOMPONENTS

The **syssecpolicycomponents** system catalog table records the components for each security policies. It has these columns.

Column	Type	Explanation
<b>secpolicyid</b>	INTEGER	Security policy ID
<b>compid</b>	INTEGER	ID of a component of the label security policy
<b>compno</b>	SMALLINT	Position of the security label component as it exists in the security policy, starting with position 1.

## SYSSECPOLICYEXEMPTIONS

The **syssecpolicyexemptions** system catalog table records the exemptions that have been given to users. It has these columns.

Column	Type	Explanation
<b>grantee</b>	CHAR(32)	The user who has this exemption
<b>secpolicyid</b>	INTEGER	ID of the policy on which the exemption is granted
<b>exemption</b>	CHAR(6)	<p>The exemption given to the user who is identified in the GRANTEE column. The six characters have the following meanings:</p> <ul style="list-style-type: none"> <li>1 = Read array</li> <li>2 = Read set</li> <li>3 = Read tree</li> <li>4 = Write array</li> <li>5 = Write set</li> <li>6 = Write tree</li> </ul> <p>Each character has one of the following values:</p>

Column	Type	Explanation
		E = Exempt D = Write down exemption U = Write up exemption – = No exemption

## SYSSEQUENCES

The **syssequences** system catalog table lists the sequence objects that exist in the database. The **syssequences** table has the following columns.

Column	Type	Explanation
<b>seqid</b>	SERIAL	Code uniquely identifying the sequence object
<b>tabid</b>	INTEGER	Identifying code of the sequence as a table object
<b>start_val</b>	INT8	Starting value of the sequence
<b>inc_val</b>	INT8	Value of the increment between successive values
<b>max_val</b>	INT8	Largest possible value of the sequence
<b>min_val</b>	INT8	Smallest possible value of the sequence
<b>cycle</b>	CHAR(1)	Zero means NOCYCLE, 1 means CYCLE
<b>restart_val</b>	INT8	Starting value of the sequence after ALTER SEQUENCE RESTART was run
<b>cache</b>	INTEGER	Number of preallocated values in sequence cache
<b>order</b>	CHAR(1)	Zero means NOORDER, 1 means ORDER

## SYSSURROGATEAUTH

The **sys surrogateauth** system catalog table stores trusted user and surrogate user information.

The **sys surrogateauth** system catalog table is populated when the GRANT SETSESSIONAUTH statement is run. Users or roles specified in the TO clause are added to **trusteduser** column. Users specified in the ON clause are added to **surrogateuser** column.

For example, consider the following statement:

```
GRANT SETSESSIONAUTH ON bill, john TO mary, peter;
```

Entries in the **sys surrogateauth** table are created as follows:

```
trusteduser  surrogateuser
mary         bill
mary         john
```

peter	bill
peter	john

The **syssurrogateauth** table has the following columns.

**Table 30. SYSSURROGATEAUTH table column descriptions**

Column	Type	Explanation
<b>trusteduser</b>	CHAR(32)	Trusted user name or role.
<b>surrogateuser</b>	CHAR(32)	Surrogate user name.

## SYSSYNONYMS

The **syssynonyms** system catalog table is unused. The **syssynonym** table describes synonyms. The **syssynonyms** system catalog table has the following columns.

**Table 31. SYSSYNONYMS table column descriptions**

Column	Type	Explanation
<b>owner</b>	VARCHAR(32)	Name of the owner of the synonym
<b>synname</b>	VARCHAR(128)	Name of the synonym
<b>created</b>	DATE	Date when the synonym was created
<b>tabid</b>	INTEGER	Identifying code of a table, sequence, or view

## SYSSYNTABLE

The **syssynonym** system catalog table outlines the mapping between each public or private synonym and the database object (table, sequence, or view) that it represents. It contains one row for each entry in the **syssynonyms** table that has a **tabtype** value of **P** or **S**. The **syssynonym** table has the following columns.

Column	Type	Explanation
<b>tabid</b>	INTEGER	Identifying code of the public synonym
<b>servername</b>	VARCHAR(128)	Name of an external database server
<b>dbname</b>	VARCHAR(128)	Name of an external database

Column	Type	Explanation
<b>owner</b>	VARCHAR(32)	Name of the owner of an external object
<b>tablename</b>	VARCHAR(128)	Name of an external table or view
<b>btabid</b>	INTEGER	Identifying code of a base table, sequence, or view

ANSI-compliant databases do not support public synonyms; their **syssytable** tables can describe only synonyms whose **syssytable.tabtype** value is `P`.

If you define a synonym for an object that is in your current database, only the **tabid** and **btabid** columns are used. If you define a synonym for a table that is external to your current database, the **btabid** column is not used, but the **tabid**, **servername**, **dbname**, **owner**, and **tablename** columns are used.

The **tabid** column maps to **systables.tabid**. With the **tabid** information, you can determine additional facts about the synonym from **systables**.

An index on the **tabid** column allows only unique values. The **btabid** column is indexed to allow duplicate values.

## SYSTABAMDATA

The **systabamdata** system catalog table stores the table-specific hashing parameters of tables that were created with a primary access method.

The **systabamdata** table has the following columns.

**Table 32. SYSTABAMDATA table column descriptions**

Column	Type	Explanation
<b>tabid</b>	INTEGER	Identifying code of the table
<b>am_pa ram</b>	LVARCHAR(81 92)	Access method parameter choices
<b>am_space</b>	VARCHAR(128)	Name of the storage space holding the data values

The **am\_param** column stores configuration parameters that determine how a primary access method accesses a given table. Each configuration parameter in the **am\_param** list has the format *keyword=value* or *keyword*.

The **am\_space** column specifies the location of the table. It might be located in a cooked file, a different database, or an sbspace within the database server.

The **tabid** column is the primary key to the **systables** table. This column is indexed and must contain unique values.



## SYSTABAUTH

The **sysabauth** system catalog table describes each set of privileges that are granted on a table, view, sequence, or synonym. It contains one row for each set of table privileges that are granted in the database; the REVOKE statement can modify a row. The **sysabauth** table has the following columns.

**Table 33. SYSTABAUTH table column descriptions**

Column	Type	Explanation
<b>grantor</b>	VARCHAR(32)	Name of the grantor of privilege
<b>grantee</b>	VARCHAR(32)	Name of the grantee of privilege
<b>tabid</b>	INTEGER	Value from <b>sysables.tabid</b> for database object
<b>tabauth</b>	CHAR(9) CHAR(8)	Pattern that specifies privileges on the table, view, synonym, or sequence:  s or S = Select u or U = Update * = Column-level privilege i or I = Insert d or D = Delete x or X = Index a or A = Alter r or R = References n or N = Under privilege

If the **tabauth** column shows a privilege code in uppercase (for example, **S** for Select), this indicates that the user also has the option to grant that privilege to others. Privilege codes listed in lowercase (for example, **s** for select) indicate that the user has the specified privilege, but cannot grant it to others.

A hyphen ( - ) indicates the absence of the privilege corresponding to that position within the **tabauth** pattern.

A **tabauth** value with an asterisk ( \* ) means column-level privileges exist; see also **syscolauth** (page [SYSINDEXES on page 41](#)). (In DB-Access, the **Privileges** option of the **Info** command for a specified table can display the column-level privileges on that table.)

A composite index on **tabid**, **grantor**, and **grantee** allows only unique values. A composite index on **tabid** and **grantee** allows duplicate values.

## SYSTABLES

The **sysables** system catalog table contains a row for each table object (a table, view, synonym, or in HCL OneDB™, a sequence) that has been defined in the database, including the tables and views of the system catalog.

**Table 34. SYSTABLES table column descriptions**

Column	Type	Explanation
<b>tabname</b>	VARCHAR(128)	Name of table, view, synonym, or sequence
<b>owner</b>	CHAR(32)	Owner of table (user <b>informix</b> for system catalog tables and <i>username</i> for database tables)
<b>partnum</b>	INTEGER	Physical storage location code
<b>tabid</b>	SERIAL	System-assigned sequential identifying number
<b>rowsize</b>	SMALLINT	Maximum row size in bytes (< 32,768)
<b>ncols</b>	SMALLINT	Number of columns in the table
<b>nindexes</b>	SMALLINT	Number of indexes on the table
<b>nrows</b>	FLOAT	Number of rows in the table
<b>created</b>	DATE	Date when table was created or last modified
<b>version</b>	INTEGER	Number that changes when table is altered
<b>tabtype</b>	CHAR(1)	Code indicating the type of table object: <ul style="list-style-type: none"> <li>• T = Table</li> <li>• E = External Table</li> <li>• V = View</li> <li>• Q = Sequence</li> <li>• P = Private synonym</li> <li>• S = Public synonym</li> </ul> (Type S is unavailable in an ANSI-compliant database.)
<b>locklevel</b>	CHAR(1)	Lock mode for the table: <ul style="list-style-type: none"> <li>• B = Page and row level</li> <li>• P = Page level</li> <li>• R = Row level</li> </ul>
<b>npused</b>	FLOAT	Number of data pages that have ever been initialized in the tablespace by the database server
<b>fextsize</b>	INTEGER	Size of initial extent (in KB)
<b>nextsize</b>	INTEGER	Size of all subsequent extents (in KB)
<b>flags</b>	SMALLINT	Codes for classifying permanent tables:

Table 34. SYSTABLES table column descriptions (continued)

Column	Type	Explanation
		<p><b>ROWID</b></p> <p>1 - Has rowid column defined</p> <p><b>UNDER</b></p> <p>2 - Table created under a supertable</p> <p><b>VIEWREMOTE</b></p> <p>4 - View is based on a remote table</p> <p><b>CDR</b></p> <p>8 - Has CDRCOLS defined</p> <p><b>RAW</b></p> <p>16 - (HCL OneDB™) RAW table</p> <p><b>EXTERNAL</b></p> <p>32- External table</p> <p><b>AUDIT</b></p> <p>64 - Audit table attribute - FGA</p> <p><b>AQT</b></p> <p>128 - View is an AQT for DWA offloading</p> <p><b>VIRTAQT</b></p> <p>256 - View is a virtual AQT</p>
<b>site</b>	VARCHAR(128)	Reserved for future use
<b>dbname</b>	VARCHAR(128)	Reserved for future use
<b>type_xid</b>	INTEGER	Code from <b>sysxdtypes.extended_id</b> for typed tables, or 0 for untyped tables
<b>am_id</b>	INTEGER	Access method code (key to <b>sysams</b> table)  NULL or 0 indicates built-in storage manager
<b>pagesize</b>	INTEGER	The pagesize, in bytes, of the dbspace (or dbspaces, if the table is fragmented) where the table data resides.
<b>ustlowts</b>	DATETIME YEAR TO FRACTION (5)	When table, row, and page-count statistics were last recorded
<b>secpolicyid</b>	INTEGER	ID of the SECURITY policy attached to the table. NULL for non-protected tables

**Table 34. SYSTABLES table column descriptions (continued)**

Column	Type	Explanation
<b>protgranularity</b>	CHAR(1)	LBAC granularity level: <ul style="list-style-type: none"> <li>• R: Row level granularity</li> <li>• C: Column level granularity</li> <li>• B: Both column and row granularity</li> <li>• Blank for non-protected tables</li> </ul>
<b>statlevel</b>	CHAR(1)	Statistics level <ul style="list-style-type: none"> <li>• T = table</li> <li>• F = fragment</li> <li>• A = automatic</li> </ul>
<b>statchange</b>	SMALLINT	For internal use only

Each table, view, sequence, and synonym recorded in the **systables** table is assigned a **tabid**, which is a system-assigned SERIAL value that uniquely identifies the object. The first 99 **tabid** values are reserved for the system catalog. The **tabid** of the first user-defined table object in a database is always 100.

The **tabid** column is indexed and contains only unique values. A composite index on the **tablename** and **owner** columns also requires unique values.

The version column contains an encoded number that is stored in **systables** when a new table is created. Portions of this value are incremented when data-definition statements, such as ALTER INDEX, ALTER TABLE, DROP INDEX, and CREATE INDEX, are performed on the table.

In the **flags** column, ST\_RAW represents a nonlogging permanent table in a database that supports transaction logging.

The setting of the SQL\_LOGICAL\_CHAR parameter is encoded into the **systables.flags** column value in the row that describes the ' **VERSION**' table. Note the leading blank space in the identifier of this system-generated table.

To determine whether the database enables the SQL\_LOGICAL\_CHAR configuration parameter, which can apply logical character semantics to the declarations of character columns, you can execute the following query:

```
SELECT flags INTO $value FROM 'informix'.systables WHERE tablename = ' VERSION';
```

Because the SQL\_LOGICAL\_CHAR setting is encoded in the two least significant bits of the " **VERSION.flags**" value, you can calculate its setting from the returned **flags** value by the following formula:

```
SQL_LOGICAL_CHAR = (value & 0x03) + 1
```

Here **&** is the bitwise **AND** operator. Any SQL\_LOGICAL\_CHAR setting greater than 1 indicates that SQL\_LOGICAL\_CHAR was enabled when the database was created, and that explicit or default maximum size specifications of character columns are multiplied by that setting.

When a prepared statement that references a database table is executed, the version value is checked to make sure that nothing has changed since the statement was prepared. If the version value has been changed by DDL operations that modified the table schema while automatic recompilation was disabled by the IFX\_AUTO\_REPREPARE setting of the SET ENVIRONMENT statement, the prepared statement is not executed, and you must prepare the statement again.

The **npused** column does not reflect the number of pages used for BYTE or TEXT data, nor the number of pages that are freed in DELETE or TRUNCATE operations.

The **nrows** column and the **npused** columns might not accurately reflect the number of rows and the number of data pages used by an external table unless the NUMROWS clause was specified when the external table was created. See the *HCL OneDB™ Administrator's Guide* for more information.

The **systables** table has two rows that store information about the database locale: GL\_COLLATE with a **tabid** of 90 and GL\_CTYPE with a **tabid** of 91. To view these rows, enter the following SELECT statement:

```
SELECT * FROM systables WHERE tabid=90 OR tabid=91;
```

## SYSTRACECLASSES

The **systraceclasses** system catalog table contains the names and identifiers of trace classes. The **systraceclasses** table has the following columns.

**Table 35. SYSTRACECLASSES table column descriptions**

Column	Type	Explanation
<b>name</b>	CHAR(18)	Name of the class of trace messages
<b>classid</b>	SERIAL	Identifying code of the trace class

A *trace class* is a category of trace messages that you can use in the development and testing of new DataBlade® modules and user-defined routines. Developers use the tracing facility by calling the appropriate DataBlade® API routines within their code.

To create a new trace class, insert a row directly into the **systraceclasses** table. By default, all users can view this table, but only users with the DBA privilege can modify it.

The database cannot support tracing unless the MITRACE\_OFF configuration parameter is undefined.

A unique index on the **name** column requires each trace class to have a unique name. The database server assigns to each class a unique sequential code. The index on this **classid** column also allows only unique values.

## SYSTRACEMSGS

The **systracemsgs** system catalog table stores internationalized trace messages that you can use in debugging user-defined routines.

The **systracemsgs** table has the following columns.

**Table 36. SYSTRACEMSGS table column descriptions**

Column	Type	Explanation
<b>name</b>	VARCHAR(128)	Name of the message
<b>msgid</b>	SERIAL	Identifying code of the message template
<b>locale</b>	CHAR(36)	Locale with which this version of the message is associated (for example, <b>en_us.8859-1</b> )
<b>seqno</b>	SMALLINT	Reserved for future use
<b>message</b>	VARCHAR(255)	The message text

DataBlade® module developers create a trace message by inserting a row directly into the **systracemsgs** table. After a message is created, the development team can specify it either by name or by **msgid** code, using trace statements that the DataBlade® API provides.

To create a trace message, you must specify its name, locale, and text. By default, all users can view the **systracemsgs** table, but only users with the DBA privilege can modify it.

The database cannot support tracing unless the MITRACE\_OFF configuration parameter is undefined.

A unique composite index is defined on the **name** and **locale** columns. Another unique index is defined on the **msgid** column.

## SYSTRIGBODY

The **systrigbody** system catalog table contains the ASCII text of the trigger definition and the linearized code for the trigger. *Linearized code* is binary data and code that is represented in ASCII format.



**Important:** The database server uses the linearized code that is stored in **systrigbody**. You must not alter the content of rows that contain linearized code.

The **systrigbody** table has the following columns.

**Table 37. SYSTRIGBODY table column descriptions**

Column	Type	Explanation
<b>trigid</b>	INTEGER	Identifying code of the trigger
<b>datakey</b>	CHAR(1)	Code specifying the type of data:  A = ASCII text for the body, triggered actions B = Linearized code for the body D = English text for the header, trigger definition

**Table 37. SYSTRIGBODY table column descriptions (continued)**

Column	Type	Explanation
		H = Linearized code for the header S = Linearized code for the symbol table
<b>seqno</b>	INTEGER	Page number of this data segment
<b>data</b>	CHAR(256)	English text or linearized code

A composite index on the **trigid**, **datakey**, and **seqno** columns allows only unique values.

## SYSTRIGGERS

The **systriggers** system catalog table contains information about the SQL triggers in the database. This information includes the triggering event and the correlated reference specification for the trigger. The **systriggers** table has the following columns.

**Table 38. SYSTRIGGERS table column descriptions**

Column	Type	Explanation
<b>trigid</b>	SERIAL	Identifying code of the trigger
<b>trigname</b>	VARCHAR(128)	Name of the trigger
<b>owner</b>	VARCHAR(32)	Name of the owner of the trigger
<b>tabid</b>	INTEGER	Identifying code of the triggering table
<b>event</b>	CHAR(1)	Code for the type of triggering event:  D = Delete trigger I = Insert trigger U = Update trigger S = Select trigger d = INSTEAD OF Delete trigger i = INSTEAD OF Insert trigger u = INSTEAD OF Update trigger
<b>old</b>	VARCHAR(128)	Name of value before update
<b>new</b>	VARCHAR(128)	Name of value after update
<b>mode</b>	CHAR(1)	Reserved for future use
<b>collation</b>	CHAR(32)	Collating order at the time when the routine was created

A composite index on the **trigname** and **owner** columns allows only unique values. An index on the **trigid** column also requires unique values. An index on the **tabid** column allows duplicate values.

## SYSUSERS

The **sysusers** system catalog table lists the authorization identifier of every individual user, or public for the PUBLIC group, who holds database-level access privileges. This table also lists the name of every role that holds access privileges on any object in the database.

This system catalog table has the following columns:

**Table 39. SYSUSERS table column descriptions**

Column	Type	Explanation
<b>username</b>	VARCHAR(32)	Name of the database user or role. An index on <b>username</b> allows only unique values. The <b>username</b> value can be the login name of a user or the name of a role.
<b>usertype</b>	CHAR(1)	Code specifying the highest database-level privilege held by <b>username</b> , where <b>username</b> is an individual user or the PUBLIC group, or a role name. The valid codes are:  D = DBA (all privileges) R = Resource (create UDRs, UDTs, permanent tables, and indexes) C = Connect (work with existing tables) G = Role U = Default role. When a user is assigned a default role, an implicit connection to the database is granted to the user. This is the role the user has before being granted a C, D, or R role.
<b>priority</b>	SMALLINT	Reserved for future use.
<b>password</b>	CHAR(16)	Reserved for future use.
<b>defaultrole</b>	VARCHAR(32)	Name of the default role.

## SYSVIEWS

The **sysviews** system catalog table describes each view in the database. Because it stores the SELECT statement that created the view, **sysviews** can contain multiple rows for each view. It has the following columns.



Column	Type	Explanation
<b>tabid</b>	INTEGER	Identifying code of the view
<b>seqno</b>	SMALLINT	Line number of the SELECT statement
<b>viewtext</b>	CHAR(256)	Actual SELECT statement used to create the view

A composite index on **tabid** and **seqno** allows only unique values.

## SYSVIOLATIONS

The **sysviolations** system catalog table stores information about constraint violations for base tables.

This table is updated when the DELETE, INSERT, MERGE, or UPDATE statement detects a violation of an enabled constraint or unique index in a database table for which the START VIOLATIONS TABLE statement of SQL has created an associated violations table (and for HCL OneDB™, a diagnostics table). For each base table that has an active violations table, the **sysviolations** table has a corresponding row, with the following columns.

Column	Type	Explanation
<b>targettid</b>	INTEGER	Identifying code of the <i>target table</i> (the base table on which the violations table and the diagnostic table are defined)
<b>viotid</b>	INTEGER	Identifying code of the violations table
<b>diatid</b>	INTEGER	Identifying code of the diagnostics table
<b>maxrows</b>	INTEGER	Maximum number of rows that can be inserted into the diagnostics table by a single insert, update, or delete operation on a target table that has a filtering mode object defined on it.

The **maxrows** column also signifies the maximum number of rows that can be inserted in the diagnostics table during a single operation that enables a disabled object or that sets a disabled object to filtering mode (provided that a diagnostics table exists for the target table). If no maximum is specified for the diagnostics or violations table, then **maxrows** contains a NULL value.

The primary key of this table is the **targettid** column. An additional unique index is also defined on the **viotid** column.

HCL OneDB™ also has a unique index on the **diatid** column.

## SYSXADATASOURCES

The **sysxdatasources** system catalog table stores XA data sources.

The **sysxdatasources** table has the following columns.

Column	Type	Explanation
<b>xa_datasrc_owner</b>	CHAR(32)	The user ID of the XA data source owner

Column	Type	Explanation
<b>xa_datasrc_name</b>	VARCHAR(128)	The name of the XA data source
<b>xa_datasrc_rmid</b>	SERIAL	Unique RMID of the XA data source
<b>xa_source_typeid</b>	INTEGER	XA data source type ID

## SYSXASOURCETYPES

The **sysxasourcetypes** system catalog table stores XA data source types.

The **sysxasourcetypes** table has the following columns.

Column	Type	Explanation
<b>xa_source_typeid</b>	SERIAL	A unique identifier for the source type
<b>xa_source_owner</b>	CHAR(32)	The user ID of the owner
<b>xa_source_name</b>	VARCHAR(128)	The name of the source type
<b>xa_flags</b>	INTEGER	
<b>xa_version</b>	INTEGER	
<b>xa_open</b>	INTEGER	UDR ID of xa_open_entry
<b>xa_close</b>	INTEGER	UDR ID of xa_close_entry
<b>xa_start</b>	INTEGER	UDR ID of xa_start entry
<b>xa_end</b>	INTEGER	UDR ID of xa_end_entry
<b>xa_rollback</b>	INTEGER	UDR ID of xa_rollback_entry
<b>xa_prepare</b>	INTEGER	UDR ID of xa_prepare_entry
<b>xa_commit</b>	INTEGER	UDR ID of xa_commit_entry
<b>xa_recover</b>	INTEGER	UDR ID of xa_recover_entry
<b>xa_forget</b>	INTEGER	UDR ID of xa_forget_entry
<b>xa_complete</b>	INTEGER	UDR ID of xa_complete_entry

## SYSXTDDESC

The **sysxtddesc** system catalog table provides a text description of each user-defined data type (UDT) defined in the database. The **sysxtddesc** table has the following columns.

Column	Type	Explanation
<b>extended_id</b>	INTEGER	Code uniquely identifying the extended data types

Column	Type	Explanation
<b>seqno</b>	SMALLINT	Value to order and identify one line of the description of the UDT  A new line is created only if the remaining text string is larger than 255 bytes.
<b>description</b>	CHAR(256)	Textual description of the extended data type

A composite index on **extended\_id** and **seqno** allows duplicate values.

## SYSXTDTYPEAUTH

The **sysxtdtypeauth** system catalog table identifies the privileges on each UDT (user-defined data type).

The **sysxtdtypeauth** table contains one row for each set of privileges granted and has the following columns:

Column	Type	Explanation
<b>grantor</b>	VARCHAR(32)	Name of grantor of privilege
<b>grantee</b>	VARCHAR(32)	Name of grantee of privilege
<b>type</b>	INTEGER	Code identifying the UDT
<b>auth</b>	CHAR(2)	Code identifying privileges on the UDT:  n or N = Under privilege u or U = Usage privilege

If the privilege code in the **auth** column is upper case (for example, 'U' for usage), a user who has this privilege can also grant it to others. If the code is in lower case, a user who has the privilege cannot grant it to others.

A composite index on **type**, **grantor**, and **grantee** allows only unique values. A composite index on the **type** and **grantee** columns allows duplicate values.

## SYSXTDTYPES

The **sysxtdtypes** system catalog table has an entry for each UDT (user-defined data type), including opaque and distinct data types and complex data types (named ROW types, unnamed ROW types, and COLLECTION types), that is defined in the database.

The **sysxtdtypes** table has the following columns.

**Table 40. SYSXTDTYPES table column descriptions**

Column	Type	Explanation
<b>extended_id</b>	SERIAL	Unique identifying code for extended data type

Table 40. SYSXTDTYPES table column descriptions

(continued)

Column	Type	Explanation
<b>domain</b>	CHAR(1)	Code for the domain of the UDT
<b>mode</b>	CHAR(1)	Code classifying the UDT: <ul style="list-style-type: none"> <li>• B = Base (opaque) type</li> <li>• C = Collection type or unnamed ROW type</li> <li>• D = Distinct type</li> <li>• R = Named ROW type</li> <li>• S = Reserved for internal use</li> <li>• T = Reserved for internal use</li> <li>• ' ' (blank) = Built-in type</li> </ul>
<b>owner</b>	VARCHAR(32)	Name of the owner of the UDT
<b>name</b>	VARCHAR(128)	Name of the UDT
<b>type</b>	SMALLINT	Code classifying the UDT
<b>source</b>	INTEGER	The <b>sysxtdtypes</b> reference (for distinct types only)  Zero (0) indicates that a distinct UDT was created from a built-in data type.
<b>maxlen</b>	INTEGER	The maximum length for variable-length data types  Zero indicates a fixed-length UDT.
<b>length</b>	INTEGER	The length in bytes for fixed-length data types  Zero indicates a variable-length UDT.
<b>byvalue</b>	CHAR(1)	'T' = UDT is passed by value  'F' = UDT is not passed by value
<b>canonhash</b>	CHAR(1)	'T' = UDT is hashable by default hash function  'F' = UDT is not hashable by default function
<b>align</b>	SMALLINT	Alignment ( = 1, 2, 4, or 8) for this UDT
<b>locator</b>	INTEGER	Locator key for unnamed ROW type

Each extended data type is characterized by a unique identifier, called an extended identifier (**extended\_id**), a data type identifier (**type**), and the length and description of the data type.

For distinct types created from built-in data types, the **type** column codes correspond to the value of the **syscolumns.coltype** column (indicating the source type) as listed on page [SYSCOLUMNS on page 24](#), but incremented by the hexadecimal value 0x0000800. The file `$ONEDB_HOME/incl/esql/sqltypes.h` contains information about **sysxdtypes.type** and **syscolumns.coltype** codes.

An index on the **extended\_id** column allows only unique values. An index on the **locator** column allows duplicate values, as does a composite index on the **name** and **owner** columns. A composite index on the **type** and **source** columns also allows duplicate values.

## Information Schema

The Information Schema consists of read-only views that provide information about all the tables, views, and columns in the current database server to which you have access. These views also provide information about SQL dialects (such as HCL OneDB™, Oracle, or Sybase) and SQL standards. Note that unlike a system catalog, whose tables describes an individual database, these views describe the HCL OneDB™ instance, rather than a single database.

This version of the Information Schema views is an X/Open CAE standard. These standards are provided so that applications developed on other database systems can obtain HCL OneDB™ system catalog information without accessing the HCL OneDB™ system catalog tables directly.



**Important:** Because the X/Open CAE standard for Information Schema views differs from ANSI-compliant Information Schema views, it is recommended that you do not install the X/Open CAE Information Schema views on ANSI-compliant databases.

The following Information Schema views are available:

- **tables**
- **columns**
- **sql\_languages**
- **server\_info**

Sections that follow contain information about how to generate and access Information Schema views and information about their structure.

## Generating the Information Schema Views

### About this task

The Information Schema views are generated automatically when you, as DBA, run the following DB-Access command:

```
dbaccess database-name $ONEDB_HOME/etc/xpg4_is.sql
```

The views display data from the system catalog tables. If tables, views, or routines exist with any of the same names as the Information Schema views, you must either rename those database objects or rename the views in the script before you can install the views. You can drop the views with the DROP VIEW statement on each view. To re-create the views, rerun the script.



**Important:** In addition to the columns specified for each Information Schema view, individual vendors might include additional columns or change the order of the columns. It is recommended that applications not use the forms SELECT \* or SELECT table-name\* to access an Information Schema view.

## Accessing the Information Schema Views

All Information Schema views have the Select privilege granted to PUBLIC WITH GRANT OPTION so that all users can query the views. Because no other privileges are granted on the Information Schema views, they cannot be updated.

You can query the Information Schema views as you would query any other table or view in the database.

## Structure of the Information Schema Views

The following Information Schema views are described in this section:

- **tables**
- **columns**
- **sql\_languages**
- **server\_info**

In order to accept long identifier names, most of the columns in the views are defined as VARCHAR data types with large maximum sizes.

### The tables Information Schema View

The **tables** Information Schema view contains one row for each table to which you have access. It contains the following columns.

Column	Data Type	Explanation
<b>table_schema</b>	VARCHAR(32)	Name of owner of table
<b>table_name</b>	VARCHAR(128)	Name of table or view
<b>table_type</b>	VARCHAR(128)	BASE TABLE for table or VIEW for view
<b>remarks</b>	VARCHAR(255)	Reserved for future use

The visible rows in the **tables** view depend on your privileges. For example, if you have one or more privileges on a table (such as Insert, Delete, Select, References, Alter, Index, or Update on one or more columns), or if privileges are granted to PUBLIC, you see the row that describes that table.

## The columns Information Schema View

The **columns** Information Schema view contains one row for each accessible column. It contains the following columns.

**Table 41. Description of the columns Information Schema View**

Column	Data Type	Explanation
<b>table_schema</b>	VARCHAR(128)	Name of owner of table
<b>table_name</b>	VARCHAR(128)	Name of table or view
<b>column_name</b>	VARCHAR(128)	Name of the column in the table or view
<b>ordinal_position</b>	INTEGER	Position of the column within its table  The <b>ordinal_position</b> value is a sequential number that starts at 1 for the first column. This is the HCL OneDB™ extension to XPG4.
<b>data_type</b>	VARCHAR(254)	Name of the data type of the column, such as CHARACTER or DECIMAL
<b>char_max_length</b>	INTEGER	Maximum length (in bytes) for character data types; NULL otherwise
<b>numeric_precision</b>	INTEGER	Uses one of the following values: <ul style="list-style-type: none"> <li>• Total number of digits for exact numeric data types (DECIMAL, INTEGER, MONEY, SMALLINT)</li> <li>• Number of digits of mantissa precision (machine-dependent) for approximate data types (FLOAT, SMALLFLOAT)</li> <li>• NULL for all other data types.</li> </ul>
<b>numeric_prec_radix</b>	INTEGER	Uses one of the following values: <ul style="list-style-type: none"> <li>• 2 = Approximate data types (FLOAT and SMALLFLOAT)</li> <li>• 10 = Exact numeric data types (DECIMAL, INTEGER, MONEY, and SMALLINT)</li> <li>• NULL for all other data types</li> </ul>
<b>numeric_scale</b>	INTEGER	Number of significant digits to the right of the decimal point for DECIMAL and MONEY data types

**Table 41. Description of the columns Information Schema View (continued)**

Column	Data Type	Explanation
		0 for INTEGER and SMALLINT types NULL for all other data types
<b>datetime_precision</b>	INTEGER	Number of digits in the fractional part of the seconds for DATE and DATETIME columns; NULL otherwise  This column is the HCL OneDB™ extension to XPG4.
<b>is_nullable</b>	VARCHAR(3)	Indicates whether a column allows NULL values; either YES or NO
<b>remarks</b>	VARCHAR(254)	Reserved for future use

## The sql\_languages Information Schema View

The **sql\_languages** Information Schema view contains a row for each instance of conformance to standards that the current database server supports. The **sql\_languages** view contains the following columns.

Column	Data Type	Explanation
<b>source</b>	VARCHAR(254)	Organization defining this SQL version
<b>source_year</b>	VARCHAR(254)	Year the source document was approved
<b>conformance</b>	VARCHAR(254)	Standard to which the server conforms
<b>integrity</b>	VARCHAR(254)	Indication of whether this is an integrity enhancement feature; either <b>YES</b> or <b>NO</b>
<b>implementation</b>	VARCHAR(254)	Identification of the SQL product of the vendor
<b>binding_style</b>	VARCHAR(254)	Direct, module, or other binding style
<b>programming_lang</b>	VARCHAR(254)	Host language for which binding style is adapted

The **sql\_languages** view is completely visible to all users.

## The server\_info Information Schema View

The **server\_info** Information Schema view describes the database server to which the application is currently connected. It contains two columns.

Column	Data Type	Explanation
<b>server_attribute</b>	VARCHAR(254)	An attribute of the database server



Column	Data Type	Explanation
<b>attribute_value</b>	VARCHAR(254)	Value of the <b>server_attribute</b> as it applies to the current database server

Each row in this view provides information about one attribute. X/Open-compliant databases must provide applications with certain required information about the database server.

The **server\_info** view includes the following **server\_attribute** information.

server_attribute	Explanation
<b>identifier_length</b>	Maximum number of bytes for a user-defined identifier
<b>row_length</b>	Maximum number of bytes in a row
<b>userid_length</b>	Maximum number of bytes in a user name
<b>txn_isolation</b>	Initial transaction isolation level for the database server:  Read Uncommitted ( = Default isolation level for databases with no transaction logging; also called Dirty Read)  Read Committed ( = Default isolation level for databases that are not ANSI-compliant, but that support explicit transaction logging)  Serializable ( = Default isolation level for ANSI-compliant databases; also called Repeatable Read)
<b>collation_seq</b>	Assumed ordering of the character set for the database server The following values are possible: ISO 8859-1 EBCDIC  The default HCL OneDB™ representation shows ISO 8859-1.

The **server\_info** view is completely visible to all users.

## Data types

Every column in a table in a database is assigned a data type. The data type precisely defines the kinds of values that you can store in that column.

These topics describe built-in and extended data types, casting between two data types, and operator precedence.

## Summary of data types

HCL OneDB™ supports the most common set of built-in data types. Additionally, an extended set of data types are supported on the database server.

You can use both *built-in* data types (which are system-defined) and *extended* data types (which you can define) in the following ways:

- Use them to create columns within database tables.
- Declare them as arguments and as returned types of routines.
- Use them as base types from which to create DISTINCT data types.
- Cast them to other data types.
- Declare and access host variables of these types in SPL and ESQL/C.

You assign data types to columns with the CREATE TABLE statement and change them with the ALTER TABLE statement. When you change an existing column data type, all data is converted to the new data type, if possible.

For information about the ALTER TABLE and CREATE TABLE statements, on SQL statements that create specific data types, that create and drop casts, and on other data type topics, see the *HCL OneDB™ Guide to SQL: Syntax*.

For information about how to create and use complex data types supported by HCL OneDB™, see the *HCL OneDB™ Database Design and Implementation Guide*. For information about how to create user-defined data types, see *HCL OneDB™ User-Defined Routines and Data Types Developer's Guide*.

Some data types can be used in distributed SQL operations, while others can be used only in SQL operations within the same database.

### Built-in data types supported in local and distributed SQL operations

The following table lists all of the built-in SQL data types that HCL OneDB™ supports. These built-in SQL data types are valid in all HCL OneDB™ SQL transactions, including data-manipulation language (DML) operations of these types:

- Operations on objects in the local database
- Cross-database operations on objects in databases of the local server instance
- Cross-server operations on objects in databases of two or more database server instances

**Table 42. Data types supported in all operations**

Data type	Explanation
<a href="#">BIGINT data type on page 87</a>	Stores 8-byte integer values from $-(2^{63}-1)$ to $2^{63}-1$
<a href="#">BIGSERIAL data type on page 87</a>	Stores sequential, 8-byte integers from 1 to $2^{63}-1$
<a href="#">BSON and JSON built-in opaque data types on page</a>	The BSON data type is the binary representation of a JSON data type format for serializing JSON documents. The JSON data type is a plain text format for entering and displaying structured data.
<a href="#">BYTE data type on page 89</a>	Stores any kind of binary data, up to $2^{31}$ bytes in length
<a href="#">CHAR(n) data type on page 90</a>	Stores character strings; collation is in code-set order

**Table 42. Data types supported in all operations (continued)**

Data type	Explanation
<a href="#">CHARACTER(n) data type on page 91</a>	Is a synonym for CHAR
<a href="#">CHARACTER VARYING(m,r) data type on page 91</a>	Stores character strings of varying length (ANSI-compliant); collation is in code-set order
<a href="#">DATE data type on page 93</a>	Stores calendar dates
<a href="#">DATETIME data type on page 93</a>	Stores calendar date combined with time of day
<a href="#">DEC data type on page 96</a>	Is a synonym for DECIMAL
<a href="#">DECIMAL on page 96</a>	Stores floating-point numbers with definable precision; if database is ANSI-compliant, the scale is zero
<a href="#">DECIMAL (p,s) Fixed Point on page 97</a>	Stores fixed-point numbers of defined scale and precision
<a href="#">DOUBLE PRECISION data types on page 99</a>	Synonym for FLOAT
<a href="#">FLOAT(n) on page 99</a>	Stores double-precision floating-point numbers corresponding to the <b>double</b> data type in C
<a href="#">INT data type on page 100</a>	Is a synonym for INTEGER
<a href="#">INT8 on page 100</a>	Stores 8-byte integer values from $-(2^{63}-1)$ to $2^{63}-1$
<a href="#">INTEGER data type on page 100</a>	Stores whole numbers from -2,147,483,647 to +2,147,483,647
<a href="#">INTERVAL data type on page 100</a>	Stores a span of time (or level of effort) in units of <i>years</i> and <i>months</i> .
<a href="#">INTERVAL data type on page 100</a>	Stores a span of time in a contiguous set of units of <i>days</i> , <i>hours</i> , <i>minutes</i> , <i>seconds</i> , and <i>fractions of a second</i>
<a href="#">MONEY(p,s) data type on page 105</a>	Stores currency amounts
<a href="#">NCHAR(n) data type on page 106</a>	Same as CHAR, but can support localized collation
<a href="#">NUMERIC(p,s) data type on page 107</a>	Synonym for DECIMAL(p,s)
<a href="#">NVARCHAR(m,r) data type on page 107</a>	Same as VARCHAR, but can support localized collation
<a href="#">REAL data type on page 108</a>	Is a synonym for SMALLFLOAT
<a href="#">SERIAL(n) data type on page 111</a>	Stores sequential integers ( $> 0$ ) in positive range of INT

**Table 42. Data types supported in all operations (continued)**

Data type	Explanation
<a href="#">SERIAL8(n) data type on page 112</a>	Stores sequential integers (> 0) in positive range of INT8
<a href="#">SMALLFLOAT on page 114</a>	Stores single-precision floating-point numbers corresponding to the <b>float</b> data type of the C language
<a href="#">SMALLINT data type on page 114</a>	Stores whole numbers from -32,767 to +32,767
<a href="#">TEXT data type on page 115</a>	Stores any kind of text data, up to 2 <sup>31</sup> bytes in length
<a href="#">VARCHAR(m,r) data type on page 116</a>	Stores character strings of varying length (up to 255 bytes); collation is in code-set order

In cross-server MERGE operations, the source table (but not the target table) can be in a database of a remote HCL OneDB™ server.

For the character data types (CHAR, CHAR VARYING, LVARCHAR, NCHAR, NVARCHAR, and VARCHAR), a data string can include letters, digits, punctuation, whitespace, diacritical marks, ligatures, and other printable symbols from the code set of the database locale. For **UTF-8** and for code sets of some East Asian locales, multibyte characters are supported within data strings.

### Built-in data types supported only in local database SQL operations

The following table lists the data types that HCL OneDB™ supports only for use in SQL operations in a local database.

**Table 43. Data types supported in a local database**

Data type	Explanation
<a href="#">BLOB data type on page 88</a>	Stores binary data in random-access chunks
<a href="#">The binary18 data type on page</a>	Stores 18 byte binary-encoded strings
<a href="#">The binaryvar data type on page</a>	Stores binary-encoded strings with a maximum length of 255 bytes
<a href="#">BOOLEAN data type on page 89</a>	Stores Boolean values true and false
<a href="#">CLOB data type on page 92</a>	Stores text data in random-access chunks
<a href="#">DISTINCT data types on page 98</a>	Stores data in a user-defined type that has the same format as a source type on which it is based, but its casts and functions can differ from those on the source type
<a href="#">Calendar data type on page</a>	Stores a calendar for a TimeSeries data type

**Table 43. Data types supported in a local database (continued)**

Data type	Explanation
<a href="#">CalendarPattern data type on page</a>	Stores the structure of the calendar pattern for a Calendar data type
<a href="#">IDSSECURITYLABEL data type on page 99</a>	Stores LBAC security label objects.
<a href="#">LIST(e) data type on page 103</a>	Stores a sequentially ordered collection of elements, all of the same data type, e; allows duplicate values
<a href="#">The lld_locator data type on page</a>	Stores a large object identifier
<a href="#">The lld_job data type on page</a>	Stores the location of a smart large object and specifies whether the object contains binary or character data
<a href="#">LVARCHAR(m) data type on page 104</a>	Stores variable-length strings of up to 32,739 bytes
<a href="#">MULTISET(e) data type on page 106</a>	Stores a non-ordered collection of values, with elements all of the same data type, e; allows duplicate values.
<a href="#">The node data type for querying hierarchical data on page</a>	Stores a combination of integers and decimal points that represents hierarchical relationships, of variable length up to 256 characters
<a href="#">OPAQUE data types on page 107</a>	Stores a user-defined data type whose internal structure is inaccessible to the database server
<a href="#">ROW data type, Named on page 108</a>	Stores a named ROW type
<a href="#">ROW data type, Unnamed on page 109</a>	Stores an unnamed ROW type
<a href="#">SET(e) data type on page 113</a>	Stores a non-ordered collection of elements, all of the same data type, e; does not allow duplicate values
<a href="#">ST_LineString data type on page</a>	Stores a one-dimensional object as a sequence of points defining a linear interpolated path
<a href="#">ST_MultiLineString data type on page</a>	Stores a collection of ST_LineString data types
<a href="#">ST_MultiPoint data type on page</a>	Stores a collection of ST_Point data types

**Table 43. Data types supported in a local database (continued)**

Data type	Explanation
<a href="#">ST_MultiPolygon data type on page</a>	Stores a collection of ST_Polygon data types
<a href="#">ST_Point data type on page</a>	Stores a zero-dimensional geometry that occupies a single location in coordinate space
<a href="#">ST_Polygon data type on page</a>	Stores a two-dimensional surface stored as a sequence of points defining its exterior bounding ring and 0 or more interior rings
<a href="#">TimeSeries data type on page</a>	Stores a collection of row subtypes

These extended data types of HCL OneDB™ are individually described in other topics. These data types are valid in local operations on databases where the data types are defined.

### Extended data types in cross-database distributed SQL transactions

Distributed operations on other databases of the same HCL OneDB™ instance can access BOOLEAN, BLOB, CLOB, and LVARCHAR data types, which are implemented as built-in opaque types. Such operations can also access DISTINCT types whose base types are built-in types, and user-defined types (UDTs), if the UDTs and DISTINCT types are explicitly cast to built-in types, and if all of the UDTs, casts, and DISTINCT types are defined in all the participating databases.

You cannot, however, reference the following extended data types in cross-database transactions that access multiple databases of the local HCL OneDB™ instance:

- UDTs that are not cast to built-in data types
- DISTINCT types that are not cast to built-in data types
- Collection data types
- Named or unnamed ROW data types

### Extended data types in cross-server distributed SQL transactions

Distributed SQL transactions and function calls that access databases of other HCL OneDB™ instances cannot return values of complex or smart large object data types, nor of most distinct or built-in opaque data types. Among the extended data types, only the following can be accessed in cross-server SQL operations:

- Any non-opaque built-in data type
- BOOLEAN
- DISTINCT of non-opaque built-in types
- DISTINCT of BOOLEAN
- DISTINCT of LVARCHAR
- DISTINCT of any of the DISTINCT types listed above

- IDSSECURITYLABEL
- LVARCHAR

A cross-server distributed SQL transaction can support DISTINCT data types only if they are cast explicitly to built-in types, and all of the DISTINCT types, their data type hierarchies, and their casts are defined exactly the same way in each database that participates in the distributed operation. For queries or other DML operations in cross-server UDRs that use the data types in the preceding list as parameters or as returned data types, the UDR must also have the same definition in every participating database.

The built-in DISTINCT data type IDSSECURITYLABEL, which stores security label objects, can be accessed in cross-server and cross-database operations on protected data by users who hold sufficient security credentials. Like local operations on protected data, distributed queries that access remote tables protected by a security policy can return only the qualifying rows that IDSLBACRULES allow, after the database server has compared the security label that secures the data with the security credentials of the user who issues the query.

## Description of Data Types

This section describes the data types that HCL OneDB™ supports.

### BIGINT data type

The BIGINT data type stores integers from  $-(2^{63}-1)$  to  $2^{63}-1$ , which is  $-9,223,372,036,854,775,807$  to  $9,223,372,036,854,775,807$ , in eight bytes.

This data type has storage advantages over INT8 and advantages for some arithmetic operations and sort comparisons over INT8 and DECIMAL data types.

### BIGSERIAL data type

The BIGSERIAL data type stores a sequential integer, of the BIGINT data type, that is assigned automatically by the database server when a new row is inserted. The behavior of the BIGSERIAL data type is similar to the SERIAL data type, but with a larger range.

The default BIGSERIAL starting number is 1, but you can assign an initial value,  $n$ , when you create or alter the table. The value of  $n$  must be a positive integer in the range of 1 to 9,223,372,036,854,775,807. If you insert the value zero (0) in a BIGSERIAL column, the value that is used is the maximum positive value that already exists in the BIGSERIAL column + 1. If you insert any value that is not zero, that value will be inserted as it is.

A table can have no more than one SERIAL column, but it can have a SERIAL column and either a SERIAL8 column or a BIGSERIAL column.

For information about:

- The SERIAL data type, see [SERIAL\(n\) data type on page 111](#)
- Using the SERIAL8 data type with the INT8 or BIGINT data type, see [Using SERIAL8 and BIGSERIAL with INT8 or BIGINT on page 88](#)

## Using SERIAL8 and BIGSERIAL with INT8 or BIGINT

All the arithmetic operators that are valid for INT8 and BIGINT (such as +, -, \*, and /) and all the SQL functions that are valid for INT8 and BIGINT (such as ABS, MOD, POW, and so on) are also valid for SERIAL8 and BIGSERIAL values.

Data conversion rules that apply to INT8 and BIGINT also apply to SERIAL8 and BIGSERIAL, but with a NOT NULL constraint on SERIAL8 or BIGSERIAL.

The value of a SERIAL8 or BIGSERIAL column of one table can be stored in INT8 or BIGINT columns of another table. In the second table, however, the INT8 or BIGINT values are not subject to the constraints on the original SERIAL8 or BIGSERIAL column.

## BLOB data type

The BLOB data type stores any kind of binary data in random-access chunks, called sbspaces. Binary data typically consists of saved spreadsheets, program-load modules, digitized voice patterns, and so on. The database server performs no interpretation of the contents of a BLOB column.

A BLOB column can be up to 4 terabytes ( $4 \times 2^{40}$  bytes) in length, though your system resources might impose a lower practical limit. The minimum amount of disk space allocated for smart large object data types is 512 bytes.

The term *smart large object* refers to BLOB and CLOB data types. Use CLOB data types (see page [CLOB data type on page 92](#)) for random access to text data. For general information about BLOB and CLOB data types, see [Smart large objects on page 122](#).

You can use these SQL functions to perform operations on a BLOB column:

- **FILETOBLOB** copies a file into a BLOB column.
- **LOTOFILE** copies a BLOB (or CLOB) value into an operating-system file.
- **LOCOPY** copies an existing smart large object to a new smart large object.

For more information about these SQL functions, see the *HCL OneDB™ Guide to SQL: Syntax*.

Within SQL, you are limited to the equality (=) comparison operation and the encryption and decryption functions for BLOB data. (The encryption and decryption functions are described in the *HCL OneDB™ Guide to SQL: Syntax*.) To perform additional operations, you must use one of the application programming interfaces (APIs) from within your client application.

You can insert data into BLOB columns in the following ways:

- With the dbload utility
- With the LOAD statement (DB-Access)
- With the FILETOBLOB function
- From BLOB (**ifx\_lo\_t**) host variables ()

If you select a BLOB column using DB-Access, only the string `<SBlob value>` is returned; no actual value is displayed.



## BOOLEAN data type

The BOOLEAN data type stores `TRUE` or `FALSE` data values as a single byte.

The following table shows internal and literal representations of the BOOLEAN data type.

Logical Value	Internal Representation	Literal Representation
TRUE	\0	't'
FALSE	\1	'f'
NULL	Internal Use Only	NULL

You can compare two BOOLEAN values to test for equality or inequality. You can also compare a BOOLEAN value to the Boolean literals `'t'` and `'f'`. BOOLEAN values are not case-sensitive; `'t'` is equivalent to `'T'` and `'f'` to `'F'`.

You can use a BOOLEAN column to store what a Boolean expression returns. In the following example, the value of `boolean_column` is `'t'` if `column1` is less than `column2`, `'f'` if `column1` is greater than or equal to `column2`, and NULL if the value of either `column1` or `column2` is unknown:

```
UPDATE my_table SET boolean_column = lessthan(column1, column2)
```

## BYTE data type

The BYTE data type stores any kind of binary data in an undifferentiated byte stream. Binary data typically consists of digitized information, such as spreadsheets, program load modules, digitized voice patterns, and so on.

The term *simple large object* refers to an instance of a TEXT or BYTE data type. No more than 195 columns of the same table can be declared as BYTE and TEXT data types.

The BYTE data type has no maximum size. A BYTE column has a theoretical limit of  $2^{31}$  bytes and a practical limit that your disk capacity determines.

You cannot use the MEDIUM or HIGH options of the UPDATE STATISTICS statement to calculate distribution statistics on BYTE columns.

### BYTE objects in DML operations

You can store, retrieve, update, or delete the contents of a BYTE column. You cannot, however, use BYTE operands in arithmetic or string operations, nor assign literals to BYTE columns with the SET clause of the UPDATE statement. You also cannot use BYTE objects in any of the following contexts in a SELECT statement:

- With aggregate functions
- With the IN clause
- With the MATCHES or LIKE clauses
- With the GROUP BY clause
- With the ORDER BY clause

BYTE operands are valid in Boolean expressions only when you are testing for NULL values with the IS NULL or IS NOT NULL operators.

You can use the following methods, which can load rows or update fields, to insert BYTE data:

- With the **dbload** utility
- With the LOAD statement (DB-Access)
- From BYTE host variables ( )

You cannot use a quoted text string, number, or any other actual value to insert or update BYTE columns.

When you select a BYTE column, you can receive all or part of it. To retrieve it all, use the regular syntax for selecting a column. You can also select any part of a BYTE column by using subscripts, as the next example, which reads the first 75 bytes of the **cat\_picture** column associated with the catalog number 10001:

```
SELECT cat_picture [1,75] FROM catalog WHERE catalog_num = 10001
```

A built-in cast converts BYTE values to BLOB values. For more information, see the *HCL OneDB™ Database Design and Implementation Guide*.

If you select a BYTE column using the DB-Access Interactive Schema Editor, only the string "<BYTE value>" is returned; no data value is displayed.



**Important:** If you try to return a BYTE column from a subquery, an error results, even if the column is not used in a Boolean expression nor with an aggregate.

## CHAR(n) data type

The CHAR data type stores any string of letters, numbers, and symbols. It can store single-byte and multibyte characters, based on the database locale.

A CHAR(*n*) column has a length of *n* bytes, where  $1 < n < 32,767$ . If you do not specify *n*, CHAR(1) is the default length. Character columns typically store alphanumeric strings, such as names, addresses, phone numbers, and so on. When a value is retrieved or stored as CHAR(*n*), exactly *n* bytes of data are transferred. If the string is shorter than *n* bytes, the string is extended with blank spaces up to the declared length. If the data value is longer than *n* bytes, a data string of length *n* that has been truncated from the right is inserted or retrieved, without the database server raising an exception.

This does not create partial characters in multibyte locales. In right-to-left locales, such as Arabic, Hebrew, or Farsi, the truncation is from the left.

Size specifications in CHAR data type declarations can be affected by the SQL\_LOGICAL\_CHAR feature that is described in the section [Logical Character Semantics in Character Type Declarations on page 119](#).

For more information about East Asian locales that support multibyte code sets, see [Multibyte Characters with VARCHAR on page 118](#).

## Treating CHAR Values as Numeric Values

If you plan to perform calculations on numbers stored in a column, you should assign a number data type to that column. Although you can store numbers in CHAR columns, you might not be able to use them in some arithmetic operations. For example, if you insert a sum into a CHAR column, you might experience overflow problems if the CHAR column is too small to hold the value. In this case, the insert fails. Numbers that have leading zeros (such as some zip codes) have the zeros stripped if they are stored as number types INTEGER or SMALLINT. Instead, store these numbers in CHAR columns.

## Sorting and Relational Comparisons

In general, the collating order for sorting CHAR values is the order of characters in the code set. (An exception is the MATCHES operator with ranges; see [Collating VARCHAR Values on page 118](#).) For more information about collation order, see the *HCL OneDB™ GLS User's Guide*.

For multibyte locales, the database supports any multibyte characters in the code set. When storing multibyte characters in a CHAR data type, make sure to calculate the number of bytes needed. For more information about multibyte characters and locales, see the *HCL OneDB™ GLS User's Guide*.

CHAR values are compared to other CHAR values by padding the shorter value on the right with blank spaces until the values have equal length, and then comparing the two values, using the code-set order for collation.

## Nonprintable Characters with CHAR

A CHAR value can include tab, newline, whitespace, and nonprintable characters. You must, however, use an application to insert nonprintable characters into host variables and the host variables into your database. After passing nonprintable characters to the database server, you can store or retrieve them. After you select nonprintable characters, fetch them into host variables and display them with your own display mechanism.

An important exception is the first value in the ASCII code set is used as the end-of-data terminator symbol in columns of the CHAR data type. For this reason, any subsequent characters in the same string cannot be retrieved from a CHAR column, because the database server reads only the characters (if any) that precede this null terminator. For example, you cannot use the following 7-byte string as a CHAR data type value with a length of 7 bytes:

```
abc\0def
```

If you try to display nonprintable characters with DB-Access your screen returns inconsistent results. (Which characters are nonprintable is locale-dependent. For more information see the discussion of code-set conversion between the client and the database server in the *HCL OneDB™ GLS User's Guide*.)

## CHARACTER(n) data type

The CHARACTER data type is a synonym for CHAR.

## CHARACTER VARYING(m,r) data type

The CHARACTER VARYING data type stores a string of letters, digits, and symbols of varying length, where *m* is the maximum size of the column (in bytes) and *r* is the minimum number of bytes reserved for that column.

The CHARACTER VARYING data type complies with ANSI/ISO standard for SQL; the non-ANSI VARCHAR data type supports the same functionality. For more information, see the description of the VARCHAR type in [VARCHAR\(m,r\) data type on page 116](#).

## CLOB data type

The CLOB data type stores any kind of text data in random-access chunks, called sbspaces. Text data can include text-formatting information, if this information is also textual, such as PostScript™, Hypertext Markup Language (HTML), Standard Graphic Markup Language (SGML), or Extensible Markup Language (XML) data.

The term *smart large object* refers to CLOB and BLOB data types. The CLOB data type supports special operations for character strings that are inappropriate for BLOB values. A CLOB value can be up to 4 terabytes ( $4 \times 2^{40}$  bytes) in length. The minimum amount of disk space allocated for smart large object data types is 512 bytes.

Use the BLOB data type (see [BLOB data type on page 88](#)) for random access to binary data. For general information about the CLOB and BLOB data types, see [Smart large objects on page 122](#).

The following SQL functions can perform operations on a CLOB column:

- FILETOCLOB copies a file into a CLOB column.
- LOTOFILE copies a CLOB (or BLOB) value into a file.
- LOCOPY copies a CLOB (or BLOB) value to a new smart large object.
- ENCRYPT\_DES or ENCRYPT\_TDES creates an encrypted BLOB value from a plain-text CLOB argument.
- DECRYPT\_BINAR or DECRYPT\_CHAR returns an unencrypted BLOB value from an encrypted BLOB argument (that ENCRYPT\_DES or ENCRYPT\_TDES created from a plain-text CLOB value).

For more information about these SQL functions, see the *HCL OneDB™ Guide to SQL: Syntax*.

No casts exist for CLOB data. Therefore, the database server cannot convert data of the CLOB type to any other data type, except by using these encryption and decryption functions to return a BLOB. Therefore, the database server cannot convert data of the CLOB type to any other data type. Within SQL, you are limited to the equality (=) comparison operation for CLOB data. To perform additional operations, you must use one of the application programming interfaces from within your client application.

## Multibyte characters with CLOB

### About this task

You can insert data into CLOB columns in the following ways:

- With the dbload utility
- With the LOAD statement (DB-Access)
- From CLOB (**ifx\_lo\_t**) host variables (ESQL/C)

For examples of CLOB types, see the *HCL OneDB™ Guide to SQL: Tutorial* and the *HCL OneDB™ Database Design and Implementation Guide*.

With GLS, the following rules apply:

- Multibyte CLOB characters must be defined in the database locale.
- The CLOB data type is collated in code-set order.
- The database server handles code-set conversions for CLOB data.

For more information about database locales, collation order, and code-set conversion, see the *HCL OneDB™ GLS User's Guide*.

## DATE data type

The DATE data type stores the calendar date. DATE data types require four bytes. A calendar date is stored internally as an integer value equal to the number of days since December 31, 1899.

Because DATE values are stored as integers, you can use them in arithmetic expressions. For example, you can subtract a DATE value from another DATE value. The result, a positive or negative INTEGER value, indicates the number of days that elapsed between the two dates. (You can use a UNITS DAY expression to convert the result to an INTERVAL DAY TO DAY data type.)

The following example shows the default display format of a DATE column:

```
mm/dd/yyyy
```

In this example, *mm* is the month (1-12), *dd* is the day of the month (1-31), and *yyyy* is the year (0001-9999). You can specify a different order of time units and a different time-unit separator than / (or no separator) by setting the **DBDATE** environment variable. For more information, see [DBDATE environment variable on page 166](#).

In non-default locales, you can display dates in culture-specific formats. The locale and the **GL\_DATE** and **DBDATE** environment variables (as described in the next chapter) affect the display formatting of DATE values. They do not, however, affect the internal storage format for DATE columns in the database. For more information, see the *HCL OneDB™ GLS User's Guide*.

## DATETIME data type

The DATETIME data type stores an instant in time expressed as a calendar date and time of day.

You select how precisely a DATETIME value is stored; its precision can range from a year to a fraction of a second.

DATETIME stores a data value as a contiguous series of fields that represents each time unit (*year, month, day*, and so forth) in the data type declaration.

Field qualifiers to specify a DATETIME data type have this format:

```
DATETIME largest_qualifier TO smallest_qualifier
```

This resembles an INTERVAL field qualifier, but DATETIME represents a point in time, rather than (like INTERVAL) a span of time. These differences exist between DATETIME and INTERVAL qualifiers:

- The DATETIME keyword replaces the INTERVAL keyword.
- DATETIME field qualifiers cannot specify a nondefault precision for the *largest\_qualifier* time unit.
- Field qualifiers of a DATETIME data type can include YEAR, MONTH, and smaller time units, but an INTERVAL data type that includes the DAY field qualifier (or smaller time units) cannot also include the YEAR or MONTH field qualifiers.

The *largest\_qualifier* and *smallest\_qualifier* of a DATETIME data type can be any of the fields that the following table lists, provided that *smallest\_qualifier* does not specify a larger time unit than *largest\_qualifier*. (The largest and smallest time units can be the same; for example, DATETIME YEAR TO YEAR.)

**Table 44. DATETIME field qualifiers**

Qualifier field	Valid entries
YEAR	A year numbered from 1 to 9,999 (A.D.)
MONTH	A month numbered from 1 to 12
DAY	A day numbered from 1 to 31, as appropriate to the month
HOUR	An hour numbered from 0 (midnight) to 23
MINUTE	A minute numbered from 0 to 59
SECOND	A second numbered from 0 - 59
FRACTION	A decimal fraction-of-a-second with up to 5 digits of scale. The default scale is 3 digits (a thousandth of a second). For <i>smallest_qualifier</i> to specify another scale, write FRACTION( <i>n</i> ), where <i>n</i> is the number of digits from 1 - 5.

The declaration of a DATETIME column need not include the full YEAR to FRACTION range of time units. It can include any contiguous subset of these time units, or even only a single time unit.

For example, you can enter a MONTH TO HOUR value in a column declared as YEAR TO MINUTE, if each entered value contains information for a contiguous series of time units. You cannot, however, enter a value for only the MONTH and HOUR; the entry must also include a value for DAY.

If you use the DB-Access TABLE menu, and you do not specify the DATETIME qualifiers, a default DATETIME qualifier, YEAR TO YEAR, is assigned.

A valid DATETIME literal must include the DATETIME keyword, the values to be entered, and the field qualifiers. You must include these qualifiers because, as noted earlier, the value that you enter can contain fewer fields than were declared for that column. Acceptable qualifiers for the first and last fields are identical to the list of valid DATETIME fields that are listed in the table [Table 44: DATETIME field qualifiers on page 94](#).

Write values for the field qualifiers as integers and separate them with delimiters. The following table lists the delimiters that are used with DATETIME values in the default US English locale. (These are a superset of the delimiters that are used in INTERVAL values.)

**Table 45. Delimiters used with DATETIME**

Delimiter	Placement in DATETIME Literal
Hyphen ( - )	Between the YEAR, MONTH, and DAY time-unit values
Blank space ( )	Between the DAY and HOUR time-unit values
Colon ( : )	Between the HOUR, MINUTE, and SECOND time-unit values
Decimal point ( . )	Between the SECOND and FRACTION time-unit values

The following illustration shows a DATETIME YEAR TO FRACTION(3) value with delimiters.

Figure 1. Example DATETIME Value with Delimiters

2003-09-23 12:42:06.001

When you enter a value with fewer time-unit fields than in the column, the value that you enter is expanded automatically to fill all the declared time-unit fields. If you leave out any more significant fields, that is, time units larger than any that you include, those fields are filled automatically with the current values for those time units from the system clock calendar. If you leave out any less-significant fields, those fields are filled with zeros (or with `1` for MONTH and DAY) in your entry.

You can also enter DATETIME values as character strings. The character string must include information for each field defined in the DATETIME column. The INSERT statement in the following example shows a DATETIME value entered as a character string:

```
INSERT INTO cust_calls (customer_num, call_dtime, user_id,
    call_code, call_descr)
VALUES (101, '2001-01-14 08:45', 'maryj', 'D',
    'Order late - placed 6/1/00');
```

If `call_dtime` is declared as DATETIME YEAR TO MINUTE, the character string must include values for the *year*, *month*, *day*, *hour*, and *minute* fields.

If the character string does not contain information for all the declared fields (or if it adds additional fields), then the database server returns an error.

All fields of a DATETIME column are two-digit numbers except for the *year* and *fraction* fields. The *year* field is stored as four digits. When you enter a two-digit value in the year field, how the abbreviated year is expanded to four digits depends on the setting of the **DBCENTURY** environment variable.

For example, if you enter `02` as the year value, whether the year is interpreted as `1902`, `2002`, or `2102` depends on the setting of **DBCENTURY** and on the value of the system clock calendar at execution time. If you do not set **DBCENTURY**, the leading digits of the current year are appended by default.

The *fraction* field requires  $n$  digits where  $1 < n < 5$ , rounded up to an even number. You can use the following formula (rounded up to a whole number of bytes) to calculate the number of bytes that a DATETIME value requires:

$$(total\ number\ of\ digits\ for\ all\ fields) / 2 + 1$$

For example, a YEAR TO DAY qualifier requires a total of eight digits (four for *year*, two for *month*, and two for *day*). According to the formula, this data value requires 5, or  $(8/2) + 1$ , bytes of storage.

The USEOSTIME configuration parameter can affect the subsecond granularity when the database server obtains the current time from the operating system in SQL statements. For details, see the *HCL OneDB™ Administrator's Reference*.

With an ESQL API, the **DBTIME** environment variable affects DATETIME formatting. Nondefault locales and settings of the **GL\_DATE** and **DBDATE** environment variables also affect the display of datetime data. They do not, however, affect the internal storage format of a DATETIME column.

If you specify a locale other than U.S. English, the locale defines the culture-specific display formats for DATETIME values. To change the default display format, change the setting of the **GL\_DATETIME** environment variable. When a database with a nondefault locale uses a nondefault **GL\_DATETIME** setting, the **USE\_DTEENV** environment variable must be set to 1 before the database server can correctly process localized DATETIME values in the following operations:

- using the LOAD or UNLOAD feature of DB-Access
- using the dbexport or dbimport migration utilities
- using DML statements of SQL on database tables or on objects that the CREATE EXTERNAL TABLE statement defined.

For more information about locales and GLS environment variables that can specify end-user DATETIME formats, see the *HCL OneDB™ GLS User's Guide*.

## DEC data type

The DEC data type is a synonym for DECIMAL.

## DECIMAL

The DECIMAL data type can take two forms: DECIMAL( $p$ ) floating point and DECIMAL( $p,s$ ) fixed point.

In an ANSI-compliant database all DECIMAL numbers are fixed point.

By default, literal numbers that include a decimal ( . ) point are interpreted by the database server as DECIMAL values.

### DECIMAL( $p$ ) Floating Point

The DECIMAL data type stores decimal floating-point numbers up to a maximum of 32 significant digits, where  $p$  is the total number of significant digits (the *precision*).

Specifying precision is optional. If you specify no precision ( $p$ ), DECIMAL is treated as DECIMAL(16), a floating-point decimal with a precision of 16 places. DECIMAL( $p$ ) has an absolute exponent range between  $10^{-130}$  and  $10^{124}$ .



If you declare a `DECIMAL(p)` column in an ANSI-compliant database, the scale defaults to `DECIMAL(p, 0)`, meaning that only integer values can be stored in this data type.

In a database that is not ANSI-compliant, a `DECIMAL(p)` is a floating-point data type of a scale large enough to store the exponential notation for a value.

For example, the following calculation shows how many bytes of storage a `DECIMAL(5)` column requires in the default locale (where the decimal point occupies a single byte):

```

1 byte for the sign of the data value
1 byte for the 1st digit
1 byte for the decimal point
4 bytes for the rest of the digits (precision of 5 - 1)
1 byte for the e symbol
1 byte for the sign of the exponent
3 bytes for the exponent
-----
12 bytes total

```

Thus, "12345" in a `DECIMAL(5)` column is displayed as "12345.00000" (that is, with a scale of 6) in a database that is not ANSI-compliant.

## DECIMAL (p,s) Fixed Point

In fixed-point numbers, `DECIMAL(p,s)`, the decimal point is fixed at a specific place, regardless of the value of the number. When you specify a column of this type, you declare its precision (*p*) as the total number of digits that it can store, from 1 to 32. You declare its *scale* (*s*) as the total number of digits in the fractional part (that is, to the right of the decimal point).

All numbers with an absolute value less than  $0.5 * 10^{-s}$  have the value zero. The largest absolute value of a `DECIMAL(p,s)` data type that you can store without an overflow error is  $10^{p-s} - 10^{-s}$ . A `DECIMAL` column typically stores numbers with fractional parts that must be stored and displayed exactly (for example, rates or percentages). In an ANSI-compliant database, all `DECIMAL` numbers must have absolute values in the range  $10^{-32}$  to  $10^{+31}$ .

## DECIMAL Storage


The database server uses one byte of disk storage to store two digits of a decimal number, plus an additional byte to store the exponent and sign, with the first byte representing a sign bit and a 7-bit exponent in excess-65 format. The rest of the bytes express the mantissa as base-100 digits. The significant digits to the left of the decimal and the significant digits to the right of the decimal are stored in separate groups of bytes. At the maximum *precision* specification, `DECIMAL(32,s)` data types can store *s*-1 decimal digits to the right of the decimal point, if *s* is an odd number.

How the database server stores decimal numbers is illustrated in the following example. If you specify `DECIMAL(6,3)`, the data type consists of three significant digits in the integral part and three significant digits in the fractional part (for instance, 123.456). The three digits to the left of the decimal are stored on 2 bytes (where one of the bytes only holds a single digit)

and the three digits to the right of the decimal are stored on another 2 bytes, as [Figure 2: Schematic that illustrates the storage of digits in a decimal \(p,s\) value on page 98](#) illustrates.

(The exponent byte is not shown.) With the additional byte required for the exponent and sign, DECIMAL(6,3) requires a total of 5 bytes of storage.

Figure 2. Schematic that illustrates the storage of digits in a decimal (p,s) value



You can use the following formulas (rounded down to a whole number of bytes) to calculate the byte storage (N) for a DECIMAL(p,s) data type (where N includes the byte that is required to store the exponent and the sign):

If the *scale* is odd:  $N = (\textit{precision} + 4) / 2$   
 If the *scale* is even:  $N = (\textit{precision} + 3) / 2$

For example, the data type DECIMAL(5,3) requires 4 bytes of storage ( $9/2$  rounded down equals 4).

There is one caveat to these formulas. The maximum number of bytes the database server uses to store a decimal value is 17. One byte is used to store the exponent and sign, leaving 16 bytes to store up to 32 digits of precision. If you specify a precision of 32 and an *odd* scale, however, you lose 1 digit of precision. Consider, for example, the data type DECIMAL(32,31). This decimal is defined as 1 digit to the left of the decimal and 31 digits to the right. The 1 digit to the left of the decimal requires 1 byte of storage. This leaves only 15 bytes of storage for the digits to the right of the decimal. The 15 bytes can accommodate only 30 digits, so 1 digit of precision is lost.

## DISTINCT data types

A DISTINCT type is a data type that is derived from a source type (called the base type).

A source type can be:

- A built-in type
- An existing DISTINCT type
- An existing named ROW type
- An existing opaque type

A DISTINCT type inherits from its source type the length and alignment on the disk. A DISTINCT type thus makes efficient use of the preexisting functionality of the database server.

When you create a DISTINCT data type, the database server automatically creates two explicit casts: one cast from the DISTINCT type to its source type and one cast from the source type to the DISTINCT type. A DISTINCT type based on a built-in source type does not inherit the built-in casts that are provided for the built-in type. A DISTINCT type does inherit, however, any user-defined casts that have been defined on the source type.

A DISTINCT type cannot be compared directly to its source type. To compare the two types, you must first explicitly cast one type to the other.

You must define a DISTINCT type in the database. Definitions of DISTINCT types are stored in the **sysxdtypes** system catalog table. The following SQL statements maintain the definitions of DISTINCT types in the database:

- The CREATE DISTINCT TYPE statement adds a DISTINCT type to the database.
- The DROP TYPE statement removes a previously defined DISTINCT type from the database.

For more information about the SQL statements mentioned above, see the *HCL OneDB™ Guide to SQL: Syntax*. For information about casting DISTINCT data types, see [Casts for distinct types on page 136](#). For examples that show how to create and register cast functions for a DISTINCT type, see the *HCL OneDB™ Database Design and Implementation Guide*.

Size specifications in declarations of DISTINCT types whose base types are built-in character types can be affected by the SQL\_LOGICAL\_CHAR feature that is described in the section [Logical Character Semantics in Character Type Declarations on page 119](#).

## DOUBLE PRECISION data types

The DOUBLE PRECISION keywords are a synonym for the FLOAT keyword.

### FLOAT(*n*)

The FLOAT data type stores double-precision floating-point numbers with up to 17 significant digits. FLOAT corresponds to IEEE 4-byte floating-point, and to the **double** data type in C. The range of values for the FLOAT data type is the same as the range of the C **double** data type on your computer.

You can use *n* to specify the precision of a FLOAT data type, but SQL ignores the precision. The value *n* must be a whole number between 1 and 14.

A column with the FLOAT data type typically stores scientific numbers that can be calculated only approximately. Because floating-point numbers retain only their most significant digits, the number that you enter in this type of column and the number the database server displays can differ slightly.

The difference between the two values depends on how your computer stores floating-point numbers internally. For example, you might enter a value of 1.1000001 into a FLOAT field and, after processing the SQL statement, the database server might display this value as 1.1. This situation occurs when a value has more digits than the floating-point number can store. In this case, the value is stored in its approximate form with the least significant digits treated as zeros.

FLOAT data types usually require 8 bytes of storage per value. Conversion of a FLOAT value to a DECIMAL value results in 17 digits of precision.

## IDSSECURITYLABEL data type

The IDSSECURITYLABEL type stores a security label in a table that is protected by a label-based access control (LBAC) security policy.

Only a user who holds the **DBSECADM** role can create, alter, or drop a column of this data type. `IDSSECURITYLABEL` is a built-in `DISTINCT OF VARCHAR(128)` data type, but because its use is restricted to databases that implement label-based access control, it is not classified as a character data type. A table that is protected by a security policy can have only one `IDSSECURITYLABEL` column. A table that is not associated with any label-based security policy cannot include an `IDSSECURITYLABEL` column. You cannot encrypt the security label in a column of type `IDSSECURITYLABEL`.

For a discussion of security policies, security components, security labels, and other concepts of label-based access control (LBAC), see the HCL OneDB™ Security Guide.

## INT data type

The `INT` data type is a synonym for `INTEGER`.

## INT8

The `INT8` data type stores whole numbers that can range in value from  $-9,223,372,036,854,775,807$  to  $9,223,372,036,854,775,807$  [or  $-(2^{63}-1)$  to  $2^{63}-1$ ], for 18 or 19 digits of precision.

The number  $-9,223,372,036,854,775,808$  is a reserved value that cannot be used. The `INT8` data type is typically used to store large counts, quantities, and so on.

HCL OneDB™ stores `INT8` data in internal format that can require up to 10 bytes of storage.

Arithmetic operations and sort comparisons are performed more efficiently on integer data than on floating-point or fixed-point decimal data, but `INT8` cannot store data with absolute values beyond  $|2^{63}-1|$ . If a value exceeds the numeric range of `INT8`, the database server does not store the value.

## INTEGER data type

The `INTEGER` data type stores whole numbers that range from  $-2,147,483,647$  to  $2,147,483,647$  for 9 or 10 digits of precision.

The number  $2,147,483,648$  is a reserved value and cannot be used. The `INTEGER` value is stored as a signed binary integer and is typically used to store counts, quantities, and so on.

Arithmetic operations and sort comparisons are performed more efficiently on integer data than on float or decimal data. `INTEGER` columns, however, cannot store absolute values beyond  $(2^{31}-1)$ . If a data value lies outside the numeric range of `INTEGER`, the database server does not store the value.

`INTEGER` data types require 4 bytes of storage per value.

## INTERVAL data type

The `INTERVAL` data type stores a value that represents a span of time. `INTERVAL` types are divided into two classes: *year-month intervals* and *day-time intervals*.

A year-month interval can represent a span of years and months, and a day-time interval can represent a span of days, hours, minutes, seconds, and fractions of a second.

An INTERVAL value is always composed of one value or a series of values that represents time units. Within a data-definition statement such as CREATE TABLE or ALTER TABLE that defines the precision of an INTERVAL data type, the qualifiers must have the following format:

```
INTERVAL largest_qualifier(n) TO smallest_qualifier
```

Here the *largest\_qualifier* and *smallest\_qualifier* keywords are taken from one of the two INTERVAL classes, as shown in the table [Table 46: Interval Classes on page 101](#).

If SECOND (or a larger time unit) is the *largest\_qualifier*, the declaration of an INTERVAL data type can optionally specify *n*, the precision of the largest time unit (for *n* ranging from 1 to 9); this is not a feature of DATETIME data types.

If *smallest\_qualifier* is FRACTION, you can also specify a scale in the range from 1 to 5. For FRACTION TO FRACTION qualifiers, the upper limit of *n* is 5, rather than 9. There are two incommensurable classes of INTERVAL data types:

- Those with a *smallest\_qualifier* larger than DAY
- Those with a *largest\_qualifier* smaller than MONTH

**Table 46. Interval Classes**

Interval Class	Time Units	Valid Entry
YEAR-MONTH INTERVAL	YEAR	A number of years
YEAR-MONTH INTERVAL	MONTH	A number of months
DAY-TIME INTERVAL	DAY	A number of days
DAY-TIME INTERVAL	HOUR	A number of hours
DAY-TIME INTERVAL	MINUTE	A number of minutes
DAY-TIME INTERVAL	SECOND	A number of seconds
DAY-TIME INTERVAL	FRACTION	A decimal fraction of a second, with up to 5 digits. The default scale is 3 digits (thousandth of a second). To specify a non-default scale, write FRACTION( <i>n</i> ), where $1 < n < 5$ .

As with DATETIME data types, you can define an INTERVAL to include only the subset of time units that you need. But because the construct of “month” (as used in calendar dates) is not a time unit that has a fixed number of days, a single INTERVAL value cannot combine months and days; arithmetic that involves operands of the two different INTERVAL classes is not supported.

A value entered into an INTERVAL column need not include the full range of time units that were specified in the data-type declaration of the column. For example, you can enter a value of HOUR TO SECOND precision into a column defined as DAY TO SECOND. A value must always consist, however, of contiguous time units. In the previous example, you cannot enter only the HOUR and SECOND values; you must also include MINUTE values.

A valid INTERVAL literal contains the INTERVAL keyword, the values to be entered, and the field qualifiers. (See the discussion of literal intervals in the *HCL OneDB™ Guide to SQL: Syntax*.) When a value contains only one field, the largest and smallest fields are the same.

When you enter a value in an INTERVAL column, you must specify the largest and smallest fields in the value, just as you do for DATETIME values. In addition, you can optionally specify the precision of the first field (and the scale of the last field if it is a FRACTION). If the largest and smallest field qualifiers are both FRACTION, you can specify only the scale in the last field.

Acceptable qualifiers for the largest and smallest fields are identical to the list of INTERVAL fields that the tab;e [Table 46: Interval Classes on page 101](#) displays.

If you use the DB-Access **TABLE** menu, but you specify no INTERVAL field qualifiers, then a default INTERVAL qualifier, YEAR TO YEAR, is assigned.

The *largest\_qualifier* in an INTERVAL value can be up to nine digits (except for FRACTION, which cannot be more than five digits), but if the value that you want to enter is greater than the default number of digits allowed for that field, you must explicitly identify the number of significant digits in the value that you enter. For example, to define an INTERVAL of DAY TO HOUR that can store up to 999 days, you can specify it the following way:

```
INTERVAL DAY(3) TO HOUR
```

INTERVAL literals use the same delimiters as DATETIME literals (except that MONTH and DAY time units are not valid within the same INTERVAL value). the following table shows the INTERVAL delimiters.

**Table 47. INTERVAL Delimiters**

Delimiter	Placement in an INTERVAL Literal
Hyphen	Between the YEAR and MONTH portions of the value
Blank space	Between the DAY and HOUR portions of the value
Colon	Between the HOUR, MINUTE, and SECOND portions of the value
Decimal point	Between the SECOND and FRACTION portions of the value

You can also enter INTERVAL values as character strings. The character string must include information for the same time units that were specified in the data-type declaration for the column. The INSERT statement in the following example shows an INTERVAL value entered as a character string:

```
INSERT INTO manufact (manu_code, manu_name, lead_time)
VALUES ('BR0', 'Ball-Racquet Originals', '160')
```

Because the **lead\_time** column is defined as INTERVAL DAY(3) TO DAY, this INTERVAL value requires only one field, the span of days required for lead time. If the character string does not contain information for all fields (or adds additional fields), the database server returns an error. For additional information about entering INTERVAL values as character strings, see the *HCL OneDB™ Guide to SQL: Syntax*.

By default, all fields of an INTERVAL column are two-digit numbers, except for the year and fraction fields. The year field is stored as four digits. The fraction field requires  $n$  digits where  $1 < n < 5$ , rounded up to an even number. You can use the following formula (rounded up to a whole number of bytes) to calculate the number of bytes required for an INTERVAL value:

$$(total\ number\ of\ digits\ for\ all\ fields) / 2 + 1$$

For example, INTERVAL YEAR TO MONTH requires six digits (four for *year* and two for *month*), and requires 4, or  $(6/2) + 1$ , bytes of storage.

For information about using INTERVAL as a constant expression, see the description of the INTERVAL Field Qualifier in the *HCL OneDB™ Guide to SQL: Syntax*.

## LIST(e) data type

The LIST data type is a collection type that can store ordered non-NULL elements of the same SQL data type.

The LIST data type supports, but does not require, duplicate element values. The elements of a LIST data type have ordinal positions. The LIST object must have a first element, which can be followed by a second element, and so on.

For unordered collection data types that do not support ordinal positions, see [MULTISET\(e\) data type on page 106](#) and [SET\(e\) data type on page 113](#). For complex data types that can store a set of values that includes different SQL data types, see [ROW Data Types on page 130](#).

No more than 97 columns of the same table can be declared as LIST data types. (The same restriction applies to SET and MULTISET collection types.)

By default, the database server inserts new elements into a LIST object at the end of the set of elements. To support the ordinal position of a LIST, the INSERT statement provides the AT clause. This clause allows you to specify the position at which you want to insert a LIST element value. For more information, see the INSERT statement in the *HCL OneDB™ Guide to SQL: Syntax*.

All elements in a LIST object have the same element type. To specify the element type, use the following syntax:

```
LIST(element_type NOT NULL)
```

The *element\_type* of a LIST can be any of the following data types:

- A built-in type, except SERIAL, SERIAL8, BIGSERIAL, BYTE, and TEXT
- A DISTINCT type
- An unnamed or named ROW type
- Another collection type
- An opaque type

You must specify the NOT NULL constraint for LIST elements. No other constraints are valid for LIST columns. For more information about the syntax of the LIST data type, see the *HCL OneDB™ Guide to SQL: Syntax*.

You can use LIST in most contexts where any other data type is valid. For example:

- After the IN predicate in the WHERE clause of a SELECT statement to search for matching LIST values
- As an argument to the CARDINALITY or `mi_collection_card()` function to determine the number of elements in a LIST column

You *cannot* use LIST values as arguments to an aggregate function such as AVG, MAX, MIN, or SUM.

Just as with the other collection data types, you must use parentheses ( `()` ) in data type declarations to delimit the set of elements of a LIST data type:

```
CREATE FUNCTION update_nums( list1 LIST (ROW (a VARCHAR(10),
                                     b VARCHAR(10),
                                     c INT) NOT NULL ));
```

In SQL expressions that include literal LIST values, however, you must use braces ( `{ }` ) to delimit the set of elements of a LIST object, as in the examples that follow.

Two LIST values are equal if they have the same elements in the same order. The following are both examples of LIST objects, but their values are not equal. :

```
LIST{"blue", "green", "yellow"}
```

```
LIST{"yellow", "blue", "green"}
```

The above expressions are not equal because the values are not in the same order. To be equal, the second statement must be:

```
LIST{"blue", "green", "yellow"}
```

## LVARCHAR(m) data type

Use the LVARCHAR data type to create a column for storing variable-length character strings whose upper limit (*m*) can be up to 32,739 bytes.

This limit is much greater than the VARCHAR data type, which is used for character strings that are no longer than 255 bytes.

The LVARCHAR data type is implemented as a built-in opaque data type. You can access LVARCHAR columns in remote tables by using distributed queries across databases of the same or different HCL OneDB™ instances.

By default, the database server interprets quoted strings as LVARCHAR types. It also uses LVARCHAR for input and output casts for opaque data types.

The LVARCHAR data type stores opaque data types in the string (external) format. Each opaque type has an input support function and cast, which convert it from LVARCHAR to a form that database servers can manipulate. Each opaque type also has an output support function and cast, which convert it from its internal representation to LVARCHAR.



**Important:** When LVARCHAR is declared (with no size specification) as the data type of a column in a database table, the default maximum size is 2 KB (2048 bytes), but you can specify an explicit maximum length of up to 32,739





bytes. When LVARCHAR is used in I/O operations on an opaque data type, however, the maximum size is limited only by the operating system.

You cannot use the MEDIUM or HIGH options of the UPDATE STATISTICS statement to calculate distribution statistics on LVARCHAR columns.

Size specifications in LVARCHAR data type declarations can be affected by the SQL\_LOGICAL\_CHAR feature that is described in the section [Logical Character Semantics in Character Type Declarations on page 119](#).

For more information about LVARCHAR, see the *HCL OneDB™ User-Defined Routines and Data Types Developer's Guide*.

## MONEY(p,s) data type

The MONEY data type stores currency amounts.

Like the DECIMAL(p,s) data type, MONEY can store fixed-point numbers up to a maximum of 32 significant digits, where *p* is the total number of significant digits (the precision) and *s* is the number of digits to the right of the decimal point (the scale).

Unlike the DECIMAL data type, the MONEY data type is always treated as a fixed-point decimal number. The database server defines the data type MONEY(*p*) as DECIMAL(*p*,2). If the precision and scale are not specified, the database server defines a MONEY column as DECIMAL(16,2).

You can use the following formula (rounded down to a whole number of bytes) to calculate the byte storage for a MONEY data type:

```
If the scale is odd: N = (precision + 4) / 2
If the scale is even: N = (precision + 3) / 2
```

For example, a MONEY data type with a precision of 16 and a scale of 2 (MONEY(16,2)) requires 10 or  $(16 + 3)/2$ , bytes of storage.

In the default locale, client applications format values from MONEY columns with the following currency notation:

- A currency symbol: a dollar sign ( \$ ) at the front of the value
- A thousands separator: a comma ( , ) that separates every three digits in the integer part of the value
- A decimal point: a period ( . ) between the integer and fractional parts of the value

To change the format for MONEY values, change the **DBMONEY** environment variable. For valid **DBMONEY** settings, see [DBMONEY environment variable on page 171](#).

The default value that the database server uses for scale is locale-dependent. The default locale specifies a default scale of two. For non-default locales, if the scale is omitted from the declaration, the database server creates MONEY values with a locale-specific scale.

The currency notation that client applications use is locale-dependent. If you specify a nondefault locale, the client uses a culture-specific format for MONEY values that might differ from the default U.S. English format in the leading (or trailing) currency symbol, thousands separator, and decimal separator, depending on what the locale files specify. For more information about locale dependency, see the *HCL OneDB™ GLS User's Guide*.

## MULTISET(e) data type

The MULTISET data type is a collection type that stores a non-ordered set that can include duplicate element values.

The elements in a MULTISET have no ordinal position. That is, there is no concept of a first, second, or third element in a MULTISET. (For a collection type with ordinal positions for elements, see [LIST\(e\) data type on page 103.](#))

All elements in a MULTISET have the same element type. To specify the element type, use the following syntax:

```
MULTISET(element_type NOT NULL)
```

The *element\_type* of a collection can be any of the following types:

- Any built-in type, except SERIAL, SERIAL8, BIGSERIAL, BYTE, and TEXT
- An unnamed or a named ROW type
- Another collection type or opaque type

You can use MULTISET anywhere that you use any other data type, unless otherwise indicated. For example:

- After the IN predicate in the WHERE clause of a SELECT statement to search for matching MULTISET values
- As an argument to the CARDINALITY or `mi_collection_card()` function to determine the number of elements in a MULTISET column

You *cannot* use MULTISET values as arguments to an aggregate function such as AVG, MAX, MIN, or SUM.

You must specify the NOT NULL constraint for MULTISET elements. No other constraints are valid for MULTISET columns. For more information about the MULTISET collection type, see the *HCL OneDB™ Guide to SQL: Syntax*.

Two multiset data values are equal if they have the same elements, even if the elements are in different positions within the set. The following examples are both multiset values but are not equal:

```
MULTISET {"blue", "green", "yellow"}
MULTISET {"blue", "green", "yellow", "blue"}
```

The following multiset values are equal:

```
MULTISET {"blue", "green", "blue", "yellow"}
MULTISET {"blue", "green", "yellow", "blue"}
```

No more than 97 columns of the same table can be declared as MULTISET data types. (The same restriction applies to SET and LIST collection types.)

## Named ROW

See [ROW data type, Named on page 108.](#)

## NCHAR(n) data type

The NCHAR data type stores fixed-length character data. The data can be a string of single-byte or multibyte letters, digits, and other symbols that are supported by the code set of the database locale.

The main difference between CHAR and NCHAR data types is the collating order.

The collation order of the CHAR data type follows the code-set order, but the collating order of the NCHAR data type can be a localized order, if **DB\_LOCALE** (or SET COLLATION) specifies a locale that defines a localized order for collation.

Size specifications in NCHAR data type declarations can be affected by the SQL\_LOGICAL\_CHAR configuration parameter that is described in the section [Logical Character Semantics in Character Type Declarations on page 119](#).

In databases that are created with the NLSCASE INSENSITIVE property, operations on NCHAR strings ignore letter case, ordering data values without respect to or preference for letter case. For example, the NCHAR string "IDS" might precede or follow "IdS" or "iDs" in the collated list that a query returns, depending on the order in which these data strings are retrieved, because all of the following NCHAR strings are treated as duplicate values:

```
"ids" "IDS" "iDs" "IDs" "IdS" "iDs" "iDS" "Ids"
```

## NUMERIC(p,s) data type

The NUMERIC data type is a synonym for fixed-point DECIMAL.

## NVARCHAR(m,r) data type

The NVARCHAR data type stores strings of varying lengths. The string can include digits, symbols, and both single-byte and (in some locales) multibyte characters.

The main difference between VARCHAR and NVARCHAR data types is the collation order. Collation of VARCHAR data follows code-set order, but NVARCHAR collation can be locale specific, if **DB\_LOCALE** (or SET COLLATION) has specified a locale that defines a localized order for collation. (The section [Collating VARCHAR Values on page 118](#) describes an exception.)

A column declared as NVARCHAR, without parentheses or parameters, has a maximum size of one byte, and a reserved size of zero.

The first parameter in NVARCHAR data type declarations can be affected by the SQL\_LOGICAL\_CHAR configuration parameter that is described in the section [Logical Character Semantics in Character Type Declarations on page 119](#).

No more than 195 columns of the same table can be NVARCHAR data types.

In databases that are created with the NLSCASE INSENSITIVE property, operations on NVARCHAR strings ignore letter case, ordering data values without respect to or preference for letter case. For example, the NVARCHAR string "IBM" might precede or follow "IbM" or "iBm" in the collated list that a query returns, depending on the order in which these data strings are retrieved, because all of the following NVARCHAR strings are treated as duplicate values:

```
"ibm" "IBM" "iBm" "IBm" "IbM" "iBM" "iBm" "Ibm"
```

## OPAQUE data types

An OPAQUE type is a data type for which you must provide information to the database server.

You must provide this information:

- A data structure for how the data values are stored on disk
- Support functions to determine how to convert between the disk storage format and the user format for data entry and display
- Secondary access methods that determine how the index on this data type is built, used, and manipulated
- User functions that use the data type
- A system catalog entry to register the OPAQUE type in the database

The internal structure of an OPAQUE type is not visible to the database server and can only be accessed through user-defined routines. Definitions for OPAQUE types are stored in the **sysxdtypes** system catalog table. These SQL statements maintain the definitions of OPAQUE types in the database:

- The CREATE OPAQUE TYPE statement registers a new OPAQUE type in the database.
- The DROP TYPE statement removes a previously defined OPAQUE type from the database.

For more information about the above-mentioned SQL statements, see the *HCL OneDB™ Guide to SQL: Syntax*. For information about how to create OPAQUE types and an example of an OPAQUE type, see *HCL OneDB™ User-Defined Routines and Data Types Developer's Guide*.

## REAL data type

The REAL data type is a synonym for SMALLFLOAT.

## ROW data type, Named

A named ROW data type must be declared with a name. This SQL identifier must be unique among data type names within the same database.

(An unnamed ROW type is a ROW type that contains fields but has no user-defined name.) Only named ROW types support data type inheritance. For more information, see [ROW Data Types on page 130](#).

### Defining named ROW types

You must declare and register in the database a new named ROW type by using the CREATE ROW TYPE statement of SQL. Definitions for named ROW types are stored in the **sysxdtypes** system catalog table.

The fields of a ROW data type can be any built-in data type or UDT, but TEXT or BYTE fields of a ROW type are valid in typed tables only. If you want to assign a ROW type to a column in the CREATE TABLE or ALTER TABLE statements, its elements cannot be TEXT or BYTE data types.

In general, the data type of a field of a ROW type can be any of these types:

- A built-in type (except for the TEXT or BYTE data types)
- A collection type (LIST, MULTISSET, or SET)
- A distinct type
- Another named or unnamed ROW type
- An opaque type

These SQL statements maintain the definitions of named ROW data types:

- The CREATE ROW TYPE statement adds a named ROW type to the database.
- The DROP ROW TYPE statement removes a previously defined named ROW type from the database.

No more than 195 columns of the same table can be named ROW types.

For details about these SQL syntax statements, see the *HCL OneDB™ Guide to SQL: Syntax*. For examples of how to create and use named ROW types, see the *HCL OneDB™ Database Design and Implementation Guide*.

## Equivalence and named ROW types

No two named ROW types can be equal, even if they have identical structures, because they have different names. For example, the following named ROW types have the same structure (the same number of fields and the same order of data types of fields within the row) but they are not equal:

```
name_t (lname CHAR(15), initial CHAR(1), fname CHAR(15))
emp_t (lname CHAR(15), initial CHAR(1), fname CHAR(15))
```

A Boolean equality condition like `name_t = emp_t` always evaluates to FALSE if both of the operands are different named ROW types.

## Named ROW types and inheritance

Named ROW types can be part of a type-inheritance hierarchy. One named ROW type can be the parent (or supertype) of another named ROW type. A subtype in a hierarchy inherits all the properties of its supertype. Type inheritance is explained in the CREATE ROW TYPE statement in the *HCL OneDB™ Guide to SQL: Syntax* and in the *HCL OneDB™ Database Design and Implementation Guide*.

## Typed tables

Tables that are part of an inheritance hierarchy must be typed tables. Typed tables are tables that have been assigned a named ROW type. For the syntax you use to create typed tables, see the CREATE TABLE statement in the *HCL OneDB™ Guide to SQL: Syntax*. Table inheritance and its relation to type inheritance is also explained in that section. For information about how to create and use typed tables, see the *HCL OneDB™ Database Design and Implementation Guide*.

## ROW data type, Unnamed

An unnamed ROW type contains fields but has no user-declared name. An unnamed ROW type is defined by its structure.

Two unnamed ROW types are equal if they have the same structure (meaning the ordered list of the data types of the fields). If two unnamed ROW types have the same number of fields, and if the order of the data type of each field in one ROW type matches the order of data types of the corresponding fields in the other ROW data type, then the two unnamed ROW data types are equal.

For example, the following unnamed ROW types are equal:

```
ROW (lname char(15), initial char(1) fname char(15))
ROW (dept char(15), rating char(1) name char(15))
```

The following ROW types have the same number of fields and the same data types, but are not equal, because their fields are not in the same order:

```
ROW (x integer, y varchar(20), z real)
ROW (x integer, z real, y varchar(20))
```

A field of an unnamed ROW type can be any of the following data types:

- A built-in type
- A collection type
- A distinct type
- Another ROW type
- An opaque type

Unnamed ROW types cannot be used in typed tables or in type inheritance hierarchies. For more information about unnamed ROW types, see the *HCL OneDB™ Guide to SQL: Syntax* and the *HCL OneDB™ Database Design and Implementation Guide*.

## Creating unnamed ROW types

### About this task

You can create an unnamed ROW type in several ways:

- You can declare an unnamed ROW type using the ROW keyword. Each field in a ROW can have a different field type. To specify the field type, use the following syntax:

```
ROW (field_name field_type, ...)
```

The *field\_name* must conform to the rules for SQL identifiers. (See the Identifier section in the *HCL OneDB™ Guide to SQL: Syntax*.)

- To generate an unnamed ROW type, use the ROW keyword as a constructor with a series of values. A corresponding unnamed ROW type is created, using the default data types of the specified values.

For example, the following declaration:

```
ROW(1, 'abc', 5.30)
```

defines this unnamed ROW data type:

```
ROW (x INTEGER, y VARCHAR, z DECIMAL)
```

- You can create an unnamed ROW type by an implicit or explicit cast from a named ROW type or from another unnamed ROW type.
- The rows of any table (except a table defined on a named ROW type) are unnamed ROW types.

No more than 195 columns of the same table can be unnamed ROW types.

## Inserting Values into Unnamed ROW Type Columns

### About this task

When you specify field values for an unnamed ROW type, list the field values after the constructor and between parentheses. For example, suppose you have an unnamed ROW-type column. The following INSERT statement adds one group of field values to this ROW column:

```
INSERT INTO table1 VALUES (ROW(4, 'abc'))
```

You can specify a ROW column in the IN predicate in the WHERE clause of a SELECT statement to search for matching ROW values. For more information, see the Condition section in the *HCL OneDB™ Guide to SQL: Syntax*.

## SERIAL(n) data type

The SERIAL data type stores a sequential integer, of the INT data type, that is automatically assigned by the database server when a new row is inserted.

The default serial starting number is 1, but you can assign an initial value, *n*, when you create or alter the table.

- You must specify a positive number for the starting number.
- If you specify zero (0) for the starting number, the value that is used is the maximum positive value that already exists in the SERIAL column + 1.

The maximum value for SERIAL is 2,147,483,647. If you assign a number greater than 2,147,483,647, you receive a syntax error. Use the SERIAL8 or BIGSERIAL data type, rather than SERIAL, if you need a larger range.

A table can have no more than one SERIAL column, but it can have a SERIAL column and either a SERIAL8 column or a BIGSERIAL column.

SERIAL values in a column are not automatically unique. You must apply a unique index or primary key constraint to this column to prevent duplicate serial numbers. If you use the interactive schema editor in DB-Access to define the table, a unique index is applied automatically to a SERIAL column.

SERIAL numbers might not be consecutive, because of concurrent users, rollbacks, and other factors.

The DEFINE *variable* LIKE *column* syntax of SPL for indirect typing declares a variable of the INTEGER data type if *column* is a SERIAL data type.

After a number is assigned, it cannot be changed. You can insert a value into a SERIAL column (using the INSERT statement) or reset a serial column (using the ALTER TABLE statement), if the new value does not duplicate any existing value in the column. To insert into a SERIAL column, your database server increments by one the previous value (or the reset value, if that is larger) and assigns the result as the entered value. If ALTER TABLE has reset the next value of a SERIAL column to a value smaller than values already in that column, however, the next value follows this formula:

```
(maximum existing value in SERIAL column) + 1
```

For example, if you reset the serial value of **customer.customer\_num** to 50, when the largest existing value is 128, the next assigned number will be 129. For more details on SERIAL data entry, see the *HCL OneDB™ Guide to SQL: Syntax*.

A SERIAL column can store unique codes such as order, invoice, or customer numbers. SERIAL data values require four bytes of storage, and have the same precision as the INTEGER data type. For details of another way to assign unique whole numbers to each row of a database table, see the CREATE SEQUENCE statement in *HCL OneDB™ Guide to SQL: Syntax*.

## SERIAL8(n) data type

The SERIAL8 data type stores a sequential integer, of the INT8 data type, that is assigned automatically by the database server when a new row is inserted.

The SERIAL8 data type behaves like the SERIAL data type, but with a larger range. For more information about how to insert values into SERIAL8 columns, see the *HCL OneDB™ Guide to SQL: Syntax*.

A SERIAL8 data column is commonly used to store large, unique numeric codes such as order, invoice, or customer numbers. SERIAL8 data values have the same precision and storage requirements as INT8 values (page [INT8 on page 100](#)).

The default serial starting number is 1, but you can assign an initial value, *n*, when you create or alter the table.

- You must specify a positive number for the starting number.
- If you specify zero (0) for the starting number, the value that is used is the maximum positive value that already exists in the SERIAL8 column + 1.

A table can have no more than one SERIAL column, but it can have a SERIAL column and either a SERIAL8 column or a BIGSERIAL column.

SERIAL8 values in a column are not automatically unique. You must apply a unique index or primary key constraint to this column to prevent duplicate serial numbers. If you use the interactive schema editor in DB-Access to define the table, a unique index is applied automatically to a SERIAL8 column.

SERIAL8 numbers might not be consecutive, because of concurrent users, rollbacks, and other factors.

The DEFINE *variable* LIKE *column* syntax of SPL for indirect typing declares a variable of the INTEGER data type if *column* is a SERIAL8 data type.

For more information, see [Assigning a Starting Value for SERIAL8 on page 112](#). For information about using the SERIAL8 data type with the INT8 or BIGINT data type, see [Using SERIAL8 and BIGSERIAL with INT8 or BIGINT on page 88](#)

## Assigning a Starting Value for SERIAL8

The default serial starting number is 1, but you can assign an initial value, *n*, when you create or alter the table. To start the values at **1** in a SERIAL8 column of a table, give the value **0** for the SERIAL8 column when you insert rows into that table.

The database server will assign the value **1** to the SERIAL8 column of the first row of the table. The largest SERIAL8 value that you can assign is  $2^{63}-1$  (9,223,372,036,854,775,807). If you assign a value greater than this, you receive a syntax error.

When the database server generates a SERIAL8 value of this maximum number, it wraps around and starts generating values beginning at **1**.



After a nonzero SERIAL8 number is assigned, it cannot be changed. You can, however, insert a value into a SERIAL8 column (using the INSERT statement) or reset the SERIAL8 value  $n$  (using the ALTER TABLE statement), if that value does not duplicate any existing values in the column.

When you insert a number into a SERIAL8 column or reset the next value of a SERIAL8 column, your database server assigns the next number in sequence to the number entered. If you reset the next value of a SERIAL8 column to a value that is less than the values already in that column, however, the next value is computed using the following formula:

$$\text{maximum existing value in SERIAL8 column} + 1$$

For example, if you reset the SERIAL8 value of the **customer\_num** column in the **customer** table to 50, when the highest-assigned customer number is 128, the next customer number assigned is 129.

For information about using the SERIAL8 data type with the INT8 or BIGINT data type, see [Using SERIAL8 and BIGSERIAL with INT8 or BIGINT on page 88](#)

## SET(e) data type

The SET data type is an unordered collection type that stores unique elements

Duplicate element values are not valid as explained in *HCL OneDB™ Guide to SQL: Syntax*. (For a collection type that supports duplicate values, see the description of MULTISET in [MULTISET\(e\) data type on page 106](#).)

No more than 97 columns of the same table can be declared as SET data types. (The same restriction also applies to MULTISET and LIST collection types.)

The elements in a SET have no ordinal position. That is, no construct of a first, second, or third element in a SET exists. (For a collection type with ordinal positions for elements, see [LIST\(e\) data type on page 103](#).) All elements in a SET have the same element type. To specify the element type, use this syntax:

```
SET(element_type NOT NULL)
```

The *element\_type* of a collection can be any of the following types:

- A built-in type, except SERIAL, SERIAL8, BIGSERIAL, BYTE, and TEXT
- A named or unnamed ROW type
- Another collection type
- An opaque type

You must specify the NOT NULL constraint for SET elements. No other constraints are valid for SET columns. For more information about the syntax of the SET collection type, see the *HCL OneDB™ Guide to SQL: Syntax*.

You can use SET anywhere that you use any other data type, unless otherwise indicated. For example:

- After the IN predicate in the WHERE clause of a SELECT statement to search for matching SET values
- As an argument to the CARDINALITY or **mi\_collection\_card()** function to determine the number of elements in a SET column

SET values are not valid as arguments to an aggregate function such as AVG, MAX, MIN, or SUM. For more information, see the Condition and Expression sections in the *HCL OneDB™ Guide to SQL: Syntax*.

The following examples declare two sets. The first statement declares a set of integers and the second declares a set of character elements.

```
SET(INTEGER NOT NULL)
SET(CHAR(20) NOT NULL)
```

The following examples construct the same sets from value lists:

```
SET{1, 5, 13}
SET{"Oakland", "Menlo Park", "Portland", "Lenexa"}
```

In the following example, a SET constructor function is part of a CREATE TABLE statement:

```
CREATE TABLE tab
(
  c CHAR(5),
  s SET(INTEGER NOT NULL)
);
```

The following set values are equal:

```
SET{"blue", "green", "yellow"}
SET{"yellow", "blue", "green"}
```

## SMALLFLOAT

The SMALLFLOAT data type stores single-precision floating-point numbers with approximately nine significant digits.

SMALLFLOAT corresponds to the **float** data type in C. The range of values for a SMALLFLOAT data type is the same as the range of values for the C **float** data type on your computer.

A SMALLFLOAT data type column typically stores scientific numbers that can be calculated only approximately. Because floating-point numbers retain only their most significant digits, the number that you enter in this type of column and the number the database displays might differ slightly depending on how your computer stores floating-point numbers internally.

For example, you might enter a value of 1.1000001 in a SMALLFLOAT field and, after processing the SQL statement, the application might display this value as 1.1. This difference occurs when a value has more digits than the floating-point number can store. In this case, the value is stored in its approximate form with the least significant digits treated as zeros.

SMALLFLOAT data types usually require four bytes of storage. Conversion of a SMALLFLOAT value to a DECIMAL value results in nine digits of precision.

## SMALLINT data type

The SMALLINT data type stores small whole numbers that range from -32,767 to 32,767. The maximum negative number, -32,768, is a reserved value and cannot be used.

The SMALLINT value is stored as a signed binary integer.

Integer columns typically store counts, quantities, and so on. Because the SMALLINT data type requires only two bytes per value, arithmetic operations are performed efficiently. SMALLINT, however, stores only a limited range of values, compared to other built-in numeric data types. If a number is outside the range of the minimum and maximum SMALLINT values, the database server does not store the data value, but instead issues an error message.

## TEXT data type

The TEXT data type stores any kind of text data. It can contain both single-byte and multibyte characters that the locale supports. The term *simple large object* refers to an instance of a TEXT or BYTE data type.

A TEXT column has a theoretical limit of  $2^{31}$  bytes (two gigabytes) and a practical limit that your available disk storage determines. No more than 195 columns of the same table can be declared as TEXT data types. The same restriction also applies to BYTE data types.

You can store, retrieve, update, or delete the value in a TEXT column.

You can use TEXT operands in Boolean expressions only when you are testing for NULL values with the IS NULL or IS NOT NULL operators.

You can use the following methods, which can load rows or update fields, to insert TEXT data:

- With the **dbload** utility
- With the LOAD statement (DB-Access)
- From TEXT host variables (ESQL)

A built-in cast exists to convert TEXT objects to CLOB objects. For more information, see the *HCL OneDB™ Database Design and Implementation Guide*.

Strings of the TEXT data type are collated in code-set order. For more information about collating orders, see the *HCL OneDB™ GLS User's Guide*.

### Selecting data in a TEXT column

When you select a TEXT column, you can receive all or part of it. To retrieve it all, use the regular syntax for selecting a column. You can also select any part of a TEXT column by using subscripts, as this example shows:

```
SELECT cat_descr [1,75] FROM catalog WHERE catalog_num = 10001
```

The SELECT statement reads the first 75 bytes of the **cat\_descr** column associated with the **catalog\_num** value **10001**.

### Loading data into a TEXT column

You can use the LOAD statement to insert data into a table. For example, the `inp.txt` file contains the following information:

```
1|aaaaa|
2|bbbbb|
3|ccccc|
```

To load this data into the `blobtab` table use the following statement:

```
LOAD FROM inp.txt INSERT INTO blobtab;
```

## Limitations

You cannot use TEXT operands in arithmetic or string expressions, nor can you assign literals to TEXT columns in the SET clause of the UPDATE statement.

You also cannot use TEXT values in any of the following ways:

- With aggregate functions
- With the IN clause
- With the MATCHES or LIKE clauses
- With the GROUP BY clause
- With the ORDER BY clause

You cannot use a quoted text string, number, or any other actual value to insert or update TEXT columns.

You cannot use the MEDIUM or HIGH options of the UPDATE STATISTICS statement to calculate distribution statistics on TEXT columns.



**Important:** An error results if you try to return a TEXT column from a subquery, even if no TEXT column is used in a comparison condition or with the IN predicate.

## Nonprintable Characters in TEXT Values

TEXT columns typically store documents, program source files, and so on. In the default U.S. English locale, data objects of type TEXT can contain a combination of printable ASCII characters and the following control characters:

- Tab (CTRL-I)
- New line (CTRL-J)
- New page (CTRL-L)

Both printable and nonprintable characters can be inserted in text columns. HCL OneDB™ products do not do any checking of data values that are inserted in a column of the TEXT data type. (Applications might have difficulty, however, in displaying TEXT values that include non-printable characters.) For detailed information about entering and displaying nonprintable characters, see [Nonprintable Characters with CHAR on page 91](#).

## Unnamed ROW

See [ROW data type, Unnamed on page 109](#).

## VARCHAR(m,r) data type

The VARCHAR data type stores character strings of varying length that contain single-byte and (if the locale supports them) multibyte characters, where *m* is the maximum size (in bytes) of the column and *r* is the minimum number of bytes reserved for that column.

A column declared as VARCHAR without parentheses or parameters has a maximum size of one byte, and a reserved size of zero.

The VARCHAR data type is the HCL OneDB™ implementation of a character varying data type. The ANSI standard data type for varying-length character strings is CHARACTER VARYING.

The size of the maximum size (*m*) parameter of a VARCHAR column can range from 1 to 255 bytes. If you are placing an index on a VARCHAR column, the maximum size is 254 bytes. You can store character strings that are shorter, but not longer, than the *m* value that you specify.

Specifying the minimum reserved space (*r*) parameter is optional. This value can range from 0 to 255 bytes but must be less than the maximum size (*m*) of the VARCHAR column. If you do not specify any minimum value, it defaults to 0. You should specify this parameter when you initially intend to insert rows with short or NULL character strings in the column but later expect the data to be updated with longer values.

For variable-length strings longer than 255 bytes, you can use the LVARCHAR data type, whose upper limit is 32,739 bytes, instead of VARCHAR.

In an index based on a VARCHAR column (or on a NVARCHAR column), each index key has a length that is based on the data values that are actually entered, rather than on the declared maximum size of the column. (See, however, [IFX\\_PAD\\_VARCHAR environment variable on page 191](#) for information about how you can configure the effective size of VARCHAR and NVARCHAR data strings that HCL OneDB™ sends or receives.)

When you store a string in a VARCHAR column, only the actual data characters are stored. The database server does not strip a VARCHAR string of any user-entered trailing blanks, nor pad a VARCHAR value to the declared length of the column. If you specify a reserved space (*r*), but some data strings are shorter than *r* bytes, some space reserved for rows goes unused.

VARCHAR values are compared to other VARCHAR values (and to other character-string data types) in the same way that CHAR values are compared. The shorter value is padded on the right with blank spaces until the values have equal lengths; then they are compared for the full length.

No more than 195 columns of the same table can be VARCHAR data types.

### Nonprintable Characters with VARCHAR

Nonprintable VARCHAR characters are entered, displayed, and treated in the same way that nonprintable characters in CHAR values are treated. For details, see [Nonprintable Characters with CHAR on page 91](#).

### Storing Numeric Values in a VARCHAR Column

When you insert a numeric value in a VARCHAR column, the stored value does not get padded with trailing blanks to the maximum length of the column. The number of digits in a numeric VARCHAR value is the number of characters that are required to store that value. For example, in the next example, the value stored in table **mytab** is 1.

```
create table mytab (col1 varchar(10));
insert into mytab values (1);
```



**Tip:** VARCHAR treats C null (binary 0) and string terminators as termination characters for nonprintable characters.

In some East Asian locales, VARCHAR data types can store multibyte characters if the database locale supports a multibyte code set. If you store multibyte characters, make sure to calculate the number of bytes needed. For more information, see the *HCL OneDB™ GLS User's Guide*.

## Multibyte Characters with VARCHAR

The first parameter in VARCHAR data type declarations can be affected by the SQL\_LOGICAL\_CHAR feature that is described in the section [Logical Character Semantics in Character Type Declarations on page 119](#).

## Collating VARCHAR Values

The main difference between the NVARCHAR and the VARCHAR data types (like the difference between CHAR and NCHAR) is the difference in collating order. In general, collation of VARCHAR (like CHAR and LVARCHAR) values is in the order of the characters as they exist in the code set.

An exception is the MATCHES operator, which applies a localized collation to NVARCHAR and VARCHAR values (and to CHAR, LVARCHAR, and NCHAR values) if you use bracket ( [ ] ) symbols to define ranges when **DB\_LOCALE** (or SET COLLATION) has specified a localized collating order. For more information, see the *HCL OneDB™ GLS User's Guide*.

## Built-In Data Types

HCL OneDB™ supports the following built-in data types.

Category	Data Types
Character	CHAR, CHARACTER VARYING, LVARCHAR, NCHAR, NVARCHAR, VARCHAR, IDSSSECURITYLABEL
Large-object	Simple-large-object types: BYTE, TEXT Smart-large-object types: BLOB, CLOB
Logical	BOOLEAN
Multirepresentational	<a href="#">BSON and JSON built-in opaque data types on page</a>
Numeric	BIGINT, BIGSERIAL, DECIMAL, FLOAT, INT8, INTEGER, MONEY, SERIAL, SERIAL8, SMALLFLOAT, SMALLINT
Time	DATE, DATETIME, INTERVAL

## Character Data Types

The character data types store string values.

## Built-in Character Types

**Table 48. Attributes of built-in character data types**

	Size (in bytes)	Default	Reserved	Collation	Length
<b>CHAR(n)</b>	1 to 32,767	1 byte	None	Code set	Fixed
<b>NCHAR(n)</b>	1 to 32,767	1 byte	None	Localized	Fixed
<b>VARCHAR(m, r)</b>	1 to 255	0 for r	0 to 255 bytes	Code set	Variable
<b>NVARCHAR(m, r)</b>	1 to 255	0 for r	0 to 255 bytes	Localized	Variable
<b>LVARCHAR(m)</b>	1 to 32,739	2048 bytes	None	Code set	Variable

The database server also uses LVARCHAR to represent the external format of opaque data types. In I/O operations of the database server, LVARCHAR data values have no upper limit on their size, apart from file size restrictions or limits of your operating system or hardware resources.

### Logical Character Semantics in Character Type Declarations

HCL OneDB™ supports a configuration parameter, `SQL_LOGICAL_CHAR`, whose setting can instruct the SQL parser to interpret the maximum size of character columns in data type declarations of the `CREATE TABLE` or `ALTER TABLE` statements as logical characters, rather than in units of bytes.

When a database is created, the current `SQL_LOGICAL_CHAR` setting for the database server is recorded in the **systables** table of the system catalog. The feature has no effect on tables that are subsequently created or altered in the database if the setting is `OFF` or `1`.

In a database where the `SQL_LOGICAL_CHAR` setting is `ON` or is a digit between 2, 3, or 4, however, the SQL parser interprets explicit and implicit size declarations as logical characters in declarations of SPL variables and declarations of columns in database tables for the following character types:

- CHAR and CHARACTER
- CHARACTER VARYING and VARCHAR
- LVARCHAR
- NCHAR
- NVARCHAR
- DISTINCT types of the data types listed above
- DISTINCT types of those DISTINCT types
- ROW data type fields of the types listed above .
- LIST, MULTISSET, and SET elements of the types listed above.

This feature has no effect on the maximum storage size limits for the character types listed in the previous table. For databases that use a multibyte locale, however, it can reduce the risk of data truncation when a string is inserted into a character column or assigned to a character variable.

For example, if 4 is the SQL\_LOGICAL\_CHAR setting for the database, then a VARCHAR(10, 5) specification is interpreted as requesting a maximum of 40 bytes of storage, with 5 of these bytes reserved, creating a VARCHAR(40, 5) data type in standard SQL notation, rather than what was specified in the declaration.

The reserve size parameters of VARCHAR and NVARCHAR data types are not affected by the SQL\_LOGICAL\_CHAR setting, because the minimum size of a multibyte character is 1 byte. In this example, the minimum size of 5 multibyte characters is 5 bytes, a size that remains unchanged.

See the description of the SQL\_LOGICAL\_CHAR configuration parameter in the *HCL OneDB™ Administrator's Reference* for more information about the effect of the SQL\_LOGICAL\_CHAR setting in databases whose **DB\_LOCALE** specifies a multibyte locale. For additional information about multibyte locales and logical characters, see the *HCL OneDB™ GLS User's Guide*.

## IDSSECURITYLABEL

HCL OneDB™ also supports the IDSSECURITYLABEL data type for systems that implement label-based access control (LBAC). This built-in data type can be formally classified as a character type, because it is defined as a DISTINCT OF VARCHAR(128) data type, but only users who hold the **DBSECADM** role can declare this data type in DDL operations. It supports the LBAC security feature, rather than functioning as a general-purpose character type.

## Data Type Promotion

For some string-manipulation operations of HCL OneDB™, the five built-in character data types listed above support data type promotion, in order to reduce the risk of string operations failing because a returned string is too large to be stored in an NVARCHAR or VARCHAR column or program variable. See the topic "Return Types from CONCAT and String Manipulation Functions" in *HCL OneDB™ Guide to SQL: Syntax* for details of data type promotion among HCL OneDB™ character types.

## National Language Support

The NCHAR and NVARCHAR types are sometimes called National Language Support data types because of their support for localized collation. Because columns of type VARCHAR or NVARCHAR have no default size, you must specify a size (no greater than 255) in their declaration. For VARCHAR or NVARCHAR columns on which an index is defined, the maximum size is 254 bytes.

## NLSCASE INSENSITIVE Databases

In databases created with the `NLSCASE INSENSITIVE` keyword option, operations on data strings of the NCHAR or NVARCHAR types makes no distinction between uppercase and lowercase variants of the same letter. Rows are stored in NCHAR or NVARCHAR columns in the letter case in which characters were loaded, but data strings that consist of the same letters in the same sequence are evaluated as duplicates, even if the case of some letters is not identical. For example, the three NCHAR strings `"ABC"` and `"AbC"` and `"abC"` are treated as duplicates. Other built-in character types, including CHAR, NVARCHAR, and VARCHAR, follow the default case-sensitive rules, so that the same three strings in a CHAR column evaluate to distinct values.

Databases with the `NLSCASE INSENSITIVE` property also ignore the letter case of DISTINCT data types whose base types are NCHAR or NVARCHAR, as well as NCHAR or NVARCHAR fields in named or unnamed ROW types, and NCHAR or NVARCHAR elements of COLLECTION data types, including LIST, SET, or MULTISET.




In a database that is insensitive to the letter case of NCHAR or NVARCHAR values, string manipulation operations might produce unexpected results, if they implicitly cast CHAR, LVARCHAR, or VARCHAR operands or arguments to NCHAR or NVARCHAR data types. In some contexts, the operation can return a duplicate string, despite letter case variations that the database server would not have treated as duplicates for the original data types.

## Large-Object Data Types

A large object is a data object that is logically stored in a table column but physically stored independent of the column. Large objects are stored separate from the table because they typically store a large amount of data. Separation of this data from the table can increase performance.

Figure 3: Large-Object Data Types on page 121 shows the large-object data types.

Figure 3. Large-Object Data Types



Only HCL OneDB™ supports BLOB and CLOB data types.

For the relative advantages and disadvantages of simple and smart large objects, see the *HCL OneDB™ Database Design and Implementation Guide*.

## Simple Large Objects

Simple large objects are a category of large objects that have a theoretical size limit of  $2^{31}$  bytes and a practical limit that your disk capacity determines. HCL OneDB™ supports these simple-large-object data types:

### BYTE

Stores binary data. For more detailed information about this data type, see the description on page [BYTE data type on page 89](#).

### TEXT

Stores text data. For more detailed information about this data type, see the description on page [TEXT data type on page 115](#).

No more than 195 columns of the same table can be declared as BYTE or TEXT data types. Unlike smart large objects, simple large objects do not support random access to the data. When you transfer a simple large object between a client application and the database server, you must transfer the entire BYTE or TEXT value. If the data cannot fit into memory, you must store the data value in an operating-system file and then retrieve it from that file.

The database server stores simple large objects in *blobspaces*. A *blobspace* is a logical storage area that contains one or more chunks that only store BYTE and TEXT data. For information about how to define blobspaces, see your *HCL OneDB™ Administrator's Guide*.

## Smart large objects

Smart large objects are a category of large objects that support random access to the data, and that are generally recoverable.

The random access feature allows you to seek and read through the smart large object as if it were an operating-system file.

Smart large objects are also useful for opaque data types with large storage requirements. (See the description of opaque data types in [Opaque Data Types on page 131](#).) They have a theoretical size limit of  $2^{42}$  bytes and a practical limit that your disk capacity determines.

HCL OneDB™ supports the following smart-large-object data types:

### **BLOB**

Stores binary data. For more information about this data type, see the description on page [BLOB data type on page 88](#).

### **CLOB**

Stores text data. For more information about this data type, see [CLOB data type on page 92](#).

HCL OneDB™ stores smart large objects in *sbspaces*. An *sbspace* is a logical storage area that contains one or more chunks that store only BLOB and CLOB data. For information about how to define sbspaces, see your *HCL OneDB™ Performance Guide*.

When you define a BLOB or CLOB column, you can determine the following large-object characteristics:

- LOG and NOLOG: whether the database server should log the smart large object in accordance with the current database logging mode.
- KEEP ACCESS TIME and NO KEEP ACCESS TIME: whether the database server should keep track of the last time the smart large object was accessed.
- HIGH INTEG and MODERATE INTEG: whether the database server should use sbspace page headers and page footers to detect data corruption (HIGH INTEG), or only use page headers (MODERATE INTEG).

Use of these characteristics can affect performance. For information, see your *HCL OneDB™ Performance Guide*.

When an SQL statement accesses a smart-large-object, the database server does not send the actual BLOB or CLOB data. Instead, it establishes a pointer to the data and returns this pointer. The client application can then use this pointer in open, read, or write operations on the smart large object.

To access a BLOB or CLOB column from within a client application, use one of the following application programming interfaces (APIs):

- From within programs, use the smart-large-object API. (For more information, see the *HCL OneDB™ ESQL/C Programmer's Manual*.)
- From within a DataBlade® module, use the Client and Server API. (For more information, see the *HCL OneDB™ DataBlade® API Programmer's Guide*.)

For information about smart large objects, see the *HCL OneDB™ Guide to SQL: Syntax* and *HCL OneDB™ Database Design and Implementation Guide*.

## Time Data Types

DATE and DATETIME data values represent zero-dimensional points in time; INTERVAL data values represent 1-dimensional spans of time, with positive or negative values. DATE precision is always an integer count of days, but various field qualifiers can define the DATETIME and INTERVAL precision. You can use DATE, DATETIME, and INTERVAL data in arithmetic and relational expressions. You can manipulate a DATETIME value with another DATETIME value, an INTERVAL value, the current time (specified by the keyword CURRENT), or some unit of time (using the keyword UNITS).

You can use a DATE value in most contexts where a DATETIME value is valid, and vice versa. You also can use an INTERVAL operand in arithmetic operations where a DATETIME value is valid. In addition, you can add two INTERVAL values and multiply or divide an INTERVAL value by a number.

An INTERVAL column can hold a value that represents the difference between two DATETIME values or the difference between (or sum of) two INTERVAL values. In either case, the result is a span of time, which is an INTERVAL value. Conversely, if you add or subtract an INTERVAL from a DATETIME value, another DATETIME value is produced, because the result is a specific time.

[Table 49: Arithmetic Operations on DATE, DATETIME, and INTERVAL Values on page 123](#) lists the binary arithmetic operations that you can perform on DATE, DATETIME, and INTERVAL operands, and the data type that is returned by the arithmetic expression.

**Table 49. Arithmetic Operations on DATE, DATETIME, and INTERVAL Values**

Operand 1	Operator	Operand 2	Result
DATE	-	DATETIME	INTERVAL
DATETIME	-	DATE	INTERVAL
DATE	+ or -	INTERVAL	DATETIME
DATETIME	-	DATETIME	INTERVAL
DATETIME	+ or -	INTERVAL	DATETIME
INTERVAL	+	DATETIME	DATETIME
INTERVAL	+ or -	INTERVAL	INTERVAL
DATETIME	-	CURRENT	INTERVAL
CURRENT	-	DATETIME	INTERVAL
INTERVAL	+	CURRENT	DATETIME
CURRENT	+ or -	INTERVAL	DATETIME
DATETIME	+ or -	UNITS	DATETIME

**Table 49. Arithmetic Operations on DATE, DATETIME, and INTERVAL Values (continued)**

Operand 1	Operator	Operand 2	Result
INTERVAL	+ or -	UNITS	INTERVAL
INTERVAL	* or /	NUMBER	INTERVAL

No other combinations are allowed. You cannot add two DATETIME values because this operation does not produce either a specific time or a span of time. For example, you cannot add December 25 and January 1, but you can subtract one from the other to find the time span between them.

## Manipulating DATETIME Values

You can subtract most DATETIME values from each other.

Dates can be in any order and the result is either a positive or a negative INTERVAL value. The first DATETIME value determines the precision of the result, which includes the same time units as the first operand.

If the second DATETIME value has fewer fields than the first, the precision of the second operand is increased automatically to match the first.

In the following example, subtracting the DATETIME YEAR TO HOUR value from the DATETIME YEAR TO MINUTE value results in a positive interval value of 60 days, 1 hour, and 30 minutes. Because minutes were not included in the second operand, the database server sets the minutes value for the second operand to 0 before performing the subtraction.

```
DATETIME (2003-9-30 12:30) YEAR TO MINUTE
- DATETIME (2003-8-1 11) YEAR TO HOUR

Result: INTERVAL (60 01:30) DAY TO MINUTE
```

If the second DATETIME operand has more fields than the first (regardless of whether the precision of the extra fields is larger or smaller than those in the first operand), the additional time unit fields in the second value are ignored in the calculation.

In the next expression (and its result), the year is not included for the second operand. Therefore, the year is set automatically to the current year (from the system clock-calendar), in this example 2005, and the resulting INTERVAL is negative, which indicates that the second date is later than the first.

```
DATETIME (2005-9-30) YEAR TO DAY
- DATETIME (10-1) MONTH TO DAY

Result: INTERVAL (-1) DAY TO DAY [assuming that the current
                                year is 2005]
```

You can compare two DATETIME values by using the `mi_datetime_compare()` function.

## Manipulating DATETIME with INTERVAL Values

INTERVAL values can be added to or subtracted from DATETIME values. In either case, the result is a DATETIME value. If you are adding an INTERVAL value to a DATETIME value, the order of values is unimportant; however, if you are subtracting, the

DATETIME value must come first. Adding or subtracting a positive INTERVAL value moves the DATETIME result forward or backward in time. The expression shown in the following example moves the date ahead by three years and five months:

```
DATETIME (2000-8-1) YEAR TO DAY
+ INTERVAL (3-5) YEAR TO MONTH

Result: DATETIME (2004-01-01) YEAR TO DAY
```



**Important:** Evaluate the logic of your addition or subtraction. Remember that months can have 28, 29, 30, or 31 days and that years can have 365 or 366 days.

In most situations, the database server automatically adjusts the calculation when the operands do not have the same precision. In certain contexts, however, you must explicitly adjust the precision of one value to perform the calculation. If the INTERVAL value you are adding or subtracting has fields that are not included in the DATETIME value, you must use the EXTEND function to increase the precision of the DATETIME value. (For more information about the EXTEND function, see the Expression segment in the *HCL OneDB™ Guide to SQL: Syntax*.)

For example, you cannot subtract an INTERVAL MINUTE TO MINUTE value from the DATETIME value in the previous example that has a YEAR TO DAY field qualifier. You can, however, use the EXTEND function to perform this calculation, as the following example shows:

```
EXTEND (DATETIME (2008-8-1) YEAR TO DAY, YEAR TO MINUTE)
- INTERVAL (720) MINUTE(3) TO MINUTE

Result: DATETIME (2008-07-31 12:00) YEAR TO MINUTE
```

The EXTEND function allows you to explicitly increase the DATETIME precision from YEAR TO DAY to YEAR TO MINUTE. This allows the database server to perform the calculation, with the resulting extended precision of YEAR TO MINUTE.

## Manipulating DATE with DATETIME and INTERVAL Values

You can use DATE operands in some arithmetic expressions with DATETIME or INTERVAL operands by writing expressions to do the manipulating, as [Table 50: Results of Expressions That Manipulate DATE with DATETIME or INTERVAL Values on page 125](#) shows.

**Table 50. Results of Expressions That Manipulate DATE with DATETIME or INTERVAL Values**

Expression	Result
DATE – DATETIME	INTERVAL
DATETIME – DATE	INTERVAL
DATE + or – INTERVAL	DATETIME

In the cases that [Table 50: Results of Expressions That Manipulate DATE with DATETIME or INTERVAL Values on page 125](#) shows, DATE values are first converted to their corresponding DATETIME equivalents, and then the expression is evaluated by the rules of arithmetic.

Although you can interchange DATE and DATETIME values in many situations, you must indicate whether a value is a DATE or a DATETIME data type. A DATE value can come from the following sources:

- A column or program variable of type DATE
- The TODAY keyword
- The DATE( ) function
- The MDY function
- A DATE literal

A DATETIME value can come from the following sources:

- A column or program variable of type DATETIME
- The CURRENT keyword
- The EXTEND function
- A DATETIME literal

The database locale defines the default DATE and DATETIME formats. For the default locale, U.S. English, these formats are 'mm/dd/yy' for DATE values and 'yyyy-mm-dd hh:MM:ss' for DATETIME values.

To represent DATE and DATETIME values as character strings, the fields in the strings must be in the required order. In other words, when a DATE value is expected, the string must be in DATE format and when a DATETIME value is expected, the string must be in DATETIME format. For example, you can use the string 10/30/2008 as a DATE string but not as a DATETIME string. Instead, you must use 2008-10-30 or 08-10-30 as the DATETIME string.

In a nondefault locale, literal DATE and DATETIME strings must match the formats that the locale defines. For more information, see the *HCL OneDB™ GLS User's Guide*.

You can customize the DATE format that the database server expects with the **DBDATE** and **GL\_DATE** environment variables. You can customize the DATETIME format that the database server expects with the **DBTIME** and **GL\_DATETIME** environment variables. For more information, see [DBDATE environment variable on page 166](#) and [DBTIME environment variable on page 176](#). For more information about all these environment variables, see the *HCL OneDB™ GLS User's Guide*.

You can also subtract one DATE value from another DATE value, but the result is a positive or negative INTEGER count of days, rather than an INTERVAL value. If an INTERVAL value is required, you can either use the UNITS DAY operator to convert the INTEGER value into an INTERVAL DAY TO DAY value, or else use EXTEND to convert one of the DATE values into a DATETIME value before subtracting.

For example, the following expression uses the **DATE( )** function to convert character string constants to DATE values, calculates their difference, and then uses the UNITS DAY keywords to convert the INTEGER result into an INTERVAL value:

```
(DATE ('5/2/2007') - DATE ('4/6/1968')) UNITS DAY
```

```
Result: INTERVAL (12810) DAY(5) TO DAY
```



**Important:** Because of the high precedence of UNITS relative to other SQL operators, you should generally enclose any arithmetic expression that is the operand of UNITS within parentheses, as in the preceding example.

If you need YEAR TO MONTH precision, you can use the EXTEND function on the first DATE operand, as the following example shows:

```
EXTEND (DATE ('5/2/2007'), YEAR TO MONTH) - DATE ('4/6/1969')

Result: INTERVAL (39-01) YEAR TO MONTH
```

The resulting INTERVAL precision is YEAR TO MONTH, because the DATETIME value came first. If the DATE value had come first, the resulting INTERVAL precision would have been DAY(5) TO DAY.

## Manipulating INTERVAL Values

You can add or subtract INTERVAL values only if both values are from the same class; that is, if both are year-month or both are day-time.

In the following example, a SECOND TO FRACTION value is subtracted from a MINUTE TO FRACTION value:

```
INTERVAL (100:30.0005) MINUTE(3) TO FRACTION(4)
- INTERVAL (120.01) SECOND(3) TO FRACTION

Result: INTERVAL (98:29.9905) MINUTE TO FRACTION(4)
```

The use of numeric qualifiers alerts the database server that the MINUTE and FRACTION in the first value and the SECOND in the second value exceed the default number of digits.

When you add or subtract INTERVAL values, the second value cannot have a field with greater precision than the first. The second INTERVAL, however, can have a field of smaller precision than the first. For example, the second INTERVAL can be HOUR TO SECOND when the first is DAY TO HOUR. The additional fields (in this case MINUTE and SECOND) in the second INTERVAL value are ignored in the calculation.

You can compare two INTERVAL values by using the `mi_interval_compare()` function.

## Multiplying or Dividing INTERVAL Values

You can multiply or divide INTERVAL values by numbers. Any remainder from the calculation is ignored, however, and the result is truncated to the precision of the INTERVAL. The following expression multiplies an INTERVAL value by a literal number that has a fractional part:

```
INTERVAL (15:30.0002) MINUTE TO FRACTION(4) * 2.5

Result: INTERVAL (38:45.0005) MINUTE TO FRACTION(4)
```

In this example,  $15 * 2.5 = 37.5$  minutes,  $30 * 2.5 = 75$  seconds, and  $2 * 2.5 = 5$  FRACTION (4). The 0.5 minute is converted into 30 seconds and 60 seconds are converted into 1 minute, which produces the final result of 38 minutes, 45 seconds, and 0.0005 of a second. The result of any calculation has the same precision as the original INTERVAL operand.

## Extended Data Types

HCL OneDB™ enables you to create *extended data types* to characterize data that cannot easily be represented with the built-in data types. (You cannot, however, use extended data types in distributed transactions that query external tables.) You can create these categories of extended data types:

- Complex data types
- Distinct data types
- Opaque data types

Sections that follow provide an overview of each of these data types.


For more information about extended data types, see the *HCL OneDB™ Database Design and Implementation Guide* and *HCL OneDB™ User-Defined Routines and Data Types Developer's Guide*.

## Complex data types

A *complex data type* can store one or more values of other built-in or extended data types.

Figure 4: [Complex Data Types of HCL OneDB on page 128](#) shows the complex types that HCL OneDB™ supports.

Figure 4. Complex Data Types of HCL OneDB™



The following table summarizes the structure of the complex data types.

**Table 51. Collection types are complex data types that are made up of elements, each of which is of the same data type.**

Collection types	Description
LIST	A group of ordered elements, each of which need not be unique within the group.
MULTISET	A group of elements, each of which need not be unique. The order of the elements is ignored.
SET	A group of elements, each of which is unique. The order of the elements is ignored.

**Table 52. ROW types are complex data types that are made up of fields.**

ROW types	Description
Named ROW type	Row types that are identified by their name.
Unnamed ROW type	Row types that are identified by their structure.



Complex data types can be nested. For example, you can construct a ROW type whose fields include one or more sets, multisets, ROW types, and lists. Likewise, a collection type can have elements whose data type is a ROW type or a collection type.

Complex types that include opaque types inherit the following support functions.

- **input**
- **output**
- **send**
- **recv**
- **import**
- **export**
- **import\_binary**
- **export\_binary**
- **assign**
- **destroy**
- **LO\_handles**
- **hash**
- **lessthan**
- **equal**
- **lessthan** (for ROW types only)

The topics that follow summarize the complex data types. For more information, see the *HCL OneDB™ Database Design and Implementation Guide*.

## Collection Data Types

A collection data type is a complex type that is made up of one or more elements, all of the same data type. A collection element can be of any data type (including other complex types) except BYTE, TEXT, SERIAL, SERIAL8, or BIGSERIAL.



**Important:** An element cannot have a *NULL* value. You must specify the *NOT NULL* constraint for collection elements. No other constraints are valid for collections.

HCL OneDB™ supports three kinds of built-in collection types: LIST, SET, and MULTiset. The keywords used to declare these collections are the names of the *type constructors* or just *constructors*. For the syntax of collection types, see the *HCL OneDB™ Guide to SQL: Syntax*. No more than 97 columns of the same table can be declared as collection data types.

When you specify element values for a collection, list the element values after the constructor and between braces ({}). For example, suppose you have a collection column with the following MULTiset data type:

```
CREATE TABLE table1
(
  mset_col MULTiset(INTEGER NOT NULL)
)
```

The next INSERT statement adds one group of element values to this column. (The word MULTiset in these two examples is the MULTiset constructor.)

```
INSERT INTO table1 VALUES (MULTiset{5, 9, 7, 5})
```

You can leave the braces empty to indicate an empty set:

```
INSERT INTO table1 VALUE (MULTiset{})
```

An empty collection is not equivalent to a NULL value for the column.

## Accessing collection data

### About this task

To access the elements of a collection column, you must fetch the collection into a collection variable and modify the contents of the collection variable. Collection variables can be either of the following types:

- Variables in an SPL routine

For more information, see the *HCL OneDB™ Guide to SQL: Tutorial*.

- Host variables in programs

For more information, see the *HCL OneDB™ ESQL/C Programmer's Manual*.

You can also use nested dot notation to access collection data. For more about accessing elements of a collection, see the *HCL OneDB™ Guide to SQL: Tutorial*.



**Important:** Collection data types are not valid as arguments to functions that are used for functional indexes.

## ROW Data Types

A ROW data type is an ordered collection of one or more elements, called *fields*. Each field has a name and a data type. The fields of a ROW are comparable to the columns of a table, but with important differences:

- A field has no default clause.
- You cannot define constraints on a field.
- You can only use fields with row types, not with tables.

Two kinds of ROW data types exist:

- *Named ROW data types* are identified by their names.
- *Unnamed ROW data types* are identified by their structure.

The *structure* of an unnamed ROW data type is the number (and the order of data types) of its fields.

No more than 195 columns of the same table can be declared as ROW data types. For more information about ROW data types, see [ROW data type, Named on page 108](#) and [ROW data type, Unnamed on page 109](#).

You can cast between named and unnamed ROW data types; this is described in the *HCL OneDB™ Database Design and Implementation Guide*.

## Distinct Data Types

A distinct data type has the same internal structure as some other source data type in the database. The source type can be a built-in or extended data type. What distinguishes a distinct type from its source type are support functions that are defined on the distinct type.

No more than approximately 97 columns of the same table can be DISTINCT of collection data types (SET, LIST, and MULTISSET). No more than approximately 195 columns of the same table can be DISTINCT types that are based on BYTE, TEXT, ROW, LVARCHAR, NVARCHAR, or VARCHAR source types. (Here 195 columns is an approximate lower limit that applies to platforms with a 2 Kb base page size. For platforms with a base page size of 4 Kb, such as Windows™ and AIX® systems, the upper limit is approximately 450 columns of these data types.) For more information, see the section [DISTINCT data types on page 98](#). See also *HCL OneDB™ User-Defined Routines and Data Types Developer's Guide*.

## Opaque Data Types

An opaque data type is a user-defined or built-in data type that is fully encapsulated. The internal structure of an opaque data type is unknown to the database server.

Except for user-defined types (UDTs) that are DISTINCT of built-in non-opaque types, UDTs whose source types are built-in types are opaque data types. Similarly, UDTs that are DISTINCT of built-in opaque types are opaque types.

### Built-in opaque data types

The built-in data types BLOB, BOOLEAN, CLOB, BSON, JSON, and LVARCHAR are implemented as opaque data types. You can access all of these in other databases of the same HCL OneDB™ instance, but you cannot access the BLOB or CLOB built-in opaque data types in cross-server distributed operations.

UDTs that are DISTINCT of built-in opaque types and that are cast to built-in types are valid in cross-server queries and other DML operations, but all the casts and all the DISTINCT OF definitions for the UDTs must be identical in every participating database.

Several system catalog tables, whose schema are shown in [Structure of the System Catalog on page 11](#), have columns of built-in opaque data types. For information on how the system catalog encodes columns of built-in opaque data types, see [SYSCOLUMNS on page 24](#).

### User-defined opaque data types

You must provide the following information to the database server for an opaque data type:

- A data structure for how the data values are stored on disk
- Support functions to determine how to convert between the disk storage format and the user format for data entry and display
- Secondary access methods that determine how the index on this data type is built, used, and manipulated
- User functions that use the data type
- A system catalog entry to register the opaque type in the database

The internal structure of an opaque type is not visible to the database server and can only be accessed through user-defined routines. Definitions for opaque types are stored in the `sysxdtypes` system catalog table. These SQL statements maintain the definitions of opaque types in the database:

- The `CREATE OPAQUE TYPE` statement registers a new opaque type in the database.
- The `DROP TYPE` statement removes a previously defined opaque type from the database.

For more information, see the section [OPAQUE data types on page 107](#). See also *HCL OneDB™ User-Defined Routines and Data Types Developer's Guide*.

## Data Type Casting and Conversion

### About this task

Occasionally, the data type that was assigned to a column with the `CREATE TABLE` statement is inappropriate. You can change the data type of a column when you are required to store larger values than the current data type can accommodate. The database server allows you to change the data type of the column or to cast its values to a different data type with either of the following methods:

- Use the `ALTER TABLE` statement to modify the data type of a column.

For example, if you create a `SMALLINT` column and later find that you must store integers larger than 32,767, you must change the data type of that column to store the larger value. You can use `ALTER TABLE` to change the data type to `INTEGER`. The conversion changes the data type of all values that currently exist in the column and any new values that might be added.

- Use the `CAST AS` keywords or the double colon (`::`) cast operator to cast a value to a different data type.

Casting does not permanently alter the data type of a value; it expresses the value in a more convenient form. Casting user-defined data types into built-in types allows client programs to manipulate data types without knowledge of their internal structure.

If you change data types, the new data type must be able to store all of the old value.

Both data-type conversion and casting depend on casts registered in the `syscasts` system catalog table. For information about `syscasts`, see [SYSCASTS on page 20](#).

A cast is either built-in or user defined. Guidelines exist for casting distinct and extended data types. For more information about casting opaque data types, see *HCL OneDB™ User-Defined Routines and Data Types Developer's Guide*. For information about casting other extended data types see, the *HCL OneDB™ Database Design and Implementation Guide*.

## Using Built-in Casts

User **informix** owns built-in casts. They govern conversions from one built-in data type to another. Built-in casts allow the database server to attempt the following data-type conversions:

- A character type to any other character type
- A character type to or from another built-in type
- A numeric type to any other numeric type

The database server automatically invokes appropriate built-in casts when required. For time data types, conversion between DATE and DATETIME data types requires explicit casts with the EXTEND function, and explicit casts with the UNITS operator are required for number-to-INTERVAL conversion. Built-in casts are unavailable for converting large (BYTE, BLOB, CLOB, and TEXT) built-in types to other built-in data types.

When you convert a column from one built-in data type to another, the database server applies the appropriate built-in casts to each value already in the column. If the new data type cannot store any of the resulting values, the ALTER TABLE statement fails.

For example, if you try to convert a column from the INTEGER data type to the SMALLINT data type and the following values exist in the INTEGER column, the database server does not change the data type, because SMALLINT columns cannot accommodate numbers greater than 32,767:

```
100    400    700    50000    700
```

The same situation might occur if you attempt to transfer data from FLOAT or SMALLFLOAT columns to INTEGER, SMALLINT, or DECIMAL columns. Errors of overflow, underflow, or truncation can occur during data type conversion.

Sections that follow describe database server behavior during certain types of casts and conversions.

## Converting from number to number

When you convert data from one number data type to another, you occasionally find rounding errors.

The following table indicates which numeric data type conversions are acceptable and what kinds of errors you can encounter when you convert between certain numeric data types. In the table, the following codes are used:

**OK**

No error

**P**

An error can occur, depending on the precision of the decimal

**E**

An error can occur, depending on the data value

**D**

No error, but less significant digits might be lost

**Table 53. Acceptable conversions and possible errors**

Target Type	SMALL INT	INTEGER	INT8	SMALL FLOAT	FLOAT	DECIMAL
SMALLINT	OK	OK	OK	OK	OK	OK
INTEGER	E	OK	OK	E	OK	P
INT8	E	E	OK	D	E	P
SMALLFLOAT	E	E	E	OK	OK	P
FLOAT	E	E	E	D	OK	P
DECIMAL	E	E	E	D	D	P

For example, if you convert a FLOAT value to DECIMAL(4,2), your database server rounds off the floating-point number before storing it as DECIMAL.

This conversion can result in an error depending on the precision assigned to the DECIMAL column.

### Converting Between Number and Character

You can convert a character column (of a data type such as CHAR, NCHAR, NVARCHAR, or VARCHAR) to a numeric column. If a data string, however, contains any characters that are not valid in a number column (for example, the letter *l* instead of the number *1*), the database server returns an error.

You can also convert a numeric column to a character column. If the character column is not large enough to receive the number, however, the database server generates an error. If the database server generates an error, it cannot complete the ALTER TABLE statement or cast, and leaves the column values as characters. You receive an error message and the statement is rolled back automatically (regardless of whether you are in a transaction).

### Converting Between INTEGER and DATE

You can convert an integer column (SMALLINT, INTEGER, or INT8) to a DATE value. The database server interprets the integer as a value in the internal format of the DATE column. You can also convert a DATE column to an integer column. The database server stores the internal format of the DATE column as an integer representing a Julian date.

## Converting Between DATE and DATETIME

You can convert DATE columns to DATETIME columns. If the DATETIME column contains more fields than the DATE column, however, the database server either ignores the fields or fills them with zeros. The illustrations in the following list show how these two data types are converted (assuming that the default date format is *mm/dd/yyyy*):

- If you convert DATE to DATETIME YEAR TO DAY, the database server converts the existing DATE values to DATETIME values. For example, the value 08/15/2002 becomes 2002-08-15.
- If you convert DATETIME YEAR TO DAY to the DATE format, the value 2002-08-15 becomes 08/15/2002.
- If you convert DATE to DATETIME YEAR TO SECOND, the database server converts existing DATE values to DATETIME values and fills in the additional DATETIME fields with zeros. For example, 08/15/2002 becomes 2002-08-15 00:00:00.
- If you convert DATETIME YEAR TO SECOND to DATE, the database server converts existing DATETIME to DATE values but drops fields for time units smaller than DAY. For example, 2002-08-15 12:15:37 becomes 08/15/2002.

## Using User-Defined Casts

Implicit and explicit casts are owned by the users who create them. They govern casts and conversions between user-defined data types and other data types. Developers of user-defined data types must create certain implicit and explicit casts and the functions that are used to implement them. The casts allow user-defined types to be expressed in a form that clients can manipulate.

For information about how to register and use implicit and explicit casts, see the CREATE CAST statement in the *HCL OneDB™ Guide to SQL: Syntax* and the *HCL OneDB™ Database Design and Implementation Guide*.

### Implicit Casts

Implicit casts allow you to convert a user-defined data type to a built-in type or vice versa. The database server automatically invokes a single implicit cast when it must evaluate and compare expressions or pass arguments. Operations that require more than one implicit cast fail.

Users can explicitly invoke an implicit cast using the CAST AS keywords or the double colon (::) cast operator.

### Explicit Casts

Explicit casts, unlike implicit casts or built-in casts, are *never* invoked automatically by the database server. Users must invoke them explicitly with the CAST AS keywords or with the double colon (::) cast operator.

## Determining Which Cast to Apply

The database server uses the following rules to determine which cast to apply in a particular situation:

- To compare two built-in types, the database server automatically invokes the appropriate built-in casts.
- The database server applies only one implicit cast per operand. If two or more casts are required to convert the operand to the specified type, the user must explicitly invoke the additional casts.

In the following example, the literal value 5.55 is implicitly cast to DECIMAL, and is then explicitly cast to MONEY, and finally to yen:

```
CREATE DISTINCT TYPE yen AS MONEY
. . .
INSERT INTO currency_tab
VALUES (5.55::MONEY::yen)
```

- To compare a distinct type to its source type, the user must explicitly cast one type to the other.
- To compare a distinct type to a type other than its source, the database server looks for an implicit cast between the source type and the specified type.

If neither cast is registered, the user must invoke an explicit cast between the distinct type and the specified type. If this cast is not registered, the database server automatically invokes a cast from the source type to the specified type.

If none of these casts is defined, the comparison fails.

- To compare an opaque type to a built-in type, the user must explicitly cast the opaque type to a data type that the database server understands (such as LVARCHAR, SENDRECV, IMPEXP, or IMPEXPBIN). The database server then invokes built-in casts to convert the results to the specified built-in type.
- To compare two opaque types, the user must explicitly cast one opaque type to a form that the database server understands (such as LVARCHAR, SENDRECV, IMPEXP, or IMPEXPBIN) and then explicitly cast this type to the second opaque type.

For information about casting and the BOOLEAN, BSON, JSON, IMPEXP, IMPEXPBIN, LVARCHAR, and SENDRECV built-in opaque data types, see *HCL OneDB™ User-Defined Routines and Data Types Developer's Guide*.

## Casts for distinct types

You define a distinct type based on a built-in type or an existing opaque type or ROW type. Although data of the distinct type has the same length and alignment and is passed in the same way as data of the source type, the two cannot be compared directly. To compare a distinct type and its source type, you must explicitly cast one type to the other.

When you create a new distinct type, the database server automatically registers two explicit casts:

- A cast from the distinct type to its source type
- A cast from the source type to the distinct type

You can create an implicit cast between a distinct type and its source type. To create an implicit cast, however, you must first drop the default explicit cast between the distinct type and its source type.

You also can use all casts that have been registered for the source type without modification on the distinct type. You can also create and register new casts and support functions that apply only to the distinct type.

For examples that show how to create a cast function for a distinct type and register the function as cast, see the *HCL OneDB™ Database Design and Implementation Guide*.



**!** **Important:** For releases of HCL OneDB™ earlier than Version 9.21, distinct data types inherited the built-in casts that are provided for the source type. The built-in casts of the source type are not inherited by distinct data types in this release.

## What Extended Data Types Can Be Cast?

The next table shows the extended data type combinations that you can cast.

**Table 54. Extended data type combinations**

Target Type	Opaque Type	Distinct Type	Named ROW Type	Unnamed ROW Type	Collection Type	Built-in Type
<b>Opaque Type</b>	Explicit or implicit	Explicit	Explicit	Not Valid	Not Valid	Explicit or implicit <sup>3</sup>
<b>Distinct Type</b>	Explicit <sup>3</sup>	Explicit	Explicit	Not Valid	Not Valid	Explicit or implicit
<b>Named ROW Type</b>	Explicit <sup>3</sup>	Explicit	Explicit <sup>3</sup>	Explicit <sup>1</sup>	Not Valid	Not Valid
<b>Unnamed ROW Type</b>	Not Valid	Not Valid	Explicit <sup>1</sup>	Implicit <sup>1</sup>	Not Valid	Not Valid
<b>Collection Type</b>	Not Valid	Not Valid	Not Valid	Not Valid	Explicit <sup>2</sup>	Not Valid
<b>Built-in Type</b>	Explicit or implicit <sup>3</sup>	Explicit or implicit	Not Valid	Not Valid	Not Valid	System defined (implicit)

<sup>1</sup> Applies when two ROW types are structurally equivalent or casts exist to handle data conversions where corresponding field types are not the same.

<sup>2</sup> Applies when a cast exists to convert between the element types of the respective collection types.

<sup>3</sup> Applies when a user-defined cast exists to convert between the two data types.

The table shows only whether a cast between a source type and a target type are possible. In some cases, you must first create a user-defined cast before you can perform a conversion between two data types. In other cases, the database server provides either an implicit cast or a built-in cast that you must explicitly invoke.

## Operator Precedence

An *operator* is a symbol or keyword that can be in an SQL expression. Most SQL operators are restricted in the data types of their operands and returned values. Some operators only support operands of built-in data types; others can support built-in and extended data types as operands.

The following table shows the precedence of the operators that HCL OneDB™ supports, in descending (highest to lowest) order of precedence. Operators with the same precedence are listed in the same row.

Operator Precedence	Example in Expression
. ( <i>membership</i> ) [] ( <i>substring</i> )	<b>customer.phone</b> [1, 3]
UNITS	<b>x</b> UNITS DAY
+ - ( <i>unary</i> )	<b>- y</b>
:: ( <i>cast</i> )	NULL::TEXT
* /	<b>x / y</b>
+ - ( <i>binary</i> )	<b>x -y</b>
( <i>concatenation</i> )	<b>customer.fname    customer.lname</b>
ANY ALL SOME	<b>orders.ship_date &gt; SOME (SELECT paid_date FROM orders)</b>
NOT	NOT <b>y</b>
< <= = > >= != <>	<b>x &gt;= y</b>
IN BETWEEN ... AND LIKE MATCHES	<b>customer.fname MATCHES y</b>
AND	<b>x AND y</b>
OR	<b>x OR y</b>

See the *HCL OneDB™ Guide to SQL: Syntax* for the syntax and semantics of these SQL operators.

## Environment variables

Various *environment variables* affect the functionality of your HCL OneDB™ products. You can set environment variables that identify your terminal, specify the location of your software and define other parameters.

Some environment variables are required; others are optional. You must either set or accept the default setting for required environment variables.

These topics describe how to use the environment variables that apply to one or more HCL OneDB™ products and shows how to set them.

## Types of environment variables

Two types of environment variables are explained in this chapter:

- Environment variables that are specific to HCL OneDB™

Set HCL OneDB™ environment variables when you want to work with HCL OneDB™ products. Each HCL OneDB™ product publication specifies the environment variables that you must set to use that product.

- Environment variables that are used with a specific operating system

HCL OneDB™ products rely on the correct setting of certain standard operating system environment variables. For example, you must always set the **PATH** environment variable.

In a UNIX™ environment, you might also be required to set the TERMCAP or TERMINFO environment variable to use some products effectively.

The GLS environment variables that support nondefault locales are described in the *HCL OneDB™ GLS User's Guide*. The GLS variables are included in the list of environment variables in [Table 55: Uses for environment variables on page 147](#).

The database server uses the environment variables that were in effect at the time when the database server was initialized.

The `onstat -g env` command lists the active environment settings.



**Tip:** Additional environment variables that are specific to your client application or SQL API might be explained in the publication for that product.



**Important:** Do not set any environment variable in the home directory of user **informix** (or in the file `.informix` in that directory) while initializing the database and creating the **sysmaster** database.

## Limitations on environment variables

### Size of a block of environment variables

At the start of a session, the client groups all the environment variables that the server will use and sends the environment variables to the server as single block. The maximum size of this block is 32K. If the block of environment variables is greater than 32K, the error -1832 is returned to the application. The text of this error is "Environment block is greater than 32K."

To resolve this error, you can either unset one or more environment variables or reduce the size of some of the environment variables.

### Using environment variables on UNIX™

You can set, unset, modify, and view environment variables. If you already use any HCL OneDB™ products, some or all of the appropriate environment variables might be set.

You can set environment variables on UNIX™ in the following places:

- At the system prompt on the command line

When you set an environment variable at the system prompt, you must reassign it the next time you log in to the system.

- In an environment-configuration file

An environment-configuration file is a common or private file where you can set all the environment variables that HCL OneDB™ products use. The use of such files reduces the number of environment variables that you must set at the command line or in a shell file.

- In a login file

Values of environment variables set in your `.login`, `.cshrc`, or `.profile` file are assigned automatically every time you log in to the system.

- In the SET ENVIRONMENT statement of SQL

Values of some environment variables can reset by the SET ENVIRONMENT statement. The scope of the new settings is generally the routine that executed the SET ENVIRONMENT statement, but it is the current session for the **OPTCOMPIND** environment variable of HCL OneDB™, as described in the section [OPTCOMPIND environment variable on page 208](#). For more information about these routines and on the SET ENVIRONMENT statement, see the *HCL OneDB™ Guide to SQL: Syntax*.

In C, you can set supported environment variables within an application with the `putenv()` system call and retrieve values with the `getenv()` system call, if your UNIX™ system supports these functions. For more information about `putenv()` and `getenv()`, see the *HCL OneDB™ ESQL/C Programmer's Manual* and your C documentation.

## Setting environment variables in a configuration file

### About this task

The common (shared) environment-configuration file that is provided with HCL OneDB™ products is located in **\$ONEDB\_HOME/etc/onedb.rc**. Permissions for this shared file must be set to `644`.

A user can override the system or shared environment variables by setting variables in a private environment-configuration file. This file must have all of the following characteristics:

- Stored in the user's home directory
- Named **.informix**
- Permissions set to readable by the user

An environment-configuration file can contain comment lines (preceded by the `#` comment indicator) and variable definition lines that set values (separated by blank spaces or tabs), as the following example shows:

```
# This is an example of an environment-configuration file
#
DBDATE DMY4-
#
# These are ESQL/C environment variable settings
#
ONEDB_C gcc
CPFIRST TRUE
```

You can use the **ENVIGNORE** environment variable, described in [ENVIGNORE environment variable \(UNIX\) on page 182](#), to override one or more entries in an environment-configuration file. Use the HCL OneDB™ **chkenv** utility, described in [Checking environment variables with the chkenv utility on page 143](#), to perform a sanity check on the contents of an environment-configuration file. The **chkenv** utility returns an error message if the file contains a bad environment variable or if the file is too large.

The first time you set an environment variable in a shell file or environment-configuration file, you must tell the shell process to read your entry before you work with your HCL OneDB™ product. If you use a C shell, **source** the file; if you use a Bourne or Korn shell, use a period (.) to execute the file.

## Setting environment variables at login time

### About this task

Add commands that set your environment variables to the appropriate login file:

#### For C shell

**.login** or **.cshrc**

#### For Bourne shell or Korn shell

**.profile**

## Syntax for setting environment variables

Use standard UNIX™ commands to set environment variables. The examples in the following table show how to set the ABCD environment variable to *value* for the C shell, Bourne shell, and Korn shell. The Korn shell also supports a shortcut, as the last row indicates. Environment variables are case-sensitive.

Shell	Command
C	<code>setenv ABCD value</code>
Bourne	<code>ABCD=value</code> <code>export ABCD</code>
Korn	<code>ABCD=value</code> <code>export ABCD</code>
Korn	<code>export ABCD=value</code>

The following diagram shows how the syntax for setting an environment variable is represented throughout this chapter. These diagrams indicate the setting for the C shell; for the Bourne or Korn shells, use the syntax illustrated in the preceding table.

```
setenvABCDvalue
```

## Unsetting environment variables

### About this task

To unset an environment variable, enter the following command.

Shell	Command
C	<code>unsetenv ABCD</code>
Bourne or Korn	<code>unset ABCD</code>

## Modifying an environment-variable setting

### About this task

Sometimes you must add information to an environment variable that is already set. For example, the **PATH** environment variable is always set on UNIX™. When you use HCL OneDB™ productd, you must add to the **PATH** setting the name of the directory where the executable files for the HCL OneDB™ products are stored.

In the following example, the **ONEDB\_HOME** is **/usr/informix**. (That is, during installation, the HCL OneDB™ products were installed in the **/usr/informix** directory.) The executable files are in the **bin** subdirectory, **/usr/informix/bin**. To add this directory to the front of the C shell **PATH** environment variable, use the following command:

```
setenv PATH /usr/informix/bin:$PATH
```

Rather than entering an explicit pathname, you can use the value of the **ONEDB\_HOME** environment variable (represented as **\$ONEDB\_HOME**), as the following example shows:

```
setenv ONEDB_HOME /usr/informix
setenv PATH $ONEDB_HOME/bin:$PATH
```

You might prefer to use this version to ensure that your **PATH** entry does not conflict with the search path that was set in **ONEDB\_HOME**, and so that you are not required to reset **PATH** whenever you change **ONEDB\_HOME**. If you set the **PATH** environment variable on the C shell command line, you might be required to include braces ( **{ }** ) with the existing **ONEDB\_HOME** and **PATH**, as the following command shows:

```
setenv PATH ${ONEDB_HOME}/bin:${PATH}
```

For more information about how to set and modify environment variables, see the publications for your operating system.

## Viewing your environment-variable settings

### About this task

After you install one or more HCL OneDB™ products, enter the following command at the system prompt to view your current environment settings.

UNIX™ version	Command
BSD UNIX™	env
UNIX™ System V	printenv

## Checking environment variables with the chkenv utility

### About this task

The **chkenv** utility checks the validity of shared or private environment-configuration files. It validates the names of the environment variables in the file, but not their values. Use **chkenv** to provide debugging information when you define, in an environment-configuration file, all the environment variables that your HCL OneDB™ products use.

```
chkenv [pathname] filename
```

#### **filename**

is the name of the environment-configuration file to be debugged.

#### **pathname**

is the full directory path in which the environment variable file is located.

File **\$ONEDB\_HOME/etc/onedb.rc** is the shared environment-configuration file. A private environment-configuration file is stored as **.informix** in the home directory of the user. If you specify no *pathname* for **chkenv**, the utility checks both the shared and private environment configuration files. If you provide a *pathname*, **chkenv** checks only the specified file.

Issue the following command to check the contents of the shared environment-configuration file:

```
chkenv onedb.rc
```

The **chkenv** utility returns an error message if it finds a bad environment-variable name in the file or if the file is too large. You can modify the file and rerun the utility to check the modified environment-variable names.

HCL OneDB™ products ignore all lines in the environment-configuration file, starting at the point of the error, if the **chkenv** utility returns the following message:

```
-33523 filename: Bad environment variable on line number.
```

If you want the product to ignore specified environment-variables in the file, you can also set the **ENVIGNORE** environment variable. For a discussion of the use and format of environment-configuration files and the **ENVIGNORE** environment variable, see page [ENVIGNORE environment variable \(UNIX\) on page 182](#).

## Rules of precedence for environment variables

When HCL OneDB™ products accesses an environment variable, normally the following rules of precedence apply:

1. Of highest precedence is the value that is defined in the environment (shell) by explicitly setting the value at the shell prompt.
2. The second highest precedence goes to the value that is defined in the private environment-configuration file in the home directory of the user (`~/.onedb`).
3. The next highest precedence goes to the value that is defined in the common environment-configuration file (`$ONEDB_HOME/etc/onedb.rc`).
4. The lowest precedence goes to the default value, if one exists.

For precedence information about GLS environment variables, see the *HCL OneDB™ GLS User's Guide*.



**Important:** If you set one or more environment variables before you start the database server, and you do not explicitly set the same environment variables for your client products, the clients will adopt the original settings.

## Using environment variables on Windows™

The following sections discuss setting, viewing, unsetting, and modifying environment variables for Windows™ applications.

### Where to set environment variables on Windows™

You can set environment variables in several places on Windows™, depending on which HCL OneDB™ application you use.

Environment variables can be set in several ways, as described in [Setting environment variables on Windows on page 144](#).

The SET ENVIRONMENT statement of SQL can set certain routine-specific environment options. For more information, see the description of SET ENVIRONMENT in the *HCL OneDB™ Guide to SQL: Syntax*.

To use client applications such as or the Schema Tools on Windows™ environment, use the Setnet32 utility to set environment variables. For information about the Setnet32 utility, see the *HCL OneDB™ Client Products Installation Guide* for your operating system.

In , you can set supported environment variables within an application with the `ifx_putenv()` function and retrieve values with the `ifx_getenv()` function, if your Windows™ system supports them. For more information about `ifx_putenv()` and `ifx_getenv()`, see the *HCL OneDB™ ESQL/C Programmer's Manual*.

## Setting environment variables on Windows™

### About this task

You can set environment variables for command-prompt utilities in the following ways:

- With the System applet in the Control Panel
- In a command-line session



## Using the system applet to change environment variables

The System applet provides a graphical interface to create, modify, and delete system-wide and user-specific variables. Environment variables that are set with the System applet are visible to all command-prompt sessions.

### About this task

#### To change environment variables with the System applet in the control panel

1. Double-click the System applet icon from the Control Panel window.
2. Click the Environment tab near the top of the window.

Two list boxes display System Environment Variables and User Environment Variables. System Environment Variables apply to an entire system, and User Environment Variables apply only to the sessions of the individual user.

3. To change the value of an existing variable, select that variable. The name of the variable and its current value are in the boxes at the bottom of the window.
4. To add a new variable, highlight an existing variable and type the new variable name in the box at the bottom of the window.
5. Next, enter the value for the new variable at the bottom of the window and click **Set**.
6. To delete a variable, select the variable and click **Delete**.

### Results



**Important:** In order to use the System applet to change System environment variables, you must belong to the Administrators group. For information about assigning users to groups, see your operating-system documentation.

## Using the command prompt to change environment variables

You can change the setting of an environment variable at a command prompt.

### About this task

The following diagram shows the syntax for setting an environment variable at a command prompt in Windows™.

```
setABCD=value
```

If no *value* is specified, the environment variable is unset, as if it did not exist.

To view your current settings after one or more HCL OneDB™ products are installed, enter the following command at the command prompt.

```
set
```

Sometimes you must add information to an environment variable that is already set. For example, the **PATH** environment variable is always set in Windows™ environments. When you use HCL OneDB™ products, you must add the name of the directory where the executable files for the HCL OneDB™ products are stored to the **PATH**.

In the following example, **ONEDB\_HOME** is `d:\informix` (that is, during installation, HCL OneDB™ products were installed in the `d:\informix` directory). The executable files are in the `bin` subdirectory, `d:\informix\bin`. To add this directory at the beginning of the **PATH** environment-variable value, use the following command:

```
set PATH=d:\informix\bin;%PATH%
```

Rather than entering an explicit pathname, you can use the value of the **ONEDB\_HOME** environment variable (represented as **%ONEDB\_HOME%**), as the following example shows:

```
set ONEDB_HOME=d:\informix
set PATH=%PATH%
```

You might prefer to use this version to ensure that your **PATH** entry does not contradict the search path that was set in **ONEDB\_HOME** and to avoid the requirement to reset **PATH** whenever you change **ONEDB\_HOME**.

For more information about setting and modifying environment variables, see your operating-system publications.

## Using `dbservername.cmd` to initialize a command-prompt environment

Each time that you open a Windows™ command prompt, it acts as an independent environment. Therefore, environment variables that you set within it are valid only for that particular command-prompt instance.

### About this task

For example, if you open one command window and set the variable, **ONEDB\_HOME**, and then open another command window and type `set` to check your environment, you will find that **ONEDB\_HOME** is not set in the new command-prompt session.

The database server installation program creates a batch file that you can use to configure command-prompt utilities, ensuring that your command-prompt environment is initialized correctly each time that you run a command-prompt session. The batch file, `dbservername.cmd`, is located in `%ONEDB_HOME%`, and is a plain text file that you can modify with any text editor. If you have more than one database server installed in `%ONEDB_HOME%`, there will be more than one batch file with the `.cmd` extension, each bearing the name of the database server with which it is associated.

To run `dbservername.cmd` from a command prompt, type `dbservername` or configure a command prompt so that it runs `dbservername.cmd` automatically at start.

## Rules of precedence for Windows™ environment variables

When HCL OneDB™ products access an environment variable, normally the following rules of precedence apply:

1. The setting in Setnet32 with the **Use my settings** box selected.
2. The setting in Setnet32 with the **Use my settings** box cleared.
3. The setting on the command line before running the application.
4. The setting in Windows™ as a user variable.
5. The setting in Windows™ as a system variable.
6. The lowest precedence goes to the default value.

An application examines the first five values as it starts. Unless otherwise stated, changing an environment variable after the application is running does not have any effect.

## Environment variables in HCL OneDB™ products

The topics that follow discuss (in alphabetic order) environment variables that HCL OneDB™ database server products and their utilities use.



**Important:** The descriptions of the following environment variables include the syntax for setting the environment variable on UNIX™. For a general description of how to set these environment variables on Windows™, see [Setting environment variables on Windows on page 144](#).

## Environment variable portal

This portal is an index of usage categories for HCL OneDB™ and UNIX™ environment variables. The portal contains links to the topics that describe the environment variables.

Because the following table contains a comprehensive list of categories with links to applicable topics. Some environment variables are applicable for more than one category.

**Table 55. Uses for environment variables**

Functional category	Environment variable
Abbreviated year values	Specify how to expand literal DATE and DATETIME values: <a href="#">DBCENTURY environment variable on page 163</a>
ANSI/ISO SQL compliance	Set the case of owner names: <a href="#">ANSIOWNER environment variable on page 157</a> Specify if you want to check for HCL OneDB™ extensions to ANSI-standard SQL syntax: <a href="#">DBANSIWARN environment variable on page 162</a> No default table or routine access privileges for PUBLIC in databases not created WITH LOG MODE ANSI: <a href="#">NODEFDAC environment variable on page 206</a>
archecker utility	Specify the full path name for the archecker configuration file: <a href="#">AC_CONFIG file environment variable on page</a>
Buffers	Manage the fetch buffer size: <a href="#">FET_BUF_SIZE environment variable on page 183</a> Manage the network size: <a href="#">IFX_NETBUF_SIZE environment variable on page 189</a> Manage the network pool size: <a href="#">IFX_NETBUF_PVTPOOL_SIZE environment variable (UNIX) on page 188</a> Manage the BYTE or TEXT data buffer: <a href="#">DBBLOBBUF environment variable on page 163</a>

**Table 55. Uses for environment variables (continued)**

Functional category	Environment variable
Cache	<p>Control the use of the shared-statement cache on a session: <a href="#">STMT_CACHE environment variable on page 217</a></p> <p>Set information for the Optical Subsystem: <a href="#">IONEDB_OPCACHE environment variable on page 199</a></p>
Client/server	<p>Specify the default database server: <a href="#">ONEDB_SERVER environment variable on page 199</a></p> <p>Specify where shared-memory segments are attached to the client process: <a href="#">ONEDB_SHMBASE environment variable (UNIX) on page 200</a></p> <p>Specify the stack size for a client process: <a href="#">ONEDB_STACKSIZE environment variable on page 201</a></p> <p>Specify locale information, including for the client and server: <a href="#">GLS-related environment variables on page</a></p>
Code-set conversion	<p>Specify locale and multibyte information: <a href="#">GLS-related environment variables on page</a></p> <p>Specify the location of the <b>concsm.cfg</b> file: <a href="#">ONEDB_CONCSMCFG environment variable</a></p> <p>Specify the filename or pathname of the C compiler: <a href="#">ONEDB_C environment variable (UNIX) on page 195</a></p> <p>Specify the pathname of the map file for C++ programs: <a href="#">ONEDB_CPPMAP environment variable on page 198</a></p> <p>Specify information for compiling multithreaded applications: <a href="#">THREADLIB environment variable (UNIX) on page 219</a></p>
Configuration	<p>Specify the name of the active that holds configuration parameters: <a href="#">ONCONFIG environment variable on page 207</a></p> <p>Ignore specified environment variable settings: <a href="#">ENVIGNORE environment variable (UNIX) on page 182</a></p> <p>Specify the default database server: <a href="#">ONEDB_SERVER environment variable on page 199</a></p> <p>Specify the dbspaces in which temporary tables are built: <a href="#">DBSPACETEMP environment variable on page 175</a></p>

**Table 55. Uses for environment variables (continued)**

Functional category	Environment variable
	<p>Manage query optimizer directives: <a href="#">IFX_DIRECTIVES</a> environment variable on page 185 and <a href="#">IFX_EXTDIRECTIVES</a> environment variable on page 185</p> <p>Modify the value of the OPTCOMPIND configuration parameter: <a href="#">OPTCOMPIND</a> environment variable on page 208</p> <p>Specify the query performance goal for the optimizer: <a href="#">OPT_GOAL</a> environment variable (UNIX) on page 209</p> <p>Specify the degree of parallelism that the database server uses: <a href="#">PDQPRIORITY</a> environment variable on page 210</p> <p>Specify the stack size that is applied to all client processes: <a href="#">ONEDB_STACKSIZE</a> environment variable on page 201</p>
Connecting	<p>Set the maximum number of <i>additional</i> connection attempts: <a href="#">CONNECT_RETRIES</a> environment variable on page 196</p> <p>Set connect time information: <a href="#">CONNECT_TIMEOUT</a> environment variable on page 197</p> <p>Specify the default database server to for connections: <a href="#">ONEDB_SERVER</a> environment variable on page 199</p> <p>Specify the location of connection information: <a href="#">ONEDB_SQLHOSTS</a> environment variable on page 200</p>
Connection Manager	<p>Specify the location of the Connection Manager configuration file: <a href="#">CMCONFIG</a> environment variable on page 158</p>
Data distributions	<p>Manage the amount of system disk space that the UPDATE STATISTICS statement can use: <a href="#">DBUPSPACE</a> environment variable on page 179</p>
Database locale	<p>Manage locale information: <a href="#">GLS-related environment variables</a> on page</p>
Database server	<p>Specify servers for connections: <a href="#">ONEDB_SERVER</a> environment variable on page 199</p> <p>Set the locale for file I/O: <a href="#">GLS-related environment variables</a> on page</p> <p>Specify the name of the active file that holds configuration parameters: <a href="#">ONCONFIG</a> environment variable on page 207</p> <p>Manage parallel sorting: <a href="#">PSORT_DBTEMP</a> environment variable on page 214 and <a href="#">PSORT_NPROCS</a> environment variable on page 215</p> <p>Manage parallelism: <a href="#">PDQPRIORITY</a> environment variable on page 210</p>

**Table 55. Uses for environment variables (continued)**

Functional category	Environment variable
	<p>Manage role separation: <a href="#">INF_ROLE_SEP</a> environment variable on page 202</p> <p>Manage shared memory: <a href="#">ONEDB_SHMBASE</a> environment variable (UNIX) on page 200</p> <p>Manage stack size: <a href="#">ONEDB_STACKSIZE</a> environment variable on page 201</p> <p>Manage temporary tables: <a href="#">DBSPACETEMP</a> environment variable on page 175, <a href="#">DBTEMP</a> environment variable on page 176, and <a href="#">PSORT_DBTEMP</a> environment variable on page 214</p> <p>Manage variable-length packets: <a href="#">IFX_PAD_VARCHAR</a> environment variable on page 191</p>
Date and time values, formats	<p>Manage date and time information: <a href="#">DBCENTURY</a> environment variable on page 163, <a href="#">DBDATE</a> environment variable on page 166, <a href="#">DBTIME</a> environment variable on page 176, <a href="#">GLS</a>-related environment variables on page <a href="#">GL_DATE</a> and <a href="#">GL_DATETIME</a>), <a href="#">USE_DTENV</a> environment variable on page <a href="#">TZ</a> environment variable on page 219</p>
DB-Access utility	<p>Manage the database server and DB-Access: <a href="#">DBANSIWARN</a> environment variable on page 162, <a href="#">DBDELIMITER</a> environment variable on page 168, <a href="#">DBEDIT</a> environment variable on page 168. <a href="#">DBFLTMASK</a> environment variable on page 169, <a href="#">DBPATH</a> environment variable on page 172, <a href="#">FET_BUF_SIZE</a> environment variable on page 183, <a href="#">ONEDB_SERVER</a> environment variable on page 199, <a href="#">ONEDB_TERM</a> environment variable (UNIX) on page 201, <a href="#">TERM</a> environment variable (UNIX) on page 218, <a href="#">TERMCAP</a> environment variable (UNIX) on page 218, and <a href="#">TERMINFO</a> environment variable (UNIX) on page 219</p>
dbexport utility	<p>Set the field delimiter: <a href="#">DBDELIMITER</a> environment variable on page 168</p>
dbinfo	<p>Specify if <code>dbinfo('dbspace', partnum)</code> raises error -727 or returns NULL for an invalid partnum: <a href="#">DBINFO_DBSPACE_RETURN_NULL_FOR_INVALID_PARTNUM</a> environment variable on page 169</p>
Delimited identifiers	<p>Set the field delimiter used with the dbexport utility and with the LOAD and UNLOAD statements: <a href="#">DBDELIMITER</a> environment variable on page 168</p>
Disk space	<p>Manage the amount of system disk space and memory that the UPDATE STATISTICS MEDIUM or HIGH statement can use: <a href="#">DBUPSPACE</a> environment variable on page 179</p>
Editor	<p>Specify the text editor to use with SQL statements and command files in DB-Access: <a href="#">DBEDIT</a> environment variable on page 168</p>

**Table 55. Uses for environment variables (continued)**

Functional category	Environment variable
Enterprise Replication	Specify information for Enterprise Replication: <a href="#">Enterprise Replication configuration parameter and environment variable reference on page</a>
ESQL/C	<p>Specify ANSI compliance: <a href="#">DBANSIWARN environment variable on page 162</a></p> <p>Specify the filename or pathname of the C compiler to use with ESQL/C: <a href="#">ONEDB_C environment variable (UNIX) on page 195</a></p> <p>Set delimited identifiers: <a href="#">DELIMIDENT environment variable on page 181</a></p> <p>Specify multibyte characters and locale information <a href="#">GLS-related environment variables on page</a> (<b>CLIENT_LOCALE, ESQLMF, and GL_USER</b>)</p> <p>Specify information for multithreaded applications: <a href="#">THREADLIB environment variable (UNIX) on page 219</a></p> <p>Specify the default compilation order: <a href="#">CPFIRST environment variable on page 158</a></p>
Executable programs	Specify the directories to search for executable programs: <a href="#">PATH environment variable on page 210</a>
Fetch buffer size	Set buffer size information: <a href="#">FET_BUF_SIZE environment variable on page 183</a>
Filenames: multibyte	<a href="#">GLS-related environment variables on page</a> ( <b>GLS8BITFSYS</b> )
Files: field delimiter	Set the field delimiter: <a href="#">DBDELIMITER environment variable on page 168</a>
Files: installation	Specify the directory that contains the subdirectories in which your product files are installed: <a href="#">ONEDB_HOME environment variable on page 198</a>
Files: locale	Specify locale information: <a href="#">GLS-related environment variables on page</a> ( <b>CLIENT_LOCALE, DB_LOCALE, and SERVER_LOCALE</b> )
Files: map for C++	Specify the pathname of the map file for C++ programs: <a href="#">ONEDB_CPPMAP environment variable on page 198</a>
Files: message	Specify the subdirectory of <b>\$ONEDB_HOME</b> or the pathname of the directory that contains the compiled message files that the database server uses: <a href="#">DBLANG environment variable on page 170</a>
Files: temporary	<a href="#">DBSPACETEMP environment variable on page 175</a>
Files: temporary	Specify a directory for temporary files: <a href="#">DBTEMP environment variable on page 176</a>
Files: temporary sorting	Specify the location of temporary files used for sorting: <a href="#">PSORT_DBTEMP environment variable on page 214</a>

**Table 55. Uses for environment variables (continued)**

Functional category	Environment variable
Files: <code>termcap</code> , <code>terminfo</code>	Specify terminal information: <a href="#">ONEDB_TERM</a> environment variable (UNIX) on page 201, <a href="#">TERM</a> environment variable (UNIX) on page 218, <a href="#">TERMCAP</a> environment variable (UNIX) on page 218, and <a href="#">TERMINFO</a> environment variable (UNIX) on page 219
Format: date and time	Define the format for date and time information: <a href="#">DBCENTURY</a> environment variable on page 163, <a href="#">DBDATE</a> environment variable on page 166, <a href="#">DBTIME</a> environment variable on page 176, GLS-related environment variables on page ( <a href="#">GL_DATE</a> and <a href="#">GL_DATETIME</a> ), The <a href="#">USE_DTENV</a> environment variable on page , and <a href="#">TZ</a> environment variable on page 219
Format: private-use characters	Set the display width for characters in Unicode Private-Use Area (PUA) ranges: <a href="#">GLS-related environment variables on page (IFX_PUA_DISPLAY_MAPPING)</a>
Format: money	Define the format for money information: <a href="#">DBMONEY</a> environment variable on page 171 and <a href="#">GLS-related environment variables on page</a>
High-availability clusters	Specify information for high-availability clusters <a href="#">IFX_SMX_TIMEOUT</a> environment variable on page 191 and <a href="#">IFX_SMX_TIMEOUT_RETRY</a> environment variable on page 192
High-Performance Loader	Specify information for the High-Performance Loader: <a href="#">DBONPLOAD</a> environment variable on page 171, , on page 211 <a href="#">PLOAD_LO_PATH</a> environment variable on page 212. and <a href="#">PLOAD_SHMBASE</a> environment variable on page 212
Identifiers	Specify field delimiters: <a href="#">DELIMIDENT</a> environment variable on page 181  Specify information for identifiers longer than 18 bytes: <a href="#">IFX_LONGID</a> environment variable on page 188  Specify information for multibyte characters: <a href="#">GLS-related environment variables on page (CLIENT_LOCALE and ESQLMF)</a>
HCL® OneDB® Storage Manager	Manage ISM: <a href="#">ISM_COMPRESSION</a> environment variable on page 203, <a href="#">ISM_DEBUG_LEVEL</a> environment variable on page 203, <a href="#">ISM_ENCRYPTION</a> environment variable on page 204, <a href="#">ISM_MAXLOGSIZE</a> environment variable on page 204, <a href="#">ISM_MAXLOGVERS</a> environment variable on page 204
Installation	Specify the directory that contains the subdirectories in which your product files are installed: <a href="#">ONEDB_HOME</a> environment variable on page 198  Specify which directories to search for executable programs: <a href="#">PATH</a> environment variable on page 210



**Table 55. Uses for environment variables (continued)**

Functional category	Environment variable
JDBC	Manage environment variables used with JDBC: <a href="#">HCL OneDB™ environment variables with the HCL OneDB™ JDBC Driver on page</a>
Language environment	Specify language and locale information: <a href="#">DBLANG environment variable on page 170</a> and <a href="#">GLS-related environment variables on page</a>
Libraries	Specify paths for libraries: <a href="#">LD_LIBRARY_PATH environment variable (UNIX) on page 205</a> , <a href="#">LIBPATH environment variable (UNIX) on page 206</a> , and <a href="#">SHLIB_PATH environment variable (UNIX) on page 216</a>
Locale	Define client, server, and database locale information: <a href="#">GLS-related environment variables on page</a> ( <b>CLIENT_LOCALE</b> , <b>DB_LOCALE</b> , and <b>SERVER_LOCALE</b> )
Lock mode	Set the default lock mode for database tables that are created without specifying the LOCKMODE PAGE or LOCKMODE ROW keywords: <a href="#">IFX_DEF_TABLE_LOCKMODE environment variable on page 184</a>
Long Identifiers	Specify information for identifiers longer than 18 bytes: <a href="#">IFX_LONGID environment variable on page 188</a>
Map file for C++	Specify the pathname of the map file for C++ programs: <a href="#">ONEDB_CPPMAP environment variable on page 198</a>
Message chaining	Enable or disable optimized message transfers (message chaining) for : <a href="#">OPTMSG environment variable on page 208</a>
Message files	Specify the directory that contains compiled message files: <a href="#">DBLANG environment variable on page 170</a>
Money format	Define the format for money information: <a href="#">DBMONEY environment variable on page 171</a> and <a href="#">GLS-related environment variables on page</a>
Multibyte characters	Specify information for multibyte characters: <a href="#">GLS-related environment variables on page</a> ( <b>CLIENT_LOCALE</b> , <b>DB_LOCALE</b> , <b>SERVER_LOCALE</b> , and <b>GL_USEGLU</b> )
Multibyte filter	Specify HCL OneDB™ ESQL/C multibyte filter information: <a href="#">GLS-related environment variables on page</a> ( <b>ESQLMF</b> )
Multithreaded applications	Specify information for compiling multithreaded applications: <a href="#">THREADLIB environment variable (UNIX) on page 219</a>
Network	Specify network information: <a href="#">DBPATH environment variable on page 172</a>
Nondefault locale	Define client, server, and database locale information: <a href="#">GLS-related environment variables on page</a> ( <b>CLIENT_LOCALE</b> , <b>DB_LOCALE</b> , and <b>SERVER_LOCALE</b> )
ON-Bar utility	Optimize the deduplication capabilities for storage managers: <a href="#">IFX_BAR_USE_DEDUP environment variable on page</a>

**Table 55. Uses for environment variables (continued)**

Functional category	Environment variable
	<p>Disable the ability to replicate, import, or export backup objects among TSM servers: <a href="#">IFX_TSM_OBINFO_OFF environment variable on page</a></p> <p>Specify information for ON-Bar to use with HCL® OneDB® Storage Manager : <a href="#">ISM_COMPRESSION environment variable on page 203</a>, <a href="#">ISM_DEBUG_LEVEL environment variable on page 203</a>, and <a href="#">ISM_ENCRYPTION environment variable on page 204</a></p>
ONCONFIG parameters	Specify the name of the file that holds configuration parameters: <a href="#">ONCONFIG environment variable on page 207</a>
oninit output (Windows™ only)	Specify a path and file for oninit output: <a href="#">ONINIT_STDOUT environment variable (Windows) on page 207</a>
Optical Subsystem	Specify the size of the memory cache: <a href="#">IONEDB_OPCACHE environment variable on page 199</a>
Optimization: directives	Manage query optimizer directives: <a href="#">IFX_DIRECTIVES environment variable on page 185</a> and <a href="#">IFX_EXTDIRECTIVES environment variable on page 185</a>
Optimization: message transfers	Enable or disable optimized message transfers (message chaining) for : <a href="#">OPTMSG environment variable on page 208</a>
Optimization: join method	Modify the value of the OPTCOMPIND configuration parameter: <a href="#">OPTCOMPIND environment variable on page 208</a>
Optimization: performance goal	Specify the query performance goal for the optimizer: <a href="#">OPT_GOAL environment variable (UNIX) on page 209</a>
OPTOFC feature	Enable optimize-OPEN-FETCH-CLOSE functionality: <a href="#">OPTOFC environment variable on page 209</a>
PAM authentication for MongoDB clients	Enable PAM authentication for MongoDB clients: <a href="#">IFMXMONGOAUTH environment variable on page 184</a>
Path name: <b>archecker</b> configuration file	Specify the full path name for the archecker configuration file: <a href="#">AC_CONFIG file environment variable on page</a>
Path name: C compiler	Specify the filename or pathname of the C compiler: <a href="#">ONEDB_C environment variable (UNIX) on page 195</a>
Path name: database files	Specify database server file and path information: <a href="#">DBPATH environment variable on page 172</a>
Path name: executable programs	Specify directories to search for executable programs: <a href="#">PATH environment variable on page 210</a>

**Table 55. Uses for environment variables (continued)**

Functional category	Environment variable
Path name: HPL smart-large-object handles	Specify the pathname for smart-large-object handles: <a href="#">PLOAD_LO_PATH environment variable on page 212</a>
Path name: installation	Specify the directory that contains the subdirectories in which your product files are installed: <a href="#">ONEDB_HOME environment variable on page 198</a>
Path name: libraries	Specify paths for libraries: <a href="#">LD_LIBRARY_PATH environment variable (UNIX) on page 205</a> , <a href="#">LIBPATH environment variable (UNIX) on page 206</a> , and <a href="#">SHLIB_PATH environment variable (UNIX) on page 216</a>
Path name: message files	Specify the directory that contains compiled message files: <a href="#">DBLANG environment variable on page 170</a> and <a href="#">GLS-related environment variables on page</a>
Path name: parallel sorting	Specify the location of temporary files for sorts: <a href="#">PSORT_DBTEMP environment variable on page 214</a>
HCL OneDB™ Primary Storage Manager	Manage the storage manager: <a href="#">PSM_ACT_LOG environment variable on page 212</a> , <a href="#">PSM_CATALOG_PATH environment variable on page 213</a> , <a href="#">PSM_DBS_POOL environment variable on page 213</a> , <a href="#">PSM_DEBUG environment variable on page 213</a> , <a href="#">PSM_DEBUG_LOG environment variable on page 214</a> , and <a href="#">PSM_LOG_POOL environment variable on page 214</a>
Preserve owner name case	Set the case of owner names: <a href="#">ANSIOWNER environment variable on page 157</a>
Printing	Specify the default printing program: <a href="#">DBPRINT environment variable on page 174</a>
Privileges	Configure role separation: <a href="#">INF_ROLE_SEP environment variable on page 202</a>
Query: optimization	<p>Manage query optimizer directives: <a href="#">IFX_DIRECTIVES environment variable on page 185</a> and <a href="#">IFX_EXTDIRECTIVES environment variable on page 185</a></p> <p>Modify the value of the OPTCOMPIND configuration parameter: <a href="#">OPTCOMPIND environment variable on page 208</a></p> <p>Specify the query performance goal for the optimizer: <a href="#">OPT_GOAL environment variable (UNIX) on page 209</a></p> <p>Specify user-defined data types can use to estimate the cost of an R-tree index for queries on UDT columns: <a href="#">RTREE_COST_ADJUST_VALUE environment variable on page 216</a></p>
Query: prioritization	Specify the degree of parallelism that the database server uses: <a href="#">PDQPRIORITY environment variable on page 210</a>
Remote shell	Specify information that overrides the default remote shell for performing remote tape operations: <a href="#">DBREMOTECMD environment variable (UNIX) on page 174</a>

**Table 55. Uses for environment variables (continued)**

Functional category	Environment variable
Role separation	Configure role separation: <a href="#">INF_ROLE_SEP environment variable on page 202</a>
Rolled-back transactions	<p>Manage what the DB-Access utility does when an error occurs: <a href="#">DBACCNOIGN environment variable on page 161</a></p> <p>Specify whether an internal rollback of a global transaction frees the transaction: <a href="#">IFX_XASTDCOMPLIANCE_XAEND environment variable on page 194</a></p>
Server locale	Define the locale of your database server: <a href="#">GLS-related environment variables on page <b>SERVER_LOCALE</b></a>
Shared memory	<p>Specify where shared-memory segments are attached to the client process: <a href="#">ONEDB_SHMBASE environment variable (UNIX) on page 200</a></p> <p>Specify the shared-memory address for High Performance Loader (HPL) processes: <a href="#">PLOAD_SHMBASE environment variable on page 212</a></p>
Shell: remote	Specify information that overrides the default remote shell for performing remote tape operations: <a href="#">DBREMOTECMD environment variable (UNIX) on page 174</a>
Shell: search path	Specify which directories to search for executable programs: <a href="#">PATH environment variable on page 210</a>
Sorting	<p>Specify the location of temporary files for sorts: <a href="#">PSORT_DBTEMP environment variable on page 214</a></p> <p>Allocate more threads for sorting: <a href="#">PSORT_NPROCS environment variable on page 215</a></p>
SQL statements	<p>Specify information for caching: <a href="#">STMT_CACHE environment variable on page 217</a></p> <p>Specify connection information: <a href="#">CONNECT_RETRIES environment variable on page 196</a>, <a href="#">CONNECT_TIMEOUT environment variable on page 197</a>, and <a href="#">ONEDB_SERVER environment variable on page 199</a></p> <p>Specify information for CREATE TEMP TABLE operations: <a href="#">DBSPACETEMP environment variable on page 175</a></p> <p>Specify information for DESCRIBE FOR UPDATE operations: <a href="#">IFX_UPDDESC environment variable on page 193</a></p> <p>Specify information for LOAD and UNLOAD operations: <a href="#">DBDELIMITER environment variable on page 168</a> and <a href="#">DBBLOBBUF environment variable on page 163</a></p>

**Table 55. Uses for environment variables (continued)**

Functional category	Environment variable
	Specify information for SELECT INTO TEMP operations: <a href="#">DBSPACETEMP environment variable on page 175</a> Specify information for SET PDQPRIORITY operations: <a href="#">PDQPRIORITY environment variable on page 210</a> Specify information for SET STMT_CACHE operations Specify information for UPDATE STATISTICS operations: <a href="#">DBUPSPACE environment variable on page 179</a>
Stack size	Define the stack size that is applied to client processes: <a href="#">ONEDB_STACKSIZE environment variable on page 201</a>
Temporary tables	Define information for temporary tables: <a href="#">DBSPACETEMP environment variable on page 175</a> , <a href="#">DBTEMP environment variable on page 176</a> , and <a href="#">PSORT_DBTEMP environment variable on page 214</a>
Terminal handling	Specify terminal information: <a href="#">ONEDB_TERM environment variable (UNIX) on page 201</a> , <a href="#">TERM environment variable (UNIX) on page 218</a> , <a href="#">TERMCAP environment variable (UNIX) on page 218</a> , and <a href="#">TERMINFO environment variable (UNIX) on page 219</a>
Time-limited software license	Set information for trial or evaluation software warning messages: <a href="#">IFX_NO_TIMELIMIT_WARNING environment variable on page 190</a>
Variables: overriding	Deactivate some specified environment variable settings: <a href="#">ENVIGNORE environment variable (UNIX) on page 182</a>
Virtual memory segments on large pages	Specify whether the database server can use large pages on platforms where the hardware and the operating system support large pages of shared memory: <a href="#">IFX_LARGE_PAGES environment variable on page 186</a>
Year values (abbreviated)	Specify how to expand DATE and DATETIME values that are entered as abbreviated year values: <a href="#">DBCENTURY environment variable on page 163</a>

## ANSIOWNER environment variable

In an ANSI-compliant database, you can prevent the default behavior of upshifting lowercase letters in owner names that are not delimited by quotation marks by setting the **ANSIOWNER** environment variable to 1.

```
setenvANSIOWNER 1
```

To prevent upshifting of lowercase letters in owner names in an ANSI-compliant database, you must set **ANSIOWNER** before you initialize HCL OneDB™.

The following table shows how an ANSI-compliant database of HCL OneDB™ stores or reads the specified name of a database object called **oblong** if you were the owner of **oblong** and your **userid** (in all lowercase letters) were **owen**:

**Table 56. Lettercase of implicit, unquoted, and quoted owner names, with and without ANSIOWNER**

Owner Format	Specification	ANSIOWNER = 1	ANSIOWNER Not Set
<b>Implicit:</b>	oblong	owen.oblong	OWEN.oblong
<b>Unquoted:</b>	owen.oblong	owen.oblong	OWEN.oblong
<b>Quoted:</b>	'owen'.oblong	owen.oblong	owen.oblong

Because they do not match the lettercase of your **userid**, any SQL statements that specified the formats that are stored as **OWEN.oblong** would fail with errors.

## CPFIRST environment variable

Use the **CPFIRST** environment variable to specify the default compilation order for all source files in your programming environment.

```
setenvCPFIRST { TRUE | FALSE }
```

When you compile programs with **CPFIRST** not set, the preprocessor runs first, by default, on the program source file and then passes the resulting file to the C language preprocessor and compiler. You can, however, compile the program source file in the following order:

1. Run the C preprocessor
2. Run the preprocessor
3. Run the C compiler and linker

To use a nondefault compilation order for a specific program, you can either give the program source file a `.ecp` extension, run the `-cp` option with the **esql** command on a program source file with a `.ec` extension, or set **CPFIRST**.

Set **CPFIRST** to `TRUE` (uppercase only) to run the C preprocessor before the preprocessor on all source files in your environment, irrespective of whether the `-cp` option is passed to the **esql** command or the source files have the `.ec` or the `.ecp` extension.

To restore the default order on a system where the **CPFIRST** environment variable has been set to `TRUE`, you can set **CPFIRST** to `FALSE`. On UNIX™ systems that support the C shell, the following command has the same effect:

```
unsetenv CPFIRST
```

## CMCONFIG environment variable

Set the **CMCONFIG** environment variable to specify the location of the Connection Manager configuration file. You use the configuration file to specify service level agreements and other Connection Manager configuration options.

```
setenvCMCONFIGpath/file_name
```

***path/file\_name***

is the full path and file name of a Connection Manager configuration file.

If the CMCONFIG environment variable is not set and the configuration file name is not specified on the oncmsm utility command line, the Connection Manager attempts to load the file from the following path and file name:

```
$ONEDB_HOME/etc/cmsm.cfg
```

**Example****Examples**

Suppose the CMCONFIG environment variable points to a valid path and file name of a Connection Manager configuration file. To reload a Connection Manager instance using the configuration file specified in the shell environment enter the following command:

```
./oncmsm -r
```

To shut down a Connection Manager instance using the configuration file specified in the shell environment:

```
./oncmsm -k
```

**CLIENT\_LABEL environment variable**

Set the **CLIENT\_LABEL** environment variable in CSDK or JDBC client to assign a character string to CSDK or JDBC client session and identify that character string on the database server. You use this for environments where same userid runs multiple instances of the same application, and there is a need to distinguish one session from the other.

```
onstat -g env ses/D
```

```
select * from sysenvses where envses_name = CLIENT_LABEL
```

**Example****CSDK Example**

Suppose the CLIENT\_LABEL is set to two different strings and the same esqlc program is executed with the session ids being 43 and 201:

```
bash-3.2$ export CLIENT_LABEL='csdk_client1'
bash-3.2$ ./myesqlc
```

```
bash-3.2$ export CLIENT_LABEL='csdk_client2'
bash-3.2$ ./myesqlc
```

**onstat**

```
onstat -g env 43
HCL OneDB -- On-Line -- Up 5 days 23:01:39 -- 210712 Kbytes
```

```
Environment for session 43:
```

```

Variable          Value [values-list]
CLIENT_LABEL     cdsk_client2
CLIENT_LOCALE    en_US.8859-1
CLNT_PAM_CAPABLE 1
↓σνιπ∅

onstat -g env 201

HCL OneDB                -- On-Line -- Up 5 days 23:02:41 -- 210712 Kbytes

Environment for session 201:

Variable          Value [values-list]
CLIENT_LABEL     cdsk_client1
CLIENT_LOCALE    en_US.8859-1
CLNT_PAM_CAPABLE 1
    
```

**sysmaster**

```

select * from sysenvses where envses_name = 'CLIENT_LABEL'

envses_sid      201
envses_id       9
envses_name     CLIENT_LABEL
envses_value    cdsk_client1

envses_sid      43
envses_id       9
envses_name     CLIENT_LABEL
envses_value    cdsk_client2

2 row(s) retrieved.

Database closed.
    
```

**Example**

**JDBC Example**

Suppose the CLIENT\_LABEL is set to two different strings in the JDBC connection URL and the same JDBC program is executed with the session ids being 232 and 234:

```

java myjdbc "jdbc:onedb://myhost:52220:user=myuser;password=mypasswd;CLIENT_LABEL=jdbc_client1"

java myjdbc "jdbc:onedb://myhost:52220:user=myuser;password=mypasswd;CLIENT_LABEL=jdbc_client2"
    
```

**onstat**

```

onstat -g env 232
HCL OneDB                -- On-Line -- Up 6 days 00:56:26 -- 210712 Kbytes

Environment for session 232:

Variable          Value [values-list]
CLIENT_LABEL     jdbc_client1
CLIENT_LOCALE    en_US.8859-1
CLNT_PAM_CAPABLE 1
    
```



```

onstat -g env 234

HCL OneDB Version 1.0.0.0                -- On-Line -- Up 6 days 00:56:59 -- 210712 Kbytes

Environment for session 234:

Variable          Value [values-list]
CLIENT_LABEL      jdbc_client2
CLIENT_LOCALE     en_US.8859-1
CLNT_PAM_CAPABLE  1

sysmaster

Database selected.

select * from sysenvses where envses_name = 'CLIENT_LABEL'

envses_sid  234
envses_id   9
envses_name CLIENT_LABEL
envses_value jdbc_client2

envses_sid  232
envses_id   9
envses_name CLIENT_LABEL
envses_value jdbc_client1

2 row(s) retrieved.

Database closed.

```

## DBACCNOIGN environment variable

Use the **DBACCNOIGN** environment variable to specify the behavior of the DB-Access utility when specified errors occurs.

The **DBACCNOIGN** environment variable affects the behavior of the DB-Access utility if an error occurs under one of the following circumstances:

- You run DB-Access in non-menu mode.
- In HCL OneDB™ only, you execute the LOAD command with DB-Access in menu mode.

Set the **DBACCNOIGN** environment variable to `1` to roll back an incomplete transaction if an error occurs while you run the DB-Access utility under either of the preceding conditions.

```
setenvDBACCNOIGN1
```

For example, assume DB-Access runs the following SQL commands:

```

DATABASE mystore
BEGIN WORK

INSERT INTO receipts VALUES (cust1, 10)
INSERT INTO receipt VALUES (cust1, 20)
INSERT INTO receipts VALUES (cust1, 30)

```

```
UPDATE customer
  SET balance =
    (SELECT (balance-60)
     FROM customer WHERE custid = 'cust1')
  WHERE custid = 'cust1
COMMIT WORK
```

Here, one statement has a misspelled table name: the **receipt** table does not exist. If **DBACCNOIGN** is not set in your environment, DB-Access inserts two records into the **receipts** table and updates the **customer** table. Now, the decrease in the **customer** balance exceeds the sum of the inserted receipts.

But if **DBACCNOIGN** is set to `1`, messages open that indicate that DB-Access rolled back all the INSERT and UPDATE statements. The messages also identify the cause of the error so that you can resolve the problem.

## LOAD statement example when DBACCNOIGN is set

You can set the **DBACCNOIGN** environment variable to protect data integrity during a LOAD statement, even if DB-Access runs the LOAD statement in menu mode.

Assume you execute the LOAD statement from the DB-Access SQL menu. Forty-nine rows of data load correctly, but the 50th row contains an invalid value that causes an error. If you set **DBACCNOIGN** to `1`, the database server does not insert the forty-nine previous rows into the database. If **DBACCNOIGN** is not set, the database server inserts the first 49 rows.

## DBANSIWARN environment variable

Use the **DBANSIWARN** environment variable to indicate that you want to check for HCL OneDB™ extensions to ANSI-standard SQL syntax.

Unlike most environment variables, you are not required to set

```
DBANSIWARN
```

to a value. You can set it to any value or to no value.

```
setenvDBANSIWARN
```

Running DB-Access with **DBANSIWARN** set is functionally equivalent to including the **-ansi** flag when you invoke DB-Access (or any HCL OneDB™ product that recognizes the **-ansi** flag) from the command line. If you set **DBANSIWARN** before you run DB-Access, any syntax-extension warnings are displayed on the screen within the SQL menu.

At runtime, the **DBANSIWARN** environment variable causes the sixth character of the **sqlwarn** array in the SQL Communication Area (SQLCA) to be set to `w` when a statement is executed that is recognized as including any HCL OneDB™ extension to the ANSI/ISO standard for SQL syntax.

For details on SQLCA, see the *HCL OneDB™ ESQ/C Programmer's Manual*.

After you set **DBANSIWARN**, HCL OneDB™ extension checking is automatic until you log out or unset **DBANSIWARN**. To turn off HCL OneDB™ extension checking, you can disable **DBANSIWARN** with this command:

```
unsetenv DBANSIWARN
```

## DBBLOBBUF environment variable

Use the **DBBLOBBUF** environment variable to control whether TEXT or BYTE values are stored temporarily in memory or in a file while being processed by the UNLOAD statement. **DBBLOBBUF** affects only the UNLOAD statement.

```
setenvDBBLOBBUFSize
```

### size

represents the maximum size of TEXT or BYTE data in KB.

If the TEXT or BYTE data size is smaller than the default of 10 KB (or the setting of **DBBLOBBUF**), the TEXT or BYTE value is temporarily stored in memory. If the data size is larger than the default or the **DBBLOBBUF** setting, the data value is written to a temporary file. For instance, to set a buffer size of 15 KB, set **DBBLOBBUF** as in the following example:

```
setenv DBBLOBBUF 15
```

Here any TEXT or BYTE value smaller than 15 KB is stored temporarily in memory. Values larger than 15 KB are stored temporarily in a file.

## DBCENTURY environment variable

Use the **DBCENTURY** environment variable to specify how to expand literal DATE and DATETIME values that are entered with abbreviated year values. To avoid problems in expanding abbreviated years, applications should require entry of 4-digit years, and should always display years as four digits.

```
setenvDBCENTURY { R | F | { C | P } }
```

When **DBCENTURY** is not set (or is set to **R**), the first two digits of the current year are used to expand 2-digit year values. For example, if today's date is 09/30/2003, then the abbreviated date 12/31/99 expands to 12/31/2099, and the abbreviated date 12/31/00 expands to 12/31/2000.

The R, P, F, and C settings determine algorithms for expanding two-digit years.

Setting	Algorithm
R = Current®	Use the first two digits of the current year to expand the year value.
P = Past	Expanded dates are created by prefixing the abbreviated year value with 19 and 20. Both dates are compared to the current date, and the most recent date that is earlier than the current date is used.
F = Future	Expanded dates are created by prefixing the abbreviated year value with 20 and 21. Both dates are compared to the current date, and the earliest date that is later than the current date is used.
C = Closest	Expanded dates are created by prefixing the abbreviated year value with 19, 20, and 21. These three dates are compared to the current date, and the date that is closest to the current date is used.

Settings are case sensitive, and no error is issued for invalid settings. If you enter **F** (for example), then the default (**R**) setting takes effect. The **P** and **F** settings cannot return the current date, which is not in the past or future.

Years entered as a single digit are prefixed with 0 and then expanded. Three-digit years are not expanded. Pad years earlier than 100 with leading zeros.

## Examples of expanding year values

The examples in this topic illustrate how various settings of **DBCENTURY** cause abbreviated years to be expanded in DATE and DATETIME values.

### DBCENTURY = P

```
Example data type: DATE
Current date: 4/6/2003
User enters: 1/1/1
Prefix with "19" expansion : 1/1/1901
Prefix with "20" expansion: 1/1/2001
Analysis: Both are prior to current date, but 1/1/2001 is closer to
current date.
```



**Important:** The effect of **DBCENTURY** depends on the current date from the system clock-calendar. Thus, 1/1/1, the abbreviated date in this example, would instead be expanded to 1/1/1901 if the current date were 1/1/2001 and **DBCENTURY = P**.

### DBCENTURY = F

```
Example data type: DATETIME year to month
Current date: 5/7/2005
User enters: 1-1
Prefix with "20" expansion: 2001-1
Prefix with "21" expansion: 2101-1
Analysis: Only date 2101-1 is after the current date, so it is chosen.
```

### DBCENTURY = C

```
Example data type: DATE
Current date: 4/6/2000
User enters: 1/1/1
Prefix with "19" expansion : 1/1/1901
Prefix with "20" expansion: 1/1/2001
Prefix with "21" expansion: 1/1/2101
Analysis: Here 1/1/2001 is closest to the current date, so it is chosen.
```

### DBCENTURY = R or DBCENTURY Not Set

```
Example data type: DATETIME year to month
Current date: 4/6/2000
User enters: 1-1
Prefix with "20" expansion: 2001-1

Example data type: DATE
Current date: 4/6/2003
User enters: 0/1/1
Prefix with "20" expansion: 2000/1
Analysis: In both examples, the Prefix with "20" algorithm is used.
```

Setting **DBCENTURY** does not affect HCL OneDB™ products when the locale specifies a non-Gregorian calendar, such as Hebrew or Islamic calendars. The leading digits of the current year are used for alternative calendar systems when the year is abbreviated.

Setting **DBCENTURY** does not affect HCL OneDB™ products when the locale specifies a non-Gregorian calendar. The leading digits of the current year are used for alternative calendar systems when the year is abbreviated.

## Abbreviated years and expressions in database objects

When an expression in a database object (including a check constraint, fragmentation expression, SPL routine, trigger, or UDR) contains a literal date or DATETIME value in which the year has one or two digits, the database server evaluates the expression using the setting that **DBCENTURY** (and other relevant environment variables) had when the database object was created (or was last modified).

If **DBCENTURY** has been reset to a new value, the new value is ignored when the abbreviated year is expanded.

For example, suppose a user creates a table and defines the following check constraint on a column named **birthdate**:

```
birthdate < '09/25/50'
```

The expression is interpreted according to the value of **DBCENTURY** when the constraint was defined. If the table that contains the **birthdate** column is created on 09/23/2000 and **DBCENTURY** =c, the check constraint expression is consistently interpreted as `birthdate < '09/25/1950'` when inserts or updates are performed on the **birthdate** column. Even if different values of **DBCENTURY** are set when users perform inserts or updates on the **birthdate** column, the constraint expression is interpreted according to the setting at the time when the check constraint was defined (or was last modified).

Database objects created on some earlier versions of HCL OneDB™ do not support the priority of creation-time settings.

### For legacy objects to acquire this feature

1. Drop the objects.
2. Recreate them (or for fragmentation expressions, detach them and then reattach them).

After the objects are redefined, date literals within expressions of the objects will be interpreted according to the environment at the time when the object was created or was last modified. Otherwise, their behavior will depend on the runtime environment and might become inconsistent if this changes.

Administration of a database that includes a mix of legacy objects and new objects might become difficult because of differences between the new and the old behavior for evaluating date expressions. To avoid this, it is recommended that you redefine any legacy objects.

The value of **DBCENTURY** and the current date are not the only factors that determine how the database server interprets date and DATETIME values. The **DBDATE**, **DBTIME**, **GL\_DATE**, and **GL\_DATETIME** environment variables can also influence how dates are interpreted. For information about **GL\_DATE** and **GL\_DATETIME**, see the *HCL OneDB™ GLS User's Guide*.

**!** **Important:** The behavior of **DBCENTURY** for HCL OneDB™ is not compatible with earlier versions.

## DBDATE environment variable

Use the **DBDATE** environment variable to specify the end-user formats of DATE values.

On UNIX™ systems that use the C shell, set **DBDATE** with this syntax.

```
setenv DBDATE { MD | DM | Y4 | Y2 } { Y4 | Y2 | MD | DM } { / | - | . | . | 0 }
```

The following formatting symbols are valid in the **DBDATE** setting:

- . /

are characters that can exist as separators in a date format.

0

indicates that no separator is displayed between time units.

**D, M**

are characters that represent the day and the month.

**Y2, Y4**

are characters that represent the year and the precision of the year.

Some East Asian locales support additional syntax for era-based dates.

**DBDATE** can specify the following attributes of the display format:

- The order of time units (the month, day, and year) in a date
- Whether the year is shown as two digits (Y2) or four digits (Y4)
- The separator between the month, day, and year time units

For the U.S. English locale, the default for **DBDATE** is `MDY4/`, where **M** represents the month, **D** represents the day, **Y4** represents a four-digit year, and slash ( / ) is the time-units separator (for example, `01/08/2011`). Other valid characters for the separator are a hyphen ( - ), a period ( . ), or a zero ( 0 ). To indicate no separator, use the zero. The slash ( / ) is used by default if you attempt to specify a character other than a hyphen, period, or zero as a separator, or if you do not include any separator in the **DBDATE** specification.

If **DBDATE** is not set on the client, any **DBDATE** setting on the database server overrides the `MDY4/` default on the client. If **DBDATE** is set on the client, that value (rather than the setting on the database server) is used by the client.

The following table shows some examples of valid **DBDATE** settings and their corresponding displays for the date 8 January, 2011:

DBDATE Setting	Representation of January 8, 2011:		DBDATE Setting	Representation of January 8, 2011:
MDY4/	01/08/2011		Y2DM.	11.08.01
DMY2-	08-01-11		MDY20	010811
MDY4	01/08/2011		Y4MD*	2011/01/08

Formats `Y4MD*` (because asterisk is not a valid separator) and `MDY4` (with no separator defined) both display the default symbol (slash) as the separator.



**Important:** If you use the Y2 format, the setting of the **DBCENTURY** environment variable can also affect how literal DATE values are evaluated in data entry.

Also, certain routines that call `DATE` can use the **DBTIME** variable, rather than **DBDATE**, to set DATETIME formats to international specifications. For more information, see the discussion of the **DBTIME** environment variable in [DBTIME environment variable on page 176](#) and in the *HCL OneDB™ ESQL/C Programmer's Manual*.

The setting of the **DBDATE** variable takes precedence over that of the **GL\_DATE** environment variable, and over any default DATE format that **CLIENT\_LOCALE** specifies. For information about **GL\_DATE** and **CLIENT\_LOCALE**, see the *HCL OneDB™ GLS User's Guide*.

End-user formats affect the following contexts:

- When you display DATE values, HCL OneDB™ products use the **DBDATE** environment variable to format the output.
- During data entry of DATE values, HCL OneDB™ products use the **DBDATE** environment variable to interpret the input.

For example, if you specify a literal DATE value in an INSERT statement, the database server expects this literal value to be compatible with the format that **DBDATE** specifies. Similarly, the database server interprets the date that you specify as the argument to the `DATE()` function to be in **DBDATE** format.

## DATE expressions in database objects

When an expression in a database object (including a check constraint, fragmentation expression, SPL routine, trigger, or UDR) contains a literal date value, the database server evaluates the expression using the setting that **DBDATE** (or other relevant environment variables) had when the database object was created (or was last modified). If **DBDATE** has been reset to a new value, the new value is ignored when the literal DATE is evaluated.

For example, suppose **DBDATE** is set to `MDY2/` and a user creates a table with the following check constraint on the column **orderdate**:

```
orderdate < '06/25/98'
```

The date of the preceding expression is formatted according to the value of **DBDATE** when the constraint is defined. The check constraint expression is interpreted as `orderdate < '06/25/98'` regardless of the value of **DBDATE** during inserts or updates on the **orderdate** column. Suppose **DBDATE** is reset to `DMY2/` when a user inserts the value `'30/01/98'` into the

**orderdate** column. The date value inserted uses the date format `DMY2/`, whereas the check constraint expression uses the date format `MDY2/`.

See [Abbreviated years and expressions in database objects on page 165](#) for a discussion of legacy objects from earlier versions of HCL OneDB™ that are always evaluated according to the runtime environment. That section describes how to redefine objects so that dates are interpreted according to environment variable settings that were in effect when the object was defined (or when the object was last modified).



**Important:** The behavior of **DBDATE** for HCL OneDB™ is not compatible with earlier versions.

## DBDELIMITER environment variable

Set the **DBDELIMITER** environment variable to specify the field delimiter used with the **dbexport** utility and with the **LOAD** and **UNLOAD** statements.

```
setenvDBDELIMITER'delimiter'
```

### *delimiter*

is the field delimiter for unloaded data files.

The *delimiter* can be any single character, except those in the following list:

- Hexadecimal digits (0 through 9, a through f, A through F)
- Newline or `CTRL-J`
- The backslash ( \ ) symbol

The vertical bar ( | = ASCII 124 ) is the default. To change the field delimiter to a plus ( + ) symbol, for example, you can set **DBDELIMITER** as follows:

```
setenv DBDELIMITER '+'
```

## DBEDIT environment variable

Use the **DBEDIT** environment variable to specify the text editor to use with SQL statements and command files in DB-Access.

If **DBEDIT** is set, the specified text editor is invoked automatically. If **DBEDIT** is not, set you are prompted to specify a text editor as the default for the rest of the session.

```
setenvDBEDITeditor
```

### *editor*

is the name of the text editor you want to use.

For most UNIX™ systems, the default text editor is **vi**. If you use another text editor, be sure that it creates flat ASCII files. Some word processors in *document mode* introduce printer control characters that can interfere with the operation of your HCL OneDB™ product.

To specify the EMACS text editor, set **DBEDIT** with the following command:



```
setenv DBEDIT emacs
```

## DBFLTMASK environment variable

The DB-Access utility displays the floating-point values of data types FLOAT, SMALLFLOAT, and DECIMAL(*p*) within a 14-character buffer. By default, DB-Access displays as many digits to the right of the decimal point as will fit into this character buffer. Therefore, the actual number of decimal digits that DB-Access displays depends on the size of the floating-point value.

To reduce the number of digits displayed to the right of the decimal point in floating-point values, set **DBFLTMASK** to the specified number of digits.

```
setenvDBFLTMASKScale
```

### *scale*

is the number of decimal digits that you want the HCL OneDB™ client application to display in the floating-point values. Here *scale* must be smaller than 16, the default number of digits displayed.

If the floating-point value contains more digits to the right of the decimal than **DBFLTMASK** specifies, DB-Access rounds the value to the specified number of digits. If the floating-point value contains fewer digits to the right of the decimal, DB-Access pads the value with zeros. If you set **DBFLTMASK** to a value greater than can fit into the 14-character buffer, however, DB-Access rounds the value to the number of digits that can fit.

## DBINFO\_DBSPACE\_RETURN\_NULL\_FOR\_INVALID\_PARTNUM environment variable

Use the **DBINFO\_DBSPACE\_RETURN\_NULL\_FOR\_INVALID\_PARTNUM** environment variable to specify if `dbinfo('dbspace', partnum)` raises an error -727 or returns NULL when an invalid partition number (`partnum`) is provided.

```
DBINFO_DBSPACE_RETURN_NULL_FOR_INVALID_PARTNUM { 0 | 1 }
```

A partition number is considered invalid if it resolves to a `dbspace` number which is not a valid `dbspace` in the instance. This includes the pseudo tables which are having partition numbers that would be associated with `dbspace 0` which is not a (real) `dbspace` in an OneDB instance. This reflects that pseudo tables do not directly have an on-disk representation but rather are state information from (shared) memory which are exposed via SQL.

In case of an invalid `partnum` the `dbinfo('dbspace', partnum)` function would result in an error '727: Invalid or NULL TBLspace number given to dbinfo(db space)'. When the environment variable **DBINFO\_DBSPACE\_RETURN\_NULL\_FOR\_INVALID\_PARTNUM** is set to 1, `dbinfo()` will not result in an error in this case, but rather does return NULL as `dbspace` name. When setting a value of '0' or not setting the environment variable, the default behavior returns an error -727 for an invalid `partnum`. In any case a NULL provided as `partnum` will result in error -727 being raised.

With SET ENVIRONMENT DBINFO\_DBSPACE\_RETURN\_NULL\_FOR\_INVALID\_PARTNUM the variable can be set dynamically at runtime. This overrides the current DBINFO\_DBSPACE\_RETURN\_NULL\_FOR\_INVALID\_PARTNUM value for the current user session only. For more information about the SET ENVIRONMENT DBINFO\_DBSPACE\_RETURN\_NULL\_FOR\_INVALID\_PARTNUM statement of SQL, see the [Guide to SQL: Syntax on page](#) .

## DBLANG environment variable

Use the **DBLANG** environment variable to specify the subdirectory of **\$ONEDB\_HOME** or the full pathname of the directory that contains the compiled message files that HCL OneDB™ products use.

```
setenvDBLANG { relative_path | full_path }
```

### **relative\_path**

is a subdirectory of **\$ONEDB\_HOME**.

### **full\_path**

is the pathname to the compiled message files.

By default, HCL OneDB™ products put compiled messages in a locale-specific subdirectory of the **\$ONEDB\_HOME/msg** directory. These compiled message files have the file extension **.iem**. If you want to use a message directory other than **\$ONEDB\_HOME/msg**, where, for example, you can store message files that you create, you must perform the following steps:

### To use a message directory other than **\$ONEDB\_HOME/msg**

1. Use the **mkdir** command to create the appropriate directory for the message files.

You can make this directory under the directory **\$ONEDB\_HOME** or **\$ONEDB\_HOME/msg**, or you can make it under any other directory.

2. Set the owner and group of the new directory to **onedb** and the access permission for this directory to 755.
3. Set the **DBLANG** environment variable to the new directory. If this is a subdirectory of **\$ONEDB\_HOME** or **\$ONEDB\_HOME/msg**, then you need only list the relative path to the new directory. Otherwise, you must specify the full pathname of the directory.
4. Copy the **.iem** files or the message files that you created to the new message directory that **\$DBLANG** specifies.

All the files in the message directory should have the owner and group **informix** and access permission 644.

HCL OneDB™ products that use the default U.S. English locale search for message files in the following order:

1. In **\$DBLANG**, if **DBLANG** is set to a full pathname
2. In **\$ONEDB\_HOME/msg/\$DBLANG**, if **DBLANG** is set to a relative pathname
3. In **\$ONEDB\_HOME/\$DBLANG**, if **DBLANG** is set to a relative pathname
4. In **\$ONEDB\_HOME/msg/en\_us/0333**
5. In **\$ONEDB\_HOME/msg/en\_us.8859-1**
6. In **\$ONEDB\_HOME/msg**
7. In **\$ONEDB\_HOME/msg/english**

For more information about search paths for messages, see the description of **DBLANG** in the *HCL OneDB™ GLS User's Guide*.

## DBMONEY environment variable

Use the **DBMONEY** environment variable to specify the display format of values in columns of smallfloat, FLOAT, DECIMAL, or MONEY data types, and of complex data types derived from any of these data types.

```
setenv DBMONEY {'$' | front | 'front'} {_ | , } [ {back | 'back'} ]
```

**\$**

is a currency symbol that precedes MONEY values in the default locale if no other *front* symbol is specified, or if **DBMONEY** is not set.

**, or .**

is a comma or period (the default) that separates the integral part from the fractional part of the FLOAT, DECIMAL, or MONEY value. Whichever symbol you do not specify becomes the thousands separator.

**back**

is a currency symbol that follows the MONEY value.

**front**

is a currency symbol that precedes the MONEY value.

The *back* symbol can be up to seven characters and can contain any character that the locale supports, except a digit, a comma ( , ), or a period ( . ) symbol. The *front* symbol can be up to seven characters and can contain any character that the locale supports except a digit, a comma ( , ), or a period ( . ) symbol. If you specify any character that is not a letter of the alphabet for *front* or *back*, you must enclose the *front* or *back* setting between single quotation ( ' ) marks.

When you display MONEY values, HCL OneDB™ products use the **DBMONEY** setting to format the output. **DBMONEY** has no effect, however, on the internal format of data values that are stored in columns of the database.

If you do not set **DBMONEY**, then MONEY values for the default locale, U.S. English, are formatted with a dollar sign ( \$ ) that precedes the MONEY value, a period ( . ) that separates the integral from the fractional part of the MONEY value, and no *back* symbol. For example, 100.50 is formatted as \$100.50.

Suppose you want to represent MONEY values as DM (deutsche mark) units, using the currency symbol **DM** and comma ( , ) as the decimal separator. Enter the following command to set the **DBMONEY** environment variable:

```
setenv DBMONEY DM,
```

Here **DM** is the *front* currency symbol that precedes the MONEY value, and a comma separates the integral from the fractional part of the MONEY value. As a result, the value 100.50 is displayed as **DM100,50**.

For more information about how **DBMONEY** formats MONEY values in nondefault locales, see the *HCL OneDB™ GLS User's Guide*.

## DBONPLOAD environment variable

Use the **DBONPLOAD** environment variable to specify the name of the database that the onpload utility of the High Performance Loader (HPL) uses.

If **DBONPLOAD** is set, **onpload** uses the specified name as the name of the database; otherwise, the default name of the database is **onpload**.

```
setenvDBONPLOADdbname
```

**dbname**

specifies the name of the database that the onpload utility uses.

For example, to specify the name **load\_db** as the name of the database, enter the following command:

```
setenv DBONPLOAD load_db
```

For more information, see the *HCL OneDB™ High-Performance Loader User's Guide*.

## DBPATH environment variable

Use the **DBPATH** environment variable to identify the database servers that contain databases. **DBPATH** can also specify a list of directories (in addition to the current directory) in which DB-Access looks for command scripts (**.sql** files).

The **CONNECT DATABASE**, **START DATABASE**, and **DROP DATABASE** statements use **DBPATH** to locate the database under two conditions:

- If the location of a database is not explicitly stated
- If the database cannot be located in the default server

The **CREATE DATABASE** statement does not use **DBPATH**.

To add a new **DBPATH** entry to existing entries, see [Modifying an environment-variable setting on page 142](#).

```
setenvDBPATH { |pathname| /servername/full_pathname | /servername }
```

**full\_pathname**

is the full path, from **root**, of a directory where **.sql** files are stored.

**pathname**

is the valid relative path of a directory where **.sql** files are stored.

**servername**

is the name of a database server where databases are stored. You cannot reference database files with a *servername*.

**DBPATH** can contain up to 16 entries. Each entry must be less than 128 characters. In addition, the maximum length of **DBPATH** depends on the hardware platform on which you set **DBPATH**.

When you access a database with the **CONNECT**, **DATABASE**, **START DATABASE**, or **DROP DATABASE** statement, the search for the database is done first in the directory or database server specified in the statement. If no database server is specified, the default database server that was specified by the **ONEDB\_SERVER** environment variable is used.

If the database is not located during the initial search, and if **DBPATH** is set, the database servers and directories in **DBPATH** are searched for in the specified database. These entries are searched in the same order in which they are listed in the **DBPATH** setting.

## Using DBPATH with DB-Access

If you use DB-Access and select the **Choose** option from the **SQL** menu without having already selected a database, you see a list of all the **.sql** files in the directories listed in your **DBPATH**. After you select a database, the **DBPATH** is not used to find the **.sql** files. Only the **.sql** files in the current working directory are displayed.

## Searching local directories

### About this task

Use a pathname without a database server name to search for **.sql** scripts on your local computer. In the following example, the **DBPATH** setting causes DB-Access to search for the database files in your current directory and then in the Joachim and Sonja directories on the local computer:

```
setenv DBPATH /usr/joachim:/usr/sonja
```

As the previous example shows, if the pathname specifies a directory name but not a database server name, the directory is sought on the computer that runs the default database server that the **ONEDB\_SERVER** specifies; see [ONEDB\\_SERVER environment variable on page 199](#). For instance, with the previous example, if **ONEDB\_SERVER** is set to **quality**, the **DBPATH** value is *interpreted*, as the following example shows, where the double slash precedes the database server name:

```
setenv DBPATH //quality/usr/joachim://quality/usr/sonja
```

## Searching networked computers for databases

### About this task

If you use more than one database server, you can set **DBPATH** explicitly to contain the database server and directory names that you want to search for databases. For example, if **ONEDB\_SERVER** is set to **quality**, but you also want to search the **marketing** database server for **/usr/joachim**, set **DBPATH** as the following example shows:

```
setenv DBPATH //marketing/usr/joachim:/usr/sonja
```

## Specifying a servername

### About this task

You can set **DBPATH** to contain only database server names. This feature allows you to locate only databases; you cannot use it to locate command files.

The database administrator must include each database server mentioned by **DBPATH** in the **\$ONEDB\_HOME/etc/sqlhosts** file. For information about communication-configuration files and dbservernames, see your *HCL OneDB™ Administrator's Guide* and the *HCL OneDB™ Administrator's Reference*.

For example, if **ONEDB\_SERVER** is set to **quality**, you can search for a database first on the **quality** database server and then on the **marketing** database server by setting **DBPATH**, as the following example shows:

```
setenv DBPATH //marketing
```

If you use DB-Access in this example, the names of all the databases on the **quality** and **marketing** database servers are displayed with the **Select** option of the DATABASE menu.

## DBPRINT environment variable

Use the **DBPRINT** environment variable to specify the default printing program.

```
setenvDBPRINTprogram
```

### **program**

Any command, shell script, or UNIX™ utility that produces standard ASCII output.

If you do not set **DBPRINT**, the default *program* is found in one of two places:

- For most BSD UNIX™ systems, the default program is **lpr**.
- For UNIX™ System V, the default program is usually **lp**.

Enter the following command to set the **DBPRINT** environment variable to specify **myprint** as the print program:

```
setenv DBPRINT myprint
```

## DBREMOTECMD environment variable (UNIX™)

Use the **DBREMOTECMD** environment variable to override the default remote shell to perform remote tape operations with the database server.

You can set **DBREMOTECMD** to a simple command or to a full path name.

```
setenvDBREMOTECMD { command | pathname }
```

### **command**

A command to override the default remote shell.

### **pathname**

A path name to override the default remote shell.

If you do not specify the full path name, the database server searches your **PATH** for the specified *command*. You should use the full path name syntax on interactive UNIX™ platforms to avoid problems with similarly named programs in other directories and possible confusion with the *restricted shell* (*/usr/bin/rsh*).

The following command sets **DBREMOTECMD** for a simple command name:

```
setenv DBREMOTECMD rcmd
```

The next command to set **DBREMOTECMD** specifies a full path name:

```
setenv DBREMOTECMD /usr/bin/remsh
```

For more information about using remote tape devices for backups, see [Specify a remote device on page](#) .

## DBSPACETEMP environment variable

The **DBSPACETEMP** environment variable specifies the dbspaces in which temporary tables are built. The list can include standard dbspaces, temporary dbspaces, or both.

```
setenvDBSPACETEMP dbspace
```

### ***dbspace***

is the name of an existing standard or temporary dbspace.

You can list dbspaces, separated by colon (:) or comma (,) symbols, to designate space for temporary tables across physical storage devices. For example, the following command to set the **DBSPACETEMP** environment variable specifies three dbspaces for temporary tables:

```
setenv DBSPACETEMP sorttmp1:sorttmp2:sorttmp3
```

**DBSPACETEMP** overrides any default dbspaces that the DBSPACETEMP parameter specifies in the configuration file of the database server. For UPDATE STATISTICS operations, DBSPACETEMP is used only when you specify the HIGH keyword option.

On UNIX™ platforms, you might have better performance if the list of dbspaces in **DBSPACETEMP** is composed of chunks that are allocated as raw devices.

The number of dbspaces is limited by the maximum size of the environment variable, as defined by your operating system. Your database server does not create a dbspace specified by the environment variable if the dbspace does not exist.

The two classes of temporary tables are *explicit* temporary tables that the user creates and *implicit* temporary tables that the database server creates. Use **DBSPACETEMP** to specify the dbspaces for both types of temporary tables.

If you create an explicit temporary table with the CREATE TEMP TABLE statement and do not specify a dbspace for the table either in the IN *dbspace* clause or in the FRAGMENT BY clause, the database server uses the settings in **DBSPACETEMP** to determine where to create the table.

If you create an explicit temporary table with the SELECT INTO TEMP statement, the database server uses the settings in **DBSPACETEMP** to determine where to create the table.

If **DBSPACETEMP** is set, and the dbspaces that it lists include both logging and non-logging dbspaces, the database server stores temporary tables that implicitly or explicitly support transaction logging in a logged dbspace, and non-logging temporary tables in a non-logging dbspace.

The database server creates implicit temporary tables for its own use while executing join operations, SELECT statements with the GROUP BY clause, SELECT statements with the ORDER BY clause, and index builds.

When it creates explicit or implicit temporary tables, the database server uses disk space for writing the temporary data. If there are conflicts among settings or statement specifications for the location of a temporary table, these conflicts are resolved in this descending (highest to lowest) order of precedence:

1. On UNIX™ platforms, the operating-system directory or directories that the environment variable **PSORT\_DBTEMP** specifies, if this is set
2. The dbspace or dbspaces that the environment variable **DBSPACETEMP** specifies, if this is set
3. The dbspace or dbspaces that the ONCONFIG parameter DBSPACETEMP specifies.
4. The operating-system file space specified by the DUMPPDIR configuration parameter
5. The directory \$ONEDB\_HOME/tmp (UNIX™) or \$ONEDB\_HOME\tmp (Windows™).



**Important:** If the **DBSPACETEMP** environment variable is set to an invalid value, the database server defaults to the root dbspace for explicit temporary tables and to **/tmp** for implicit temporary tables, rather than to the setting of the DBSPACETEMP configuration parameter. In this situation, the database server might fill **/tmp** to the limit and eventually bring down the database server or kill the file system.

## DBTEMP environment variable

The **DBTEMP** environment variable is used by DB-Access. **DBTEMP** resembles **DBSPACETEMP**, specifying the directory in which to place temporary files and temporary tables.

```
setenv DBTEMP pathname
```

### *pathname*

The full path name of the directory for temporary files and tables.

For DB-Access to work correctly on Windows™ platforms, **DBTEMP** should be set to \$ONEDB\_HOME/infxtmp.

The following example sets **DBTEMP** to the path name `usr/magda/mytemp` for UNIX™ systems that use the C shell:

```
setenv DBTEMP usr/magda/mytemp
```



**Important:** **DBTEMP** can point to an NFS-mounted directory only if the vendor of that NFS device is certified by HCL.

If **DBTEMP** is not set, the database server creates temporary files in the `/tmp` directory and temporary tables in the **DBSPACETEMP** directory. See [DBSPACETEMP environment variable on page 175](#) for the default if **DBSPACETEMP** is not set. Similarly, if you do not set **DBTEMP** on the client system, temporary files (such as those created for scroll cursors) are created in the `/tmp` directory.

You might experience unexpected behavior or failure in operations on values of large or complex data types, such as BYTE or ROW, if **DBTEMP** is not set.

## DBTIME environment variable

The **DBTIME** environment variable specifies a formatting mask for the display and data-entry format of DATETIME values.



The **DBTIME** environment variable is useful in contexts where the DATETIME data values to be formatted by **DBTIME** have the same precision as the specified **DBTIME** setting. You might encounter unexpected or invalid display formats for DATETIME values that are declared with a different DATETIME qualifier.

```
setenv DBTIME ' { / literal | % [ { - | o } ] [min] [ .precision ] special } ' .
```

**literal**

is a literal white space or any printable character.

**min**

is a literal integer, setting the minimum number of characters in the substring for the value that *special* specifies.

**precision**

is the number of digits for the value of any time unit, or the maximum number of characters in the name of a month.

**special**

is one of the placeholder characters that are listed following.

These terms and symbols are described in the pages that follow.

This quoted string can include literal characters and placeholders for the values of individual time units and other elements of a DATETIME value. **DBTIME** takes effect only when you call certain DATETIME routines. (For details, see the *HCL OneDB™ ESQL/C Programmer's Manual*.) If **DBTIME** is not set, the behavior of these routines is undefined, and "YYYY-MM-DD hh:mm:ss.ffffff" is the default display and input format for DATETIME YEAR TO FRACTION(5) literal values in the default locale.

The percentage ( %) symbol gives special significance to the *special* placeholder symbol that follows. Without a preceding % symbol, any character within the formatting mask is interpreted as a literal character, even if it is the same character as one of the placeholder characters in the following list. Note also that the *special* placeholder symbols are case sensitive.

The following characters within a **DBTIME** format string are placeholders for time units (or for other features) within a DATETIME value.

**%b**

is replaced by the abbreviated month name.

**%B**

is replaced by the full month name.

**%d**

is replaced by the day of the month as a decimal number [01,31].

**%Fn**

is replaced by a fraction of a second with a scale that the integer *n* specifies. The default value of *n* is 2; the range of *n* is  $0 < n < 5$ .

**%H**

is replaced by the hour (24-hour clock).

**%I**

is replaced by the hour (12-hour clock).

**%M**

is replaced by the minute as a decimal number [00,59].

**%m**

is replaced by the month as a decimal number [01,12].

**%p**

is replaced by A.M. or P.M. (or the equivalent in the locale file).

**%S**

is replaced by the second as a decimal number [00,59].

**%y**

is replaced by the year as a four-digit decimal number.

**%Y**

is replaced by the year as a four-digit decimal number. User must enter a four-digit value.

**%%**

is replaced by % (to allow a literal % character in the format string).

For example, consider this display format for DATETIME YEAR TO SECOND:

```
Mar 21, 2013 at 16 h 30 m 28 s
```

If the user enters a two-digit year value, this value is expanded to 4 digits according to the **DBCENTURY** environment variable setting. If **DBCENTURY** is not set, then the string `19` is used by default for the first two digits.

Set **DBTIME** as the following command line (for the C shell) shows:

```
setenv DBTIME '%b %d, %Y at %H h %M m %S s'
```

The default **DBTIME** produces the following ANSI SQL string format:

```
2001-03-21 16:30:28
```

You can set the default **DBTIME** as the following example shows:

```
setenv DBTIME '%Y-%m-%d %H:%M:%S'
```

An optional field width and precision specification (*w.p*) can immediately follow the percent (%) character. It is interpreted as follows:

**w**

Specifies the minimum field width. The value is right-justified with blank spaces on the left.

**-w**

Specifies the minimum field width. The value is left-justified with blank spaces on the right.

**0w**

Specifies the minimum field width. The value is right-justified and padded with zeros on the left.

**p**

Specifies the precision of `d`, `H`, `I`, `m`, `M`, `S`, `y`, and `Y` time unit values, or the maximum number of characters in `b` and `B` month names.

The following limitations apply to field-width and precision specifications:

- If the data value supplies fewer digits than *precision* specifies, the value is padded with leading zeros.
- If a data value supplies more characters than *precision* specifies, excess characters are truncated from the right.
- If no field width or precision is specified for `d`, `H`, `I`, `m`, `M`, `S`, or `y` placeholders, `0.2` is the default, or `0.4` for the `Y` placeholder.
- A *precision* specification is significant only when converting a DATETIME value to an ASCII string, but not vice versa.

The `F` placeholder does not support this field-width and precision syntax.



**Important:** Any separator character between the `%S` and `%F` directives for DATETIME user formats must be explicitly defined. Specifying `%S%F` concatenates the digits that represent the integer and fractional parts of the seconds value.

Like `DBDATE`, `GL_DATE`, or `GL_DATETIME`, or `USE_DTENV`, the `DBTIME` setting controls only the character-string representation of data values. It cannot change the internal storage format of the DATETIME column. (For additional information about formatting DATE values, see the discussion of `DBDATE` in the topic [DBDATE environment variable on page 166](#).)

### DBTIME formats in nondefault locales

If you specify a locale other than U.S. English, the locale defines the culture-specific display formats for DATETIME values. To change the default display format, change the setting of `DBTIME`, or of the `GL_DATETIME` and `USE_DTENV` environment variables.

In East Asian locales that support era-based dates, `DBTIME` can also specify Japanese eras. See *HCL OneDB™ GLS User's Guide* for details of additional placeholder symbols for setting `DBTIME` to display era-based DATETIME values, and for descriptions of the `GL_DATETIME`, `GL_DATE`, and `USE_DTENV` environment variables.

### DBUPSPACE environment variable

Use the `DBUPSPACE` environment variable to specify the amount of system disk space and the amount of memory that the UPDATE STATISTICS MEDIUM and UPDATE STATISTICS HIGH statement can use when it reads and sorts column values to

construct column distributions. The **DBUPSPACE** setting can also request SET EXPLAIN output to describe the execution path for calculating the statistical distributions.

```
setenv DBUPSPACE { 1024 | [ disk ] } { : 15 | [ : memory ] } [ : directive ]
```

#### **disk**

is an unsigned integer, specifying the disk space (in KiB) to allocate for sorting in UPDATE STATISTICS MEDIUM and HIGH operations.

#### **memory**

is an unsigned integer, specifying the maximum amount of sorting memory (in MiB, in the range from 4 to 50 megabytes) to allocate without using PDQ.

#### **directive**

is an unsigned integer, encoding one of the following directives for the UPDATE STATISTICS execution plan:

- 1: Do not use any indexes for sorting. Print the entire plan for update statistics in the `sqexplain.out` file.
- 2: Do not use any indexes for sorting. Do not print the plan for update statistics.
- 3 or greater: Use available indexes for sorting. Print the entire plan for update statistics in explain output file.

For example, to set **DBUPSPACE** to 2,500 KiB of disk space and 1 megabyte of memory, enter this command:

```
setenv DBUPSPACE 2500:1
```

After you set this value, the database server will attempt to use no more than 2,500 KiB of disk space during the execution of an UPDATE STATISTICS MEDIUM or HIGH statement. If a table requires 5 megabytes of disk space for sorting, then UPDATE STATISTICS accomplishes the task in two passes; the distributions for one half of the columns are constructed with each pass. For a table of a given storage size, this parameter determines the number of passes, but no pass can write less than a full column.

If you do not set **DBUPSPACE**, the default setting is 1 megabyte (1,024 KiB) for *disk*, and 15 megabytes for *memory*. If you attempt to set the first **DBUPSPACE** parameter to any value less than 1,024 KiB, it is automatically set to 1,024 KiB, but no error message is returned. If this *disk* value is not large enough to allow more than one distribution to be constructed at a time, at least one distribution is done, even if the amount of disk space required to do this is more than what **DBUPSPACE** specifies. That is, regardless of the *disk* parameter setting for **DBUPSPACE**, the largest individual column storage requirement of a table determines the actual upper limit on disk space for a single pass in any UPDATE STATISTICS HIGH or MEDIUM operation.

## DEFAULT\_ATTACH environment variable

The **DEFAULT\_ATTACH** environment variable supports the legacy behavior of Version 7.x of HCL OneDB™, in which the pages of nonfragmented B-tree indexes on nonfragmented tables were stored, by default, in the same dbspace partition as the data pages. (The name "**DEFAULT\_ATTACH**" derives from an obsolete definition of an *attached index*, a term that now refers to

an index whose fragmentation strategy is the same as the fragmentation strategy of its table. Do not confuse the obsolete Version 7.x definition with this current definition.)

```
setenvDEFAULT_ATTACH1
```

If the **DEFAULT\_ATTACH** environment variable is set to 1, then by default, the pages of nonfragmented B-tree indexes on nonfragmented tables are stored in the same partition (and in the same dbSPACE) that stores data pages of the table. The **IN TABLE** keywords of the **CREATE INDEX** statement are not required (but do not return an error).

Setting **DEFAULT\_ATTACH** to 1 has no effect, however, on any other types of indexes, whose pages are always stored in separate partitions from the data pages of the indexed table. These index types whose storage distribution is always different from that of their table include

- R-tree indexes,
- functional indexes,
- forest of trees indexes,
- fragmented indexes,
- and indexes on fragmented tables.

Index storage in the same partition as the data pages is supported only for nonfragmented B-tree indexes on nonfragmented tables.

If **DEFAULT\_ATTACH** is not set, then by default, any **CREATE INDEX** statement that does not specify **IN TABLE** as its Storage Options clause creates an index whose pages are stored in partitions separate from the data pages. This release of HCL OneDB™ can support existing indexes that were created by Version 7.x of HCL OneDB™.



**Important:** Future releases of HCL OneDB™ might not continue to support **DEFAULT\_ATTACH**. Developing new applications that depend on this deprecated feature is not recommended.

## DELIMIDENT environment variable

The **DELIMIDENT** environment variable specifies that strings enclosed between double quotation ( " ) marks are delimited database identifiers.

The **DELIMIDENT** environment variable is also supported on client systems, where it can be set to **y**, to **n**, or to no setting.

- **y** specifies that client applications must use single quotation ( ' ) symbols to delimit character strings, and must use double quotation ( " ) symbols only around delimited SQL identifiers, which can support a larger character set than is valid in undelimited identifiers. Letters within delimited strings or delimited identifiers are case-sensitive. This is the default value for OLE DB and .NET.
- **n** specifies that client applications can use double quotation ( " ) or single quotation ( ' ) symbols to delimit character strings, but not to delimit SQL identifiers. If the database server encounters a string delimited by double or single quotation symbols in a context where an SQL identifier is required, it issues an error. An owner name that qualifies an SQL identifier can be delimited by single quotation ( ' ) symbols. You must use a pair of the same quotation symbols to delimit a character string.

This is the default value for ESQL/C, JDBC, and ODBC. APIs that have ESQL/C as an underlying layer, such as the DataBlade® API (LIBDMI), and the C++ API, behave as ESQL/C, and use 'n' as the default if no value for DELIMIDENT is specified on the client system.

- Specifying the DELIMIDENT environment variable with no value on the client system requires client applications to use the DELIMIDENT setting that is the default for their application programming interface (API).

```
setenvDELIMIDENT
```

No value is required; **DELIMIDENT** takes effect if it exists, and it remains in effect while it is on the list of environment variables. Removing DELIMIDENT when it is set at the server level requires restarting the server.

Delimited identifiers can include white space (such as the phrase "**Vitamin E**") or can be identical to SQL keywords, (such as "**TABLE**" or "**USAGE**"). You can also use them to declare database identifiers that contain characters outside the default character set for SQL identifiers (such as "**Column #6**"). In the default locale, this set consists of letters, digits, and the underscore ( `_` ) symbol.

Even if DELIMIDENT is set, you can use single quotation ( `'` ) symbols to delimit authorization identifiers as the owner name component of a database object name, as in the following example:

```
RENAME COLUMN 'Owner'.table2.collum3 TO column3;
```

This example is an exception to the general rule that when **DELIMIDENT** is set, the SQL parser interprets character strings delimited by single quotation symbols as string literals, and interprets character strings delimited by double quotation symbols ( `"` ) as SQL identifiers.

*Database identifiers* (also called *SQL identifiers*) are names for database objects, such as tables and columns. *Storage identifiers* are names for storage objects, such as dbspaces, blobspaces, and sbspaces. You cannot use **DELIMIDENT** to declare storage identifiers that contain characters outside the default SQL character set.

Delimited identifiers are case sensitive. To use delimited identifiers, applications in must set **DELIMIDENT** at compile time and at run time.



**Important:** If **DELIMIDENT** is not already set, you should be aware that setting it can cause the failure of existing `.sql` scripts or client applications that use double ( `"` ) quotation marks in contexts other than delimiting SQL identifiers, such as delimiters of string literals. You must use single ( `'` ) rather than double quotation marks for delimited constructs that are not SQL identifiers if **DELIMIDENT** is set.

On UNIX™ systems that use the C shell and on which **DELIMIDENT** has been set, you can disable this feature (which causes anything between double quotation symbols to be interpreted as an SQL identifier) by the command:

```
unsetenv DELIMIDENT
```

## ENVIGNORE environment variable (UNIX™)

The **ENVIGNORE** environment variable can deactivate specified environment variable settings in the common (shared) configuration file, `onedb.rc`, and private environment-configuration file, `.informix`.

```
setenvENVIGNORE variable
```

**variable**

The name of an environment variable to be deactivated.

Use colon (:) symbols between consecutive *variable* names. For example, to ignore the **DBPATH** and **DBMONEY** entries in the environment-configuration files, enter the following command:

```
setenv ENVIGNORE DBPATH:DBMONEY
```

The common environment-configuration file is stored in `$ONEDB_HOME/etc/onedb.rc`.

The private environment-configuration file is stored in the home directory of the user as `.informix`.

For information about creating or modifying an environment-configuration file, see [Setting environment variables in a configuration file on page 140](#).

**ENVIGNORE** itself cannot be set in an environment-configuration file.

## FET\_BUF\_SIZE environment variable

The **FET\_BUF\_SIZE** environment variable can override the default setting for the size of the fetch buffer for all data types except BYTE and TEXT values. For ANSI databases, you must set transactions to READ ONLY for the **FET\_BUF\_SIZE** environment variable to improve performance, otherwise rows are returned one by one.

```
setenvFET_BUF_SIZESize
```

**size**

is a positive integer that is larger than the default buffer size, but no greater than 2147483648 (2GB), specifying the size (in bytes) of the fetch buffer that holds data retrieved by a query.

For example, to set a buffer size to 5,000 bytes on a UNIX™ system that uses the C shell, set **FET\_BUF\_SIZE** by entering the following command:

```
setenv FET_BUF_SIZE 5000
```

When **FET\_BUF\_SIZE** is set to a valid value, the new value overrides the default value (or any previously set value of **FET\_BUF\_SIZE**). The default setting for the fetch buffer is dependent on row size.

The processing of BYTE and TEXT values is not affected by **FET\_BUF\_SIZE**.

No error is raised if **FET\_BUF\_SIZE** is set to a value that is less than the default size or is larger than 2147483648 (2GB). In these cases, however, the invalid fetch buffer size is ignored, and the default size is in effect.

A valid **FET\_BUF\_SIZE** setting is in effect for the local database server and for any remote database server from which you retrieve rows through a distributed query in which the local server is the coordinator and the remote database is subordinate. The greater the size of the buffer, the more rows can be returned, and the less frequently the client application must wait while the database server returns rows. A large buffer can improve performance by reducing the overhead of filling the client-side buffer.

## IFMXMONGOAUTH environment variable

Set the **IFMXMONGOAUTH** environment variable to enable PAM authentication for MongoDB clients through the wire listener.

You can set the **IFMXMONGOAUTH** environment variable to any value or to no value.

```
setenv IFMXMONGOAUTH 1
```

Setting the **IFMXMONGOAUTH** environment variable is a prerequisite to configuring PAM authentication for MongoDB clients.

You can disable the **IFMXMONGOAUTH** environment variable with this command:

```
unsetenv IFMXMONGOAUTH
```

## IFX\_DEF\_TABLE\_LOCKMODE environment variable

The **IFX\_DEF\_TABLE\_LOCKMODE** environment variable can specify the default lock mode for database tables that are subsequently created without explicitly specifying the **LOCKMODE PAGE** or **LOCKMODE ROW** keywords. This feature is convenient if you must create several tables of the same lock mode. UNIX™ systems that use the C shell support the following syntax:

```
setenv IFX_DEF_TABLE_LOCKMODE { PAGE | ROW }
```

### PAGE

The default lock mode is page-level granularity. This value disables the LAST COMMITTED feature of COMMITTED READ.

### ROW

The default lock mode is row-level granularity.

Similar functionality is available by setting the **DEF\_TABLE\_LOCKMODE** parameter of the **ONCONFIG** file to **PAGE** or **ROW**. When a table is created or modified, any conflicting lock mode specifications are resolved according to the following descending (highest to lowest) order of precedence:

1. Explicit **LOCKMODE** specification of **CREATE TABLE** or **ALTER TABLE**
2. **IFX\_DEF\_TABLE\_LOCKMODE** environment variable setting
3. **DEF\_TABLE\_LOCKMODE** parameter setting in the **ONCONFIG** file
4. The system default lock mode (= page mode)

To make the **DEF\_TABLE\_LOCKMODE** setting the default mode (or to restore the system default if **DEF\_TABLE\_LOCKMODE** is not set) use the command:

```
unsetenv IFX_DEF_TABLE_LOCKMODE
```

If **IFX\_DEF\_TABLE\_LOCKMODE** is set in the environment of the database server before running **oninit**, then its scope is all sessions of the database server (just as if **DEF\_TABLE\_LOCKMODE** were set in the **ONCONFIG** file). If



**IFX\_DEF\_TABLE\_LOCKMODE** is set in the shell, or in the **\$HOME/.onedb** or **\$ONEDB\_HOME/etc/onedb.rc** files, then the scope is restricted to the current session (if you set it in the shell) or to the individual user.



**Important:** This has no effect on existing tables. If you specify *ROW* as the lock mode, the database will use this to restore, recover, or copy data. For tables that were created in *PAGE* mode, this might cause lock-table overflow or performance degradation.

## IFX\_DIRECTIVES environment variable

The **IFX\_DIRECTIVES** environment variable setting determines whether the optimizer allows query optimization directives from within a query. The **IFX\_DIRECTIVES** environment variable is set on the client.

You can specify either **ON** and **OFF** or **1** and **0** to set the environment variable.

```
setenvIFX_DIRECTIVES { 1 | 0 }
```

**1**

Optimizer directives accepted

**0**

Optimizer directives not accepted

The setting of the **IFX\_DIRECTIVES** environment variable overrides the value of the **DIRECTIVES** configuration parameter that is set for the database server. If the **IFX\_DIRECTIVES** environment variable is not set, however, then all client sessions will inherit the database server configuration for directives that the **ONCONFIG** parameter **DIRECTIVES** determines. The default setting for the **IFX\_DIRECTIVES** environment variable is **ON**.

For more information about the **DIRECTIVES** parameter, see the *HCL OneDB™ Administrator's Reference*. For more information about the performance impact of directives, see your *HCL OneDB™ Performance Guide*.

## IFX\_EXTDIRECTIVES environment variable

The **IFX\_EXTDIRECTIVES** environment variable specifies whether the query optimizer allows external query optimization directives from the **sysdirectives** system catalog table to be applied to queries in existing applications.

You have two options for setting the **IFX\_EXTDIRECTIVES** environment variable:

- Global, for all users:

On the server, set **IFX\_EXTDIRECTIVES** in the environment as user **informix** and then run the **oninit** command.

- Client specific:

On the client, set **IFX\_EXTDIRECTIVES** in the environment. When **IFX\_EXTDIRECTIVES** is set in the client environment, the client setting are used regardless of the server (global) setting.

You can determine the server setting using the **onstat -g env** command.

You can specify either `ON` and `OFF` or `1` and `0` to set the environment variable.

```
setenvIFX_DIRECTIVES { 1 | 0 }
```

**1**

External optimizer directives accepted

**0**

External optimizer directives not accepted

Queries within a given client application can use external directives if both the `EXT_DIRECTIVES` parameter in the configuration file of the database server and the `IFX_EXTDIRECTIVES` environment variable setting on the client system are both set to 1 or ON. If `IFX_EXTDIRECTIVES` is not set, external directives are supported only if the `ONCONFIG` parameter `EXT_DIRECTIVES` is set to 2. The following table summarizes the effect of valid `IFX_EXTDIRECTIVES` and `EXT_DIRECTIVES` settings on support for external optimizer directives.

**Table 57. Effect of `IFX_EXTDIRECTIVES` and `EXT_DIRECTIVES` settings on external directives**

	<code>EXT_DIRECTIVES = 0</code>	<code>EXT_DIRECTIVES = 1</code>	<code>EXT_DIRECTIVES = 2</code>
<b>IFX_EXTDIRECTIVES No setting</b>	OFF	OFF	ON
<b>IFX_EXTDIRECTIVES0 = OFF</b>	OFF	OFF	OFF
<b>IFX_EXTDIRECTIVES1 = ON</b>	OFF	ON	ON

The database server interprets any `EXT_DIRECTIVES` setting besides 1 or 2 (or no setting) as equivalent to OFF, disabling support for external directives. Any value of `IFX_EXTDIRECTIVES` other than 1 has the same effect for the client.

For information about how to define external optimizer directives, see the description of the `SAVE EXTERNAL DIRECTIVES` statement of SQL in the *HCL OneDB™ Guide to SQL: Syntax*. For more information about the `EXT_DIRECTIVES` configuration parameter, see the *HCL OneDB™ Administrator's Reference*. For more information about the performance impact of directives, see your *HCL OneDB™ Performance Guide*.

## IFX\_LARGE\_PAGES environment variable

The `IFX_LARGE_PAGES` environment variable specifies whether the database server can use large pages on platforms where the hardware and the operating system support large pages of shared memory. If this is enabled in the server environment, HCL OneDB™ can use the large pages for non-message shared memory segments that are located in physical memory.

The `IFX_LARGE_PAGES` environment variable is supported only on AIX®, Solaris, and Linux™ operating systems. The setting of `IFX_LARGE_PAGES` has no effect on HCL OneDB™ if the operating system does not support large pages, or if large pages are not configured on the system.

You can specify either `1` or `0` to set this environment variable.

```
setenvIFX_LARGE_PAGES { 1 | 0 }
```

**0**

The use of large pages is disabled. This is the default on AIX® systems.

**1**

The use of large pages is enabled. This is the default on Solaris and Linux™ systems.

The DBSA must use operating system commands to configure the large pages. See the operating system documentation for the configuration procedures.

HCL OneDB™ can use large pages for non-message shared memory segments that are locked in physical memory, if sufficient large pages are configured and available. The RESIDENT configuration parameter controls whether a shared memory segment is locked in physical memory, so that the segment cannot be swapped. If there are insufficient large pages to hold a segment, the segment might contain a mixture of large pages and regular pages.

On AIX® the large pages used by HCL OneDB™ are 16 MB in size.

On Linux™ x86\_64 the large pages used by HCL OneDB™ are defined by the `Hugepagesize` entry in the `/proc/meminfo` file.

HCL OneDB™ aligns the segment address and rounds up to the segment size automatically. In addition to messages regarding rounding, the server prints an informational message to the server log file whenever it attempts to use large pages to store a segment.

When **IFX\_LARGE\_PAGES** is enabled, the use of large pages can offer significant performance benefits in large memory configurations.

## IFX\_LOB\_XFERSIZE environment variable

Use the **IFX\_LOB\_XFERSIZE** environment variable to specify the number of bytes in a CLOB or BLOB data type to transfer from a client application to the database server before checking whether an error has occurred.

The error check occurs each time the specified number of bytes is transferred. If an error occurs, the remaining data is not sent and an error is reported. If no error occurs, the file transfer will continue until it finishes.

For example, if the value of **IFX\_LOB\_XFERSIZE** is set to 10485760 (10 MB), then error checking will occur after every 10485760 bytes of the CLOB or BLOB data is sent. If **IFX\_LOB\_XFERSIZE** is not set, the error check occurs after the entire BLOB or CLOB data is transferred.

The valid range for **IFX\_LOB\_XFERSIZE** is from 1 to 9223372036854775808 bytes. The **IFX\_LOB\_XFERSIZE** environment variable is set on the client.

```
setenvIFX_LOB_XFERSIZEvalue
```

### **value**

the number of bytes in a CLOB or BLOB to transfer from a client application to the database server before checking whether an error has occurred

You should adjust the value of **IFX\_LOB\_XFERSIZE** to suit your environment. Set **IFX\_LOB\_XFERSIZE** low enough so that transmission errors of large BLOB or CLOB data types are detected early, but not so low that excessive network resources are used.

## IFX\_LONGID environment variable

The **IFX\_LONGID** environment variable setting and the version number of the client application determine whether a given client application is capable of handling long identifiers. (Older versions of HCL OneDB™ restricted SQL identifiers to 18 or fewer bytes; *long identifiers* can have up to 128 bytes when **IFX\_LONGID** is set.) Valid **IFX\_LONGID** values are `1` and `0`.

```
setenvIFX_LONGID { 1 | 0 }
```

**1**

Client supports long identifiers.

**0**

Client cannot support long identifiers.

When **IFX\_LONGID** is set to zero, applications display only the first 18 bytes of long identifiers, without indicating (by `+`) that truncation has occurred.

If **IFX\_LONGID** is unset or is set to a value other than `1` or `0`, the determination is based on the internal version of the client application. If the (server-based) version is not less than 9.0304, or is in the (CSDK-based) range  $2.90 \leq version < 4.0$ , the client is considered capable of handling long identifiers. Otherwise, the client application is considered incapable.

The **IFX\_LONGID** setting overrides the internal version of the client application. If the client cannot handle long identifiers despite a newer version number, set **IFX\_LONGID** to `0`. If the client version can handle long identifiers despite an older version number, set **IFX\_LONGID** to `1`.

If you set **IFX\_LONGID** on the client, the setting affects only that client. If you start the database server with **IFX\_LONGID** set, all client applications use that setting by default. If **IFX\_LONGID** is set to different values on the client and on the database server, however, the client setting takes precedence.



**Important:** ESQL executables that have been built with the `-static` option using the `libos.a` library version that does not support long identifiers cannot use the **IFX\_LONGID** environment variable. You must recompile such applications with the new `libos.a` library that includes support for long identifiers. Executables that use shared libraries (no `-static` option) can use **IFX\_LONGID** without recompilation provided that they use the new `libifos.so` that provides support for long identifiers. For details, see your ESQL product publication.

## IFX\_NETBUF\_PVTPPOOL\_SIZE environment variable (UNIX™)

Use the **IFX\_NETBUF\_PVTPPOOL\_SIZE** environment variable to specify the maximum size of the free (unused) private network buffer pool for each database server session.

```
setenvIFX_NETBUF_PVTPPOOL_SIZECOUNT
```

**count**

an integer specifying the number of units (buffers) in the pool.

The default size is 1 buffer. If **IFX\_NETBUF\_PVTPPOOL\_SIZE** is set to 0, then each session obtains buffers from the free global network buffer pool. You must specify the value in decimal form.

**IFX\_NETBUF\_SIZE environment variable**

Use the **IFX\_NETBUF\_SIZE** environment variable to configure the network buffers to the optimum size. This environment variable specifies the size of all network buffers in the free (unused) global pool and the private network buffer pool for each database server session.

```
setenvIFX_NETBUF_SIZESize
```

**size**

is the integer size (in bytes) for one network buffer.

The default size is 4 KB (4,096 bytes). The maximum size is 64 KB (65,536 bytes) and the minimum size is 512 bytes. You can specify the value in hexadecimal or decimal form.



**Tip:** You cannot set a different size for each session.

**IFX\_NO\_SECURITY\_CHECK environment variable (UNIX™)**

The **IFX\_NO\_SECURITY\_CHECK** environment variable allows user **informix** or **root** to complete operations with a database server instance even when the HCL OneDB™ utilities detect that the \$ONEDB\_HOME path is not secure. Do not use this environment variable unless your system setup makes it absolutely necessary to do so.

The purpose of **IFX\_NO\_SECURITY\_CHECK** is for environments where the database server started but while running it detects that the runtime path is not secure anymore. In this case, a superuser might be required to stop the database server in order to remedy the security flaw. With this environment variable, either user **informix** or **root** can use the onmode utility to shut down a nonsecure HCL OneDB™ instance, which would otherwise not be possible because key programs do not run when the \$ONEDB\_HOME path is not secure.

There is some risk in using this environment variable, but in some circumstances it might be necessary to remedy a bigger security problem. The requirement that only user **informix** or **root** can invoke **IFX\_NO\_SECURITY\_CHECK** makes it unlikely that an illegitimate user would be able to run it.

To use this environment variable, set it to any non-empty string.

```
setenvIFX_NO_SECURITY_CHECK1
```

1

Any value entered here when running this environment variable disables the onsecurity utility.



**Important:** Turn off this environment variable after you have finished troubleshooting the security problem.

## IFX\_NO\_TIMELIMIT\_WARNING environment variable

Trial or evaluation versions of HCL OneDB™ software products, which cease to function when some time limit has elapsed since the software was installed, by default issue warning messages that tell users when the license will expire. If you set the **IFX\_NO\_TIMELIMIT\_WARNING** environment variable, however, the time-limited software does not issue these warning messages.

```
setenvIFX_NO_TIMELIMIT_WARNING
```

For users who dislike viewing warning messages, this feature is an alternative to redirecting the error output. Setting **IFX\_NO\_TIMELIMIT\_WARNING** has no effect, however, on when a time-limited license expires; the software ceases to function at the same point in time when it would if this environment variable had not been set. If you do set **IFX\_NO\_TIMELIMIT\_WARNING**, users will not see potentially annoying warnings about the impending license expiration, but some users might be annoyed at you when the database server (or whatever software has a time-limited license) ceases to function without any warning.

## IFX\_NODBPROC environment variable

The **IFX\_NODBPROC** environment variable lets you prevent the database server from running the `sysdbopen()` or `sysdbclose()` procedure. These procedures cannot be run if this environment variable is set to any value.

```
setenvIFX_NODBPROCstring
```

### *string*

Any value prevents the database server from running `sysdbopen()` or `sysdbclose()`.

## IFX\_NOT\_STRICT\_THOUS\_SEP environment variable

HCL OneDB™ requires the thousands separator to have 3 digits following it. For example, 1,000 is considered correct, and 1,00 is considered wrong. In previous releases, both formats were considered correct.

```
setenvIFX_NOT_STRICT_THOUS_SEPn
```

### *n*

Set *n* to 1 for the behavior in previous releases, which is that the thousands separator can have fewer than three digits following it.

## IFX\_ONTAPE\_FILE\_PREFIX environment variable

When `TAPEDEV` and `LTAPEDEV` specify directories, use the **IFX\_ONTAPE\_FILE\_PREFIX** environment variable to specify a prefix for backup file names that replaces the `hostname_servernum` format. If no value is set, file names are `hostname_servernum_Ln` for levels and `hostname_servernum_Lognnnnnnnnnn` for log files.

If you set the value of `IFX_ONTAPE_FILE_PREFIX` to `My_Backup`, the backup file names have the following names:

- My\_Backup\_L0
- My\_Backup\_L1
- My\_Backup\_L2
- My\_Backup\_Log0000000001
- My\_Backup\_Log0000000002

```
setenvIFX_ONTAPE_FILE_PREFIXString
```

### **string**

The prefix to use for the names of backup files.

## IFX\_PAD\_VARCHAR environment variable

The **IFX\_PAD\_VARCHAR** environment variable setting controls how the database server sends and receives VARCHAR and NVARCHAR data values. Valid **IFX\_PAD\_VARCHAR** values are **1** and **0**.

```
setenvIFX_PAD_VARCHAR { 1 | 0 }
```

### **1**

Transmit the entire structure, up to the declared *max* size.

### **0**

Transmit only the portion of the structure containing data.

For example, to send the string "ABC" from a column declared as NVARCHAR(255) when **IFX\_PAD\_VARCHAR** is set to 0 would send 3 bytes.

If the setting were 1 in the previous example, however, the number of bytes sent would be 255 bytes.

The effect **IFX\_PAD\_VARCHAR** is context-sensitive. In a low-bandwidth network, a setting of 0 might improve performance by reducing the total volume of transmitted data. But in a high-bandwidth network, a setting of 1 might improve performance, if the CPU time required to process variable-length packets were greater than the time required to send the entire character stream. In cross-server distributed operations, this setting has no effect, and padding characters are dropped from VARCHAR or NVARCHAR values that are passed between database servers.

## IFX\_SMX\_TIMEOUT environment variable

Use the **IFX\_SMX\_TIMEOUT** environment variable to specify the maximum number of seconds for a high-availability replication (HDR), remote stand-alone (RS) or shared disk (SD) secondary server to wait for a message from the primary server in a Server Multiplexer Group (SMX) connection.

```
setenvIFX_SMX_TIMEOUTvalue
```

### **value**

Any positive numeric value for the number of seconds or **-1** to disable this environment variable. There is no upper limit to the number of seconds that you can specify.

**default value**

10

For example, to indicate that the secondary server should wait for no more than 60 seconds, specify:

```
setenv IFX_SMX_TIMEOUT 60
```

If the secondary server does not receive any message after the number of seconds specified in the IFX\_SMX\_TIMEOUT environment variable and after the number of cycles specified in the IFX\_SMX\_TIMEOUT\_RETRY environment variable have completed, the secondary server will print the error message in the **online.log** and close the SMX connection. If an SMX timeout message is in the **online.log**, you might need to increase the IFX\_SMX\_TIMEOUT value, the IFX\_SMX\_TIMEOUT\_RETRY value, or both of these values and restart secondary node.

This environment variable applies only to secondary servers. If you set this environment variable on the primary server, it will become effective only if the primary server becomes a secondary server after a failure.

## IFX\_SMX\_TIMEOUT\_RETRY environment variable

Use the IFX\_SMX\_TIMEOUT\_RETRY environment variable to specify the number of times that the high-availability replication (HDR), remote standalone (RS) or shared disk (SD) secondary server will repeat the wait cycle specified by the IFX\_SMX\_TIMEOUT environment variable if a response from the primary server has not been received.

```
setenv IFX_SMX_TIMEOUT_RETRY value
```

**value**

Any positive numeric value

**default value**

6

For example, to indicate that the amount of time specified in the IFX\_SMX\_TIMEOUT configuration parameter should be repeated up to 20 times if a response from the primary server has not been received, specify:

```
setenv IFX_SMX_TIMEOUT_RETRY 20
```

If the secondary server does not receive any message after the number of seconds specified in the IFX\_SMX\_TIMEOUT environment variable and after the number of cycles specified in the IFX\_SMX\_TIMEOUT\_RETRY environment variable have completed, the secondary server will print the error message in the **online.log** and close the SMX connection. If an SMX timeout message is in the **online.log**, you might need to increase the IFX\_SMX\_TIMEOUT value, the IFX\_SMX\_TIMEOUT\_RETRY value, or both of these values and restart secondary node.

This environment variable applies only to secondary servers. If you set this environment variable on the primary server, it will become effective only if the primary server becomes a secondary server after a failure.



## IFX\_UNLOAD\_EILSEQ\_MODE environment variable

Use the IFX\_UNLOAD\_EILSEQ\_MODE environment variable to help migrate databases from HCL OneDB™ Version 10 to Version 11.50 or 11.70, where character data might be encoded with a codeset that is different than the codeset used to create the Version 10 database.

In earlier versions of HCL OneDB™, it was possible to load character data into a database that did not match the locale and codeset of the database. For example you could load Chinese data into a database created with the DB\_LOCALE=en\_US.8859-1 codeset. In newer versions of HCL OneDB™, to insert Chinese data you would need a database created with the Chinese (DB\_LOCALE=zh\_tw.big5 locale and codeset).

**!** **Important:** For databases created with Version 10 and Client SDK 2.4, when you attempt to unload the invalid character data an error occurs unless you have set this environment variable. The IFX\_UNLOAD\_EILSEQ\_MODE environment variable enables DB-Access, dbexport, and High Performance Loader (HPL) to unload character and bypass the GLS validation that normally occurs when you unload data by using the Version 11.50 and 11.70 tools.

To use this environment variable, set it to any non-empty string.

```
setenv IFX_UNLOAD_EILSEQ_MODE value
```

### value

Any alpha or numeric value. For example: yes, true, or 1.

This environment variable takes effect when character data is being fetched or retrieved from the database.

```
setenv IFX_UNLOAD_EILSEQ_MODE 1
setenv IFX_UNLOAD_EILSEQ_MODE yes
setenv IFX_UNLOAD_EILSEQ_MODE on
```

This environment variable is similar to setting the EILSEQ\_COMPAT\_MODE configuration parameter in the ONCONFIG file. The configuration parameter affects character data that is inserted into the database, whereas the IFX\_UNLOAD\_EILSEQ\_MODE environment variable affects character data that is unloaded from the database.

## IFX\_UPDDESC environment variable

You must set the IFX\_UPDDESC environment variable at execution time before you can do a DESCRIBE of an UPDATE statement.

```
setenv IFX_UPDDESC value
```

### value

is any non-NULL value.

A NULL value (here meaning that IFX\_UPDDESC is not set) disables the describe-for-update feature. Any non-NULL value enables the feature.

## IFX\_XASTDCOMPLIANCE\_XAEND environment variable

In earlier releases of HCL OneDB™, an internal rollback of a global transaction freed the transaction. In releases later than Version 9.40, however, the default behavior after an internal rollback is not to free the global transaction until an explicit rollback, as required by the X/Open XA standard. By setting the `DISABLE_B162428_XA_FIX` configuration parameter to 1, you can restore the legacy behavior as the default for all sessions.

The `IFX_XASTDCOMPLIANCE_XAEND` environment variable can override the configuration parameter for the current session, using the following syntax. Valid `IFX_XASTDCOMPLIANCE_XAEND` values are 1 and 0.

```
setenvIFX_XASTDCOMPLIANCE_XAEND { 1 | 0 }
```

0

Frees global transactions only after an explicit rollback

1

Frees global transactions after any rollback

This environment variable can be particularly useful when the server instance is disabled for new behavior by the `DISABLE_B162428_XA_FIX` configuration parameter, but one client requires the new behavior. Setting this environment variable to zero supports the new behavior in the current session.

## IFX\_XFER\_SHMBASE environment variable

An alternative base address for a utility to attach the server shared memory segments.

```
setenvIFX_XFER_SHMBASE { address }
```

**address**

Valid address in hexadecimal

After the database server allocates shared memory, the database server might allocate multiple contiguous OS shared memory segments. The client utility that connects to shared memory must attach all those OS segments contiguously also. The utility might have some other shared objects (for example, the `xbsa` library in `onbar`) loaded at the address where the server has shared memory segment attached. To workaround this situation, you can specify a different base address in the environment variable `IFX_XFER_SHMBASE` for the utility to attach the shared memory segments. The `onstat`, `onmode`, and `oncheck` utilities must attach to exact same shared memory base as `oninit`. Setting `IFX_XFER_SHMBASE` is not an option for these utilities.

## IFXRESFILE environment variable (Linux™)

Set the `IFXRESFILE` environment variable to the path and name of your response file before running an RPM-method installation command. If you want to accept the default HCL OneDB™ installation settings, do not use this environment variable.

```
setenvIFXRESFILEpath_filename.ini
```

***path\_filename***

Specifies the path and name of the response file (*.ini* file) in which you changed the default installation settings of the `bundle.ini` file shipped with the installation media

For information about creating a response file by customizing the `bundle.ini` file, see the *HCL OneDB™ Installation Guide* *HCL OneDB™ Installation Guide for UNIX™, Linux™, and Mac OS X*.

**ONEDB\_C environment variable (UNIX™)**

The **ONEDB\_C** environment variable specifies the filename or pathname of the C compiler to be used to compile files that generates. The setting takes effect only during the C compilation stage.

If **ONEDB\_C** is not set, the default compiler on most systems is **cc**.



**Tip:** On Windows™, you pass either `-mcc` or `-bcc` options to the `esql` preprocessor to use either the Microsoft™ or Borland C compilers.

```
setenv ONEDB_C { compiler | pathname }
```

***compiler***

The file name of the C compiler.

***pathname***

The full path name of the C compiler.

For example, to specify the GNU C compiler, enter the following command:

```
setenv ONEDB_C gcc
```



**Important:** If you use **gcc**, be aware that the database server assumes that strings are writable, so you must compile by using the `-fwritable-strings` option. Failure to do so can produce unpredictable results, possibly including core dumps.

**ONEDB\_CMNAME environment variable**

If the Connection Manager raises an event alarm, the **ONEDB\_CMNAME** environment variable is used to store the name of the Connection Manager instance that raised the alarm. The environment variable is set automatically by the Connection Manager.

The **ONEDB\_CMNAME** environment variable corresponds to the **NAME** parameter in the Connection Manager configuration file. The environment variable is used by the `CMALARMPROGRAM` program to determine the Connection Manager instance responsible for the event alarm. You can also use the environment variable in your own Connection Manager event alarm handler.

The environment variable is set automatically by the Connection Manager and should not be modified.

## ONEDB\_CMCONUNITNAM environment variable

If the Connection Manager raises an event alarm, the ONEDB\_CMCONUNITNAM environment variable is used to store the name of the Connection Manager connection unit that raised the alarm. The environment variable is set automatically by the Connection Manager.

The ONEDB\_CMCONUNITNAM environment variable corresponds to the connection unit name parameter in the Connection Manager configuration file. The environment variable is used by the CMALARMPROGRAM program to determine the Connection Manager instance responsible for the event alarm. You can also use the environment variable in your own Connection Manager event alarm handler.

The environment variable is set automatically by the Connection Manager and should not be modified.

## CONNECT\_RETRIES environment variable

The **CONNECT\_RETRIES** environment variable sets a limit on the maximum number of connection attempts that can be made to each database server by the client after the initial connection attempt fails. These attempts are made within the time limit that the **CONNECT\_TIMEOUT** setting specifies.

```
setenvCONNECT_RETRIESCOUNT
```

### *count*

The number of additional attempts to connect to each database server after the initial connection attempt fails.

For example, the following command sets **CONNECT\_RETRIES** to specify three connection attempts after the initial attempt:

```
setenv CONNECT_RETRIES 3
```

The default value for **CONNECT\_RETRIES** is one attempt after the initial connection attempt.

## Order of precedence among **CONNECT\_RETRIES** settings

When you specify a setting for the **CONNECT\_RETRIES** client environment variable, it overrides any **CONNECT\_RETRIES** configuration parameter setting in the `onconfig` file.

If the SET ENVIRONMENT statement specifies a setting for the **CONNECT\_RETRIES** session environment option, however, the SQL statement setting overrides the **CONNECT\_RETRIES** client environment variable setting for subsequent connection attempts during the current session. The SET ENVIRONMENT **CONNECT\_RETRIES** setting has no effect on other sessions.

In summary, this is the ascending order (lowest to highest) of the methods for setting a limit on attempts for a connection to a database server:

- **CONNECT\_RETRIES** configuration parameter
- **CONNECT\_RETRIES** client environment variable
- SET ENVIRONMENT **CONNECT\_RETRIES** statement of SQL

The **CONNECT\_TIMEOUT** setting takes precedence over the **CONNECT\_RETRIES** setting. Connection attempts can end after the **CONNECT\_TIMEOUT** value is exceeded, but before the **CONNECT\_RETRIES** value is reached. For more information about

restricting the time available to establish a connection to a database server, see [CONNECT\\_TIMEOUT environment variable on page 197](#)

## CONNECT\_TIMEOUT environment variable

The **CONNECT\_TIMEOUT** environment variable specifies the number of seconds the **CONNECT** statement attempts to establish a connection to a database server before returning an error. If you set no value, the default of 60 seconds can typically support a few hundred concurrent client connections. However, some systems might encounter few connection errors with a value as low as 15. The total distance between nodes, hardware speed, the volume of traffic, and the concurrency level of the network can all affect what value you should set to optimize **CONNECT\_TIMEOUT**.

The **CONNECT\_TIMEOUT** and **CONNECT\_RETRIES** environment variables let you configure your client-side connection capability to retry the connection instead of returning a **-908** error.

```
setenvCONNECT_TIMEOUTseconds
```

### **seconds**

Represents the minimum number of seconds spent in attempts to establish a connection to a database server.

For example, enter this command to set **CONNECT\_TIMEOUT** to 60 seconds:

```
setenv CONNECT_TIMEOUT 60
```

If **CONNECT\_TIMEOUT** is set to 60 and **CONNECT\_RETRIES** is set to 3, attempts to connect to the database server (after the initial attempt at 0 seconds) are made at 20, 40, and 60 seconds, if necessary, before aborting. This 20-second interval is the result of **CONNECT\_TIMEOUT** divided by **CONNECT\_RETRIES**. If you set the **CONNECT\_TIMEOUT** value to zero, the database server automatically uses the default value of 60 seconds.

If the **CONNECT** statement must search **DBPATH**, the **CONNECT\_RETRIES** setting specifies the number of additional connection attempts that can be made for each database server entry in **DBPATH**.

- All appropriate servers in the **DBPATH** setting are accessed at least once, even if the **CONNECT\_TIMEOUT** value is exceeded. Thus, the **CONNECT** statement might take longer than the **CONNECT\_TIMEOUT** time limit to return an error that indicates connection failure or that the database was not found.
- The **CONNECT\_TIMEOUT** value is divided among the number of database server entries that are specified in **DBPATH**. Thus, if **DBPATH** contains numerous servers, increase the **CONNECT\_TIMEOUT** value accordingly. For example, if **DBPATH** contains three entries, to spend at least 30 seconds attempting each connection, set **CONNECT\_TIMEOUT** to 90.

The **CONNECT\_TIMEOUT** and **CONNECT\_RETRIES** environment variables can be modified with the **onutil** SET command, as in the following example:

```
% onutil
1> SET CONNECT_TIMEOUT 120;
Dynamic Configuration completed successfully
2> SET CONNECT_RETRIES 10;
Dynamic Configuration completed successfully
```

## Order of precedence among CONNECT\_TIMEOUT settings

When you specify a setting for the **CONNECT\_TIMEOUT** client environment variable, it overrides the **CONNECT\_TIMEOUT** configuration parameter settings in the `onconfig` file for the current session.

If the SET ENVIRONMENT statement specifies a setting for the **CONNECT\_RETRIES** session environment option, however, the SQL statement setting overrides the **CONNECT\_RETRIES** client environment variable setting for subsequent connection attempts during the current session. The SET ENVIRONMENT **CONNECT\_RETRIES** setting has no effect on other sessions.

In summary, this is the ascending order (lowest to highest) of the methods for setting an upper limit on the amount of time that a **CONNECT** statement can spend attempting to connect to a database server:

- **CONNECT\_TIMEOUT** configuration parameter
- **CONNECT\_TIMEOUT** client environment variable
- SET ENVIRONMENT **CONNECT\_TIMEOUT** statement of SQL.

**CONNECT\_TIMEOUT** takes precedence over the **CONNECT\_RETRIES** setting. Connection attempts can end after the **CONNECT\_TIMEOUT** value is exceeded, but before the **CONNECT\_RETRIES** value is reached.

## ONEDB\_CPPMAP environment variable

Set the **ONEDB\_CPPMAP** environment variable to specify the fully qualified pathname of the map file for C++ programs. Information in the map file includes the database server type, the name of the shared library that supports the database object or value object type, the library entry point for the object, and the C++ library for which an object was built.

```
setenvONEDB_CPPMAPpathname
```

### **pathname**

The directory path where the C++ map file is stored.

The map file is a text file that can have any filename. You can specify several map files, separated by colons (:) on UNIX™ or semicolons (;) on Windows™.

On UNIX™, the default map file is `$ONEDB_HOME/etc/c++map`. On Windows™, the default map file is `%ONEDB_HOME%\etc\c++map`.

## ONEDB\_HOME environment variable

The **ONEDB\_HOME** environment variable specifies the directory that contains the subdirectories in which your product files are installed. You must always set **ONEDB\_HOME**. Verify that **ONEDB\_HOME** is set to the full pathname of the directory in which you installed your database server. If you have multiple versions of a database server, set **ONEDB\_HOME** to the appropriate directory name for the version that you want to access. For information about when to set **ONEDB\_HOME**, see your *HCL OneDB™ Installation Guide*.

```
setenvONEDB_HOME\pathname
```

### **pathname**

is the directory path where the product files are installed.

To set **ONEDB\_HOME** to **usr/informix/**, for example, as the installation directory, enter the following command:

```
setenv ONEDB_HOME /usr/informix
```

## IONEDB\_OPCACHE environment variable

The **IONEDB\_OPCACHE** environment variable can specify the size of the memory cache for the staging-area blob space of the client application.

```
setenv IONEDB_OPCACHE kilobytes
```

### *kilobytes*

Specifies the value you set for the optical memory cache.

Set the **IONEDB\_OPCACHE** environment variable by specifying the size of the memory cache in KB. The specified size must be equal to or smaller than the size of the system-wide configuration parameter, **OPCACHEMAX**.

If you do not set **IONEDB\_OPCACHE**, the default cache size is 128 kilobytes or the size specified in the configuration parameter **OPCACHEMAX**. The default for **OPCACHEMAX** is 0. If you set **IONEDB\_OPCACHE** to a value of 0, Optical Subsystem does not use the cache.

## ONEDB\_SERVER environment variable

The **ONEDB\_SERVER** environment variable specifies the default database server to which an explicit or implicit connection is made by an SQL API client, the DB-Access utility, or other HCL OneDB™ products.

This environment variable must be set before you can use HCL OneDB™ client products. It has the following syntax.

```
setenv ONEDB_SERVER dbservername
```

### *dbservername*

is the name of the default database server.

The value of **ONEDB\_SERVER** can be a local or remote server, but must correspond to a valid *dbservername* entry in the **\$ONEDB\_HOME/etc/sqlhosts** file on the computer running the application. The *dbservername* must begin with a lower-case letter and cannot exceed 128 bytes. It can include any printable characters except uppercase characters, field delimiters (blank space or tab), the newline character, and the hyphen (or minus) symbol.

For example, this command specifies the **coral** database server as the default:

```
setenv ONEDB_SERVER coral
```

**ONEDB\_SERVER** specifies the database server to which an application connects if the **CONNECT DEFAULT** statement is executed. It also defines the database server to which an initial implicit connection is established if the first statement in an application is not a **CONNECT** statement.



**Important:** You must set **ONEDB\_SERVER** even if the application or DB-Access does not use implicit or explicit default connections.

## ONEDB\_SHMBASE environment variable (UNIX™)

The **ONEDB\_SHMBASE** environment variable affects only client applications connected to HCL OneDB™ databases that use the interprocess communications (IPC) shared-memory (**ipcshm**) protocol.



**Important:** Resetting **ONEDB\_SHMBASE** requires a thorough understanding of how the application uses memory. Normally you do not reset **ONEDB\_SHMBASE**.

**ONEDB\_SHMBASE** specifies where shared-memory communication segments are attached to the client process so that client applications can avoid collisions with other memory segments that it uses. If you do not set **ONEDB\_SHMBASE**, the memory address of the communication segments defaults to an implementation-specific value such as `0x800000`.

```
setenv ONEDB_SHMBASE value
```

### **value**

is an integer (in KB) used to calculate the memory address.

The database server calculates the memory address where segments are attached by multiplying the value of **ONEDB\_SHMBASE** by `1,024`. For example, on a system that uses the C shell, you can set the memory address to the value `0x800000` by entering the following command:

```
setenv ONEDB_SHMBASE 8192
```

For more information, see your *HCL OneDB™ Administrator's Guide* and the *HCL OneDB™ Administrator's Reference*.

## ONEDB\_SQLHOSTS environment variable

The **ONEDB\_SQLHOSTS** environment variable specifies where the SQL client or the database server can find connectivity information.

```
setenv ONEDB_SQLHOSTS pathname
```

### **pathname**

The full path name of the connectivity information file.



**UNIX:** Default = `ONEDB_HOME/etc/sqlhosts`



**Windows server:** Default = `ONEDB_HOME%\etc\sqlhosts.%ONEDB_SERVER%`

For example, the following command overrides the default location and specifies that the `mysqlhosts` file is in the `/work/envt` directory:

```
setenv ONEDB_SQLHOSTS /work/envt/mysqlhosts
```



**Windows™ client: Windows™:** The `ONEDB_SQLHOSTS` environment variable points to the computer whose registry contains the `SQLHOSTS` subkey. For example, the following command instructs the Windows™ client to look for connectivity information in the registry of a computer named **arizona**:

```
set ONEDB_SQLHOSTS = \\arizona
```

## ONEDB\_STACKSIZE environment variable

The `ONEDB_STACKSIZE` environment variable specifies the stack size (in KB) that is applied to all client processes. Any value that you set for `ONEDB_STACKSIZE` in the client environment is ignored by the database server.

```
setenvONEDB_STACKSIZEsize
```

### size

is an integer, setting the stack size (in KB) for SQL client threads.

For example, to decrease the `ONEDB_STACKSIZE` to 20 KB, enter the following command:

```
setenv ONEDB_STACKSIZE 20
```

If `ONEDB_STACKSIZE` is not set, the stack size is taken from the database server configuration parameter `STACKSIZE` or else defaults to a platform-specific value. The default stack size value for the primary thread of an SQL client is 32 KB for nonrecursive database activity.



**Warning:** For instructions on setting this value, see the *HCL OneDB™ Administrator's Reference*. If you incorrectly set the value of `ONEDB_STACKSIZE`, it can cause the database server to fail.

## ONEDB\_TERM environment variable (UNIX™)

The `ONEDB_TERM` environment variable specifies whether DB-Access should use the information in the `terminfo` directory or the `termcap` file.

On character-based systems, the `terminfo` directory and `termcap` file determine terminal-dependent keyboard and screen capabilities, such as the operation of function keys, color and intensity attributes in screen displays, and the definition of window borders and graphic characters.

```
setenvONEDB_TERM { terminfo | termcap }
```

If `ONEDB_TERM` is not set, the default setting is `terminfo`.

The `terminfo` directory contains a file for each terminal name that has been defined. The `terminfo` setting for `ONEDB_TERM` is supported only on computers that provide full support for the UNIX™ System V `terminfo` library. For details, see the machine notes file for your product.

When DB-Access is installed on your system, a `termcap` file is placed in the `etc` subdirectory of `$ONEDB_HOME`. This file is a superset of an operating-system `termcap` file. You can use the `termcap` file that the database server supplies, the system `termcap` file, or a `termcap` file that you create. You must set the `TERMCAP` environment variable if you do not use the

default `termcap` file. For information about setting the **TERMCAP** environment variable, see [TERMCAP environment variable \(UNIX\) on page 218](#).

## INF\_ROLE\_SEP environment variable

The **INF\_ROLE\_SEP** environment variable configures the security feature of role separation when the database server is installed or reinstalled on UNIX™ systems. Role separation enforces separating administrative tasks by people who run and audit the database server. After the installation is complete, **INF\_ROLE\_SEP** has no effect. If **INF\_ROLE\_SEP** is not set, then user **informix** (the default) can perform all administrative tasks.

```
setenvINF_ROLE_SEP/n
```

*n*

is any positive integer.

On Windows™, the install process asks whether you want to enable role separation regardless of the setting of **INF\_ROLE\_SEP**. To enable role separation for database servers on Windows™, select the role-separation option during installation.

If **INF\_ROLE\_SEP** is set when HCL OneDB™ is installed on a UNIX™ platform, role separation is implemented and a separate group is specified to serve each of the following responsibilities:

- The Database Server Administrator (DBSA)
- The Audit Analysis Officer (AAO)
- The standard user

On UNIX™, you can establish role separation by changing the group that owns the `aaodir`, `dbssadir`, or `etc` directories at any time after the installation is complete. You can disable role separation by resetting the group that owns these directories to **informix**. You can have role separation enabled, for example, for the Audit Analysis Officer (AAO) without having role separation enabled for the Database Server Administrator (DBSA).

For more information about the security feature of role separation, see the *HCL OneDB™ Security Guide*. To learn how to configure role separation when you install your database server, see your *HCL OneDB™ Installation Guide*.

## INTERACTIVE\_DESKTOP\_OFF environment variable (Windows™)

This environment variable lets you prevent interaction with the Windows™ desktop when an SPL routine executes a `SYSTEM` command.

```
setenvINTERACTIVE_DESKTOP_OFF { 1 | 0 }
```

If **INTERACTIVE\_DESKTOP\_OFF** is `1` and an SPL routine attempts to interact with the desktop (for example, with the `notepad.exe` or `cmd.exe` program), the routine fails unless the user is a member of the **Administrators** group.

The valid settings (`1` or `0`) have the following effects:

**1**

Prevents the database server from acquiring desktop resources for the user executing the stored procedure

**0**

SYSTEM commands in a stored procedure can interact with the desktop. This is the default value.

Setting **INTERACTIVE\_DESKTOP\_OFF** to **1** allows an SPL routine that does not interact with the desktop to execute more quickly. This setting also allows the database server to simultaneously call a greater number of SYSTEM commands because the command no longer depends on a limited operating- system resource (Desktop and WindowStation handles).

## ISM\_COMPRESSION environment variable

Use the ISM\_COMPRESSION environment variable to specify whether the HCL® OneDB® Storage Manager should use a data-compression algorithm to store and retrieve data.

```
setenvISM_COMPRESSION { TRUE | FALSE }
```

If **ISM\_COMPRESSION** is set to **TRUE** in the environment of the ON-Bar process that makes a request, the ISM server uses the data-compression algorithm.

If **ISM\_COMPRESSION** is set to **FALSE** or is not set, the ISM server does not use compression.

## ISM\_DEBUG\_FILE environment variable

Use the **ISM\_DEBUG\_FILE** environment variable in the HCL® OneDB® Storage Manager server environment to specify where to write XBSA messages.

```
setenvISM_DEBUG_FILEpathname
```

### *pathname*

Specifies the location of the XBSA message log file.

If you do not set **ISM\_DEBUG\_FILE**, the XBSA message log is located in the `$ONEDB_HOME/ism/applogs/xbsa.messages` directory on UNIX™, or in the `c:\nsr\applogs\xbsa.messages` directory on Windows™ systems.

## ISM\_DEBUG\_LEVEL environment variable

Use the **ISM\_DEBUG\_LEVEL** environment variable in the ON-Bar environment to control the level of reporting detail recorded in the XBSA messages log. The XBSA shared library writes to this log.

```
setenvISM_DEBUG_LEVELvalue
```

### *value*

specifies the level of reporting detail, where  $1 \leq value \leq 9$ .

If **ISM\_DEBUG\_LEVEL** is not set, has a null value, or has a value outside this range, the default detail level is **1**. A detail level of **0** suppresses all XBSA debugging records. A detail level of **1** reports only XBSA failures.

## ISM\_ENCRYPTION environment variable

Use the **ISM\_ENCRYPTION** environment variable in the ON-Bar environment to specify whether HCL® OneDB® Storage Manager uses data encryption.

```
setenv ISM_ENCRYPTION { XOR | NONE | TRUE }
```

Three settings of **ISM\_ENCRYPTION** are supported:

**XOR**

Uses encryption.

**NONE**

Does not use encryption.

**TRUE**

Uses encryption.

If **ISM\_ENCRYPTION** is set to NONE or is not set, the ISM server does not use encryption.

If the **ISM\_ENCRYPTION** is set to TRUE or XOR in the environment of the ON-Bar process that makes a request, the ISM server uses encryption to store or retrieve the data specified in that request.

## ISM\_MAXLOGSIZE environment variable

Use the **ISM\_MAXLOGSIZE** environment variable in the HCL® OneDB® Storage Manager server environment to specify the size threshold of the ISM activity log.

```
setenv ISM_MAXLOGSIZE size
```

**size**

Specifies the size threshold (in megabytes) of the activity log.

If **ISM\_MAXLOGSIZE** is not set, then the default size limit is 1 megabyte. If **ISM\_MAXLOGSIZE** is set to a null value, then the threshold is 0 bytes.

## ISM\_MAXLOGVERS environment variable

Use the **ISM\_MAXLOGVERS** environment variable in the HCL® OneDB® Storage Manager server environment to specify the maximum number of activity-log files to be preserved by the ISM server.

```
setenv ISM_MAXLOGVERS value
```

**value**

specifies the number of files to be preserved.

If **ISM\_MAXLOGVERS** is not set, then the default number of files is four. If the setting is a null value, then the ISM server preserves no activity log files.

## JAR\_TEMP\_PATH environment variable

Set the **JAR\_TEMP\_PATH** variable to specify a non-default local file system location where jar management procedures such as **install\_jar()** and **replace\_jar()** can store temporary **.jar** files of the Java™ virtual machine.

```
setenv JAR_TEMP_PATH pathname
```

### *pathname*

specifies a local directory for temporary **.jar** files.

This directory must have read and write permissions for the user who starts the database server. If the **JAR\_TEMP\_PATH** environment variable is not set, temporary copies of **.jar** files are stored in the **/tmp** directory of the local file system for the database server.

## JAVA\_COMPILER environment variable

You can set the **JAVA\_COMPILER** environment variable in the Java™ virtual machine environment to disable JIT compilation.

```
setenv JAVA_COMPILER { none | NONE }
```

The **NONE** and **none** settings are equivalent. On UNIX™ systems that support the C shell and on which **JAVA\_COMPILER** has been set to **NONE** or **none**, you can enable the JIT compiler for the JVM environment by the following command:

```
unset JAVA_COMPILER
```

## JVM\_MAX\_HEAP\_SIZE environment variable

The **JVM\_MAX\_HEAP\_SIZE** environment variable can set a non-default upper limit on the size of the heap for the Java™ virtual machine.

```
setenv JVM_MAX_HEAP_SIZE size
```

### *size*

is a positive integer that specifies the maximum size (in megabytes).

For example, the following command sets the maximum heap size at 12 MB:

```
set JVM_MAX_HEAP_SIZE 12
```

If you do not set **JVM\_MAX\_HEAP\_SIZE**, 16 MB is the default maximum size.

## LD\_LIBRARY\_PATH environment variable (UNIX™)

The **LD\_LIBRARY\_PATH** environment variable tells the shell on Solaris systems which directories to search for client or shared HCL OneDB™ general libraries. You must specify the directory that contains your client libraries before you can use the product.

```
setenv LD_LIBRARY_PATH $PATH: pathname
```

### *pathname*

Specifies the search path for the library.

For INTERSOLV DataDirect ODBC Driver on AIX®, set **LIBPATH**. For INTERSOLV DataDirect ODBC Driver on HP-UX, set **SHLIB\_PATH**.

The following example sets the **LD\_LIBRARY\_PATH** environment variable to the directory:

```
setenv LD_LIBRARY_PATH
${ONEDB_HOME}/lib:${ONEDB_HOME}/lib/esql:${LD_LIBRARY_PATH}
```

## LIBPATH environment variable (UNIX™)

The **LIBPATH** environment variable tells the shell on AIX® systems which directories to search for dynamic-link libraries for the INTERSOLV DataDirect ODBC Driver. You must specify the full path name for the directory where you installed the product.

```
setenvLIBPATH pathname
```

### *pathname*

Specifies the search path for the libraries.

On Solaris, set **LD\_LIBRARY\_PATH**. On HP-UX, set **SHLIB\_PATH**.

## NODEFDAC environment variable

Enabling NODEFDAC applies the ANSI-compliant restrictions on default access privileges for the PUBLIC group when tables or Owner-mode user-defined routines are created in databases that are not ANSI-compliant.

In a database that is not ANSI-compliant, when the **NODEFDAC** environment variable enabled by setting it to yes,

- the database server withholds default table access privileges from PUBLIC when a new table is created,
- and also withholds the default Execute privilege from PUBLIC when an owner-privileged UDR is created.

```
setenvNODEFDAC { yes }
```

### *yes*

prevents default table privileges (Select, Insert, Update, and Delete) from being granted to PUBLIC on new tables in a database that is not ANSI-compliant. This setting also prevents the Execute privilege from being granted to PUBLIC by default when a new user-defined routine is created in Owner mode.

The *yes* setting is case sensitive, and is also sensitive to leading and trailing blank spaces. Including uppercase letters or blank spaces in the setting is equivalent to leaving **NODEFDAC** unset. When **NODEFDAC** is not set, or if it is set to any value besides *yes*, default privileges on tables and Owner-mode UDRs are granted to PUBLIC by default when the table or UDR is created in a database that is not ANSI-compliant. The setting YES, for example, disables **NODEFDAC**.

Enabling **NODEFDAC** has no effect in an ANSI-compliant databases.

**!** **Important:** Enabling **NODEFDAC** withholds default table or routine privileges from PUBLIC when the object is created, but the **NODEFDAC** setting cannot prevent the PUBLIC group from being granted the same privileges by a user who holds the necessary access privileges on the new table or on the new UDR.

## ONCONFIG environment variable

The **ONCONFIG** environment variable specifies the name of the active file, called the `onconfig` file, which holds the configuration parameters for the database server.

This file is read as input during the initialization procedure. After you prepare your `onconfig` configuration file, set the **ONCONFIG** environment variable to the name of this file.

```
setenvONCONFIGfilename
```

### *filename*

is the name of your `onconfig` file in the `%ONEDB_HOME%\etc\%ONCONFIG%` or `$ONEDB_HOME/etc/$ONCONFIG` directory

This file contains the configuration parameters for your database.

To prepare the `onconfig` file, make a copy of the `onconfig.std` file and modify the copy. Name the `onconfig` file so that it can easily be related to a specific database server. If you have multiple instances of a database server, each instance *must* have its own uniquely named `onconfig` file.

If the **ONCONFIG** environment variable is not set, the database server reads the configuration values from the `onconfig` file during initialization.

## ONINIT\_STDOUT environment variable (Windows™)

The **ONINIT\_STDOUT** environment variable specifies a path and file name in which output from the `oninit` command is stored.

While it is not generally necessary to view output from the `oninit` command, it might be necessary in certain situations, such as when using the `-v` (verbose) option or when you want to see output from an unhandled exception in a process launched within a virtual processor. When the value of **ONINIT\_STDOUT** is set to the name of a file, output from the `oninit` command is written to the file.

```
setONINIT_STDOUT\path\filename
```

You can set the **ONINIT\_STDOUT** environment variable as a system variable in **Control Panel > System > Advanced > Environment Variables**. If the HCL OneDB™ service is configured to log on as user **informix**, start the service using the `starts` command after setting the environment variable. Note, however, that because environment variables are read from the system when the service is started, if the service is set to log on as the local system user, you must restart your computer for the environment variable to take effect. Because the local system user is effectively logged on at all times, environment variables are refreshed only when the operating system is restarted.

For example, if the environment variable set to `C:\temp\oninit_out.txt`, you can start the server with the verbose option with the following command:

```
starts %ONEDB_SERVER% -v
```

The oninit messages are saved to the `C:\temp\oninit_out.txt` file.

**!** **Important:** Only a single instance of the database can run on a Windows™ machine if the **ONINIT\_STDOUT** environment variable is set.

## OPTCOMPIND environment variable

You can set the **OPTCOMPIND** environment variable so that the optimizer can select the appropriate join method.

```
setenv OPTCOMPIND { 2 | 1 | 0 }
```

**0**

A nested-loop join is preferred, where possible, over a sort-merge join or a hash join.

**1**

When the isolation level is *not* Repeatable Read, the optimizer behaves as in setting **2**; otherwise, the optimizer behaves as in setting **0**.

**2**

Nested-loop joins are not necessarily preferred. The optimizer bases its decision purely on costs, regardless of transaction isolation mode.

When **OPTCOMPIND** is not set, the database server uses the OPTCOMPIND value from the ONCONFIG configuration file. When neither the environment variable nor the configuration parameter is set, the default value is **2**.

On HCL OneDB™, the SET ENVIRONMENT OPTCOMPIND statement can set or reset **OPTCOMPIND** dynamically at runtime. This overrides the current **OPTCOMPIND** value (or the ONCONFIG configuration parameter OPTCOMPIND) for the current user session only. For more information about the SET ENVIRONMENT OPTCOMPIND statement of SQL see the *HCL OneDB™ Guide to SQL: Syntax*.

For more information about the ONCONFIG configuration parameter OPTCOMPIND, see the *HCL OneDB™ Administrator's Reference*. For more information about the different join methods that the optimizer uses, see your *HCL OneDB™ Performance Guide*.

## OPTMSG environment variable

Set the **OPTMSG** environment variable at runtime before you start the application to enable (or disable) optimized message transfers (message chaining) for all SQL statements in the application.

```
setenv OPTMSG { 0 | 1 }
```

**0**

disables optimized message transfers.



1

enables optimized message transfers and implements the feature for any subsequent connection.

The default value is 0 (zero), which explicitly disables message chaining. You might want, for example, to disable optimized message transfers for statements that require immediate replies, for debugging, or to ensure that the database server processes all messages before the application terminates.

When you set **OPTMSG** within an application, you can activate or deactivate optimized message transfers for each connection or within each thread. To enable optimized message transfers, you must set **OPTMSG** before you establish a connection.

For more information about setting **OPTMSG** and defining related global variables, see the *HCL OneDB™ ESQL/C Programmer's Manual*.

## OPTOFC environment variable

Use the **OPTOFC** environment variable to enable optimize-OPEN-FETCH-CLOSE functionality in applications or other APIs (such as JDBC, ODBC, OLE DB, LIBDMI, and Lib C++) that use DECLARE and OPEN statements to establish a cursor.

```
setenvOPTOFC { 0 | 1 }
```

0

disables **OPTOFC** for all threads of the application.

1

enables **OPTOFC** for every cursor in every thread of the application.

The default value is 0 (zero).

You can set the **OPTOFC** environment variable on the client or server. If this environment variable is set on the server, then any application that does not explicitly set this environment variable uses the value that is set on the server.

The **OPTOFC** environment variable reduces the number of message requests between the application and the database server.

If you set **OPTOFC** from the shell, you must set it before you start the application. For more information about enabling **OPTOFC** and related features, see the *HCL OneDB™ ESQL/C Programmer's Manual*.

## OPT\_GOAL environment variable (UNIX™)

Set the **OPT\_GOAL** environment variable in the user environment, before you start an application, to specify the query performance goal for the optimizer.

```
setenvOPT_GOAL { 0 | -1 }
```

0

Specifies user-response-time optimization.

-1

Specifies total-query-time optimization.

The default behavior is for the optimizer to use query plans that optimize the total query time.

You can also specify the optimization goal for individual queries with optimizer directives or for a session with the SET OPTIMIZATION statement.

Both methods take precedence over the **OPT\_GOAL** environment variable setting. You can also set the OPT\_GOAL configuration parameter for the HCL OneDB™ system; this method has the lowest level of precedence.

For more information about optimizing queries for your database server, see your *HCL OneDB™ Performance Guide*. For information about the SET OPTIMIZATION statement, see the *HCL OneDB™ Guide to SQL: Syntax*.

## PATH environment variable

The UNIX™ **PATH** environment variable tells the shell which directories to search for executable programs. You must add the directory containing your HCL OneDB™ product to your **PATH** setting before you can use the product.

```
setenvPATH$PATH: pathname
```

### *pathname*

Specifies the search path for the executable files.

Include a colon (:) separator between the path names on UNIX™ systems. (Use the semicolon (;) separator between path names on Windows™ systems.)

You can specify the search path in various ways. The **PATH** environment variable tells the operating system where to search for executable programs. You must include the directory that contains your HCL OneDB™ product in your **path** setting before you can use the product. This directory should be located before \$ONEDB\_HOME/bin, which you must also include.

For additional information about how to modify your path, see [Modifying an environment-variable setting on page 142](#).

## PDQPRIORITY environment variable

The **PDQPRIORITY** environment variable determines the degree of parallelism that the database server uses and affects how the database server allocates resources, including memory, processors, and disk reads.

```
setenvPDQPRIORITY { HIGH | LOW | OFF | resources }
```

### *resources*

Is an integer in the range 0 to 100. The value 1 is the same as LOW, and 100 is the same as HIGH. Values lower than 0 are set to 0 (OFF), and values greater than 100 are set to 100 (HIGH).

Value 0 is the same as OFF (for HCL OneDB™ only).

Here the HIGH, LOW, and OFF keywords have the following effects:

**HIGH**

When the database server allocates resources among all users, it gives as many resources as possible to the query.

**LOW**

Data values are fetched from fragmented tables in parallel.

**OFF**

PDQ processing is turned off (for HCL OneDB™ only).

Usually, the more resources a database server uses, the better its performance for a given query. If the server uses too many resources, however, contention for the resources can take resources away from other queries, resulting in degraded performance. For more information about performance considerations for **PDQPRIORITY**, see the *HCL OneDB™ Performance Guide*.

An application can override the setting of this environment variable when it issues the SQL statement SET PDQPRIORITY, as the *HCL OneDB™ Guide to SQL: Syntax* describes.

## Using PDQPRIORITY with HCL OneDB™

The *resources* value specifies the query priority level and the amount of resources that the database server uses to process the query.

When **PDQPRIORITY** is not set, the default value is `OFF`.

When **PDQPRIORITY** is set to HIGH, HCL OneDB™ determines an appropriate value to use for **PDQPRIORITY** based on several criteria. These include the number of available processors, the fragmentation of tables queried, the complexity of the query, and additional factors.

## PLCONFIG environment variable

The **PLCONFIG** environment variable specifies the name of the configuration file that the High Performance Loader (HPL) uses. This file must be located in the `$ONEDB_HOME/etc` directory. If the **PLCONFIG** environment variable is not set, then `$ONEDB_HOME/etc/plconfig` is the default configuration file.

```
setenv PLCONFIG filename
```

**filename**

Specifies the simple file name of the configuration file that the High-Performance Loader uses.

For example, to specify the `$ONEDB_HOME/etc/custom.cfg` file as the configuration file for the High-Performance Loader, enter the following command:

```
setenv PLCONFIG custom.cfg
```

For more information, see the *HCL OneDB™ High-Performance Loader User's Guide*.

## PLOAD\_LO\_PATH environment variable

The **PLOAD\_LO\_PATH** environment variable lets you specify the pathname for smart-large-object handles (which identify the location of smart large objects such as BLOB and CLOB data types).

```
setenvPLOAD_LO_PATHpathname
```

### **pathname**

specifies the directory for the smart-large-object handles.

If **PLOAD\_LO\_PATH** is not set, the default directory is **/tmp**.

For more information, see the *HCL OneDB™ High-Performance Loader User's Guide*.

## PLOAD\_SHMBASE environment variable

The **PLOAD\_SHMBASE** environment variable lets you specify the shared-memory address at which the High Performance Loader (HPL) onload processes will attach. If **PLOAD\_SHMBASE** is not set, the HPL determines which shared-memory address to use.

```
setenvPLOAD_SHMBASEvalue
```

### **value**

Used to calculate the shared-memory address.

If the onload utility cannot attach, an error is issued, and you must specify a new value.

The onload utility tries to determine at which address to attach, as follows in the following (descending) order:

1. Attach at the same address (SHMBASE) as the database server.
2. Attach beyond the database server segments.
3. Attach at the address specified in **PLOAD\_SHMBASE**.



**Tip:** It is recommended that you let the HPL decide where to attach and that you set **PLOAD\_SHMBASE** only if necessary to avoid shared-memory collisions between onload and the database server.

For more information, see the *HCL OneDB™ High-Performance Loader User's Guide*.

## PSM\_ACT\_LOG environment variable

Use the **PSM\_ACT\_LOG** environment variable to specify the location of the HCL OneDB™ Primary Storage Manager activity log for your environment, for example, for a single session.

```
setenvPSM_ACT_LOG pathname
```

### **pathname**

The full path name for the location of the `$ONEDB_HOME/psm_act.log`. If you specify a file name only, the storage manager creates the activity log in the working directory in which you started the storage manager.

The **PSM\_ACT\_LOG** environment variable overrides the value of the PSM\_ACT\_LOG configuration parameter.

## PSM\_CATALOG\_PATH environment variable

Use the **PSM\_CATALOG\_PATH** environment variable to specify the location of the HCL OneDB™ Primary Storage Manager catalog tables for your environment, for example, for a single session.

```
setenvPSM_CATALOG_PATH pathname
```

### *pathname*

The full path name for the location of the catalog table, which contain information about the pools, devices, and objects managed by the storage manager.

The **PSM\_CATALOG\_PATH** environment variable overrides the value of the PSM\_CATALOG\_PATH configuration parameter.

## PSM\_DBS\_POOL environment variable

Use the **PSM\_DBS\_POOL** environment variable to change the name of the pool in which the HCL OneDB™ Primary Storage Manager places backup and restore dbspace data for your environment, for example, for a single session.

```
setenvPSM_DBS_POOL pool_name
```

### *pool\_name*

The name of the storage manager pool.

The **PSM\_DBS\_POOL** environment variable overrides the value of the PSM\_DBS\_POOL configuration parameter.

## PSM\_DEBUG environment variable

Use the **PSM\_DEBUG** environment variable to specify the amount of debugging information that prints in the OneDB® Primary Storage Manager debug log for your environment, for example, for a single session.

```
setenvPSM_DEBUG value
```

### *value*

- 0 = No debugging messages.
- 1 = Prints only internal errors.
- 2 = Prints information about the entry and exit of functions and prints internal errors.
- 3 = Prints the information specified by 1-2 with additional details.
- 4 = Prints information about parallel operations and the information specified by 1-3.
- 5 = Prints information about internal states in the OneDB® Primary Storage Manager.
- 6 = Prints the information specified by 1-5 with additional details.

7 = Prints information specified by 1-6 with additional details.

8 = Prints information specified by 1-7 with additional details.

9 = Prints all debugging information.

The **PSM\_DEBUG** environment variable overrides the value of the PSM\_DEBUG configuration parameter.

## PSM\_DEBUG\_LOG environment variable

Use the **PSM\_DEBUG\_LOG** environment variable to specify the location of the HCL OneDB™ Primary Storage Manager debug log for your environment, for example, for a single session.

```
setenvPSM_DEBUG_LOG pathname
```

### *pathname*

The full path name for the location of the `$ONEDB_HOME/psm_debug.log`. If you specify a file name only, the storage manager creates the debug log in the working directory in which you started the storage manager.

The **PSM\_DEBUG\_LOG** environment variable overrides the value of the PSM\_DEBUG\_LOG configuration parameter.

## PSM\_LOG\_POOL environment variable

Use the **PSM\_LOG\_POOL** environment variable to change the name of the pool in which the HCL OneDB™ Primary Storage Manager places backup and restore log data for your environment, for example, for a single session.

```
setenvPSM_LOG_POOL pool_name
```

### *pool\_name*

The name of the storage manager log pool.

The **PSM\_LOG\_POOL** environment variable overrides the value of the PSM\_LOG\_POOL configuration parameter.

## PSORT\_DBTEMP environment variable

The **PSORT\_DBTEMP** environment variable specifies the location where the database server writes the temporary files that the **PSORT\_NPROCS** environment variable uses to perform a sort.

```
setenvPSORT_DBTEMP pathname
```

### *pathname*

The name of the UNIX™ directory used for intermediate writes during a sort.

To set the **PSORT\_DBTEMP** environment variable to specify the directory (for example, `/usr/leif/tempSORT`), enter the following command:

```
setenv PSORT_DBTEMP /usr/leif/tempSORT
```

For maximum performance, specify directories that are located in file systems on different disks.

You might also want to consider setting the environment variable **DBSPACETEMP** to place temporary files used in sorting in dbspaces rather than operating-system files. See the discussion of the **DBSPACETEMP** environment variable in [DBSPACETEMP environment variable on page 175](#).

The database server uses the directory that **PSORT\_DBTEMP** specifies, even if the environment variable **PSORT\_NPROCS** is not set. For additional information about the **PSORT\_DBTEMP** environment variable, see your *HCL OneDB™ Administrator's Guide* and your *HCL OneDB™ Performance Guide*.

## PSORT\_NPROCS environment variable

The **PSORT\_NPROCS** environment variable enables the database server to improve the performance of the parallel-process sorting package by allocating more threads for sorting.

Before the sorting package performs a parallel sort, make sure that the database server has enough memory for the sort.

```
setenv PSORT_NPROCS threads
```

### *threads*

is an integer, specifying the maximum number of threads to be used to sort a query. This value cannot be greater than 10.

The following command sets **PSORT\_NPROCS** to 4:

```
setenv PSORT_NPROCS 4
```

To disable parallel sorting, enter the following command:

```
unsetenv PSORT_NPROCS
```

It is recommended that you initially set **PSORT\_NPROCS** to 2 when your computer has multiple CPUs. If subsequent CPU activity is lower than I/O activity, you can increase the value of **PSORT\_NPROCS**.



**Tip:** If the **PDQPRIORITY** environment variable is not set, the database server allocates the minimum amount of memory to sorting. This minimum memory is insufficient to start even two sort threads. If you have not set **PDQPRIORITY**, check the available memory before you perform a large-scale sort (such as an index build) to make sure that you have enough memory.

## Default PSORT\_NPROCS values for detached indexes

If the **PSORT\_NPROCS** environment variable is set, the database server uses the specified number of sort threads as an upper limit for ordinary sorts. If **PSORT\_NPROCS** is not set, parallel sorting does not take place. If you have a single-CPU virtual processor, the database server uses one thread for the sort. If **PSORT\_NPROCS** is set to 0, the database server uses three threads for the sort.

## Default PSORT\_NPROCS values for attached indexes

The default number of threads is different for attached indexes.

If the **PSORT\_NPROCS** environment variable is set, you get the specified number of sort threads for each fragment of the index that is being built.

If **PSORT\_NPROCS** is not set, or if it is set to 0, you get two sort threads for each fragment of the index unless you have a single-CPU virtual processor. If you have a single-CPU virtual processor, you get one sort thread for each fragment of the index.

For additional information about the **PSORT\_NPROCS** environment variable, see your *HCL OneDB™ Administrator's Guide* and your *HCL OneDB™ Performance Guide*.

## RTREE\_COST\_ADJUST\_VALUE environment variable

The **RTREE\_COST\_ADJUST\_VALUE** environment variable specifies a coefficient that support functions of user-defined data types can use to estimate the cost of an R-tree index for queries on UDT columns.

```
setenvRTREE_COST_ADJUST_VALUEvalue
```

### *value*

is a floating-point number, where  $1 \leq \text{value} \leq 1000$ , specifying a multiplier for estimating the cost of using an index on a UDT column.

For spatial queries, the I/O overhead tends to exceed by far the CPU cost, so by multiplying the uncorrected estimated cost by an appropriate *value* from this setting, the database server can make better cost-based decisions on how to implement queries on UDT columns for which an R-tree index exists.

## SHLIB\_PATH environment variable (UNIX™)

The **SHLIB\_PATH** environment variable tells the shell on HP-UX systems which directories to search for dynamic-link libraries. This is used, for example, with the INTERSOLV DataDirect ODBC Driver. You must specify the full pathname for the directory where you installed the product.

```
setenvSHLIB_PATH$PATH:pathname
```

### *pathname*

Specifies the search path for the libraries.

On Solaris systems, set **LD\_LIBRARY\_PATH**. On AIX® systems, set **LIBPATH**.

## SRV\_FET\_BUF\_SIZE environment variable

Use the **SRV\_FET\_BUF\_SIZE** environment variable to specify the size of the fetch buffer that the local database server uses in distributed DML transactions across database servers.

```
setenvSRV_FET_BUF_SIZEsize
```

### *size*

is a positive integer that is no greater than 1048576 (1 MiB), specifying the size (in bytes) of the fetch buffer that holds data retrieved by a cross-server distributed query.



For example, to set a buffer size to 5,000 bytes on a UNIX™ system that uses the C shell, set **SRV\_FET\_BUF\_SIZE** by entering the following command:

```
setenv SRV_FET_BUF_SIZE 5000
```

When **SRV\_FET\_BUF\_SIZE** is set to a valid value, the new value overrides the default value (or any previously set value) of **SRV\_FET\_BUF\_SIZE**. The setting takes effect only when it is set in the starting environment of the database server.

When **SRV\_FET\_BUF\_SIZE** is not set, the default setting for the fetch buffer is dependent on row size.

No error is raised if **SRV\_FET\_BUF\_SIZE** is set to a value that is less than the default size, or that is greater than 1048576 (1MiB). If you specify a size for **SRV\_FET\_BUF\_SIZE** that is greater than 1048576, the value is set to 1048576. In older 11.70 releases, up to and including 11.70.xC4, the upper limit is 32767.

A valid **SRV\_FET\_BUF\_SIZE** setting is in effect only in cross-server DML transactions in which the local database server participates as the coordinator or as a subordinate database server.

- It has no effect, however, on queries that access only databases of the local server instance, and it does not affect the size of the fetch buffer in client-to-local-server communication.
- The processing of BYTE and TEXT objects is not affected by the **SRV\_FET\_BUF\_SIZE** setting.
- Setting **SRV\_FET\_BUF\_SIZE** for the environment of the local database server does not reset the fetch buffer size of remote server instances that coordinate or participate in cross-server DML transactions with the local server instance.

The greater the size of the buffer, the more rows can be returned, and the less frequently the local server must wait while the database server returns rows. A large buffer can improve performance when transferring a large amount of data between servers.

## STMT\_CACHE environment variable

Use the **STMT\_CACHE** environment variable to control the use of the shared-statement cache on a session.

This feature can reduce memory consumption and can speed query processing among different user sessions. Valid **STMT\_CACHE** values are 1 and 0.

```
setenvSTMT_CACHE { 1 | 0 }
```

1

enables the SQL statement cache.

0

disables the SQL statement cache.

Set the **STMT\_CACHE** environment variable for applications that do not use the SET STMT\_CACHE statement to control the use of the SQL statement cache. By default, a statement cache is disabled, but can be enabled through the STMT\_CACHE parameter of the `onconfig.std` file or by the SET STMT\_CACHE statement.

This environment variable has no effect if the SQL statement cache is disabled through the configuration file setting. Values set by the SET STMT\_CACHE statement in the application override the **STMT\_CACHE** setting.

## TERM environment variable (UNIX™)

The **TERM** environment variable is used for terminal handling. It lets DB-Access (and other character-based applications) recognize and communicate with the terminal that you are using.

```
setenvTERMtype
```

### **type**

Specifies the terminal type.

The terminal type specified in the **TERM** setting must correspond to an entry in the `termcap` file or `terminfo` directory.

Before you can set the **TERM** environment variable, you must obtain the code for your terminal from the database administrator.

For example, to specify the vt100 terminal, set the **TERM** environment variable by entering the following command:

```
setenv TERM vt100
```

## TERMCAP environment variable (UNIX™)

The **TERMCAP** environment variable is used for terminal handling. It tells DB-Access (and other character-based applications) to communicate with the `termcap` file instead of the `terminfo` directory.

```
setenvTERMCAPpathname
```

### **pathname**

Specifies the location of the `termcap` file.

The `termcap` file contains a list of various types of terminals and their characteristics. For example, to provide DB-Access terminal-handling information, which is specified in the `/usr/informix/etc/termcap` file, enter the following command:

```
setenv TERMCAP /usr/informix/etc/termcap
```

You can use set **TERMCAP** in any of the following ways. If several `termcap` files exist, they have the following (descending) order of precedence:

1. The `termcap` file that you create
2. The `termcap` file that the database server supplies (that is, `$ONEDB_HOME/etc/termcap`)
3. The operating-system `termcap` file (that is, `/etc/termcap`)

If you set the **TERMCAP** environment variable, be sure that the **ONEDB\_TERM** environment variable is set to `termcap`.

If you do not set the **TERMCAP** environment variable, the `terminfo` directory is used by default.

## TERMINFO environment variable (UNIX™)

The **TERMINFO** environment variable is used for terminal handling.

The environment variable is supported only on platforms that provide full support for the `terminfo` libraries that System V and Solaris UNIX™ systems provide.

```
setenvTERMINFO/usr/lib/terminfo
```

**TERMINFO** tells DB-Access to communicate with the `terminfo` directory instead of the `termcap` file. The `terminfo` directory has subdirectories that contain files that pertain to terminals and their characteristics.

To set **TERMINFO**, enter the following command:

```
setenv TERMINFO /usr/lib/terminfo
```

## THREADLIB environment variable (UNIX™)

Use the **THREADLIB** environment variable to compile multithreaded applications. A multithreaded application lets you establish as many connections to one or more databases as there are threads. These connections can remain active while the application program executes.

The **THREADLIB** environment variable indicates which thread package to use when you compile an application. Currently only the Distributed Computing Environment (DCE) is supported.

```
setenvTHREADLIBDCE
```

The **THREADLIB** environment variable is checked when the `-thread` option is passed to the script when you compile a multithreaded application. When you use the `-thread` option while compiling, the script generates an error if **THREADLIB** is not set, or if **THREADLIB** is set to an unsupported thread package.

## TZ environment variable

The **TZ** environment variable is used for setting the time zone. It is used by various time functions to compute times relative to Coordinated Universal Time (UTC), formerly known as Greenwich Mean Time (GMT). The format is specified by the operating system.

```
setenvTztzn [{ + | - } hh [ :mm [ :ss ] ] [ dzn ]
```

### **tzn**

Three-letter time zone name, such as PST. You must specify the correct offset from local time to UTC (Universal Time Coordinated).

### **hh**

A one- or two-digit difference in hours between UTC and local time. Optionally signed.

### **mm**

Two-digit difference in minutes between UTC and local time.

**ss**

Two-digit difference in seconds between UTC and local time.

**dzn**

Three-letter daylight-saving-time zone, such as PDT. If daylight saving time is never in effect in the locality, set **TZ** without a value for *dzn*.

For example, if you use Pacific Standard Time with Pacific daylight savings time, set the **TZ** environment variable to `PST8PDT`. For more information on setting the **TZ** environment variable, see your operating system documentation.

## USETABLENAME environment variable

The **USETABLENAME** environment variable can prevent users from using a synonym to specify the *table* in ALTER TABLE or DROP TABLE statements. Unlike most environment variables, **USETABLENAME** is not required to be set to a value. It takes effect if you set it to any value, or to no value.

```
setenvUSETABLENAME
```

By default, ALTER TABLE or DROP TABLE statements accept a valid synonym for the name of the *table* to be altered or dropped. (In contrast, RENAME TABLE issues an error if you specify a synonym, as do the ALTER SEQUENCE, DROP SEQUENCE, and RENAME SEQUENCE statements, if you attempt to substitute a synonym for the *sequence* name in those statements.)

If you set **USETABLENAME**, an error results if a synonym is in ALTER TABLE or DROP TABLE statements. Setting **USETABLENAME** has no effect on the DROP VIEW statement, which accepts a valid synonym for the view.

## Appendixes

### The stores\_demo Database

The **stores\_demo** database contains a set of tables that describe an imaginary business and many of the examples in the HCL OneDB™ documentation are based on this database.

The **stores\_demo** database uses the default (U.S. English) locale and is not ANSI-compliant.

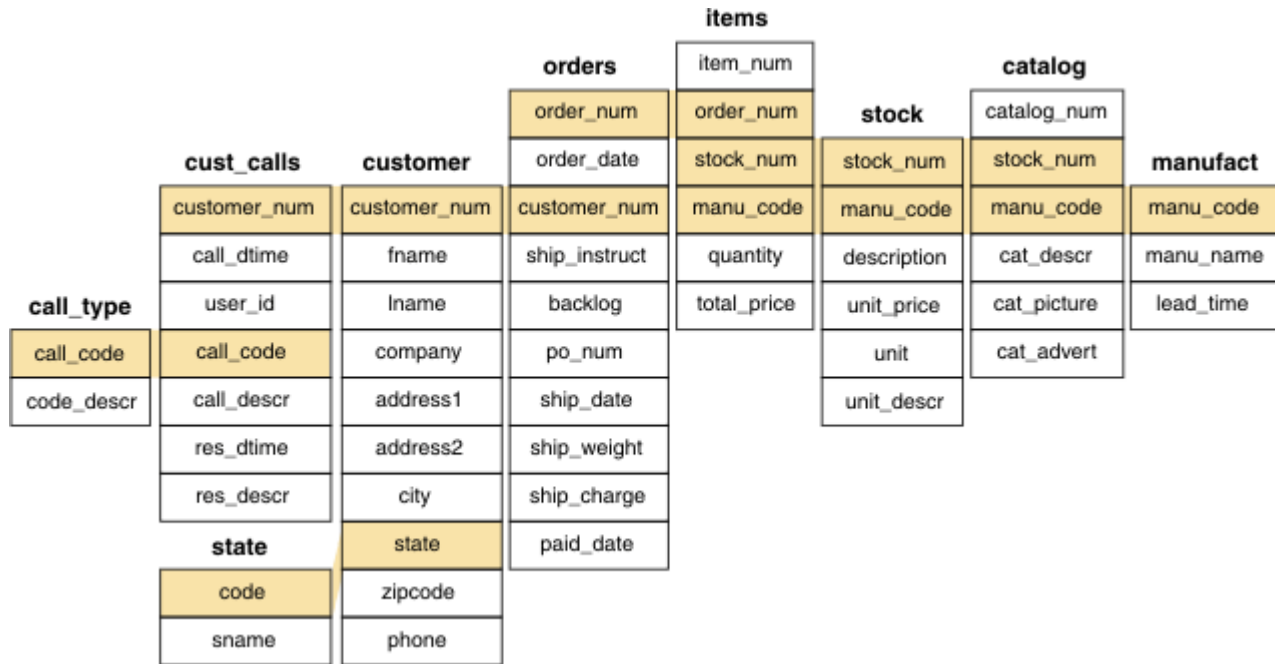
For information about how to create and populate the **stores\_demo** database, see the *HCL OneDB™ DB-Access User's Guide*. For information about how to design and implement a relational database, see the *HCL OneDB™ Database Design and Implementation Guide*.

### The stores\_demo Database Map

Some of the tables in the **stores\_demo** database have relationships between them.

The following illustration displays the joins in the **stores\_demo** database between customers, catalog orders, and customer calls. The shading that connects a column in one table to a column with the same name in another table indicates the relationships, or *joins*, between tables.

Figure 5. Joins between customers and catalog orders



The following illustration displays the joins in the **stores\_demo** database between customers, electricity meter data, and location. The **Customer\_ts\_data**, **ts\_data**, and **ts\_data\_location** tables contain time series data. You can prevent the creation of these time series tables when you create the demonstration database.

The **stores\_demo** database also contains tables that hold electricity meter data for customers. The **Customer\_ts\_data** and **ts\_data** tables contain time series data. The **customer** table and the **Customer\_ts\_data** table are joined by the **customer\_num** column. The **Customer\_ts\_data** table is joined to the **ts\_data** table by the **loc\_esi\_id**, **measure\_unit**, and **direction** columns. You can prevent the creation of these time series tables when you create the demonstration database.

Figure 6. Joins between customers, electricity usage data, and location

	Customer_ts_data	ts_data	ts_data_location
	loc_esi_id	loc_esi_id	loc_esi_id
	measure_unit	measure_unit	longlat
	direction	direction	
	customer_num	multiplier	
	meter_type	raw_reads	
<b>customer</b>			
customer_num			
fname			
lname			
company			
address1			
address2			
city			
state			
zipcode			
phone			

## The superstores\_demo database

The **superstores\_demo** database illustrates an object-relational schema.

SQL files and user-defined routines (UDRs) that are provided with DB-Access let you derive the **superstores\_demo** object-relational database.

The **superstores\_demo** database uses the default locale and is not ANSI-compliant.

For information about how to create and populate the demonstration databases, including relevant SQL files, see the *HCL OneDB™ DB-Access User's Guide*. For conceptual information about demonstration databases, see the *HCL OneDB™ Database Design and Implementation Guide*.

## Structure of the superstores\_demo Tables

Although many of the tables in the **superstores\_demo** database have the same name as **stores\_demo** tables, they are different.

The **superstores\_demo** database includes the following tables. The tables are listed alphabetically, not in the order in which they are created.

- **call\_type**
- **catalog**
- **cust\_calls**
- **customer**

- retail\_customer
- whlsale\_customer
- items
- location
  - location\_non\_us
  - location\_us
- manufact
- orders
- region
- sales\_rep
- state
- stock
- stock\_discount
- units

## User-defined routines and extended data types

The **superstores\_demo** database uses user-defined routines (UDRs) and extended data types.

A UDR is a routine that you define that can be invoked within an SQL statement or another UDR. A UDR can either return values or not.

The data type system of HCL OneDB™ is an extensible and flexible system that supports the creation of following kinds of data types:

- Extensions of existing data types by, redefining some of the behavior for data types that the database server provides
- Definitions of customized data types by a user

For information about creating and using UDRs and extended data types, see *HCL OneDB™ User-Defined Routines and Data Types Developer's Guide*.

The **superstores\_demo** database creates the *distinct* data type, percent, in a UDR, as follows:

```
CREATE DISTINCT TYPE percent AS DECIMAL(5,5);
DROP CAST (DECIMAL(5,5) AS percent);
CREATE IMPLICIT CAST (DECIMAL(5,5) AS percent);
The superstores_demo database creates the following named row types:
```

- **location** hierarchy:
  - location\_t
  - loc\_us\_t
  - loc\_non\_us\_t
- **customer** hierarchy:
  - name\_t
  - customer\_t

- retail\_t
- whlsale\_t
- orders table
  - ship\_t

### location\_t definition

<b>location_id</b>	SERIAL
<b>loc_type</b>	CHAR(2)
<b>company</b>	VARCHAR(20)
<b>street_addr</b>	LIST(VARCHAR(25) NOT NULL)
<b>city</b>	VARCHAR(25)
<b>country</b>	VARCHAR(25)

### loc\_us\_t definition

<b>state_code</b>	CHAR(2)
<b>zip</b>	ROW(code INTEGER, suffix SMALLINT)
<b>phone</b>	CHAR(18)

### loc\_non\_us\_t definition

<b>province_code</b>	CHAR(2)
<b>zipcode</b>	CHAR(9)
<b>phone</b>	CHAR(15)

### name\_t definition

<b>first</b>	VARCHAR(15)
<b>last</b>	VARCHAR(15)

### customer\_t definition

<b>customer_num</b>	SERIAL
<b>customer_type</b>	CHAR(1)
<b>customer_name</b>	name_t
<b>customer_loc</b>	INTEGER
<b>contact_dates</b>	LIST(DATETIME YEAR TO DAY NOT NULL)
<b>cust_discount</b>	percent
<b>credit_status</b>	CHAR(1)

### retail\_t definition

<b>credit_num</b>	CHAR(19)
<b>expiration</b>	DATE

### whlsale\_t definition

<b>resale_license</b>	CHAR(15)
<b>terms_net</b>	SMALLINT

### ship\_t definition

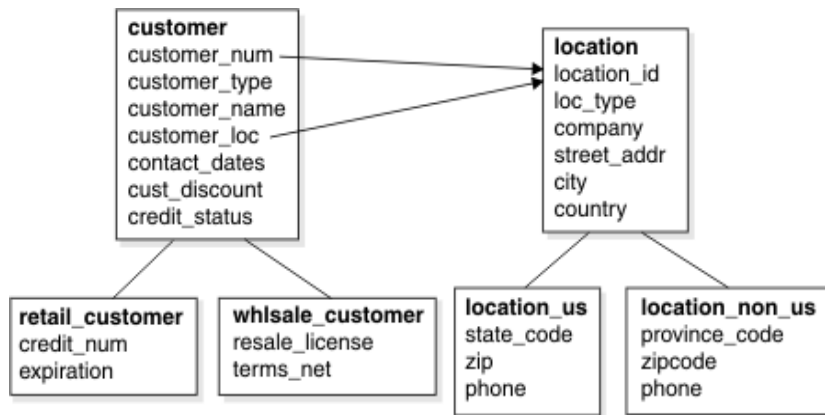
<b>date</b>	DATE
<b>weight</b>	DECIMAL(8,2)
<b>charge</b>	MONEY(6,2)
<b>instruct</b>	VARCHAR(40)



## Table Hierarchies

The following illustration shows how the hierarchical tables of the **superstores\_demo** database are related. The foreign key and primary relationships between the two tables are indicated by shaded arrows that point from the **customer.custnum** and **customer.loc** columns to the **location.location\_id** columns.

Figure 7. Hierarchies of superstores\_demo Tables



# Index

## Special Characters

- ( [ ] ), brackets
  - substring operator 89, 137
- ' VERSION' table 65
- ( \_ ), underscore
  - in SQL identifiers 181
- ( ; ), semicolon
  - list separator 198, 210
- ( : ), colon
  - cast ( :: ) operator 135, 137
  - DATETIME delimiter 93
  - INTERVAL delimiter 100
  - list separator 175, 182, 198, 205, 210
- ( != ), not equal to
  - relational operator 137
- ( ' ), single quotation
  - string delimiter 171
- ( ' ), single quotation symbols
  - string delimiter 181
- ( " ), double quotation marks
  - string delimiter 103
- ( " ), double quotation symbols
  - delimited SQL identifiers 181
  - string delimiter 81, 106, 113
- ( ( ) ), parentheses
  - delimiters in expressions 127
- ( { } ), braces
  - collection delimiters 103, 106
  - pathname delimiters 142
- ( / ), slash
  - DATE separator 93, 126, 166
  - division operator 123, 137
  - pathname delimiter 143, 172, 205
- ( \ ), backslash
  - invalid as delimiter 168
  - pathname delimiter 145, 200
- ( # ), sharp
  - comment indicator 140
- ( % ), percentage
  - DBTIME escape symbol 176
  - pathname indicator 175
- ( < ), less than
  - angle ( < > ) brackets 89
  - relational operator 137, 167
- ( > ), greater than
  - angle ( < > ) brackets 89
  - relational operator 10, 137
- ( | ), vertical bar
  - absolute value delimiter 100
  - concatenation ( || ) operator 137
  - field delimiter 168
- ( \$ ), dollar sign
  - currency symbol 105, 171
  - pathname indicator 210
- ( - ), hyphen
  - DATE separator 166
  - DATETIME delimiter 93
  - INTERVAL delimiter 100
  - subtraction operator 123, 137
  - symbol in syscolauth 5, 23
  - symbol in sysfragauth 36
  - symbol in systabauth 65
  - unary operator 124, 137
- ( , ), comma
  - decimal point 171
  - list separator 106, 109, 175
  - thousands separator 105

- ( . ), period
    - DATE separator 166, 166
    - DATETIME delimiter 93
    - decimal point 97, 105, 171
    - execution symbol 140
    - INTERVAL delimiter 100
    - membership operator 137
    - nested dot notation 130
  - ( ), blank space
    - DATETIME delimiter 93
    - INTERVAL delimiter 100
    - padding CHAR values 91
    - padding VARCHAR values 116
  - ( \* ), asterisk
    - multiplication operator 88, 123, 127, 137
    - systabauth value 5, 65
    - wildcard symbol 21, 77
  - ( + ), plus sign
    - addition operator 123, 137
    - truncation indicator 188
    - unary operator 137
  - ( = ), equality
    - assignment operator 145
    - relational operator 21, 88, 92, 137
  - ( ~ ), tilde
    - pathname indicator 143
- ## A
- Abbreviated year values 93, 163, 165, 166, 176
  - ACCESS keyword 122
  - Access method
    - B-tree 15, 15, 41, 180
    - built-in 15, 15
    - primary 15, 64, 64
    - R-Tree 180
    - secondary 15, 29, 43, 107
    - sysams data 15
    - sysindices data 43
    - sysopclasses data 48
    - systabamdata data 64
  - ACCESS\_METHOD keyword 15
  - Activity-log files 204
  - Addition (+) operator 123, 137
  - Aggregate functions 115
    - built-in 103, 106, 113
    - no BYTE argument 89
    - no collection arguments 103, 106, 113
    - sysaggregates data 14
    - user-defined 14
  - AIX operating system 186, 206
  - Alias of a table 5
  - Alignment of data type 75
  - Alignment of data types 20
  - ALL operator 137
  - ALTER OPTICAL CLUSTER statement 49
  - Alter privilege 5, 65, 78
  - ALTER SEQUENCE statement 220
  - ALTER TABLE statement
    - casting effects 133
    - changing data types 81
    - lock mode 184
    - next extent size 10
    - SERIAL columns 111
    - SERIAL8 columns 112
    - synonyms 220
  - am\_beginscan() function 15
  - am\_close() function 15
  - am\_getnext() function 15

- am\_insert() function 15
- am\_open() function 15
- AND operator 21, 137
- ANSI compliance
  - ansi flag 162
  - DATETIME literals 176
  - DBANSIWARN environment variable 162
  - DECIMAL range 97
  - DECIMAL(p) data type 96
  - Information Schema views 77
  - isolation level 80
  - public synonyms 63, 65
- ANSIOWNER environment variable 157
- ANY operator 137
- Arabic locales 90
- Archiving
  - setting DBREMOTECMD 174
- Arithmetic
  - DATE operands 93, 125
  - DATETIME operands 124
  - integer operands 88, 100, 100, 114
  - INTERVAL operands 100, 124
  - operators 137
  - string operands 91
  - time operands 123
- AS keyword 135, 135
- assign() support function 128
- AT keyword 103
- Attached index 180
- Attached indexes 39, 165, 215
- Audit Analysis officer 202, 202
- Authorization identifier 72, 80
- AUTO\_STAT\_MODE configuration
  - parameter 32, 37
  - AUTO\_STAT\_MODE session environment setting 32, 37

## B

- B-tree access method 15, 41, 180
- B-tree index 39
- Backslash ( \ ) symbol 168
- Backup
  - file prefix 190
- Bandwidth 191
- BETWEEN operator 137
- BIGINT data type 87
  - coltype code 24
  - length (syscolumns) 27
- BIGSERIAL data type 87
  - coltype code 24
  - length (syscolumns) 27
- bin subdirectory 142
- Binding style 80
- BLOB data type
  - casting unavailable 88
  - defined 88
  - inserting data 88
  - syscolattns data 22
- Blobspaces
  - defined 121
  - memory cache for staging 199
  - names 181
  - sysblobs data 20
- BOOLEAN data type
  - defined 89
- Boolean expression
  - with BOOLEAN data type 89
  - with BYTE data type 89

- Boolean expression with TEXT data type 115
- Borland C compiler 195
- Bourne shell 140, 141
- Bracket ( [ ] ) symbols 115
- brackets substring 115
- Buffers
  - BYTE or TEXT storage (DBBLOBBUF) 163
  - fetch buffer (FET\_BUFFER\_SIZE) 183
  - fetch buffer (SRV\_FET\_BUFFER\_SIZE) 216
  - floating-point display (DBFLTMASK) 169
  - network buffer (IFX\_NETBUF\_SIZE) 189
  - private network buffer pool 188
- Built-in access method 15, 15
- Built-in aggregates 14, 103, 106, 113
- Built-in casts 20, 133
- Built-in data types
  - casts 133, 137
  - listed 118
  - syscolumns.coltype code 24
  - sysdistrib.type code 32
  - sysxdtypes data 75
- Built-in opaque data types 135
- BY clause 115
- BY keyword 89, 115
- BY ORDER 115
- BYTE data type
  - casting to BLOB 89
  - coltype code 24
  - defined 89
  - increasing buffer size 163
  - inserting values 89
  - restrictions
    - in Boolean expression 89
    - systables.npused 65
    - with GROUP BY 89
    - with LIKE or MATCHES 89
    - with ORDER BY 89
  - selecting from BYTE columns 89
  - setting buffer size 163
  - sysblobs data 20, 20
  - syscolumns data 27
  - sysfragments data 39
  - sysopclstr data 49, 49

## C

- C compiler
  - default name 195
  - ONEDB\_C setting 195
  - thread package 219
- C shell 140
  - .cshrc file 141
  - .login file 141
- C++ map file 198
- CARDINALITY() function 103, 106, 113
- Cascading deletes 56
- Case-insensitive databases 11, 106, 107
- Cast ( :: ) operator 135, 137
- CAST AS keywords 135
- casting to CLOB 115
- Casts 132, 137
  - built-in 20, 133, 135
  - distinct data type 136
  - explicit 20, 135, 135
  - from BYTE to BLOB 89
  - implicit 20, 135, 135
  - rules of precedence 135
  - syscasts data 20
  - user-defined (UDCs) 20
- Casts from TEXT 115
- CHAR data type
  - built-in casts 134

- collation 90, 91
  - defined 90
  - nonprintable characters 91, 91
  - storing numeric values 91
- CHARACTER data type 91
- Character data types
  - Boolean comparisons 116
  - casting between 133
  - data strings 81
  - listed 118
- Character string
  - CHAR data type 90
  - CHARACTER VARYING data type 91
  - CLOB data type 92
  - DATETIME literals 93, 126, 176
  - INTERVAL literals 100
  - LVARCHAR data type 104
  - NCHAR data type 106
  - NVARCHAR data type 107
  - VARCHAR data type 116
  - with DELIMIDENT set 181
- CHARACTER VARYING data type
  - defined 91
- Character-based applications 201, 218
- Check constraints
  - creation-time value 165, 167
  - syschecks data 21
  - syscheckudrdep data 21
  - syscoldepend data 23
  - sysconstraints data 29
- chkenv utility 140
  - error message 143
  - syntax 143
- Chunks 121
- CLIEN\_LABEL environment variable 159
- CLIENT\_LOCALE environment variable 166
- Client/server
  - DataBlade API 122
  - default database 199
  - ONEDB\_SQLHOSTS environment variable 200
  - shared memory communication segments 200
  - stacksize for client session 201
- CLOB data type
  - casting unavailable 92
  - code-set conversion 92
  - collation 92
  - defined 92
  - inserting data 92
  - multibyte characters 92
  - syscolattrs data 22
- CLOB TEXT 115
- CLOSE statement 209
- Clustering 15, 39, 43
- CMCONFIG environment variable 158
- Code sets
  - conversion 147, 147
  - East Asian 91, 176
  - EBCDIC 80
  - ISO 8859-1 34
- Collation 115
  - CHAR data type 90, 91
  - CLOB data type 92
  - GL\_COLLATE table 65
  - NCHAR data type 106
  - NVARCHAR data type 107
  - server\_attribute data 80
- Collection data type
  - casting matrix 137
  - defined 129

- empty 129
- LIST 103
- MULTISET 106
- SET 113
  - sysattrtypes data 18
  - sysxtdesc data 74
  - sysxdtypes data 75, 75
- COLLECTION data type
  - coltype code 24
- collection delimiters 113, 129
- Colon
  - cast ( :: ) operator 135
  - DATETIME delimiter 93
  - INTERVAL delimiter 100
- Color and intensity screen attributes 201
- Column-level privileges
  - sysabauth data 5
  - sysabauth table 65
- Columns
  - changing data type 81, 132
  - constraints (sysconstraints) 29
  - default values (sysdefaults) 30
  - hashed 39
  - in superstores\_demo database 222
  - inserting BLOB data 88
  - range of values 29
  - syscolumns data 24
  - columns Information Schema view 77
- Combine function 14
- Comment indicator 140
- Comment lines 140
- Committed read 80
- Commutator function 53
- Compiling
  - ESQL/C programs 158
  - JAVA\_COMPILER setting 205
  - multithreaded ESQL/C applications 219
  - ONEDB\_C setting 195
- Complex data type 128, 131
  - collection types 129
  - ROW types 130
  - sysattrtypes data 18
- Compliance
  - ANSI/ISO standard for SQL 77, 162
  - sql\_languages.conformance 80
  - X/Open CAE standards 77
  - XPG4 standard 79
- Composite index 41
- Concatenation ( || ) operator 137
- Confidence level 37
- Configuration file
  - .cshrc file 141
  - .informix 140, 143, 182, 184
  - .login file 141
  - .profile file 141
  - for connectivity 199, 200
  - for database servers 182, 207
  - for High-Performance Loader 211
  - for terminal I/O 201
- Configuration parameters
  - DBSPACETEMP 175
  - DEF\_TABLE\_LOCKMODE 184
  - DIRECTIVES 185
  - DISABLE\_B162428\_XA\_FIX 194
  - EXT\_DIRECTIVES 31, 185
  - MITRACE\_OFF 69, 69
  - OPCACHEMAX 199
  - OPT\_GOAL 209
  - OPTCOMPIND 208
  - RESIDENT 186
  - shared memory base 194

- SQL\_LOGICAL\_CHAR 65
- STACKSIZE 201
- STMT\_CACHE 217
- USEOSTIME 93
- CONNECT DEFAULT statement 199
- Connect privilege 10, 72
- CONNECT statement 172, 197, 199
- CONNECT\_RETRIES environment variable 196
- CONNECT\_TIMEOUT environment variable 197
- Connections
  - CONNECT\_RETRIES environment variable 196
  - CONNECT\_TIMEOUT environment variable 197
  - ONEDB\_SERVER environment variable 199
- Connectivity information 200
- Constraints
  - check
    - creation-time value 167
    - syschecks data 21
    - syscheckudrdep data 21
    - syscoldepend data 23
  - column
    - sysconstraints data 29
  - not null
    - collection data types 106, 113, 129
  - NOT NULL
    - collection data types 103
    - syscoldepend data 23
    - syscolumns data 24
    - sysconstraints data 29
  - object mode 47
  - primary key
    - sysconstraints data 29
    - sysreferences data 56
    - unique SERIAL values 111
    - unique SERIAL8 values 112
  - referential
    - sysconstraints data 29
    - sysreferences data 56
  - table
    - sysconstraints data 29
  - unique
    - sysconstraints data 29
    - sysviolations data 73
    - violations 73
- Constructors 113, 129
- Converting data types
  - DATE and DATETIME 135
  - INTEGER and DATE 134
  - number and string 134
  - number to number 133
  - retyping a column 132
- CPFIRST environment variable 158
- CPU cost 216
- CREATE ACCESS\_METHOD statement 15
- CREATE CAST statement 20, 135
- CREATE DATABASE statement 172
- CREATE DISTINCT TYPE statement 75, 98, 223
- CREATE EXTERNAL TABLE statement 35, 36
- CREATE FUNCTION statement 58
- CREATE IMPLICIT CAST statement 223
- CREATE INDEX statement 41, 43, 65, 180
  - storage options 180
- CREATE OPAQUE TYPE statement 107
- CREATE OPERATOR CLASS statement 48
- CREATE OPTICAL CLUSTER statement 49
- CREATE PROCEDURE statement 58, 206
- CREATE ROLE statement 57
- CREATE ROUTINE FROM statement 58, 206

- CREATE ROW TYPE statement 24, 108
- CREATE SCHEMA statement 5
- CREATE SEQUENCE statement 62
- CREATE SYNONYM statement 63, 63
- CREATE TABLE statement
  - assigning data types 81
  - default lock mode 184
  - default privileges 206
  - SET constructor 113
  - typed tables 108
- CREATE TEMP TABLE statement 175
- CREATE TRIGGER statement 71
- CREATE VIEW statement 5, 72
- CREATE XDATASOURCE statement 73
- CREATE XDATASOURCETYPE statement 74
- Currency symbol 105, 171
- Current date 30, 163
- CURRENT keyword 123

## D

- Data corruption 10, 22
- Data dependencies
  - syscheckudrdep data 21
  - syscoldepend data 23
  - sysdepend data 31
- Data dictionary 4
- Data distributions 10, 32, 179
- Data integrity 80
- Data pages 22, 41, 65
- data type collation 115
- data type restrictions 115
- data type restrictions in Boolean expression 115
- data type UPDATE statements 115
- Data types
  - approximate 79
  - BIGINT 87
  - BIGSERIAL 87
  - BLOB 88
  - BOOLEAN 89
  - BYTE 89
  - casting 132, 137
  - CHAR 90
  - CHARACTER 91
  - CHARACTER VARYING 91
  - classified by category 81
  - CLOB 92
  - collection 129
  - complex 128
  - conversion 132
  - DATE 93
  - DATETIME 93
  - DEC 96
  - DECIMAL 96
  - distinct 131
  - DISTINCT 98
  - DOUBLE PRECISION 99
  - exact numeric 79
  - extended 128
  - fixed point 97
  - FLOAT 99
  - floating-point 96, 99, 114
  - IDSSECURITYLABEL 99, 118
  - inheritance 108
  - INT 100
  - INT8 100
  - INTEGER 100
  - internal 81
  - INTERVAL 100
  - length (syscolumns) 27
  - LIST 103
  - LVARCHAR 104
  - MONEY 105
  - MULTISET 106
  - named ROW 108
  - NCHAR 106, 106
  - NUMERIC 107
  - NVARCHAR 107
  - opaque 131
  - OPAQUE 107
  - Opaque data types
    - smart large objects 122
  - REAL 108
  - ROW 108, 109
  - sequential integer 112
  - SERIAL 111
  - SERIAL8 112
  - SET 113
  - simple large object 121
  - SMALLFLOAT 114
  - SMALLINT 114
  - smart large object 122
  - summary list 81
  - unique numeric value 112
  - unnamed ROW 109
  - VARCHAR 116
- Data-type promotion 118
- Database identifiers 181
- Database server administrator (DBSA) 5
- Database Server Administrator (DBSA) 202
- Database servers
  - attributes in Information Schema view 80
  - code set 80
  - default connection 199
  - default isolation level 80
  - optimizing queries 209
  - pathname for 172
  - remote 183
  - role separation 202
  - server name 30, 172
- DATABASE statement 172
- Databases
  - data types 81
  - Databases
    - superstores\_demo 222
  - demonstration databases
    - superstores\_demo 222
  - identifiers 181
  - joins in stores\_demo 220
  - object-relational 222
  - objects, sysobjstate data 47
  - privileges 72
  - stores\_demo 220
  - superstores\_demo
    - demonstration database 222
  - syscmd 5
  - sysmaster 5
  - sysutils 5
  - sysuid 5
- DataBlade modules
  - Client and Server API 122
  - data types (sysbuiltintypes) 5
  - trace messages (systracemsgs) 69, 69
  - user messages (syserrors) 34
- DATE data type
  - abbreviated year values 163
  - casting to integer 134
  - coltype code 24
  - converting to DATETIME 135
  - defined 93
  - display format 166
  - in expressions 123, 125

- international date formats 93
- source data 125
- DATE() function 126, 166
- DATETIME data type
  - abbreviated year values 163
  - coltype code 24
  - converting to DATE 134, 135
  - defined 93
  - display format 176, 176
  - EXTEND function 125
  - extending precision 124
  - field qualifiers 93
  - in expressions 123, 127
  - international formats 93, 93, 100
  - length (syscolumns) 27
  - literal values 93
  - localized values 93
  - precision and size 93
  - source data 126
  - two-digit year values and DBDATE variable 93
  - year to fraction example 93
- DAY keyword
  - DATETIME qualifier 93
  - INTERVAL qualifier 100
  - UNITS operator 93, 126
- DB-Access utility 10, 77, 144, 169, 172, 176, 199
- DBA privilege 34, 69, 69, 72
- DBA routines 53
- DBACCNOIGN environment variable 161, 162
- DBANSIWARN environment variable 162
- DBBLOBBUF environment variable 163, 163
- DBCENTURY environment variable
  - defined 163
  - effect on functionality of DBDATE 166
  - expanding abbreviated years 93, 164
- DBDATE environment variable 93, 93, 166
- DBDELIMITER environment variable 168
- DBEDIT environment variable 168
- dbexport utility 168
- DBFLTMASK environment variable 169
- dbinfo utility 169
- DBINFO\_DBSPACE\_RETURN\_NULL\_FOR\_INVALID\_PARTNUM environment variable 169
- DBLANG environment variable 170, 170
- dbload utility 88, 89, 115, 168
- DBMONEY environment variable 105, 171
- DBONPLOAD environment variable 171
- DBPATH environment variable 172
- DBPRINT environment variable 174
- DBREMOTECMD environment variable 174
- dbschema utility 53
- DBSECADM role 99, 118
- dbservername.cmd batch file 146
- dbspace
  - for BYTE or TEXT values 20
  - for system catalog 5
  - for table fragments 36
  - for temporary tables 175
  - name 181
- DBSPACETEMP configuration parameter 175
- DBSPACETEMP environment variable 175
- DBTEMP environment variable 176
- DBTIME environment variable 93, 176, 176
- DBUPSPACE environment variable 179
- DEC data type 96
- DECIMAL data type
  - built-in casts 133, 134
  - coltype code 24
  - defined 96
- disk storage 97
- display format 169, 171
- fixed point 97
- floating point 96
  - length (syscolumns) 27
- Decimal digits, display of 169
- Decimal point
  - DBFLTMASK setting 169
  - DBMONEY setting 171
  - DECIMAL radix 97
- Decimal separator 171
- DECLARE statement 209
- DECRYPT\_BINARY function 92
- DECRYPT\_CHAR function 92
- DEF\_TABLE\_LOCKMODE configuration parameter 184
- Default database locale 11
- DEFAULT\_ATTACH environment variable 180
- Defaults
  - C compiler 195
  - century 163, 176
  - CHAR length 90
  - character set for SQL identifiers 181
  - compilation order 158
  - configuration file 207
  - connection 199
  - data type 110
  - database server 172, 199
  - DATE display format 93
  - DATE separator 166
  - DATETIME display format 93
  - DECIMAL precision 96
  - disk space for sorting 179
  - fetch buffer size 183
  - heap size 205
  - index storage location 180
  - isolation level 80
  - join method 208
  - level of parallelism 211
  - lock mode 184
  - message directory 170
  - MONEY scale 105
  - operator class 15, 48
  - printing program 174
  - query optimizer goal 209
  - sysdefaults.default 30
  - table privileges 206
  - temporary dbspace 175
  - terminfo direcotry 218
  - text editor 168
- DEFINE statement of SPL 111, 112
- defined Data types 115
- Delete privilege 36, 65, 206
- DELETE statement 73
- DELETE statements 10
- Delete trigger 71
- DELIMIDENT environment variable 181
- Delimited identifiers 181, 181
- Delimiter
  - for DATETIME values 93
  - for fields 168
  - for identifiers 181
  - for INTERVAL values 100
- demonstration databases
  - stores\_demo 220
- Demonstration databases
  - tables 222
- Descending index 41
- DESCRIBE statement 193
- Describe-for-updates 193
- destroy() support function 128

- Detached index 180
- Deutsche mark (DM) currency symbol 171
- Diagnostics table 73
- DIRECTIVES configuration parameter 185
- Directives for query optimization 185, 208, 209
- Disabled database objects 73
- Disk space
  - for data distributions 179
  - for temporary data 175
- Distinct data types
  - casts 136
  - sysxdtypes data 75
- DISTINCT data types
  - defined 98
  - sysxtdesc data 74
  - sysxdtypes data 75, 98
- Distributed Computing Environment (DCE) 219
- Distributed queries 128, 183
- Dollar (\$) sign 105, 171
- double (C) data type 99
- Double-precision floating-point number 99
- DROP CAST statement 223
- DROP DATABASE statement 172
- DROP FUNCTION statement 53
- DROP INDEX statement 65
- DROP OPTICAL CLUSTER statement 49
- DROP PROCEDURE statement 53
- DROP ROUTINE statement 53
- DROP ROW TYPE statement 108
- DROP SEQUENCE statement 220
- DROP TABLE statement 220
- DROP TYPE statement 98, 107
- DROP VIEW statement 77, 220

## E

- EBCDIC collation 80
- Editor, DBEDIT setting 168
- EMACS text editor 168
- Empty set 129
- ENCRYPT\_DES function 92
- ENCRYPT\_TDES function 92
- Enterprise Replication 5
- env utility 142
- ENVIGNORE environment variable
  - defined 140, 182
  - relation to chkenv utility 143
- Environment configuration file
  - debugging with chkenv 143
  - setting environment variables in UNIX 139, 140
- Environment variables
  - ANSIOWNER 157
  - CLIENT\_LABEL 159
  - CLIENT\_LOCALE 166, 166
  - CMCONFIG 158
  - Colon
    - pathname separator 205
  - command-line utilities 145
  - CONNECT\_RETRIES 196
  - CONNECT\_TIMEOUT 197
  - CPFIRST 158
  - DBACCNOIGN 161, 162
  - DBANSIWARN 162
  - DBBLOBBUF 163, 163
  - DBCENTURY 163
  - DBDATE 93, 93, 166
  - DBDELIMITER 168
  - DBEDIT 168
  - DBFLTMASK 169
  - DBINFO\_DBSPACE\_RETURN\_NULL\_FOR\_INVALID\_PARTNUM 170
  - DBLANG 170

DBMONEY 105, 171  
 DBONPLOAD 171  
 DBPATH 172  
 DBPRINT 174  
 DBREMOTECMD 174  
 DBSPACETEMP 175, 175  
 DBTEMP 176  
 DBTIME 93, 176  
 DBUPSPACE 179  
 DEFAULT\_ATTACH 180  
 DELIMIDENT 181  
 displaying current settings 142, 145  
 ENVIGNORE 182  
 FET\_BUF\_SIZE 183  
 GL\_DATE 93, 93, 165  
 GL\_DATETIME 93, 165  
 how to set  
     in Bourne shell 141  
     in C shell 141  
     in Korn shell 141  
 how to set in Bourne shell 141  
 how to set in Korn shell 141  
 IFXMONGOAUTH 184  
 IFX\_DEF\_TABLE\_LOCKMODE 184  
 IFX\_DIRECTIVES 185  
 IFX\_EXTDIRECTIVES 31, 185  
 IFX\_LARGE\_PAGES 186  
 IFX\_LOB\_XFERSIZE 187  
 IFX\_LONGID 188  
 IFX\_NETBUF\_PVTPOOL\_SIZE 188  
 IFX\_NETBUF\_SIZE 189  
 IFX\_NO\_SECURITY\_CHECK 189  
 IFX\_NO\_TIMELIMIT\_WARNING 190  
 IFX\_NOBPROC 190  
 IFX\_NOT\_STRICT\_THOUS\_SEP 190  
 IFX\_ONTAPE\_FILE\_PREFIX 190  
 IFX\_PAD\_VARCHAR 191  
 IFX\_SMX\_TIMEOUT 191, 192  
 IFX\_SMX\_TIMEOUT\_RETRY 192  
 IFX\_UNLOAD\_EILSEQ\_MODE 193  
 IFX\_UPDESC 193  
 IFX\_XASTDCOMPLIANCE\_XAEND 194  
 IFX\_XFER\_SHMBASE 194  
 IFXRESFILE 194  
 INF\_ROLE\_SEP 202, 202  
 INTERACTIVE\_DESKTOP\_OFF 202  
 IONEDB\_OPCACHE 199  
 ISM\_COMPRESSION 203  
 ISM\_DEBUG\_FILE 203  
 ISM\_DEBUG\_LEVEL 203  
 ISM\_ENCRYPTION 204  
 ISM\_MAXLOGSIZE 204  
 ISM\_MAXLOGVERS 204  
 JAR\_TEMP\_PATH 205  
 JAVA\_COMPILER 205  
 JVM\_MAX\_HEAP\_SIZE 205  
 LD\_LIBRARY\_PATH 205  
 LIBPATH 206  
 limitations 139  
 listed by topic 147  
 manipulating in Windows environments 144  
 modifying settings 142  
 NODEFDAC 206  
 ONCONFIG 207  
 ONEDB\_CMCONUNITNAM 196  
 ONEDB\_CMNAME 195  
 ONEDB\_CPPMAP 198  
 ONEDB\_SHMBASE 200  
 ONEDB\_SQLHOSTS 200, 200  
 ONEDB\_STACKSIZE 201

ONEDB\_TERM 201  
 ONEDB\_C 195  
 ONEDB\_HOME 198  
 ONEDB\_SERVER 199  
 ONINIT\_STDOUT 207  
 OPT\_GOAL 209  
 OPTCOMPIND 208  
 OPTMSG 208  
 OPTOFC 209  
 overriding a setting 140, 182  
 PATH 210, 210  
 Pathname  
     for client or shared libraries 205  
 PDQPRIORITY 210  
 PLCONFIG 211  
 PLOAD\_LO\_PATH 212  
 PLOAD\_SHMBASE 212  
 PSM\_ACT\_LOG 212, 213  
 PSM\_DBS\_POOL 213  
 PSM\_DEBUG 213  
 PSM\_DEBUG\_LOG 214  
 PSM\_LOG\_POOL 214  
 PSORT\_DBTEMP 214  
 PSORT\_NPROCS 215  
 RTREE\_COST\_ADJUST\_VALUE 216  
 rules of precedence in UNIX 143  
 rules of precedence in Windows 146  
 scope of reference 145  
 setting 144  
     at the command line 139  
     in a configuration file 139  
     in a login file 139  
     in a shell file 141  
     in Windows environments 144  
     with the System applet 145  
 setting in autoexec.bat 145  
 SHLIB\_PATH 216  
 SRV\_FET\_BUF\_SIZE 216  
 standard UNIX system 138  
 STMT\_CACHE 217  
 TERM 218  
 TERMCAP 218  
 TERMINFO 219  
 THREADLIB 219  
 types of 138  
 unsetting 142, 145, 181  
 USE\_DTENV 93, 93  
 USETABLENAME 220  
 view current setting 142  
 where to set 141  
 equal() support function 128  
 Equality (=) operator 92  
 Era-based dates 176  
 Error message files 170  
 esql command 158, 195  
 ESQL/C  
     DATETIME routines 176  
     esqlc command 158  
     long identifiers 188  
     message chaining 208  
     multithreaded applications 219  
     program compilation order 158  
 Exact numeric data types 79  
 Executable programs 210  
 Execute privilege 51, 206  
 explain output file 179  
 Explicit cast 20, 135  
 Explicit pathnames 145, 173  
 Explicit temporary tables 175  
 Exponent 97  
 Exponential notation 96

export utility 141  
 export\_binary() support function 128  
 export() support function 128  
 Expression-based fragmentation 37, 39, 165, 167  
 EXT\_DIRECTIVES configuration parameter 31, 185  
 EXTEND function 125  
 Extended data types 75, 128, 128, 223  
 Extensible Markup Language (XML) 92  
 Extension checking (DBANSIWARN) 162  
 Extents, changing size 10  
 External database 63  
 External directives for query optimization 185  
 External routines 53  
 External tables  
     sysextcols data 35  
     sysextfiles data 35  
     sysexternal data 36  
     systables data 65  
 External view 63  
 extspace 15

## F

FALSE setting  
     BOOLEAN value 89  
 Farsi locales 90  
 FET\_BUF\_SIZE environment variable 183  
 Fetch buffer 183  
 Fetch buffer size 183  
 FETCH statement 209  
 Field delimiter  
     DBDELIMITER 168  
     Statements of SQL  
         LOAD 168  
         UNLOAD 168  
     Utilities  
         dbexport 168  
 Field of a ROW data type 130  
 Field qualifier  
     DATETIME values 93  
     EXTEND function 125  
     INTERVAL values 100  
 Fields of a ROW data type 130  
 File extensions  
     .a 188  
     .cmd 146  
     .ec 158  
     .ecp 158  
     .iem 170  
     .jar 205  
     .rc 140, 143, 182, 184  
     .so 188  
     .sql 77, 172, 172, 181  
     .std 207, 217  
 Files  
     environment configuration files 143  
     installation directory 198  
     permission settings 140  
     shell 141  
     temporary 175, 176, 214  
     temporary for SE 176  
     termcap, terminfo 201, 218, 219  
 FILETOBLOB function 88  
 FILETOCLOB function 92  
 Filtering mode 47, 73  
 Finalization function 14  
 Fixed point decimal 97, 105, 171  
 Fixed-length opaque data types 24  
 Fixed-length UDT 75  
 FLOAT data type

- built-in casts 133, 134
- coltype code 24
- defined 99
- display format 169, 171
- Floating-point decimal 96, 99, 114, 169
- Formatting
  - DATE values with DBDATE 166
  - DATE values with GL\_DATE 176
  - DATETIME values with DBTIME 176
  - DATETIME values with GL\_DATETIME 176
  - DATETIME values with USE\_DTENV 176
  - DECIMAL(p) values with DBFLTMASK 169
  - FLOAT values with DBFLTMASK 169
  - MONEY values with DBMONEY 171
  - SMALLFLOAT values with DBFLTMASK 169
- Formatting mask
  - with DBDATE 166
  - with DBFLTMASK 169
  - with DBMONEY 171
  - with DBTIME 176
  - with GL\_DATE 176
  - with GL\_DATETIME 176
  - with USE\_DTENV 176
- FRACTION keyword
  - DATETIME qualifier 93
- FRAGMENT BY clause 175
- Fragment-level statistics 37
- Fragmentation
  - distribution strategy 39
  - encrypted distribution 37
  - expression 37, 39, 165, 167
  - fragment statistics 37
  - list 39
  - PDQPRIORITY environment variable 211
  - PSORT\_NPROCS environment variable 216
  - round robin 37, 39
  - setting priority levels for PDQ 210
  - sysfragauth data 36
  - sysfragdist data 37
  - sysfragments data 39
- FROM keyword 10, 21
- Function keys 201
- Functional index 41, 130, 180
- Functions
  - for BLOB columns 88
  - for CLOB columns 92
  - for MULTISSET columns 106
  - support for complex types 128
- fwritable gcc option 195

## G

- gcc compiler 195
- Generic B-trees 41
- GET DIAGNOSTICS statement 34
- getenv utility 139
- GL\_COLLATE table 65
- GL\_CTYPE table 65
- GL\_DATE environment variable 93, 93, 165, 166
- GL\_DATETIME environment variable 93, 165, 166
- Global network buffer pool 188
- GLS environment variables 143
- GNU C compiler 195
- GRANT statement 57, 65
- Graphic characters 201
- GROUP BY clause 89, 115, 175
- GROUP BY TEXT 115
- Group informix 170

## H

- Hash-join 208
- hash() support function 128

- Hashed columns 39
- Hashing parameters 64
- HCL OneDB
  - ESQL/C
    - 158, 166, 176, 188, 208
  - HCL OneDB
    - extension checking (DBANSIWARN)
      - 162
    - HCL
      - OneDB
        - Storage Manager
          - 204
        - Heap size 205
        - Hebrew locales 90
        - Hexadecimal digits 168
      - HIGH INTEG keywords
        - ALTER TABLE statement 122
        - CREATE TABLE statement 122
      - HIGH keyword
        - in UPDATE STATISTICS statement 37
        - PDQPRIORITY 210
        - UPDATE STATISTICS 10, 32
      - High-Performance Loader 171, 211
      - Histogram 32
      - Host language 80
      - Host variable 88, 89, 115, 130
      - HOUR keyword
        - DATETIME qualifier 93
        - INTERVAL qualifier 100
      - HP-UX operating system 216
      - HTML (Hypertext Markup Language) 92
      - Hyphen
        - DATETIME delimiter 93
        - INTERVAL delimiter 100
    - I/O overhead 216
    - IDSSECURITYLABEL data type
      - definition 99
    - IFMXMONGOAUTH environment variable 184
    - IFX\_DEF\_TABLE\_LOCKMODE environment variable 184
    - IFX\_DIRECTIVES environment variable 185
    - IFX\_EXTDIRECTIVES environment variable 31, 185
    - IFX\_LARGE\_PAGES environment variable 186
    - IFX\_LOB\_XFERSIZE environment variable 187
    - IFX\_LONGID environment variable 188
    - IFX\_NETBUF\_PVTPOOL\_SIZE environment variable 188
    - IFX\_NETBUF\_SIZE environment variable 189
    - IFX\_NO\_SECURITY\_CHECK environment variable 189
    - IFX\_NO\_TIMELIMIT\_WARNING environment variable 190
    - IFX\_NODBPROC environment variable 190
    - IFX\_NOT\_STRICT\_THOUS\_SEP environment variable 190
    - IFX\_ONTAPE\_FILE\_PREFIX environment variable 190
    - IFX\_PAD\_VARCHAR environment variable 191
    - IFX\_SMX\_TIMEOUT environment variable 191, 192
    - IFX\_SMX\_TIMEOUT\_RETRY environment variable 192
    - IFX\_UNLOAD\_EILSEQ\_MODE environment variable 193
    - IFX\_UPDDESC environment variable 193
    - IFX\_XASTDCOMPLIANCE\_XAEND environment variable 194

- IFX\_XFER\_SHMBASE environment variable 194
- IFXRESFILE environment variable 194
- IMPEXP data type 135
- IMPEXPBIN data type 135
- Implicit cast 20, 135
- Implicit connection 199
- Implicit temporary tables 175
- import\_binary() support function 128
- import() support function 128
- IN clause 175
- IN keyword 89, 106, 110, 113, 137
- IN TABLE storage option 180
- Index
  - attached 39, 165, 180, 215
  - B-tree 41, 180
  - clustered 41, 43
  - composite 41, 41
  - default values for attached 215
  - descending 41
  - detached 180
  - distribution scheme 180
  - forest of trees 180
  - fragmented 37, 39
  - functional 41, 130, 180
  - nonfragmented 180, 180
  - of data types 81
  - of environment variables 147
  - of system catalog tables 11
  - R-Tree 180
  - sysindexes data 41
  - sysindext data 43
  - sysobjstate data 47
  - threads for sorting 215
  - unique 29, 41, 111, 112
- Index key structure 43
- Index privilege 65
- Indirect typing 111, 112
- Industry standards, compliance with 80
- INF\_ROLE\_SEP environment variable 202, 202
- Information Schema views
  - accessing 78
  - columns 79
  - defined 77, 77
  - generating 77
  - server\_info 80
  - sql\_languages 80
  - tables 78
- Informational messages 34
- informix owner name 10, 20, 32, 41, 43, 65, 170, 202
- Inheritance hierarchy 46, 109
- Initialization function 14, 58
- Input support function 104
- input() support function 128
- Insert privilege 36, 65, 206
- INSERT statements 69, 73, 93, 129, 161, 166
- Insert trigger 71
- Installation directory 198
- INSTEAD OF trigger 71
- INT data type 100
- INT8 data type
  - built-in casts 133, 134
  - coltype code 24
  - defined 100
  - using with SERIAL8 88
- INTEG keyword 122
- INTEGER data type
  - built-in casts 133, 134
  - coltype code 24
  - defined 100

- length (syscolumns) 27
- Intensity attributes 201
- INTERACTIVE\_DESKTOP\_OFF environment variable 202
- Internationalized trace messages 69
- Interprocess communications (IPC) 200
- INTERVAL data type
  - coltype code 24
  - defined 100
  - field delimiters 100
  - in expressions 123, 123, 127, 127
  - length (syscolumns) 27
- IONEDB\_OPCACHE environment variable 199
- ipcshm protocol 200
- IS NULL operator 89
- ISM\_COMPRESSION environment variable 203
- ISM\_DEBUG\_FILE environment variable 203
- ISM\_DEBUG\_LEVEL environment variable 203
- ISM\_ENCRYPTION environment variable 204
- ISM\_MAXLOGSIZE environment variable 204
- ISM\_MAXLOGVERS environment variable 204
- ISO 8859-1 code set 80
- Isolation level 80, 208
- Iterator functions 14

## J

- Japanese eras 176
- Jar management procedures 205
- JAR\_TEMP\_PATH environment variable 205
- Java virtual machine (JVM) 157, 205, 205, 205
- JAVA\_COMPILER environment variable 205
- JIT compiler 205
- Join methods 208
- Join operations 10, 175
- JVM\_MAX\_HEAP\_SIZE environment variable 205

## K

- KEEP ACCESS TIME keywords
  - ALTER TABLE statement 122
  - CREATE TABLE statement 122
- Key
  - primary 29, 56, 56, 73, 222
- Key scan 15
- Keyboard I/O
  - ONEDB\_TERM setting 201
  - TERM setting 218
  - TERMCAP setting 218
  - TERMINFO setting 219
- keyword MATCHES 115
- Korn shell 140, 141

## L

- Label-based access control (LBAC) 99, 118
- Language
  - C 58, 158, 195
  - C++ 198
  - CLIENT\_LOCALE setting 166
  - DBLANG setting 170
  - Extensible Markup Language (XML) 92
  - HCL OneDB
  - ESQL/C
    - 122, 130, 219
  - Hypertext Markup Language (HTML) 92
  - Java 157, 205, 205
  - sql\_languages information schema view 80
  - Stored Procedure Language (SPL) 130, 165, 167
  - syslangauth data 46
  - sysroutinelangs data 58
- Large pages for virtual memory segments 186
- Large-object data type

- defined 121
- listed 118
- LD\_LIBRARY\_PATH environment variable 205
- Leaf pages 39
- libos.a library 188
- LIBPATH environment variable 206
- LIKE 115
- LIKE keyword of SPL 111, 112
- LIKE operator 89, 137
- Linearized code 70
- List
  - of data types 81
  - of environment variables, by topic 147
  - of system catalog tables 11
- LIST data type
  - coltype code 24
- LIST data type, defined 103
- LO\_handles() support function 128
- LOAD statement 88, 89, 115, 168
- Locales
  - collation order 65
  - multibyte 91
  - of trace messages 69
  - right-to-left 90
  - specifying 147, 147
- Localized collation 118
- Lock-table overflow 184
- LOCKMODE keyword 184
- LOCOPY function 88, 92
- LOG keyword
  - ALTER TABLE statement 122
  - CREATE TABLE statement 122
- Logging mode 22
- Logical characters 118
- Long identifiers
  - client version 188
  - IFX\_LONGID setting 188
  - Information Schema views 78
- LOTOFILE function 88, 92
- LOW keyword
  - PDQPRIORITY 210
  - UPDATE STATISTICS 32
- Lowercase mode codes 53
- Lowercase privilege codes 5, 23, 65
- LVARCHAR data type
  - casting opaque types 135
  - coltype code (for client) 24
  - defined 104

## M

- Machine notes 201
- Machine-independent integer types 27
- Magnetic storage media 20
- Mantissa precision 79, 97
- Map file for C++ programs 198
- MATCHES 115
- MATCHES operator 89, 137
- MEDIUM keyword 10, 29, 32
- MEDIUM keyword, in UPDATE STATISTICS statement 37
- Membership operator 137
- Memory cache, for staging blobspace 199
- MERGE statement 73
- Message file
  - specifying subdirectory with DBLANG 170
- Messages
  - chaining 208
  - error in syserrors 34
  - optimized transfers 208
  - reducing requests 209
  - trace message template 69

- warning in syserrors 34
- mi\_collection\_card() function 103, 106, 113
- mi\_db\_error\_raise() function 34
- Microsoft C compiler 195
- MINUTE keyword
  - DATETIME qualifier 93
  - INTERVAL qualifier 100
- MITRACE\_OFF configuration parameter 69, 69
- mkdir utility 170
- MODERATE INTEG keywords
  - ALTER TABLE statement 122
  - CREATE TABLE statement 122
- Modifiers
  - CLASS 53
  - COSTFUNC 53
  - HANDLESNULLS 53
  - INTERNAL 53
  - NEGATOR 53
  - NOT VARIANT 53
  - PARALLELIZABLE 53
  - SELCONST 53
  - STACK 53
  - VARIANT 53
- MODIFY NEXT SIZE keywords 10
- MONEY data type
  - built-in casts 134
  - coltype code 24
  - defined 105
  - display format 171
  - international money formats 105
  - length (syscolumns) 27
- MONTH keyword
  - DATETIME qualifier 93
  - INTERVAL qualifier 100
- Multibyte characters
  - CLOB data type 92
- MULTISET data type
  - coltype code 24
  - constructor 129
  - defined 106

## N

- N setting
  - sysroleauth.is\_grantable 57
- Named ROW data type
  - casting permitted 137
  - defined 108
  - defining 108
  - equivalence 108
  - inheritance 46, 108
  - typed tables 108
- Namer ROW data type
  - coltype code 24
- National Language Support (NLS) 118
- NCHAR data type
  - collation order 106
  - coltype code 24
  - defined 106
  - multibyte characters 106
- Negator functions 53
- Nested dot notation 130
- Nested-loop join 208
- Network buffers 189
- Network environment variable, DBPATH 172
- NFS directory 176
- NLS data types
  - in system catalog tables 11
- NO KEEP ACCESS TIME keywords
  - ALTER TABLE statement 122
  - CREATE TABLE statement 122
- no setting of NODEFDAC 206



NODEFDAC environment variable 206  
 NOLOG keyword  
     ALTER TABLE statement 122  
     CREATE TABLE statement 122  
 Non-default database locales 11  
 NONE setting  
     JAVA\_COMPILER 205  
 Nonfragmented index 180  
 Nonprintable characters  
     CHAR data type 91  
     TEXT data type 116  
     VARCHAR data type 116  
 NOT NULL 115  
 NOT NULL constraint  
     collection elements 103, 106, 113, 129  
     syscoldepend data 23  
     sysconstraints data 29  
 NOT NULL keywords 89, 103  
 NOT operator 137  
 NOT VARIANT routine 53  
 NULL data type  
     coltype code 24  
 NULL value  
     allowed or not allowed 14, 24  
     BOOLEAN literal 89  
     BYTE data type 89  
 Numeric data types  
     casting between 133  
     casting to character types 134  
     listed 118  
 NVARCHAR data type  
     collation order 107  
     coltype code 24  
     defined 107  
     multibyte characters 107

## O

Object mode of database objects 47  
 Object-relational schema 222  
 ODBC driver 205, 216  
 OFF setting  
     IFX\_DIRECTIVES 185, 185  
     PDQPRIORITY 210  
 ON setting  
     IFX\_DIRECTIVES 185, 185  
 ONCONFIG environment variable 207  
 onconfig.std file 217  
 ONEDB\_CMCONUNITNAM environment variable 196  
 ONEDB\_CMNAME environment variable 195  
 ONEDB\_CPPMAP environment variable 198  
 ONEDB\_SHMBASE environment variable 200  
 ONEDB\_STACKSIZE environment variable 201  
 ONEDB\_TERM environment variable 201  
 ONEDB\_C environment variable 195  
 ONEDB\_HOME environment variable 198  
 ONEDB\_SERVER environment variable 199  
 onedb.rc file 140, 143, 184  
 oninit command 184  
 ONINIT\_STDOUT environment variable 207  
 Online transaction processing (OLTP) 39  
 onload utility 115  
 onpload utility 171, 212  
 onsecurity utility 189  
 onstat utility 138  
 Opaque data types  
     cast matrix 137  
     comparing 135  
     storage 104  
     sysxtddesc data 74  
     sysxtotypes data 75

OPAQUE data types  
     defined 107  
 OPCACHEMAX configuration parameter 199  
 OPEN statement 209  
 Operator class  
     sysams data 15  
     sysindices data 43  
     sysopclasses data 48  
 operator LIKE 115  
 Operator precedence 137  
 operator TEXT 115  
 OPT\_GOAL configuration parameter 209  
 OPT\_GOAL environment variable 209  
 OPTCOMPIND configuration parameter 208  
 OPTCOMPIND environment variable 208  
 Optical cluster  
     IONEDB\_OPCACHE setting 199  
     sysblobs.type column 20  
     sysopclstr data 49  
 Optimizer  
     setting IFX\_DIRECTIVES 185  
     setting IFX\_EXTDIRECTIVES 185  
     setting OPT\_GOAL 209  
     setting OPTCOMPIND 208  
     setting OPTOFC 209  
 Optimizer directives  
     sysdirectives data 31  
 OPTMSG environment variable 208  
 OPTOFC environment variable 209  
 OR operator 137  
 ORDER 115  
 ORDER BY clause 89, 175  
 Ordinal positions 103  
 Output support function 104  
 output() support function 128  
 Overflow error 97  
 Owner routines 53, 206

## P

Page footers in sbspaces 122  
 Page headers in sbspaces 122  
 PAGE lock mode 65, 184  
 Parallel distributed queries, setting with  
     PDQPRIORITY 210  
 Parallel sorting, setting with  
     PSORT\_NPROCS 214  
 Partial characters 90  
 Partial-column index 43  
 PATH environment variable 210, 210  
 Pathname  
     Configuration file  
         for terminal I/O 218  
     for C compiler 195, 195  
     for C++ map file 198  
     for connectivity information 200  
     for database server 172  
     for dynamic-link libraries 206, 216  
     for environment-configuration file 143  
     for executable programs 210  
     for installation 198  
     for message files 170, 170  
     for parallel sorting 214  
     for remote shell 174  
     for smart-large-object handles 212  
     for temporary .jar files 205  
     for termcap file 218  
     for terminfo directory 219  
     separator symbols 210  
 PDQ  
     OPTCOMPIND environment variable 208  
     PDQPRIORITY environment variable 210

Percentage (%) symbol 176  
 Period  
     DATE delimiter 166  
     DATETIME delimiter 93  
     INTERVAL delimiter 100  
 Permissions 140, 170  
 PLCONFIG environment variable 211  
 plconfig file 211  
 PLOAD\_LO\_PATH environment variable 212  
 PLOAD\_SHMBASE environment variable 212  
 PostScript 92  
 Precedence rules  
     for casts 135  
     for lock mode 184  
     for SQL operators 137  
     for UNIX environment variables 143  
     for Windows environment variables 146  
 Precision  
     of currency values 105  
     of numbers 79, 96, 99, 100, 100, 114  
     of time values 93, 100, 124, 127  
 PREPARE statement 65  
 Prepared statement 65  
 Primary access method 15, 64  
 Primary key 29, 56, 73, 111, 112, 222  
 Primary thread 201  
 printenv utility 142  
 Printing with DBPRINT 174  
 Private environment-configuration file 143, 182  
 Private network buffer pool 188, 189  
 Private synonym 65  
 Privilege  
     default table privileges 206  
     on columns (syscolauth table) 23  
     on procedures and functions (sysprocauth table) 51  
     on table fragments (sysfragauth table) 36  
     on tables (systabauth table) 65  
     on the database (sysusers table) 72  
     on UDTs and named row types (sysxdttypeauth) 75  
 Protected routines 53  
 Protected rows 99, 118  
 Pseudo-machine code (p-code) 51  
 PSM\_ACT\_LOG environment variable 212  
 PSM\_CATALOG\_PATH environment variable 213  
 PSM\_DBS\_POOL environment variable 213  
 PSM\_DEBUG environment variable 213  
 PSM\_DEBUG\_LOG environment variable 214  
 PSM\_LOG\_POOL environment variable 214  
 PSORT\_DBTEMP environment variable 214  
 PSORT\_NPROCS environment variable 215  
 Public synonym 63, 65  
 public user name 78  
 Purpose functions 15  
 putenv utility 139

## Q

Qualifier field  
     DATETIME 93  
     EXTEND 127  
     INTERVAL 100  
     UNITS 126  
 Query optimizer  
     directives 185, 185  
     sysdistrib data 32  
     sysprocplan data 56  
     updating distribution data 10  
 Quoted string  
     DATE and DATETIME literals 126

- DELIMIDENT setting 181
- INTERVAL literals 100
- invalid with BYTE 89
- LVARCHAR data type 104
- Quoted string invalid with TEXT 115

## R

- R-tree index 180, 216
- Read committed 80
- Read uncommitted 80
- recv() support function 128
- References privilege 23, 65
- Referential constraint 29, 56, 73
- Relational operators 91, 137
- Remote database server 63, 183
- Remote shell 174
- Remote tape devices 174
- RENAME SEQUENCE statement 220
- Repeatable read 208
- Replica identifier 39
- RESIDENT configuration parameter 186
- Resource contention 210
- Resource Grant Manager (RGM) 39
- Resource privilege 10
  - Role
    - sysusers data 72
  - System catalog
    - authorization identifiers 72
- REVOKE statement 65
- Right-to-left locales 90
- Role
  - default role 72
  - INF\_ROLE\_SEP setting 202
  - sysroleauth data 57
- Role separation 202
- Rolling-window fragmentation 39
- Round-robin fragmentation 37, 39
- Routines
  - DataBlade API routine 69
  - DATETIME formatting 176
  - identifier 53
  - owner 53
  - privileges 51
  - protected 53
  - restricted 53
  - Stored Procedure Language (SPL) 130
  - syserrors data 34
  - syslangauth data 46
  - sysprocauth data 51
  - sysprocbody data 51
  - sysprocedures data 53
  - sysprocplan data 56
  - sysroutinelangs data 58
  - systraceclasses data 69
  - sysracemsgs data 69
  - trigger 53
- ROW data types 130
  - casting permitted 137
  - equivalence 108
  - fields 18, 130
  - inheritance 46, 108
  - inserting values 110
  - named 108, 130
  - sysattrtypes data 18
  - sysxtddesc data 74
  - sysxtotypes data 75, 75
  - unnamed 109, 130
- ROW lock mode 65, 184
- ROWIDS 15
- RTNPARAMTYPES data type 53

- RTREE\_COST\_ADJUST\_VALUE environment variable 216
- Runtime
  - warnings (DBANSIWARN) 162

## S

- Sample size 32
- Sampling data 37
- SAVE EXTERNAL DIRECTIVES statement 185
- SBSPECNAME configuration parameter 32, 37
- sbspaces
  - defined 92, 122
  - name 181
  - sysams data 15
  - syscolattns data 22
  - sysstabamdata data 64
- Scale of numbers 79, 97, 169
- Scan cost 15
- Schema Tools 144
- SECOND keyword
  - DATETIME qualifier 93
  - FRACTION keyword
    - INTERVAL qualifier 100
  - INTERVAL qualifier 100
- Secondary-access methods 15, 29, 43, 48, 107
- Security policy 99
- SELECT INTO TEMP statement 175
- Select privilege 23, 65, 78, 206
- SELECT statements 10, 32
- SELECT triggers 71
- Selectivity constant 53, 53
- Self-join 5
- send() support function 128
- SENDRECQ data type 135
- Sequence
  - syssequences data 62
  - sys synonyms data 63
  - sysstntable data 63
  - sysstabauth data 65
  - sysstables data 65
- Sequential integers
  - am\_id code 15
  - classid code 69
  - constrid code 29
  - extended\_id code 75
  - langid code 58
  - msgid code 69
  - opclassid code 48
  - planid code 56
  - procid code 53, 53
  - seqid code 62
  - SERIAL data type 111
  - SERIAL8 data type 112
  - tabid code 5, 62, 65
- SERIAL data type
  - coltype code 24
  - defined 111
  - inserting values 111
  - length (syscolumns) 27
  - resetting values 111
- SERIAL8 data type
  - assigning a starting value 112
  - coltype code 24
  - defined 112
  - inserting values 112
  - length (syscolumns) 27
  - resetting values 112
  - using with INT8 88
- Serializable transactions 80
- server\_info Information Schema view 77

- SET data type
  - coltype code 24
- SET data type, defined 113
- SET ENVIRONMENT IFX\_AUTO\_REPREPARE statement 65
- SET ENVIRONMENT statement 139, 144, 208
- SET OPTIMIZATION statement 209, 209
- SET PDQPRIORITY statement 210
- SET SESSION AUTHORIZATION statement 53
- SET STMT\_CACHE statement 217, 217
- set utility 145
- setenv utility 142
- Setnet32 146
- Setnet32 utility 144
- Setting environment variables
  - in UNIX 139
  - in Windows 144
- SGML (Standard Graphic Markup Language) 92
- Shared environment-configuration file 143
- Shared libraries 188
- Shared memory
  - ONEDB\_SHMBASE 200
  - PLOAD\_SHMBASE 212
- Shell
  - remote 174
  - search path 210
  - setting environment variables in a file 141
  - specifying with DBREMOTECMD 174
- SHLIB\_PATH environment variable 216
- simple large object
  - defined 89
- Simple large objects
  - defined 121
  - location (sysblobs) 20
- Single-precision floating-point number 108, 114
- SMALLFLOAT data type
  - built-in casts 133, 134
  - coltype code 24
  - defined 114
  - display format 169, 171
- SMALLINT data type
  - built-in casts 133, 134
  - coltype code 24
  - defined 114
  - length (syscolumns) 27
- Smart large objects
  - defined 122
  - syscolattns data 22
- Smart-large-object handles 212
- Solaris operating system 186
- SOME operator 137
- Sort-merge join 208
- Sorting
  - DBSPACETEMP environment variable 175
  - PSORT\_DBTEMP environment variable 214
  - PSORT\_NPROCS environment variable 215
- Space
  - DATETIME delimiter 93
  - INTERVAL delimiter 100
- Spatial queries 216
- SPL routines 53, 130, 165, 167
- SPL variables 130
- SQL (Structured Query Language) 162
- SQL character set 181
- SQL Communications Area 162
- sql\_languages Information Schema view 77
- SQL\_LOGICAL\_CHAR configuration parameter 65, 65, 118
- sqlhosts file 199, 200

SQLHOSTS subkey 200  
 SQLSTATE values 34  
 sqltypes.h file 24  
 SQLWARN array 162  
 SRV\_FET\_BUF\_SIZE environment variable 216  
 Stack size 53, 201  
 STACKSIZE configuration parameter 201  
 Staging-area blobspace 199  
 Standard Graphic Markup Language (SGML) 92  
 START DATABASE statement 172  
 START VIOLATIONS TABLE statement 73  
 STAT data type 32  
 STATCHANGE configuration parameter 32, 37  
 STATCHANGE table attribute 32, 37  
 Statement cache 217  
 Statements of SQL
 

- ALTER INDEX 43
- ALTER OPTICAL CLUSTER 49
- ALTER SEQUENCE 62, 220
- ALTER TABLE 10, 56, 65, 220
- CLOSE 209
- CONNECT 172, 172, 197, 199
- CREATE ACCESS\_METHOD 15
- CREATE AGGREGATE 14
- CREATE CAST 20, 135
- CREATE DATABASE 172
- CREATE DISTINCT TYPE 75, 98, 223
- CREATE EXTERNAL TABLE 35, 36
- CREATE FUNCTION 58, 206
- CREATE IMPLICIT CAST 223
- CREATE INDEX 5, 41, 43, 65, 180, 180
- CREATE OPAQUE TYPE 75, 107
- CREATE OPERATOR CLASS 48
- CREATE OPTICAL CLUSTER 49, 49
- CREATE PROCEDURE 51, 58
- CREATE ROLE 57, 72
- CREATE ROUTINE FROM 58
- CREATE ROW TYPE 75, 108
- CREATE SCHEMA AUTHORIZATION 5
- CREATE SEQUENCE 62
- CREATE SYNONYM 63
- CREATE TABLE 30, 56, 64
- CREATE TRIGGER 71
- CREATE VIEW 72
- CREATE XDATASOURCE 73
- CREATE XDATASOURCETYPE 74
- DATABASE 172
- DECLARE 209
- DELETE 10, 56, 73, 73
- DESCRIBE 193
- DROP CAST 223
- DROP DATABASE 172
- DROP FUNCTION 53
- DROP INDEX 65
- DROP OPTICAL CLUSTER 49
- DROP PROCEDURE 53
- DROP ROUTINE 53
- DROP ROW TYPE 108
- DROP SEQUENCE 220
- DROP TABLE 220
- DROP TYPE 98, 107
- DROP VIEW 77, 220
- FETCH 209
- GET DIAGNOSTICS 34
- GRANT 36, 57, 65, 65, 78
- INSERT 73, 129, 161, 166
- LOAD 89, 162, 162
- MERGE 73
- OPEN 209
- PREPARE 65
- RENAME SEQUENCE 220
- RENAME TABLE 220
- REVOKE 65, 72
- SELECT 10, 32, 56, 175
- SET ENVIRONMENT 208
- SET OPTIMIZATION 209
- SET PDQPRIORITY 210
- SET SESSION AUTHORIZATION 53
- SET STMT\_CACHE 217
- START DATABASE 172
- START VIOLATIONS TABLE 73
- UNLOAD 163
- UPDATE 161
- UPDATE STATISTICS 10, 43, 179
- UPDATE STATISTICS FOR PROCEDURE 56
- UPDATE STATISTICS FOR TABLE 29

 Statements of SQL LOAD 115  
 Statements of SQL UPDATE 115  
 static option of ESQL/C 188  
 STATLEVEL table attribute 37  
 STMT\_CACHE configuration parameter 217  
 STMT\_CACHE environment variable 217  
 STMT\_CACHE keyword 217  
 Storage identifiers 181  
 Stored procedure language (SPL) 53, 130, 165  
 stores\_demo database 220
 

- join columns 220

 strings option of gcc 195  
 Structured Query Language (SQL) 162  
 Subscripts 89  
 Subscripts ([ ]), 115  
 SUBSTRING function 10  
 Subtable 37, 39, 46, 225  
 Subtype 46, 108  
 Summary
 

- of data types 81
- of environment variables, by topic 147
- of system catalog tables 11

 superstores\_demo database
 

- structure of tables 222

 Supertable 46, 225  
 Supertype 46, 108  
 Support functions
 

- DISTINCT data types 131
- OPAQUE data types 107, 128
- routine identifier 53

 Symbol table 53, 53  
 Synonym
 

- sys synonyms data 63
- sys syntax data 63
- systables data 65
- USETABLENAME setting 220

 sysaggregates system catalog table 14  
 sysams system catalog table 15  
 sysattrtypes system catalog table 18  
 sysautolocate system catalog table 19  
 sysblobs system catalog table 20  
 sysbuiltintypes table 5  
 syscasts system catalog table 20, 132  
 syschecks system catalog table 21  
 syscheckudrdep system catalog table 21  
 syscolattns system catalog table 22  
 syscolauth system catalog table 23  
 syscoldepend system catalog table 23  
 syscolumns system catalog table 24  
 sysconstraints system catalog table 29  
 syscrd database 5  
 sysdbclose
 

- disabling with IFX\_NODBPROC 190

 sysdbopen
 

- disabling with IFX\_NODBPROC 190

 sysdefaults system catalog table 30  
 sysdepend system catalog table 31  
 sysdirectives system catalog table 31  
 sysdistrib system catalog table 32  
 sysdomains system catalog view 34  
 syserrors system catalog table 34  
 sysextcols system catalog table 35  
 sysextdfiles system catalog table 35  
 sysexternal system catalog table 36  
 sysfragauth system catalog table 36  
 sysfragdist system catalog table 37  
 sysfragments system catalog table 39  
 sysindexes system catalog table 41  
 sysindexes system catalog tables 43  
 sysinherits system catalog table 46  
 syslangauth system catalog table 46  
 syslogmap system catalog table 47  
 sysmaster database 5
 

- contrasted with system catalog tables 5
- initialization 138

 sysobjstate system catalog table 47  
 sysopclasses system catalog table 48  
 sysopclstr system catalog table 49  
 sysprocauth system catalog table 51  
 sysprocbody system catalog table 51  
 sysproccolumns system catalog table 52  
 sysprocedures system catalog table 53  
 sysproplan system catalog table 56  
 sysreferences system catalog table 56  
 sysroleauth system catalog table 57  
 sysroutinelangs system catalog table 58  
 sysseclabelauth system catalog table 58  
 sysseclabelcomponentelements system catalog table 59  
 sysseclabelcomponents system catalog table 58  
 sysseclabelnames system catalog table 59  
 sysseclabels system catalog table 60  
 syssecpolicies system catalog table 60  
 syssecpolicycomponents system catalog table 61  
 syssecpolicyexemptions system catalog table 61  
 syssequences system catalog table 62  
 sysssurrogateauth system catalog table 62  
 sysssynonyms system catalog table 63  
 sysssyntax system catalog table 63  
 systabamdata system catalog table 64  
 systabauth system catalog table 65  
 systables system catalog table 65  
 System administrator (DBA) 5  
 System applet 145  
 System catalog
 

- access methods 15, 64
- access privileges 23, 36
- accessing 10
- altering contents 10
- casts 20
- columns 24
- complex data types 18, 75
- constraint violations 73
- constraints 21, 23, 29
- data distributions 32
- database tables 65
- default values 30
- defined 4
- dependencies 31
- discretionary access privileges 65
- drvurity policies 60
- example 5
- external directives 31

- external tables 35, 35, 36
- fragment distributions 37
- fragment privileges 36
- fragments 39
- indexes 41, 43
- inheritance 46
- list of tables 11
- messages 34
- operator classes 48
- optical clusters 49
- privileges 72, 75
- programming languages 46, 58
- referential constraints 29, 56, 73
- roles 57
- routine parameters 52
- routines 51, 53, 56
- security label components 58
- sequence objects 62
- simple large objects 20
- smart large objects 22
- synonyms 63
- text of routines 51
- trace classes 69
- trace messages 69
- triggers 70, 71
- updating 10
- use by database server 5
- user-defined aggregates 14
- user-defined data types 74, 75
- views 65, 72
- XA data source types 74
- XA data sources 73

System catalog tables

- synonyms 63
- sysaggregates 14
- sysams 15
- sysattrtypes 18
- sysautolocate 19
- sysblobs 20
- syscasts 20
- syschecks 21
- syscheckudrdep 21
- syscolattrs 22
- syscolauth 23
- syscoldepend 23
- syscolumns 24
- sysconstraints 29
- sysdefaults 30
- sysdepend 31
- sysdirectives 31
- sysdistrib 32
- sysdomains 34
- syserrors 34
- sysextcols 35
- sysextdfiles 35
- sysexternal 36
- sysfragauth 36
- sysfragdist 37
- sysfragments 39
- sysindexes 41
- sysindices 43
- sysinherits 46
- syslangauth 46
- syslogmap 47
- sysobjstate 47
- sysopclasses 48
- sysopclstr 49
- sysprocauth 51
- sysprocbody 51
- sysproccolumns 52
- sysprocedures 53

- sysprocplan 56
- sysreferences 56
- sysroleauth 57
- sysroutinelangs 58
- sysseclabelauth 58
- sysseclabelcomponentelements 59
- sysseclabelcomponents 58
- sysseclabelnames 59
- sysseclabels 60
- syssecpolicies 60
- syssecpolicycomponents 61
- syssecpolicyexemptions 61
- syssequences 62
- sys surrogateauth 62
- sys synonyms 63
- sys syntable 63
- sys tabamdata 64
- sys tabauth 65
- sys tables 65
- sys traceclasses 69
- sys tracemsgs 69
- sys trigbody 70
- sys triggers 71
- sys users 72
- sys views 72
- sys violations 73
- sys xadatasources 73
- sys xasourcetypes 74
- sys xtddesc 74
- sys xtdtypeauth 75
- sys xtdtypes 75

SYSTEM() command, on NT 202

- sys traceclasses system catalog table 69
- sys tracemsgs system catalog table 69
- sys trigbody system catalog table 70
- sys triggers system catalog table 71
- sys users system catalog table 72
- sys utils database 5
- sys uuid database 5
- sys views system catalog table 72
- sys violations system catalog table 73
- sys xadatasources system catalog table 73
- sys xasourcetypes system catalog table 74
- sys xtddesc system catalog table 74
- sys xtdtypeauth system catalog table 75
- sys xtdtypes system catalog table 75, 107, 108

## T

- tabid 5, 65
- Table
  - changing a column data type 132
  - dependencies, in sysdepend 31
  - diagnostic 73
  - extent size 65
  - fragmented 37, 39
  - hashing parameters 64
  - hierarchy 37, 39, 46, 108, 225
  - inheritance, sysinherits data 46
  - lock mode 65, 184
  - nonfragmented 180
  - separate from large object storage 121
  - structure in superstores\_demo database 222
  - synonyms in sys syntable 63
  - sys tables data 65
  - system catalog tables 14
  - temporary 175, 176
  - temporary in SE 176
  - untyped, and unnamed ROW 110
  - version value 65
  - violations 73
- Table-based fragmentation 39
- Table-level privileges
  - PUBLIC 78
  - sys fragauth data 36
  - sys tabauth data 5, 65
- tables Information Schema view 77
- Tape management
  - setting DBREMOTECMD 174
- Temporary dbspace 175
- Temporary files 176
  - in SE, specifying directory with DBTEMP 176
  - setting DBSPACETEMP 175
  - setting PSORT\_DBTEMP 214
- Temporary tables 175
  - in SE, specifying directory with DBTEMP 176
  - specifying dbspace with DBSPACETEMP 175
- TERM environment variable 218
- TERMCAP environment variable 218
- termcap file
  - setting ONEDB\_TERM 201
  - setting TERMCAP 218
- Terminal handling
  - setting ONEDB\_TERM 201
  - setting TERM 218
  - setting TERMCAP 218
  - setting TERMINFO 219
- terminfo directory 201, 219
- TERMINFO environment variable 219
- TEXT 115
- TEXT argument 115
- TEXT Character string TEXT 115
- TEXT data type 115, 115
  - coltype code 24
  - increasing buffer size 163
  - length (syscolumns) 27
  - nonprintable characters 116
  - setting buffer size 163
  - sysblobs data 20
  - sysfragments data 39
  - with control characters 116
- TEXT data type IS NULL 115
- TEXT data type restrictions 115
- Text editor 168
- thousands separator 190
- Thousands separator 105
- thread flag of ESQL/C 219
- THREADLIB environment variable 219
- Time data types
  - arithmetic 123
  - length (syscolumns) 27
  - listed 118
- Time values
  - DBCENTURY setting 163
  - DBDATE setting 166
  - DBTIME setting 176
  - GL\_DATETIME settings 176
  - USEOSTIME configuration parameter 93
- Time-limited licenses
  - (IFX\_NO\_TIMELIMIT\_WARNING) 190
- Timezone
  - setting TZ 219
- TO keyword
  - DATETIME qualifier 93
  - EXTEND function 125
  - INTERVAL qualifier 100
- TODAY operator 30
- Trace class 69
- Trace messages 69

- Trace statements 69
- Transaction isolation level 80, 208
- Transaction logging 22, 80
- Trigger routines 53
- Triggers
  - creation-time value 165, 167
  - sysobjstate data 47
  - systrigbody data 70, 70
  - systriggers data 71
- TRUE setting
  - BOOLEAN values 89
  - sysams table 15, 15, 15, 15, 15
- Truncation 90
- TYPE keyword 109
- TZ environment variable 219

## U

- UDT indexes 216
- Unary arithmetic operators 137
- Uncommitted read 80
- Under privilege 65
- Unique constraint 73, 111, 112
- Unique index 41, 111
- Unique keys 15
- Unique numeric values
  - SERIAL data type 111
  - SERIAL8 data type 112
- UNITS operator 93, 123, 126, 137
- UNIX
  - BSD, default print utility 174
  - environment variables 138
  - PATH environment variable 210
  - System V
    - default print utility 174
    - terminfo libraries 201, 219
  - temporary files 214
  - TERM environment variable 218
  - TERMCAP environment variable 218
  - TERMINFO environment variable 219
- UNLOAD statement 163, 168
- Unnamed ROW data type
  - coltype code 24
  - declaring 110
  - defined 109
  - inserting values 110
- unset utility 142
- unsetenv utility 142
- Unsetting an environment variable 142
- Untyped table 65
- Update privilege 23, 36, 65, 206
- UPDATE statement 73
- UPDATE statements 193
- UPDATE STATISTICS FOR PROCEDURE statement 56
- UPDATE STATISTICS statement 43, 179
  - and DBUPSPACE environment variable 179
  - effect on sysdistrib table 32
  - sysindices (index statistics) 49
  - sysindices data 43
  - updating system catalog tables 10
- Update trigger 71
- Uppercase mode codes 53
- Uppercase privilege codes 5, 23, 65
- USE\_DTEENV environment variable 93, 93
- USEOSTIME configuration parameter 93
- User environment variable 146
- User informix 10, 20, 133
- User name 80
- User privileges
  - syscolauth data 23
  - sysfragauth data 36

- syslangauth data 46
- sysprocauth data 51
- sysstabauth data 65
- sysusers data 72
- sysxdttypeauth data 75
- User-defined aggregates 14
- User-defined casts 135
- User-defined casts (UDCs) 20
- User-defined data types
  - casting 135
  - casting into built-in type 132
  - opaque 131
  - sysxtddesc data 74
  - sysxtdtype data 75, 75
- User-defined routines
  - casts (syscasts) 20
  - check constraints (syscheckudrdep) 21
  - error messages (syserrors) 34
  - for OPAQUE data types 107
  - functional index 180
  - language authorization (syslangauth) 46
  - privileges 51, 206
  - protected 53
  - secondary access method 29
  - sysprocedures data 53
- USETABLENAME environment variable 220
- Utilities
  - chkenv 140, 143
  - DB-Access 10, 77, 144, 162, 169, 199
  - dbload 88, 89
  - dbschema 53, 53
  - env 142
  - export 141
  - gcc 195
  - getenv 139
  - ifx\_getenv 144
  - ifx\_putenv 144
  - lp 174
  - lpr 174
  - oninit 184
  - onload 89
  - onpload 171, 212
  - onsecurity 189
  - printenv 142
  - putenv 139
  - set 145
  - setenv 142
  - Setnet32 144
  - source 140
  - unset 142
  - unsetenv 142, 181
  - vi 168
- Utilities dbload 115

## V

- VARCHAR data type
  - ([]), brackets
    - MATCHES range delimiters 116
  - CHAR data type
    - collation 116
  - Code sets
    - collation order 116
    - East Asian 116
  - Collation
    - VARCHAR data type 116
  - coltype code 24
  - defined 116
  - Locales
    - collation order 116
  - MATCHES operator 116
  - Multibyte characters

- VARCHAR data type 116
  - nonprintable characters 116
  - SQL\_LOGICAL\_CHAR configuration parameter 116
  - storing numeric values 116
  - VARCHAR data type
    - collation 116
    - multibyte characters 116
  - Zero (0)
    - C null as terminator 116
- Variable-length opaque data types 24
- Variable-length packets 191
- Variable-length UDT 75
- VARIANT routine 53
- Version of a table 65
- vi text editor 168
- View
  - columns view 79
  - Information Schema 77
  - server\_info view 80
  - sql\_languages view 80
  - sysdepend data 31
  - sysindexes view 43
  - sys synonyms data 63
  - sysstable data 63
  - sysstabauth data 65
  - sysables data 65
  - sysviews data 72
  - tables view 78
- Violations
  - sysobjstate data 47
  - sysviolations data 73
- Virtual machine 157, 205
- Virtual processors 216

## W

- Warning message 34, 162
- WHERE 115
- WHERE keyword 10, 21
- Whitespace in identifiers 181
- Window borders 201
- Windows environments
  - manipulating environment variables 144
  - setting environment variables 144

## X

- X setting
  - sysams.am\_sptype 15
  - sysstabauth.tabauth 65
- X/Open
  - compliance 80
  - server\_info view 80
- X/Open CAE standards 77
- XA data source types 74
- XA data sources 73
- XML (Extensible Markup Language) 92
- XPG4 standard 79, 79

## Y

- Y setting
  - DBDATE 166
  - DBTIME 176
  - sysroleauth.is\_grantable 57
- Year 2000 163
- YEAR keyword
  - DATETIME qualifier 93
  - EXTEND function 125
  - INTERVAL qualifier 100
- Year values, two and four digit 93, 163, 166, 176
- yes setting
  - NODEFDAC 206

YES setting  
columns.is\_nullable 79  
sql\_languages.integrity 80

## Z

Zero  
extent size encoding 43  
Zero (0)  
DBDATE separator 166  
DECIMAL scale 96  
hexadecimal digit 168  
IFX\_DIRECTIVES setting 185, 185  
IFX\_LARGE\_PAGES setting 186  
IFX\_LONGID setting 188  
IFX\_NETBUF\_PVTPOOL\_SIZE setting 188  
integer scale 79, 96  
IONEDB\_OPCACHE setting 199  
OPTCOMPIND setting 208  
OPTMSG setting 208  
padding of 1-digit years 163  
padding with DBFLTMASK 169  
padding with DBTIME 176  
PDQPRIORITY setting 210  
PSORT\_NPROCS setting 216  
STMT\_CACHE setting 217  
sysams values 15, 15, 15, 15, 15  
sysfragments.hybdpos 39  
sysindices.nrows 43  
systables.type\_xid 65  
sysxdtypes values 75