

## Reading Sample

*In Chapter 3, you'll find out how SAP HANA changes the way data is stored. Learn the technical details behind row and column storage, and get hands-on instructions for creating tables and data marts for an SAP HANA project.*



**"Data Storage in SAP HANA"**



**Contents**



**Index**



**The Authors**

Jonathan Haun, Chris Hickman, Don Loden, Roy Wells

### Implementing SAP HANA

860 Pages, 2015, \$79.95/€79.95

ISBN 978-1-4932-1176-0



[www.sap-press.com/3703](http://www.sap-press.com/3703)

*This chapter helps you understand how data is stored most effectively in memory so you can get the best results in both compression and performance.*

## 3 Data Storage in SAP HANA

In this chapter, we'll go into great detail on how data is stored in SAP HANA. Understanding data storage in SAP HANA is an important foundation because data storage differs from traditional database management systems in a number of ways. First, we'll start with an overview of data storage in SAP HANA to highlight these differences, and then we'll move into all of the components that make this possible (Section 3.1 and Section 3.2, respectively). We'll then discuss physical data modeling for SAP HANA in Section 3.3 to draw clear differences between traditional database systems and techniques and tools that are available in SAP HANA, and why it makes sense to actually think backward about a data model in certain cases. This chapter ends in Section 3.4 with a case study for data modeling using our sample organization, AdventureWorks Cycle Company.

### 3.1 OLAP and OLTP Data Storage

Storing data in SAP HANA is quite different from doing so in a traditional disk-based database. The first and most obvious point is that SAP HANA is a *relational database management system* (RDBMS), where data is stored entirely in memory, instead of relational data being stored entirely on spinning disks.

Storing data entirely in memory was once a revolutionary concept that first had its detractors making statements such as, "Data for an entire application or data warehouse structure would never all fit into memory." In fact, it was such an unconventional idea that it took some time to gain ground. However, many leading vendors now have in-memory solutions and are touting both the in-memory platform and stance for the same reason SAP sought to use this strategy in the first place—unbelievable performance. Data loaded into SAP HANA and consumed by

external applications performs at an unbelievable speed—almost as if the data were staged for a demonstration. The response times are simply too fast.

SAP HANA Real-World Performance: Exhibit A

In our lab at Decision First Technologies, we took data from a customer paired with the SQL produced by an SAP BusinessObjects Web Intelligence report and placed the supporting data in SAP HANA. We then took the underlying query provided by the Web Intelligence report and ran it at the command line against the SQL Server database. The original SQL Server–based query runtime? More than an hour. The query was tuned, and the data was optimized in the SQL Server database, but the query was, frankly, quite complex, and the data volume was large. The report was critical to the customer’s business, so more than an hour of runtime was simply too long to wait for the data.

As a proof of concept, we moved the data to SAP HANA for the customer, used the same exact SQL from the Web Intelligence report. We did not tune the database tables or structures for SAP HANA; we merely ported the data from SQL Server to SAP HANA. We did not tune the query. This was simply a copy-and-paste exercise. The new SAP HANA query runtime? Four seconds.

Although we did absolutely nothing to the data or the report, the runtime was immediate. Needless to say, this was a compelling story for the customer, even before we invoked the modeling techniques that exploit the storage and engine processes in SAP HANA (we’ll discuss these later in this chapter).

The example in the preceding box is a real-world result that this particular customer would benefit from immediately just by porting its data to SAP HANA. These are the incredible performance benefits of in-memory computing that SAP has not been shy about touting—and rightfully so.

However, as with any great software platform, a developer must consider the needs of the platform and embrace techniques that envelop all of its strengths. This is where a gap has existed in the discussion of SAP HANA. SAP HANA simply performs so well that it allows some sloppiness in the design and still performs at an incredible pace. We believe that you can avoid this sloppiness by merely taking a step back and catering the pillars of the development effort to the needs and special characteristics native to the SAP HANA platform. As you weigh design considerations at the onset of the project, begin by considering how you want to store the data in the architecture that is unique to SAP HANA. In this section, we’ll prepare you for these considerations by introducing you to the *spinning disk* problem, and then talk about how this problem can be combated with some of the unique features that SAP HANA brings to the development effort.

3.1.1 The Spinning Disk Problem

Spinning disks have been a performance bottleneck ever since they were introduced. The closer the disk is to the CPU, the faster data is rendered, searched, sorted, and processed; in SAP HANA, you take the physically spinning disk completely out of the equation to fully maximize this realization. Take, for instance, the process flow of information in a typical system and database:

1. Data is collected from an application via user input from a screen or form.
2. Data is passed to the database in a process known as an input/output (or I/O) transfer of information.
3. Data may be written to or read from a cache in memory on the database server.
4. Data is finally stored on a spinning disk.

I/O transfers performed without the use of a cache can take much longer to process. Factors that contribute to extra time include physical disk platter spinning rates, time needed to move mechanical components of the drive heads to read the disk platter, and numerous other factors that are inherent to this disk-based process and that add additional latency. This is a rather archaic process that hasn’t changed greatly since the onset of computing. Conventional database systems try to improve on this by targeting specific systems that provide disk caching controllers.

*Caching data* is a method used to speed up this process of data access from a spinning disk, and all of the major database vendors work closely with the disk manufacturers to tune the needs and specific requirements of the database I/O processing needs. In most cases, the database vendors seek to exploit caching techniques to limit that final disk interaction as much as possible. This is simply to avoid the native issues present with disk seek and write times by using the various optimizations of the caching controllers. This is all an effort to work around the slowness of the disk, whose performance can be maximized only so far.

3.1.2 Combating the Problem

Many technologies that we rely on today were invented to work around the inherent slowness caused by the disk. Take, for instance, *online analytical processing* (OLAP) technologies (which enable faster read performance by physically restructuring the data), *online transaction processing* (OLTP) technologies (whose goal is to make writing data to disk as fast as possible), or, finally, column storage

technologies (whose goal is compression to both minimize storage and increase the speed of access to the data). The important thing to keep in mind is that all of these technologies, at their core, were designed around the spinning disk and its native challenges. We'll introduce each of these technologies briefly and then talk about how they all fit into SAP HANA.

OLTP Storage Methods

An OLTP, or relational database, stores data in a normalized fashion at its core. Data is normalized to reduce redundant data and data storage patterns to optimize precious disk space and make the writing of that data to disk as fast as possible. Without techniques to minimize the storage factor, relational databases, by nature, use lots of space to store these redundant values. Consider Figure 3.1, which shows a typical normalized RDBMS table structure that's been designed to reduce redundant data storage.

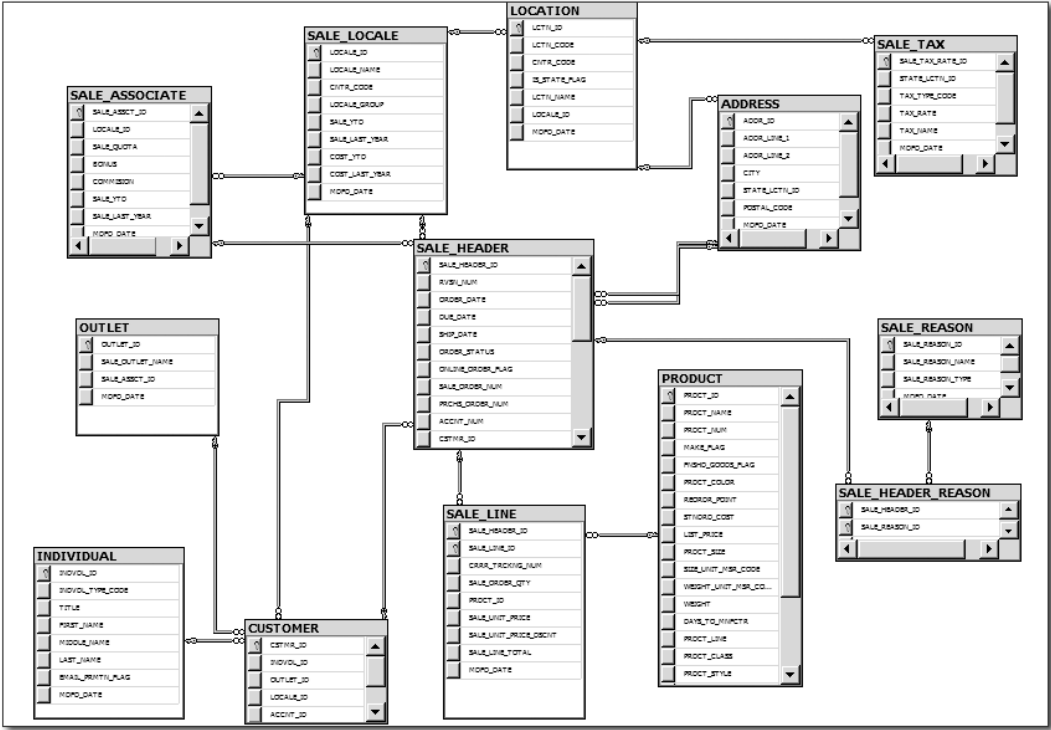


Figure 3.1 Normalized RDBMS Table Structure

Data is normalized or reduced into multiple tables so that repeating values are removed into multiple tables to store repeating values once and contain a pointer to those repeating values. For example, in Figure 3.1, SALE\_HEADER records are normalized into their own table instead of just storing the columns into the SALE\_HEADER table. This concept is the pinnacle of an OLTP system. This is simply the design principal on which OLTP systems are based.

There is nothing wrong with this design for inserting or storing data in *conventional* RDBMS systems. In fact, for this purpose, it's quite good. (There is a reason this methodology is the way the world stores its data!) However, there is one fundamental problem with this system: getting data out.

Retrieving data from an OLTP system requires multiple joins and combinations of various related tables. This is expensive in terms of processing in these database designs. Often, reporting in these systems is certainly an afterthought. It is problems like this one—combined with the slowness and natural speed impediment—that many technologies evolve to solve. Techniques such as OLAP technologies were invented to solve this problem.

OLAP Storage Methods

OLAP data storage methods were conceived to combat slowness caused by both data access to disk and the way that data was stored in conventional relational databases, as just described. Technologies such as OLAP data storage physically store the data in a different way because traversing a relational database on disk isn't exactly the fastest solution for reading or retrieving data. Figure 3.2 shows this alternative data storage in an OLAP database, in a typical *star schema* (named so because of the shape the related tables resemble).

In an OLAP database, data is organized into concepts called *facts* and *dimensions*. The facts and dimensions are just standard tables, but their names denote what they store. Facts are the heart of the star schema or dimensional data model. For example, FACT\_SALE is the fact table in Figure 3.2. Fact tables store all of the measures or values that will be used as metrics to measure or describe facts about a business concept. Fact tables may also contain foreign keys to the date dimension tables to allow pivoting or complex date metrics. Fact tables will be arranged with differing granularities. Fact tables could have a high granularity and be at an aggregate level, aggregating measures by calendar week or a product line, for instance, or a fact table could be at the lowest level of granularity: a transaction

line from a source system or combined source systems. Fact tables also contain foreign keys that refer back to dimension tables by the primary key of the dimension table. A fact is the “many” side of the relationship.

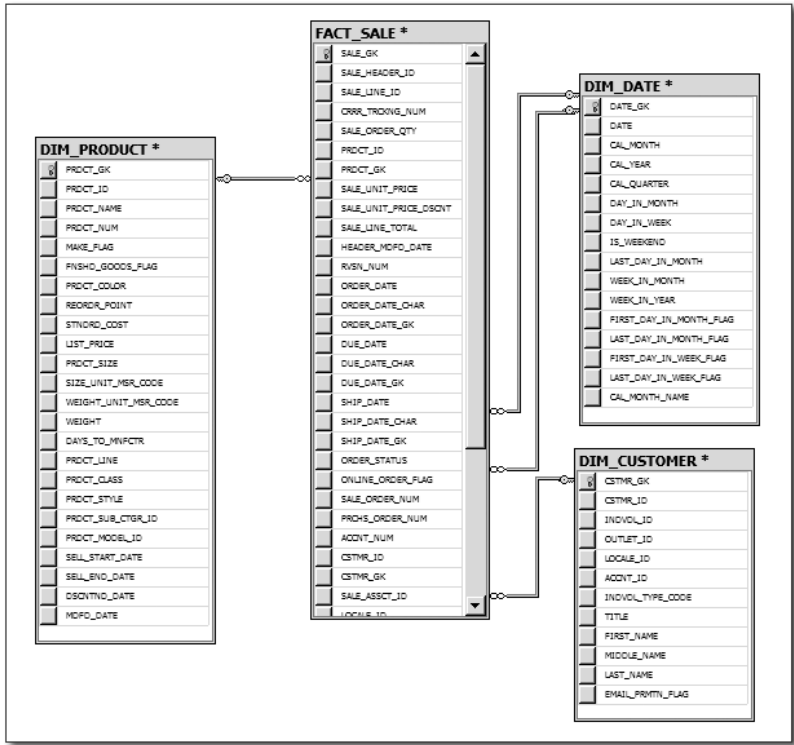


Figure 3.2 Data Stored in an OLAP Database

Dimension tables are the ancillary tables prefixed with “DIM\_” in Figure 3.2. Dimension tables are somewhat the opposite of fact tables because dimensions contain descriptions of the measures in the form of accompanying text to describe the data set for analysis by labeling the data, or the dimensions are often used to query or filter the data quickly. In Figure 3.2, the DIM\_CUSTOMER table provides details about customer data or attributes and is used to filter and query sales from the prospect of customer data. The same can be said for DIM\_PRODUCT.

This is a dramatic solution because an entirely different table structure had to be established and created. If the modeling task symbolized in Figure 3.2 isn’t

enough, another element adds to the complexity: a batch-based process is created out of necessity.

A batch-based process is needed to both load and transform the data from the OLTP normalized data structure into the denormalized OLAP structure needed for fast querying. That batch process is typically called *extract, transform, and load* (ETL). An ETL process physically transforms the data to conform to this OLAP storage method.

Typical ETL Process Workflow

1.

After data is extracted from one or multiple source systems, the data loads to a staging database, where multiple transformations occur.

2.

Staging is a key layer where the data loses the mark of the source system and is standardized into business concepts.

3.

Data is loaded into the data warehouse tables and finalized into an OLAP structure to allow for both high-performing reads and flexibility in analytical methods and ad hoc data access.

SAP’s solution for ETL data integration is SAP Data Services. SAP Data Services is a fully functional ETL and data quality solution that makes building very complex processes relatively straightforward. SAP Data Services is used to extract and transform data through complex transforms with many powerful, built-in functions. Because it’s the primary means to provision non-SAP data into SAP HANA, SAP Data Services plays a pivotal role in setting up data models and data storage the right way for maximum performance in SAP HANA. We’ll discuss this tool’s capabilities at length later in this book.

OLAP data structures like those shown in Figure 3.2 are the focus and primary use case of ad hoc analytical tools such as SAP BusinessObjects BI. The OLAP or star schema structure allows the tool to be quite flexible with the data in terms of drilling if hierarchies exist in the data or if you are using a date dimension (in the preceding example, this is DIM\_DATE) to not only search and sort but also effortlessly provide running calculations or more complex, date-based aggregations. Analytic activities like these would be quite difficult to address in a fully normalized OLTP system. Certainly, this data storage and system design eases the burden placed by the slowness of the platform, as well as adding nice features for analytics.



Columnar Data Storage Technologies

One final data storage technology, and the one most relevant to SAP HANA, is the *columnar database architecture*. Columnar databases also take on the problem of working around the slowness of the disk by changing the way that data is stored on the disk. We'll walk through the SAP HANA specifics later in this chapter, but it's important to have a basic understanding of columnar architectures now.

Columnar databases have been around for quite some time, and the concept was certainly not invented with SAP HANA. Rather, this was a design methodology that was integrated into SAP HANA for the unique features and data storage aspects that a columnar database brings to the data storage equation. Columnar databases still use tables to store data as logical components, but the way that the data is laid out on the disk differs considerably from standard, row-store tables. Data values are gathered and stored in columns instead of rows. A very simple example is a product table with colors and product descriptions. In Figure 3.3, the data is stored in rows as it's represented in the logical tables in the database.



Figure 3.3 Data Storage in Rows in a Table

Data is organized into rows in the physical storage of the data on disk. This is a great design for OLTP systems and is the standard for most of the world's data. So, data in a column-store table would be arranged quite differently on the disk. Data is arranged by the columns of the data in Figure 3.4.




Figure 3.4 Data Stored as Columns in a Column-Store Table

Notice that the repeating values are stored only once, to minimize the physical footprint of the data on the disk.

Note

Column-store tables can still be relational tables and data. The difference lies in the way the data is arranged on the disk.

Columnar databases have traditionally been used for OLAP applications, wherein reads are very important because data can be read much more efficiently from this type of storage structure. Data is read and presented quickly to a consuming application or report. Other challenges can arise when you insert data into disk-based column-store tables. For example, `UPDATE` operations are quite expensive for column-store data structures compared to their row-store cousins.

Inserts in Disk-Based Column-Store Tables

In our lab at Decision First Technologies, we recently ported data for a data warehouse OLAP structure from SQL Server to SAP (Sybase) IQ to take advantage of the superior compression and read technology available in columnar SAP (Sybase) IQ tables. However, we did notice some considerations that should be made in this port. These considerations are somewhat alleviated by the in-memory storage in SAP HANA, but they are still worth considering because they are in the domain of a column-based database:

- ▶ `SELECT` statements or reads are much faster than with a conventional row-based database. The data then loads to a staging database, where multiple transformations occur.
- ▶ Using `BULK INSERT` uploading data is considerably faster and should be used whenever possible, especially with large record sets.
- ▶ `UPDATES` or `MERGE` target operations are considerably slower than a conventional row-based database.
- ▶ `DELETE` inserts are faster when updates are needed.

The main takeaway is that `SELECT` SQL statements or reading data for operations such as a report do not need to be altered too much, but the ETL process will most likely require `INSERTS`, `UPDATES`, and `DELETES` to be altered, especially for delta or change-based loads.

For reasons like these, porting an existing structure to a columnar form—while not an insurmountable task—certainly has more considerations than simply moving the data over to a different platform. As mentioned, SAP HANA mitigates some of these issues because in memory storage is so much faster. In a sense, SAP HANA masks some of these issues, but you should still consider them when

you're porting very large existing data warehouse structures that require some type of ETL process with, most often, non-SAP data.

### Solutions Used by SAP HANA

We've discussed OLTP, OLAP, and columnar data storage methods and the reasons they were introduced, and SAP HANA is unique in the sense that it can be a bit of a chameleon. SAP HANA can act as any of these platforms by first physically storing data in both row and column fashions; however, even more than that, it can also act as an OLAP structure and even process data by interpreting multi-dimensional expressions (MDX query language). It also has a common and conventional SQL interface.

In essence, SAP takes advantage of the best of all of these platforms natively. This adaptable nature has been great for SAP because it allows SAP HANA to quickly and seamlessly be addressed under many conventional applications. If a multidimensional, cube-based application, such as SAP BW or SAP Business Planning and Consolidation (SAP BPC), needs MDX to interface data, then no problem. SAP HANA has an interface layer to behave just like a cube to the application. Most applications interact with a database via SQL, and SAP HANA is just as comfortable interfacing as SQL.

It's important to note that most of these technologies were invented to combat the slowness of disk-based data access. But SAP HANA is different. Even though it can masquerade as any of these technologies, it's taking on the performance problems directly. Instead of working around the issues of the platform, SAP HANA simply changes the platform altogether. It skips the disk, and data is stored directly in memory close to the CPU, where it performs better. That SAP HANA works natively as any of these technologies is merely a product-related strategy to foster adoption of SAP HANA as a platform capable of replacing existing, underlying database technologies while offering developers new and exciting ways to both access and model data. We'll cover both accessing and modeling data in later chapters of this book.

SAP HANA goes even further in rethinking the developer's platform by moving the various data processing layers in an application so that a developer must re-envision what he or she is trying to achieve. It's truly a revolutionary platform.

## 3.2 Data Storage Components

To begin using SAP HANA, you must first load or provision your data into SAP HANA, but to do this, you need a *persistent layer of data storage*. This persistent layer (also known as a persistent model) is made up of basic data storage and organizational components that are actually quite common concepts to database-savvy professionals. The first two organizational components are schemas and users. From there, the components start to diverge and take on a much more SAP HANA-specific dialect: *row-store tables* and *column-store tables*.

Let's wade further into the organizational components mentioned above: schemas and users, column-store tables, and row-store tables. We'll conclude our discussion with a comparison of use cases for row- and column-store tables. All of the storage components mentioned in this chapter are found in SAP HANA Studio under the ADMINISTRATION CONSOLE and MODELER perspectives.

### 3.2.1 Schemas and Users

Recall that SAP HANA has many conventional components that make database administrators and database developers quickly feel at home. These are mostly organizational components that facilitate administrative tasks. At a very high level, a *user* is used to connect to and authenticate SAP HANA, and a *schema* is used to group and classify various database objects, such as tables or views. Because these aren't new concepts for SAP HANA, we will assume basic knowledge of what they mean, and will instead focus our discussion on what is required to provide a foundation for further discussion of SAP HANA-specific topics and for building the physical database-level objects for the case study examples.

#### Schemas

Schemas are similar to concepts that exist in other conventional database platforms. Most database platforms use schemas as a subdividing tool, and SAP HANA is no exception. In SAP HANA, schemas are used to divide the larger database installation into multiple sub-databases to organize the database objects into logical groupings. You use schemas to logically group objects such as tables, views, and stored procedures. A schema in SAP HANA is essentially a database within the larger database or catalog. (We'll go into specific details about how to create a schema in Section 3.4.1, as part of the case study for this chapter.)

Figure 3.5 shows the BOOK\_USER schema in SAP HANA, from which all of the case study examples in this book will be crafted. The BOOK\_USER schema is the only *user-defined* schema that is visible. The rest of the schemas visible in the figure—SYS, \_SYS\_BI, \_SYS\_BIC, and \_SYS\_REPO—are all default *system-generated* schemas.

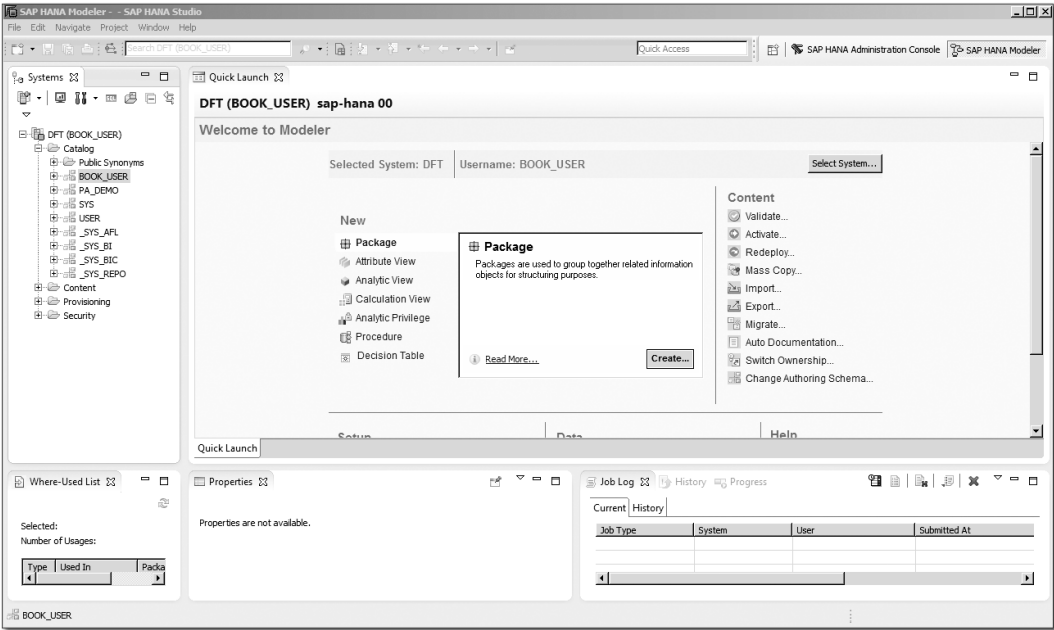


Figure 3.5 The BOOK\_USER Schema in SAP HANA

Users and User Authentication

SAP HANA users are no different from users in any other conventional database in the sense that, if you want to work in SAP HANA, you must have a user name to log on to the system. After logging on to SAP HANA, your user must have privileges to perform certain tasks. Much like schemas, users feel quite standard in concept to most savvy database administrators.

SAP HANA also supports the concept of a *role*, which is a superset of privileges. Roles are granted to database users and inherit the privileges assigned to the role the user belongs to.

When SAP HANA is installed, a database user called SYSTEM is created as the default admin user. This user has superior system-level privileges to create users, access system tables, and so on. As a best practice, you should not use the system

user for normal administration activities or assign roles to this user. Use the SYSTEM user to create database users with roles with the minimum set of responsibilities to perform the user's duties.

Operating System Administrator User

Aside from the SYSTEM database user, it's also important to note that an operating system administrator user (<sid>adm) is also created on the SAP HANA system upon installation. This user exists to provide a context or linkage to the base operating system in SAP HANA.

This user has unlimited access to all local system resources and owns all SAP HANA files and all operating system processes. Within SAP HANA Studio, this user's credentials are required to perform advanced operations, such as stopping or starting a database process or executing a recovery. This isn't to be confused with a database user because the <sid>adm user is concerned with the operating system on only the SAP HANA machine.

Users in SAP HANA exist only as database users to map to the privileges discussed earlier, and for internal authentication, this is the only means available.

Additional References

For additional information about user authorizations, roles, and best practices on SAP HANA security, please consult Chapter 2 and Chapter 11 of this book.

3.2.2 Column-Store Tables

Because SAP HANA is optimized, or tuned, for storing data in columns over storing data in rows, you should use *column-store tables* whenever possible. Reading data is much faster in column-based tables; from a data storage perspective, columnar storage and compression are two of SAP HANA's best offerings. In a column-store table, data simply compresses at higher rates.

As discussed earlier in this chapter, columnar storage allows repeating values to be expressed only once in storage, which allows the physical storage required to compress. In SAP HANA, this compression is due to run-length encoding or the storage of sorted data where there is a high probability of two or more values being stored contiguously or in the same spatial locale. Run-length encoding counts the repeating values as the same value, which is achieved by storing the original column as a two-column list.

This sophisticated system of reducing redundancy is an important concept of column-based storage for financial reasons. SAP HANA is licensed and priced by



memory blocks, so the more memory you need to store your data, the more expensive your SAP HANA solution will be. However, pricing and cost are only one side of the equation.

Compression is also an important aspect of high-performing queries in SAP HANA. When data is compressed, it can be loaded into the CPU cache faster. The limiting factor is the distance between memory and the CPU cache, and this performance gain will exceed the additional computing needed to decompress the data. One factor that enables compression is run-length encoding, which stores values as a two-column list, while repeated values are stored only once in one column, with another column as an index or pointer to the repetitious storage. One would think this would cause a latency in performance, but the two-column list's equality check on the index column is based on a higher-performing integer value for the equality comparison—which is why proper compression can speed up aggregate operations or table scans considerably. These are the operations that stand to benefit the most from compressed data in SAP HANA.

It's easy to create a table as a column-store table in SAP HANA. To create a column-store table, just use the ADMINISTRATION CONSOLE perspective in SAP HANA Studio (as shown in Figure 3.6), and select COLUMN STORE under the TYPE menu. Now, you have a column-store table that is ready for use!

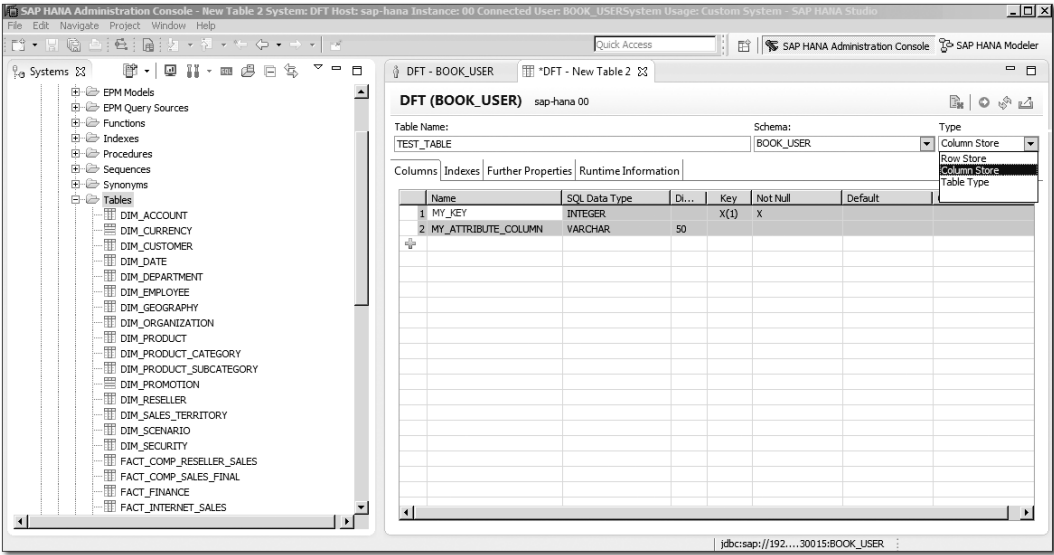


Figure 3.6 Creating a Column-Store Table

When you're deciding between a row- and column-store table, consider how the data is going to be used. For example, column-store tables are a good choice because some features, such as partitioning, are available to only column-based tables. So if partitioning is required in your application, your decision of whether to use column- or row-based tables has already been made.

You should also weigh column-based storage in terms of updates and inserts. Bulk updates, or bulk operations in general, perform well against large tables with column storage. Column-store tables are great choices for large tables with lots of read-based operations or `SELECT` statements—especially when you're performing aggregate operations. A number of aggregate functions exist natively in the column engine. Consider the list of SAP HANA functions that are available as native column functions by using column engine expressions as arguments. Thus, columnar tables simply perform better because they're able to use functions built directly into this column engine rather than having to switch the processing and physically move the data to the row engine.

The following functions use column-engine expressions as arguments:

- ▶ **Numeric functions:** `TO_DECIMAL`, `TO_NUMBER`, `TO_BIGINT`, `TO_REAL`, `TO_DOUBLE`, `TO_CHAR`, `TO_NCHAR`, `TO_DATE`, `TO_TIMESTAMP`, and `BINTOHEX/HEXTOBIN`.
- ▶ **String functions:** `LTRIM`, `RTRIM`, `TRIM`, `LENGTH`, `SUBSTR`, `INSTR`, and `LOWER`, `UPPER`.
- ▶ **Date and time functions:** `WEEKDAY`, `DAYS_BETWEEN`, `SECONDS_BETWEEN`, `ADD_DAYS`, `UTCTOLOCAL`, `LOCALTOUTC`, `ISOWEEK`, and `QUARTER`.
- ▶ **Mathematical functions:** `LN`, `LOG`, `EXP`, `POWER`, `SQRT`, `SIN`, `COS`, `TAN`, `ASIN`, `ACOS`, `ATAN`, `SINH`, `COSH`, `FLOOR`, and `CEIL`.
- ▶ **Logic driving functions:** `NULLIF`, `NVL`, `NVL2`, and `COALESCE`.
- ▶ **Date extract function:** `EXTRACT ( YEAR /MONTH FROM <column engine expression> )*`.

Three more specific advantages to column-store tables will never be achieved in row-store tables. The first of these advantages is that columnar storage with proper compression eliminates the need for additional indexing. The columnar scans of the column-store tables, especially for run-length encoded tables, allow very high-performing reads. In most cases, these reads are fast enough that an index, with its additional overhead of metadata in terms of both storage and maintenance, is simply not necessary. It's basically an obsolete concept for a column-store table in many cases. Without having a need to index, not only does SAP

HANA gain storage due to compression, but you also don't need to account for extra storage space or time in terms of jobs and scheduled offline tasks necessary to maintain indexes to speed data retrieval as you would in a conventional database. In a sense, you're actually gaining performance while simplifying the physical model because you don't have to maintain separate index structures.

The second advantage is that the nature of the column-store structure makes parallelization possible; that is, data is already vertically partitioned into columns. With that partitioned structure of divided columns, SAP HANA easily allows parallel operations to be assigned to multiple cores on the CPU. This way, multiple columns can be searched and aggregated at the same time by different cores. The portioning that requires extra thought and maintenance—much like the indexing structures—is both redundant and unnecessary with column-store tables and column-engine processing in SAP HANA.

The final advantage is the elimination of physical aggregate data structures. Traditional BI applications and designs often call for aggregation in the database models at the presentation layer simply to deal with reporting or retrieving data against large and cumbersome record sets. This is often to work around the fact that the platform and disk-based data access bind I/O operations and simply prove negative performance implications when performing complex aggregations or queries across larger data sets. To solve this problem in a traditional RDBMS, data is physically persisted into aggregate tables that roll up the data to a higher level of granularity.

In Figure 3.7, we see an example of an aggregate table where transaction-level sales data has been aggregated to raise the granularity of the data to records totaled by period, year, and product. This table would need to be created for analysis in a traditional RDBMS if the sales transaction table contained lots of history and the analysis was mostly done at the year level of granularity. This would eliminate the performance problem while still addressing the reporting need.

	PERIOD	YEAR	PRODUCT	PERIOD_QTY	PERIOD_SOLD...
►	1	2013	BLUE WAGON	25	250.00
	2	2013	RED WAGON	30	300.00
	12	2012	RED WAGON	20	200.00
*	NULL	NULL	NULL	NULL	NULL

Figure 3.7 Example of an Aggregate Table

Deriving this aggregate table is relatively straightforward; it's just a SUM of the quantity and amount column in the transactional source. This means that

```
Select PERIOD, YEAR, PRODUCT, QTY, SOLD_AMT From Table_A
```

would become

```
Select PERIOD, YEAR, PRODUCT, SUM(QTY), SUM(SOLD_AMT)
From Table_A
Group By PERIOD, YEAR, PRODUCT
```

This is a very simple example; the logic from moving from transactional granularity to an aggregate byproduct isn't terribly difficult to derive or design. However, you would need an ETL process to physically transform the structure and move the data over to this new structure. So, even with this one simple example, we've added quite a bit of complexity in terms of more data and more processes to be maintained.

On top of this complexity, this model introduces another problem: inherent latency. The data in the aggregate will never be real time because the aggregate will be handled by either an ETL process (by definition, a batch-based process) or the database layer (which may introduce concurrency issues with updates in terms of locking operations that could potentially block reads during rebuilds). So the important point to take away about a column-store table in SAP HANA operating using column-engine native functions is that *it isn't necessary!*

This layer can be completely removed. SAP HANA can scan the data and perform the simple or complex aggregation at runtime in memory with similar speeds as a conventional architecture performing against aggregates. This is all happening in real time against the base transaction-level data; there is no need to have a latent, batch-based process. When you have this level of performance natively, you simply don't need these additional layers. Because the data has not persisted, storage needs and costs actually diminish with the support of these column-store structures in SAP HANA.

This is a very simple example, but you can see how this might grow as the needs for multiple views of aggregated data produce more duplicated, redundant data with more processes to maintain. By removing these layers, you dramatically simplify the data model, thus simplifying the interaction of querying the data. With this single-layer model, there is no need to hop from reading an aggregate view of the data to reading the base transactional view of the data. You use the same SQL from the clause and base statement and add in function calls when necessary. This type of simplification is a major benefit of using SAP HANA and one of the ways SAP HANA is transforming the data landscape.

3.2.3 Row-Store Tables

Row-store tables are exactly what they sound like; data is stored in memory but in a row fashion. Because these tables, at the base storage level, are very similar to traditional database structures and constructs found in conventional databases, we won't go into the level of detail in this book to discuss row base-level components and data storage methodology as we did with column-store tables.

However, one item to pay particular attention to with row-store tables is that there is virtually no additional compression occurs when using a row-store table. So, what is a proper use case for a row-store table?

Row-store tables were included in the SAP HANA platform to first and foremost offer the ability for SAP HANA to be used as a valid and suitable OLTP platform (i.e., as a basis for SAP Business Suite). A large part of enabling that possibility is that row-store tables and a row engine exist to process row-by-row data access requests.

The backbone of any OLTP system that involves data entry is rapid, row-by-row access to complete or mostly complete records. These aren't cases in which one SQL statement is returning, parsing, and aggregating millions of records on just a few columns. An OLTP design requires one customer record to be looked up and written into the application layer quickly, in real time, while a sales transaction is being established in the system. This response time needs to be instantaneous, and, in most cases, the entire row of the record is needed to satisfy the application.

This type of data access is effectively the complete opposite of the OLAP style of churning through complex data sets to group, sort, and aggregate on just a few columns. Because of needs like this, SAP needed to include both platforms and engines. This inclusion of both sides (both row and column) of the data processing house makes SAP HANA truly unique, and presenting a viable row-store option fosters rapid adoption of SAP HANA as much more than a valid BI- or OLAP-serving platform. By serving the row needs, SAP HANA is the new, remarkable, multifaceted platform built and scaled to handle complex and sizable applications, such as the SAP Business Suite.

In Summary

In short, use a row-store table if you're developing a transactional interactive system, such as a row-based system or an OLTP design. Row-store tables will suit this purpose well. The bottom line is:

- ▶ If your table will be used mostly for getting data in through a user input-driven design, use a row-store table.
- ▶ If your table will mostly be used for retrieval or aggregate-based operations, use a column-store table.

It's easy to create a table as a row-store table in SAP HANA Studio. The process is much like the one outlined earlier to create a column-store table. Just use the ADMINISTRATION CONSOLE perspective, and select ROW STORE under the TYPE menu, as shown in Figure 3.8. After performing this step, you now have a row-store table that is ready to use.

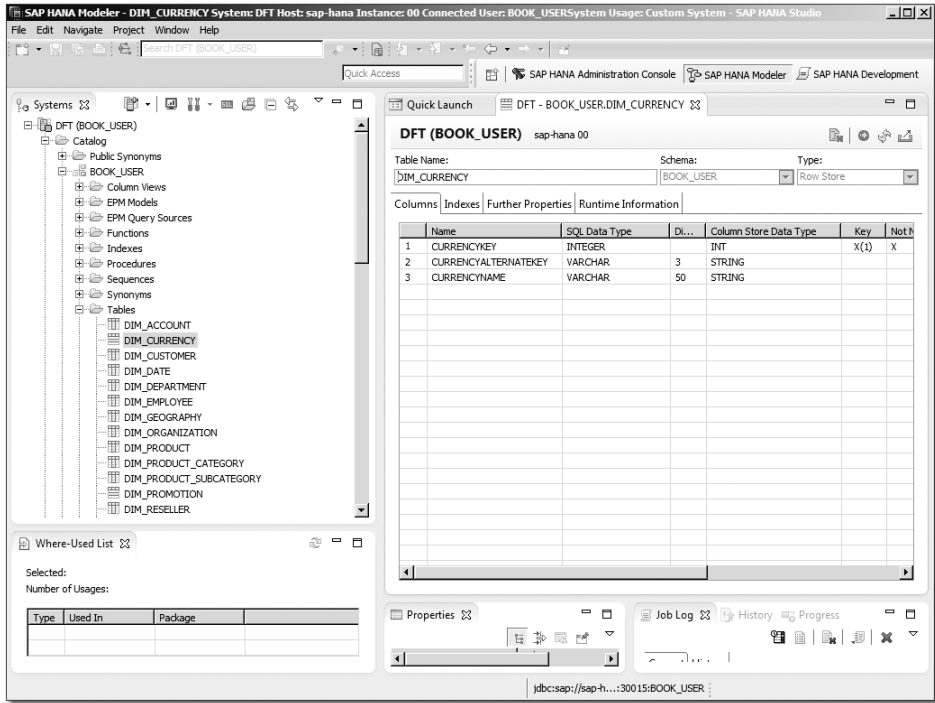


Figure 3.8 Row-Store Table in SAP HANA Studio

3.2.4 Use Cases for Both Row- and Column-Store Tables

Because they are primarily suited for most tasks in SAP HANA, column-store tables are generally the reflexive first choice for an application developer. However, as shown earlier, row stores certainly have their place for developers, as

well. Though we've already touched upon some of the reasons for row- vs. column-storage, we'll conclude this section with a succinct list that will help you decide the correct type of table to create.

If you find yourself in the following scenarios, use column-based storage:

- ▶ Tables are very large in size, and bulk operations will be used against the data. There are two primary examples that fit this scenario:
  - ▶ Data warehouse tables
  - ▶ Historical tables with large record sets
- ▶ Data is primarily staged for reads or `SELECT` statements. There are two primary examples that fit this scenario:
  - ▶ Data warehouse tables or data mart tables for BI reporting
  - ▶ Application-based tables that will serve as the basis for reports or getting data out
- ▶ Aggregate functions will be used on a small number of columns with each `SELECT` or read operation.
- ▶ Table will be searched by values in one or a few columns.
- ▶ Table will be used to perform operations that require scanning the entire set of data, such as average calculations. Searches like this are quite slow, even with proper indexing in conventional or row-based structures. The columnar constructs of SAP HANA are quite good at this type of analysis.
- ▶ High compression can be achieved by large tables that contain columns with few distinct values in relation to the record count.
- ▶ Complex aggregations or calculations will be needed often on one or a few columns in the table.

If you find yourself in the following scenarios, use row-based storage:

- ▶ Table is relatively small in size or record count, making low compression rates less of an issue.
- ▶ Table will be used for returning one record quite often. A classic use case for this is an OLTP application table. This is probably the most important point and will ultimately be the best overall use case.
- ▶ Row-store tables in SAP HANA will be the backbone of the OLTP application base.
- ▶ Aggregations aren't required.

- ▶ Writing data one record at a time is required.
- ▶ Fast searching isn't required.

When considering these criteria, you'll notice clear patterns that emerge regarding which type of data is best for each storage method. If your application requires record-by-record OLTP-style data interaction, you'll need to use row-based tables. Be cautious with these tables because when they become large, they offer virtually no compression. This will bloat the licensed memory required to store the data. Column-based storage is best used for applications that have many complex read-based or `SELECT` operations, such as OLAP or data warehousing structures. Column-based table structures compress nicely, and properly modeled physical data structures will take advantage of all of the sophisticated functions that are available only to column-store tables.

### 3.3 Modeling Tables and Data Marts

When considering modeling data for SAP HANA, we'll limit our discussion to modeling the data needed to fuel and power the column-store tables and engine for maximum performance and processing. To examine row-store data modeling in this book would overlap too much with conventional data modeling books because modeling data in a row-store table follows a conventional normalized playbook. The column-store tables and the compression that is offered in the SAP HANA platform are what expand this playbook into something that exists outside of the conventional normalized data constructs.

The SAP HANA data modeling playbook offers ideas that initially seem contrary to conventional data logic and wisdom. However, this is with good reason. It's only when considering the SAP HANA platform and storage paradigm, as discussed in detail earlier in this chapter (Section 3.2.2), that these ideas begin to converge, resonate, and ultimately become conventional.

In this section, we'll review the modeling of tables and data marts that take advantage of SAP HANA-specific functionality that ultimately prepares SAP HANA for any type of data consumption. We'll start with modeling for a traditional OLAP approach and then see how this evolves for SAP HANA. We'll then move on to a discussion of how to denormalize data, which is an especially important part of the data modeling process for SAP HANA.



3.3.1 Legacy Relational OLAP Modeling

To compare and contrast data modeling techniques for SAP HANA, it's valuable to start with a basic understanding of legacy relational data modeling techniques. Legacy relational OLAP modeling is when data is arranged into a series of tables of both facts and dimensions for performance for reporting, as well as to organize data effectively into data marts. A *data mart* is a collection of one or more relational OLAP fact tables and dimensions that are unified in purpose. For this chapter, we'll use two example data marts: the first for sales data and the second for financial data.

Both structures are simplistic in nature; it's plain to see that the focal point of the design resides around speed of access to the data. The fact table in each case contains fields that are used for measuring data. Usually amounts or counts will be used as attributes in a fact table. The other fields present in a fact table will be foreign key fields related to a primary key field in a dimension.

This series of one-to-many relationships of dimensions to facts gives expansive querying abilities to the dimensions that will, in many cases, search, sort, and pivot the data effectively. The dimensions are used to describe the facts and grant a means for actively querying the facts.

It's important to remember that all of these tables are just regular database tables. Fact tables will always be the heart of the dimensional data model. Dimensions can also be conformed or shared across multiple fact tables or data marts. Conforming dimensions, which is the overlap of tables, is a common practice in relational OLAP data models. Dimensions are usually conformed because there is no need to store the table more than once; they will be used across data marts. These conformed dimensions are logically represented in a logical data model of each mart, but the data is physically persisted in only one table to reduce storage.

Figure 3.9 shows an example of a financial data mart. In this mart, you have ACCOUNT, DATE (TIME), DEPARTMENT, SCENARIO, and some type of ORGANIZATION structure dimensions—all relating to a simple fact table containing the AMOUNT values. The data is modeled into these tables because these subjects of data are commonly used concepts for financial applications. The DATE dimension allows for flexible time-based reporting, and this dimension will be conformed across the next data mart shown in Figure 3.10: the sales data mart.

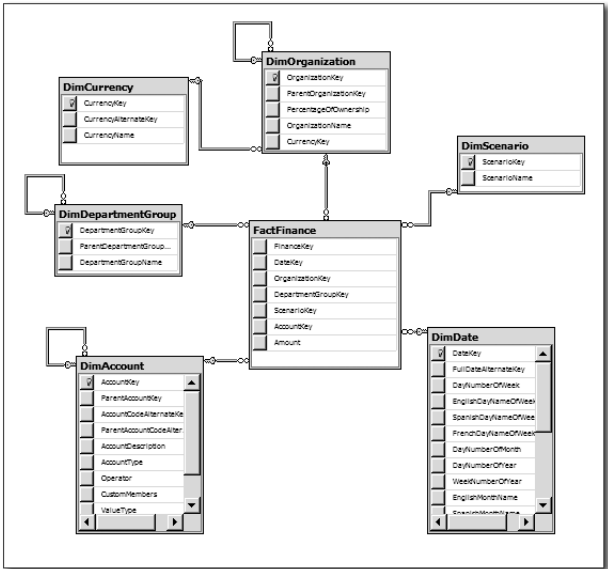


Figure 3.9 Typical Relational OLAP Data Model for a Financial Data Mart

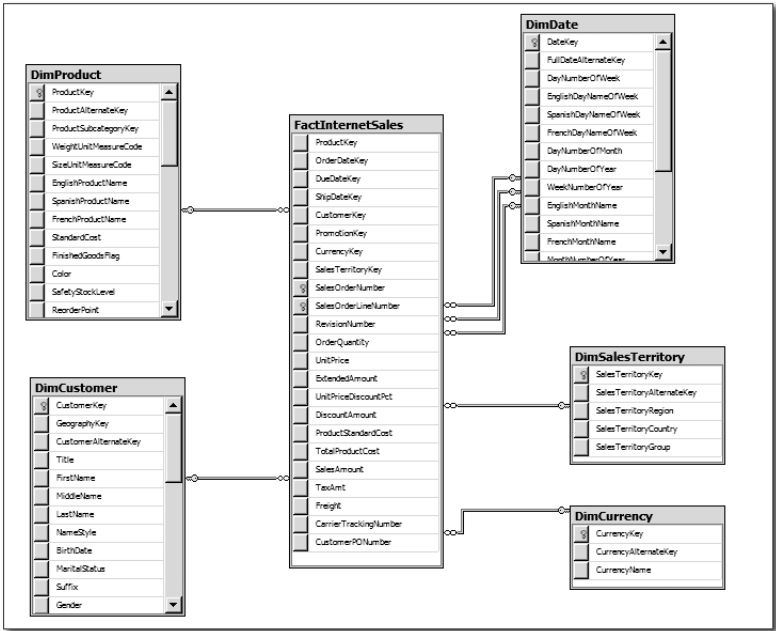


Figure 3.10 Typical Relational OLAP Data Model for a Sales Data Mart

In a typical sales data mart, you have dimensions such as PRODUCT, DATE (TIME), CUSTOMER, CURRENCY, and some type of ORGANIZATION (TERRITORY in this case)—all relating to a simple fact table containing the AMOUNT values. The DATE dimension is used for the same flexible time reporting principles outlined in the financial data mart, but all of the other dimensions represent subjects that are needed for sales analysis and reporting.

Notice that the DATE dimension and the CURRENCY dimension are repeated or shared across the two data marts. Physically, at the database level, the tables aren't repeated; there is only one DIM\_CURRENCY table and one DIM\_DATE table, so it's merely a logical repetition for organizational purposes. This repetition is a perfect example of a conformed dimension. We'll use conformed dimensions while building out the examples in the case study in this chapter (Section 3.4) to save on table space for the data marts in SAP HANA.

Note

Data marts can contain more than one fact table, but for the sake of simplicity and clear explanation of concepts in SAP HANA, we'll keep the dimensional model relatively simple.

Recall that data marts are segregated in terms of the data that they measure and describe. Considering Figure 3.9 and Figure 3.10, it's easy to see why there are two different data marts. There is little similarity with respect to the fact tables between these two data marts. The finance data mart simply has one amount field and all of the dimensional foreign keys. This allows quick pivoting and aggregations of the amount data by any of the dimensions. The sales data mart contains more complex facts because there are many more facets of a sale to measure, but the important point to note is that all sales measures are centrally located in the fact table so that all of the dimensional pivoting, querying, and sorting is done effortlessly across any or all of the dimensions. This is the primary concept that is central to a relational OLAP model.

The other concept central to a data mart is the grain of the data that is stored in the fact table or tables that make up the data mart. The grain of the data—sometimes known as the granularity of the data mart—is specific to the logical key structure of the fact table. The logical key structure is the single attribute (or combination of attributes) that makes the row of data in the fact unique.

Take, for example, the sales data mart in Figure 3.10. The FactInternetSales table has two primary key columns that make a composite key or a combined logical key: SalesOrderNumber and SalesOrderLineNumber. Data in this example is stored at the line level of the transaction; this is the lowest level of granularity, as this is at the line item of the sale. However, data is often repeated and stored in other fact tables at a higher level of granularity, such as by calendar week or product. This pre-aggregation, or materialized aggregation, is often necessary as a performance boost or method to realize a performance gain when the database platform runs out of tuning capabilities.

In this example, the fact table is at the line level, although it also contains header information in a denormalized fashion. *Denormalizing data* is merely the process of optimizing the data for reads by grouping redundant attributes together into one structure, rather than splitting the redundant attributes into multiple normalized tables. In an OLTP-normalized model, both the header-level data and the line-level data would be in two separate tables. Denormalizing occurs to ensure that one read from the FactInternetSales table obtains the necessary data, rather than reading two tables and reconstructing the data logically in the database engine with a join. This principle of denormalizing is crucial to optimizing performance in an OLAP model.

SAP HANA takes this principle of denormalizing data further to a level that at first seems contrary to performance and storage considerations; however, upon closer inspection, we find that, in many cases, denormalized data performs much faster in a column column-store table in SAP HANA using native column column-store functions than a normalized table structure in SAP HANA. In a normalized design, in many cases, SAP HANA may need to use the slower row engine to join the data. So even though this may seem counterintuitive at first, it's a core principle that will make an already fast SAP HANA system even faster!

Key Points about Data Marts

- ▶ Fact tables are the central element, or heart, of the data mart.
- ▶ Fact tables are surrounded by dimension tables.
- ▶ Fact tables contain only the measures and foreign keys back to the dimension tables.
- ▶ Dimension tables describe the facts.
- ▶ Denormalizing data or materializing aggregated data are techniques that are often used to boost performance when performance gains are no longer available from the database platform.

3.3.2 SAP HANA Relational OLAP Modeling

Many of the concepts and even physical structures translate over to SAP HANA directly from the conventional legacy OLAP counterparts. However, some distinctions specific to SAP HANA emerge. Let's now explore these distinctions in detail; these will drive the focus and the discussion to draw attention to the techniques to exploit SAP HANA for maximum performance benefit.

The baseline for relational OLAP modeling in SAP HANA is exactly what we've just described. It is best practice to lay out a solid relational OLAP model as a starting point for running BI operations against SAP HANA, but then a major deviation occurs: *further denormalizing*.

In SAP HANA, joins are sometimes more costly than one read against a compressed column-store data set containing the columns needed for any and all aggregate operations. Because SAP HANA has some pretty sophisticated built-in functions available in the column engine, we recommend that you flatten, or denormalize, dimensional columns into the facts so that you have data between two tables with high cardinality, or a high degree of uniqueness within the column. This way, there is one read against the table that is used to get all the data you need. (We'll go into further detail around denormalizing techniques in Section 3.3.3.)

It's true that SAP HANA can render hierarchies against denormalized or flattened data natively, but to maximize reusability in the SAP HANA analytic model, we still recommend that you keep the core attributes of a dimension in a true dimensional structure residing in a separate table. This way, if attributes from one dimension are all that is needed—for an attribute view, for example—then there is less work needed if a change needs to occur in the base table's columns. This allows for greater reusability when the base data is distributed in a more standard fashion.

Another technique that works well in an SAP HANA relational OLAP data model is adding aggregate columns directly into the fact tables rather than only storing the components of aggregations. For example, if you often multiply quantity by price to store an extended price in a fact table, consider storing the extended price as a calculation. The calculation will happen just as fast as if you stored the calculated value in a separate column. Storing the calculation is also always faster than reassembling the calculation at runtime in either the column or row engine. The final benefit of not storing the calculated data is that there is no redundant data occupying valuable space in memory in the SAP HANA database.

Of course, keep in mind that after you have calculated columns stored in SAP HANA tables, you'll need to explicitly state the columns when you need to insert data to avoid inadvertently setting the calculated columns by mistake.

Table 3.1 shows a scenario in which you have a simple example of a conventional query and table storage structure that stores the components of quantity and price as a stored, calculated value in the faster query block. This is unnecessary with SAP HANA because the calculation can be stored instead of a value that needs updates. This speeds the query because the calculation happens in real time and is always updated.

Scenario	SQL Needed
Conventional	<pre>select * from my_table where quantity * price = 100;</pre>
Faster query	<pre>select * from my_table where extended price = 100;</pre>
Supporting DDL for calculated extended_price column	<pre>alter table fact_sale add (extended_price decimal(10,2) GENERATED ALWAYS AS quantity * price);</pre>

Table 3.1 Calculated Column for a Faster Query

Paying attention to dates and time is also important for dimensional modeling in SAP HANA. For a best practice much like the denormalizing examples listed earlier, keep a date or time dimension separated from your facts for drilling or range-based date manipulation, just as with a standard dimensional model. However, SAP HANA offers built-in time dimension tables, as shown in Table 3.2.

SAP HANA Generated Time Dimension Table	Description
_SYS_BI.M_TIME_DIMENSION_YEAR	Time series with a year focus
_SYS_BI.M_TIME_DIMENSION_MONTH	Time series with a month focus
_SYS_BI.M_TIME_DIMENSION_WEEK	Time series with a week focus
_SYS_BI.M_TIME_DIMENSION	General time series generated data

Table 3.2 SAP HANA–Generated Time Dimension Tables

These tables will be shown in detail later in the analytic modeling sections. To take advantage of this built-in functionality, it's important to note that dates should be stored in your fact tables with sister `varchar(8)` columns in a format of `YYYYMMDD`. You'll notice in the case study at the end of this chapter that, for each date column listed in the `FactInternetSales` table, there are both `datefieldnameKEY` columns and `datefieldname_CHAR` columns present. The second column is simply a `varchar(8)` representation of the date, needed to facilitate joins to the `varchar(8)` fields in the `M_TIME_DIMENSION` tables referenced in Table 3.2. This may or may not be convenient for your data, depending on how your date columns are stored in the source data. This is very convenient for native SAP Business Suite data because this is how this data is stored in SAP. So, you can take advantage of this functionality without modification when you are running SAP Business Suite on SAP HANA, but with non-SAP data, you may very well have to transform the base date values to this `varchar()` format.

As an example value for both types of date fields, `2013-01-01 00:00:00` in your table would also be stored in a `varchar(8)` column (`20130101`) to take advantage of the built-in date and time attribute tables present in SAP HANA. Note that this functionality really serves its best use with SAP-native date formats in SAP sources. For source-agnostic BI, we recommend that you use a custom date dimension table because this adds the most flexibility and deals with dates in all formats.

A final item worth considering when you're constructing your data model in SAP HANA is the fact that you must ensure data type support for all of your aggregate operations. Take a simple example of a numeric column with a data type `Decimal(8,4)` containing a value `1111.1111`. If this number is multiplied by 10, you have a value of `11110.1111`. This value is now out of range in the base column of the SAP HANA table. You must always store your data at the greatest precision required for the max operation that will occur on that data.

This requires some thinking in advance about the types of calculations that will occur on the data you're using, even while choosing data types up front for your tables. Please keep in mind that the maximum declaration that is currently allowed is `Decimal(34,4)`. No matter what you're going to use, if you're sticking with a `DECIMAL` data type, this is the maximum value allowed. A best practice to avoid this behavior is to simply convert all decimal columns to float to handle overflow for division operations ratios; this always avoids overflow issues because SAP HANA doesn't handle the conversion automatically.

Key Points about SAP HANA-Specific Data Marts

- ▶ As with non-SAP HANA data marts, fact tables are the central element, or heart, of the data mart.
- ▶ As with non-SAP HANA data marts, fact tables are surrounded by dimension tables.
- ▶ As with non-SAP HANA data marts, fact tables contain the measures and foreign keys back to the dimension tables, but also specific denormalized attributes where cardinality is high between tables.
- ▶ SAP HANA data marts also contain a time dimension, but have dates stored both as date values and character data types for use against SAP HANA time dimension tables when SAP application data is primarily used in SAP HANA.
- ▶ Denormalizing data to extremes that would cripple a conventional database is often the way to get optimal performance in SAP HANA.
- ▶ Materializing aggregated data is simply not necessary with SAP HANA because performance is considerable in column-store tables and in the column engine.
- ▶ Numeric types that will be used in aggregate calculations require particular attention. You must cover the size of the resulting value from a calculation in the base numeric column. Remember that `decimal(34,4)` is the maximum allowed decimal type.

3.3.3 Denormalizing Data in SAP HANA

In column-store tables within SAP HANA, denormalized data is something that you should always find in some area of a dimensional data model. Recall our recommendation that you flatten or denormalize dimensional columns into the facts where you have data between two tables that have high cardinality or a high degree of uniqueness within the column. This is because joining data from a dimension with high cardinality is often more costly in terms of performance than just storing the attributes from the dimension that will be used often for querying or aggregations directly into the fact.

Two important principles are addressed by denormalizing data in SAP HANA. First, avoid the join in the engine and (most times) stay entirely in the column engine, where processing is much faster. Second, the penalty for the redundant data, from a storage perspective, isn't too severe because the column-store table stores only the values that repeat once, anyway, due to the nature of compression in the column-store table. Normalization is something that typically occurs in a relational database to increase performance and decrease storage; however, in columnar tables in SAP HANA, this idea is turned on its head



because compression helps with both the speed of access and limiting the extra footprint of the data in memory.

Take, for instance, a product dimension and a sales fact table. These tables are often used together in SQL queries for reporting. Maybe you want to filter on attributes such as color, class, or style, or you need to see standard cost as an aggregated value to be used in calculations such as price or sold quantity. These are combinations that will occur quite often in typical sales analysis scenarios. Product data will have a high degree of uniqueness or cardinality, as well, because data is often stored at the SKU or UPC level. A record in the product dimension will be a record of unique product attributes and is the perfect candidate for denormalizing aspects of the dimension into the fact.

To start, you must identify the attributes in the table that will be the subject of querying, filtering, or aggregations. For this example, we selected the highlighted attributes from the DIM\_PRODUCT table shown in Figure 3.11 to store as denormalized attributes in the fact table.

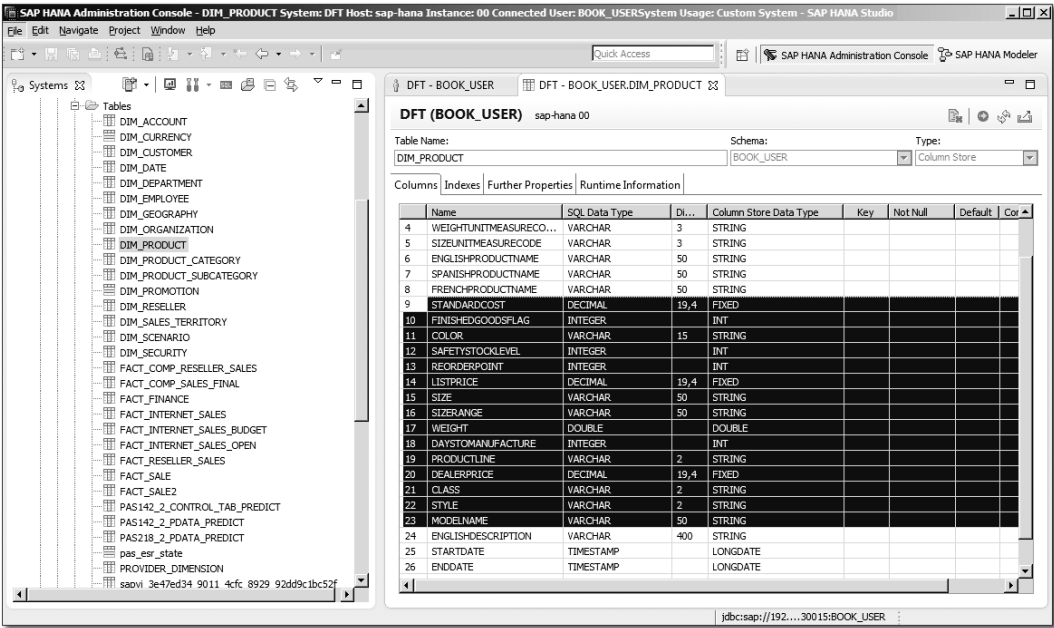


Figure 3.11 Columns from DIM\_PRODUCT Added to Reduce Frequent Joins

To create the columns in the FactInternetSales table, you now need to write an ALTER TABLE SQL statement to add the new columns. Listing 3.1 shows an example of the ALTER TABLE statement that is used to add the columns to the FactInternetSales table.

```
SQL Used to Add the Columns to the Fact Table: FactInternetSales
-- Add DIM_Product columns to FACT_INTERNET_SALES
-- due to high cardinality.
-- Don Loden
-- 02.15.2013
```

```
alter table fact_internet_sales add
(
  "DIM_PRD__STANDARDCOST" DECIMAL(19,4) CS_FIXED null,
  "DIM_PRD__FINISHEDGOODSFLAG" INTEGER CS_INT null,
  "DIM_PRD__COLOR" VARCHAR(15) null,
  "DIM_PRD__SAFETYSTOCKLEVEL" INTEGER CS_INT null,
  "DIM_PRD__REORDERPOINT" INTEGER CS_INT null,
  "DIM_PRD__LISTPRICE" DECIMAL(19,4) CS_FIXED null,
  "DIM_PRD__SIZE" VARCHAR(50) null,
  "DIM_PRD__SIZERANGE" VARCHAR(50) null,
  "DIM_PRD__WEIGHT" DOUBLE CS_DOUBLE null,
  "DIM_PRD__DAYSTOMANUFACTURE" INTEGER CS_INT null,
  "DIM_PRD__PRODUCTLINE" VARCHAR(2) null,
  "DIM_PRD__DEALERPRICE" DECIMAL(19,4) CS_FIXED null,
  "DIM_PRD__CLASS" VARCHAR(2) null,
  "DIM_PRD__STYLE" VARCHAR(2) null,
  "DIM_PRD__MODELNAME" VARCHAR(50) null
);
```

Listing 3.1 SQL Data Definition Language (DDL) Used to Create FactInternetSales

In Figure 3.12, you can see what the FactInternetSales table looks like after you execute the SQL to add the columns. All of the denormalized columns are ready for use in the fact table. Notice that the columns were not removed from DIM\_PRODUCT to foster reusability and ease maintenance for analytic modeling, as you'll see later in the book.

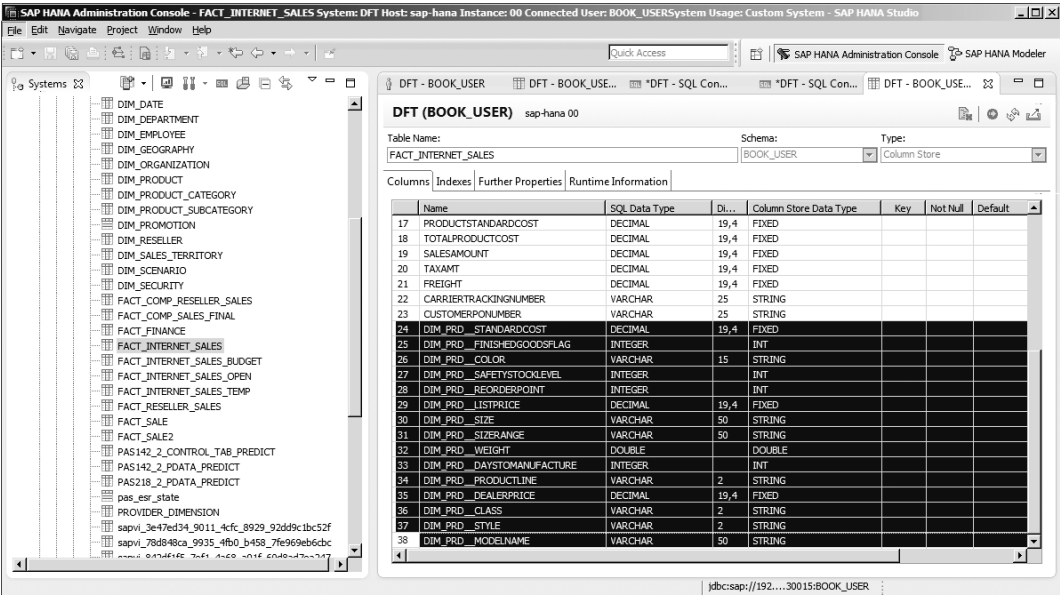


Figure 3.12 FactInternetSales Table after Replicating the Columns from DIM\_PRODUCT

3.4 Case Study: Creating Data Marts and Tables for an SAP HANA Project

To illustrate various presentation options and use cases, this book uses a case study to follow a project from the ground up by starting with the data model; then provisioning the data, creating the analytic model; and, finally, fully realizing the BI capabilities with the consumption of the data using the SAP BusinessObjects BI tools. To perform all of these actions, we'll be using the sample Microsoft AdventureWorks data model for a fictitious company called AdventureWorks Cycle Company. We chose this data and model because it's a readily available sample schema with data that is familiar to many developers.

Currently, this SAP HANA system is a blank slate containing nothing but a bare install. So, first, you'll need to create a schema to house and organize the tables that you'll create. Then, you'll finally create the tables and model them to follow the best practices in an SAP HANA data model.

3.4.1 Creating a Schema for the Data Mart

Before you can begin building tables in SAP HANA Studio using SQL or a tool such as SAP Data Services, you need a schema created to house and organize your tables. To create the schema in SAP HANA, you must have a user created that can authenticate to SAP HANA. For all of the connections in the case study for this book, you'll be using the user BOOK\_USER.

To create the schema using BOOK\_USER, perform the following steps:

- 1. Open SAP HANA Studio and connect using the BOOK\_USER user, as shown in Figure 3.13.
- If you're currently connected as a different user and need to change the user, you may do so in the pop-up menu. Get to this menu by right-clicking your connected SAP HANA system.
- 2. Open the DEVELOPMENT perspective.
- 3. Open the PROJECT EXPLORER view.

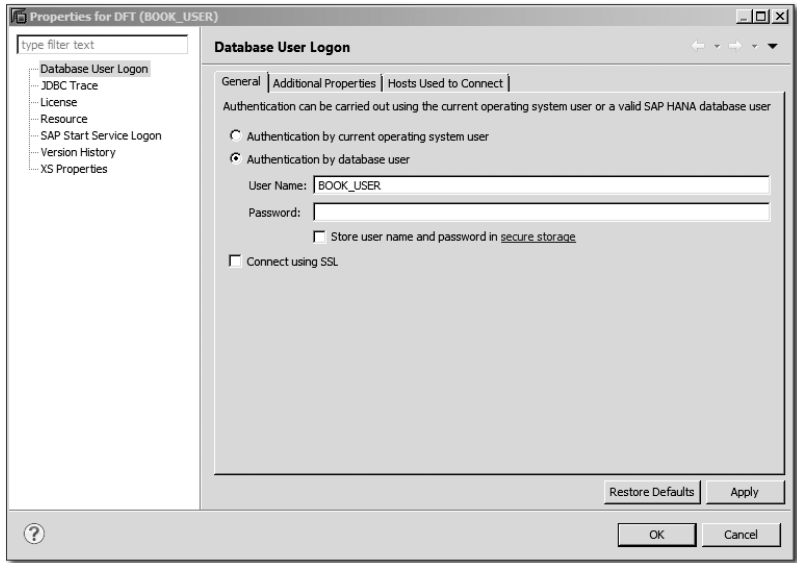


Figure 3.13 Choosing a User Name to Sign In

- 4. Browse in the Project Workspace to the folder where you want to create your schema definition file and right-click the folder. A menu pops up with a field

where you can specify the name of the schema. For our example, use BOOK\_USER.hdbschema. Then, choose FINISH to save the schema.

Caution!

If you want your schema to be a design-time object, you'll need to create the schema as a file to be saved in the repository.

- 5. Define the schema name by opening the file you just created in the previous step by inserting this code: schema\_name = "BOOK\_USER";.
- 6. Save and activate the schema file.
  - ▶ Commit the schema to the repository by right-clicking the BOOK\_USER schema and choosing TEAM • COMMIT.
  - ▶ Activate the schema by right-clicking the BOOK\_USER schema.
  - ▶ Choose TEAM • COMMIT.

By performing these steps, you've now both created and activated a schema in SAP HANA, as shown in Figure 3.14. This schema is ready for use. In the next section of the case study, you'll begin to create the column-store tables. These tables will be the foundation of all the rest of the examples in this book.

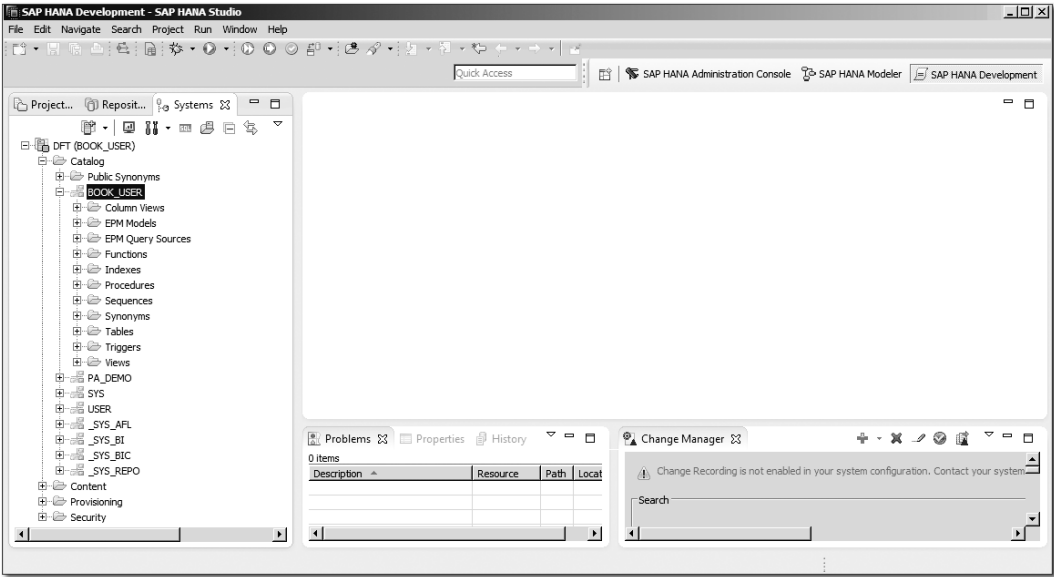


Figure 3.14 Finished Schema Ready for Use in SAP HANA

3.4.2 Creating the Fact Table and Dimension Tables in SAP HANA

We'll show you a few different ways to create the fact and dimension tables in SAP HANA, especially during the data provisioning sections using the unique features of SAP Data Services. However, for this chapter, to focus on creating the tables and the underlying model of the tables, you'll create the tables using SQL in the SAP HANA Studio.

To create the tables using SQL in the SAP HANA Studio, perform the following steps:

- 1. Open SAP HANA Studio and connect using the BOOK\_USER user.
- 2. Open the MODELER perspective.
- 3. Open the PROJECT EXPLORER view.
- 4. Browse in the Project Workspace to select the tables folder under the BOOK\_USER schema that you created earlier (shown in Figure 3.14).
- 5. Click the SQL button, indicated by the arrow in Figure 3.15.

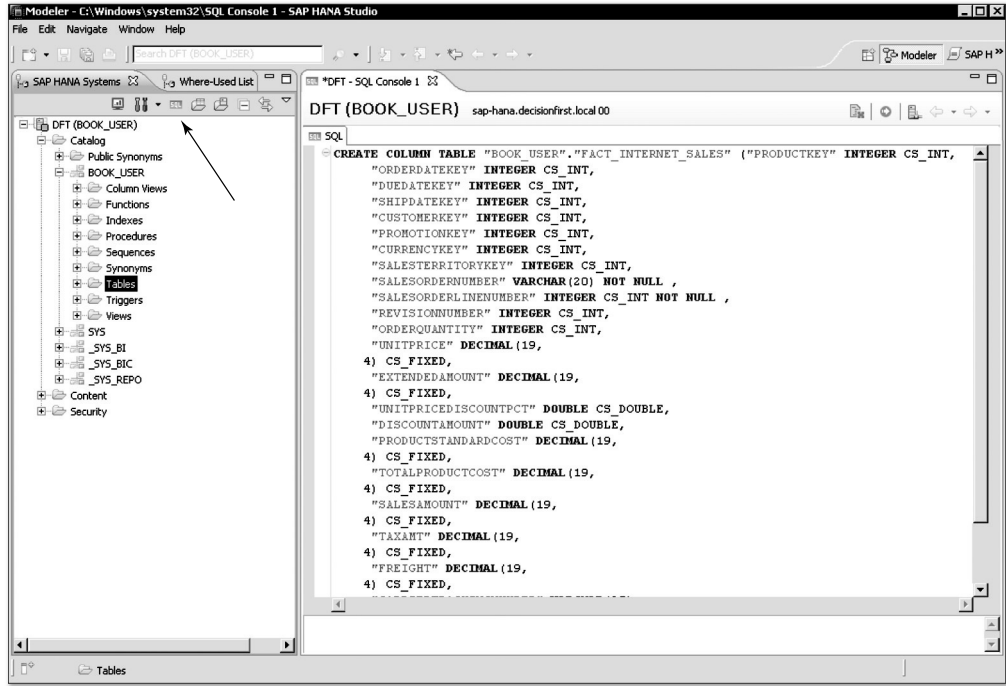


Figure 3.15 Opening the SQL Editor for the Current Session to Create the Tables

6. Type each of the following SQL statements—Listing 3.2 for FactInternetSales, Listing 3.3 for DIM\_PRODUCT, Listing 3.4 for DIM\_CUSTOMER, and Listing 3.5 for DIM\_DATE into the SQL Editor, as shown in Figure 3.15.

Listing 3.2 is the main fact table with Internet sales measures. The only things differentiating this fact table from a standard fact table are the extra varchar() date columns for SAP HANA functions and denormalized columns from product dimension.

```
CREATE COLUMN TABLE "BOOK_USER"."FACT_INTERNET_SALES" ("PRODUCTKEY"
INTEGER CS_INT,
  "ORDERDATEKEY" INTEGER CS_INT,
  "ORDERDATE_CHAR" VARCHAR(8), --SUPPORTS HANA DATE FUNCTIONS
  "DUEDATEKEY" INTEGER CS_INT,
  "DUEDATE_CHAR" VARCHAR(8), --SUPPORTS HANA DATE FUNCTIONS
  "SHIPDATEKEY" INTEGER CS_INT,
  "SHIPDATE_CHAR" VARCHAR(8), --SUPPORTS HANA DATE FUNCTIONS
  "CUSTOMERKEY" INTEGER CS_INT,
  "PROMOTIONKEY" INTEGER CS_INT,
  "CURRENCYKEY" INTEGER CS_INT,
  "SALESTERRITORYKEY" INTEGER CS_INT,
  "SALESORDERNUMBER" VARCHAR(20) NOT NULL ,
  "SALESORDERLINENUMBER" INTEGER CS_INT NOT NULL ,
  "REVISIONNUMBER" INTEGER CS_INT,
  "ORDERQUANTITY" INTEGER CS_INT,
  "UNITPRICE" DECIMAL(19,
4) CS_FIXED,
  "EXTENDEDAMOUNT" DECIMAL(19,
4) CS_FIXED,
  "UNITPRICEDISCOUNTPCT" DOUBLE CS_DOUBLE,
  "DISCOUNTAMOUNT" DOUBLE CS_DOUBLE,
  "PRODUCTSTANDARDCOST" DECIMAL(19,
4) CS_FIXED,
  "TOTALPRODUCTCOST" DECIMAL(19,
4) CS_FIXED,
  "SALESAMOUNT" DECIMAL(19,
4) CS_FIXED,
  "TAXAMT" DECIMAL(19,
4) CS_FIXED,
  "FREIGHT" DECIMAL(19,
4) CS_FIXED,
  "CARRIERTRACKINGNUMBER" VARCHAR(25),
  "CUSTOMERPONUMBER" VARCHAR(25),
  "DIM_PRD__STANDARDCOST" DECIMAL(19,
4) CS_FIXED,
  "DIM_PRD__FINISHEDGOODSFLAG" INTEGER CS_INT,
  "DIM_PRD__COLOR" VARCHAR(15),
```

```
"DIM_PRD__SAFETYSTOCKLEVEL" INTEGER CS_INT,
"DIM_PRD__REORDERPOINT" INTEGER CS_INT,
"DIM_PRD__LISTPRICE" DECIMAL(19,
4) CS_FIXED,
"DIM_PRD__SIZE" VARCHAR(50),
"DIM_PRD__SIZERANGE" VARCHAR(50),
"DIM_PRD__WEIGHT" DOUBLE CS_DOUBLE,
"DIM_PRD__DAYSTOMANUFACTURE" INTEGER CS_INT,
"DIM_PRD__PRODUCTLINE" VARCHAR(2),
"DIM_PRD__DEALERPRICE" DECIMAL(19,
4) CS_FIXED,
"DIM_PRD__CLASS" VARCHAR(2),
"DIM_PRD__STYLE" VARCHAR(2),
"DIM_PRD__MODELNAME" VARCHAR(50),
PRIMARY KEY ("SALESORDERNUMBER",
"SALESORDERLINENUMBER"))
```

Listing 3.2 SQL DDL (Fact Table) for FactInternetSales

The standard product dimension describes product-level attributes. Notice that certain columns have been repeated in the fact table in Listing 3.3, yet they still exist here for reusability in the SAP HANA analytic model.

```
CREATE COLUMN TABLE "BOOK_USER"."DIM_PRODUCT" ("PRODUCTKEY" INTEGER
CS_INT NOT NULL ,
  "PRODUCTALTERNATEKEY" VARCHAR(25),
  "PRODUCTSUBCATEGORYKEY" INTEGER CS_INT,
  "WEIGHTUNITMEASURECODE" VARCHAR(3),
  "SIZEUNITMEASURECODE" VARCHAR(3),
  "ENGLISHPRODUCTNAME" VARCHAR(50),
  "SPANISHPRODUCTNAME" VARCHAR(50),
  "FRENCHPRODUCTNAME" VARCHAR(50),
  "STANDARDCOST" DECIMAL(19,
4) CS_FIXED,
  "FINISHEDGOODSFLAG" INTEGER CS_INT,
  "COLOR" VARCHAR(15),
  "SAFETYSTOCKLEVEL" INTEGER CS_INT,
  "REORDERPOINT" INTEGER CS_INT,
  "LISTPRICE" DECIMAL(19,
4) CS_FIXED,
  "SIZE" VARCHAR(50),
  "SIZERANGE" VARCHAR(50),
  "WEIGHT" DOUBLE CS_DOUBLE,
  "DAYSTOMANUFACTURE" INTEGER CS_INT,
  "PRODUCTLINE" VARCHAR(2),
  "DEALERPRICE" DECIMAL(19,
4) CS_FIXED,
  "CLASS" VARCHAR(2),
```



```
"STYLE" VARCHAR(2),
"MODELNAME" VARCHAR(50),
"ENGLISHDESCRIPTION" VARCHAR(400),
"STARTDATE" LONGDATE CS_LONGDATE,
"ENDDATE" LONGDATE CS_LONGDATE,
"STATUS" VARCHAR(7),
PRIMARY KEY ("PRODUCTKEY"))
```

Listing 3.3 SQL DDL (Dimension Table) for DIM\_PRODUCT

This standard customer dimension table describes customer-level attributes (Listing 3.4). The CUSTOMERKEY field has a foreign key that relates this table to the fact table.

```
CREATE COLUMN TABLE "BOOK_USER"."DIM_CUSTOMER" ("CUSTOMERKEY"
INTEGER CS_INT NOT NULL ,
"GEOGRAPHYKEY" INTEGER CS_INT,
"CUSTOMERALTERNATEKEY" VARCHAR(15),
"TITLE" VARCHAR(8),
"FIRSTNAME" VARCHAR(50),
"MIDDLENAME" VARCHAR(50),
"LASTNAME" VARCHAR(50),
"NAMESTYLE" INTEGER CS_INT,
"BIRTHDATE" DAYDATE CS_DAYDATE,
"MARITALSTATUS" VARCHAR(1),
"SUFFIX" VARCHAR(10),
"GENDER" VARCHAR(1),
"EMAILADDRESS" VARCHAR(50),
"YEARLYINCOME" DECIMAL(19,
4) CS_FIXED,
"TOTALCHILDREN" INTEGER CS_INT,
"NUMBERCHILDRENATHOME" INTEGER CS_INT,
"ENGLISHEDUCATION" VARCHAR(40),
"SPANISHEDUCATION" VARCHAR(40),
"FRENCHEDUCATION" VARCHAR(40),
"ENGLISHOCCUPATION" VARCHAR(100),
"SPANISHOCCUPATION" VARCHAR(100),
"FRENCHOCCUPATION" VARCHAR(100),
"HOUSEOWNERFLAG" VARCHAR(1),
"NUMBERCARSOWNED" INTEGER CS_INT,
"ADDRESSLINE1" VARCHAR(120),
"ADDRESSLINE2" VARCHAR(120),
"PHONE" VARCHAR(20),
"DATEFIRSTPURCHASE" DAYDATE CS_DAYDATE,
"COMMUTEDISTANCE" VARCHAR(15),
PRIMARY KEY ("CUSTOMERKEY"))
```

Listing 3.4 SQL DDL (Dimension Table) for DIM\_CUSTOMER

The standard time dimension describes date attributes (Listing 3.5). The DATE-KEY field has a foreign key that relates this table to the fact table on multiple date attributes. The basic concept is that the date dimension will be related back on any date column in the fact table to allow for flexibility on any type of date- or time-based reporting.

```
CREATE COLUMN TABLE "BOOK_USER"."DIM_DATE" ("DATEKEY" INTEGER CS_INT
NOT NULL ,
"FULDATEALTERNATEKEY" DAYDATE CS_DAYDATE,
"DAYNUMBEROFWEEK" INTEGER CS_INT,
"ENGLISHDAYNAMEOFWEEK" VARCHAR(10),
"SPANISHDAYNAMEOFWEEK" VARCHAR(10),
"FRENCHDAYNAMEOFWEEK" VARCHAR(10),
"DAYNUMBEROFMONTH" INTEGER CS_INT,
"DAYNUMBEROFYEAR" INTEGER CS_INT,
"WEEKNUMBEROFYEAR" INTEGER CS_INT,
"ENGLISHMONTHNAME" VARCHAR(10),
"SPANISHMONTHNAME" VARCHAR(10),
"FRENCHMONTHNAME" VARCHAR(10),
"MONTHNUMBEROFYEAR" INTEGER CS_INT,
"CALENDARQUARTER" INTEGER CS_INT,
"CALENDARQUARTERYEAR" VARCHAR(8),
"CALENDARYEAR" INTEGER CS_INT,
"CALENDARYEARMONTH" VARCHAR(15),
"CALENDARYEARWEEK" VARCHAR(20),
"CALENDARSEMESTER" INTEGER CS_INT,
"FISCALQUARTER" INTEGER CS_INT,
"FISCALYEAR" INTEGER CS_INT,
"FISCALSEMESTER" INTEGER CS_INT,
PRIMARY KEY ("DATEKEY"))
```

Listing 3.5 SQL DDL (Dimension Table) for DIM\_DATE

7. Press **F8** to execute the queries.

After executing all four SQL statements, you have one fact table and three dimension tables. These tables form the core of the data mart that will be used in the subsequent sections of the case study, and this data set will remain the base data for all of the examples present in this book. You'll also notice a financial structure consisting of the following data mart tables:

- ▶ FACT\_FINANCE
- ▶ DIM\_CURRENCY
- ▶ DIM\_ORGANIZATION

- ▶ DIM\_SCENARIO
- ▶ DIM\_DATE
- ▶ DIM\_ACCOUNT
- ▶ DIM\_DEPARTMENT\_GROUP

Note that DIM\_DATE is a conformed dimension across the financial mart and sales mart. DIM\_DATE references the same table that was created in this section. These financial data mart tables are created in the same manner as the sales data mart. The descriptions were given only to limit redundant descriptions for the case study.

### 3.5 Summary

SAP HANA is a tremendously powerful and flexible platform, in part because it truly has the ability to act as a chameleon and masquerade as multiple platforms. SAP HANA is unique in the sense that it can easily replace many of these platforms quickly because it shares the common, conventionally approved language for data access: SQL. This makes SAP HANA a plug-and-play fit for replacing the data and analytic architecture for many applications with a far more sophisticated and well-thought-out development platform. The fact that SAP HANA can also interpret MDX queries natively speaks to the same rapid integration and replacement of conventional cube-based technologies.

Native support for MDX was one reason it was no surprise that SAP undertook the task of moving SAP BW to SAP HANA so quickly. For an application such as SAP BW, moving to SAP HANA was merely another database port. This ease of movement and transport goes a long way toward SAP's no-disruption model. Now that the SAP Business Suite is also certified to run on SAP HANA, the sky is the limit in terms of possibilities on a mature and robust platform that really does do it all.

Now that you have an understanding of how SAP HANA stores data and what is needed for high-performing data in SAP HANA, we can look toward Part II of the book, which focuses on the data provisioning process. We call out the word *process* because you shouldn't just load your data into SAP HANA. Before you provision data into SAP HANA, there are some things that need to be addressed with a thorough pre-provisioning process. We'll examine this pre-provisioning in process in detail in Chapter 4.

# Contents

Acknowledgments .....	17
Preface .....	19

## PART I Introduction

<b>1</b>	<b>SAP HANA, SAP BusinessObjects BI, and SAP Data Services .....</b>	<b>27</b>
1.1	What Is SAP HANA? .....	28
1.1.1	Software Layers and Features .....	28
1.1.2	Hardware Layers and Features .....	32
1.2	Business Intelligence Solutions with SAP HANA .....	38
1.2.1	SAP BW on SAP HANA .....	38
1.2.2	Native Implementation of SAP HANA for Analytics .....	42
1.3	SAP Business Suite on SAP HANA .....	57
1.4	Traditional EIM with SAP Data Services .....	60
1.4.1	Align IT with the Business .....	60
1.4.2	Establish Processes to Manage the Data .....	61
1.4.3	Source System Analysis .....	61
1.4.4	Develop a Data Model .....	62
1.4.5	Load the Data .....	62
1.5	Traditional Business Intelligence with SAP BusinessObjects BI .....	62
1.5.1	The Semantic Layer (Universe) .....	63
1.5.2	Ad Hoc Reporting .....	64
1.5.3	Self-Service BI .....	64
1.5.4	IT-Developed Content .....	66
1.6	Solution Architectural Overview .....	66
1.6.1	SAP Data Services .....	67
1.6.2	SAP BusinessObjects BI .....	70
1.6.3	SAP HANA .....	73
1.7	Summary .....	76
<b>2</b>	<b>Securing the SAP HANA Environment .....</b>	<b>77</b>
2.1	Configuring the SAP HANA Environment for Development .....	78
2.1.1	Introduction to the SAP HANA Repository .....	79
2.1.2	Configuring SAP HANA Studio .....	80

2.1.3	Setting Up Packages and Development Projects .....	87
2.1.4	Setting up Schemas in SAP HANA .....	95
2.2	SAP HANA Authorizations .....	100
2.2.1	Types of SAP HANA Privileges .....	101
2.2.2	Granting of Privileges and the Life Cycle of a Grant .....	106
2.3	User and Role Provisioning .....	108
2.3.1	Creating Roles (the Traditional Approach) .....	109
2.3.2	Creating Roles as Repository Objects .....	110
2.3.3	Preventing Rights Escalation Scenarios .....	115
2.3.4	Common Role Scenarios and Their Privileges .....	116
2.3.5	User Provisioning .....	126
2.4	SAP HANA Authentication .....	131
2.4.1	Internal Authentication with User Name and Password ....	133
2.4.2	Kerberos Authentication .....	135
2.4.3	SAML Authentication .....	136
2.4.4	Other Web-Based Authentication Methods for SAP HANA XS .....	137
2.4.5	Summary and Recommendations .....	138
2.5	Case Study: An End-to-End Security Configuration .....	139
2.5.1	Authentication Plan .....	139
2.5.2	Authorization Plan .....	147
2.5.3	User Provisioning Plan .....	150
2.6	Summary .....	152
<b>3</b>	<b>Data Storage in SAP HANA .....</b>	<b>155</b>
3.1	OLAP and OLTP Data Storage .....	155
3.1.1	The Spinning Disk Problem .....	157
3.1.2	Combating the Problem .....	157
3.2	Data Storage Components .....	165
3.2.1	Schemas and Users .....	165
3.2.2	Column-Store Tables .....	167
3.2.3	Row-Store Tables .....	172
3.2.4	Use Cases for Both Row- and Column-Store Tables .....	173
3.3	Modeling Tables and Data Marts .....	175
3.3.1	Legacy Relational OLAP Modeling .....	176
3.3.2	SAP HANA Relational OLAP Modeling .....	180
3.3.3	Denormalizing Data in SAP HANA .....	183
3.4	Case Study: Creating Data Marts and Tables for an SAP HANA Project .....	186
3.4.1	Creating a Schema for the Data Mart .....	187

3.4.2	Creating the Fact Table and Dimension Tables in SAP HANA .....	189
3.5	Summary .....	194

**PART II Getting Data Into SAP HANA**

<b>4</b>	<b>Preprovisioning Data with SAP Data Services .....</b>	<b>197</b>
4.1	Making the Case for Source System Analysis .....	197
4.2	SSA Techniques in SAP Data Services .....	202
4.2.1	Column Profiling .....	205
4.2.2	Relationship Profiling .....	211
4.3	SSA: Beyond Tools and Profiling .....	215
4.3.1	Establishing Patterns .....	217
4.3.2	Looking Across Sources .....	219
4.3.3	Treating Disparate Systems as One .....	219
4.3.4	Mapping Your Data .....	220
4.4	Summary .....	222
<b>5</b>	<b>Provisioning Data with SAP Data Services .....</b>	<b>223</b>
5.1	Provisioning Data Using SAP Data Services Designer .....	223
5.1.1	Metadata .....	225
5.1.2	Datastores .....	227
5.1.3	Jobs .....	231
5.1.4	Workflows .....	234
5.1.5	Data Flows .....	244
5.1.6	Transforms .....	255
5.1.7	Built-In Functions .....	277
5.1.8	Custom Functions and Scripts .....	281
5.1.9	File Formats .....	285
5.1.10	Real-Time Jobs .....	288
5.2	Introduction to SAP Data Services Workbench .....	290
5.2.1	Building a Data Flow .....	293
5.2.2	Moving Data from an Existing Data Warehouse .....	297
5.2.3	Porting Data with the Quick Replication Wizard .....	297
5.2.4	Modifying Data Flows and Jobs .....	304
5.3	Data Provisioning via Real-Time Replication .....	305
5.3.1	SAP Data Services ETL-Based Method (ETL and DQ) .....	306
5.3.2	SAP Landscape Transformation .....	307
5.4	Summary .....	308



<b>6</b>	<b>Loading Data with SAP Data Services .....</b>	<b>309</b>
6.1	Loading Data in a Batch .....	309
6.1.1	Steps .....	309
6.1.2	Methods .....	319
6.1.3	Triggers .....	328
6.2	Loading Data in Real Time .....	335
6.3	Case Study: Loading Data in a Batch .....	340
6.3.1	Initialization .....	344
6.3.2	Staging .....	345
6.3.3	Mart .....	374
6.3.4	End Script .....	387
6.4	Case Study: Loading Data in Real Time .....	389
6.5	Summary .....	395

**PART III Multidimensional Modeling in SAP HANA**

<b>7</b>	<b>Introduction to Multidimensional Modeling .....</b>	<b>399</b>
7.1	Understanding Multidimensional Models .....	400
7.2	Benefits of SAP HANA Multidimensional Modeling .....	404
7.2.1	Business Benefits .....	404
7.2.2	Technology Benefits .....	408
7.3	Summary .....	411

<b>8</b>	<b>Tools and Components of Multidimensional Modeling .....</b>	<b>413</b>
8.1	SAP HANA Studio .....	413
8.1.1	Systems View .....	417
8.1.2	Quick Launch View .....	418
8.2	Schemas .....	420
8.3	Packages .....	423
8.4	Summary .....	426

<b>9</b>	<b>Creating SAP HANA Information Views .....</b>	<b>427</b>
9.1	Attribute Views .....	427
9.1.1	Creating an Attribute View .....	429
9.1.2	Defining Properties of an Attribute View .....	431

9.1.3	Creating Hierarchies .....	439
9.1.4	Saving and Activating the Attribute View .....	442
9.2	Analytic Views .....	444
9.2.1	Creating an Analytic View .....	445
9.2.2	Defining Properties of an Analytic View .....	447
9.2.3	Saving and Activating the Analytic View .....	459
9.3	Calculation Views .....	460
9.3.1	Creating a Calculation View .....	461
9.3.2	Defining a Graphical Calculation View .....	466
9.3.3	Defining a Script-Based Calculation View .....	474
9.4	Summary .....	478

**10 Multidimensional Modeling in Practice ..... 479**

10.1	Data Processing in SAP HANA .....	479
10.1.1	Normalized Data versus Denormalized Data .....	480
10.1.2	Data Modeling versus Multidimensional Modeling .....	485
10.1.3	Managing Normalized Data in SAP HANA .....	487
10.2	Case Study 1: Modeling Sales Data to Produce Robust Analytics ....	490
10.2.1	Creating the Supporting Attribute Views .....	490
10.2.2	Creating Analytic Views .....	508
10.3	Case Study 2: Building Complex Calculations for Executive-Level Analysis .....	515
10.3.1	Creating the Package .....	516
10.3.2	Creating the Calculation View .....	518
10.3.3	Defining the Calculation View .....	520
10.4	Summary .....	533

**11 Securing Data in SAP HANA ..... 535**

11.1	Introduction to Analytic Privileges .....	536
11.1.1	What are Analytic Privileges? .....	536
11.1.2	Types of Analytic Privileges .....	537
11.1.3	Dynamic vs. Static Value Restrictions .....	538
11.2	Creating Analytic Privileges .....	540
11.2.1	Traditional Analytic Privileges .....	540
11.2.2	SQL-Based Analytic Privileges .....	556
11.3	Applying Analytic Privileges .....	561
11.3.1	Applying Analytic Privileges to Information Views .....	561
11.3.2	Interaction of Multiple Analytic Privileges and Multiple Restrictions .....	563

11.3.3	Interaction of Multiple Information Views with Analytic Privileges .....	564
11.4	Case Study: Securing Sales Data with Analytic Privileges .....	567
11.4.1	Overview and Requirements .....	568
11.4.2	Implementation Strategy .....	569
11.4.3	Implementation Examples .....	570
11.5	Summary .....	581

**PART IV Integrating SAP HANA with SAP Business Intelligence Tools**

<b>12</b>	<b>Building Universes for SAP HANA .....</b>	<b>585</b>
12.1	SAP HANA and the Universe .....	587
12.1.1	When to Use a Universe with SAP HANA .....	590
12.1.2	Connecting Universes to SAP HANA .....	592
12.2	Manually Building UNX Universes for SAP HANA .....	600
12.2.1	Creating Relational Connections .....	601
12.2.2	Creating OLAP Connections .....	610
12.2.3	Testing Connections Using the Local or Server Middleware .....	612
12.2.4	Creating Projects .....	613
12.2.5	Designing the Data Foundation .....	617
12.2.6	Designing the Business Layer .....	632
12.2.7	Publishing the Universe .....	639
12.3	Automatically Generating UNX Universes for SAP HANA .....	642
12.3.1	Creating a Local Connection .....	642
12.3.2	Selecting Information Views .....	645
12.3.3	Reviewing the Data Foundation and Business Layer .....	648
12.3.4	How SAP HANA Metadata Impacts the Process .....	655
12.4	The SAP HANA Engines in Universe Design .....	656
12.4.1	SAP HANA Join Engine .....	658
12.4.2	SAP HANA OLAP Engine .....	659
12.4.3	SAP HANA Calculation Engine .....	660
12.5	Case Study: Designing a Universe to Support Internet Sales Data ...	662
12.5.1	Creating the Universe Connection and Project .....	662
12.5.2	Designing the Data Foundation .....	663
12.5.3	Designing the Business Layer .....	671
12.5.4	Publishing the Universe .....	688
12.6	Summary .....	689

<b>13</b>	<b>Predictive Analytics with SAP HANA .....</b>	<b>691</b>
13.1	Predictive Analysis and SAP HANA: The Basics .....	692
13.1.1	The Predictive Analysis Process .....	696
13.1.2	When to Use Predictive Analytics .....	703
13.1.3	Predictive Tools Available in SAP HANA .....	706
13.2	Integrating with SAP HANA .....	712
13.2.1	Installing the Application Function Libraries .....	712
13.2.2	Deploying Rserve .....	712
13.2.3	Leveraging R and PAL to Produce Predictive Results .....	713
13.2.4	Installing SAP Predictive Analysis .....	714
13.2.5	User Privileges and Security with SAP Predictive Analysis .....	714
13.3	Integrating with SAP BusinessObjects BI .....	716
13.3.1	Exporting Scored Data Back to Databases .....	716
13.3.2	Exporting Algorithms .....	717
13.4	Case Study 1: Clustering Analysis .....	719
13.4.1	Preparing the Data .....	720
13.4.2	Performing Clustering Analysis .....	724
13.4.3	Implementing the Model .....	732
13.5	Case Study 2: Product Recommendation Rules .....	738
13.5.1	Preparing the Data .....	738
13.5.2	Performing Apriori Analysis .....	738
13.5.3	Implementing the Model .....	745
13.6	Summary .....	750
<b>14</b>	<b>Professionally Authored Dashboards with SAP HANA .....</b>	<b>751</b>
14.1	SAP HANA as a Data Source for SAP BusinessObjects Dashboards .....	754
14.2	SAP HANA as a Data Source for SAP BusinessObjects Design Studio .....	759
14.2.1	Connecting to SAP BW on SAP HANA .....	760
14.2.2	Connecting Directly to SAP HANA Data Sources .....	760
14.2.3	Connecting to the SAP HANA XS Engine .....	761
14.2.4	Consuming the SAP HANA Connections .....	763
14.3	Case Study: Exploring Data with SAP BusinessObjects Design Studio on Top of SAP HANA .....	764
14.3.1	Gathering Requirements .....	764
14.3.2	Laying Out the Components .....	765

14.3.3	Connecting to SAP HANA .....	766
14.4	Summary .....	769
<b>15</b>	<b>Data Exploration and Self-Service Analytics with SAP HANA .....</b>	<b>771</b>
15.1	SAP HANA as a Data Source for SAP BusinessObjects Explorer .....	772
15.1.1	Exploring and Indexing .....	773
15.1.2	Connecting SAP BusinessObjects Explorer to SAP HANA .....	776
15.1.3	Creating an Information Space on SAP HANA .....	778
15.2	SAP HANA as a Data Source for SAP Lumira .....	780
15.2.1	Online Connectivity .....	783
15.2.2	Offline Connectivity .....	784
15.3	Case Study: Exploring Sales Data with SAP Lumira on Top of SAP HANA .....	789
15.3.1	Business Requirements .....	790
15.3.2	Planned Solution .....	790
15.4	Summary .....	796
<b>16</b>	<b>SAP BusinessObjects Web Intelligence with SAP HANA .....</b>	<b>797</b>
16.1	Connecting SAP BusinessObjects Web Intelligence to SAP HANA .....	797
16.2	Report Optimization Features with SAP HANA .....	800
16.2.1	Usage of JOIN_BY_SQL .....	800
16.2.2	Merged Dimensions versus Analytic/Calculation Views .....	803
16.2.3	Query Drill .....	804
16.2.4	Query Stripping .....	806
16.3	Case Study: Exploring Sales Data with SAP BusinessObjects Web Intelligence on Top of SAP HANA .....	808
16.4	Summary .....	814
<b>17</b>	<b>SAP Crystal Reports with SAP HANA .....</b>	<b>815</b>
17.1	SAP HANA as a Data Source for SAP Crystal Reports .....	817
17.1.1	Configuring ODBC and JDBC Connections .....	819
17.1.2	Using SAP BusinessObjects IDT Universes .....	822
17.1.3	Using SAP BusinessObjects BI Relational Connections .....	824
17.1.4	Direct OLAP Connectivity to Analytic and Calculation Views .....	826

17.2	Case Study: Exploring Data with SAP Crystal Reports on Top of SAP HANA .....	831
17.2.1	Connecting to Data .....	832
17.2.2	Designing the Query .....	833
17.2.3	Limiting Query Results with Filter .....	834
17.2.4	Formatting the Report Display .....	835
17.3	Summary .....	836
	<b>Appendices .....</b>	<b>837</b>
A	Source System Analysis with SAP Information Steward Data Insight .....	837
A.1	Column Profiling .....	841
A.2	Address Profiling .....	843
A.3	Dependency Profiling .....	845
A.4	Redundancy Profiling .....	846
A.5	Uniqueness Profiling .....	847
A.6	Summary .....	848
B	The Authors .....	849
	Index.....	851

# Index

## A

- ACID compliance, 34
- Active Directory, 135, 139
  - configuration in *BOBJ*, 143
  - role mapping, 151
- Ad hoc reporting, 63, 64
- Administration Console perspective, 168
- Advanced analytics, 693
- AdventureWorks Cycle Company, 139, 340, 490, 515, 567, 764, 765, 789, 808, 831
  - analytic privileges, 569
  - fact table, 385
  - real-time loading, 389
  - SAP BusinessObjects Web Intelligence, 808
  - SAP Crystal Reports, 831
  - SAP Lumira, 790
  - tables in SAP HANA, 340
- Aggregate, 410, 781
  - all nodes, 441
  - awareness, 624
  - table, 170
- Algorithm
  - predictive, 692
  - segmentation, 719
  - supervised learning, 695
  - unsupervised learning, 695
- Analytic privilege, 102, 105, 536, 606
  - configure, 567
  - create traditional, 540, 541
  - dynamic values, 538, 546
  - Editor, 543
  - information views, 561, 564
  - SQL-based, 538, 556
  - static values, 538
  - traditional, 537
- Analytic view, 444, 623, 777
  - calculated column, 453
  - column parameter, 454
  - copy from, 447
  - create, 445, 508
  - Data Foundation node, 447, 448
  - Derived from table parameter, 454
  - direct parameter, 454
  - hidden column, 456
- Analytic view (Cont.)
  - input parameter, 457
  - Label column, 456
  - local attribute, 456
  - Logical Join node, 447, 450
  - measure, 444
  - properties, 447
  - save and activate, 459
  - Semantics node, 447, 455
  - static list parameter, 454
  - variable, 457
  - variable defined as, 458
- Application Function Libraries (AFL), 75, 706, 712
- Apriori algorithm, 707, 738, 741
- ArcGIS, 781
- Association algorithm, 707, 738, 741
- Attribute, 401
- Attribute view, 427, 428, 490, 627
  - calculated column, 434
  - create, 429, 493
  - Data Foundation node, 431
  - define, 496, 501, 502
  - derived, 507
  - filter, 434
  - hierarchy, 438
  - join types, 432
  - key column, 435
  - level hierarchy, 440
  - parent-child hierarchy, 440
  - primary key, 427
  - properties, 430, 431
  - save and activate, 442
  - Semantics node, 431, 438
- Attunity Connector, 203
- Authentication, 131, 137, 166, 605
  - AdventureWorks, 139
  - Kerberos, 135
  - methods, 132
  - SAML, 136
  - user name and password, 133
- Authorization, 100, 137
  - AdventureWorks, 147
  - data, 100
  - functional, 100



B

Batch data loading, 309  
    *business rules validation stage*, 358  
    *driver stage*, 348  
    *end*, 310  
    *end script*, 318, 387  
    *full data set comparison target-based CDC*, 321  
    *initialization*, 310, 311, 344  
    *lookup stage*, 352  
    *mart*, 310, 316, 374  
    *methods*, 319  
    *parsing stage*, 351  
    *source-based CDC*, 326  
    *staging*, 310, 313, 345  
    *standard target-based CDC*, 324  
    *steps*, 310  
    *triggers*, 328  
    *truncate and reload*, 320  
BEx Web Application Designer, 759  
Big data, 776, 817  
Bulk loading, 252, 319  
Business Function Library (BFL), 29, 408, 409, 707  
    *security*, 715

C

Calculation engine, 31, 657, 660  
Calculation view, 460, 624, 777  
    *add analytic view*, 520  
    *attribute columns*, 470  
    *calculated columns*, 471  
    *configuration options*, 462  
    *counters*, 471  
    *create*, 461, 518  
    *define*, 520  
    *Details pane*, 519  
    *example workflow*, 465  
    *final output*, 529  
    *graphical*, 466  
    *input parameters*, 471  
    *measure columns*, 470  
    *Output columns*, 519  
    *package*, 516  
    *Palette pane*, 519

Calculation view (Cont.)  
    *project analytic view*, 521  
    *Properties window*, 519  
    *restricted columns*, 471  
    *script-based*, 474  
    *scripting*, 474  
    *Tools Palette*, 467  
    *Union transform*, 525  
Cardinality, 347, 353  
CDC, 50, 319, 327  
    *source-based*, 324, 326, 347  
    *target-based*, 324  
CE\_ functions, 474  
Central Management Console (CMC), 610, 777  
Central processing unit (CPU), 32  
Classification algorithm, 707  
Clustering algorithm, 707, 719  
Columnar database architectures, 162  
Column-store table, 30, 42, 162, 163, 165, 167, 173, 628  
Compile server, 75  
Compression, 32, 167, 168, 314  
CPU  
    *cores*, 32  
CSV, 780  
Cube, 402, 445

D

Data  
    *caching*, 157  
    *discovery*, 701  
    *quality*, 53  
    *standardization*, 347  
Data flow, 244  
    *business rules validation stage*, 251, 348  
    *driver stage*, 246, 348  
    *lookup stage*, 248, 348  
    *parsing stage*, 247, 348  
Data Foundation node, 431, 432, 436, 448, 508, 617, 798  
Data loading, 309  
    *real time*, 335, 336  
Data mart, 54, 59, 62, 175, 176, 178, 179  
    *modeling*, 175, 186  
    *SAP HANA-specific*, 183

Error, 232  
    *handling*, 254  
    *logging*, 358  
ESRI, 781  
ETL, 38, 52, 53, 161, 171, 217, 306, 316  
Excel, 760, 780  
Exponential smoothing algorithm, 707  
Extract, transform, and load (see ETL)

F

Fact, 62  
Fact table, 159, 176  
    *create*, 189  
    *FactInternetSales example*, 185  
Federated Identity System, 136  
Flat file, 55  
Forecasting  
    *BFL algorithm*, 707  
    *PAL algorithm*, 707, 708  
Function  
    *date and time*, 169  
    *date extract*, 169  
    *logic driving*, 169  
    *mathematical*, 169  
    *numeric*, 169  
    *string*, 169

G

Geography, 781  
Geospatial engines, 29  
Grant, 106  
    *life cycle*, 106, 107  
    *privilege*, 106  
    *roles to users*, 130  
    *SQL-based analytic privilege*, 559  
    *statements*, 113  
Granularity, 178, 239

H

Hierarchy, 400, 401  
    *aggregate all nodes*, 441  
    *create*, 439  
    *level*, 440

Data mining, 693  
Data model, 53, 62, 175  
    *techniques*, 176  
Data provisioning, 45, 223  
    *business rules validation stage*, 250  
    *data flow*, 244  
    *data load stage*, 252  
    *datastores*, 227  
    *driver stage*, 246  
    *jobs*, 231  
    *lookup stage*, 248  
    *metadata*, 225  
    *parsing stage*, 247  
    *transform*, 255  
    *workflows*, 234  
Data source name (DSN), 593, 605  
Data warehouse, 53, 54, 59, 62  
Datastore, 227, 231  
    *configuration properties*, 228  
    *connection parameters*, 228  
Decision tree algorithm, 707  
Delta load process, 50  
Denormalization, 179, 180, 183, 347, 479, 483  
    *logical*, 484  
    *physical*, 484  
Descriptive analytics, 694  
DIM\_CUSTOMER, 379  
DIM\_DATE, 380  
DIM\_PRODUCT, 376  
DIM\_SALES\_TERRITORY, 380  
Dimension, 62, 400, 401, 428, 799  
Dimension table, 159, 160, 181, 242  
    *create*, 189  
Direct Extract Connector (DXC), 45, 49  
Disaster recovery (DR), 35  
Disparate systems, 219  
Document properties, 807, 813  
DRAM, 32  
Drilling, 805

E

EFFECTIVE\_PRIVILEGES, 567  
Enterprise information management (EIM), 28, 60

Hierarchy (Cont.)  
    *node style*, 440  
    *parent-child*, 440  
    *properties*, 440, 441  
    *view*, 626  
    *with root node*, 442  
High availability (HA), 35

I

IBM DB2, 46, 200, 203  
Information Design Tool (IDT), 777, 798, 824  
Information view, 96, 427, 485, 620, 628  
    *analytic privileges*, 561, 564  
In-memory database, 29, 36  
Input parameter, 454  
Ivy Bridge, 33

J

Java database connectivity (see JDBC)  
JDBC, 592, 798, 819, 822  
JDE One World, 203  
JDE World, 203  
Job  
    *batch*, 232, 233  
    *real time*, 233, 288  
Join  
    *inner*, 432, 451  
    *left outer*, 433, 451  
    *referential*, 433, 451  
    *right outer*, 433, 451  
    *star*, 473  
    *temporal*, 452  
    *text*, 434, 452  
Join engine, 31, 657, 658  
JOIN\_BY\_SQL, 800, 811

K

Kerberos, 132, 135, 137, 606  
    *AdventureWorks*, 146  
    *C:/WINNT*, 141  
Key performance indicator (KPI), 752, 764

L

Latency, 802  
Local connection, 615  
Logical Join node, 450

M

Machine learning, 694  
MDX, 164, 194, 408  
Measure, 400, 799  
Merged, 53  
Metadata, 54, 63, 225, 227, 230, 402  
Microsoft SQL Server, 46, 203  
Modeler perspective, 416  
MOLAP, 403  
Multidimensional model, 29, 42, 49, 399,  
    400, 402, 403, 479, 485  
    *benefits*, 404  
    *cube*, 402  
    *dimension*, 400  
    *hierarchy*, 400  
    *measure*, 400  
    *metadata*, 402  
    OLAP, 400  
    *star schema*, 400  
    *table*, 422  
    *transaction*, 400  
MySQL, 200

N

Name server, 75  
Nearest Neighbor algorithm, 707  
Node style, 440  
Normalized data, 159, 479, 480, 487  
    *challenges*, 481  
    *options*, 488

O

Object Linking and Embedding Database for  
    OLAP (OLE DB for OLAP), 408, 592  
ODBC, 592, 593, 798, 819

OLAP, 155, 157, 172, 175, 176, 771, 806  
    *connection*, 610  
    *data storage*, 159  
    *engine*, 31, 657, 659  
    *modeling*, 176, 180  
    *modeling in SAP HANA*, 180  
    *multidimensional OLAP*, 403  
OLTP, 155, 157, 172, 175, 248, 479, 480  
Online analytic processing (see OLAP)  
Online transaction processing (see OLTP)  
Open database connectivity (see ODBC)  
Operational reporting, 815, 816  
    *use case*, 815  
Oracle, 203  
Oracle Enterprise Edition, 46

P

Package, 423  
    *create*, 90, 424  
    *custom*, 89  
    *default*, 88  
    *example structure*, 90  
    *hierarchy*, 423  
    *properties*, 425  
    *root*, 492  
    *setup*, 87  
    *structure*, 568  
    *sub*, 492  
Parallel processing, 54  
Parallelization, 170  
PeopleSoft, 203  
Persistent layer, 165  
Planning engines, 29  
Polestar, 772  
Predictive analysis, 692, 703  
    *business case*, 705  
    *data discovery*, 701  
    *implementation*, 698, 702, 716, 717, 732,  
        745  
    *model development*, 702  
    *model evaluation*, 702  
    *process*, 696  
    *strategy*, 697  
    *tools*, 699, 706

Predictive Analysis Library (PAL), 29, 408,  
    409, 707  
    *implementation*, 719  
    *installation*, 712  
    *security*, 715  
Predictive model, 692, 702  
    *maintenance*, 702  
Predictive Model Markup Language (PMML),  
    718  
Preprocessing algorithm, 707  
Preprocessor server, 75  
Privilege, 101  
    *analytic*, 102, 105  
    *application*, 102, 104  
    *object*, 102  
    *on users*, 102  
    *package*, 102, 103  
    *system*, 102  
Profiling  
    *ad hoc*, 203  
    *advanced column*, 210  
    *basic column*, 208  
    *column*, 203, 205, 207, 219  
    *connection options*, 203  
    *relationship*, 203, 211, 214  
Project  
    *create*, 92  
    *setup*, 87  
    *share*, 94  
Push down, 225, 253

Q

QaaWS Designer, 73  
Query as a Service, 746  
Query Browser panel, 756, 757  
Query panel, 755  
Query processing engine, 31  
Quick Launch view, 416, 429

R

R language, 408, 410, 708  
    *implementation*, 728  
    *installation*, 712, 714  
    *security*, 715

Random access memory (RAM), 33  
RDBMS, 38, 57, 155, 170, 479  
Real time, 46, 234, 816  
Regression algorithm, 707  
Relational connection, 610  
Relational database management system (see RDBMS)  
Relational OLAP (ROLAP), 403  
Relationship profile, 212  
Replication  
    *real time*, 305  
Repositories view  
    *configuring*, 85  
Repository connection, 615  
Repository workspace, 86  
Restricted column, 453  
Role, 109, 166  
    *AdventureWorks*, 149  
    *as repository object*, 110  
    *CONTENT\_ADMIN*, 117  
    *create*, 109  
    *custom functional*, 118  
    *developer*, 120  
    *Editor*, 113  
    *end user*, 119  
    *ETL service account*, 125  
    *MODELING*, 117  
    *MONITORING*, 117  
    *PUBLIC*, 117  
    *RESTRICTED\_USER\_JDBC\_ACCESS/RESTRICTED\_USER\_ODBC\_ACCESS*, 117  
    *rights escalation*, 115  
    *SAP\_INTERNAL\_HANA\_SUPPORT*, 118  
    *scenarios*, 116  
    *security admin*, 122  
    *system admin*, 124  
Root node, 442  
Row engine, 31  
Row-store table, 162, 165, 172  
    *use case*, 173  
Rserve, 709  
    *installation*, 712  
Rules engines, 29

S

SAML, 136  
    *configuration*, 145

SAP (Sybase) IQ, 203, 789, 791  
    *Hilo.db*, 789  
SAP Assertion Tickets, 133  
SAP Business Suite, 57  
SAP Business Suite on SAP HANA, 47, 57, 58  
SAP BusinessObjects Analysis for OLAP, 610  
SAP BusinessObjects BI, 27, 28, 44, 59, 62, 70, 772  
    *Central Management Console*, 71  
    *database repositories*, 72  
    *Java web application server*, 71  
    *server architecture layer*, 71  
    *universe*, 63  
SAP BusinessObjects BI Launch Pad, 71  
SAP BusinessObjects BI Scheduler, 329  
SAP BusinessObjects Credential Mapping, 605, 777  
SAP BusinessObjects Dashboards, 66  
    *connect to SAP HANA*, 754  
    *direct binding*, 755  
SAP BusinessObjects Design Studio, 66, 759  
    *connect to SAP HANA*, 759  
SAP BusinessObjects Explorer, 772, 780  
    *connect to SAP HANA*, 772, 776  
    *development background*, 772  
    *exploration view set*, 772  
    *facet*, 773  
    *index*, 773  
    *index storage*, 775  
    *index structure and storage*, 775  
    *indexing*, 773  
    *indexing on SAP HANA*, 773  
    *Information Space*, 772, 773, 777, 778, 779  
    *Information Space on SAP HANA*, 778  
SAP BusinessObjects Web Intelligence, 797, 808  
    *connect to SAP HANA*, 797  
    *merged dimensions*, 803  
    *optimization for SAP HANA*, 800  
    *predictive model implementation*, 746  
    *query drilling*, 804, 813  
    *query stripping*, 637, 806, 813  
SAP BW, 30, 59, 760  
SAP BW Accelerator (BWA), 38  
SAP BW on SAP HANA, 38, 39  
SAP Crystal Reports, 66, 815, 817, 818, 821, 832  
    2013, 816  
    2013 vs. *Enterprise*, 817

SAP Crystal Reports (Cont.)  
    *64-bit architecture*, 817  
    *analytic view*, 826  
    *application database*, 816  
    *BEx query*, 819  
    *calculation view*, 826  
    *connect to SAP HANA*, 817  
    *connecting to data*, 832  
    *database middleware*, 819  
    *designing a query*, 833  
    *direct connection to SAP HANA*, 826  
    *externally facing report*, 816  
    *highly formatted report*, 815  
    *JDBC*, 818, 819, 821  
    *ODBC*, 818, 819, 820  
    *OLAP connectivity*, 826  
    *OLAP data sources*, 828  
    *relational connection*, 818, 819  
    *SAP HANA connection options*, 817, 818  
    *SAP HANA relational connections*, 824  
    *SAP HANA universe connections*, 822  
SAP Data Services, 27, 28, 44, 52, 59, 60, 67, 223, 399  
    *auto correct load*, 365, 377  
    *batch job*, 310, 341  
    *built-in functions*, 244, 277, 278, 280, 281  
    *business rules validation stage*, 246  
    *Case transform*, 263, 355  
    *Central Management Console*, 67  
    *checkpoint*, 239  
    *column profile results*, 202  
    *column profiling*, 205, 207  
    *custom functions*, 281  
    *data flow*, 244  
    *data flow processing*, 245, 247, 248, 251  
    *data load stage*, 246  
    *data loading*, 309  
    *Data\_Cleanse*, 270, 272, 337  
    *database repositories*, 68  
    *datastore*, 227  
    *driver stage*, 245  
    *Exec() function*, 393  
    *file format*, 285, 286  
    *information platform services*, 68  
    *Java web application server*, 67  
    *job*, 231  
    *job execution controls*, 312, 313  
    *job recovery*, 236, 237, 240

SAP Data Services (Cont.)  
    *job server*, 68  
    *job structure best practices*, 341  
    *Key\_Generation*, 365, 373  
    *lookup stage*, 245  
    *Lookup\_ext()*, 372  
    *Management Console*, 332  
    *Map\_Operation*, 265  
    *Match Editor*, 275  
    *Match transform*, 273, 274, 339  
    *merge*, 358  
    *metadata*, 225  
    *parallel execution*, 235  
    *parallel operations*, 236  
    *parsing stage*, 245  
    *preprovisioning data*, 197  
    *Query transform*, 255  
    *real-time job*, 288, 290, 332, 390  
    *relationship profiling*, 211, 213  
    *reusable object*, 244  
    *scheduling*, 328, 329  
    *script*, 311, 318  
    *script object*, 284  
    *series execution*, 240  
    *single-use object*, 243  
    *staging*, 313, 314, 315, 345, 346  
    *system configuration*, 227  
    *Table\_Comparison*, 261, 365, 366, 368, 373  
    *transform*, 244, 255, 257, 258, 260, 262, 263, 265, 267, 268  
    *Try/Catch*, 342  
    *Validation transform*, 267, 360  
    *workflow*, 234  
SAP Data Services BI Scheduler, 329  
SAP Data Services Designer, 69, 204, 212, 223, 224  
SAP Data Services Repository Manager, 69  
SAP Data Services Scheduler, 328  
SAP Data Services Server Manager, 69  
SAP Data Services Workbench, 69, 290, 291, 293, 299  
    *data flow*, 293  
    *existing data warehouse*, 297  
    *modifying data flows and jobs*, 304  
    *porting data*, 297, 303  
    *Quick Replication Wizard*, 303  
    *supported transforms*, 295  
SAP Event Stream Processor (ESP), 45, 51

SAP Governance, Risk, and Compliance (GRC), 126

SAP GUI, 760

SAP HANA, 28, 42, 73, 751, 754, 771, 772, 780, 815, 822

- analytics appliance*, 28
- architecture*, 74
- calculation engine*, 453, 626
- client*, 592
- complex transformations*, 45
- custom solutions*, 45
- direct connection*, 760
- Extended Application Services*, 410
- hardware*, 36
- join engine*, 626, 629
- middleware*, 760
- multinode*, 35
- native performance*, 45
- OLAP connection*, 610
- OLAP engine*, 434, 626
- online*, 783
- performance*, 156
- R integration*, 708, 720
- real-time replication*, 45
- repository*, 79
- schema*, 420
- security*, 77
- sizing*, 30
- software layers*, 28
- table*, 420, 422
- third-party data*, 44
- web application server*, 37

SAP HANA Live, 47, 57

SAP HANA Studio, 45, 55, 413, 419, 781

- content*, 418
- Development perspective*, 81, 82
- hierarchies*, 442
- Modeler perspective*, 416
- Navigator view*, 416
- perspectives*, 414, 415
- Quick Launch view*, 416, 418
- security*, 80, 418
- supported OSs*, 414
- system object*, 417
- Systems view*, 417

SAP HANA XS, 74, 137, 410, 411

SAP InfiniteInsight, 63, 691

SAP Information Steward, 269

- address profiling*, 843
- column profiling*, 841
- data cleansing*, 843
- dependency profiling*, 845
- metadata*, 837
- profiling*, 840
- redundancy profiling*, 846
- uniqueness profiling*, 847
- validation rule*, 839

SAP Landscape Transformation (SLT), 45, 201

SAP Logon Ticket, 133

SAP Lumira, 709, 772, 780

- AdventureWorks*, 790
- connect to SAP HANA*, 780
- connectivity options*, 785
- forecast*, 781
- linear regression*, 781
- LUMS*, 788
- offline*, 783, 784
- online*, 783
- time-related hierarchy*, 792

SAP MaxDB, 46

SAP Predictive Analysis, 63, 700, 709, 720, 738

- clustering analysis*, 724
- installation*, 714
- SAP HANA security*, 714

SAP Replication Server, 45, 52

SAP Sybase, 788

- iqsrv15.exe*, 788

SAPUI5, 88

Scheduling

- execution command*, 332
- third party*, 334

Schema, 95, 165, 187, 188, 316, 317, 420

- as repository object*, 97
- column views*, 422
- create*, 99, 187
- create with user*, 96
- procedures*, 422
- properties*, 95
- setup*, 95
- tables*, 421
- views*, 421

Scope of analysis, 813

Script server, 75

Scrum, 751

Security, 77

- adding system*, 83
- authentication*, 131
- authorizations*, 100
- Business Function Libraries (BFL)*, 715
- configuring*, 78
- credentials*, 132
- data security*, 535
- Predictive Analysis Library (PAL)*, 715
- SAP HANA-R integration*, 715
- SAP Predictive Analysis*, 714
- SSL*, 144
- user and role provisioning*, 108

Security Assertion Markup Language (see SAML)

Self-service analytics, 63, 64, 771

Semantics node, 431, 438, 455, 459, 473, 514

Siebel, 203

Single sign-on (SSO), 132, 606, 623

Smart Data Access, 31

Source system analysis (SSA), 61, 197, 198, 200, 837

- example*, 198
- field misuse*, 218
- field overuse*, 218
- mapping*, 220
- mapping document*, 215, 216, 220, 221
- mappings*, 200
- multiple sources*, 219
- pattern*, 217
- technique*, 202

Spinning disk, 157

SQL, 409, 780

- analytic privileges*, 538
- engine*, 629
- Procedure Editor*, 550
- statement*, 811
- wizard*, 548

SQL Data Definition Language (DDL), 185

SQLScript, 49, 474

- Application Function Libraries*, 707

Star schema, 159, 400

Statistical models, 694

Stored procedure, 49

\_SYS\_BIC, 409, 537, 620, 623, 798

\_SYS\_REPO, 96, 111

SYSTEM, 96, 111, 148, 150

- disable*, 150

System Landscape Transformation (SLT), 45, 307, 308, 816

Systems view, 429

SysWOW64 directory, 820

## T

Table, 175

- modeling*, 175, 186

Tailored Datacenter Integration, 36

Teradata, 203

Time series algorithm, 707, 708

Transaction, 400

Transform, 53, 251, 255

- Aggregate*, 468
- Case*, 263
- data quality*, 257
- Data\_Cleanse*, 269
- Join*, 467
- Map\_Operation*, 265
- Match*, 273
- platform*, 257
- Projection*, 467
- Query*, 226, 326
- Table\_Comparison*, 261, 325
- text data processing*, 258
- Union*, 468
- Validation*, 267

## U

Universe, 63, 772, 776, 780, 822

- business layer*, 632, 798
- Query panel*, 833
- UNIX*, 818

Universe Designer, 72

UNIX, 819

User

- restricted*, 127

User provisioning, 126

- AdventureWorks*, 150
- automation*, 128
- manual*, 129



V

---

Variable, 318  
Visualization, 710

W

---


Web service, 331, 746  
    *scheduling*, 330

Workflow  
    *parallel execution*, 234  
    *reusability*, 242  
    *series execution*, 240

X

---

X509, 133



Jonathan Haun, Chris Hickman, Don Loden, Roy Wells

# Implementing SAP HANA

860 Pages, 2015, \$79.95/€79.95  
ISBN 978-1-4932-1176-0

 [www.sap-press.com/3703](http://www.sap-press.com/3703)



**Jonathan Haun** currently serves as the lead SAP HANA consultant and consulting manager with Decision First Technologies. He is an SAP Certified Application Associate and SAP Certified Technology Associate for SAP HANA 1.0.



**Chris Hickman** is a certified SAP BusinessObjects BI consultant and principal consultant at Decision First Technologies. His specific areas of expertise include reporting, analysis, dashboard development, and visualization techniques.



**Don Loden** is a principal consultant at Decision First Technologies with full lifecycle data warehouse and information governance experience in multiple verticals. He is an SAP Certified Application Associate on SAP BusinessObjects Data Integrator and is very active in the SAP community, speaking globally at numerous SAP and ASUG conferences and events.



**Roy Wells** is a principal consultant at Decision First Technologies, where he uses his 15 years of experience in system and application architecture to lead clients in the successful implementation of end-to-end BI solutions.

*We hope you have enjoyed this reading sample. You may recommend or pass it on to others, but only in its entirety, including all pages. This reading sample and all its parts are protected by copyright law. All usage and exploitation rights are reserved by the author and the publisher.*