

HCL OneDB 2.0.1

HCL OneDB Data Warehouse Guide



Contents

Chapter 1. Dimensional databases.....	3
Dimensional databases.....	3
Overview of data warehousing.....	3
Dimensional databases.....	
What is dimensional data?.....	6
Design a dimensional data model.....	8
Concepts of dimensional data modeling.....	8
Building a dimensional data model.....	12
Handle common dimensional data-modeling problems.....	30
Implement a dimensional database.....	33
Implement the sales_demo dimensional database.....	34
Moving data from relational tables into dimensional tables by using external tables.....	43
Performance tuning dimensional databases.....	44
Query execution plans.....	44
Data distribution statistics.....	45
Index.....	52

Chapter 1. Dimensional databases

The *HCL OneDB™ Data Warehouse Guide* provides information to help you design, implement, and manage dimensional databases, and describes the tools that you can use to create data warehouses and optimize your data warehouse queries.

These topics are of interest to the following users:

- Database administrators
- System administrators
- Performance engineers

These topics are written with the assumption 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 dimensional databases, relational databases, or exposure to database concepts
- Some experience with database server administration, operating-system administration, or network administration

Dimensional databases

A dimensional database is a relational database that uses a *dimensional data model* to organize data. This model uses fact tables and dimension tables in a star or snowflake schema.

A dimensional database is the optimal type of database for data warehousing.

The availability and reliability of the HCL OneDB™ database server includes a full active-active cluster solution for high availability and low cost scalability. You can use HCL OneDB™ to manage workload distribution across multiple read-only or full-transaction nodes. You can dynamically add different types of nodes into your cluster environment to scale out or increase availability in the most demanding environments.

Warehouse workloads have the flexibility to work on the same database with operational data, running real-time on a separate node in the cluster. Data can also be replicated in real-time using Enterprise Replication, or copied to a separate data warehouse server. With HCL OneDB™, you have the flexibility to design the system to meet your needs and to make the most of your existing infrastructure.

Overview of data warehousing

Data warehouse databases provide a decision support system (DSS) environment in which you can evaluate the performance of an entire enterprise over time.

In the broadest sense, the term *data warehouse* is used to refer to a database that contains very large stores of historical data. The data is stored as a series of snapshots, in which each record represents data at a specific time. By analyzing these snapshots you can make comparisons between different time periods. You can then use these comparisons to help make important business decisions.

Data warehouse databases are optimized for data retrieval. The duplication or grouping of data, referred to as *database denormalization*, increases query performance and is a natural outcome of the dimensional design of the data warehouse.

By contrast, traditional online transaction processing (OLTP) databases automate day-to-day transactional operations. OLTP databases are optimized for data storage and strive to eliminate data duplication. Databases that achieve this goal are referred to as *normalized* databases.

An enterprise data warehouse (EDW) is a data warehouse that services the entire enterprise. An *enterprise data warehousing environment* can consist of an EDW, an operational data store (ODS), and physical and virtual data marts.

A data warehouse can be implemented in several different ways. You can use a single data management system, such as HCL OneDB™, for both transaction processing and business analytics. Or, depending on your system workload requirements, you can build a data warehousing environment that is separate from your transactional processing environment.

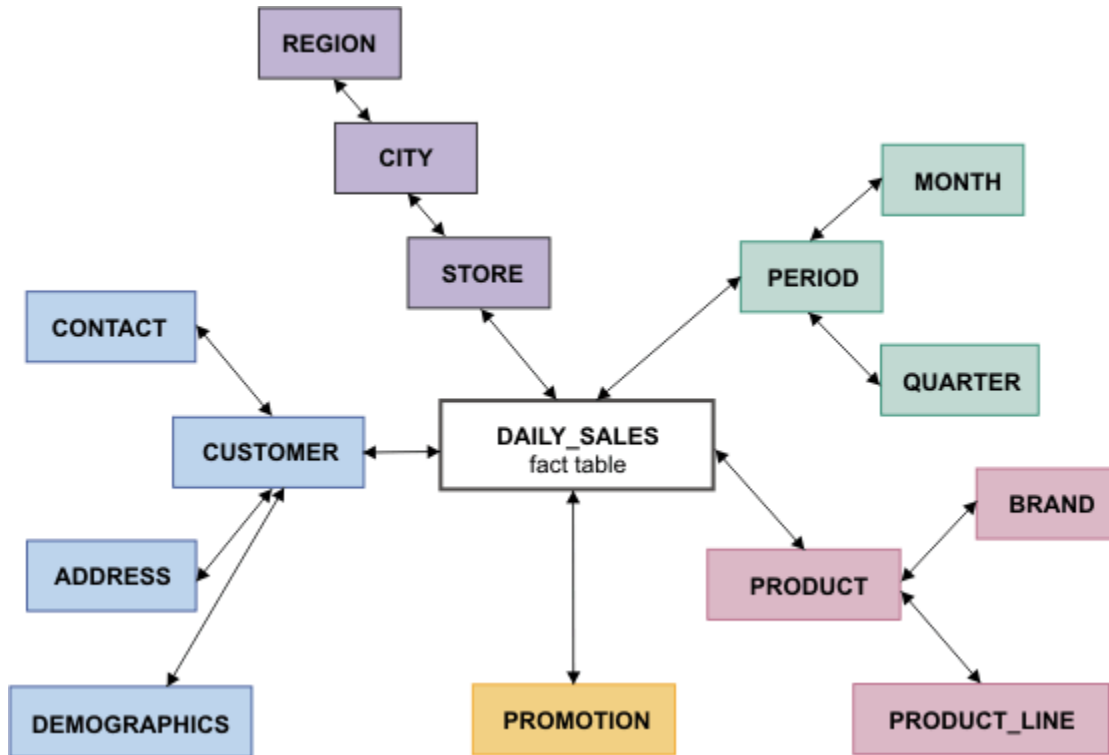
HCL OneDB™ uses the umbrella terms *data warehousing* and *data warehousing environment* to encompass any of the following forms that you might use to store your data:

Data warehouse

A database that is optimized for data retrieval to facilitate reporting and analysis. A data warehouse incorporates information about many subject areas, often the entire enterprise. Typically you use a dimensional data model to design a data warehouse. The data is organized into dimension tables and fact tables using star and snowflake schemas. The data is denormalized to improve query performance. The design of a data warehouse often starts from an analysis of what data already exists and how to collect it in such a way that the data can later be used. Instead of loading transactional data directly into a warehouse, the data is often integrated and transformed before it is loaded into the warehouse.

The primary advantage of a data warehouse is that it provides easy access to and analysis of vast stores of information on many subject areas.


Figure 1. A sample snowflake schema which has the DAILY_SALES table as the fact table.



Data mart

A database that is oriented towards one or more specific subject areas of a business, such as tracking inventories or transactions, rather than an entire enterprise. A data mart is used by individual departments or groups. Like a data warehouse, you typically use a dimensional data model to build a data mart. For example the data mart might use a single star schema comprised of one fact table and several dimension tables. The design of a data mart often starts with an analysis of what data the user needs rather than focusing on the data that already exists.

Figure 2. A data mart with the DAILY_SALES fact table



Operational data store

A subject-oriented system that is optimized for looking up one or two records at a time for decision making. An operational data store (ODS) is a hybrid form of data warehouse that contains timely, current, integrated information. Including the ODS in the data warehousing environment enables access to more current data more quickly, particularly if the data warehouse is updated by one or more batch processes rather than updated continuously. The data typically is of a higher level granularity than the transaction. You can use an ODS for clerical, day-to-day decision making. This data can serve as the common source of data for data warehouses.

What is dimensional data?

Traditional relational databases, such as OLTP databases, are organized around a list of records. Each record contains related information that is organized into attributes (fields). The **customer** table of the **stores_demo** demonstration database, which includes fields for name, company, address, phone, and so forth, is a typical example. While this table has several

fields of information, each row in the table pertains to only one customer. If you wanted to create a two-dimensional matrix with customer name and any other field, for example, phone number), you would realize that there is only a one-to-one correspondence. The following table is an example of a database table with fields that have only a one-to-one correspondence.

Table 1. A table with a one-to-one correspondences between fields

Customer	Phone number --->		
Ludwig Pauli	408-789-8075	-----	-----
Carole Sadler	-----	415-822-1289	-----
Philip Currie	-----	-----	414-328-4543

You could put any combination of fields from the preceding **customer** table in this matrix, but you would always end up with a one-to-one correspondence, which shows that this table is not multidimensional and would not be well suited for a dimensional database.

However, consider a relational table that contains more than a one-to-one correspondence between the fields of the table. Suppose you create a table that contains sales data for products sold in each region of the country. For simplicity, the company has three products that are sold in three regions. The following table shows how you might store this data in a table, using a normalized data model. This table lends itself to multidimensional representation because it has more than one product per region and more than one region per product.

Table 2. A simple table with a many-to-many correspondence

Product	Region	Unit Sales
Football	East	2300
Football	West	4000
Football	Central	5600
Tennis racket	East	5500
Tennis racket	West	8000
Tennis racket	Central	2300
Baseball	East	10000
Baseball	West	22000
Baseball	Central	34000

Although this data can be forced into the three-field relational table, the data fits more naturally into the two-dimensional matrix in the following table. This matrix better represents the many-to-many relationship of product and region data shown in the previous table.

Table 3. A simple two-dimensional example

	Region	Central	East	West
Product	Football	5600	2300	4000
	Tennis Racket	2300	5500	8000
	Baseball	34000	10000	22000

The performance advantages of the dimensional model over the normalized model can be great. A dimensional approach simplifies access to the data that you want to summarize or compare. For example, using the dimensional model to query the number of products sold in the West, the database server finds the **West** column and calculates the total for all row values in that column. To perform the same query on the normalized table, the database server has to search and retrieve each row where the **Region** column equals 'West' and then aggregate the data. In queries of this kind, the dimensional table can total all values of the **West** column in a fraction of the time it takes the relational table to find all the 'West' records.

Design a dimensional data model

To build a dimensional database, you start by designing a dimensional data model for your business.

You will learn how a dimensional model differs from a transactional model, what fact tables and dimension tables are and how to design them effectively. You will learn how to analyze the business processes in your organization where data is gathered and use that analysis to design a model for your dimensional data.

HCL OneDB™ includes several demonstration databases that are the basis for many examples in HCL OneDB™ publications, including examples in the *HCL OneDB™ Data Warehouse Guide*. The **stores_demo** database illustrates a relational schema with information about a fictitious wholesale sporting-goods distributor. You will use SQL and the data in the **stores_demo** database to populate a new dimensional database. The dimensional database is based on the simple dimensional data model that you learned about.

To understand the concepts of dimensional data modeling, you should have a basic understanding of SQL and relational database theory. This section provides only a summary of data warehousing concepts and describes a simple dimensional data model.


Concepts of dimensional data modeling

To build a dimensional database, you start with a dimensional data model. The dimensional data model provides a method for making databases simple and understandable. You can conceive of a dimensional database as a database *cube* of three

or four dimensions where users can access a slice of the database along any of its dimensions. To create a dimensional database, you need a model that lets you visualize the data.

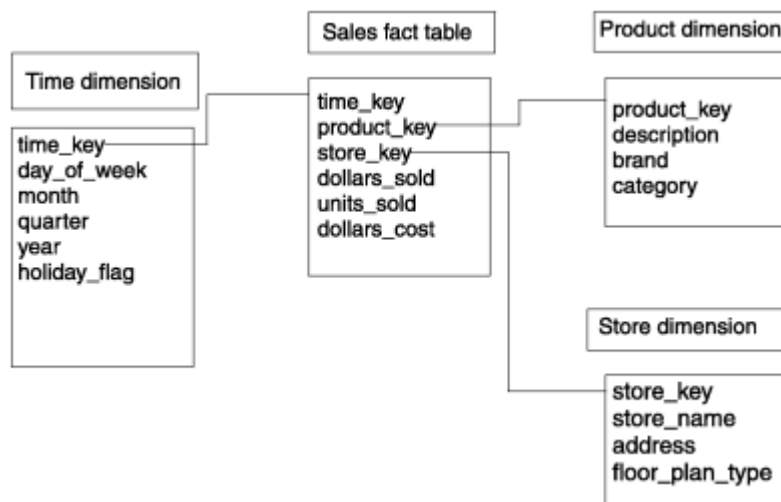
Suppose your business sells products in different markets and you want to evaluate the performance over time. It is easy to conceive of this business process as a cube of data, which contains dimensions for time, products, and markets. The following figure shows this dimensional model. The various intersections along the lines of the cube would contain the *measures* of the business. The measures correspond to a particular combination: product, market, and time data.

Figure 3. A dimensional model of a business that has time, product, and market dimensions



Another name for the dimensional model is the *star schema*. The database designers use this name because the diagram for this model looks like a star with one central table around which a set of other tables are displayed. The central table is the only table in the schema with multiple joins connecting it to all the other tables. This central table is called the *fact table* and the other tables are called *dimension tables*. The dimension tables all have only a single join that attaches them to the fact table, regardless of the query. The following figure shows a simple dimensional model of a business that sells products in different markets and evaluates business performance over time.

Figure 4. A typical dimensional model



The fact table

The fact table stores the measures of the business and points to the key value at the lowest level of each dimension table. The *measures* are quantitative or factual data about the subject.

The measures are generally numeric and correspond to the "how much" or "how many" aspects of a question. Examples of measures are price, product sales, product inventory, revenue, and so forth. A measure can be based on a column in a table or it can be calculated.

The following table shows a fact table whose measures are sums of the units sold, the revenue, and the profit for the sales of that product to that account on that day.

Table 4. A fact table with sample records

Product Code	Account code	Day code	Units sold	Revenue	Profit
1	5	32104	1	82.12	27.12
3	17	33111	2	171.12	66.00
1	13	32567	1	82.12	27.12

Before you design a fact table, you must determine the *granularity* of the fact table. The granularity corresponds to how you define an individual low-level record in that fact table. The granularity might be the individual transaction, a daily snapshot, or a monthly snapshot. The fact table shown contains one row for every product sold to each account each day. Thus, the granularity of the fact table is expressed as *product by account by day*.

Dimensions of the data model

A *dimension* represents a single set of objects or events in the real world. Each dimension that you identify for the data model gets implemented as a dimension table. Dimensions are the qualifiers that make the measures of the fact table meaningful, because they answer the what, when, and where aspects of a question. For example, consider the following business questions, for which the dimensions are italicized:

- What *accounts* produced the highest revenue last *year*?
- What was our profit by *vendor*?
- How many units were sold for each *product*?


In the preceding set of questions, revenue, profit, and units sold are measures (not dimensions), as each represents quantitative or factual data.

Dimension elements

A dimension can define multiple *dimension elements* for different levels of summation.

For example, all the elements that relate to the structure of a sales organization might comprise one dimension. The following figure shows the dimension elements that the **Accounts** dimension defines.

Figure 5. Dimension elements in the accounts dimension



Dimensions are made up of hierarchies of related elements. Because of the hierarchical aspect of dimensions, users are able to construct queries that access data at a higher level (*roll up*) or lower level (*drill down*) than the previous level of detail. The figure shows the hierarchical relationships of the dimension elements:

- The account elements roll up to the territory elements
- The territory elements roll up to the region elements

Users can query at different levels of the dimension, depending on the data they want to retrieve. For example, users might perform a query against all regions and then drill down to the territory or account level for detailed information.

Dimension elements are usually stored in the database as numeric codes or short character strings to facilitate joins to other tables.

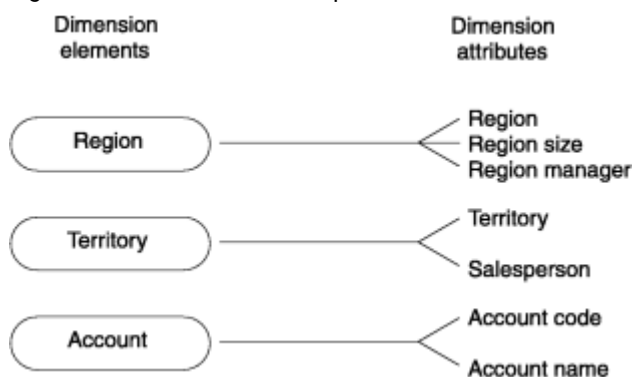
Each dimension element can define multiple dimension attributes, in the same way dimensions can define multiple dimension elements.

Dimension attributes

A *dimension attribute* is a column in a dimension table. Each attribute describes a level of summary within a dimension hierarchy.

The dimension elements define the hierarchical relationships within a dimension table. The dimension attributes describe the dimension elements in terms that are familiar to users. The following figure shows the dimension elements and corresponding attributes of the **Account** dimension.

Figure 6. Attributes that correspond to the dimension elements



Because dimension attributes describe the items in a dimension, they are most useful when they are text.



Tip: Sometimes during the design process, it is unclear whether a numeric data field from a production data source is a measured fact or an attribute. Generally, if the numeric data field is a measurement that changes each time you sample it, the field is a fact. If field is a discretely valued description of something that is more or less constant, it is a dimension attribute.

Dimension tables

A *dimension table* is a table that stores the textual descriptions of the dimensions of the business. A dimension table contains an element and an attribute, if appropriate, for each level in the hierarchy.

The lowest level of detail that is required for data analysis determines the lowest level in the hierarchy. Levels higher than this base level store redundant data. This denormalized table reduces the number of joins that are required for a query and makes it easier for users to query at higher levels and then drill down to lower levels of detail. The term *drilling down* means to add row headers from the dimension tables to your query. The following table shows an example of a dimension table that is based on the **Account** dimension.

Table 5. An example of a dimension table

Acct code	Account name	Territory	Salesman	Region	Region size	Region manager
1	Javier's Mfg.	101	B. Gupta	Asia-Pacific	Over 50	T. Sent
2	TBD Sales	101	B. Gupta	Asia-Pacific	Over 50	T. Sent
3	Tariq's Wares	101	B. Gupta	Asia-Pacific	Over 50	T. Sent
4	The Golf Co.	201	S. Chiba	Asia-Pacific	Over 50	T. Sent

Building a dimensional data model

About this task

To build a dimensional data model, you need a methodology that outlines the decisions you need to make to complete the database design. This methodology uses a top-down approach because it first identifies the major processes in your organization where data is collected. An important task of the database designer is to start with the existing sources of data that your organization uses. After the processes are identified, one or more fact tables are built from each business process. The following steps describe the methodology you use to build the data model.

A dimensional database can be based on multiple business processes and can contain many fact tables. However, to focus on the concepts, the data model that this section describes is based on a single business process and has one fact table.

To build a dimensional database:

1. Choose the business processes that you want to use to analyze the subject area to be modeled.
2. Determine the granularity of the fact tables.
3. Identify dimensions and hierarchies for each fact table.
4. Identify measures for the fact tables.
5. Determine the attributes for each dimension table.
6. Get users to verify the data model.

A business process

A *business process* is an important operation in your organization that some legacy system supports. You collect data from this system to use in your dimensional database.

The business process identifies what end users are doing with their data, where the data comes from, and how to transform that data to make it meaningful. The information can come from many sources, including finance, sales analysis, market analysis, customer profiles. The following list shows different business processes you might use to determine what data to include in your dimensional database:

- Sales
- Shipments
- Inventory
- Orders
- Invoices

Summary of a business process

Suppose your organization wants to analyze customer buying trends by product line and region so that you can develop more effective marketing strategies. In this scenario, the subject area for your data model is **sales**.

After many interviews and thorough analysis of your sales business process, your organization collects the following information:

- Customer-base information has changed.

Previously, sales districts were divided by city. Now the customer base corresponds to two regions: Region 1 for California and Region 2 for all other states.

- The following reports are most critical to marketing:
 - Monthly revenue, cost, net profit by product line from each vendor
 - Revenue and units sold by product, by region, and by month
 - Monthly customer revenue
 - Quarterly revenue from each vendor
- Most sales analysis is based on monthly results, but you can choose to analyze sales by week or accounting period (at a later date).
- A data-entry system exists in a relational database.

To develop a working data model, you can assume that the relational database of sales information has the following properties:

- The **stores_demo** database provides much of the revenue data that the marketing department uses.
- The product code that analysts use is stored in the **catalog** table by the catalog number.
- The product line code is stored in the **stock** table by the stock number. The product line name is stored as description.
- The product hierarchies are somewhat complicated. Each product line has many products, and each manufacturer has many products.
- All the cost data for each product is stored in a flat file named **costs.lst** on a different purchasing system.
- Customer data is stored in the **stores_demo** database.

The region information has not yet been added to the database.

An important characteristic of the dimensional model is that it uses business labels familiar to end users rather than internal tables or column names. After the analysis of the business process is completed, you should have all the information you need to create the measures, dimensions, and relationships for the dimensional data model. This dimensional data model is used to implement the **sales_demo** database that the section [Implement a dimensional database on page 33](#) describes.

The **stores_demo** demonstration database is the primary data source for the dimensional data model that this section builds. For detailed information about the data sources that are used to populate the tables of the **sales_demo** database, see [Mapping data from data sources to the database on page 36](#).

Determine the granularity of the fact table

After you gather all the relevant information about the subject area, the next step in the design process is to determine the granularity of the fact table.

To do this you must decide what an individual low-level record in the fact table should contain. The components that make up the granularity of the fact table correspond directly with the dimensions of the data model. Therefore, when you define the granularity of the fact table, you identify the dimensions of the data model.

How granularity affects the size of the database

The granularity of the fact table also determines how much storage space the database requires.

For example, consider the following possible granularities for a fact table:

- Product by day by region
- Product by month by region

The size of a database that has a granularity of product by day by region would be much greater than a database with a granularity of product by month by region. The database contains records for every transaction made each day as opposed to a monthly summation of the transactions. You must carefully determine the granularity of your fact table because too fine a granularity could result in an astronomically large database. Conversely, too coarse a granularity could mean the data is not detailed enough for users to perform meaningful queries against the database.

Use the business process to determine the granularity

A careful review of the information gathered from the business process should provide what you need to determine the granularity of the fact table. To summarize, your organization wants to analyze customer-buying trends by product line and region so that you can develop more effective marketing strategies.

Customer by product level granularity

The granularity of the fact table should always represent the lowest level for each corresponding dimension.

When you review the information from the business process, the granularity for customer and product dimensions of the fact table are apparent. Customer and product cannot be reasonably reduced any further. These dimensions already express the lowest level of an individual record for the fact table. In some cases, product might be further reduced to the level of product component because a product could be made up of multiple components.

Customer by product by district level granularity

Because the customer buying trends that your organization wants to analyze include a geographical component, you still need to decide the lowest level for the region information.

The business process indicates that in the past, sales districts were divided by city, but now your organization distinguishes between two regions for the customer base:

- Region 1 for California
- Region 2 for all other states

Nonetheless, at the lowest level, your organization still includes sales district data. The district represents the lowest level for geographical information and provides a third component to further define the granularity of the fact table.

Customer by product by district by day level granularity

Customer-buying trends always occur over time, so the granularity of the fact table must include a time component.

Suppose your organization decides to create reports by week, accounting period, month, quarter, or year. At the lowest level, you probably want to choose a base granularity of day. This granularity allows your business to compare sales on Tuesdays with sales on Fridays, compare sales for the first day of each month, and so forth. The granularity of the fact table is now complete.


The decision to choose a granularity of day means that each record in the **time** dimension table represents a day. In terms of the storage requirements, even 10 years of daily data is only about 3,650 records, which is a relatively small dimension table.

Identify the dimensions and hierarchies

After you determine the granularity of the fact table, it is easy to identify the primary dimensions for the data model because each component that defines the granularity corresponds to a dimension.

The following figure shows the relationship between the granularity of the fact table and the dimensions of the data model.

Figure 7. The granularity of the fact table corresponds to the dimensions of the data model

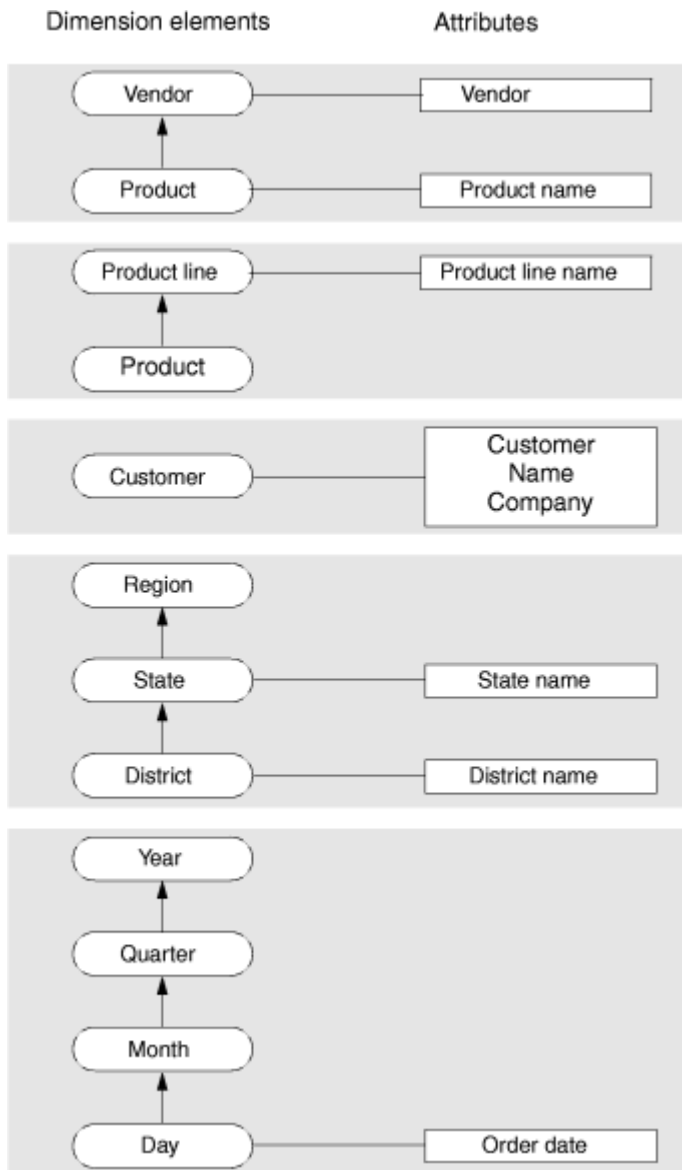


With the dimensions (customer, product, geography, time) for the data model in place, the schema diagram begins to take shape.

i Tip: At this point, you can add additional dimensions to the primary granularity of the fact table, where the new dimensions take on only a single value under each combination of the primary dimensions. If you see that an additional dimension violates the granularity because it causes additional records to be generated, then you must revise the granularity of the fact table to accommodate the additional dimension. For this data model, no additional dimensions need to be added.

You can now map out dimension elements and hierarchies for each dimension. The following figure shows the relationships among dimensions, dimension elements, and the inherent hierarchies.

Figure 8. The relationships between dimensions, dimension elements, and the inherent hierarchies



In most cases, the dimension elements need to express the lowest possible granularity for each dimension, not because queries need to access individual low-level records, but because queries need to cut through the database in precise ways. In other words, even though the questions that a data warehousing environment poses are usually broad, these questions still depend on the lowest level of product detail.

Product dimension

The dimension elements for the **product** dimension are product, product line, and vendor:

- Product has a roll-up hierarchical relationship with product line and with vendor. Product has an attribute of product name.
- Product line has an attribute of product line name.
- Vendor has an attribute of vendor.

Customer dimension

The dimension element for the customer dimension is customer, which has attributes of customer, name, and company.

Geography dimension

The dimension elements for the geography dimension are district, state, and region:

- District has a roll-up hierarchical relationship with state, which has a roll-up hierarchical relationship with region.
- District has an attribute of district name.
- State has an attribute of state name.

Time dimension

The dimensional elements for the time dimension are day, month, quarter, and year.

- Day has a roll-up hierarchical relationship with month, which has a roll-up hierarchical relationship with quarter, which has a roll-up hierarchical relationship with year.
- Day has an attribute of order date.

Establish referential relationships

For the database server to support the dimensional data model, you must define logical dependencies between the fact tables and their dimension tables.

These logical dependencies should be reflected in the columns and indexes that you include in the schema of each table, and in the referential constraints that you define between each fact table and the associated dimension tables. For the large fragmented tables in typical data warehousing operations, these logical dependencies can be the basis for:

- Fragment-key expressions
- Join conditions
- Query predicates for fragment elimination

These query components can significantly improve the performance and throughput of the data warehouse.

A referential constraint enforces a one-to-one relationship between the values in referencing columns (of the foreign key) and the referenced columns (of the primary key or unique constraint). The relationship between the referenced table with the primary key constraint and the referencing table with the foreign key constraint is sometimes called a *parent-child relationship*. The corresponding columns of the parent and child tables can have the same identifiers, but having the same identifiers is not a requirement. There can also be a many-to-one relationship between the referencing table (with the foreign key) and the referenced table (with the primary key, or with the unique constraint).

In the dimensional model, a primary key constraint or a unique constraint in the fact table corresponds to a foreign key constraint in the dimension table. These constraints are specified in the CREATE TABLE or ALTER TABLE statements of SQL that defines the schema of the tables. Because the tables in the primary key and foreign key constraints must be in the same database, the database schema must include the dimension tables of each fact table.

The same data values can appear in the constrained columns of both tables. As a result, the index on which these referential constraints are defined can be used in queries as join predicates to join the fact table and the dimensional table.

For tables that are fragmented by expression or fragmented by list, you can use the foreign key as the fragmentation key for the dimension tables. If you use the foreign key as the fragmentation key, you can use the equality operator or MATCHES operator with the primary key and foreign key values as the join predicate in queries and other data manipulation operations. The join predicate will be `TRUE` for only a subset of the fact table fragments. As a result, the query optimizer can use fragment elimination to process only the fact table partitions that contain qualifying rows.

Resisting normalization

Efforts to normalize a dimensional database can actually prohibit an efficient dimensional design.

If the four foreign keys of the fact table are tightly administered consecutive integers, you could reserve as little as 16 bytes for all four keys (4 bytes each for time, product, customer, and geography) of the fact table. If the four measures in the fact table were each 4-byte integer columns, you would need to reserve only another 16 bytes. Thus, each record of the fact table would be only 32 bytes. Even a billion-row fact table would require only about 32 gigabytes of primary data space.

With its compact keys and data, such a storage-lean fact table is typical for dimensional databases. The fact table in a dimensional model is by nature highly normalized. You cannot further normalize the extremely complex many-to-many relationships among the four keys in the fact table because no correlation exists between the four dimension tables. Virtually every product is sold every day to all customers in every region.

The fact table is the largest table in a dimensional database. Because the dimension tables are usually much smaller than the fact table, you can ignore the dimension tables when you calculate the disk space for your database. Efforts to normalize any of the tables in a dimensional database solely to save disk space are pointless. Furthermore, normalized dimension tables undermine the ability of users to explore a single dimension table to set constraints and choose useful row headers.

Choose the attributes for the dimension tables

After you complete the fact table, you can decide the dimension attributes for each of the dimension tables. To illustrate how to choose the attributes, consider the **time** dimension. The data model for the sales business process defines a granularity of day that corresponds to the time dimension, so that each record in the **time** dimension table represents a day. Keep in mind that each field of the table is defined by the particular day the record represents.

The analysis of the sales business process also indicates that the marketing department needs monthly, quarterly, and annual reports, so the time dimension includes the elements: day, month, quarter, and year. Each element is assigned an attribute that describes the element and a code attribute, to avoid column values that contain long character strings. The following table shows the attributes for the **time** dimension table and sample values for each field of the table.

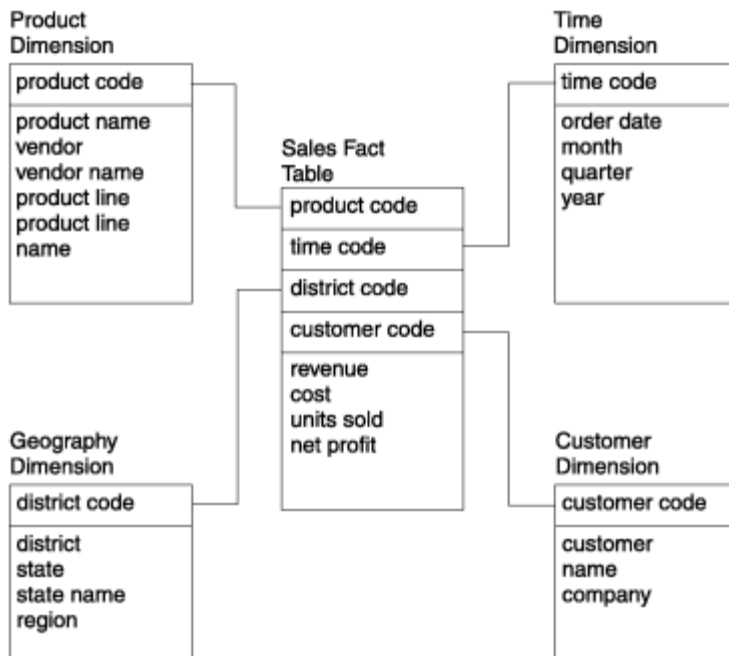
Table 6. Attributes for the time dimension

time code	order date	month code	month	quarter code	quarter	year
35276	07/31/2010	7	july	3	third q	2010
35277	08/01/2010	8	aug	3	third q	2010
35278	08/02/2010	8	aug	3	third q	2010

The previous table shows that the attribute names you assign should be familiar business terms that make it easy for end users to form queries on the database.

The following figure shows the completed data model for the sales business process with all the attributes defined for each dimension table. The elements of the Sales fact table are: product code, time code, district code, customer code, revenue, cost, units sold, and net profit. Some of these elements join the Sales fact table to the dimension tables. Additional elements for each dimension table have been identified.

Figure 9. The completed dimensional data model for the sales business process



Product dimension table

The product code element joins the Sales fact table to the Product dimension table. The additional elements in the Product dimension table are: product name, vendor, vendor name, product line, and product line name.

Time dimension table

The time code element joins the Sales fact table to the Time dimension table. The additional elements in the Time dimension table are: order date, month, quarter, and year.

Geography dimension table

The district code element joins the Sales fact table to the Geography dimension table. The additional elements in the Geography dimension table are: district, state, state name, and region.

Customer dimension table

The customer code element joins Sales fact table to Customer dimension table. The additional elements in the Customer dimension table are: customer name and company.

Fragmentation: Storage distribution strategies

The performance of data warehousing applications can typically benefit from distributed storage allocation designs for partitioning a database table into two or more fragments. Each fragment has the same schema as the table, and stores a subset of the rows in the table (rather than a subset of its columns).

The fragments of a table can be stored in dbspaces on different devices, or in dbspaces on the same physical storage device. The fragments can also be stored in named partitions within a single dbspace.

A database can include both fragmented and nonfragmented tables. Index storage can also be fragmented, either in the same storage spaces as their table (called *attached indexes*) or in a different storage distribution scheme (*detached indexes*).

Potential performance and security advantages of distributed storage include these:

- For frequently-accessed tables, fragmentation can reduce the overhead of I/O contention for data that resides on a single storage device.
- The GRANT FRAGMENT and REVOKE FRAGMENT statements of SQL can specify the access privileges that users, roles, or the PUBLIC group hold on specified fragments of the table. With appropriate fragmentation strategies, these statements can selectively restrict user access to subsets of the records in a table.
- For databases that enables parallel-database queries (PDQ), multiple scan threads require less time to scan the fragments than to scan the same rows in a nonfragmented table.
- Input operations that distribute new rows across multiple fragments run more quickly (using multiple INSERT threads) than if a single table extent stores the same rows.
- For fragmentation strategies where the storage allocation of rows is correlated with data values, query execution plans can ignore fragments that are logically excluded by predicates in the query. Defining fragments to improve selectivity is called *fragment elimination*.
- In cluster environments, fragmentation can reduce the time required for recovery from hardware failure, because restoring only a subset of the fragments imposes a smaller data load than restoring the entire table.
- For tables fragmented by interval, the database server create new fragments automatically, simplifying management of the data.



Note: Do not confuse table fragmentation strategies, which can improve the efficiency and throughput of database operations, with the various pejorative meanings of *fragmentation* in reference to file systems that waste storage



space or increase retrieval time through inefficient storage algorithms, or through insufficient use of defragmentation tools to store files in contiguous disk partitions.

HCL OneDB™ fragmentation options

HCL OneDB™ supports the following storage fragmentation strategies that can be applied to database tables:

By Round-robin

A specified number of fragments is defined for the table. Inserted rows are automatically distributed for storage in these fragments, without regard to data values in the row, in order to balance the number of rows in each fragment. Such fragments are called *round-robin fragments*.

By Expression

Each fragment is defined by a Boolean expression that can be evaluated for one or more columns of the table. Inserted rows are stored in a fragments for which the expression that defines the fragment is true for the data in that row. Rows that match the expression for more than one fragment are stored in the first matching fragment within the ordered list of fragments that the system catalog maintains for the table. Such fragments are called *expression fragments*.

By List

Each fragment is defined by a list of one or more constant values that correspond to one or more columns in the table. No two fragments can share the same value in their lists. These values must be categories on a nominal scale that has no quantified order within the set of categories. Inserted rows are stored in the fragment that matches the data value of one or more columns. Such fragments are called *list fragments*.

By Interval

At least one fragment must be defined for values less than a numeric, DATE, or DATETIME column in the table. An interval size, specifying the range of fragment key values assigned to a single fragment, must also be defined. You can optionally specify a list of dbspaces to store interval fragments. The fragments created by the user when the fragmentation strategy is defined are called *range fragments*. The database server automatically creates new fragments of the same interval size to store rows whose fragment key values are outside the range of the user-defined range fragments. Fragments created by the database server are called *interval fragments*.

Each user-defined permanent or temporary database table can either be *nonfragmented* or else can have exactly one fragmentation scheme. You cannot, for example, define a table in which some fragments use a round-robin strategy, and other fragments use a list or interval strategy.

You can use the ALTER FRAGMENT statement of SQL, however, to modify the fragmentation scheme of a table in various ways, including these:

- to change the fragmentation strategy of a fragmented table,
- to define a fragmentation strategy for a nonfragmented table,
- to change a fragmented table to a nonfragmented table,
- to add another fragment to an existing fragmented table,
- to combine two tables that have identical structures into a single fragmented table,

- to drop one or more dbspaces from the list of dbspaces that store interval fragments.
- to detach one fragment from a fragmented table and store the rows in a new nonfragmented table.

For more information about the ALTER FRAGMENT statement and some of the tasks that it can accomplish in data warehousing operations, see the [Change the storage distribution strategy on page 40](#).

Storage fragmentation terms

The following terms are useful for understanding and using the various strategies available for the distributed storage of table and index fragments.

Fragment key

The column or a set of columns on which the table or index is fragmented. Depending on the chosen fragmentation strategy, the fragment key can be a column, or a single column expression, or a multi-column expression. For a row inserted into a table for which a fragment key is defined, the value of the column (or the set for values in the fragment key columns) determines which fragment stores the row. A synonym for fragment key is *partitioning key*. Tables partitioned by round-robin have no fragment key.

Fragment list

An ordered list of the fragments that the database server maintains for every fragmented table or index. By default, the ordinal positions of each fragment on this list reflects the sequence in which the fragments were created. The system catalog stores this integer value in the **sysfragments.evalpos** column of the row that describes the fragment. Queries that do not use fragment elimination read the fragments in ascending order of their **evalpos** values. The database server automatically updates **evalpos** values to reflect changes to the fragment list. Updates to the list are required, for example, when the database server creates an interval fragment, or when the ALTER FRAGMENT statement of SQL adds new fragments, or drops or modifies existing fragments.

Fragment expression

An expression that defines a specific fragment. For example, if the fragment key is **colA** of data type SMALLINT, a fragment could be defined by the expression `colA <=8 OR colA IN (9,10,21,22,23)` in an expression based fragmentation strategy.

- Expression-based fragments are defined by a Boolean expression.
- List-based fragments are defined by one or more constant expressions.
- Range fragments (in interval fragmentation) are defined by a range expression. The only valid operator in the range expression is the less-than (`<`) operator. (For example, `VALUES < 100`).
- System-defined interval fragments (in interval fragmentation) are defined by a system-generated expression that includes the greater-than-or-equal `>=` relational operator, the `AND` Boolean operator, and the less-than (`<`) relational operator. (For example, `VALUES >= 100 AND VALUES < 300` specifies an interval that includes fragmentation key values ranging from 100 to the (non-inclusive) upper limit of 300.)

Tables partitioned by round-robin have no fragment expressions.

NULL fragment

A fragment that stores NULL values (either because its range fragment or list fragment expression is `IS NULL`, or because a list-based or expression-based fragment is defined with NULL as its fragment expression). For all fragmentation strategies except round-robin, the database server returns an exception if you insert a row whose fragment key value is missing, but no NULL fragment is defined (and for list or expression strategies, no REMAINDER fragment is defined). You do not need to define a NULL fragment if the fragment key column enforces a NOT NULL constraint.

REMAINDER fragment

A fragment that stores any row whose fragment key value does not match the fragment expression of any fragment. If you attempt to insert a row that does not match any fragment key value for a table or index that is fragmented by expression or by list, and no REMAINDER fragment is defined, the database server issues an exception. You cannot define a REMAINDER fragment for tables fragmented by a round-robin or interval strategy.

Transition fragment

In an interval fragmentation scheme, the range fragment whose upper limit in its VALUES clause is larger than the upper limit for any other range fragment. If no interval fragments have been created for the table, inserting a row whose fragment-key value exceeds that upper limit requires the database server to create a new interval fragment. The upper limit of the transition fragment VALUES clause is called the *transition value* for the table.

The MODIFY INTERVAL TRANSITION option to the ALTER FRAGMENT statement can increase the transition value for a table. This can result in a different fragment becoming the new transition fragment. This and other ALTER FRAGMENT operations can cause changes to column values in the **sysfragments** system catalog table for the transition fragment, including these:

- its position relative to other fragments (the **evalpos** column),
- its fragment expression (the **exprtext** and **exprbin** columns),
- and its name (the **partition** column).

Fragmentation by ROUND ROBIN

For a table that uses a round-robin distribution scheme, the rows that the database server stores in an insert or load operation are distributed cyclically among a user-defined number of fragments, so that the number of rows inserted into each fragment is approximately the same (± 100).

Round-robin distributions are also called *even* distributions, because the design goal of this strategy is for an evenly balanced distribution among the fragments. To that end, a newly added round-robin fragment will be favored exclusively by inserts and loads until it no longer has the fewest number of rows among the table's fragments.

The syntax for defining round-robin interval fragmentation requires that you specify at least two round-robin fragments in one of two forms. This form defines round-robin fragments and declares a name for each fragment:

```
FRAGMENT BY ROUND ROBIN
PARTITION partition IN dbspace,
. . .
```



```
PARTITION partition IN dbspace
```

As in other fragmentation schemes, each `PARTITION partition` specification declares the name of a fragment, which must be unique among the names of fragments of the same table. The `dbspace` specification can be different for each fragment, or some fragments (or all of the fragments) can be stored in separate named partitions of the same `dbspace`. Each `partition` is the name of a round-robin fragment.

This alternative form defines round-robin fragments with no explicit name:

```
FRAGMENT BY ROUND ROBIN IN dbspace_list
```

Here the `dbspace_list` specification is a comma-separated list of at least 2 (but no more than 2048) `dbspaces`, each of which stores a single round-robin fragment. No `dbspace` can appear more than once in this list. (In the system catalog, the `sysfragments.partition` column stores the identifier of the fragment. For fragments defined without the `PARTITION` keyword, the `partition` value is the identifier of the `dbspace` where the fragment is stored. For this reason, a repeated `dbspace` in `dbspace_list` violates a uniqueness requirement for names of fragments of the same table.)

A round-robin distribution scheme must be defined by only one or the other of these two syntax forms.

A table that is fragmented by round-robin has no fragment key, no fragment expressions, and no `REMAINDER` fragment. (An alternative description is that every round-robin fragment resembles a remainder fragment, because no fragment expressions are defined to match a fragment key for the inserted rows. But the `REMAINDER` keyword is not valid in the SQL syntax to define a round-robin distribution strategy.)

Because no fragment expressions are evaluated when the database server loads new rows into round-robin fragments, this strategy provides the best performance for insert operations.

Only tables, not indexes, can be defined with round-robin fragmentation. For performance reasons, any indexes that you define on a table that is fragmented by round-robin should be nonfragmented indexes.

Because a round-robin distribution strategy has no fragment key and no fragment expressions, you cannot explicitly define a `NULL` round-robin fragment. When rows with missing data are loaded into a table by round-robin, the rows with `NULL` values are stored wherever the database server happens to insert them as it approximately equalizes the number of inserted rows for every fragment.

By design, the `GRANT FRAGMENT` and `REVOKE FRAGMENT` statements of SQL cannot reference round-robin fragments. Because each fragment stores a quasi-random subset of the rows, the DBA cannot predict which rows will be stored in a given round-robin fragment. If some rows contain unencrypted sensitive information, table-level (rather than fragment-level) is a more appropriate granularity for granting or withholding discretionary access privileges in databases that do not implement label-based (LBAC) security policies.

Because round-robin fragments are uncorrelated with data values, queries of tables that are fragmented by round-robin cannot benefit from fragment elimination. Round-robin distribution schemes are useful for balancing the rows in a set of table fragments across multiple devices, but other storage distribution schemes are typically used in data warehouse applications that query dimensional tables, because the performance advantages of round-robin in loading data are more than offset by slower data retrieval from round-robin fragments.

Fragmentation by EXPRESSION

For a table that uses an expression-based distribution scheme, the rows that the database server stores in an insert or load operation are distributed among a user-defined number of fragments, in which each fragment is defined by a Boolean expression for the fragment key.

The fragment expression must be a column expression. This can be the same column (or the same set of columns) for all of the fragments, or different fragments can be defined with different keys. The expression can only reference columns in the table that is being fragmented. Subqueries or calls to user-defined routines are not valid.

The syntax for defining an expression fragmentation strategy defines one or more expression fragments of this form:

```
FRAGMENT BY EXPRESSION
  PARTITION partition expression IN dbspace,
  . . .
  PARTITION partition expression IN dbspace,
  PARTITION partition VALUES (NULL) IN dbspace,
  PARTITION partition REMAINDER IN dbspace
```

As in other fragmentation schemes, each `PARTITION partition` specification declares the unique name of a fragment. The `expression` specification defines the fragment expression, and the `IN dbspace` specification defines the storage location for the fragment. You can optionally define a NULL fragment by specifying `NULL` as the `expression`.

You also can optionally define a REMAINDER fragment for rows that match none of the specified fragment expressions. For some queries, the REMAINDER fragment might be difficult to eliminate, and for some tables, the REMAINDER fragment might become quite large, but the database server issues an exception if the fragment key value for an inserted row matches no fragment expression, and no REMAINDER fragment is defined.

You can optionally define a NULL fragment to stores rows in which the fragment key value is missing.

During an insert into a table that is fragmented by expression, the database server takes these actions:

1. The fragment key value for the row is evaluated.
2. The fragment expression for each fragment is evaluated and compared to the fragment key value for the row, beginning with the fragment whose `sysfragments.ivalpos` value in the system catalog is lowest.
3. If there is no match, the previous step is repeated for the fragment with next highest `sysfragments.ivalpos` value.
4. This continues until the first match is found between the fragment key value and a fragment expression, after which the row is stored in the matching fragment.
5. If no match is found in the entire list of fragments, the row is stored in the REMAINDER fragment. (In this case of a row with an unmatched fragment key, if no REMAINDER is defined, an exception is issued.)

For expression-based fragmentation schemes that define overlapping fragment expressions, the storage location of rows that match the fragment expression of more than one fragment is dependent on the `ivalpos` value for that fragment. You can avoid this dependency by only defining non-overlapping fragment expressions.

The `ivalpos` value of a fragment is determined by its position in the initial fragment list within the `FRAGMENT BY EXPRESSION` or `PARTITION BY EXPRESSION` clause that defined the storage distribution of the table. Any new fragments

added by ALTER FRAGMENT operations are assigned, by default, the next higher **evalpos** value (and will therefore be evaluated last during INSERT operations) unless you explicitly specify a position with the BEFORE or AFTER keyword. In this case, the **evalpos** value for the new fragment will be the ordinal position where was inserted into the fragment list. For tables that are fragmented by expression into a large number of fragments, you can achieve greater efficiency in INSERT and LOAD operations when fragments that are more likely to match fragment key values have relatively low **evalpos** values within the fragment list.

Fragmentation by expressions that creates nonoverlapping fragments on a single column can be an effective strategy for supporting fragment elimination in queries. The database server can eliminate fragments, for example, for queries with range expressions as well as queries with equality expressions if the query predicates correspond to fragment expressions. Expressions with relational operators and logical operators (or with both) can similarly be used for fragment expressions that match query filters.

Fragmentation by LIST

A list fragmentation strategy partitions data into a set of fragments that are each defined by a list of discrete values of the fragment key. Every expression must be a quoted string or a literal value. Each value in the list must be unique among the lists for fragments of the same object.

Fragmenting by list resembles fragmentation by expression (where the fragment expressions include the IN operator or the logical OR operator) in these respects:

- Every non-REMAINDER fragment stores rows for which the fragment key values matches the fragment expression.
- You can optionally define a REMAINDER fragment.
- You can optionally define a NULL fragment.

As the name implies, however, fragmentation by list defines each fragment by a list of fragment expressions, rather than restricting each fragment to a single expression.

The syntax for defining a list fragmentation strategy requires one or more list fragments of the following form.

```
FRAGMENT BY LIST
  PARTITION partition VALUES (expression_list) IN dbspace,
  . . .
  PARTITION partition VALUES (expression_list) IN dbspace,
  PARTITION partition VALUES (NULL) IN dbspace,
  PARTITION partition REMAINDER IN dbspace
```

Here the last two partitions (whose expressions define a NULL fragment and a REMAINDER fragment) are not required.

As with other fragmentation schemes, each `PARTITION partition` specification declares the unique name of a fragment. The `(expression_list)` specification is the comma-separated list of one or more constant expressions that defines each list fragment, and the `IN dbspace` specification identifies the storage location of the fragment.

You can optionally define a NULL fragment by specifying `NULL` as the only expression in the `expression_list`. You cannot include `NULL` in an expression list with other values that define the same fragment.

An alternative syntax notation for defining the NULL fragment is `VALUES IS NULL` (with no delimiting parentheses) as the only expression for a fragment. The digit `0` is not equivalent to the NULL or IS NULL keywords.

Just as in expression-based fragmentation, you can optionally define a REMAINDER fragment for rows that match none of the specified fragment expressions. If you define a REMAINDER fragment but no NULL fragment, rows with the fragment key value missing are stored in the REMAINDER fragment. The database server issues an exception for INSERT operations if the fragment key value for an inserted row matches no fragment expression, and no REMAINDER fragment is defined. An exception is similarly issued if data is missing from the fragment key column, but the fragment list includes no NULL fragment and no REMAINDER fragment.

When you use the CREATE INDEX statement to define an index on a table that is fragmented by list, it is not necessary to include the FRAGMENT BY or PARTITION BY clause to create indexes that use the same list fragmentation strategy as their table. By default, the database server partitions the index by the same list fragmentation strategy as its table, and declares for each index fragment the same name that you specified after the PARTITION keyword for the corresponding table fragment.

The most important difference between fragmentation by list and fragmentation by expression is that every value in the list for each fragment must be unique among all the expression lists that define fragments of the same table or index. The database server issues an error if the lists of expressions for two list fragments include the same fragment key value. This uniqueness requirement for fragment expressions simplifies fragment elimination in queries, if the fragment expressions correspond to query predicates and filters that support fragment elimination.

A list fragmentation strategy is most effective when the fragment key for a table has finite set of values, and queries on the table specify equality predicates on the fragment key. For a table whose fragment key is a numeric or time data type with a range of possible values that resembles a continuum, an interval fragmentation scheme is recommended, rather than list fragmentation.

Fragmentation by INTERVAL

An interval fragmentation strategy partitions data into fragments based on an interval value of the fragment key. The interval value must be a column expression that references a single column of a numeric, DATE, or DATETIME data type.

This type of distribution scheme is sometimes called a *range interval* distribution because:

- The RANGE and INTERVAL keywords are required in the DDL syntax that defines this strategy.
- The initial user-defined fragments are called *range fragments* to distinguish these fragments from system-defined fragments. The database server creates system-defined fragments automatically when a row is inserted whose fragment key value does not match the expression that defines any existing fragment.

The INTERVAL distribution strategy is useful when all possible fragment key values in a growing table are not known, and the DBA does not want to allocate fragments for data rows that are not yet loaded. For example, by using a DATE column as a fragment key could define a fragment for every month, or a BIGSERIAL column as a fragment key could define a fragment for every million customer records. The automatic creation of interval fragments avoids the need for a REMAINDER fragment (with its associated fragment-elimination difficulties) and can also reduce the maintenance workload of the DBA.

Defining an interval distribution strategy

The definition of an interval distribution scheme can include several required or optional parameters:

Fragment key

This must specify a column expression referencing a single numeric, DATE, or DATETIME column of the table.

Interval value expression

This constant expression defines an interval size within the range of fragment key values for system-generated interval fragments.

Storage location for interval fragments

This is a list of dbspaces where interval fragments will be stored.

Range fragment list

You must declare the name and define the fragment expression and the storage location for at least one range fragment.

The syntax for defining these parameters of a range interval distribution has this general form:

```
FRAGMENT BY RANGE (column_expr)
  INTERVAL (interval_size) STORE IN (dbspace_list)
  PARTITION partition VALUES < upper_bound IN dbspace,
  . . .
  PARTITION partition VALUES < upper_bound IN dbspace,
  PARTITION partition VALUES IS NULL IN dbspace
```

In this template, the syntax tokens that are not keywords specify the following parameters of the storage distribution scheme:

RANGE(*column_expr*)

This column expression, referencing a single column and delimited by parentheses, defines the fragment key. This clause is required.

INTERVAL(*interval_size*)

This interval value expression, delimited by parentheses, defines the interval size (within the range of values of the fragment key) for system-generated interval fragments. This clause is required.

STORE IN(*dbspace_list*)

This is a list of dbspaces where interval fragments will be stored. If you specify more than one dbspace, the database server creates successive new interval fragments in these dbspaces, in round robin fashion. This clause is optional.

If you omit this clause, the database server stores interval fragments in the dbspace that stores the range fragment. (If you define two or more range fragments and store them in different dbspaces, the database server stores each new interval fragment in one of these dbspaces, assigning successive interval fragments to these dbspaces in round robin fashion.)

PARTITION *partition*

This declares the name of a range or NULL fragment. You must define at least one range fragment. The NULL fragment is optional, but no more than one NULL fragment can be defined.

If you define more than one fragment, their names must conform to the rules for SQL identifiers, and must be unique among the fragments of the same table or index.

VALUES < *upper_bound*

This defines the fragment expression. Unlike list fragments, which can be defined in an arbitrary order, if you define more than one range fragment, their expressions must be defined in ascending order of the *upper_bound*. This clause is required.

The last range fragment that you define (which can be the first, if you define only one), is called the *transition fragment*, and its upper bound is called the *transition value* for the fragmented object. Any inserted rows with a larger fragment key value must be stored in an interval fragment.

VALUES IS NULL

This is the fragment expression to define the NULL fragment. Whether you define a NULL fragment is optional. The NULL fragment is not a range fragment, because NULL indicates the absence of a fragment key value. The database server issues an exception if a DML operation attempts to insert a row that has no fragment key value into a fragmented table for which no NULL fragment is defined.

If you define a NULL fragment, it can be listed in any position within the PARTITION specifications. The database server, rather than the sequence in which you declare the fragments, internally determines the order of each fragment within the fragment list of a table or index that uses an interval fragmentation scheme. The NULL fragment, if it exists, is always the first on this list, as indicated by its **sysfragments.evalpos** value in the system catalog.

When a row is inserted whose fragment key value is outside the range of any existing range or interval fragments, the database server will automatically create a new interval fragment based on the *interval_size* value and the transition value, without DBA intervention.

This kind of fragmentation strategy is useful when all possible fragment key values in a growing table are not known and the DBA does not want to allocate fragments for data that is not yet loaded.

Handle common dimensional data-modeling problems

The dimensional model that the previous sections describe illustrates only the most basic concepts and techniques of dimensional data modeling. The data model you build to address the business needs of your enterprise typically involves additional problems and difficulties that you must resolve to achieve the best possible query performance from your database. This section describes various methods you can use to resolve some of the most common problems that arise when you build a dimensional data model.

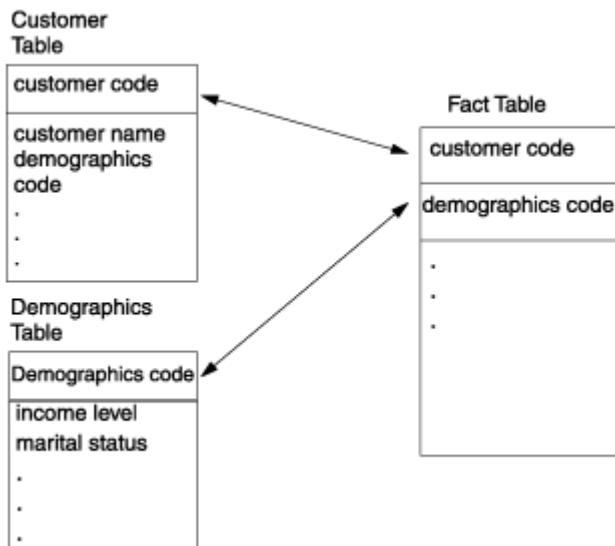
Minimize the number of attributes in a dimension table

Dimension tables that contain customer or product information might easily have 50 to 100 attributes and many millions of rows. However, dimension tables with too many attributes can lead to excessively wide rows and poor performance. For this reason, you might want to separate out certain groups of attributes from a dimension table and put them in a separate table called a *minidimension* table. A minidimension table consists of a small group of attributes that are separated out from a larger dimension table. You might choose to create a minidimension table for attributes that have either of the following characteristics:

- The fields are rarely used as constraints in a query.
- The fields are frequently compared together.

The following figure shows a minidimension table for demographic information that is separated out from a **customer** table.

Figure 10. A Minidimension Table for Demographics Information




In the **demographics** table, you can store the demographics key as a foreign key in both the fact table and the **customer** table, which allows you to join the demographics table directly to the fact table. You can also use the demographics key directly with the **customer** table to browse demographic attributes.

Dimensions that occasionally change

In a dimensional database where updates are infrequent (as opposed to OLTP systems), most dimensions are relatively constant over time, because changes in sales districts or regions, or in company names and addresses, occur infrequently. However, to make historical comparisons, these changes must be handled when they do occur. The following figure shows an example of a dimension that has changed.

Figure 11. A dimension that changes



You can use three ways to handle changes that occur in a dimension:

Change the value stored in the dimension column

In the previous figure, the record for Bill Adams in the **customer** dimension table is updated to show the new address `Arlington Heights`. All of this customer's previous sales history is now associated with the district of Arlington Heights instead of Des Plaines.

Create a second dimension record with the new value and a generalized key

This approach effectively partitions history. The **customer** dimension table would now contain two records for Bill Adams. The old record with a key of 101 remains, and records in the fact table are still associated with it. A new record is also added to the **customer** dimension table for Bill Adams, with a new key that might consist of the old key plus some version digits (101.01, for example). All subsequent records that are added to the fact table for Bill Adams are associated with this new key.

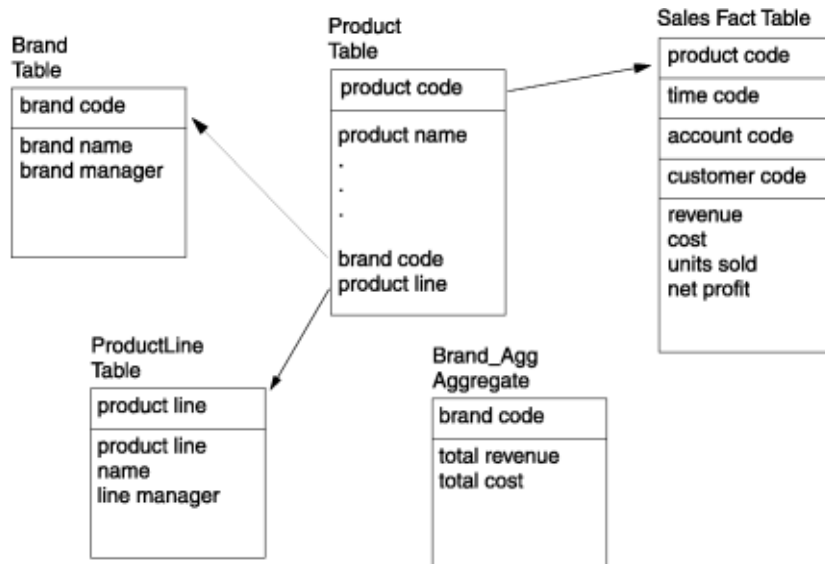
Add a new field in the customer dimension table for the affected attribute and rename the old attribute

This approach is rarely used unless you need to track old history in terms of the new value and vice-versa. The **customer** dimension table gets a new attribute named **current address**, and the old attribute is renamed **original address**. The record that contains information about Bill Adams includes values for both the original and current address.

Use the snowflake schema for hierarchical dimension tables

A *snowflake schema* is a variation on the star schema, in which very large dimension tables are normalized into multiple tables. Dimensions with hierarchies can be decomposed into a snowflake structure when you want to avoid joins to big dimension tables when you are using an aggregate of the fact table. For example, if you have brand information that you want to separate out from a **product** dimension table, you can create a brand snowflake that consists of a single row for each brand and that contains significantly fewer rows than the **product** dimension table. The following figure shows a snowflake structure for the brand and product line elements and the **brand_agg** aggregate table.

Figure 12. An example of a snowflake schema



If you create an aggregate table, **brand_agg**, that consists of the brand code and the total revenue per brand, you can use the snowflake schema to avoid the join to the much larger **sales** table. For example, you can use the following query on the **brand** and **brand_agg** tables:

```
SELECT brand.brand_name, brand_agg.total_revenue
FROM brand, brand_agg
WHERE brand.brand_code = brand_agg.brand_code
AND brand.brand_name = 'Anza'
```

Without a snowflaked dimension table, you use a `SELECT UNIQUE` or `SELECT DISTINCT` statement on the entire **product** table (potentially, a very large dimension table that includes all the brand and product-line attributes) to eliminate duplicate rows.

While snowflake schemas are unnecessary when the dimension tables are relatively small, a retail or mail-order business that has customer or product dimension tables that contain millions of rows can use snowflake schemas to significantly improve performance.

If an aggregate table is not available, any joins to a dimension element that was normalized with a snowflake schema must now be a three-way join, as the following query shows. A three-way join reduces some of the performance advantages of a dimensional database.

```
SELECT brand.brand_name, SUM(sales.revenue)
FROM product, brand, sales
WHERE product.brand_code = brand.brand_code
AND brand.brand_name = 'Alltemp'
GROUP BY brand_name
```

Implement a dimensional database

You will learn the SQL statements that you need to implement the dimensional data model

This section shows you the SQL statements required to implement the dimensional database that is described in the section [Design a dimensional data model on page 8](#). Remember that this database serves only as an illustrative example of a data-warehousing environment. For the sake of the example, it is translated into SQL statements.

This section describes the **sales_demo** database.

Implement the sales_demo dimensional database

This section shows the SQL statements that you can use to create a dimensional database from the data model that you learned about in [Design a dimensional data model on page 8](#). You can use interactive SQL to write the individual statements that create the database or you can run a script that automatically executes all the statements that you need to implement the database. The CREATE DATABASE and CREATE TABLE statements create the data model as tables in a database. After you create the database, you can use the LOAD and INSERT statements to populate the tables.

Create the dimensional database

You must create the dimensional database before you can create any of the tables or other objects that the database must contain.

When you use the HCL® OneDB® database server to create a database, the server sets up records that show the existence of the database and its mode of logging. The database server manages disk space directly, so these records are not visible to operating-system commands.


The following statement shows the syntax that you use to create a database that is called **sales_demo**:

```
CREATE DATABASE sales_demo
```

The CREATE TABLE statement for the dimension and fact tables

This section includes the CREATE TABLE statements that you use to create the tables of the **sales_demo** dimensional database.

Referential integrity is, of course, an important requirement for dimensional databases. However, the following schema for the **sales_demo** database does not define the primary and foreign key relationships that exist between the fact table and its dimension tables. The schema does not define these primary and foreign key relationships because data-loading performance improves dramatically when the database server does not enforce constraint checking. Given that data warehousing environments often require that tens or hundreds of gigabytes of data are loaded within a specified time, data-load performance should be a factor when you decide how to implement a database in a warehousing environment. Assume that if the **sales_demo** database is implemented as a live data mart, some data extraction tool (rather than the database server) is used to enforce referential integrity between the fact table and dimension tables.

 **Tip:** After you create and load a table, you can add primary key and foreign key constraints to the table with the ALTER TABLE statement to enforce referential integrity. This method is required only for express load mode. If the constraints and indexes are necessary and costly to drop before a load, then deluxe load mode is the best option.

The following statements create the **time**, **geography**, **product**, and **customer** tables. These tables are the dimensions for the **sales** fact table. A SERIAL field serves as the primary key for the **district_code** column of the **geography** table.

```
CREATE TABLE time
(
time_code      INT,
order_date     DATE,
month_code     SMALLINT,
month_name     CHAR(10),
quarter_code   SMALLINT,
quarter_name   CHAR(10),
year          INTEGER
);

CREATE TABLE geography
(
district_code  SERIAL,
district_name  CHAR(15),
state_code     CHAR(2),
state_name     CHAR(18),
region        SMALLINT
);

CREATE TABLE product (
product_code   INTEGER,
product_name   CHAR(31),
vendor_code    CHAR(3),
vendor_name    CHAR(15),
product_line_code  SMALLINT,
product_line_name  CHAR(15)
);

CREATE TABLE customer (
customer_code  INTEGER,
customer_name  CHAR(31),
company_name   CHAR(20)
);
```

The **sales** fact table has pointers to each dimension table. For example, **customer_code** references the customer table, **district_code** references the geography table, and so forth. The **sales** table also contains the measures for the units sold, revenue, cost, and net profit.

```
CREATE TABLE sales
(
customer_code  INTEGER,
district_code  SMALLINT,
time_code     INTEGER,
product_code   INTEGER,
units_sold    SMALLINT,
revenue       MONEY(8,2),
cost          MONEY(8,2),
```

```
net_profit    MONEY(8,2)
);
```

i Tip: The most useful measures (facts) are numeric and additive. Because of the great size of databases in data-warehousing environments, virtually every query against the fact table might require thousands or millions of records to construct a result set. The only useful way to compress these records is to aggregate them. In the **sales** table, each column for the measures is defined on a numeric data type, so you can easily build result sets from the **units_sold**, **revenue**, **cost**, and **net_profit** columns.

For your convenience, the file called `createdw.sql` contains all the preceding CREATE TABLE statements.

Mapping data from data sources to the database

The **stores_demo** demonstration database is the primary data source for the **sales_demo** database.

The following table shows the relationship between data warehousing business terms and the data sources. It also shows the data source for each column and table of the **sales_demo** database.

Table 7. The relationship between data warehousing business terms and data sources

Business Term	Data Source	Table.Column Name
Sales Fact Table:		
product code		sales.product_code
customer code		sales.customer_code
district code		sales.district_code
time code		sales.time_code
revenue	stores_demo:items.total_price	sales.revenue
units sold	stores_demo:items.quantity	sales.units_sold
cost	costs.lst (per unit)	sales.cost
net profit	calculated: revenue minus cost	sales.net_profit
Product Dimension Table:		
product	stores_demo:catalog.catalog_num	product.product_code
product name	stores_demo:stock.manu_code and stores_demo:stock.description	product.product_name
product line	stores_demo:orders.stock_num	product.product_line_code
product line name	stores_demo:stock.description	product.product_line_name

Table 7. The relationship between data warehousing business terms and data sources

(continued)

Business Term	Data Source	Table.Column Name
vendor	stores_demo:orders.manu_code	product.vendor_code
vendor name	stores_demo:manufact.manu_name	product.vendor_name
Customer Dimension Table:		
customer	stores_demo:orders.customer_num	customer.customer_code
customer name	stores_demo:customer.fname plus stores_demo:customer.lname	customer.customer_name
company	stores_demo:customer.company	customer.company_name
Geography Dimension Table:		
district code	generated	geography.district_code
district	stores_demo:customer.city	geography.district_name
state	stores_demo:customer.state	geography.state_code
state name	stores_demo.state.sname	geography.state_name
region	derived: If state = "CA" THEN region = 1, ELSE region = 2	geography.region
Time Dimension Table:		
time code	generated	time.time_code
order date	stores_demo:orders.order_date	time.order_date
month	derived from order date generated	time.month_name time.month.code
quarter	derived from order date generated	time.quarter_name time.quarter_code
year	derived from order date	time.year

Several files with a `.unl` suffix contain the data that is loaded into the **sales_demo** database. The files that contain the SQL statements that create and load the database have a `.sql` suffix.

If your database server runs on UNIX™, you can access the `*.sql` and `*.unl` files from the directory `$ONEDB_HOME/demo/dbaccess`.

If your database server runs on Windows™, you can access the `*.sql` and `*.unl` files from the directory `%ONEDB_HOME%\demo\dbaccess`.

Load data into the dimensional database

An important step when you implement a dimensional database is to develop and document a load strategy. This section shows the LOAD and INSERT statements that you can use to populate the tables of the **sales_demo** database.

i Tip: In a live data warehousing environment, you typically do not use the LOAD or INSERT statements to load large amounts of data to and from HCL® OneDB® databases.

HCL® OneDB® database servers provide different features for loading and unloading of data.

For information about loading, see your *HCL OneDB™ Administrator's Guide*.

The following statement loads the **time** table with data first so that you can use it to determine the time code for each row that is loaded into the **sales** table:

```
LOAD FROM 'time.unl' INSERT INTO time
```

The following statement loads the **geography** table. After you load the **geography** table, you can use the district code data to load the **sales** table.

```
INSERT INTO geography(district_name, state_code, state_name)
SELECT DISTINCT c.city, s.code, s.sname
  FROM stores_demo:customer c, stores_demo:state s
  WHERE c.state = s.code
```

The following statements add the region code to the **geography** table:

```
UPDATE geography
  SET region = 1
  WHERE state_code = 'CA'

UPDATE geography
  SET region = 2
  WHERE state_code <> 'CA'
```

The following statement loads the **customer** table:

```
INSERT INTO customer (customer_code, customer_name, company_name)
SELECT c.customer_num, trim(c.fname) || ' ' || c.lname, c.company
FROM stores_demo:customer c
```

The following statement loads the **product** table:

```
INSERT INTO product (product_code, product_name, vendor_code,
  vendor_name, product_line_code, product_line_name)
SELECT a.catalog_num,
  trim(m.manu_name) || ' ' || s.description,
  m.manu_code, m.manu_name,
  s.stock_num, s.description
FROM stores_demo:catalog a, stores_demo:manufact m,
  stores_demo:stock s
  WHERE a.stock_num = s.stock_num
  AND a.manu_code = s.manu_code
  AND s.manu_code = m.manu_code;
```

The following statement loads the **sales** fact table with one row for each product, per customer, per day, per district. The cost from the **cost** table is used to calculate the total cost (cost * quantity).

```
INSERT INTO sales (customer_code, district_code, time_code,
  product_code, units_sold, cost, revenue, net_profit)
SELECT
  c.customer_num, g.district_code, t.time_code,
  p.product_code, SUM(i.quantity),
  SUM(i.quantity * x.cost), SUM(i.total_price),
  SUM(i.total_price) - SUM(i.quantity * x.cost)
FROM stores_demo:customer c, geography g, time t,
  product p, stores_demo:items i,
  stores_demo:orders o, cost x
WHERE c.customer_num = o.customer_num
  AND o.order_num = i.order_num
  AND p.product_line_code = i.stock_num
  AND p.vendor_code = i.manu_code
  AND t.order_date = o.order_date
  AND p.product_code = x.product_code
  AND c.city = g.district_name
GROUP BY 1,2,3,4;
```

Test the dimensional database

After you create the tables and load the data into the database, you should test the dimensional database.

You can create SQL queries to retrieve the data necessary for the standard reports listed in the business-process summary (see the [Summary of a business process on page 13](#)). Use the following ad hoc queries to test that the dimensional database was properly implemented.

The following statement returns the monthly revenue, cost, and net profit by product line for each vendor:

```
SELECT vendor_name, product_line_name, month_name,
  SUM(revenue) total_revenue, SUM(cost) total_cost,
  SUM(net_profit) total_profit
FROM product, time, sales
WHERE product.product_code = sales.product_code
  AND time.time_code = sales.time_code
GROUP BY vendor_name, product_line_name, month_name
ORDER BY vendor_name, product_line_name;
```

The following statement returns the revenue and units sold by product, by region, and by month:

```
SELECT product_name, region, month_name,
  SUM(revenue), SUM(units_sold)
FROM product, geography, time, sales
WHERE product.product_code = sales.product_code
  AND geography.district_code = sales.district_code
  AND time.time_code = sales.time_code
GROUP BY product_name, region, month_name
ORDER BY product_name, region;
```

The following statement returns the monthly customer revenue:

```
SELECT customer_name, company_name, month_name,
  SUM(revenue)
FROM customer, time, sales
```

```
WHERE customer.customer_code = sales.customer_code
      AND time.time_code = sales.time_code
GROUP BY customer_name, company_name, month_name
ORDER BY customer_name;
```

The following statement returns the quarterly revenue per vendor:

```
SELECT vendor_name, year, quarter_name, SUM(revenue)
FROM product, time, sales
WHERE product.product_code = sales.product_code
      AND time.time_code = sales.time_code
GROUP BY vendor_name, year, quarter_name
ORDER BY vendor_name, year
```

Change the storage distribution strategy

Use the ALTER FRAGMENT statement to change the storage distribution strategy of the data rows that are being loaded into an existing database table.

You should adjust the storage distribution strategy when the volume or distribution of the data is different than what was originally expected when the storage distribution plan was first implemented. The ALTER FRAGMENT can also be used as part of the workflow of a data warehouse. If a table is partitioned with a fragment key that is based on values in a DATE or DATETIME column, the fragments can be detached from the table. As new fragments are added to the table older fragments, that store rows from earlier time periods, can be detached from the table.

The ALTER FRAGMENT statement supports the following six options for table fragments. Some ALTER FRAGMENT options are valid for nonfragmented tables or for index fragments.



Note: The following summary ignores tables that are fragmented by ROUND ROBIN, because other fragmentation strategies are more often used in data warehousing applications.

ADD

Adds a new fragment in the list of fragments that are part of a table that has been fragmented by any fragmentation scheme.

For LIST or EXPRESSION fragments, you can add a NULL fragment or a REMAINDER fragment, if none of these types of fragments have already been defined. You can use the BEFORE or AFTER keyword to specify the ordinal position of the new fragment in the list of expressions or list of fragments.

For a table that has been fragmented with the INTERVAL option, you can use the ADD option can add new storage spaces to the list of dbspaces where the database server creates new INTERVAL fragments.

ATTACH

Combines two or more tables that have identical structures into a fragmentation strategy. All of the consumed tables specified in the ALTER FRAGMENT ATTACH statement must have the same structure as the surviving table. The number, names, data types, and relative positions of the columns must be identical. The consumed tables must be nonfragmented, and must be stored in a different dbspace from the surviving table. The ATTACH option does not support index fragments.

For LIST or EXPRESSION fragments, you can attach a NULL fragment or a REMAINDER fragment, if none of these types of fragments have already been defined. You can use the BEFORE or AFTER keyword to specify the ordinal position of the new fragment in the list of fragments.

For a table that has been fragmented by INTERVAL, the ATTACH option can attach new RANGE fragments. However you cannot attach new INTERVAL fragments, and you cannot use the BEFORE or AFTER keyword to specify the ordinal position of the new RANGE fragment.

When a new EXPRESSION fragment is attached to table that is fragmented by LIST or by INTERVAL, the rows from the consumed table and the affected fragments in the surviving table are scanned and moved into appropriate partitions. These strategies are not overlapping.

You can also include the ONLINE keyword in ALTER FRAGMENT ATTACH statements with interval partitioning. Specifying this keyword can improve concurrency for other sessions that attempt to access the tables on which the ALTER FRAGMENT ONLINE statement operates.

DETACH

Removes a table fragment from a distribution scheme and places the contents into a new, nonfragmented table. The DETACH option does not support index fragments.

The table from which the fragment was detached remains fragmented, unless it is fragmented by LIST or by EXPRESSION and had only two fragments before the DETACH operation.

The new table does not inherit any indexes, constraints, or discretionary access privileges of the table from which it was detached. The new table has the default access privileges of any new table.

The ALTER FRAGMENT DETACH statement cannot remove a fragment from a table that is the parent of a referential constraint, or from a table on which a ROWID column is defined.

You can also include the ONLINE keyword in ALTER FRAGMENT DETACH statements with interval partitioning.

DROP

Drops a fragment from a table or index that is fragmented by LIST or by EXPRESSION. Using the DROP option requires, however, that any rows currently stored in the fragment can be moved to another existing fragment. For LIST fragments, the existing fragment can only be the REMAINDER fragment, because of the uniqueness requirement for LIST fragment expressions.

For a table or index that is fragmented by INTERVAL, you can use the DROP option to drop one or more dbspaces from the list of dbspaces that store INTERVAL fragments. No new INTERVAL fragments will be created in the specified dbspaces.

ALTER TABLE DROP operations that result in moving a large number of rows can fail if insufficient log space or disk space is available. You might be able to complete the operation by dividing it into a series of smaller operations. If insufficient log space causes the failure, an alternative is to temporarily turning off logging. Then retry the ALTER TABLE operation and turn transaction-logging back on after the operation completes. To perform ALTER TABLE DROP operations requires a backup of the **root** dbspace.

INIT

Defines, modifies, or replaces the fragmentation strategy or the storage location of an existing table or an existing index.

For an index, you can accomplish these tasks:

- Change an existing fragmented index to a nonfragmented index.
- Change the interval value of the interval distribution scheme for a fragmented index.
- Change the interval fragment key of the interval distribution scheme for a fragmented index.
- Fragment an existing index that is not fragmented without redefining the index.
- Change the distribution scheme of an existing fragmented index to another distribution scheme of the same expression, list, or interval type, or to a different type of distribution scheme.

For a nonfragmented table, you can accomplish these tasks:

- Move a nonfragmented table from one dbspace to another dbspace.
- Move a nonfragmented table from one dbspace to a named fragment.
- Change a nonfragmented table to a fragmented table.

For a fragmented table, you can accomplish these tasks:

- Convert a fragmented table to a nonfragmented table.
- Replace the current fragmentation scheme with a different fragmentation scheme of the same type or of a different type
- Change the expression associated with an existing list-based or expression-based fragment
- Add a new rowid column to a fragmented table. This column stores a unique integer that cannot be updated. The database server automatically creates an index on the rowid column. With this column, the database server can find the physical location of any row.

If the table is not empty when you convert an existing storage fragmentation strategy to another strategy, the database server discards the existing fragmentation strategy and moves data rows to fragments that you define in the new fragmentation strategy. Data movement also occurs when you convert a nonfragmented index to a fragmented index, and when you convert a fragmented index to a nonfragmented index. For large tables, data movement can cause significant logging, or the transaction might approach the long-transaction high-watermark, and a relatively long exclusive lock might be held on the affected tables. Use these ALTER FRAGMENT INIT options when they do not interfere with day-to-day operations.

MODIFY

Change the current fragmentation list of a table or of an index.

For LIST or EXPRESSION fragments, the MODIFY option can accomplish these tasks:

- Move an existing fragment from one dbspace to a different dbspace.
- Change the expression associated with an existing list-based or expression-based fragment.
- Rename one or more existing fragments.

For a table that is fragmented by INTERVAL, the MODIFY option can accomplish these tasks:

- Modify the expression that defines a range fragment.
- Increase the value of the expression that defines the transition value of the table.
- Enable or disable the automatic creation of interval fragments.
- Replace the list of dbspaces where system-generated interval fragments will be created. Existing fragments in the old dbspaces are not moved, and new rows that match their fragment expressions will be inserted into those fragments.
- Move a range fragment or an interval fragment to a different dbspace.
- Rename one or more existing fragments.

When you change the expression that defines a range fragment, the replacement expression cannot cross adjacent fragment boundaries.

You cannot modify the system-generated expression for any INTERVAL fragment, and you cannot decrease the transition value of a table that is fragmented by INTERVAL.

You can also include the ONLINE keyword in ALTER FRAGMENT ON TABLE INTERVAL TRANSITION statements.

Moving data from relational tables into dimensional tables by using external tables

Use SQL statements to unload data from relational tables into external tables, which are data files that are in table format, and then load the data from the data files into the dimensional tables.

Before you begin

Before beginning, document a strategy for mapping data in the relational database to the dimensional database.

About this task

To unload data from the relational database into external tables and then load the data into the dimensional database:

1. Unload the data from a relational database to external tables.

Repeat the following steps to create as many external tables as are required for the data that you want to move.

- a. Use the CREATE EXTERNAL TABLE statement to describe the location of the external table and the format of the data.

Example

The following sample CREATE EXTERNAL TABLE statement creates an external table called `emp_ext`, with data stored in a specified fixed format:

```
CREATE EXTERNAL TABLE emp_ext
( name CHAR(18) EXTERNAL CHAR(18),
  hiredate DATE EXTERNAL CHAR(10),
```

```
address VARCHAR(40) EXTERNAL CHAR(40),
empno INTEGER EXTERNAL CHAR(6) )
USING (
FORMAT 'FIXED',
DATAFILES
("DISK:/work2/mydir/employee.unl")
);
```

- b. Use the INSERT...SELECT statement to map the relational database table to the external table.

Example

The following sample INSERT statement loads the employee database table into the external table called emp_ext:

```
INSERT INTO emp_ext SELECT * FROM employee
```

The data from the employee database table is stored in a data file called employee.unl.

2. If necessary, copy or move the data files to the system where the dimensional database is located.
3. Load the data from the data files to the dimensional database.

Repeat the following steps to load all the data files that you created in the previous steps.

- a. Use the CREATE EXTERNAL TABLE statement to describe the location of the data file and the format of the data.

Example

The following code is a sample CREATE EXTERNAL TABLE statement:

```
CREATE EXTERNAL TABLE emp_ext
( name CHAR(18) EXTERNAL CHAR(18),
hiredate DATE EXTERNAL CHAR(10),
address VARCHAR(40) EXTERNAL CHAR(40),
empno INTEGER EXTERNAL CHAR(6) )
USING (
FORMAT 'FIXED',
DATAFILES
("DISK:/work3/mydir/employee.unl")
);
```

- b. Use the INSERT...SELECT statement to map the data from the data file to the table in the dimensional database.

Example

The following sample INSERT statement loads the employee data file into the employee database table:

```
INSERT INTO employee SELECT * FROM emp_ext
```

Performance tuning dimensional databases

This section describes how to tune the performance of your queries and to understand data distribution statistics.

Query execution plans

When a SELECT statement or other DML operation is submitted to the database server, the query execution optimizer designs a query execution plan. The query execution optimizer is often referenced as the *query optimizer*.

To design a query execution plan, and estimate the costs of candidate query plans, the query optimizer considers a wide range of information including:

- Specifications that identify the database objects, predicates, filters, joins, and other operations in the SQL syntax that defines the query operation
- System catalog information about indexes and constraints on the tables, views, and columns that are referenced or implied in the query
- Data distribution statistics for the tables and indexes, or for their fragments, that are referenced or implied in the query
- Optimizer directives that are specified inline or as external optimizer directives that favor or avoid subsets of the potential query plans
- Information in the database server environment or in the session environment that affects the query execution optimizer

Data distribution statistics

Data distribution statistics are stored in the system catalog for use by the query optimizer when it designs query execution plans. These statistics, together with other information, enable the optimizer to estimate the relative costs among the execution plans that the optimizer is considering for a specific query. Distribution statistics that the optimizer examines for tables that are referenced in queries can include column distribution statistics for the table and for its indexes, as well as fragment-level statistics, if the database server has gathered statistics for individual table or index fragments.

The following system catalog tables store data distribution information that is available to the query optimizer:

SYSDISTRIB

Stores data distribution information for tables and indexes.

SYSFRAGDIST

Stores fragment-level data distribution information for fragments of tables and indexes.

The following system catalog tables store information pertaining to changes to rows since the most recent update to table, index, or fragment statistics.

SYSDISTRIB

Counts the number of rows changed by DML operations since table statistics were last updated, the date and time of that update, and the time required to build column distributions.

SYSFRAGDIST

Counts the number of rows changed by DML operations since fragment-level statistics were last updated, and the date and time of that update.

SYSFRAGMENTS

Counts the number of rows changed by DML operations since fragment-level statistics were last updated.

SYSINDICES

Counts the number of rows changed by DML operations since index statistics were last updated, the date and time of that update, and time required to build low level distributions for the lead column of the index.

The following configuration parameters can affect the database server behavior for the calculation, display, or other operations on data distribution statistics for tables or for fragments that can be used in query plans:

AUTO_STAT_MODE

Enable or disable the detection (and selective refreshing) of stale statistics during UPDATE STATISTICS operations. You can override the setting of this parameter by using the **onmode -wm** or **onmode -wf** command-line utilities, or SQL administration API function calls, or (for the current session) by the SET ENVIRONMENT AUTO_STAT_MODE statement of SQL.

EXPLAIN_STAT

Enable or disable the inclusion of a Query Statistics section in the explain output file. This is enabled by default.

SYSSBSPACENAME

Specifies the name of the sbspace in which the database server stores data-distribution statistics (as smart large objects) that the UPDATE STATISTICS statement collects for certain user-defined data types. Because the data distributions for UDTs can be large, you have the option to store them in an sbspace instead of in the **sysdistrib** system catalog table (for table-level statistics) or in the **sysfragdist** system catalog table (for fragment-level statistics), where distribution statistics are stored by default.

STATCHANGE

Specifies a positive integer as a change threshold to identify table or fragment distribution statistics that need to be updated. This is the default threshold for refreshing distribution statistics on tables for which no specific threshold has been specified as a table or session attribute. If no value is specified, the default is 10. While selective refreshing of data distribution statistics enabled (by default, or by the AUTO_STAT_MODE setting, or by the AUTO keyword of the UPDATE STATISTICS statement, UPDATE STATISTICS operations only refresh stale or missing statistics. The default value of 10 restricts recalculation to only the tables or fragments in which DML, load, or TRUNCATE operations have changed more than 10% of the rows since data distribution statistics were most recently calculated.

Fragment-level statistics

For tables and indexes that have been partitioned according to fragment key values, the distribution statistics in the system catalog for some fragments might closely approximate current data distributions in those fragments, despite subsequent DELETE, INSERT, UPDATE, or MERGE operations that have caused the statistics for other fragments to become stale. For large tables that contain millions of rows, substantial resources of the database server can be conserved by updating only the subset of fragments with stale statistics, rather than recalculating distribution statistics for every fragment.

The STATLEVEL table attribute

For tables and indexes that are partitioned into multiple fragments by a distributed storage scheme, you can specify the granularity of its data distribution statistics, and you can specify the criteria by which stale statistics are defined. This can be accomplished by specifying keyword options of the Statistics Options clause in either of two DDL statements:

- in the CREATE TABLE statement (when defining a new fragmented table)
- in the ALTER TABLE statement (when changing the statistics granularity of an existing fragmented table).

In both cases, your options for specifying the granularity of the distribution statistics are the same:

STATLEVEL AUTO

Specifies that the database server apply the following criteria at runtime to determine if fragment-level distributions should be created:

- The table is fragmented by EXPRESSION, by LIST, or by INTERVAL.
- The table has more than 1,000,000 rows.

Unless both of these criteria are satisfied, table-level distributions are created. AUTO is the default setting in the CREATE TABLE statement, if you specify no explicit STATLEVEL setting.

STATLEVEL FRAGMENT

Data distributions will be created and maintained for each fragment. The FRAGMENT option is not valid for nonfragmented tables, or for tables that use a round robin storage distribution scheme.

STATLEVEL TABLE

All data distributions for the table will be created at the table level. This emulates the legacy behavior of HCL OneDB™ servers earlier than version 11.70.

To support fragment level data distribution statistics, you must specify the name of an sbspace as the setting of the SYSSBSPACENAME configuration parameter, and you must also declare the name and allocate storage for that sbspace by using the **-c -S** option of the **onspaces** utility. For any table whose STATLEVEL attribute is set to FRAGMENT, the database server returns an error if SYSSBSPACENAME is not set, or if the sbspace to which is SYSSBSPACENAME is set is not properly allocated. For each fragment, the most recently calculated data distribution statistics are stored as a BLOB object in the **sysfragdist.ncndist** column in the system catalog.

Data distribution statistics gathered at the fragment level can be aggregated to provide table level statistics from the constituent fragment statistics.

The STATCHANGE threshold for refreshing data distribution statistics

The same Statistics Options clause of the CREATE TABLE or ALTER TABLE statement can also specify a change threshold for data distribution statistics. The database server applies this STATCHANGE attribute of a fragmented table to all of

the fragments of the table. The STATCHANGE table attribute can be set to an integer value, or you can specify the AUTO keyword:

integer

This defines an integer change threshold between 0 and 100 which defines how much table or fragment data is allowed to change before its statistics are considered stale in UPDATE STATISTICS operations that selectively update only stale distribution statistics.

AUTO

The threshold is the value of the STATCHANGE configuration parameter (or else 10, if no value is set for the STATCHANGE parameter). If the SET ENVIRONMENT statement has set a different value for the current session, that value overrides the default or explicit STATCHANGE configuration parameter setting.

AUTO is the default setting in the CREATE TABLE statement, if you specify no explicit STATCHANGE setting.

For the table and index fragments for which data distribution statistics are already stored in the system catalog, the STATCHANGE setting specifies the percentage of rows in the fragment that have been deleted, inserted, or modified by DML operations since its distribution statistics were most recently updated. (This is the same significance that STATCHANGE has for table-level statistics.)

Automatic management of data distribution statistics

The HCL OneDB™ database server supports several mechanisms for automating some of the tasks that are involved in gathering, dropping, and refreshing data distribution statistics for tables, indexes, table fragments, and index fragments.

Automatic detection and refreshing of stale statistics during UPDATE STATISTICS operations

You can set the AUTO_STAT_MODE configuration parameter to enable the HCL OneDB™ database server to automatically detect which table and index statistics are stale, and only refresh the stale statistics when the UPDATE STATISTICS statement is run. The data distribution statistics that are automatically detected and refreshed are calculated at the table, fragment, or index level, not at the individual column level. If you set no value for this parameter, the automatic statistics mode is enabled by default. When automatic mode is enabled, the default threshold that defines stale statistics is reached when at least 10% of the rows in the table or fragment are changed by DML, LOAD, or TRUNCATE operations since the most recent calculation of data distribution statistics.

You can set another configuration parameter, STATCHANGE, to specify a nondefault change threshold for refreshing distribution statistics when automatic statistics mode is enabled. For example, if you set the STATCHANGE value to 15, statistics are refreshed if 15% of the rows in the table or fragment are changed. If the STATCHANGE parameter is not set, the system default value for STATCHANGE is 10.

You can override the STATCHANGE or AUTO_STAT_MODE configuration parameter setting for the current session by using the SET ENVIRONMENT statement to set session environment variables of the same names. The DBA can include SET ENVIRONMENT statements in the **sysdbopen** routine to enable or disable automatic statistics mode, or to change the stale distribution threshold (or both) at connection time. These settings are applied to UPDATE STATISTICS statements that are issued in the current session.

A table can be created with its own `STATCHANGE` table attribute, whose value overrides the setting of the `STATCHANGE` session environment variable or configuration parameter. For fragmented tables whose distribution statistics are calculated for each fragment, the value of its `STATCHANGE` attribute determines whether statistics are refreshed for individual fragments. The `ALTER TABLE` statement of SQL can reset the `STATCHANGE` attribute of a table.

You can also use (or disable for your current operation) the explicit or default `AUTO_STAT_MODE` and `STATCHANGE` settings during `UPDATE STATISTICS` statements that include the `AUTO` or the `FORCE` keyword:

AUTO

This keyword puts the `UPDATE STATISTICS` statement in automatic mode for detecting tables and fragments whose statistics are stale. Distribution statistics are not refreshed for tables or fragments whose `STATCHANGE` value is below the specified threshold.

FORCE

This keyword refreshes the statistics for all tables and columns within the specified scope. If automatic mode for detecting stale statistics is enabled, the `FORCE` keyword overrides automatic mode, so that values of the `STATCHANGE` attributes of tables and fragments are ignored, and statistics are recalculated for all database objects within the scope of the `FOR TABLE` specification.

The scope of `AUTO` or `FORCE` is limited to the `UPDATE STATISTICS` statement in which the keyword is specified. `UPDATE STATISTICS` statements that include neither of these keywords use the current `AUTO_STAT_MODE` setting of the database server (or for their session environment, if that is different). If `AUTO_STAT_MODE` is enabled, the `STATCHANGE` value is determined in the following (descending) order of precedence:

1. The value of the `STATCHANGE` attribute of the table, if `AUTO` is not the specified value.
2. The value that is set by the most recent `SET ENVIRONMENT STATCHANGE` statement in the same session.
3. The explicit setting of the `STATCHANGE` configuration parameter.
4. The system default `STATCHANGE` value is 10.

Automatic statistics maintenance in DDL operations

The HCL OneDB™ database server automatically creates, updates, or drops data distribution statistics during certain operations that create, alter, or destroy database objects.

ALTER FRAGMENT ATTACH operations

If the automatic mode for detecting stale distribution statistics is enabled, and the table that is being attached to has fragmented distribution statistics, the database server calculates the distribution statistics of the new fragment. Stale distribution statistics of existing fragments are also recalculated. This recalculation of fragment statistics runs in the background. After the database server calculates the fragment statistics, it merges them to form table distribution statistics, and stores the results in the system catalog.

Distribution statistics are not recalculated, however, unless explicit or default value of the `AUTO_STAT_MODE` configuration parameter or the `AUTO_STAT_MODE` session environment setting enables the automatic mode for detecting stale data distribution statistics.

ALTER FRAGMENT DETACH operations

Some ALTER FRAGMENT DETACH statements to attach a fragment can cause the database server to update the index structure. When an index is rebuilt in those cases, the database server also recalculates the associated column distributions, and these statistics are available to the query optimizer when it designs query plans for the table from which the fragment was detached:

- For an indexed column (or for a set of columns) on which ALTER FRAGMENT DETACH automatically rebuilds a B-tree index, the recalculated column distribution statistics are equivalent to distributions created by the UPDATE STATISTICS statement in HIGH mode.
- If the rebuilt index is not a B-tree index, the automatically recalculated statistics correspond to distributions created by the UPDATE STATISTICS statement in LOW mode.

If the automatic mode for detecting stale distribution statistics is enabled, and the table from which the fragment is being detached has fragment-level distribution statistics, the database server takes the following actions:

- Uses the distribution statistics of the detached fragment to form a new table distribution.
- Merges the distribution statistics of the remaining fragments to calculate distribution statistics for the surviving table
- Stores the statistics that result from these operations in the system catalog.

This recalculation of fragment statistics runs in the background.

Distribution statistics are not recalculated, however, unless an explicit or default value of the AUTO_STAT_MODE configuration parameter or the AUTO_STAT_MODE environment setting enables the automatic mode for detecting stale data distribution statistics.

ALTER TABLE ADD CONSTRAINT operations

ALTER TABLE ADD CONSTRAINT statements that use the Single Column Constraint format to implicitly create an index on a non-opaque column also automatically calculate the distribution of the specified column. Similarly, if the Multiple-Column Constraint format specifies a list of columns as the scope of the new constraint, the database server implicitly creates an index on the same non-opaque column or set of columns as the referential constraint, distribution statistics are automatically calculated on the specified column, or on the lead column of a multiple-column constraint.

These distribution statistics are available to the query optimizer when it designs query plans for the table on which the constraint is defined:

- For columns on which the new constraint is implemented as a B-tree index, the recalculated column distribution statistics are equivalent to distributions created by the UPDATE STATISTICS statement in HIGH mode.
- If the new constraint is not implemented as a B-tree index, the automatically recalculated statistics correspond to distributions created by the UPDATE STATISTICS statement in LOW mode.

These distribution statistics are available to the query optimizer when it designs query plans for the table on which the new constraint was created.

The automatic calculation of column distribution statistics in ALTER TABLE MODIFY operations that define a constraint on a non-opaque column is not dependent on whether AUTO_STAT_MODE is enabled or disabled.

ALTER TABLE MODIFY operations

ALTER TABLE MODIFY statements that use the Single Column Constraint format or Multiple Column Constraint format to define constraints similarly cause the database server to calculate data distribution statistics for the indexes that are implicitly created to enforce the constraints. These distribution statistics have the same attributes as the statistics that are calculated automatically for an index on a non-opaque column, and that are also automatically calculated during ALTER TABLE ADD CONSTRAINT operations. These statistics are available to the query optimizer when it designs query plans for the table on which the constraints are defined.

The automatic calculation of column distribution statistics in ALTER TABLE MODIFY operations that define a constraint on a non-opaque column is not dependent on whether AUTO_STAT_MODE is enabled or disabled.

CREATE INDEX operations

The database server automatically calculates index statistics, equivalent to the statistics gathered by UPDATE STATISTICS in LOW mode, when you create a B-tree index on a UDT column of an existing table, or if you create a functional index or a virtual index interface (VII) index on a column of an existing table. Statistics that are collected automatically by this feature are stored in the system catalog and are available to the query optimizer, without the need for running the UPDATE STATISTICS statement manually. When B-tree indexes are created, column statistics are collected on the first index column, equivalent to what UPDATE STATISTICS generates in HIGH mode, with a resolution is 1% for tables of fewer than a million rows, and 0.5% for larger tables. (Tables with more than 1 million rows have a better resolution because they have more bins for statistics.)

The automatic calculation of column distribution statistics in CREATE INDEX operations is not dependent on whether AUTO_STAT_MODE is enabled or disabled.

Auto Update Statistics (AUS) maintenance system

This uses a combination of Scheduler sensors, tasks, thresholds, and tables to evaluate and update data distribution statistics. The system provides as built-in input criteria a set of configuration parameter values. The system administrator can modify these to reflect current requirements and workloads. The AUS system combines these criteria with information from the **sysmaster** database to automatically identify tables whose distributions are becoming stale, and generates the text of UPDATE STATISTICS statements to refresh the distribution statistics for those tables.

The list of generated UPDATE STATISTICS statements is run automatically each week at a designated period of low throughput, to update as many table distributions as can be recalculated during the designated maintenance period. Any UPDATE STATISTICS statements that do not complete are retained on the list for the next maintenance period.

Index

A

- Access privileges 21
- ALTER FRAGMENT ADD CONSTRAINT statement 48
- ALTER FRAGMENT ATTACH statement 48
- ALTER FRAGMENT DETACH statement 48
- ALTER FRAGMENT statement 40
- ALTER TABLE MODIFY statement 48
- ALTER TABLE statement 46
- Attached index 21
- AUS_CHANGE configuration parameter 48
- Auto Update Statistics (AUS) maintenance system 48
- AUTO_STAT_MODE configuration parameter 45, 48

B

- business process defined 13

C

- CREATE DATABASE statement 34
 - dimensional data model 34
- CREATE INDEX statement 48
- CREATE TABLE statement 46
 - dimensional data model 34
- CREATE TABLE statements 34

D

- Data mart
 - description 3
- data modeling
 - dimensional 3
- Data models
 - dimensional 8, 12
- Data warehouse
 - denormalization 3
 - description 3
- Databases
 - demonstration
 - sales_demo 13
- Detached index 21
- Dimension table
 - changing dimensions 31
 - choosing attributes 19
 - description 12
- Dimension tables
 - creating 34
- dimensional data model
 - creating fact tables 34
- Dimensional data model
 - building 12
 - changing dimensions 31
 - creating dimension tables 34
 - denormalization 19
 - designing 8
 - dimension attributes 11
 - dimension elements 10
 - dimension tables 12
 - dimensions 10
 - fact table 9
 - implementing 34
 - measures, definition 9
 - minidimension tables 31
- Dimensional database 34, 34
 - loading data 38
 - loading from external tables 43
 - snowflake schema 32

- testing 39
- dimensional database model 3
- Dimensional databases
 - mapping data sources 36
- Dimensional table
 - identifying granularity 15
 - loading data 38
- Discretionary access privileges 21
- Distributed storage designs 21

E

- EXPLAIN_STAT configuration parameter 45

F

- Fact table
 - description 8, 9
 - determining granularity 14
 - dimensions 15
 - granularity 9
- Fact tables
 - creating 34
- Foreign key 18
- Fragment elimination 21
- Fragment expression 21
- Fragment key 21
- Fragment list 21
- Fragmentation key 18
- Fragmentation strategies
 - by expression 26
 - by interval 28
 - by list 27
 - by round-robin 24

G

- granularity
 - data distribution statistics 46
 - identifying dimensions 15
- Granularity, fact table 9

I

- I/O contention 21

L

- Logging table
 - creation 34

N

- Nonlogging tables
 - creation 34
- NULL fragment 27

O

- Operational data store
 - description 3

P

- Parallel-database queries (PDQ) 21
- Primary key 18

Q

- query execution plans
 - considerations 44
- query optimizer 44

R

- Range fragments 28
- Range interval distribution 28
- Referential constraints
 - foreign key 18
 - primary key 18
- REMAINDER fragment 27

S

- sales_demo database
 - creating 34
 - data model 13
 - data sources for 36
 - loading 38
- SBSPACENAME configuration parameter 45
- SET ENVIRONMENT AUTO_STAT_MODE statement 48
- SET ENVIRONMENT STATCHANGE statement 46, 48
- Snowflake schema
 - example 32
- Star schema
 - description 8
- STATCHANGE configuration parameter 45, 46, 48
- STATCHANGE table attribute 48
- STATCHANGE table property 46
- STATLEVEL table property 46
- SYSDISTRIB system catalog table 45
- SYSFRAGDIST system catalog table 45, 46
- SYSFRAGMENTS system catalog table 45
- SYSINDICES system catalog table 45
- SYSSBSPACENAME configuration parameter 46

T

- Transition fragment 21
- Transition value 21

U

- Unique constraints 18
- UPDATE STATISTICS statement 48
- Utility program
 - dbload 34