HA200

SAP HANA 2.0 SPS02 - Installation and Administration

PARTICIPANT HANDBOOK
INSTRUCTOR-LED TRAINING

Course Version: 14
Course Duration: 5 Day(s)
e-book Duration: 36 Hours 50 Minutes
Material Number: 50144114

SAP Copyrights and Trademarks

Typographic Conventions

American English is the standard used in this handbook.

The following typographic conventions are also used.

This information is displayed in the instructor's presentation	
Demonstration	
Procedure	
Warning or Caution	
Hint	
Related or Additional Information	
Facilitated Discussion	
User interface control	Example text
Window title	Example text

Contents

Course Overview

TARGET AUDIENCE

This course is intended for the following audiences:

- System Administrator

UNIT 1 SAP HANA Introduction

UNIT OBJECTIVES

- Understand the need for SAP HANA
- Find SAP HANA Information Sources

Unit 1
Lesson 1

SAP HANA Introduction

LESSON OVERVIEW

In this lesson, you will learn about SAP HANA components.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Understand the need for SAP HANA

SAP HANA as a Platform

SAP HANA is an in-memory, column-oriented, relational database management system. It combines online analytical processing (OLAP) and online transaction processing (OLTP) in one database system. Its primary function as database server, is to store and retrieve data as requested by the applications. In addition, it performs the following advanced analytical functions:

- Predictive analytics,
- Spatial data processing
- Text analytics
- Text search
- Streaming analytics
- Graph data processing

SAP HANA also includes ETL capabilities as well as an application server.

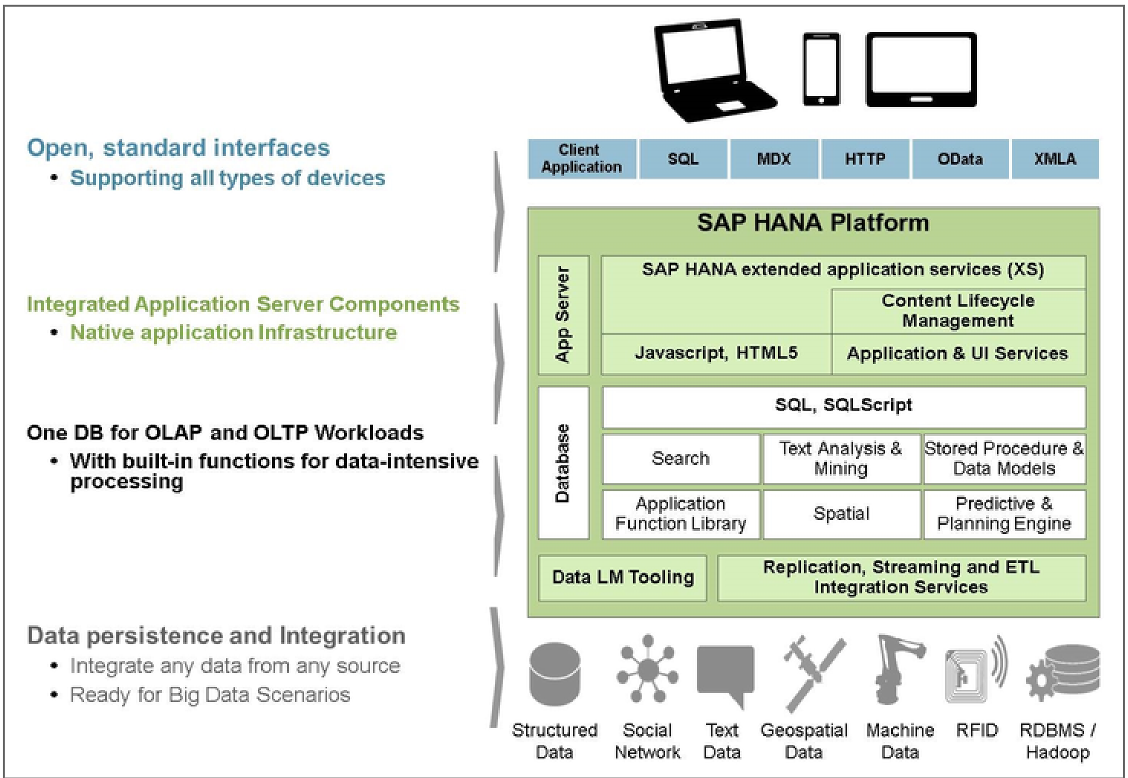

Figure 1: One Platform for Any Kind of Application

SAP HANA Design

SAP HANA is different by design. It stores all data in-memory and in a compressed columnar format.

Because SAP HANA is so fast, it does not require sums, indexes, materialized views, or aggregates. This reduces the database footprint. Everything is calculated on-demand and in main memory. This process makes it possible for companies to run online transaction processing (OLTP) and analytics applications (OLAP) on the same instance at the same time. It allows for any type of real time, specific queries, and analyses.

In addition, SAP has built solutions to all of the problems of columnar databases, such as concurrency (SAP HANA uses MVCC), and row-level insert and update performance (SAP HANA uses various mechanisms, such as a delta store). SAP also added engines inside SAP HANA to provide the following functions:

- Virtual online analytical processing (OLAP)
- Data virtualization
- Text analysis
- Search
- Geospatial
- Graph
- Web

SAP HANA supports the following open standards:

- Representational State Transfer (REST)
- JavaScript Object Notation (JSON)
- OLE DB for OLAP (ODBO)
- MultiDimensional eXpressions (MDX)
- Open Database Connectivity (ODBC)
- Java Database Connectivity (JDBC

Daily Irritations

The existing technology was not designed for modern challenges and use cases, as shown in the figure, Daily Irritations . Long-running transactions cannot keep pace with the speed of information.

First and foremost, users need a new technology platform. For example, they need a unified, low latency, and low complexity platform that supports real time business requirements.

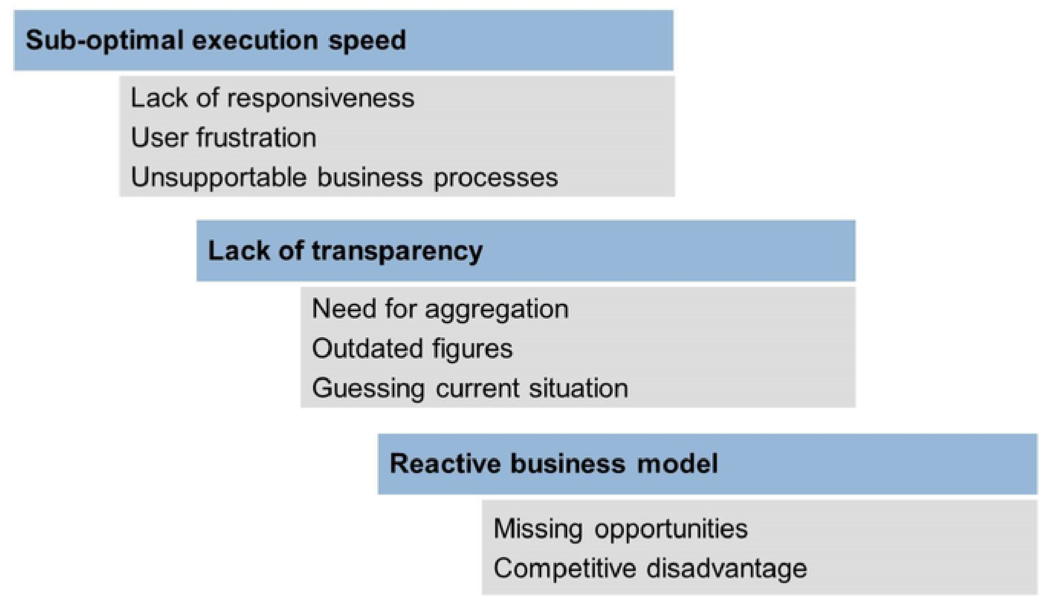

Figure 2: Daily Irritations

In-Memory Computing

Improved hardware economics and software innovations have made it possible for SAP to deliver on its vision of the Real Time Enterprise with in-memory business applications.

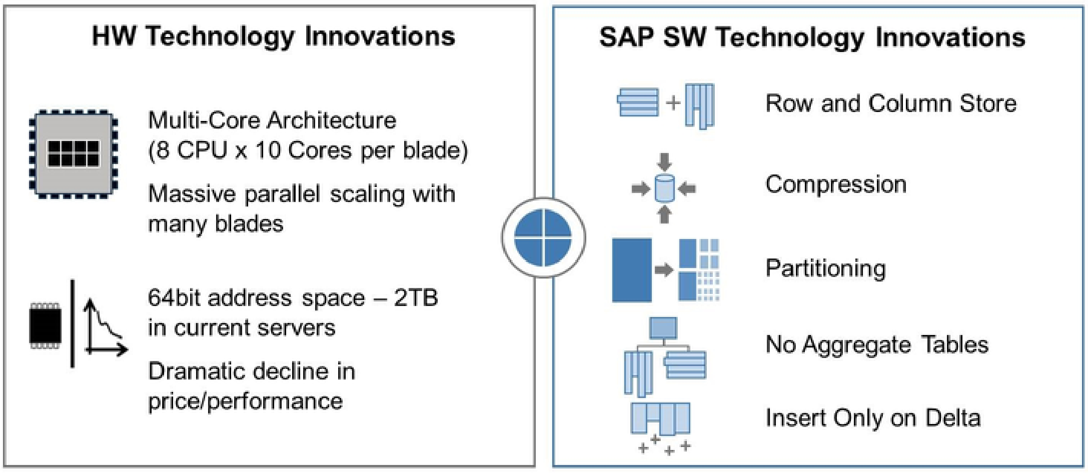

Figure 3: What Is In-Memory Computing?

The Past Disk-Centric, Singular Processing Platforms

Increased data volume causes major bottlenecks in data transfers. For example, input/output transfer rates from storage disks to servers have not kept up with data volumes. Disk-centric computing creates significant bottlenecks in data management. As a result, users are experiencing slow online transactions and batch processes.

The Future: Low Latency Computing Driven by In-Memory Technology

To overcome performance bottlenecks, IT systems have added complex deployment architectures that have compromised business user flexibility and added significant cost.

It is necessary, therefore, to use recent innovation to build software that takes key characteristics into its design principles.

The following list outlines some unique features of in-memory technology:

- Stores considerable amounts of information compressed in main memory
- Uses parallel processing on multiple cores
- Moves data-intensive calculations from the applications layer into the database layer for even faster processing

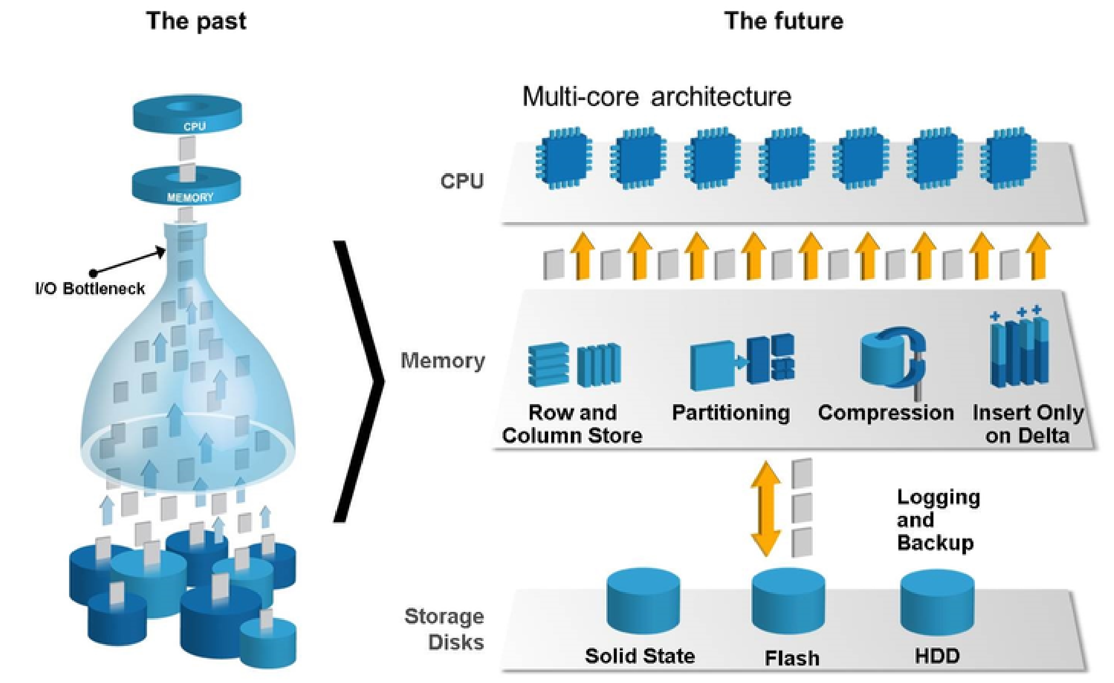

Figure 4: The Future: Low Latency Computing Driven by In-Memory Technology

Because all the detailed data is available in main memory and processed without interruption, there is no need for aggregated information and materialized views. This approach simplifies the architecture and reduces latency, complexity, and cost. In addition, with new multicore multithreaded processors, 64-bit address space, and advancement in parallel data processing, scalability is significantly improved.

Row vs Column Store

SAP HANA is an ACID-compliant, in-memory database. ACID is an acronym that means the database can support Atomicity, Consistency, Isolation, and Durability (ACID). This is a data requirement of a database, which must prove that it is 100% reliable for mission critical applications. The database must guarantee data accuracy and integrity even when there are lots of simultaneous updates across multiple tables.

The traditional database systems focus on one workload OLTP or OLAP. With SAP HANA this has changed, as it handles transactional and analytical workload very well. The SAP HANA database stores the data in a columnar way, therefore organizing the data in DRAM in an optimal way for the CPU to access.

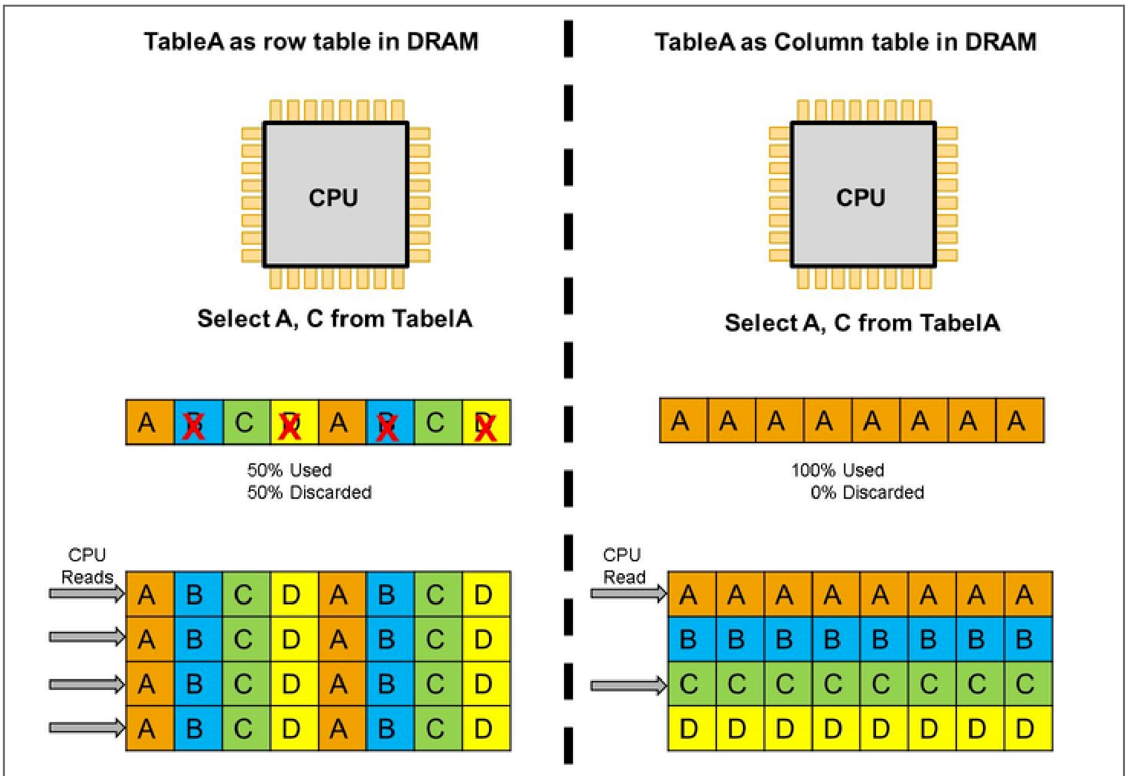

Figure 5: Row vs Column Store Tables RAM

Queries from analytical applications that are sent to the database often require only a subset of the overall data in the table. Usually only limited columns are required. With the column store, only the required columns are searched, so you avoid unneeded searching thru memory.

With column store, SAP HANA scans columns of data so quickly that additional indexes are usually not required. This helps to reduce the complexity by avoiding the need to constantly create drop and rebuilding indexes. Column store tables are optimal for parallel processing, as each core is able to work on a different column.

The column store is seen as optimal for analytical processing, but with inclusion of the delta store, the column store has also optimal performance for transactional processing. A delta store is added to every column table and is write optimized. In this way a columnar table is fast for read and write operations.

SAP HANA Deployment View

SAP HANA can be deployed or installed as a Single Host or a Multi Host system. A Single Host system delivers all the SAP HANA features and performance capabilities, but has no high availability features. To make SAP HANA more resilient, install the SAP HANA system as a Multi Host system using multiple compute nodes and one or more standby nodes.

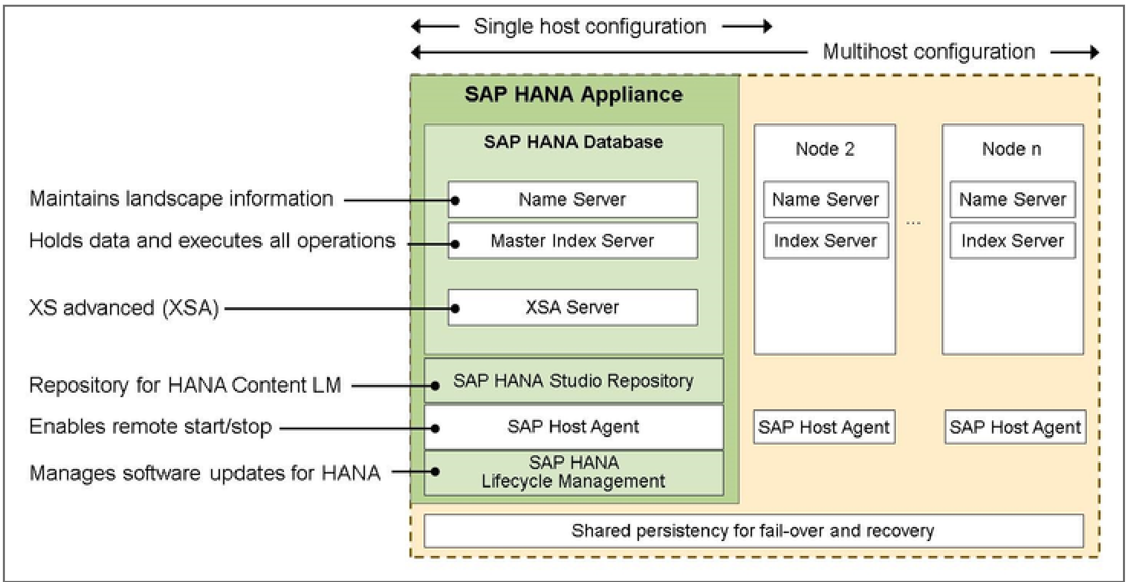

Figure 6: SAP HANA Deployment View

Component Architecture View: Application Function Libraries (AFL)

The SAP in-memory computing engine offers various algorithms for in-memory computing. It provides several application libraries for developers, partners, and customers who develop applications that run on SAP HANA. The libraries are linked to the SAP HANA database kernel.

The Business Function Library (BFL) is one of these application libraries. It contains prebuilt, parameter-driven functions in the financial area. The functions are implemented by C++. The functions include, but are not restricted to, the following:

- Forecast functions
- Max value
- Inflate cash flow function

> Note:
> For more detailed information on the supported functions see the Business Function Library (BFL) guide.

This library helps you to develop compound business algorithms that are fully compliant with the SAP HANA calculation engine. It offers you the flexibility and efficiency to develop HANA-based applications with incredible performance.

For a complete list of the available Application Function Library, see SAP Note 2198403 : AFL and corresponding components (AFL) components.

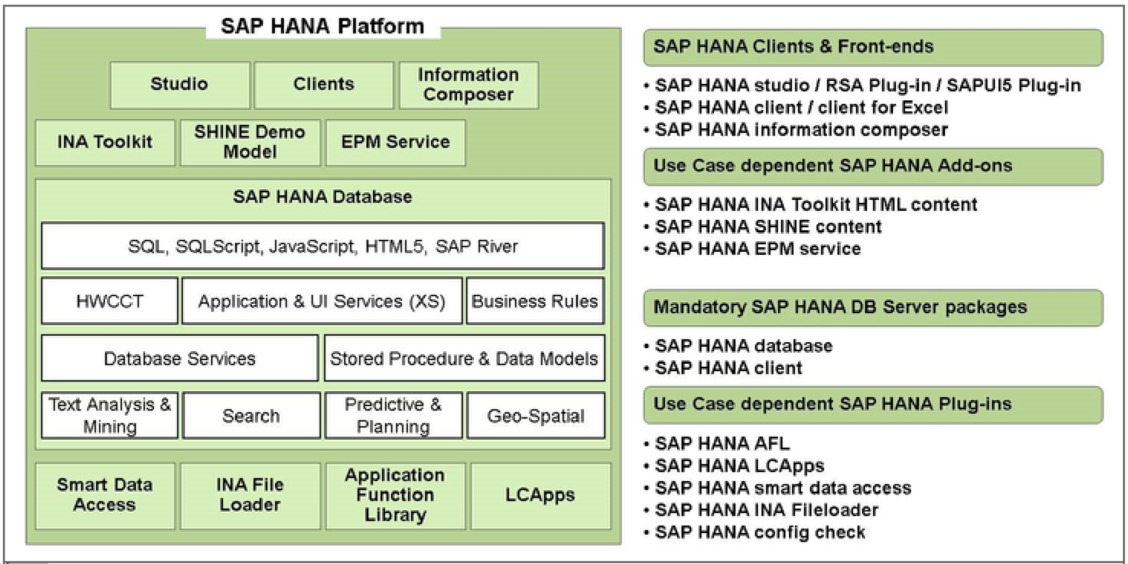

Figure 7: Component Architecture View

SAP HANA Info AccessToolkit HTML Content

The SAP HANA Info Access toolkit is a predefined, single-purpose HTML5 template that is configured to access data through attribute views. It provides a somewhat static search UI. The toolkit exposes a limited range of configuration options, but has no options for programmatic extensions. As such, it is most suitable for rapid prototyping and demonstrating, or for feature exploration purposes.

The toolkit enables a freestyle search of an SAP HANA attribute view, and it displays and analyzes the result set. The toolkit provides UI elements (widgets), such as a search and filter box, a result list with a detailed view, and charts for basic analytics on the result set. The widgets are interconnected and they adapt immediately to user entries and selections.

The toolkit is based on HTML5 and JavaScript libraries such as JQuery/JQueryUI, d3 (Data Driven Documents), and Tempo. The widgets use the SAP HANA info access HTTP service. You do not need an additional layer to run the UI: SAP HANA and a Web browser are sufficient.

SAP HANA Interactive (SHINE)

SAP HANA Interactive Education , or SHINE, is a demonstration application that teaches users how to build native SAP HANA applications. The demonstration application, delivered with SAP HANA in a special delivery unit (DU), includes the following features:

- Sample data and design-time developer objects for the application's database tables
- Data views
- Stored procedures
- OData
- A user interface

For more information about SHINE, see the following file: https://help.sap.com/hana/SAP_HANA_Interactive_Education_SHINE_en.pdf .

Enterprise Procurement Model (EPM)

The Enterprise Procurement Model is a framework developed by SAP. It includes all the data models, tables, views, dashboards, and so on, with a real enterprise use case.

Application Function Library (AFL)

The Application Function Library includes the Predictive Analysis Library (PAL) and the Business Function Library (BFL).

Predictive Analysis Library (PAL)

The Predictive Analysis Library (PAL) defines functions that can be called from within SQLScript procedures to perform analytic algorithms. This release of PAL includes classic and universal predictive analysis algorithms in the following data-mining categories:

- Clustering
- Classification
- Regression
- Association
- Time Series
- Preprocessing
- Statistics
- Social Network Analysis
- Miscellaneous

File Loader

The file loader is a set of HTTP services that you can use to develop your own applications for searching in file contents. The file loader package also contains a basic example application with monitoring and statistical information about the current file loader schedule.

SAP HANA Hardware Configuration Check

The SAP HANA hardware configuration check tool allows you to measure the performance of your hardware components so that hey meet the criteria for running SAP HANA.

To download the latest version of the tool as a SAR file from the SAP Service Marketplace, follow the instructions in SAP Note [1943937: Hardware Configuration Check Tool - Central Note].

SAP S/4HANA

x BW410_14: represents SAP Business Suite 4 SAP HANA, where "S" stands for both Simple and Suite. The "4" stands for fourth generation.

SAP S/4HANA is the next generation business suite. It is a new product, fully built on the SAP HANA platform and designed with the SAP Fiori user experience. SAP S/4HANA delivers significant simplifications in adoption, data model, user experience, decision-making, and business processes. It also provides innovations for the Internet of Things, big data, business networks, and mobile first, which will help to reinvent businesses. SAP S/4HANA brings the next series of innovation to customers, similar to the transition from SAP R/2 to SAP R/3.

SAP S/4HANA vs SAP Business Suite Powered by SAP HANA

What is new in SAP S/4HANA compared to SAP Business Suite powered by SAP HANA?

With SAP Business Suite powered by SAP HANA, SAP is the only software vendor allowing SAP Business Suite customers to bring together transactions and analysis into a single in-memory platform. This innovation has been extremely successful: In less than 2 years, over

1,850 existing and new customers acquired SAP Business Suite powered by SAP HANA to run their business in real time. This made it one of the fastest-growing products in the history of SAP.

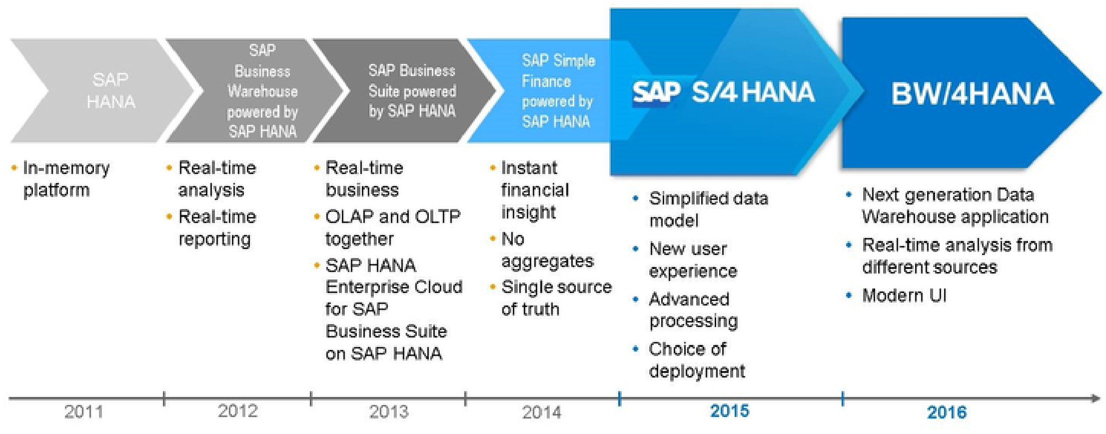

Figure 8: S4 HANA Overview

With SAP Business Suite powered by SAP HANA, the product approach has been to port the applications on the SAP HANA platform and optimize the code. This allowed customers to improve performance in their mission-critical business processes and reporting activities, and improved performance on relational databases. SAP HANA represented a new database alternative for existing customers, with a simple database migration required to drive the entire business in real time.

SAP S/4HANA creates unique opportunities to run your day-to-day business in real-time, with industry best practices, and also to reinvent business models and drive new revenues.

SAP S/4HANA the New Business Suite

SAP S/4HANA is a new business suite of applications. It is built to drive instant value across enterprises, industries, lines of business, data, regions, and deployments, with greater simplicity: The following list outlines the key features of the SAP S/4HANA application suite:

- **Reimagined User Experience**

  SAP have built the S/4HANA platform to suit the way business users work, with a role-based, consistent user experience available on any device.

- **Reimagined business models**

  SAP S/4HANA simplifies real time connections to people, devices, and business networks, so that you can deliver new experiences and value to customers. The Internet of Things and big data become accessible to any business, eliminating the need for more complex business collaboration and interactions.

- **Reimagined business processes**

  SAP S/4HANA provides simplicity to focus on the essential tasks in real time and to gain flexibility and agility so that you can change business processes as needed for new efficiencies.

- **Reimagined Data Model**

  SAP S/4HANA is purposefully built to offer the best level of responsiveness and performance with the lowest data footprint.

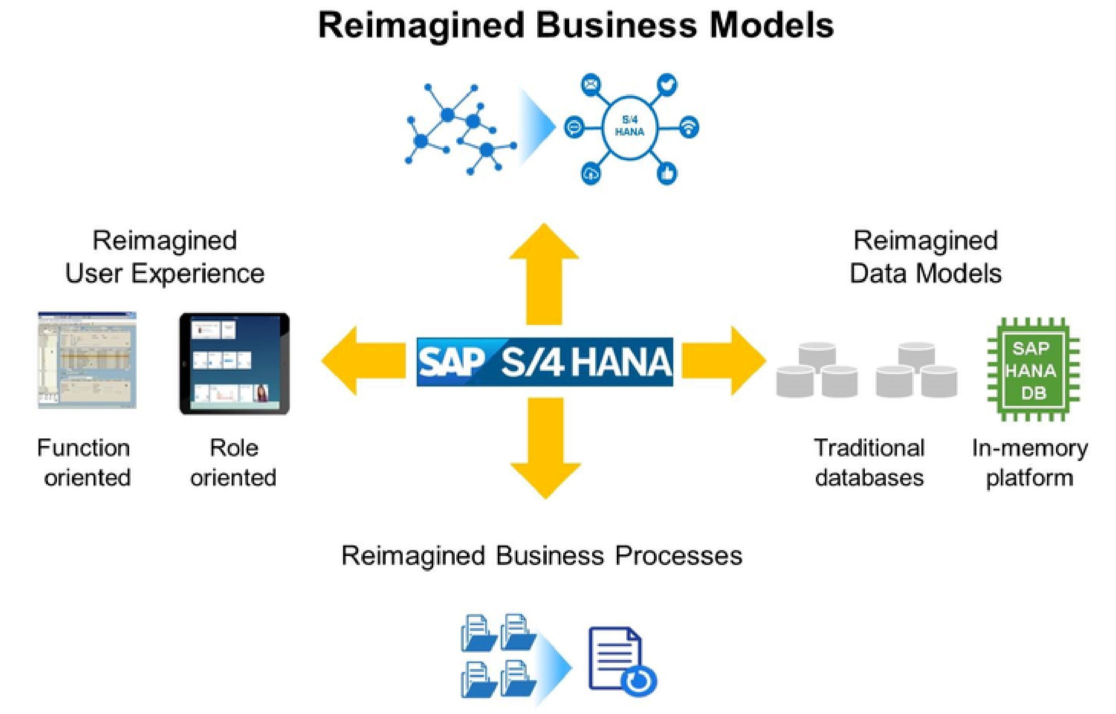

Figure 9: SAP S/4HANA the New Business Suite

From an IT value perspective, this means that SAP S/4HANA creates unique opportunities to simplify the landscape and reduce the total cost of ownership, with SAP HANA as the main driver.

Enterprises can now significantly reduce their data footprint and work with larger data sets in one system (for example, SAP ERP, SAP CRM, SAP SRM, SAP SCM, and SAP PLM in one system). This saves hardware costs, operational costs, and time. Enterprises no longer face discrepancies between different systems, with one common source of live data on one system.

Innovation is also simplified by an open platform (the SAP HANA Cloud Platform) that drives advanced applications, such as predicting, recommending, and simulating, while also protecting existing investments.

Business users can use a simple, role-based user experience based on modern design principles, to minimize training efforts while also increasing productivity. Note that clients are also supported with simple configuration, such as setting up the system, and so on, during usage.

Enterprises have a choice of deployment to drive easy adoption: cloud, on-premise, and even hybrid.

SAP S/4HANA is built on SAP HANA because only the SAP HANA platform can deliver such a significant level of simplifications and innovations.

Example: SAP S/4HANA Simple Finance

This new type of architecture has been piloted in Simple Finance and adopted for Simple Logistics. It works and performs for SAP HANA, which is the target architecture for all other simplified components.

Previously, indices and total tables were created to prevent systems from continuous calculation. This ensured overall system performance, but increased complexity and inflexibility. Both have now been eliminated from the system.

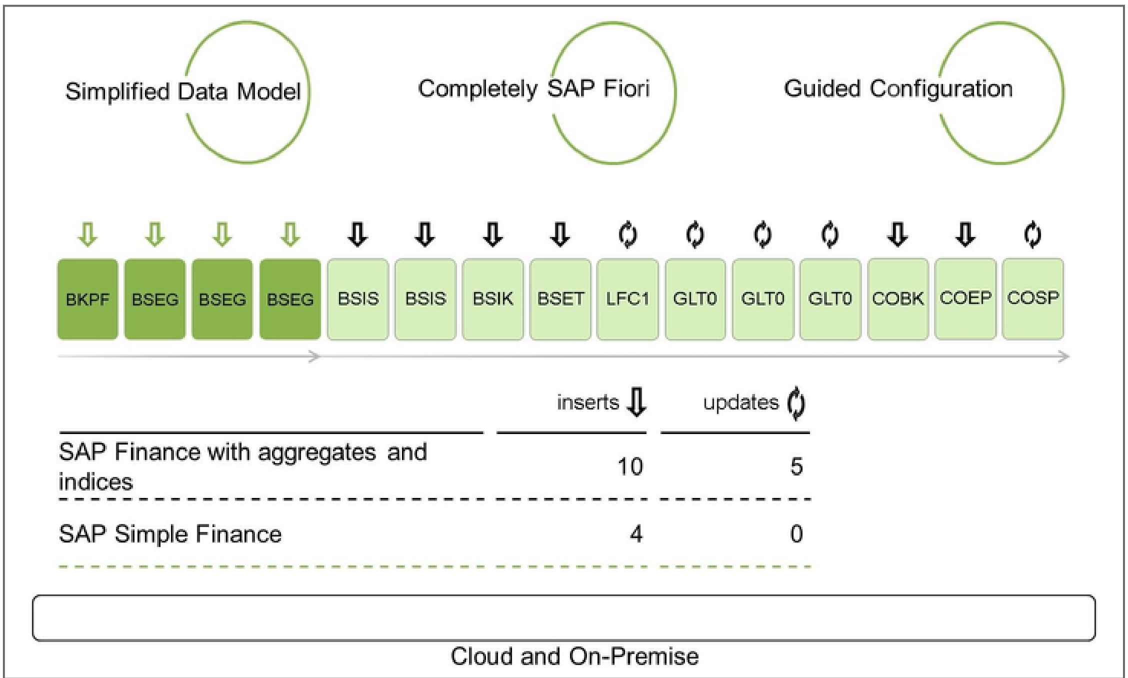

Figure 10: SAP S/4HANA Simplified Data Model

A typical booking in Financials touched 15 tables; now it is 4, working on document level.

On the data side, the underlying data structure is simplified to produce a highly performing and optimized system. The underlying data structure is based on no indices, aggregates or redundancies.

SAP S/4 HANA Benefits

SAP S/4 HANA includes the following key benefits:

- Smaller total data footprint
- Higher throughput
- Faster analytics and reporting
- Fewer process steps
- No locking, parallelism
- Actual data (25%) and historical (75%)
- Predict, recommend, simulate
- SAP HANA Cloud Platform extensions
- SAP HANA multitenancy
- All data: social, text, geo, graph, processing
- New SAP Fiori User Experiences for any device (mobile, desktop, tablet)

SAP HANA and SAP BW

Does SAP HANA Replace SAP BW? No, they complement each other. There is no plan to retire SAP BW.

SAP BW is better on SAP HANA, because SAP BW is free. There is plenty of prebuilt content for SAP BW, and you can have instant, certified solutions on top of SAP BW.

Many SAP BW customers have SAP Business Warehouse Accelerator to increase speeds in the SAP BW disk-based, relational database management system. SAP HANA provides a much simpler landscape, reducing total cost of ownership and complexity. It reduces your hardware footprint dramatically. For example, to accelerate 5 TB of SAP BW data, requires 21 blades versus one server in HANA. It has the added benefit of no third party database, because SAP HANA is the single persistent database.

SAP HANA is many things (a database for SAP BW, a high-performance analytical appliance, a platform for new applications), but matching the entire system, known as SAP BW, point-for-point is a huge project for any company.

LESSON SUMMARY

You should now be able to:

- Understand the need for SAP HANA

Unit 1
Lesson 2

SAP HANA Information Sources

LESSON OVERVIEW

In this lesson, you learn how to find information sources and guidelines for SAP HANA.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Find SAP HANA Information Sources

The Most Important Information Sources for SAP HANA

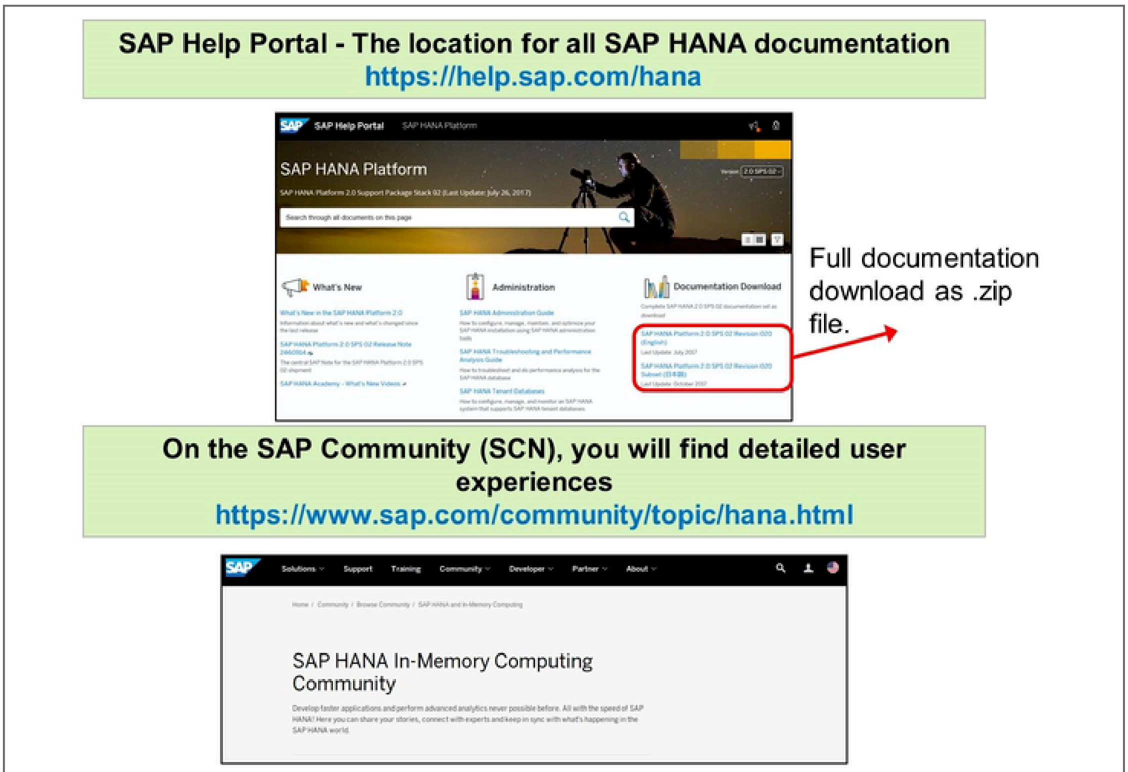

Figure 11: The Most Important Information Sources for SAP HANA

The SAP documentation website https://help.sap.com/hana contains the official SAP HANA documentation. This is the best entry point when looking for information regarding SAP HANA. The SAP community network website https://www.sap.com/community/topic/hana.html contains lots of SAP HANA-related information, written by SAP and non-SAP administrators, developers, and consultants.

A full set of the current SAP HANA documentation can be downloaded from https://help.sap.com/hana as a zip file.

Administration guides

All the important system administration guides, including those for installing and updating SAP HANA, can be found on https://help.sap.com/hana. The following guides are a selection of what can be found on the SAP documentation website:

- The **SAP HANA Master Guide** is the entry point for planning the installation of your SAP HANA system landscape.
- The **SAP HANA Server Installation and Update Guide** explains how to install and update an SAP HANA system with the SAP HANA lifecycle management tools.
- The **SAP HANA Administration Guide** explains how to configure, manage, maintain, and optimize your SAP HANA installation using SAP HANA administration tools.
- The **SAP HANA Security Guide** explains how to ensure the security of the SAP HANA platform and its components.
- The **SAP HANA SQL and System Views Reference guide** describes the SQL features supported by SAP HANA, including option-specific SQL used by various SAP HANA options.
- The **SAP HANA Developer Guide** for SAP HANA explains how to build and deploy applications that run in the SAP HANA XS Advanced Model runtime.

SAP HANA Documentation of Previous Releases

On https://help.sap.com/hana, you can also find the documentation for older SAP HANA releases. To reach the download location of the older SAP HANA documentation, follow the path shown in the figure, Finding the Older SAP HANA Documentation.

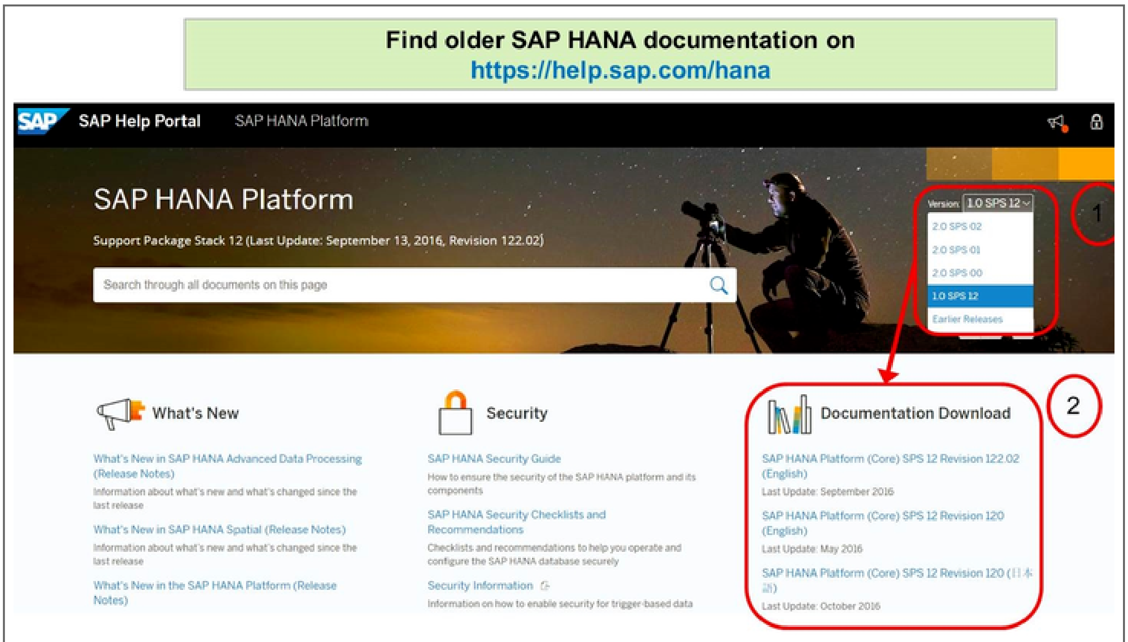

Figure 12: SAP HANA Guides

To Find Older SAP HANA Documentation

To find the SAP HANA documentation for previous releases, proceed as follows:

1. From the Version dropdown, choose the SAP HANA 1.0 SPS12 version.
2. Locate the Documentation Download section, and select the version you require.

SAP HANA Content Location

Table 1: SAP HANA Content Location

Content	Location
Latest documentation for SAP HANA	https://help.sap.com/hana https://help.sap.com/hana_platform
SAP HANA Quick Sizer with sizing guidelines.	https://service.sap.com/sizing
The SAP HANA Certified Hardware Directory.	https://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/index.html
Collection of all security guides of SAP applications	https://service.sap.com/securityguide
The SAP Software Download Center	https://support.sap.com/swdc
SAP Community Network	https://www.sap.com/community/topic/hana.html
Deployment Options	https://hana.sap.com/deployment.html
A good entry point for information on SAP HANA	https://hana.sap.com/

SAP HANA Hardware Directory

The SAP HANA hardware directory lists all hardware that have been certified or are supported under the following scenarios:

- Hardware that has been certified within the SAP HANA hardware certification program
- Previously validated hardware based on Westmere technology, as reflected earlier in the Product Availability Matrix (PAM).
- Supported Intel ®Systems: Only single node systems with minimum 8 cores per supported Intel architecture and valid for particular SAP HANA SPS releases.

> Note:
> The SAP HANA hardware directory can be found at: https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/index.html.

The certification is valid for the period stipulated in the Integration Certification Agreement. The exact validity date of the certification is indicated in the certification letter that is issued upon successful completion of testing.

The certification confirms the existence of product features in accordance with SAP certification procedures. SAP only certifies successful integration of the product with SAP solutions. The vendor is responsible for the product itself and its error-free operation.

The hardware is required to have a valid SAP HANA hardware certification at the point of purchase by the customer. Once the validity date of the certification has passed, the hardware continues to be supported by the Partner until the end of maintenance, as indicated by the Partner.

SAP HANA setup scenarios

Consistent documentation for the group of appliances is applicable for all SAP HANA setup scenarios, including single node, scale up, and scale out, as follows:

- **Scale Up**

  BWoH/BW4H/DM/SoH/S4H includes hardware approved for all single server configuration scenarios for SAP BW, powered by SAP HANA, S/4HANA .

  SoH/S4H includes additional single server configurations specific for SAP Business Suite powered by SAP HANA and S/4HANA, not covered by Scale-up: BWoH/BW4H/DM/SoH/S4H .

- **Scale Out**

  BWoH/BW4H/DM includes multi-server configuration scenario for SAP BW powered by SAP HANA

  S4H includes multi-server configurations scenario for S/4HANA (see SAP note 2408419 - SAP S/4HANA - Multi-Node Support)

> Note:
> BWoH = BW on HANA
>
> BW4H = BW for HANA
>
> DM = Data Mart
>
> SoH = Suite on HANA
>
> S4H = SAP S/4HANA

SAP Business One includes single server configurations specific for SAP Business One .

Additional CPUs for SAP Business One include the following: Haswell EP architecture: Intel E5-2670 v3 ,Broadwell EP architecture: Intel E5-2650 v4 or E5-26## v4 with higher specification; System size less than or equal to 256 GB is supported with E5-2630 v4 .

The supported entry level systems are valid for specific service packs. The hardware was tested by the hardware partner with SAP LinuxLab . The systems are supported for SAP HANA.

SAP HANA Curriculum

Additional courses, which focus on specific aspects of SAP HANA, are available. They include the following topics:

- SAP HANA: Administration and Installation
- SAP HANA: Modeling
- SAP: Development
- SAP BW to SAP BW/4HANA

> Note:
> The curricula for the different roles in SAP HANA can be found at https://training.sap.com/de/en/courses-and-curricula/hana

Other SAP HANA courses cover the following topics:

- SAP HANA: Monitoring and performance tuning (HA215)
- Authorization, Security, and Scenarios (HA240)
- Migration to SAP HANA using DMO (HA250)
- SAP HANA Data Provisioning (HA350)
- SAP HANA Smart Data Integration (HA355)
- SAP HANA Vora (HA500)
- SAP BW on SAP HANA (BW305H , BW310H , and BW362)

More information about these courses can be found at https://training.sap.com/de/en/.

LESSON SUMMARY

You should now be able to:

- Find SAP HANA Information Sources

Learning Assessment

1. What software innovations did SAP deliver with SAP HANA?

   Choose the correct answers.

   - [ ] **A** Row and Delta store.
   - [ ] **B** Compression and Decompression.
   - [ ] **C** Partitioning.
   - [ ] **D** Insert only on Column store.

2. Which of the following websites are SAP HANA documentation websites?

   Choose the correct answer.

   - [ ] **A** https://help.sap.com/hana
   - [ ] **B** https://support.sap.com/swdc
   - [ ] **C** https://www.sap.com/community/topic/hana.html

Unit 1

Learning Assessment - Answers

1. What software innovations did SAP deliver with SAP HANA?

   Choose the correct answers.

   [X] **A** Row and Delta store.

   [ ] **B** Compression and Decompression.

   [X] **C** Partitioning.

   [ ] **D** Insert only on Column store.

   Correct! SAP HANA provides a Row and Column store to store database tables. To increase row-level insert and update performance a database table has a Main store and a Delta store. SAP HANA can partition the in-memory tables to speed up query execution. SAP HANA compresses the data stored in memory and can work with the compressed data without decompressing it. For human readability SAP HANA performs a tuple reconstruction on the result set. To speed up row-level insert and update performance, SAP HANA provided insert only on the Delta store of a column table. In the Main store of a a table records will be deleted during a Delta Merge . Read more on this in the lesson SAP HANA Introduction (Unit 1, Lesson 1) of the course HA200_14.

2. Which of the following websites are SAP HANA documentation websites?

   Choose the correct answer.

   [X] **A** https://help.sap.com/hana

   [ ] **B** https://support.sap.com/swdc

   [ ] **C** https://www.sap.com/community/topic/hana.html

   Correct! At https://help.sap.com/hana you can find the latest documentation for SAP HANA. The support website is the SAP Download Center. The community website is where you can share your stories, connect with experts and keep in synch with what's happening in the SAP Hana world. Read more on this in the lesson SAP HANA Information Sources (Unit 1, Lesson 2) of the course HA200_14.

UNIT 2 Installation Preparations

UNIT OBJECTIVES

- Perform SAP HANA Sizing
- Understand the System Requirements for SAP HANA

Unit 2
Lesson 1

SAP HANA Sizing

LESSON OVERVIEW

The goal of this lesson is to understand what needs to be considered for a correct sizing of SAP HANA for an SAP HANA appliance or an SAP HANA Tailored Datacenter Integration (TDI) approach.

Business Example

Your company has decided that all SAP Business Suite and SAP BW systems will be migrated to the SAP HANA database. It is your task to investigate what is the best method to deploy the SAP HANA database, that is, to deploy as an SAP In-Memory Appliance (SAP HANA) or to deploy following the SAP HANA Tailored Datacenter Integration (TDI) approach.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Perform SAP HANA Sizing

General Concept of Sizing SAP HANA

The SAP HANA database can be deployed as an SAP In-Memory Appliance (SAP HANA) or following the SAP HANA tailored data center integration approach.

> Note:
> The following information refers solely to the sizing of the SAP HANA database server. Depending on the scenario, the sizing of other components, such as the application server, must be considered separately.

When sizing an SAP HANA system, every SAP HANA customer must perform a memory sizing. Depending on the SAP HANA deployment, the sizing approach differs as follows:

- For new SAP HANA implementations, it is necessary to size the memory for an SAP HANA system using the SAP Quick Sizer in Related Information.
- For systems that are migrating to SAP HANA , we recommend the following:
  - If the migration is from an SAP NetWeaver-based system, use the sizing report on the source database.
  - If the migration is from a non-SAP NetWeaver data source, use the sizing as in SAP Note 1514966 .
- Any system that is large or complex requires sizing from an SAP sizing expert.

Sizing of the SAP HANA database is based mainly on the required main memory size. The memory size is determined by the amount of data that is to be stored in memory.

SAP HANA compresses the data.. Because the compression factor depends on the used scenario, you cannot estimate the amount of memory needed. The memory sizing must always be performed using the Quick Sizer for SAP HANA , or the SAP HANA Sizing Reports and SAP Notes .

Hint:
If you are interested in how SAP HANA compression works, see SAP Note 2112604 : FAQ - SAP HANA Compression .

SAP HANA Sizing Scenarios

Depending on the used scenario, the following table gives the recommended sizing approach.

Table 2: SAP HANA Sizing Scenarios

Scenario	Sizing SAP BW/4HANA	Sizing Suite on SAP HANA	Sizing Non-NetWeaver
New Implementation	Quick Sizer for SAP BW/4HANA	Quick Sizer for Suite on SAP HANA	Sizing guide as available
Database migration	SAP Note 1637145 : Sizing for SAP BW on HANA SAP Note 1736976 : Sizing Report for SAP BW on HANA SAP Note 2121330 - FAQ: SAP BW on HANA Sizing Report SAP Note 2296290 - New Sizing Report for BW/4HANA	SAP Note 1514966 - SAP HANA 1.0 : Sizing SAP HANA SAP Note 1793345 : Sizing for SAP Suite on SAP HANA SAP Note 1872170 : Suite on SAP HANA memory sizing report SAP Note 2367595 - Suite on HANA memory Sizing	SAP Note 1514966 : SAP HANA: Sizing SAP In-Memory Database

Hardware recommendation

The result of the memory sizing is the basis for the hardware recommendation for an SAP HANA system. If you decide to buy the in-memory appliance (SAP HANA), you have a selection of certified appliances from certified hardware partners. Check the SAP certified and supported SAP HANA hardware or hardware that matches your memory sizing results. Ivy Bridge customers should check the SAP Community Network (SCN). For an in-memory appliance, you don’t need to consider storage and CPU sizing, because they are included in the certified appliance offering.

Caution:
Applications other than the SAP HANA database software must not be installed on the database server, except for scenarios that are explicitly supported by SAP. This is discussed in the lesson, Describing Deployment Options .

Sizing Approach for the SAP In-Memory Appliance

If you decide to buy the In-Memory Appliance, you have a selection of appliances from hardware partners available to you. Check the SAP HANA Hardware Directory for hardware that matches your memory sizing results.

> Note:
> For an In-Memory Appliance, you do not need to consider storage and CPU sizing, because they are included in the appliance offering.

To calculate the memory requirements, use the following SAP notes:

- SAP Quick Sizer: https://service.sap.com/quicksizer
- SAP Note 1514966 : SAP HANA- Sizing SAP In-Memory Database
- SAP Note 1637145 : Sizing for SAP BW/4HANA
- SAP Note 1736976 : SAP BW/4HANA sizing report
- SAP Note 1793345 : Sizing for SAP Suite on SAP HANA
- SAP Note 1872170 : Suite on HANA memory sizing report

SAP HANA appliances

An overview of the available and certified SAP HANA appliances can be found on the SAP HANA Hardware Directory . The list is displayed on the URL https://global.sap.com/community/ebook/2014-09-02-hana-hardware/enen/index.html .

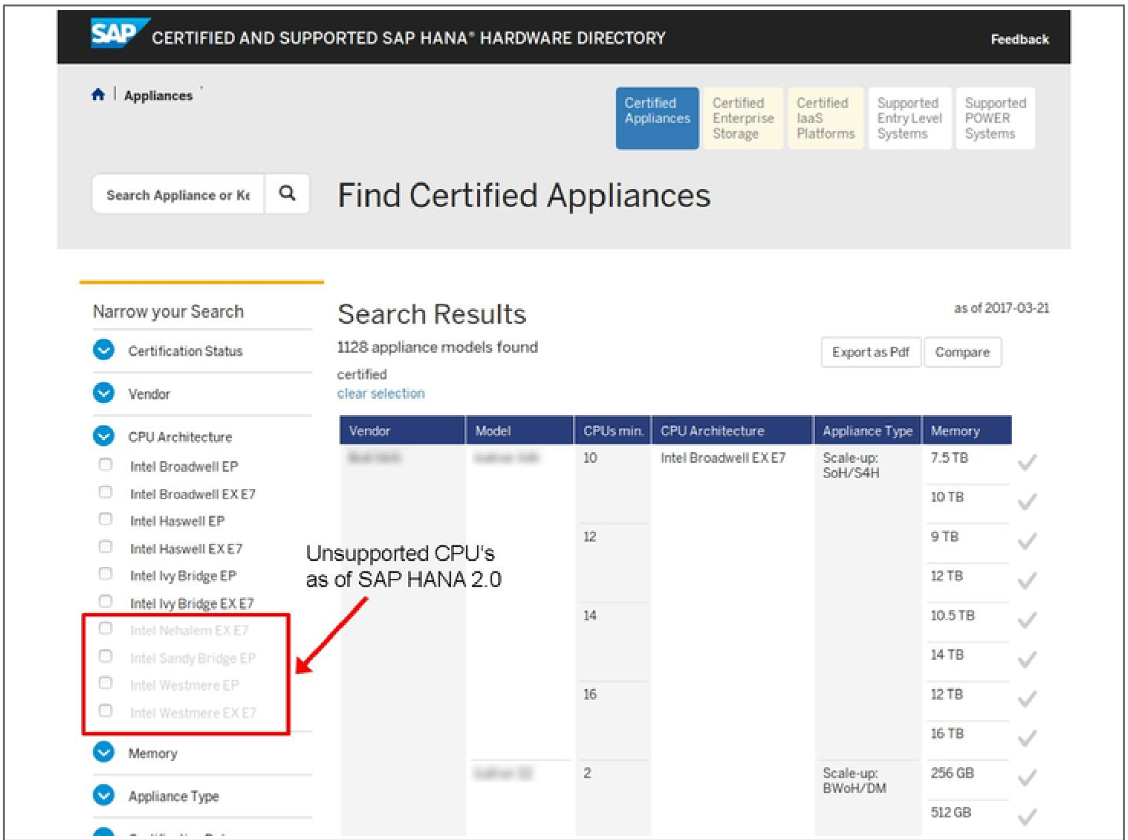

Figure 13: The SAP HANA Hardware Directory

> Caution:
> It is mandatory to begin with SAP HANA 2.0 computing nodes with at least Intel Haswell CPU or later. See SAP Note 2399995 : Hardware requirement for SAP HANA 2.0.

Sizing recommendations apply for certified hardware only. Contact your hardware vendor for suitable hardware configuration. The SAP HANA development team is constantly optimizing the SAP HANA database. We recommend that you always check the latest documentation and SAP Notes when preforming an SAP HANA memory sizing.

Sizing Approach for the SAP HANA Tailored Datacenter Integration (TDI) System

Next to the black-box appliance approach, a customer can increase flexibility according to the IT landscape, and optimize for special requirements by choosing from the following options:

- To use available hardware
- To reuse hardware and save costs
-

If you decide to build the SAP HANA system based on the SAP HANA Tailored Data Center Iintegration approach, you must become certified in tailored data center integration. For storage sizing recommendations, see the SAP HANA Storage Requirements whitepaper on the SAP Documents website (https://www.sap.com/docs/download/2015/03/74cdb554-5a7c-0010-82c7-eda71af511fa.pdf).

> Note:
> IBM provides a process to support mapping of the SAP sizing to a hardware or partition configuration that meets the sizing needs of the customer. For more information, see SAP Note 2055470 : SAP HANA on POWER Planning and Installation Specifics - Central Note .

Sizing an SAP HANA Tailored Data Center Integration Setup

Sizing an SAP HANA Tailored Data Center Integration setup consists of three main steps, as follows:

- RAM sizing for static and dynamic data
- Disk sizing for the persistence storage
- CPU sizing for the queries and calculations

The three sizing steps are explained in the following sections.

Main Memory Sizing

The customers first step to sizing an SAP HANA Tailored Data Center Integration system is to perform a memory sizing. Depending on the use case, shown in the figure, Example Sizing Report Result for Suite on SAP HANA Memory , you would use the following:

- SAP Quick Sizer: http://service.sap.com/quicksizer
- SAP Note 1514966 : SAP HANA: Sizing SAP In-Memory Database

- SAP Note 1637145 : Sizing for SAP BW/4HANA
- SAP Note 1736976 : SAP BW/4HANA Sizing Report
- SAP Note 1793345 : Sizing for SAP Suite on HANA
- SAP Note 1872170 : Suite on SAP HANA Memory Sizing Report
- SAP Note 2055470 : SAP HANA on POWER Planning and Installation Specifics - Central Note

Using the SAP Quick Sizer or SAP Notes helps you to determine the SAP HANA memory size.

Sizing for Suite on SAP HANA Using SAP Notes 1793345 and 1872170

These two SAP Notes describe the approach of sizing a Business Suite system on SAP HANA and SAP S/4HANA . There is also a sizing script attached to these notes.

```
RESULTS OF SUITE ON HANA SIZING IN GB
Based on the selected table(s), the anticipated maximum memory requirement is

for Suite on HANA:
- Total memory requirement                                          2.412,3
- Net data size on disk                                             1.333,8

for Suite on HANA with Simple Finance 2.0:
- Total memory requirement                                          2.356,2
- Net data size on disk                                             1.354,0

for optimal set-up of Suite on HANA:
- Total memory requirement                                          1.975,8
- Net data size on disk                                             1.182,8
```

Figure 14: Example Sizing Report Result for Suite on SAP HANA Memory

Sizing for SAP BW/4HANA Using SAP Notes 1637145 and 1736976

The sizing notes describe the memory sizing (column, row store, and additional components), the disk sizing for data and log files, and the CPU sizing. There is also a sizing script attached to these notes.

```
MINIMUM SIZING RECOMMENDATION - CURRENT
=======================================

                                                    minimum      recomm.
                         Phys. memory per node:     512 GB       2048 GB
------------------------------------------------------------------------------
Memory Requirement (Minimum Total):                 1099 GB   |   999 GB
Disk Space Requirement - data (Minimum Total):       992 GB   |   992 GB
Disk Space Requirement - logs (Minimum Total):       512 GB   |   512 GB
Number of Nodes incl. master  (Minimum Total):         3      |     1
------------------------------------------------------------------------------
NOTE:
- Please carefully read documentation attached to SAP NOTE 1736976
  for a detailed description of the sizing procedure and its results!
- Disk space requirement calculation no longer includes space for backups
  dumps, etc. This space has to be provided on additional disk volumes. The
  guidance for sizing these additional requirements is described within the
  documentation attached to SAP Note 1736976.
------------------------------------------------------------------------------


SIZING DETAILS
==============

 (For  2048  GB nodes)  data [GB]     total [GB]     total [GB]     util.
                                      incl. tmp.     (non-act.)
  Row Store                    23             40             40
  Master Column Store          34             68             68
  Worker Column Store         439            853            841
  Caches / Services            50             50             50
 ---------------------------------------------------------------
  TOTAL (All Servers)         546           1011            999     50 %
```

Figure 15: Example Sizing Report Result for SAP BW/4HANA

Note:
The ZNEWHDB_SIZE report runs with a low system load, depending on the size of your suite on the SAP HANA system, it takes up to 8 to 12 hours or more. Therefore, we recommend that you test the report in your consolidation system before loading.

Sizing SAP HANA for the Non-NetWeaver Approach Using SAP Note 1514966

This SAP Note describes the sizing of SAP HANA in a non-NetWeaver scenario, for example when the data is coming from an external data source and SAP HANA is used to model and analyze that data. Do not use these sizing rules for sizing SAP BW/4HANA, Business Suite on SAP HANA systems, or S/4HANA.

General Sizing Approach :

Determine static + dynamic RAM requirement :

- Calculate "uncompressed" data volume to be loaded into HANA (tools provided)
- Apply compression factor (from attachment to sizing notes)
- Multiply result by 2 because we recommend

RAM dynamic = RAM static

Figure 16: General SAP HANA Main Memory Sizing Approach

Static vs. Dynamic Memory Requirements

Distinguish between the static and the dynamic RAM requirement as follows:

Static data memory requirements

The static RAM requirements refer to the amount of main memory that is used for holding the table data. Static memory sizing of SAP HANA is determined by the amount of data that is to be stored in memory. Specifically, this figure relates to the amount of disk space covered by the corresponding database tables, excluding their associated indexes.

Note that, if the database supports compression, the space of the uncompressed data is needed. Based on this amount of data, a compression factor is applied to determine the size of the RAM needed for SAP HANA.

Dynamic data memory requirements

Additional memory is required for objects that are created dynamically when new data is loaded or queries are executed. Because we recommend reserving as much memory for dynamic objects as for static ones, calculate the total RAM by multiplying the static RAM by two.

Note:
The following SAP Notes provide information on SAP HANA support for virtualized environments and the Business Suite on SAP HANA:

- SAP Note 1788665 : SAP HANA Support for virtualized or partitioned (multitenant) environments
- SAP Note 1995460 : SAP HANA supported production scenarios on VMware
- SAP Notes 1781986 , 1825774 , and 1950470 : Support for Business Suite on SAP HANA

Disk Sizing

SAP HANA is an in-memory database, which stores and processes the bulk of its data in-memory. Additionally, it provides protection against data loss by saving the data in persistent storage locations.

Persistent storage distinguishes between the data volume and the log volume. In the data volume of SAP HANA, a copy of the in-memory data persists by writing changed data to the data volume. The log volume ensures the recovery of the database with zero data loss in case of faults. SAP HANA records each transaction in the form of a redo log entry.

Disk Space Required for the Data Volume

Whenever you create a Savepoint or a Snapshot, or perform a delta merge, data persists from memory to the data volume under `/hana/data/<SID>`. The recommended size of the data volume for a given SAP HANA system is equal to the calculated results from the sizing reports. Use the value net data size on the disk plus an additional free space of 20%. The figure, Determining the SAP HANA Data Volume Size, shows an example sizing report result for Suite on HANA. The sizing report shows the Net data size on disk. To determine the required SAP HANA data volume size, add 20%.

```
RESULTS OF SUITE ON HANA SIZING IN GB
Based on the selected table(s), the anticipated maximum memory requirement is

for Suite on HANA:
- Total memory requirement                                    2.412,3
- Net data size on disk                                       1.333,8
```

$$Size_{data} = 1.2x \text{ Net data size on disk}$$

Figure 17: Determining the SAP HANA Data Volume Size

During the migration of a non-SAP HANA database to SAP HANA, the system may temporarily need more disk space for data than calculated in the sizing phase. With Enterprise Storage, this is not considered relevant for the overall storage sizing, because the storage system can provide that additional space, if required.

Disk Space Required for the Log Volume

The minimum size of the log volume depends on the number of data changes occurring between two SAP HANA Savepoints which, by default, are created every 5 minutes. The more data changes that are executed by write transactions in that period of time, the more redo log segments that are written to the log volume under `/hana/log/<SID>`. When sizing the log volume, consider the following points:

- The redo log must not be overwritten before a Savepoint entry is available in the data volume, otherwise, the SAP HANA database may be unable to restart.

  Situations may occur where the writing of a Savepoint is delayed. For example, delays occur if a high workload must be processed during a database migration process in an environment with slow input/output between the source and target (SAP HANA) database. In such cases, as long as the Savepoint has not been written to the data volume, the amount of redo logs in the log volume continue to grow until all log segments are full.

- If `LOG_MODE = NORMAL` is set, the redo log must not be overwritten before a backup takes place. Therefore, keep some extra space available for situations where incidents or faults may interrupt the backup process. That extra space allows system administrators to fix and finish the backup process before the log volume runs full.

Determine Log Volume Size

There is no direct correlation between the SAP HANA database size and the required log volume size. Nevertheless, we recommend using the formula in the figure, Determine Log Volume Size , because it is based on best practice and experiences with productive SAP HANA installations. Unlike the formula for the data volume, it is calculated depending on the total memory requirement ("RAM").

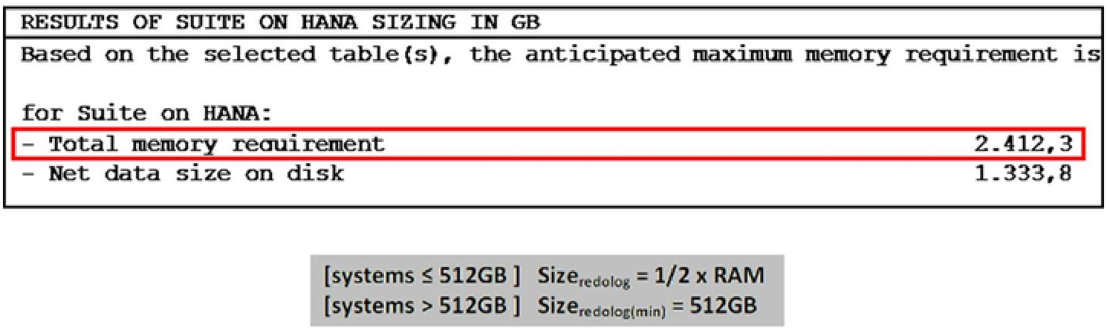

Figure 18: Determining Log Volume Size

Examples: Log Volume Size

Examples of log volume size are as follows:

- 128 GB system ≥Size log volume = 64 GB
- 256 GB system ≥Size log volume = 128 GB
- 512 GB system ≥Size log volume = 256 GB
- 1 TB system ≥Size log volume $_{min}$ = 512 GB
- 2 TB system ≥Size log volume $_{min}$ = 512 GB
- 4 TB system ≥Size log volume $_{min}$ = 512 GB

> Note:
> For systems with more than 512 GB of in-memory database size, the previous formula represents a minimum value. As of today, based on the experience made with productive SAP-internal SAP HANA installations, this value is considered sufficient for each SAP HANA use case. Nevertheless, as described previously , as the amount of data stored in the log volume depends on the workload processed, there may be situations where this value is not sufficient for log volume sizing.

Disk Space Required for SAP HANA Installation

All binary, trace, and configuration files are stored on a shared file system that is exposed to all hosts of a system under `/hana/shared/<SID>`. Therefore, additional space is required for the traces written by the compute node or nodes of the SAP HANA database. Experience with productive SAP HANA installations shows that the bigger the size of the SAP HANA database, the more traces that are written. Therefore, the calculation is based on the total memory requirement (RAM).

For single-node SAP HANA systems, the recommended disk space for `/hana/shared/<SID>` is shown in the figure, Determine Share Single-Host Size .

$$Size_{installation(single\text{-}node)} = MIN(1 \text{ x } RAM; 1 \text{ TB})$$

Figure 19: Determine Share Single-Host Size

Examples: Single Size node

The following are examples of single-node SAP HANA installation sizes:

- Single-node 128 GB $\geq Size_{installation}$ = 128 GB
- Single-node 256 GB $\geq Size_{installation}$ = 256 GB
- Single-node 512 GB $\geq Size_{installation}$ = 512 GB
- Single-node 1 TB $\geq Size_{installation}$ = 1 TB
- Single-node 2 TB $\geq Size_{installation}$ = 1 TB
- Single-node 4 TB $\geq Size_{installation}$ = 1 TB

Determine Share Scale-Out Size

For scale-out SAP HANA systems, the recommended disk space for `/hana/shared/<SID>` depends on the number of worker nodes. A disk space of 1x RAM per worker is recommended for every four worker nodes of a given scale-out system.

$$Size_{installation(scale\text{-}out)} = 1 \text{ x } RAM\_of\_worker \text{ per 4 worker nodes}$$

Figure 20: Determine Share Scale-Out Size

Examples: Scale out HANA installation

The following are examples of scale-out SAP HANA installation sizes:

- 3+1 system, 512 GB per node $\geq Size_{installation}$ = 1x 512 GB = 512 GB
- 4+1 system, 512 GB per node $\geq Size_{installation}$ = 1x 512 GB = 512 GB
- 5+1 system, 512 GB per node $\geq Size_{installation}$ = 2x 512 GB = 1 TB
- ...
- 9+1 system, 512 GB per node $\geq Size_{installation}$ = 3x 512 GB = 1.5 TB
- ...
- 3+1 system, 1 TB per node $\geq Size_{installation}$ = 1x 1 TB = 1 TB

- 4+1 system, 1 TB per node $\geq Size_{installation} = 1x\ 1\ TB = 1\ TB$
- 5+1 system, 1 TB per node $\geq Size_{installation} = 2x\ 1\ TB = 2\ TB$
- 9+1 system, 1 TB per node $\geq Size_{installation} = 3x\ 1\ TB = 3\ TB$

Disk Space Required for Backups

A complete data backup contains the entire payload of all data volumes. The size required by the backup directory not only depends on the total size of the data volumes, but also on the number of backup generations kept on disk, and on the frequency with which data is changed in the SAP HANA database. For example, if the backup policy requires you to perform complete data backups on a daily basis and to keep those backups for one week, the size of the backup storage must be seven times the size of the data area.

In addition to data backups, backup storage for log backups must be reserved to provide the possibility for a point-in-time database recovery. The number and size of log backups to be written depend on the number of change operations in the SAP HANA database.

$$Size_{backups} \geq Size_{data} + Size_{redolog}$$

Figure 21: Determine Backup Size

Technically, it is possible to store the backups of several SAP HANA databases in a central shared backup storage. But if several backup or recovery processes run in parallel, this impacts the overall data throughput of the given backup storage. That is, if the backup storage cannot guarantee a constant level of data throughput once the number of parallel processes exceeds a certain number, backup and recovery processes can slow down significantly.

Disk Space Required for Exports

Sometimes the database content is needed for a root cause analysis of problems. For this purpose, sufficient disk space must be provided to hold the binary exports. In most cases it is not necessary to export the entire database content for root cause analysis. Therefore it is sufficient to reserve storage space of about two times the size of the largest database table.

DB CPU Sizing

The CPU requirements for migrating to SAP HANA standalone are difficult to anticipate, as there is no real reference against which to compare. Therefore, the sizing referred to previously has the following formula: 300 SAPS per active user / 0.65 for a CPU utilization buffer. An active user is one that consumes CPU power at a given point in time. In sizing, customers often overestimate the (overlapping) activity patterns of their end users. Some end users also may perform more or less intensive calculations on the database level.

Sizing approach similar to user based CPU sizing of BW and BWA

Maximize query throughout by multiuser scenarios with queries of different complexity out of delivered content, 10 - 20 million records

Assumptions :

- Three different query complexity classes
- Three different user profiles (click rate, query complexity)

Normalization to query throughout per core resp. active per core :

CPU : 300 SAPS / active user

Note that the CPU sizing has to be adjusted so that the server load does not exceed 65 % in average.

Figure 22: CPU Sizing Approach

Consider this recommendation as an initial estimate that needs verification. The more users there are on the system, the less likely it is that this formula will be accurate. The decision of whether you invest time into further CPU analysis depends upon the risk of reaching CPU limits. SAP HANA servers with two sockets, for example, deliver round about 60,000 SAPS.

If you want to verify the CPU requirements, a test with the top 5 to 10 SAP HANA transactions can be helpful, either within a single user test or a load test.

Sizing SAP HANA Using the Quick Sizer

SAP HANA Database can also be sized using SAP Quick Sizer . For more information, see http://service.sap.com/quicksizer .

The Quick Sizer calculates the following:

- CPU
- Disk
- Memory
- Input/output resource categories

It calculates these based on throughput numbers and the number of users working with the different SAP solutions in a hardware and database-independent format. Sizing is an iterative process that continuously brings together customers, hardware vendors, and SAP. So, for example, direct links to SAP hardware vendors facilitate the tendering procedure.

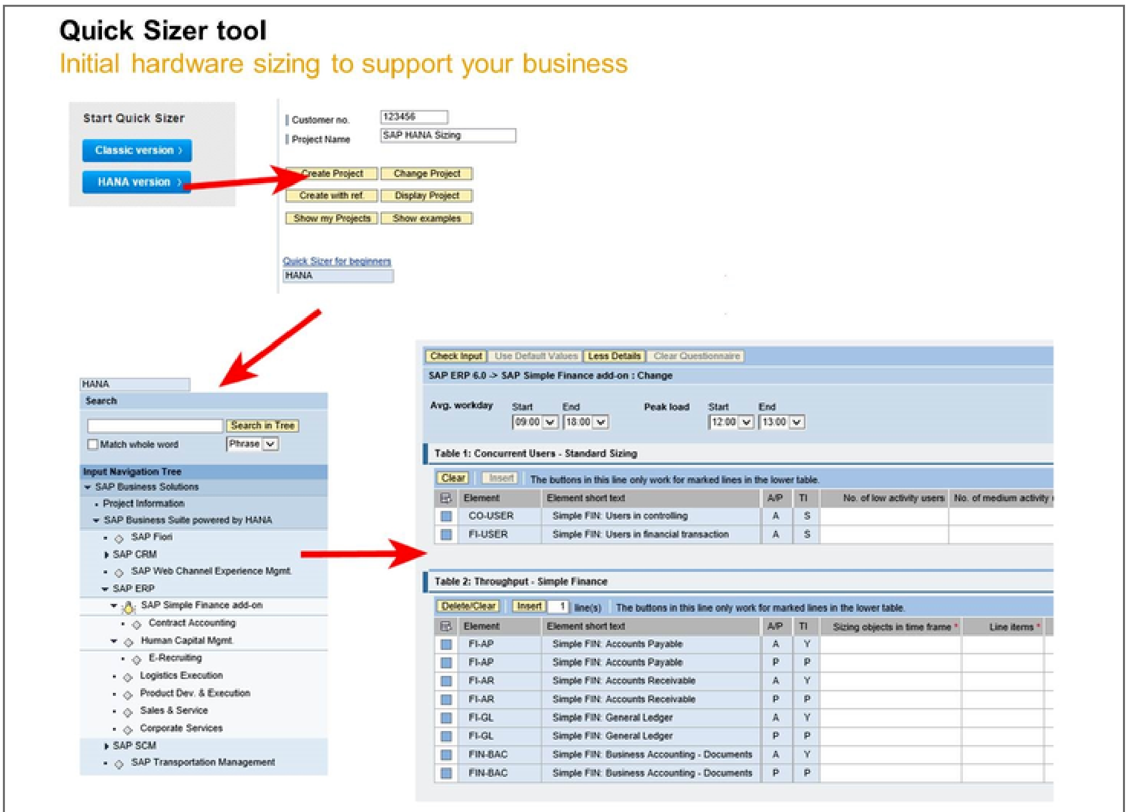

Figure 23: SAP Quick Sizer Tool

For an initial sizing recommendation using the SAP Quick Sizer, follow the steps shown in the figure, AP Quick Sizer Tool. Sample configurations, which can be checked at http://www.sap.com/benchmark.

In SAP Quick Sizer, multiple predefined scenarios can be selected. For example, the following scenarios can be selected:

- SAP Business Suite powered by SAP HANA
- SAP BW/4HANA
- Standalone SAP HANA

For each of the scenarios, the expected compression of the data is different.

Additional Remarks

For various SAP HANA scenarios, native and third party technologies provide features to displace data not frequently used either for the SAP HANA persistence or for other database management systems. If such a technology is used, this is considered in the main memory sizing. The following are examples:

- Nonactive data concept for SAP BW/4HANA (SAP Note 1767880) and Nearline Storage Solutions

  Large SAP BW systems contain large amounts of data that are no longer, or rarely, used. However, they remain in the system, for example, historical data, keeping data for legal reasons, and so on. This data is called nonactive data. An implementation for SAP BW/4HANA allows for the displacement of non-active data if the main memory bottlenecks use a last-recently-used concept. This concept improves main memory resource management, which has positive effects on hardware sizing for a large amount of nonactive data. For more information, see SAP Note 1736976 . In addition, nearline storage solutions could be used to store cold data, which can also help to reduce the memory amount.

- SAP HANA Smart Data Access (SAP Note 1879294)

  SAP HANA smart data access enables access to remote data as if it was stored in local tables. Since the data is not copied to SAP HANA, you no longer need to consider it for the main memory sizing of the SAP HANA server.

- SAP HANA Dynamic Tiering (SAP Note 2225582 : SAP HANA Dynamic Tiering SPS 11 Release Note

  SAP HANA dynamic tiering is a native big data solution for SAP HANA. Dynamic tiering adds smart, disk-based extended storage to your SAP HANAdatabase. Dynamic tiering enhances SAP HANA with large volume, warm data management capability. By using dynamic tiering to place hot data in SAP HANA in-memory tables, and warm data in extended tables, the highest value data remains in-memory, and cooler less-valuable data is saved to the extended store. This can reduce the size of your in-memory database.

LESSON SUMMARY
You should now be able to:

- Perform SAP HANA Sizing

Unit 2
Lesson 2

Linux Operating System Requirements

LESSON OVERVIEW

This lesson describes the Linux operating system requirements that have to be fulfilled before you can start the installation of an SAP HANA system. This lesson does not replace the "SAP HANA Server Installation and Update Guide" and the SAP HANA installation SAP Notes.

Business Example

You need to set up the Linux operating system so that all of the SAP HANA requirements are fulfilled and you can start the SAP HANA installation.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Understand the System Requirements for SAP HANA

System Requirements for SAP HANA

Before we discuss the SAP HANA requirements for the Linux operating system, this lesson describes some terms that are often used in the SAP HANA documentation. SAP HANA systems are available in the following two types, known as system types:

- A single-host system is the simplest system installation type. The SAP HANA system runs entirely on one host and the server needs to handle the full query load.
- A multi-host system is a system with more than one host, which can be configured as active worker hosts or idle standby hosts. This means that load can be balanced between different hosts.

Single-Host System

The figure, Single-Host System, shows the file system layout for a single-host system.

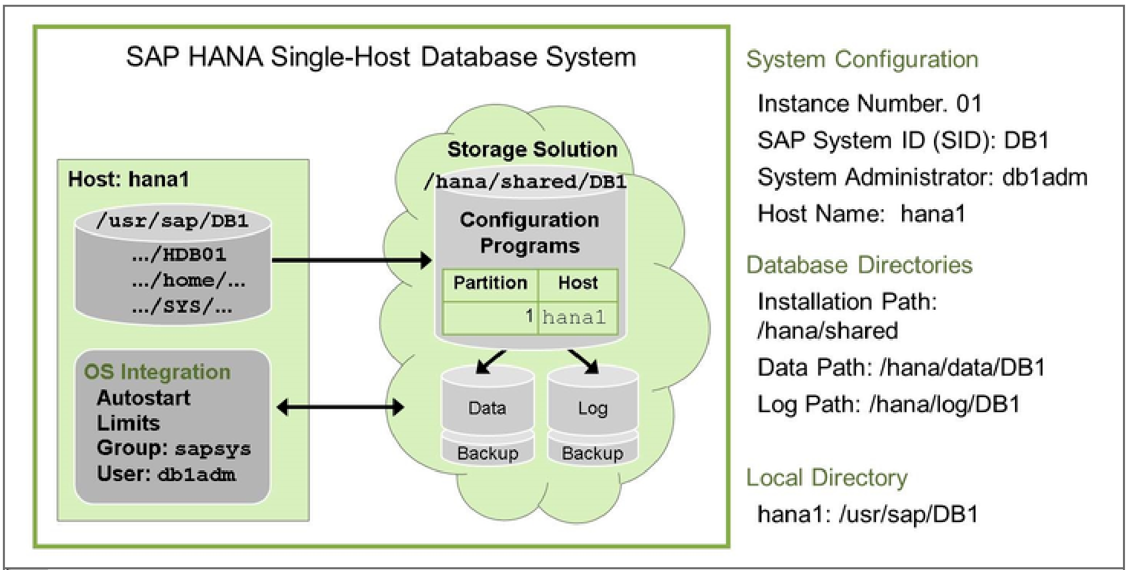

Figure 24: Single-Host System

Multi-Host System

The figure, Multi-Host System , shows the file system layout for a multi-host system.

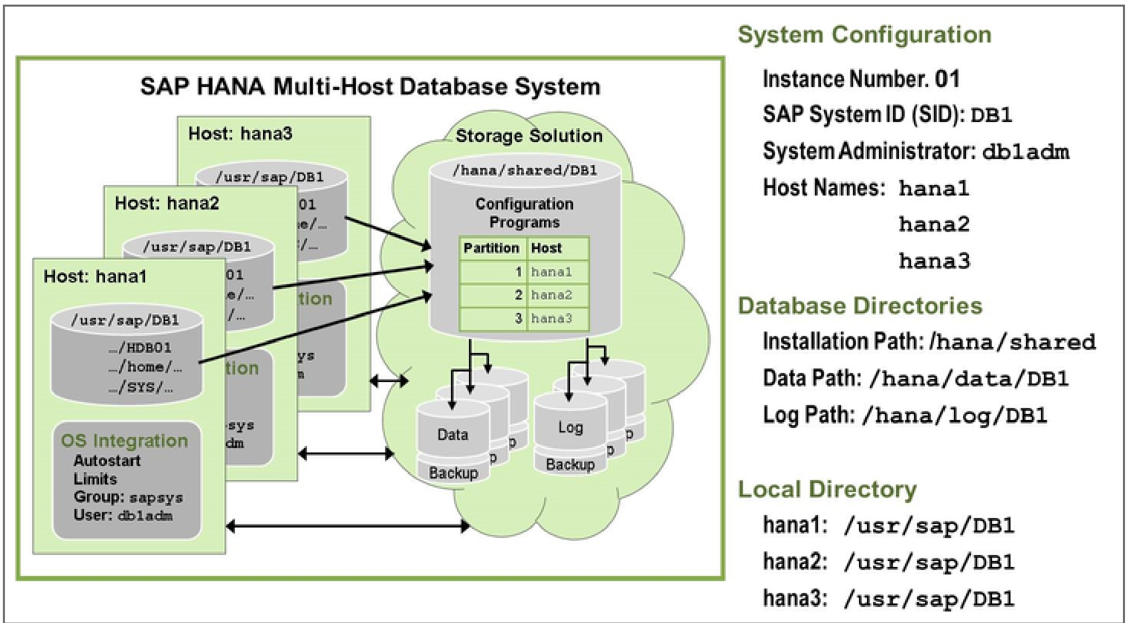

Figure 25: Multi-Host System

Note:
SAP HANA certified hardware partners or owners of a **C_HANATEC_11** or newer, can install an SAP HANA system. In both cases, the hardware running SAP HANA must be certified by SAP.

Basic Components of an SAP HANA System

An SAP HANA system is composed of three main components, as follows:

- The host
- The SAP HANA system
- The SAP HANA instance

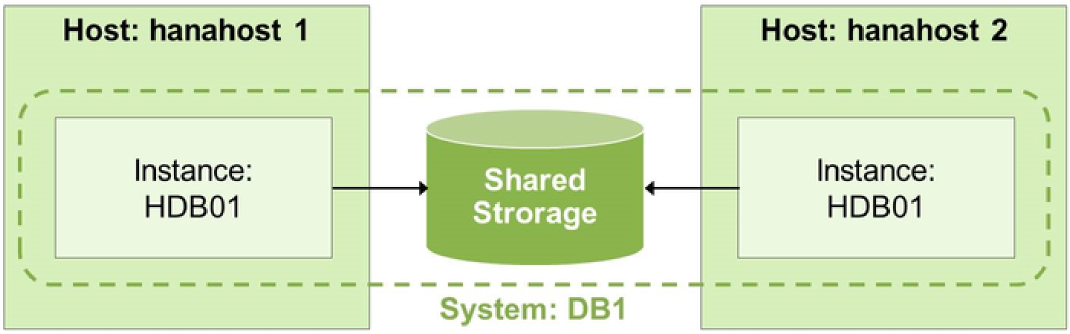

Figure 26: Basic Components of an SAP HANA System

Terms in SAP HANA System

It is important to understand the following terms as they apply to the SAP HANA system:

Host

A host is the hardware and operating environment in which the SAP HANA database runs. SAP HANA is supported on SUSE Linux Enterprise Server and Red Hat Enterprise Server. The host provides all the resources and services (CPU, memory, network, and storage) that the SAP HANA database requires. The storage for an installation does not have to be on the host; it can be shared storage as well. Multi-host SAP HANA systems require shared storage or storage that is accessible on-demand from all hosts.

System

A system is one or more instances with the same SAP system ID and instance number. The term system is interchangeable with the term SAP HANA database. If an SAP HANA system has more than one instance, it is distributed over several hosts. The SAP system ID (SAPSID or SID) is the identifier for the SAP HANA system.

Instance

An SAP HANA instance is the set of SAP HANA system components that are installed on one host. A system can be distributed as several instances among several hosts, but each instance in a multi-host system must have the same instance number.

Standalone SAP HANA System with Single-SID and Multiple-SID Installations

The figure, Standalone SAP HANA System with Single-SID and Multiple-SID Installations, shows two possible single-host SAP HANA system configurations.

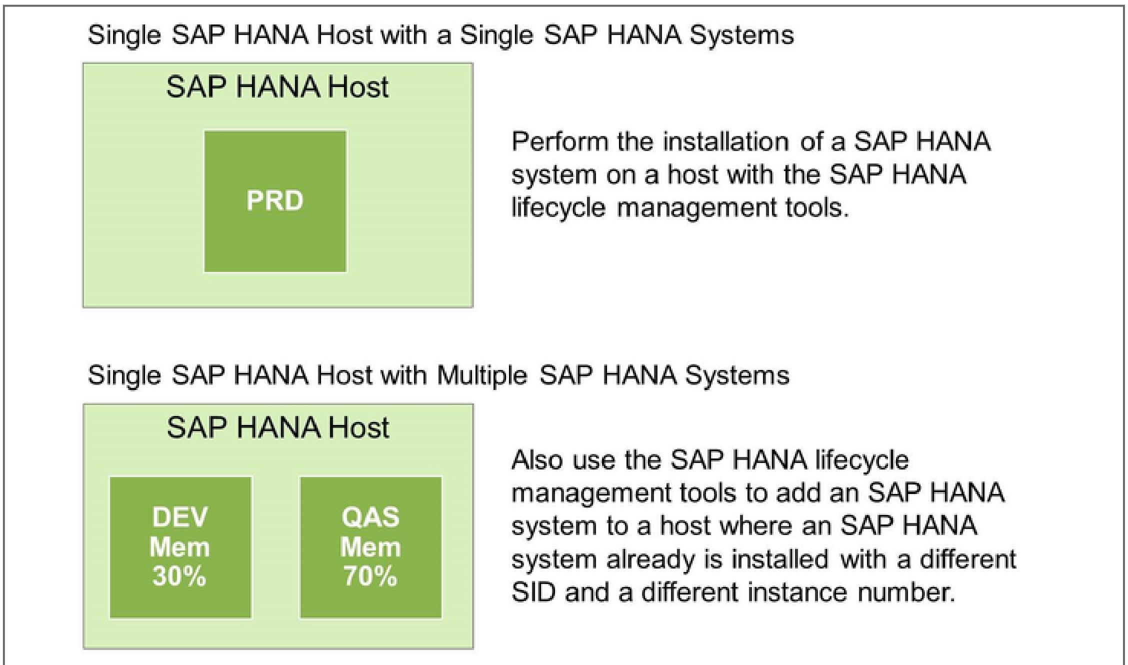

Figure 27: Standalone SAP HANA System with Single-SID and Multiple-SID Installations

Supported Operating Systems for SAP HANA 2.0 on Intel-Based Hardware Platforms

SAP HANA is available on the SUSE Linux and Red Hat Linux. Check that you are using the SAP-supported version.

Before the SAP HANA installation can be started, configure the Linux system according to the recommended operating system settings for SUSE Linux Enterprise Server (SLES) and Red Hat Enterprise Linux (RHEL).

> Note:
> For an overview of all of the supported Linux versions, see SAP Note 2235581 : SAP HANA: Supported Operating Systems .

For an SAP HANA system on Intel-based hardware platforms, the following operating systems are available for SAP HANA 2.0 :

- SUSE Linux Enterprise Server (SLES) for SAP Applications 12 SP1
  - SAP Note 2205917 : SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12
  - SAP Note 1984787 : SUSE LINUX Enterprise Server 12: Installation notes
- Red Hat Enterprise Linux for SAP HANA (RHEL for SAP HANA) 7.2
  - SAP Note 2292690 : SAP HANA DB: Recommended OS settings for RHEL 7.2

Operating System for SAP HANA 2.0 on IBM Power Servers

For an SAP HANA system on IBM Power servers, the following operating system is available for SAP HANA 2.0 :

- SUSE Linux Enterprise Server (SLES) for SAP Applications 12 SP1

- SAP Note 2205917 : SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12
- SAP Note 1984787 : SUSE LINUX Enterprise Server 12: Installation notes
- SAP Note 2055470 : HANA on POWER Planning and Installation Specifics - Central Note

For migration information for IBM Power Systems (Big-Endian to Little-Endian), see the document `SAP_HANA_System_Migration_en.pdf` attached to the SAP Note 2380257 : SAP HANA Platform 2.0 SPS 00 Release Note .

> Note:
> See SAP Note 2188482 : SAP HANA on IBM Power Systems: Allowed Hardware .

Hardware Requirements

For a new installation, you must have at least 20 GB RAM in total for the software, 15 GB for the basic software plus 5 GB for programs, as well as some space for trace files. The additional memory required for data and log volumes varies according to your requirements. You also need this space for an update, because the old software version is not deleted.

> Note:
> During an installation or update of the SAP HANA database, a hardware check is performed to ensure that the hardware in use is supported.
>
> The hardware check is a script that is automatically called by the SAP HANA installation tool. It aborts the installation process if any unsupported hardware is detected. The certified SAP HANA configurations have been designed and tested together with our hardware partners to ensure that the SAP HANA database runs optimally on the used hardware. SAP HANA performance and stability cannot be guaranteed when using unsupported hardware.

Hardware Requirements for SAP HANA Network Connection

For efficient data replication, use a dedicated server network communication of 10 GBit/s between the SAP HANA landscape and the source system.

Disk Space Requirements Single-Host system

The figure, Disk Space Requirements Single-Host System , shows a list of the important file systems to include on an SAP HANAhost.

File System **Default Path**	Description	Minimum space required
Root **/**	The root partition	**10 GB**
Installation path (sapmnt) **/hana/shared**	All binary, trace and configuration files are stored on a shared file system that is exposed to all hosts of a SAP HANA system under /hana/shared/<sid>.	**Min. 1 * RAM** **Max. 1 TB**
System instance **/usr/sap**	The file system for local SAP system instance directories	At least **50 GB**
Data Volume **/hana/data/<SID>**	The default path to the data directory depends on the system ID of the SAP HANA host.	At least **3 * RAM (Appliance)** **1.2 * RAM (TDI)**
Log volume **/hana/log/<SID>**	The default path to the log directory depends on the system ID of the SAP HANA host	**Min. 0.5 * RAM** **Max. 512 GB**

Figure 28: Disk Space Requirements Single-Host System

Note:
For patching, you must have approximately 3GB in your working directory.

File System Structure for a Multi-SID TDI Installation

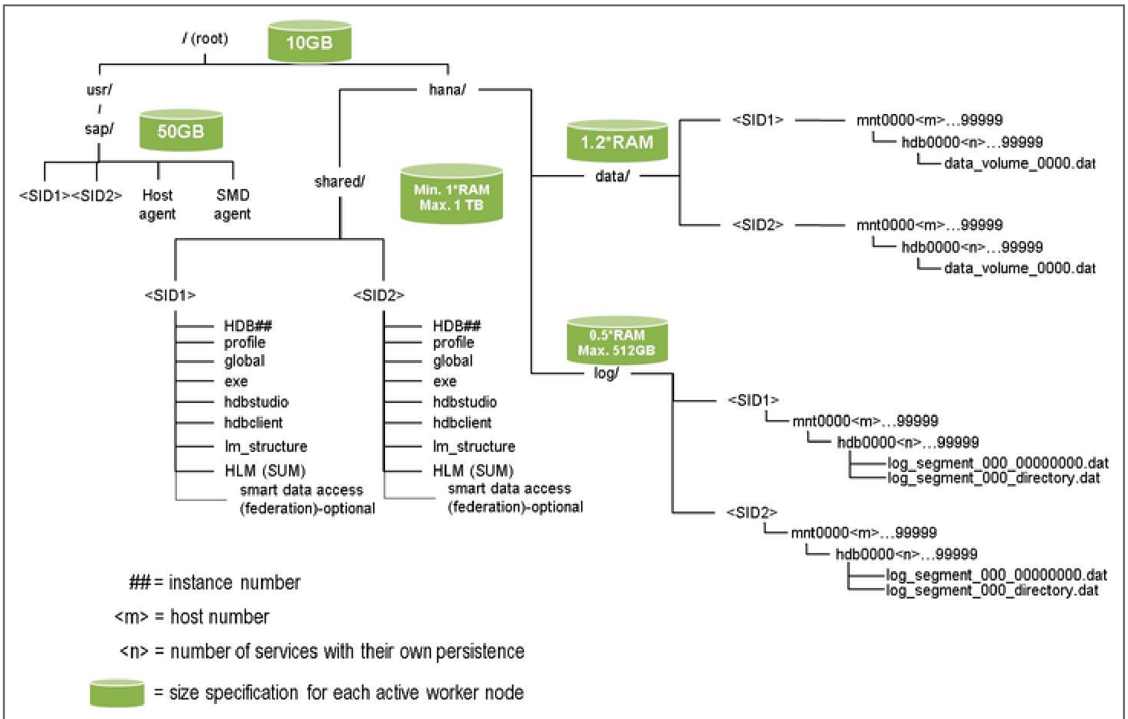

Figure 29: File System Structure for a Multi-SID TDI Installation

Note:
The hdblcm(gui) tools supports the installation of multi-SID scenarios.

An SAP HANA system in a production environment must not share any infrastructure with another SAP HANA system. Hosts running more than one SAP HANA system (sometimes referred to as multi-SID installations) should only be used for non-production purposes such as development, quality assurance, or testing.

SAP supports running multiple SAP HANA systems on a single host in production. This is restricted to single host and scale-up scenarios only. Note that multi-SID requires significant attention to various detailed tasks related to the system administration and performance management. For more information, see SAP Note 1681092 : Multiple SAP HANA DBMSs (SIDs) on one SAP HANA system .

For production systems with high availability, it is possible to share some temporarily unused resources from the standby hosts. As soon as the standby resources are needed, they become exclusively available for the production system, and are no longer shared. For more details, see the high availability information in the SAP HANA Administration Guide .

Caution:
SAP strongly recommends that you keep the data volumes on different disks.

LESSON SUMMARY

You should now be able to:

- Understand the System Requirements for SAP HANA

Learning Assessment

1. Which components for an In-Memory Appliance do you need to consider for sizing?

   Choose the correct answer.

   - [ ] **A** Storage
   - [ ] **B** Memory
   - [ ] **C** CPU

2. What is the minimum space requirement for the Data Volume?

   Choose the correct answers.

   - [ ] **A** 1.2 * RAM (TDI)
   - [ ] **B** Min. 1 * RAM
   - [ ] **C** Min. 0.5 * RAM
   - [ ] **D** 3 * RAM (Appliance)

Unit 2

Learning Assessment - Answers

1. Which components for an In-Memory Appliance do you need to consider for sizing?

   Choose the correct answer.

   [ ] **A** Storage

   [X] **B** Memory

   [ ] **C** CPU

   Correct! You have a selection of certified appliances from certified hardware partners. Check the SAP HANA Hardware Directory for hardware that matches your memory sizing results. You do not need to consider storage or CPU. They are included in the certified appliance offering. Read more on this in the lesson SAP HANA Sizing (Unit 2, Lesson 1) of the course HA200_14.

2. What is the minimum space requirement for the Data Volume?

   Choose the correct answers.

   [X] **A** 1.2 * RAM (TDI)

   [ ] **B** Min. 1 * RAM

   [ ] **C** Min. 0.5 * RAM

   [X] **D** 3 * RAM (Appliance)

   Correct! The sizing report shows the net data size on disk. To determine the required SAP HANA data volume size, add 20%. For an appliance approach the minimum data volume size is three times of the RAM size. Min. 1 * RAM is the minimum size for the /hana/shared/<sid> directory for all binary, trace and configuration files. Min. 0.5 * RAM is the minimum size for the log directory, having a system size <= 512 GB. Read more on this in the lesson Linux Operating System Requirements (Unit 2, Lesson 2) of the course HA200_14.

UNIT 3 SAP HANA Installation

UNIT OBJECTIVES

- Explain SAP HANA lifecycle management tools
- Explain the use of the command line options
- Explain a distributed system installation

Unit 3
Lesson 1

Introducing SAP HANA Lifecycle Management Tools

LESSON OVERVIEW

This lesson explains the various SAP HANA Lifecycle Management tools for installing the SAP HANA system.

Business Example

You want to install an SAP HANA single-host system and are investigating which SAP HANA Lifecycle Management tools are the best to use.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explain SAP HANA lifecycle management tools

Introduction to SAP HANA Lifecycle Management Tools

The release of SAP HANA SPS09 means that the SAP HANA lifecycle management (HDBLCM) tools replace all the other tools from previous releases. The SAP HANA unified installer, the on-site configuration tool, SAP SUM for HANA, hdbinst, and the SAP HANA lifecycle manager tools are all replaced by the SAP HANA lifecycle management tools.

> Note:
> SAP HANA installations are performed using the SAP HANA database lifecycle manager (HDBLCM).

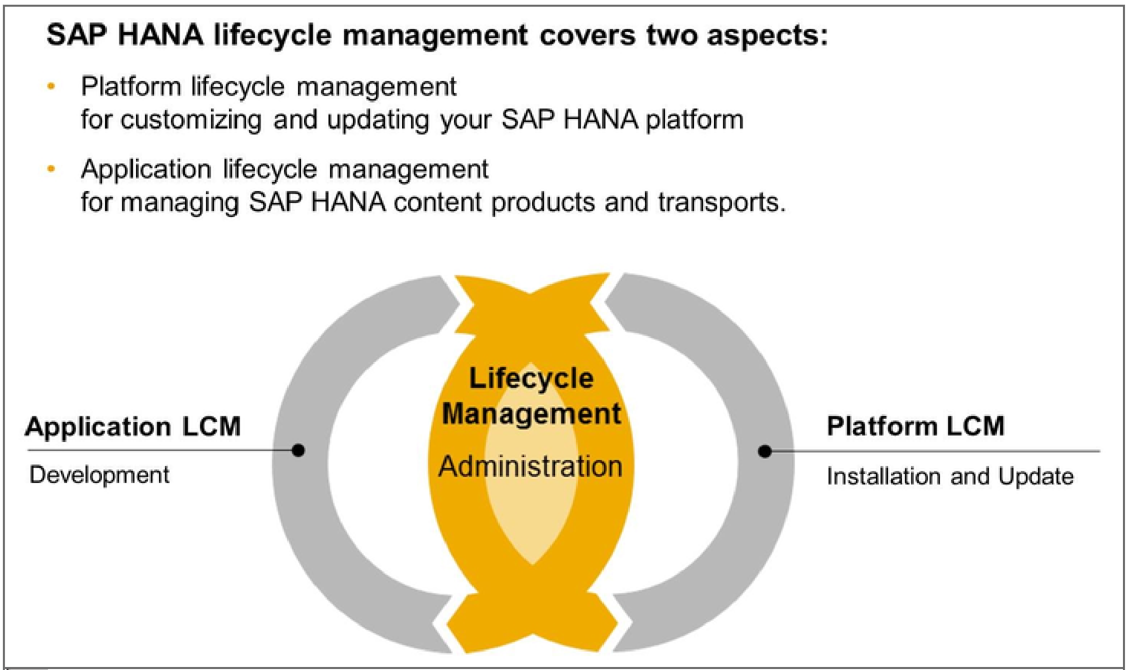

Figure 30: SAP HANA Lifecycle Management Aspects

Platform Lifecycle Management Aspects

The platform lifecycle management tasks on your SAP HANA system can be performed by using one of the three SAP HANA database lifecycle manager tool user interfaces.

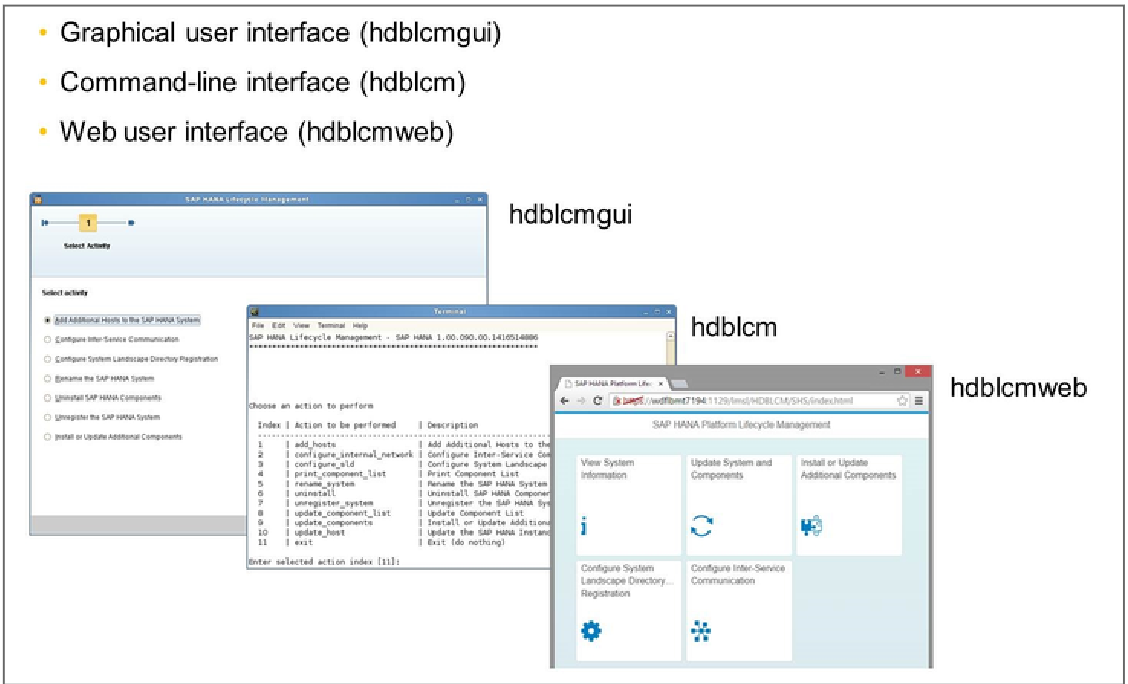

Figure 31: SAP HANA Database Lifecycle Manager User Interfaces

Note:
The Web interface can be used as a stand-alone via a Web browser or started from within SAP HANA studio.

Overview of HDBLCM Tool Tasks

SAP HANA platform lifecycle management tools can be used to install, configure, and update an SAP HANA server, adding both mandatory components and additional components. The tools can also be used to perform post-installation configuration tasks.

Note:
In general, installation and updates are carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM tool.

HDBLCM on installation media:

- Installation SAP HANA and its components
- Update SAP HANA and its components

Resident HDBLCM:

- Add Host Roles
- Add Hosts to the SAP HANA Database System
- Configure Inter-Service Communication
- Configure System Landscape Directory Registration
- Convert to Multitenant Database Containers
- Remove Host Roles
- Rename the SAP HANA Database System
- Uninstall SAP HANA Database Components
- Unregister the SAP HANA Database System
- Install or Update Additional Components

Figure 32: Overview of HDBLCM Tool Tasks

Location of the HDBLCM Tools

Different task are performed by the different HDBLCM tools.

Depending on the task, you need to select the correct HDBLCM tool. The figure, Location of the HDBLCM Tools, provides an overview of the tools and their specific tasks.

File location for the SAP HANA database lifecycle management tools

Tool	Usage	Location
SAP HANA HDBLCM	Installation and update	/<DVD>/DATA_UNITS/HDB_LCM_LINUX_X86_64 (Intel-Based) /<DVD>/DATA_UNITS/HDB_LCM_LINUX_PPC64 (IBM Power)
SAP HANA resident HDBLCM	Configuration tasks	/<sapmnt>/<SID>/hdblcm For example, /hana/shared/<SID>/hdblcm

Figure 33: Location of the HDBLCM Tools

Application Lifecycle Management (ALM) Aspects

SAP HANA application lifecycle management (ALM) tasks can also be performed using different user interfaces. The available interfaces are as follows:

- A Web interface
- A command-line tool (hdbalm)
- ALM integrated in SAP HANA studio

SAP HANA application lifecycle management provides support in all phases of the lifecycle of an SAP HANA application or add-on product. This includes modeling your product structure, application development, transport, assembly, installing, and updating products that you have downloaded from SAP Service Marketplace or that you have assembled yourself.

System administrators use SAP HANA application lifecycle management mainly to install and update SAP HANA applications or add-on products.

SAP HANA Platform Lifecycle Management Tools

The SAP HANA database lifecycle manager (HDBLCM) performs tasks such as installing, updating, and configuring an SAP HANA system. The SAP HANA database lifecycle manager helps hardware partners and administrators to perform their tasks efficiently.

The SAP HANA database lifecycle manager can be run with a graphical user interface, a command-line interface, a Web user interface in a browser, or from the SAP HANA Studio . It replaces the old tools completely.

First, decide which SAP HANA database lifecycle manager (HDBLCM) interface type you prefer to use. You can change the default behavior of the lifecycle management tools by using parameters. Parameters can be modified in a number of ways, for example, in the entry field of a graphical interface, as a call option with the program call, or in a configuration file. These options can be mixed and matched depending on the parameters that you need to use and the program interaction mode that you choose.

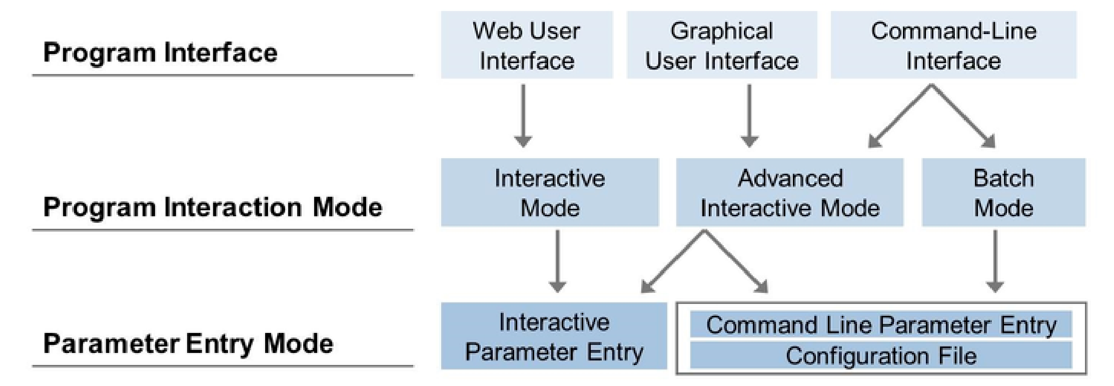

Figure 34: Many Paths Leading to a Common Goal

Once you have chosen the graphical user, command-line, or Web user interface, you can enter parameter values interactively. Alternatively, you can provide all of the required parameters with the call to the platform lifecycle management tool, and let it run unattended to completion.

Program Interaction Mode

Interactive mode is available for all user interfaces, and is the default mode for program interaction. To use the interactive mode, call the SAP HANA database lifecycle manager user interface, and enter the parameter values as they are requested by the program. Advanced interactive mode involves entering some parameter values interactively and providing some parameter values as call options or in a configuration file. This is the recommended interaction mode if you would like to modify parameter default values that are not requested in interactive mode.

Batch mode is an advanced interaction method for platform lifecycle management, as the call to the lifecycle management program on the command line must include all required parameters. Batch mode is designed for large-scale platform lifecycle management tasks, which would be time consuming to perform interactively.

Platform lifecycle management parameters can be entered interactively, although this feature is only available for interactive mode or advanced interactive mode. Alternatively, they can be entered as a call option on the command line, or via a configuration file. If you are performing platform LCM tasks in advanced interactive mode, you can choose any of the three parameter entry methods, or use more than one. If you are using batch mode, enter parameter values, either as call options to the SAP HANA database lifecycle manager, or from a configuration file.

Platform Lifecycle Management Tasks on the Graphical User Interface

SAP HANA platform lifecycle management tasks can be performed from a graphical interface. In the figure, HDBLCMGUI Versus Resident HDBLCMGUI , you see an example of the user interface.

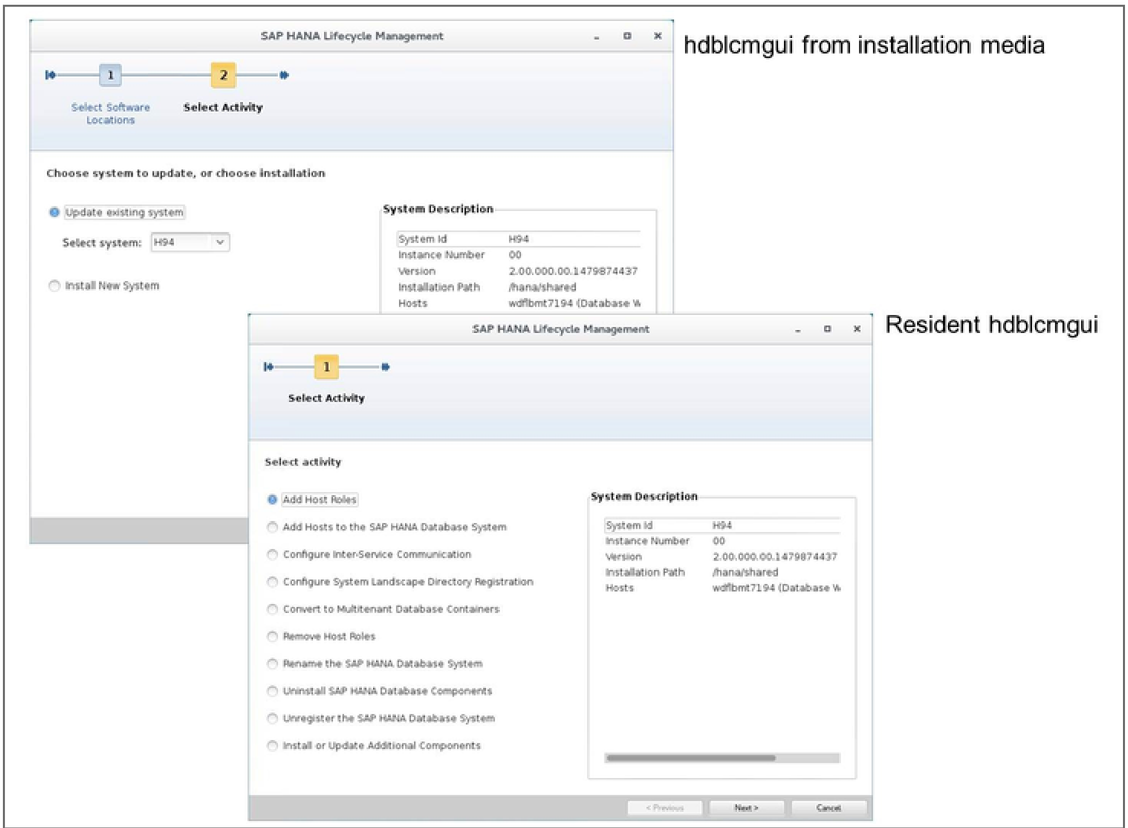

Figure 35: HDBLCMGUI Versus Resident HDBLCMGUI

In general, the installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM.

Start the graphical SAP HANA platform lifecycle management tool `hdblcmgui` from the appropriate directory.

Platform Lifecycle Management Tasks on the Command-Line Interface

SAP HANA platform lifecycle management tasks can be performed from the command line.

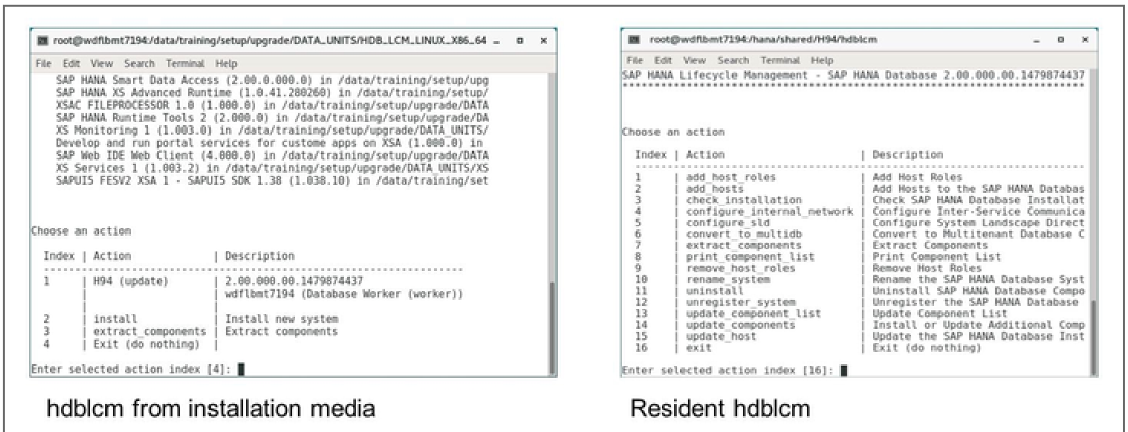

Figure 36: HDBLCM Versus Resident HDBLCM

In general, installation and updates are carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM.

Start the SAP HANA platform lifecycle management command line tool `HDBLCM` from the appropriate directory.

Platform Lifecycle Management Tasks on the Web User Interface

The SAP HANA database lifecycle manager (HDBLCM) can be accessed as a Web user interface using a HTML15 compatible browser.

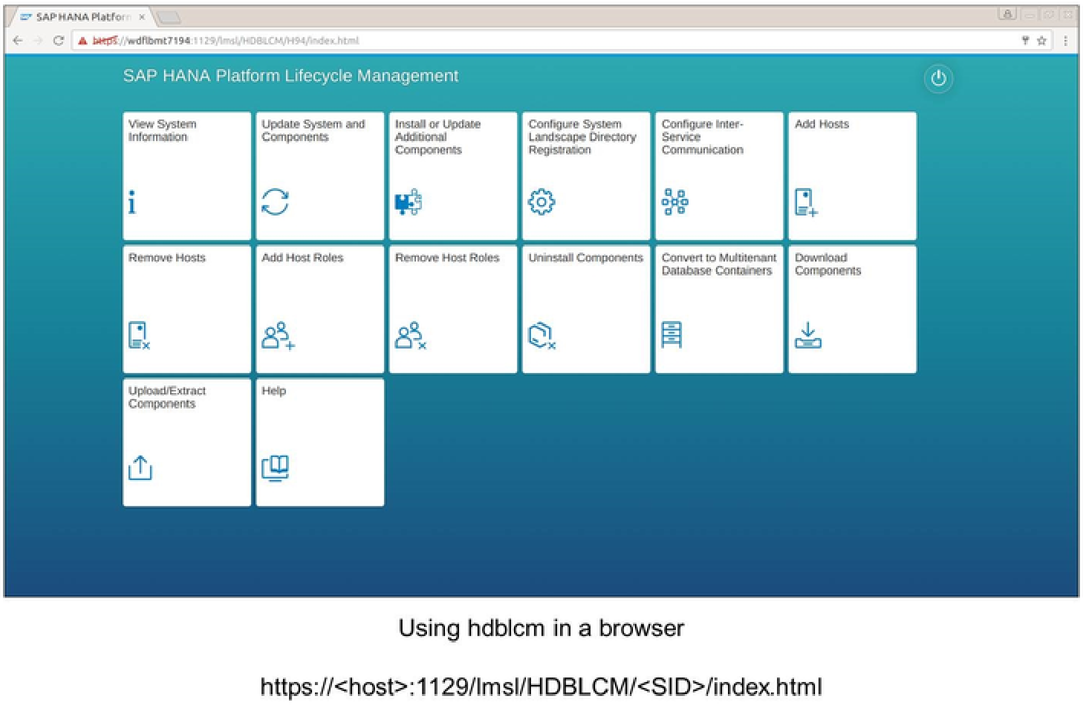

Figure 37: Web User Interface for HDBLCM

The prerequisites for using the Web user interface for HDBLCM are as follows:

- The SAP HANA database is revision 90 or higher.
- The communication port 1129 is open.

Several browsers are supported when using the Web user interface. The following Web browser are supported:

- Internet Explorer - Version 9 or higher
- Mozilla Firefox - Latest version and Extended Support Release
- Google Chrome - Latest version
- Safari 5.1 or higher on Mac OS

You have two options to start the Web interface, depending on whether you use SAP HANA Studio or a browser.

In SAP HANA Studio, open the Context menu of your system and choose Lifecycle Management → Platform Lifecycle Management → SAP HANA Platform Lifecycle Management.

The SAP HANA database lifecycle manager (hdblcmweb) Web user interface is hosted by the SAP Host Agent, which is installed on the SAP HANA host. To access the Web user interface, in the browser, open the URL https://<hostname>:1129/lmsl/HDBLCM/<SID>/index.html.

> Note:
> Do not start hdblcmweb manually. The executable hdblcmweb is started automatically by the SAP Host Agent as soon as an action is triggered from the Web user interface.

LESSON SUMMARY

You should now be able to:

- Explain SAP HANA lifecycle management tools

Unit 3
Lesson 2

Describing Advanced Installation Options

LESSON OVERVIEW

This lesson explains the various advanced installation methods of an SAP HANA system.

Business Example

You want to install several SAP HANA systems and need insight to the advanced, batch oriented installation methods that are available for installing multiple SAP HANA systems.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explain the use of the command line options

Advanced Installation Options

Installation automation is designed for those who are familiar with SAP HANA, and are installing it regularly in various production environments. It refers to the installation of SAP HANA systems using batch mode, with a combination of a configuration file and call options passed on the command line.

To provide flexibility, you can install the same SAP HANA system in several ways. The differences between the installation methods are best shown by a one-to-one comparison of the same system installed with each available method. The figure, Advanced Installation Examples , shows the specifications for an installation that illustrate the differences between the installation methods.

Advanced installation examples:

- Using the command line with parameters.
- Batch installation using the command line with parameters.

The goal is to install a SAP HANA single-host system using the command line.
The installation has the following specifications.

Single-host system name: P01

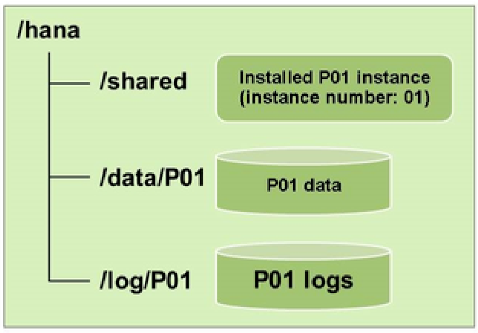

Definition	Value
System name	P01
Instance number	01
Installation path	/hana/shared
Data path	/hana/data/P01
Log path	/hana/log/P01
Install components	Server and client

Figure 38: Advanced Installation Examples

Using the Command Line

The hdblcm tool on the command line can be used in the two different ways. Both methods can be used in a batch mode, which doesn’t require any user interaction, as the following options show:

- Command line options (in batch mode)
- Configuration file (in batch mode)

Using the Command line with Parameters

Installing an SAP HANA system using the Command Line with parameters

```
./hdblcm --action=install --sid=P01 --number=01 --components=server,client
```

Definition	Specification	Simplified Command Line Parameter
System name	P01	`--sid=P01`
Instance number	01	`--number=01`
Installation path	/hana/shared	`<default>`
Data path	/hana/data/P01	`<default>`
Log path	/hana/log/P01	`<default>`
Install components	server and client	`--components=server,client`

Info: For all other input parameters, use the default value.

Figure 39: Using the Command Line with Parameters

As shown in the figure, Using the Command Line with Parameters , the SAP HANA installation is performed semiautomatically. The parameters specified on the command line are used, but the installer must still confirm the default parameter values. This can be improved by adding the command line option `-- batch`. This is shown in the figure, Using the Command Line in Batch Mode . In batch mode, the default values for unspecified parameters are used without confirmation. The mandatory parameters are still requested.

Caution:
The password parameters are mandatory, so they must be provided.

Using the Command Line in Batch Mode

Batch installing a SAP HANA system using the command line with parameters

```
./hdblcm --batch --action=install --sid=P01 --number=01 --components=server,client
--password=Abcd1234 --system_user_password=1234abcD
```

Definition	Specification	Command line parameter
System name	P01	`--sid=P01`
Instance number	01	`--number=01`
Installation path	/hana/shared	`<default>`
Data path	/hana/data/P01	`<default>`
Log path	/hana/log/P01	`<default>`
Install components	server and client	`--components=server,client`
Password <sid>adm	Abcd1234	`--password=Abcd1234`
Password system user	1234abcD	`--system_user_password=1234abcD`

Info: For all other input parameters, the default value are used.

Figure 40: Using the Command Line in Batch Mode

In the figure, Using the Command Line in Batch Mode , the mandatory parameters for the passwords are provided also. However, avoid providing the password on the command line like this because the statement is stored in the history file on Linux. Instead, create a configuration file that holds all the required parameters.

Using a Configuration File

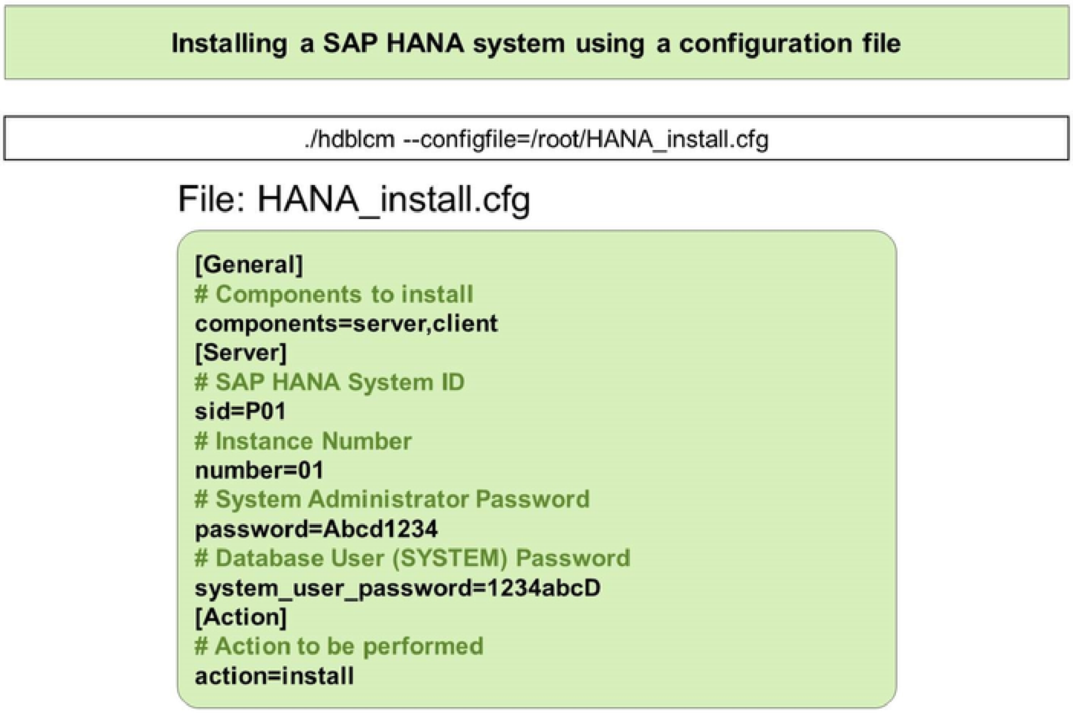

Figure 41: Using the Configuration File

Using a Configuration File and the Batch Mode

To perform an automated installation with the SAP HANA lifecycle management tool hdblcm, you must combine the configuration file and the batch mode.

Note that, until now, you had to enter passwords interactively or specify them on the command line. The batch mode is designed to automate the installation process.

The batch mode runs the installer without asking for any confirmation or parameter entry. This allows the installation to run to completion without any user interaction. It can be started from the command line with the use of a configuration file.

Installing a SAP HANA system using a configuration file and the batch mode

```
./hdblcm --batch --configfile=/root/HANA_install.cfg
```

File: HANA_install.cfg

```
[General]
# Components to install
components=server,client
[Server]
# SAP HANA System ID
sid=P01
# Instance Number
number=01
# System Administrator Password
password=Abcd1234
# Database User (SYSTEM) Password
system_user_password=1234abcD
[Action]
# Action to be performed
action=install
```

Figure 42: Using Batch Mode and the Configuration File

With the configuration file and the batch mode, the SAP HANA installation is installed completely without user interaction. This is useful if you want to set up many systems with a standard setup, or if you want to redeploy a system on a weekly basis because of system copies.

Changeable Installation Parameters

Unless you choose to change them, the SAP HANA database lifecycle manager (HDBLCM) uses default values during installation. Some default values are based on the predefined values on the current host.

Several of the default values are shown in the figures, Default Parameters, and Additional Default Parameters . This list is not complete because it is updated with every SAP HANAsupport package stack.

> Caution:
> In a multiple-host system, manually check the mandatory values on each host before installation.

Parameter	System Default Value	Interactive Mode Availability
autostart	`0` (off)	■
certificates_hostmap	**<current host>**	■
client_path	`${sapmnt}/${SID}/hdbclient`	◆
components	`client, server, studio` (dependent on the installer finding installation sources for the component)	■
copy_repository	`/hana/shared/${SID} /hdbstudio_update`	◆
datapath	`/hana/data/${SID}`	■
db_mode	`single_container`	■
groupid	`79`	◆
home	`/usr/sap/${SID}/home`	■
hostname	**<current host>**	■
install_hostagent	`y` (on)	◆
install_ssh_key	`y` (on)	◆

Figure 43: Default Parameters

Additional Default Parameters

Parameter	System Default Value	Interactive Mode Availability
logpath	`/hana/log/${SID}`	■
max_mem		■
number	**<next successive un-used instance number on the host>**	■
remote_execution	`ssh`	◆
restrict_max_mem	(off)	■
root_user	`root`	◆
sapmnt	`/hana/shared`	■
shell	`/bin/sh`	■
studio_path	`${sapmnt}/${SID}/hdbstudio`	◆
studio_repository	`1` (on)	◆
system_usage	`custom`	■
userid	**<next successive un-used user ID on the host>**	■

Figure 44: Additional Default Parameters

> Hint:
> For the complete list of changeable parameters, see the SAP HANA Installation Guide .

Users Created During Installation

Table 3: User Descriptions

During Installation the following users are created automatically:

User	Description
<sid>adm	The operating system administrator. The user <sid>adm is the operating system user required for administrative tasks such as starting and stopping the system. The user ID of the <sid> adm user is defined during the system installation. The user ID and group ID of this operating system user must be unique and identical on each host of a multiple-host system. The password of the <sid>adm user is set during installation with the password parameter.
sapadm	The SAP Host Agent administrator. If there is no SAP Host Agent available on the installation host, it is created during the installation, along with the user sapadm. If the SAP Host Agent is already available on the installation host, it is not modified by the installer. The sapadm user and password are also not modified. The password of the sapadm user is set during installation with the sapadm_password parameter.
SYSTEM	The database superuser. Initially, the SYSTEM user has all system permissions. Additional permissions can be granted and revoked again. However the initial permissions can never be revoked. The password of the SYSTEM user is set during installation with the system_user_password parameter.

Troubleshooting of Failed Installations

Refer to troubleshooting if the installation fails for an unknown reason, or for work-arounds in special circumstances.

Checking the Log Files

The SAP HANA lifecycle management tools hdblcm and hdblcmgui write log files during installation. The most recent log file is always available under `/var/tmp/hdblcm.log` or `/var/tmp/hdblcmgui.log`. Additionally, a copy of the log files is archived in the directory `hdb_<SID>_hdblcm_<action>_<date>`.

Because the SAP HANA lifecycle management tools hdblcm and hdblcmgui are wrappers for underlying component installers, you can also check the component logs. Review and analyze the logs for the SAP HANA lifecycle management tools hdblcm and hdblcmgui first. Once you limit the source of the problem to a specific component, then you can analyze the component logs further.

The component log files are stored in the following path: /var/tmp/hdb_<SID>_<action>_<time_stamp> where <action> refers to install, update, addhost, uninstall, and so on.

The following log files are written while performing the action:

- <hdbcommand>.log: Can be read using a text editor
- <hdbcommand>.msg: XML format for the display in the installation tool with the GUI
- <hostname>_tracediff.tgz: Provides a delta analysis of the original trace files, and makes a detailed analysis easier

You can also view the last three log files in the SAP HANA studio using the administration function Diagnosis Files .

Installer Trace Enablement

If the SAP HANA installer crashes or loops, trace the installer until the problem is found. To switch on the installer trace, set the environment variable HDB_INSTALLER_TRACE_FILE to <tracefilename>. The directory containing the trace file must already exist.

After the trace is generated, you can open it and check the trace file for error messages. If needed, open an SAP support ticket on http://support.sap.com, and attach the trace file for further analysis.

Location of SAP HANA File System Components

In addition to the main components installed in the default file systems, you can also locate the temporary files from the SAP HANA system. They can be found in the directories shown in the following figure.

SAP HANA database files

- /var/lib/hdb - IPC data (volatile)
- /var/tmp - Installer log files, HDB_alive_*
- /tmp - .hdb_*_lock (volatile)

The sapstartsrv files

- /etc/init.d/sapinit* - Boot script
- /tmp/.sap* - Sockets, lock files (volatile)

Figure 45: Locating All SAP HANA File System Components

Access to Underlying Installer Components

Because hdblcm and hdblcmgui are wrapper tools, in some troubleshooting cases, you can pass component options on to the underlying component tools (hdbinst or hdbupd) in combination with the call to the hdblcm or hdblcmgui SAP HANA lifecycle management tools.

To view the available underlying component parameters as extended help output, use the pass_through_help parameter. Specify the action parameter and --help or -h in combination with pass_through_help.

LESSON SUMMARY
You should now be able to:

- Explain the use of the command line options

Unit 3
Lesson 3

Explaining a Distributed System

LESSON OVERVIEW

This lesson describes what you need to install a distributed system, and explains how to perform this installation.

The installation of a distributed system is described in the SAP HANA Server Installation Guide.

Business Example

A distributed landscape consisting of multiple hosts provides more memory and more CPU power beyond the limitation of a single physical hardware box.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explain a distributed system installation

Multi-Host System Installation

It is important to understand the multi-host system concepts, such as host grouping and storage options, before installing a multi-host system.

When configuring a multi-host system, define the additional hosts as worker nodes or standby nodes. The worker option is the default option. The following host types can have high availability:

- **Worker** nodes process data. They are also called compute nodes.
- **Standby** nodes wait to take over processes of a failed worker node.

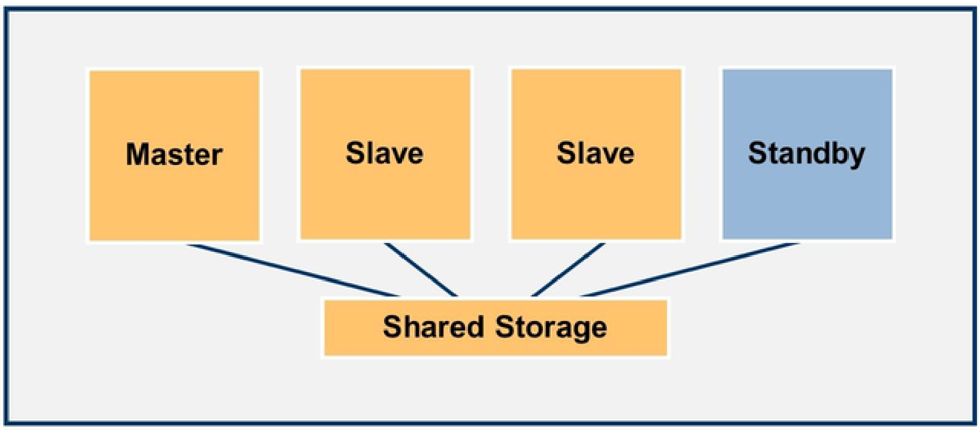

Figure 46: Scale Out

Another important term is the server role. There are two types of server roles, as follows:

- MASTER role

  The master index server is assigned on the same host as the name server with the actual role MASTER. The actual index server role of this host is MASTER. The master index server provides metadata for the other active index servers (that is, those with actual index server role SLAVE).

- SLAVE role

  The index server role of the remaining hosts (except those configured as standby hosts) is SLAVE. These are active index servers and are assigned to one volume. If an active index server fails, the active master name server assigns its volume to one of the standby hosts.

> Note:
> SAP recommends that all servers have the same size.

A Typical Configuration for an Eight-Nodes Distributed System

Host	Name Server Configured Role	Name Server Actual Role	Index Server Configured Role	Index Server Actual Role
Initial host	Master 1	Master	Worker	Master
1st host added	Master 2	Slave	Worker	Slave
2nd host added	Master 3	Slave	Worker	Slave
3rd host added	Slave	Slave	Worker	Slave
4th host added	Slave	Slave	Worker	Slave
5th host added	Slave	Slave	Worker	Slave
6th host added	Slave	Slave	Worker	Slave
7th host added	Slave	Slave	Standby	Standby

Figure 47: A Typical Configuration for an Eight-Nodes Distributed System

Host grouping does not affect the load distribution among worker hosts; the load is distributed among all workers in an SAP HANA system.

Host Grouping

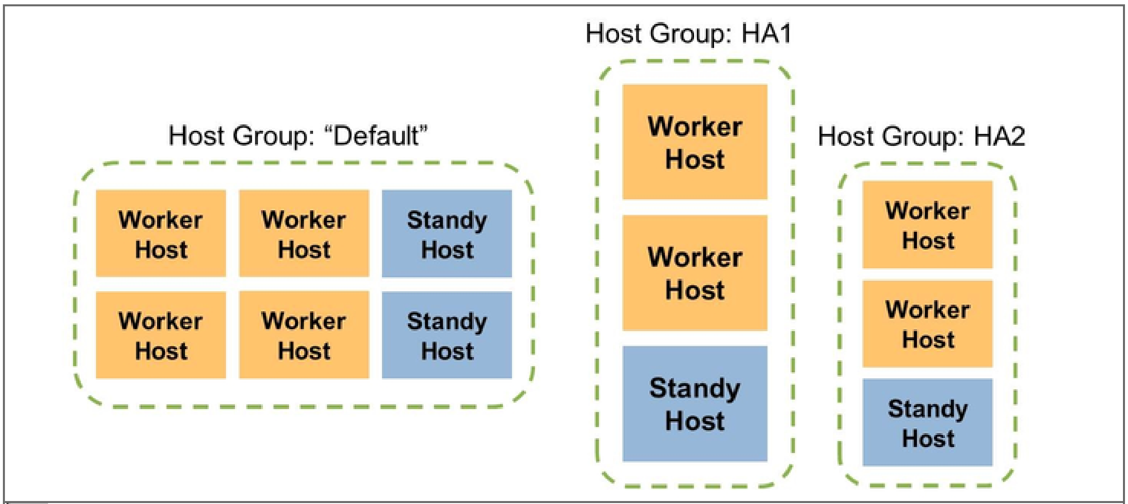

Figure 48: Host Grouping

If there are multiple standby hosts in a system, consider host grouping. Host grouping decides the allocation of standby resources if a worker machine fails. If no host group is specified, then all hosts are allocated to one host group called "default" . If the number of standby hosts in one host group increases, more will fail over security.

Note:
The installer distinguishes between the following types of groups:

- SAP system group (sapsys group): This defines all hosts in a system. Therefore, all hosts in a multi-host system must have the same SAP system group ID, which is the default configuration with hdblcm.
- Host group: This contains hosts that share the same standby resources only. Therefore, if the multi-host system has one standby host, leave all hosts in the same host group ("default"). This ensures that all hosts have access to the standby host in case a worker host fails.

Distributed Systems and Scale Out

Note the following information:

- SAP HANA is the name for multiple connected nodes of an SAP HANA database that use the same server software installation.
- Every system has a unique SAP system ID. This is called the `<SID>`.

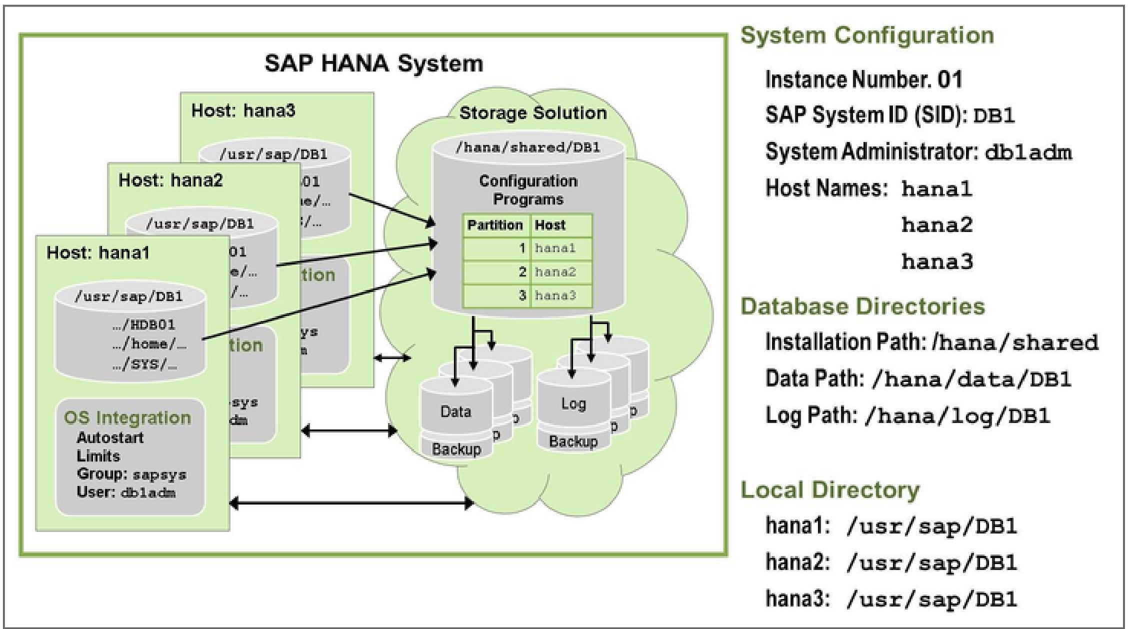

Figure 49: Multi-Host System

Both the hdblcm and hdblcmgui SAP HANA lifecycle management tools can install an SAP HANA multi-host system in one of the installer modes and with a combination of parameter specification methods.

Creating a Multi-Host System During Installation

The SAP HANA lifecycle management tools, hdblcm and hdblcmgui, can build a multi-host system during installation in interactive mode, in batch mode, and with the available parameter specification methods: interactively, using command line options, or with the configuration file.

The prerequisite for creating a multi-host system is that the shared file systems for the data files and log files are configured so that they are present and mounted on all hosts, including the primary host.

The suggested locations for the file systems are as follows:

- `/hana/shared`
- `/hana/data/<sid>`
- `/hana/log/<sid>`

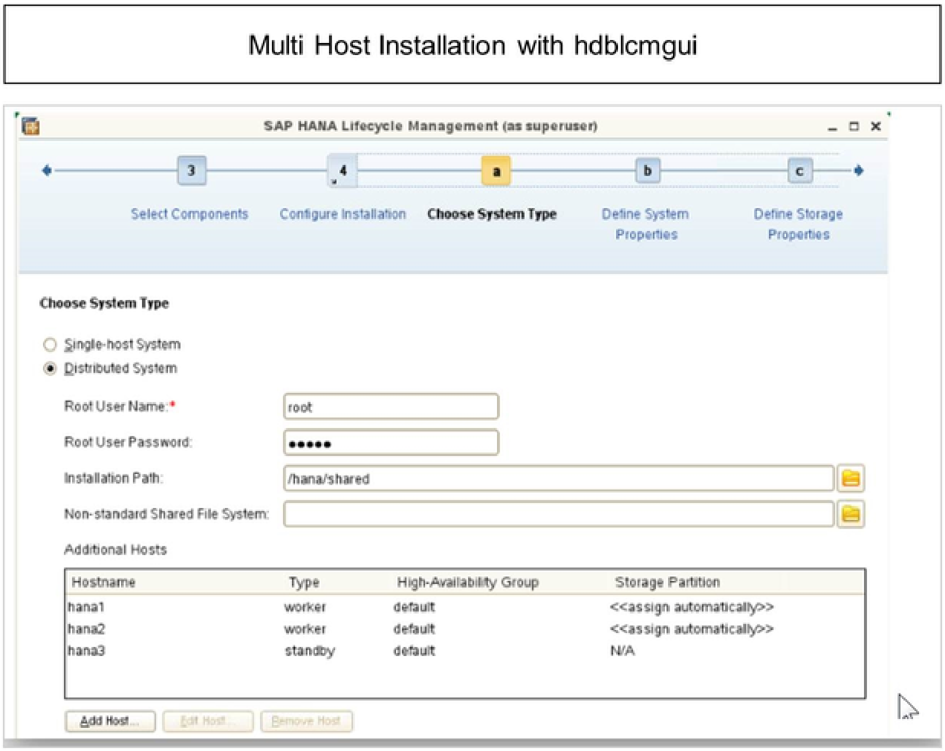

Figure 50: Multi-Host Installation

Test and Simulation

For testing and debugging, you can copy a scale-out landscape to a single node. To set up the testing and debugging system, perform a system copy. Another unit in the course explains the methods that can be used.

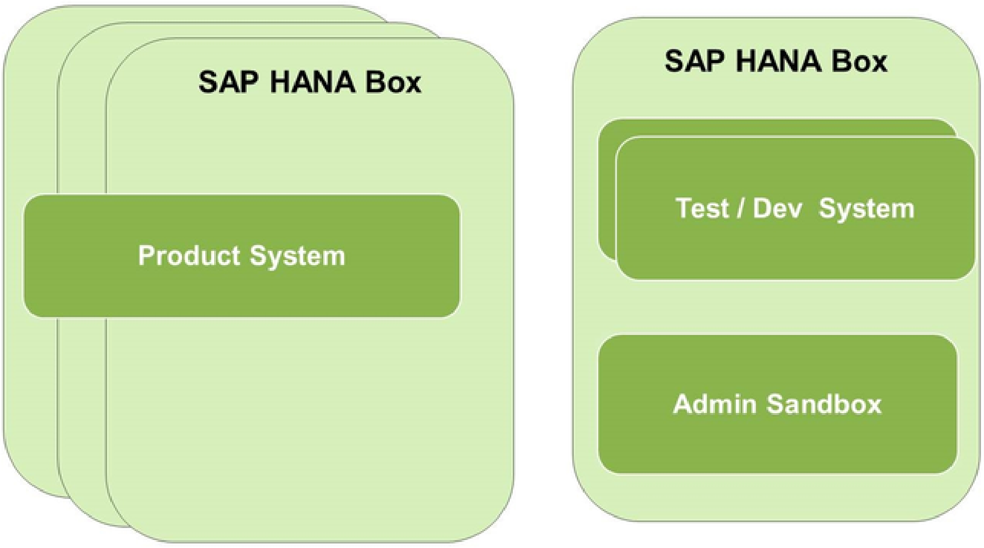

Figure 51: Test And Simulation

Hint:
Use a separate sandbox system to test all of the administrator tasks, such as backup and recovery.

Storage Options

In single-host SAP HANA systems, you can use plain attached storage devices, such as Small Computers System Interface hard drives, solid-state drives (SSDs), or storage area networks (SANs). However, to build a multi-host system with failover capabilities, the storage must ensure the following:

- The standby host has file access.
- The failed worker host no longer has access to write to files, called fencing.

There are two fundamentally different storage configurations that meet these two conditions: shared storage devices, and separate storage devices with failover reassignment. A shared storage subsystem, such as Network File System (NFS) or IBM's General Parallel File System (GPFS), is the commonly used storage option because it is easy to ensure that the standby host has access to all active host files in the system.

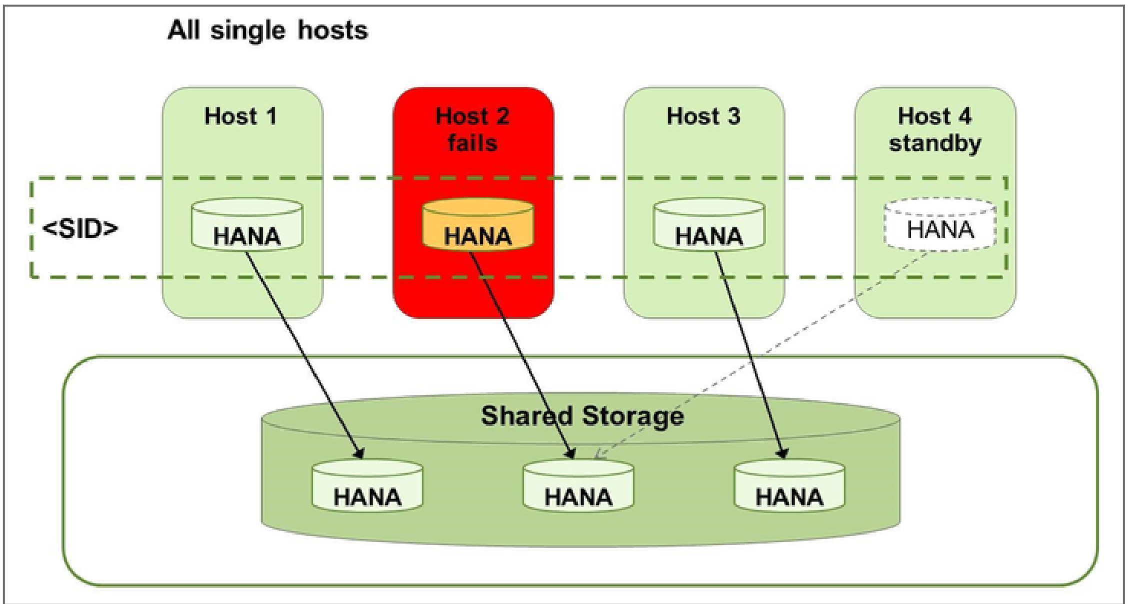

Figure 52: Shared Storage

Shared Storage System

In a shared storage solution, the externally attached storage subsystem devices can provide dynamic mount points for hosts. Because shared storage subsystems vary in their handling of fencing, it is the responsibility of the hardware partner and their storage partners to develop a corruption-safe failover solution.

A shared storage system can be configured as shown in the figure, Shared Storage System .
However, mounts can differ among hardware partners and their configurations.

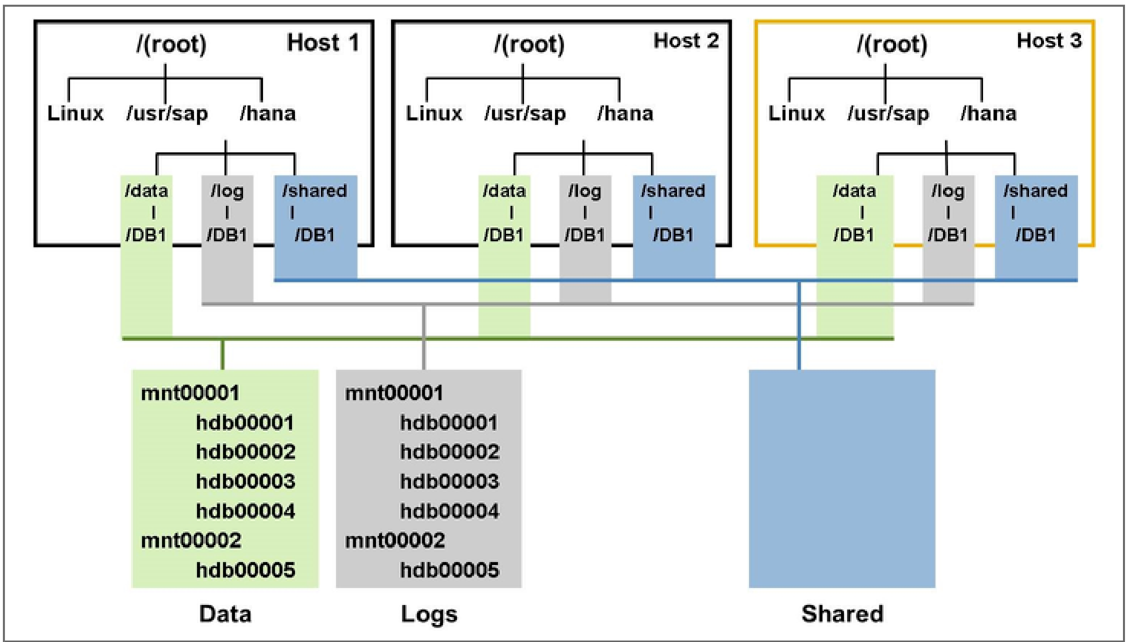

Figure 53: Shared Storage System

LESSON SUMMARY

You should now be able to:

- Explain a distributed system installation

Learning Assessment

1. Which of the following SAP HANA Lifecycle Management Tools can be used for configuration works?

   Choose the correct answer.

   - [ ] **A** hdblcm
   - [ ] **B** hdblcmgui
   - [ ] **C** Resident hdblcm
   - [ ] **D** hdblcmweb

2. During the Installation, the following users are created automatically. Which of them is the SAP Host Agent administrator?

   Choose the correct answer.

   - [ ] **A** <sid>adm user
   - [ ] **B** sapadm user
   - [ ] **C** SYSTEM user

3. The shared file systems for data and log files are configured and mounted on all hosts, including the primary host automatically, if you choose to build a multi-host system during the installation.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

Unit 3

Learning Assessment - Answers

1. Which of the following SAP HANA Lifecycle Management Tools can be used for configuration works?

   Choose the correct answer.

   [ ] **A** hdblcm

   [ ] **B** hdblcmgui

   [X] **C** Resident hdblcm

   [ ] **D** hdblcmweb

   Correct! Resident hdblcm is the tool used for after installation activities like configuring the system, adding or removing hosts, converting to multitenant database containers, and so on. hdblbm is the command-line interface for installation or upgrade tasks. hdblcmgui is the graphical user interface for installation or upgrade tasks. hdblcmweb is the Web user interface for installation or upgrade tasks. Read more on this in the lesson Introducing SAP HANA Lifecycle Management Tools (Unit 3, Lesson 1) of the course HA200_14.

2. During the Installation, the following users are created automatically. Which of them is the SAP Host Agent administrator?

   Choose the correct answer.

   [ ] **A** <sid>adm user

   [X] **B** sapadm user

   [ ] **C** SYSTEM user

   Correct! sapadm is the SAP Host Agent administrator. The password is set during installation with the “sapadm_password” parameter. The <sid>adm user is the operating system administrator. The SYSTEM user is the database superuser. Read more on this in the lesson Describing Advanced Installation Options (Unit 3, Lesson 2) of the course HA200_14.

3. The shared file systems for data and log files are configured and mounted on all hosts, including the primary host automatically, if you choose to build a multi-host system during the installation.

   Determine whether this statement is true or false.

   [ ] True

   [X] False

   Correct! The prerequisite for creating a multi-host system is that the shared file systems for the data files and log files are configured so that they are present and mounted on all hosts, including the primary host. The precondition is to have the shared files systems set up and ready to use before the installation of a multi-host system installation starts. The suggested locations for the file systems are: “/hana/shared”, “/hana/data/<sid>” and “/hana/log/<sid>”. Read more on this in the lesson Explaining a Distributed System (Unit 3, Lesson 3) of the HA200_14.

UNIT 4 SAP HANA Architecture

UNIT OBJECTIVES

- Explain the SAP HANA Architecture and Technology
- Explaining SAP HANA Memory Management and Data Persistence

Unit 4
Lesson 1

SAP HANA Architecture and Technology

LESSON OVERVIEW

This lesson gives a brief overview of the architecture of multitenant database containers.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explain the SAP HANA Architecture and Technology

Overview of Architecture and Technology

An SAP HANA system is a database management system (DBMS); that is, it is a server component that manages a database model. A database management system captures and analyzes data by interacting with the user, other applications, and the database itself. A general-purpose DBMS allows the definition, creation, querying, update, and administration of databases.

A DBMS performs the following functions:

- Manages large amounts of data in a multiuser environment so that many users can concurrently access the same data
- Maintains relationships between data
- Provides secure access to data using the user authorization concept
- Recovers data automatically to the most recent consistent status after a system failure
- Delivers high performance for processing data requests

Database Terminology

A database is an organized collection of data and metadata. It is the collection of schemata, tables, queries, reports, views, and other objects. Because the database is only a passive part of a database server, it requires some processes and memory structures to access the data and manage the database. The combination of processes and memory buffers is called a database instance.

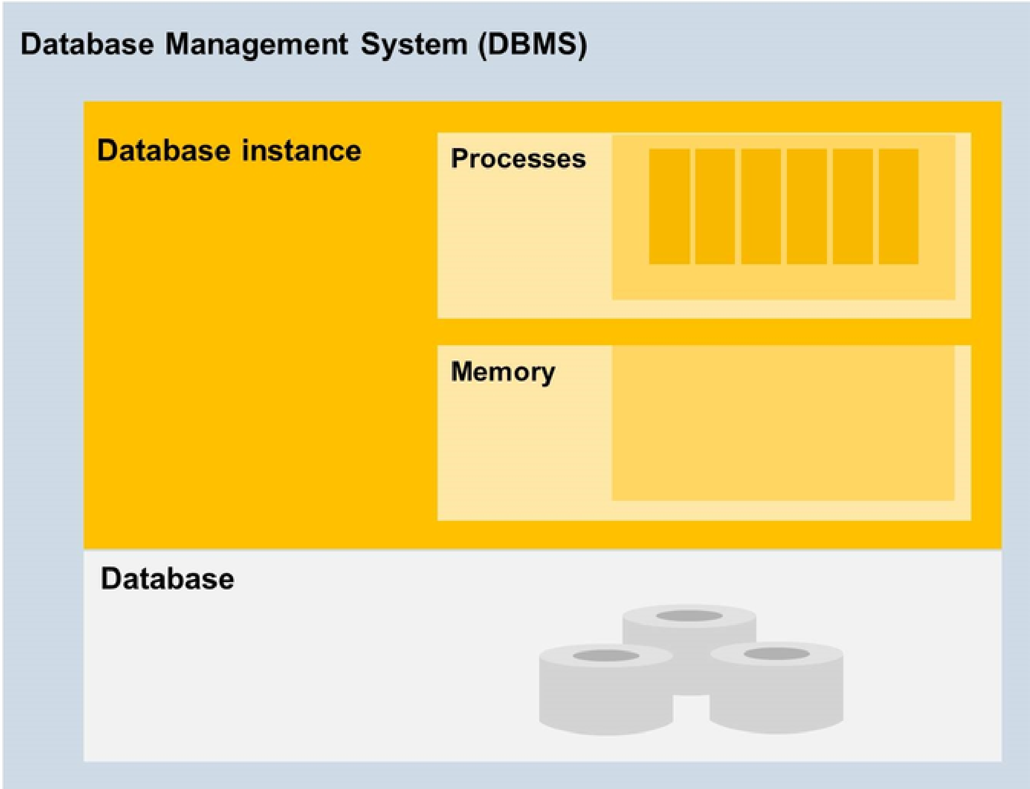

Figure 54: Database Terminology

Single-Container Versus Multiple-Container System

An SAP HANA system can be installed as a single-container system or as a multiple-container system. A single-container system consists of one database that is managed by the SAP HANAdatabase management system. An SAP HANA system installed in multiple-container mode can contain more than one multitenant database container.

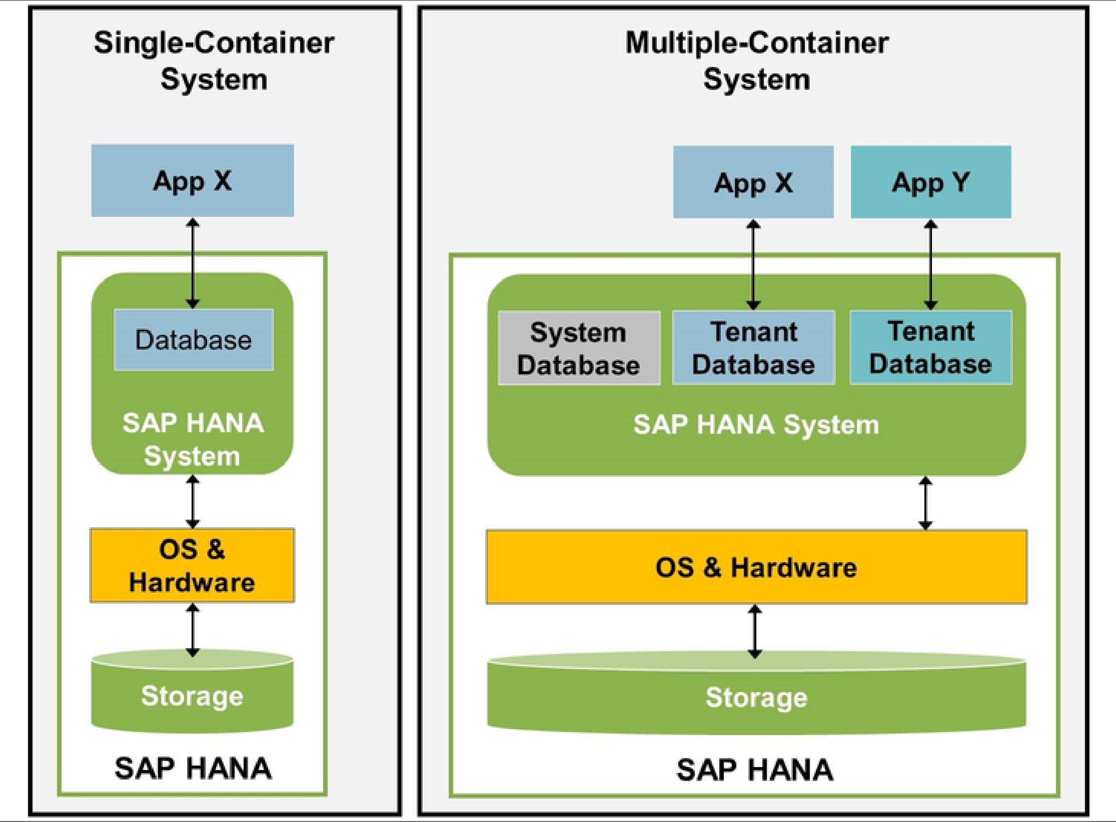

Figure 55: Single-Container Versus Multiple-Container System

The concept of a multitenant database container system is based on a single SAP HANA system, or database management system, with a single system ID. This single system ID contains at least one tenant database, in addition to a system database. The system database keeps the system-wide landscape information, and provides system-wide configuration and monitoring. Users of one tenant database cannot connect to other tenant databases. They also cannot access application data there, unless the system is enabled for cross database access. The tenant databases are, by default, isolated from each other in terms of application data and user management. Each tenant database can be backed up and recovered independently from one another. Because all tenant databases are part of the same SAP HANA database management system, they all run with the same SAP HANA version (revision number). In addition, the defined high availability disaster recovery scenario applies to all tenant databases.

A multiple-container system always has exactly one system database, used for central system administration, and any number of multitenant database containers (including zero), also called tenant databases. An SAP HANA system installed in multiple-container mode is identified by a single system ID (SID). Databases are identified by a SID and a database name. From the administration perspective, there is a distinction between tasks performed at system level and those performed at database level. Database clients, such as the SAP HANA studio, connect to specific databases.

In a multiple-container system, only the system database runs the name server. The name server contains landscape information about the system as a whole, including which tenant databases exist. It also provides indexserver functionality for the system database. Unlike the name server in a single-container system, the name server of the system database in a multiple-container system does not own topology information. That is, it does not own information about the location of tables and table partitions in databases. Database-related topology information is stored in the relevant tenant database catalog.

Tenant databases require only their own index server. Servers that do not persist data, such as the compile server and the preprocessor server, run on the system database and serve all databases. The server for SAP HANA extended application services runs embedded in the (master) index server of the tenant database by default. However, it can be added as a separate service, if necessary.

> Note:
> As of SAP HANA 2.0 SPS 01 , the multi-container database mode is the only database mode. By default, a single tenant database is created during installation. The upgraded system will have one tenant database that corresponds to the old single container. You can add additional tenant databases later using the SAP HANA cockpit.

A single-container system will automatically convert to a tenant database system during the update. The database of a single-container system is converted into a system database and a tenant database. The upgraded system will comprise one tenant database that corresponds to the old single container. The system database (SYSTEMDB) creates a new user (SYSTEM). During the update, a password must be given for this user. The database superuser (SYSTEM) of the single-container system becomes the SYSTEM user of the tenant database. You may have to adapt your operations concept to include the new system database.

Features of SAP Web Dispatcher

The SAP Web Dispatcher runs as a separate database service on the system database. It routes incoming HTTP requests from clients to the correct server for SAP HANA extended application services, based on virtual host names. This is part of network configuration.

All the databases in a multiple-container system share the same installation of database system software, the same computing resources, and the same system administration. They have the following features:

- Shared installation of database system software
- Tenant databases identified by name or port
- Additive sizing for all tenant database
- Strong isolation features, so that each tenant database has the following:
  - Database administration and end users, database catalog, repository, persistence, backups, traces, and logs
  - Tenants memory sizing and CPU consumption, which can be configured independently
- Integration with SAP HANA data center operation procedures, housekeeping, backups, and so on

Multitenant Database Containers: Architecture

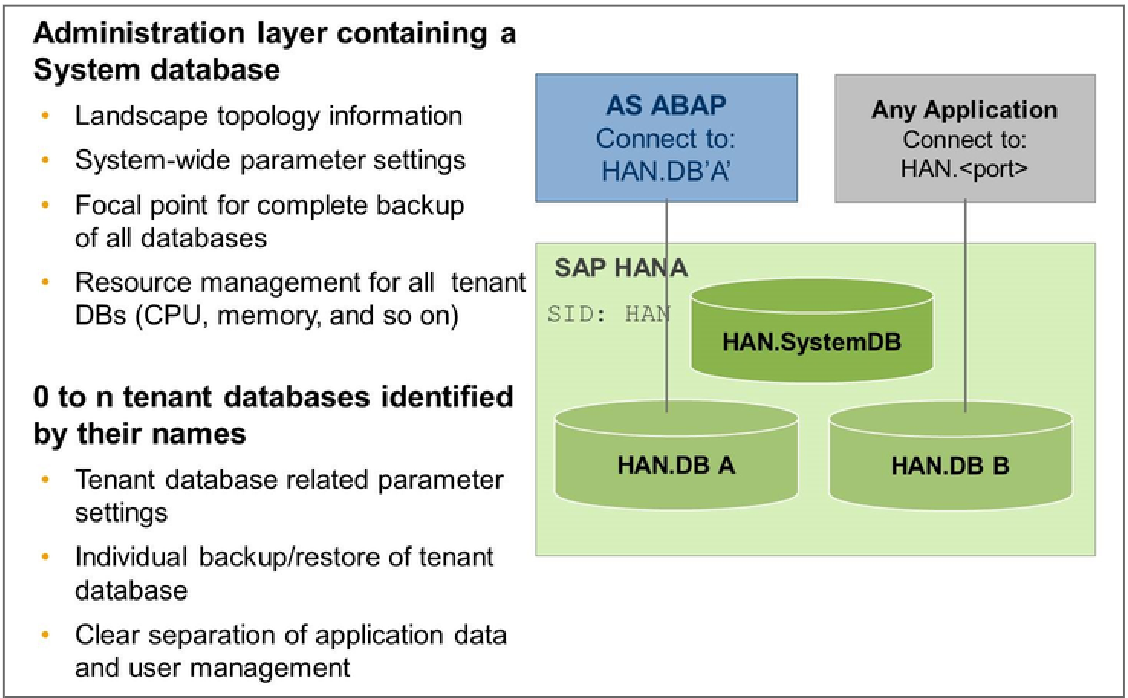

Figure 56: Multitenant Database Containers: Architecture

During the installation of a multiple-container system, the system database is created . It contains information about the system as a whole and all tenant databases. It is used for central system administration.

If you use a multiple-container system, you have one system database and any number of tenant databases. Multiple applications run in different tenant databases. You can use this deployment option to replace existing MCOS on-premise scenarios.

MDC: User and Administration Layers

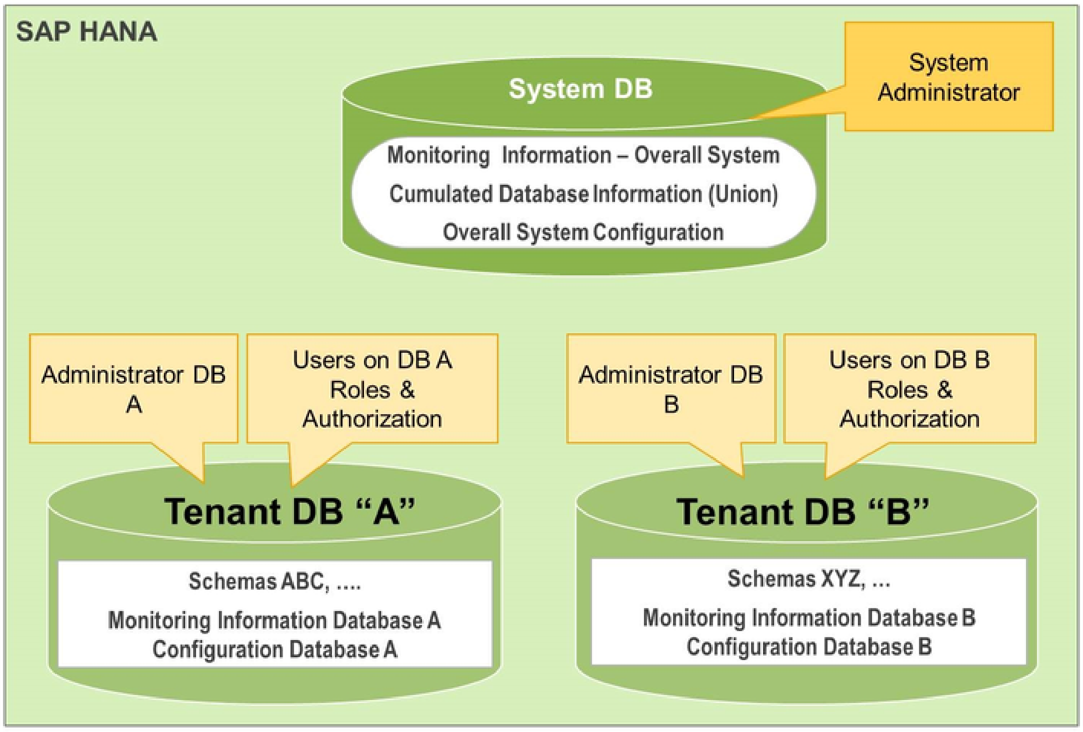

Figure 57: MDC: User and Administration Layers

Administration tasks performed in the system database apply to the system as a whole and to all of its databases (for example, system-level configuration settings). Alternatively, it can target specific tenant databases (for example, backup of a tenant database).

The System Database

System database creation occurs during either installation of a multiple-container system or during the conversion from a single-container system to a multiple-container system. The system database contains information about the system as a whole, as well as all its tenant databases. It is used for central system administration.

A multiple-container system has exactly one system database. It is created during system installation or during the migration from a single-container system. It contains the data and users for system administration. System administration tools, such as the SAP HANA cockpit or the SAP HANA Studio , can connect to this database. The system database stores overall system landscape information, including information on the tenant databases that exist in the system. However, it doesn't own database-related topology information, that is, information about the location of tables and table partitions in databases. Database-related topology information is stored in the relevant tenant database catalog.

Administration tasks performed in the system database apply to the system as a whole and to all of its databases (for example, system-level configuration settings). Alternatively, it can target specific tenant databases (for example, backup of a tenant database).

Points to Note About the System Database

Note the following information about the system database:

- The system database is not a database with full SQL support.
- The system database cannot be distributed across multiple hosts: scale-out is not possible.
- If you need a full-featured SAP HANA database in a multiple-container system, create at least one tenant database.
- The system database can show monitoring data from tenant databases (views in the schema SYS_DATABASES), but it can never show actual content from tenant databases.

MDC: Database Isolation

The Database Isolation specifies the isolation of the tenant databases on the operating system level for multitenant database container SAP HANA systems. By default, all database processes in a multiple-container system run under the OS user <sid>adm. If you want to mitigate against cross-database attacks through OS mechanisms, configure the system for high isolation. In this way, the processes of individual tenant databases must run under dedicated OS users belonging to dedicated OS groups. Database-specific data on the file system is then protected with standard OS file and directory permissions.

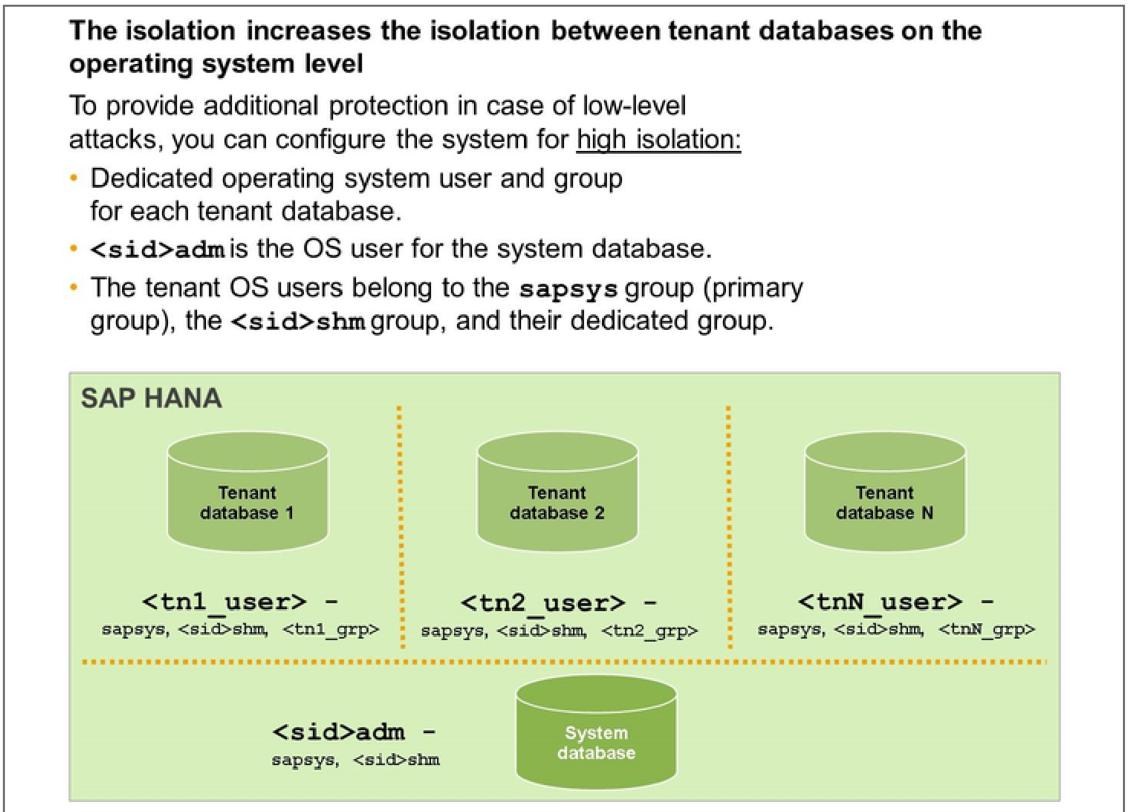

Figure 58: MDC: Database Isolation

Properties of MCD Database Isolation

The properties of a system with a high isolation level are as follows:

- Processes of individual tenant databases run under the dedicated OS users belonging to dedicated OS groups.

- Database-specific data on the file system is protected using OS file and directory permissions. Note that <sid>adm does not have OS access to tenant data volumes, log volumes, or backups, but it can access tenant-specific trace and configuration files.
- Operations that require OS access are restricted to users with the correct permissions. This feature adds another layer of protection between tenants: Tenant administrators with access to the OS cannot access other tenants or the system database with OS commands.

Port Assignment in Tenant Databases

Every tenant database in a multiple-container system has dedicated ports for SQL- and HTTP-based client communication, as well as for internal icommunication. However, there are no standard port number assignments. Port numbers are assigned automatically from the available port number range according to availability at the time of database creation or when a service is added. Administrators can also specify which port numbers to use when they create a tenant database or add a service. The only exception to this practice is the tenant database. This database is created automatically when you convert a single-container system to a multiple-container system. This database retains the port numbers of the original single-container system.

The default port number range for tenant databases is 3<instance>40—3<instance>99. This means that the maximum number of tenant databases that can be created per instance is 20. However, you can increase this figure by reserving the port numbers of further instances. To do this, configure the [multidb] reserved_instance_numbers property in the `global.ini` file. The default value of this property is 0. If you change the value to 1, the port numbers of one further instance are available (for example, 30040—30199, if the first instance is 00). If you change it to 2, the port numbers of two further instances are available (for example, 30040—30299, if the first instance is 00).

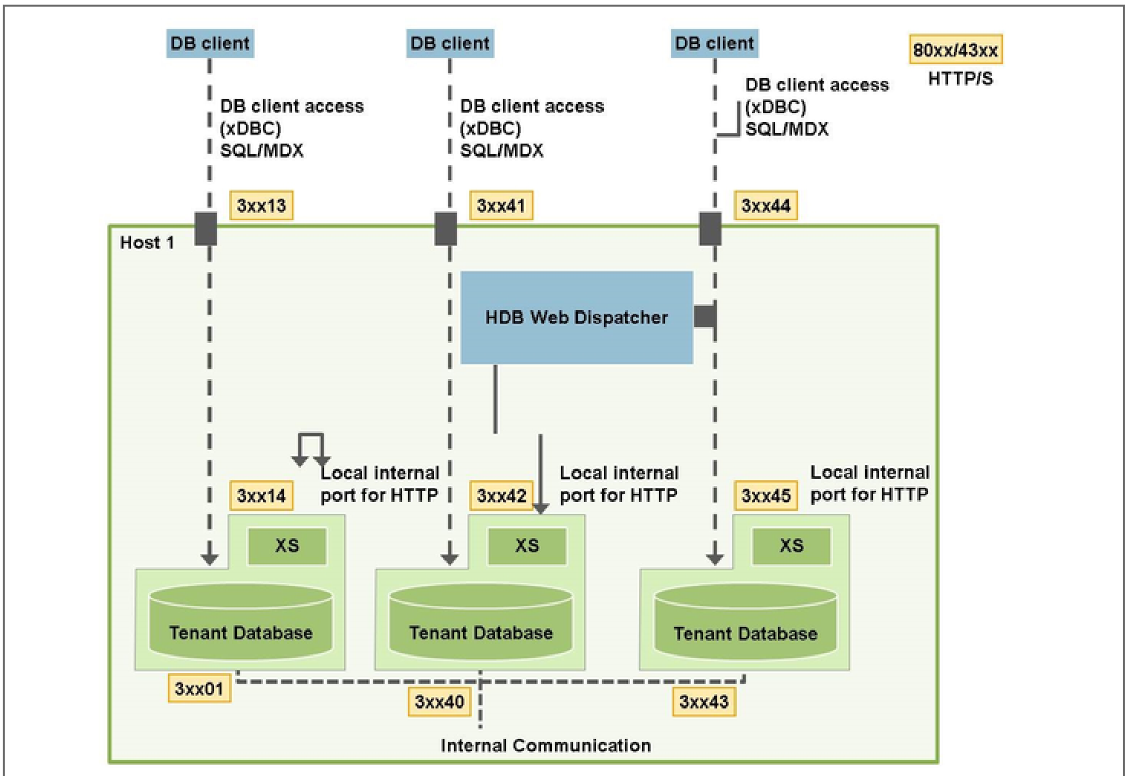

Figure 59: MDC Port Assignment

HTTPS Client Access

The server for SAP HANA extended application services allows Web-based applications to access SAP HANA via HTTPs. The internal Web Dispatcher of the SAP HANA system manages these incoming HTTPs requests. To allow applications to send requests to specific databases in a multiple-container system, every tenant database needs an alias hostname. Requests to the alias hostname are then forwarded to the server of the corresponding tenant database. Requests with the physical hostname in the HTTP host header are forwarded to the server running on the system database.

The default HTTP ports are used in all cases, that is, 80<instance> (HTTP) and 43<instance> (HTTPs). Alias hostnames are mapped to internal HTTPs ports so that incoming requests can be routed to the correct database.

To configure the internal SAP Web Dispatcher, specify the URLs by which tenant databases are publicly accessible in the `xsengine.ini` file of each individual tenant database. You don't need to specify the URL of the system database, because this is done automatically.

Scale-Out-Scenario

A system with multitenant database containers can be distributed across several hosts. To ensure availability, an instance of the system database runs on all hosts (worker and standby) in a configuration of a single master and multiple workers. Tenant databases can be created on worker hosts, and existing databases can be scaled out through the addition of services. If a host fails, the standby instance will fail over all active databases and their services.

The figure, MDC: Scale-Out-Scenario, shows a scale-out-scenario for a multiple-container system with three tenant databases distributed across four hosts (three worker and one standby). If host 2 goes down, the standby host becomes active. The tenant databases normally running on host 2 then become active on the standby host.

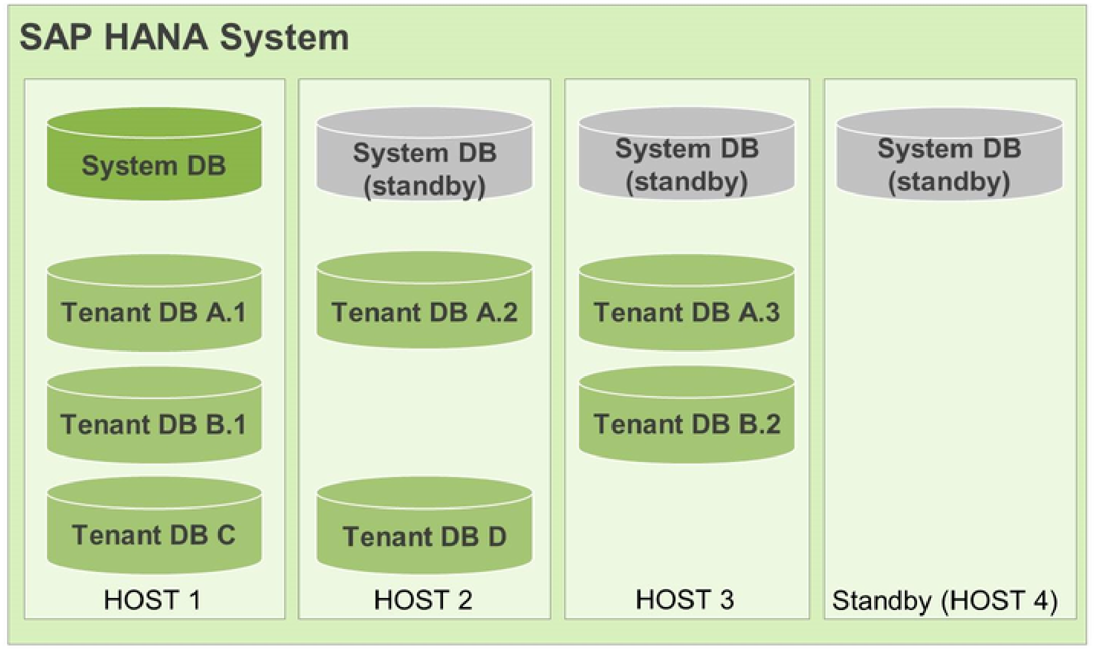

Figure 60: MDC: Scale-Out-Scenario

Migration of a Single Database to a Multitenant Database System

You can migrate an SAP HANA single database system to a multitenant database system. This step is irrevocable.

When you migrate a single database to a multitenant database system, the following occurs:

- System database is generated
- Single database is converted into a tenant database automatically
- No changes to application or customer data
- Migration must be explicitly triggered

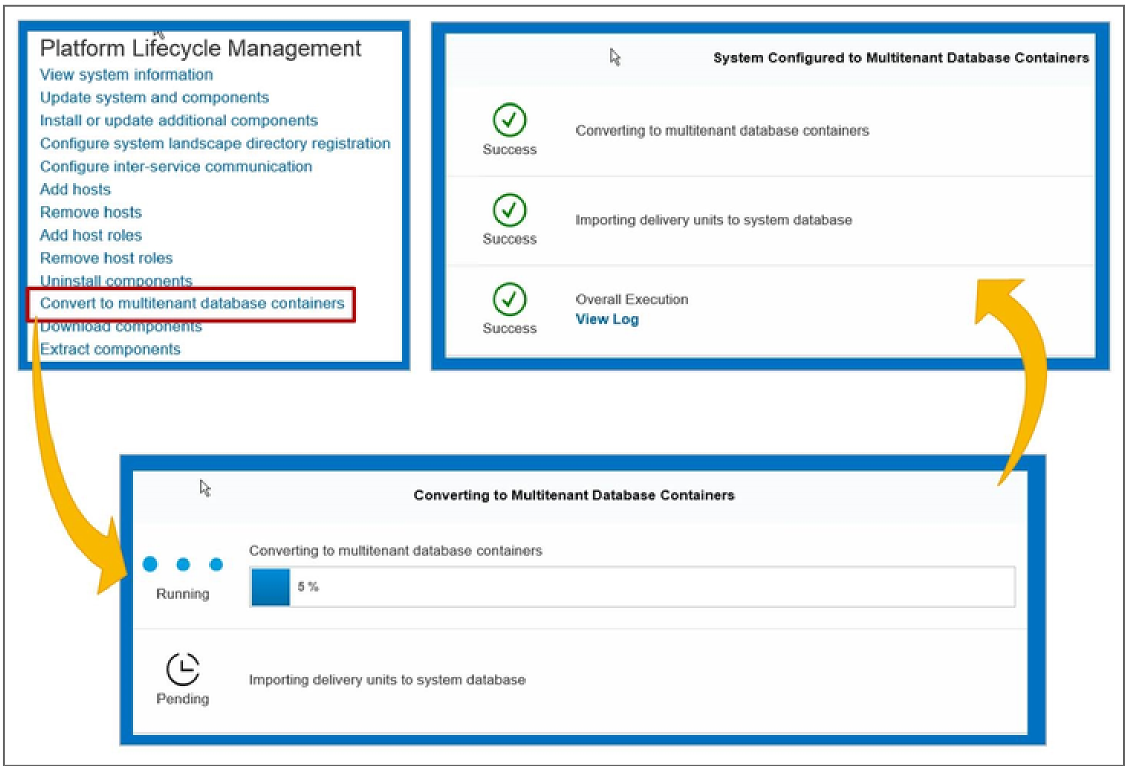

Figure 61: Migration of a Single Database to a Multitenant Database System

LESSON SUMMARY

You should now be able to:

- Explain the SAP HANA Architecture and Technology

Unit 4
Lesson 2

SAP HANA Memory Management and Data Persistence

LESSON OVERVIEW

Although SAP HANA is often referred to as an in-memory database management system, data is not only kept in the RAM. Rather, it is also durably maintained in data and log volumes. This lesson describes the memory management and persistence.

Business Example

For monitoring purposes, you want to understand the SAP HANA memory usage and allocation behavior in detail, and learn about optimization potential.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explaining SAP HANA Memory Management and Data Persistence

SAP HANA Server Components

An SAP HANA system contains all the server components for an installation of SAP HANA. An SAP HANA system consists of several servers, the most important of which is the index server. The index server contains the actual data stores and the engines for processing the data.

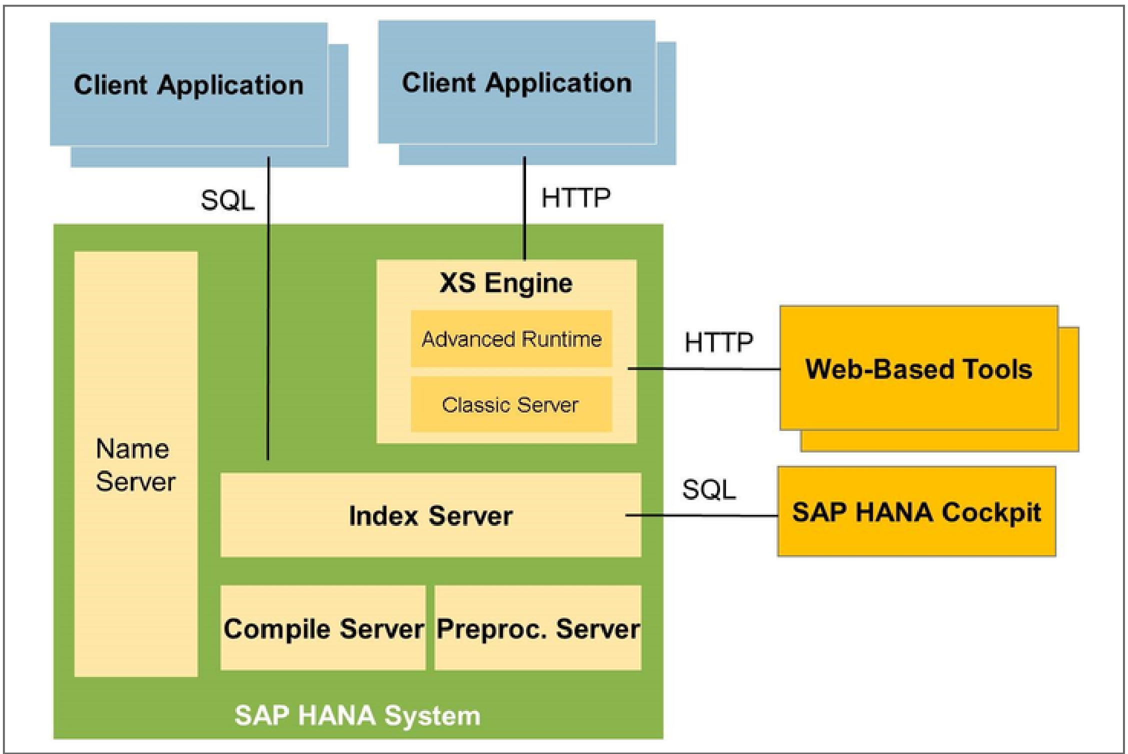

Figure 62: Main Components of the SAP HANA System

Core Processes on an SAP HANA Single-Node Instance

The SAP HANA database functions are implemented in different services. These are briefly described in the figure, Core Processes on an SAP HANA Single-Node Instance . Following the concept of a shared architecture, each of the processes maintains data in the corresponding data and log volumes independently.

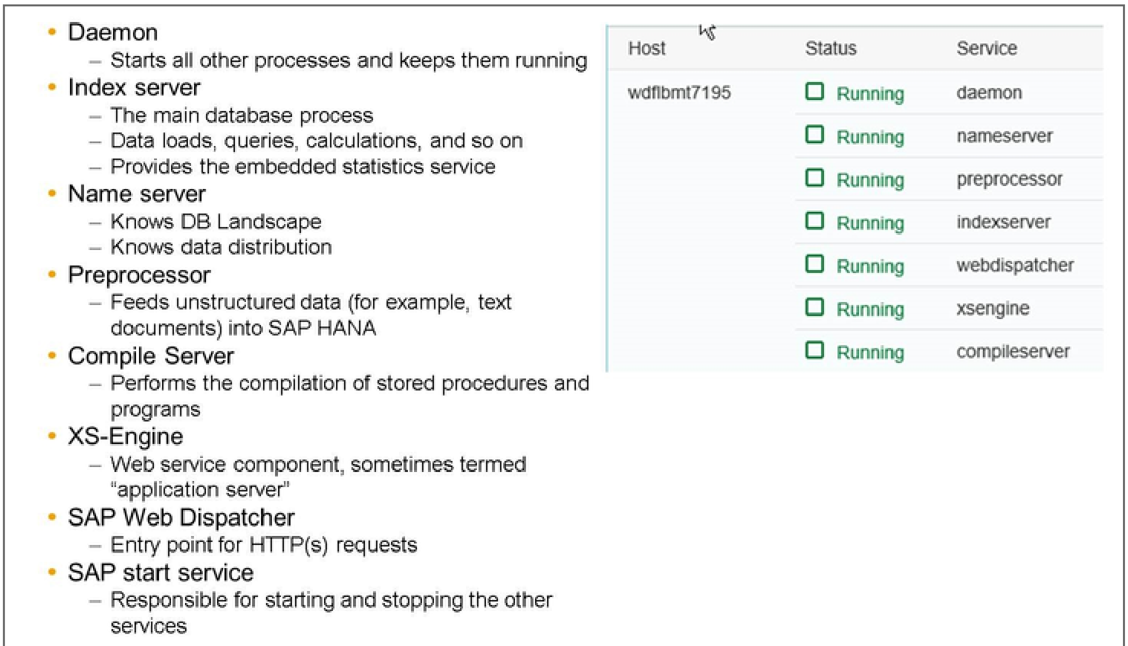

Figure 63: Core Processes on an SAP HANA Single-Node Instance

Hint:
Note that some of the services are optional. For example, the xsengine service can be deactivated and removed if not required. For details, see SAP Note 1867324 .

The SAP HANA Extended Application Services Engine

SAP HANA extended application services is the application server for native SAP HANA-based web applications. It is installed with the SAP HANA system and allows developers to write and run SAP HANA-based applications without the need to run an additional application server. SAP HANA extended application services is also used to run web-based tools that come with SAP HANA, for instance for administration, lifecycle management, and development.

Server on the Classic Model of SAP HANA Extended Application Services

The classic model of SAP HANA extended application services is the original implementation of SAP HANA extended application services. The server on the classic model of SAP HANA extended application services can run as a separate server process or embedded within the index server.

Runtime for the Advanced Model of SAP HANA Extended Application Services

Since SAP HANA 1.0 SPS 11 , SAP HANA includes an additional runtime environment for application development: SAP HANA extended application services, advanced model. The advanced model of SAP HANA extended application services represents an evolution of the application server architecture within SAP HANA. It builds on the strengths and expands the scope of the classic model of SAP HANA extended application services.

The runtime for the advanced model of SAP HANA extended application services consists of several processes for platform services and for executing applications. For more information about the individual services, see the SAP HANA Administration Guide .

The runtime for the advanced model of SAP HANA extended application services runs either on dedicated hosts or together with other SAP HANA components on the same host.

Note:
SAP recommends that customers and partners who want to develop new applications use the advanced model of SAP HANA extended application services. If you want to migrate existing applications from the classic model of SAP HANA extended application services to run in the new advanced runtime environment, first check the features available with the installed version of SAP HANAextended application services, advanced model. If the features of the advanced model of SAP HANA extended application services match the requirements of the classic application that you want to migrate, then you can start the migration process. For more information, see the SAP HANA Extended Application Services Advanced Migration Guide .

Architecture of the SAP HANA Indexserver

Because it keeps the tables in main memory and executes requests, the indexserver process is most relevant for this lesson. It is described in detail in the following section.

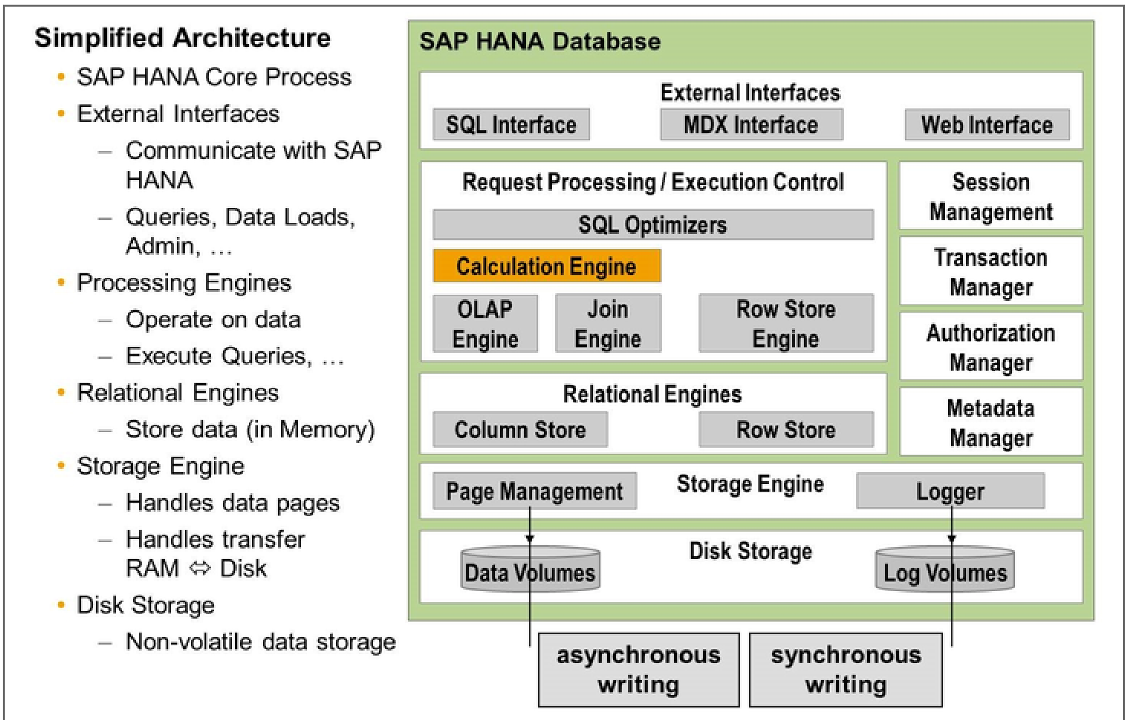

Figure 64: Architecture of the SAP HANA Indexserver

From an architectural standpoint, the SAP HANA Indexserver consists of several components that implement various features, as follows:

- External Interfaces

  Structured Query Language (SQL), Multidimensional Expressions (MDX), and Web interfaces allow clients to connect and communicate with the SAP HANA database.

- Request Processing and Execution Control

  Depending on the interface and the statement, different components for processing can be used. For example, SQL script implementations are executed within the Calculation Engine.

- Relational Engines

  The table data in SAP HANA is kept in two different relational stores: Row Store and Column Store. Each of these stores shows significant differences with regard to the main memory management.

- Storage Engine and Disk Storage

  To gain consistency and maintain durable changes, a Storage Engine with Page Management and Logger is used. This ensures that the database can be restored to the most recent committed state after a restart. It also ensures that transactions are either completely executed or completely undone. Disk Storage is divided in Data Volumes and Log Volumes. While changes need to be written to the log area before a successful commit of a transaction (synchronous writing), the data area contains the complete main memory content at a specific point in time and is written asynchronously.

Persistence

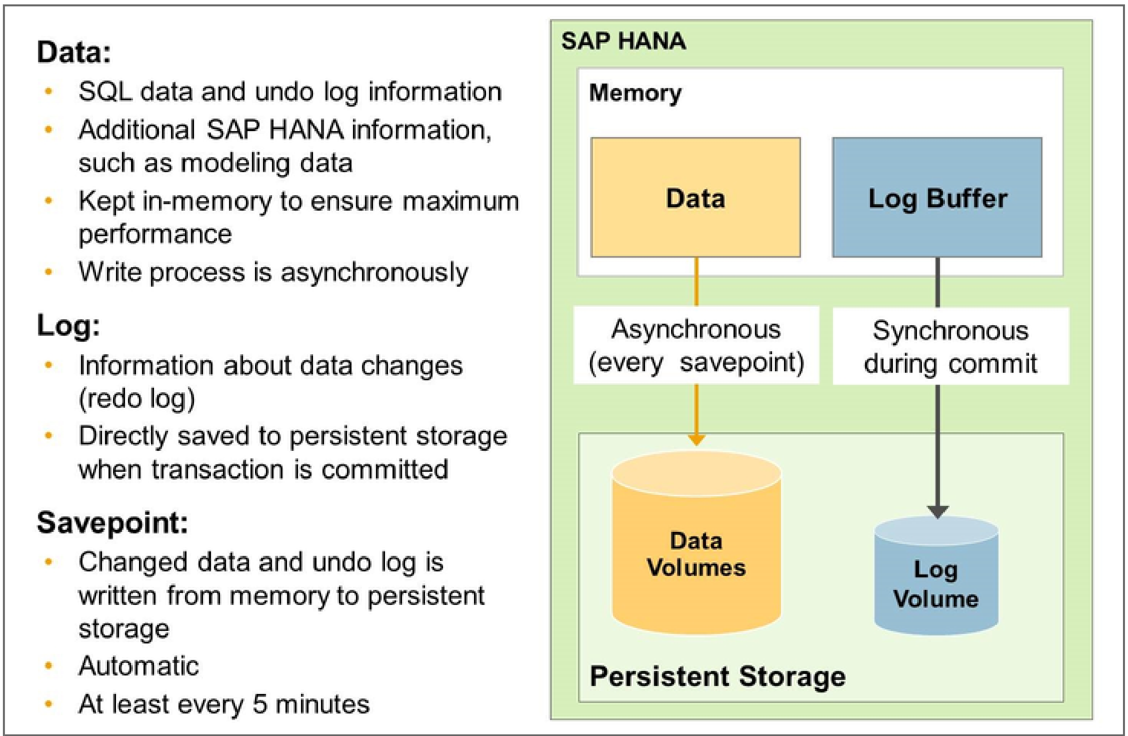

Figure 65: In-Memory Data Is Regularly Saved to Disk

Disk storage is still required to allow restarts in case of power failure and for permanent persistency. The SAP HANA persistency layer stores data in persistent disk volumes that are organized in pages. It is divided in both the log and data area, as follows:

- Data changes such as insert, delete, and update are saved on disk immediately in the logs (synchronously). This is required to make a transaction durable. It is not necessary to keep the entire data, but the transaction log can be used to replay changes after a crash or database restart.
- In customizable intervals (standard: every five minutes), a new savepoint is created. That is, all of the pages that were changed are refreshed in the data area of the persistence.

Whether or not disk access can become to a performance bottleneck depends on the usage. Because changes are written to the Data Volumes asynchronously, the user or application does not need to wait for this. When data that already resides in the main memory is read, there is no need to access the persistent storage. However, when applying changes to data, the transaction cannot be successfully committed before the changes are persisted to the log area.

To optimize the performance, fast storage is used for the log area. For example, it uses solid-state drives (SSDs) or Fusion-io drives (see also, certified hardware configurations in the Product Availability Matrix).

Storing Data in Data Volumes: Details

Data Volumes are located in file systems

Per instance one data volume

- Growing until disk or LUN is full
- Logical Volume Manger (LVM) or similar needed on OS level to extend the file systems or dedicated partition/LUN
- Growing with number of data volumes currently not available (unlike MaxDB)
- Different Page Sizes (4k, 16k, ...16M), which are arranged in superblocks of 64M

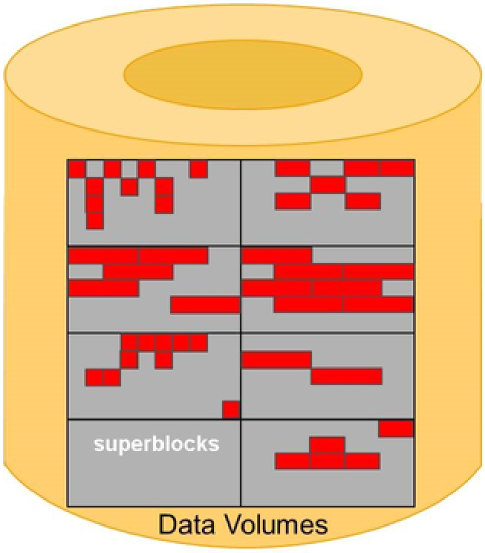

Figure 66: Storing Data in Data Volumes: Details

Like many modern database management system, SAP HANA can use the file abstraction layer of the host operating system.

Each data volume contains one file in which data is organized into pages ranging in size from 4KB to 16MB (page size class). Data is written to and loaded from the data volume by page. Over time, pages are created, changed, overwritten, and deleted. The size of the data file is increased automatically as more space is required. However, it is not decreased automatically when less space is required. This means that, at any given time, the actual payload of a data volume (that is the cumulative size of the pages currently in use) may be less than its total size.

This is not necessarily significant; it simply means that the amount of data in the file is currently less than at some point in the past (for example, after a large data load). If a data volume has a considerable amount of free space, it might be appropriate to shrink the data volume. However, a data file that is excessively large for its typical payload can also indicate a more serious problem with the database. SAP support can help to analyze the situation.

Shadow Paging Concept

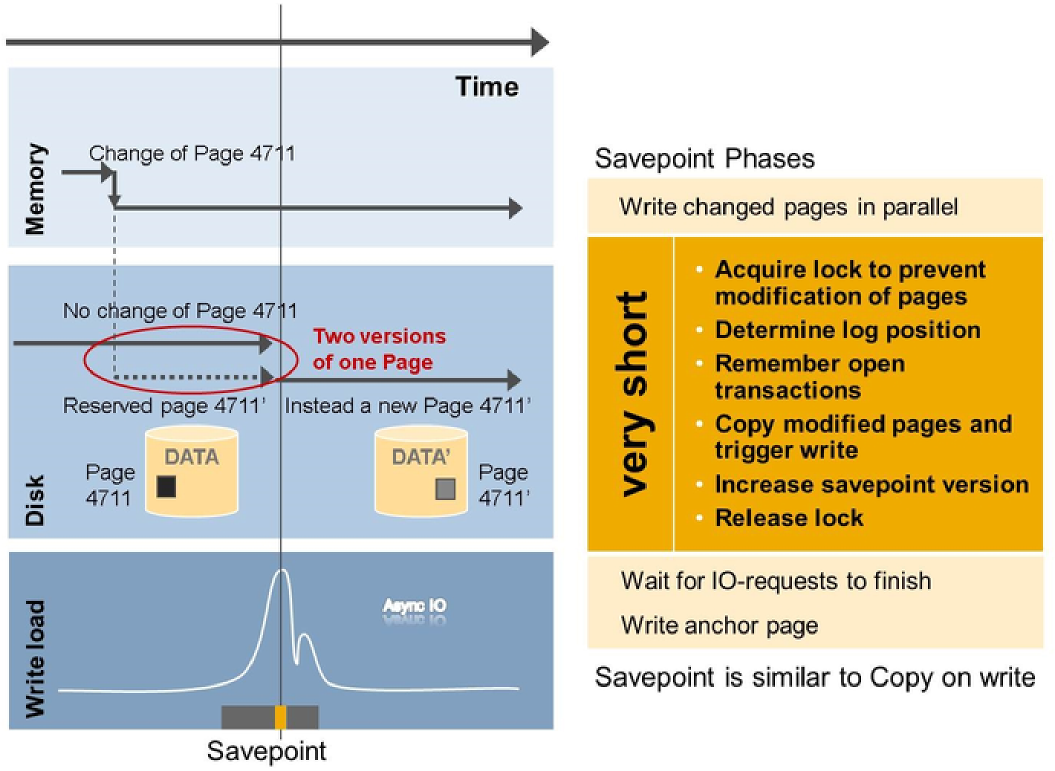

Figure 67: Shadow Paging Concept

While (redo) log entries are written synchronously, changed data in data volumes is periodically copied to disk in a savepoint operation. During the savepoint operation, the SAP HANA database flushes all changed data from memory to the data volumes. The data belonging to a savepoint represents a consistent state of the data on disk. It remains so until the next savepoint operation is completed.

Note:
The frequency for savepoint creation can be configured. This is described in detail later in this course. Savepoints are also triggered automatically by a number of other operations such as data backup, and database shutdown and restart. You can trigger a savepoint manually by executing the statement `ALTER SYSTEM SAVEPOINT`.

The phases of the savepoint operation are shown in the figure, Shadow Paging Concept . SAP HANA uses a "Shadow Paging Concept" . This means that write operations write to new physical pages and the previous savepoint version is still kept in shadow pages. So, if a system crashes during a savepoint operation, it can still be restored from the last savepoint.

Restart Process

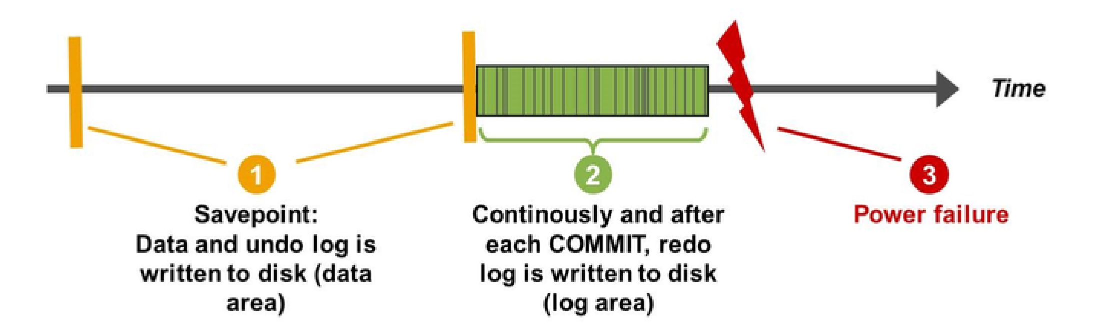

Figure 68: Restart Process

During a database restart (for example after a crash), the data from the last completed savepoint can be read from the data volumes and the redo log entries written to the log volumes since the last savepoint. This allows the database to be restored to the last committed state.

In-Memory Computing Security

- The SAP in-memory database holds the bulk of its data in-memory for maximum performance. However, it still uses persistent storage to provide a fallback in case of failure. The log captures all changes by database transactions (redo logs).
- Data and undo log information (part of data) are automatically saved to disk at regular savepoints.
- The log is also saved to disk continuously and synchronously after each COMMIT of a database transaction (waiting for end of disk write operation).
- After a power failure, the database can be restarted like a disk-based database:
  - The system is normally restarted (“lazy” reloading of tables to keep the restart time short)
  - The system returns to its last consistent state (by replaying the redo log since the last savepoint)

> Note:
> After a system restart, by default, not all tables are loaded into the main memory immediately.

Start of SAP HANA Database

The SAP HANA system restart sequence restores the system to a fully operational state quickly.

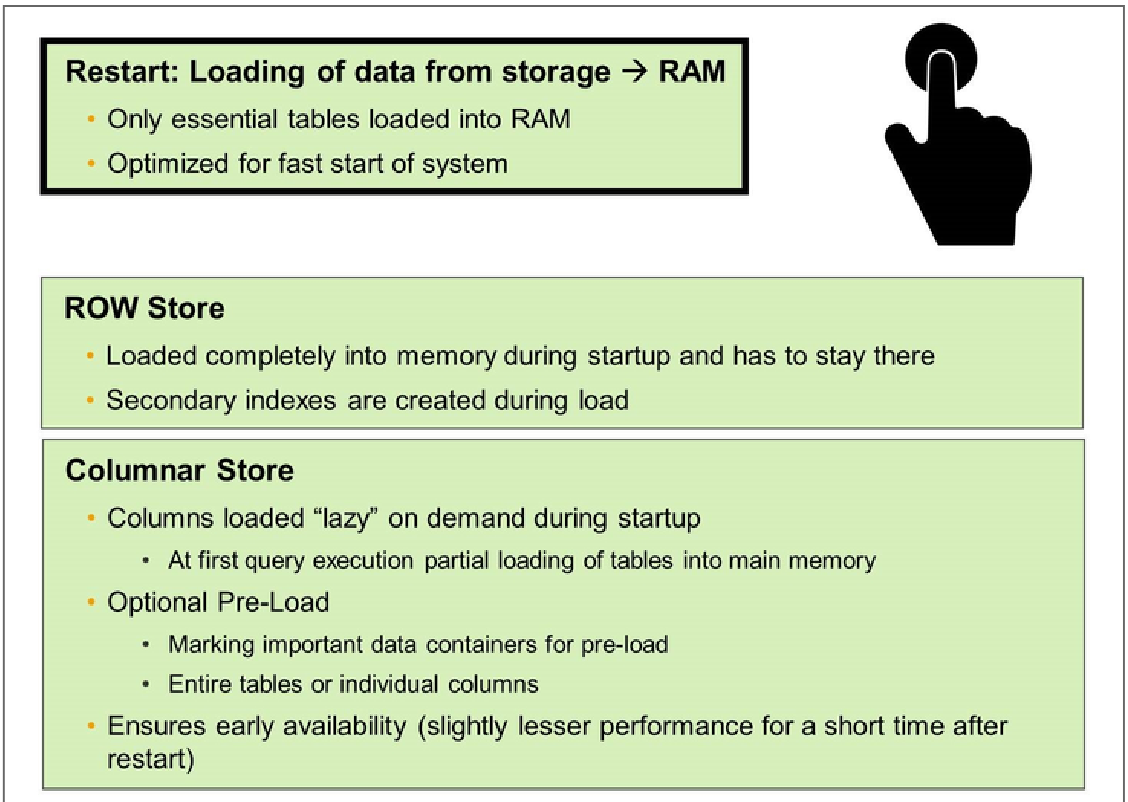

Figure 69: Start of SAP HANA Database

Start-Up Process: Persistence Layer Activities

When you restart an SAP HANA system, the following activities are executed by the restart agent of the persistence layer:

1. The data volume of each service is accessed in order to read and load the restart record.
2. The list of open transactions is read into memory.
3. The row tables are loaded into memory.
4. Open transactions are processed using the redo log, as follows:
   a. Write transactions that were open when the database was stopped are rolled back.
   b. Changes of committed transactions that were not written to the data area are rolled forward. The first column tables are reloaded into memory because they are accessed for rollforward
   c. After this step, the database is technically available and logon is possible.
5. Aborted transactions are determined and rolled back.
6. A savepoint is performed with the restored consistent state of the database.
7. Column tables, which are marked for preload, and their attributes are asynchronously loaded in the background, if they have not already been loaded as part of log replay.

   The preload parameter is configured in the metadata of the table. This feature makes certain tables and columns that were used by important business processes available more quickly.

8. Column tables that were loaded before restart, and their attributes, start reloading asynchronously in the background, if they have not already been loaded as part of log replay or because they are marked for preload.

During normal operation, the system tracks the tables currently in use. This list is used as basis for reloading tables after a restart.

> Note:
> Because a regular or soft shutdown writes a savepoint, there are no replay log entries to be processed.

Start-Up Process

Restart Sequence:

1. The data volume of each service is accessed.
2. The list of open transactions is read into memory.
3. Row tables are loaded into memory.
4. Open transactions are recovered.
5. Aborted transactions are determined and rolled back.
6. A save point is performed with the restored consistent state of the database.
7. Column tables that are marked for preload are asynchronously loaded.
8. Column tables that were loaded before restart start reloading asynchronously

Important factors for startup

- Remaining Log to be rolled forward
- I/O performance of data and log disks
- Separate log, data, and backup disk areas not only logically, but also physically for the best performance

Figure 70: Start-Up Process

While the row store is always loaded entirely, only those columns of essential column tables are loaded into memory. The other columns are loaded, if requested.

For example, if a query only uses some of the fields (columns) of a table, only those fields are loaded into the memory at the time of query execution. All row-based tables (usually system tables) are available in the main memory. Their size significantly influences the time required to start the database. Other factors that influence startup time are mentioned in the figure, Start-Up Process .

Start-Up Process: Tables

During the normal operation, SAP HANA tracks a list of column tables that are currently loaded (once per day). This list is now the basis of loading the necessary tables into main memory during restart. Reloading column tables in this way restores the database to a fully operational state more quickly. However, it does create performance overhead, and may not be necessary in nonproductive systems. You can deactivate the reload feature in the `indexserver.ini` file by setting the reload_tables parameter in the sql section to false . In addition, you can configure the number of tables whose attributes are loaded in parallel using the tables_preloaded_in_parallel parameter in the parallel section of `indexserver.ini`. This parameter also determines the number of tables that are preloaded in parallel.

> Note:
> You can mark individual columns as well as entire column tables for preload.

When the Preload checkbox is selected, tables are loaded into memory automatically after an index server start. The current status of the Preload checkbox is visible in the system table TABLES in the PRELOAD column. Possible values are **`FULL`**, **`PARTIALLY`**, and **`NO`**. Also, in the system table TABLE_COLUMNS in the PRELOAD column, the possible values are **`TRUE`** or **`FALSE`**.

> Note:
> When fields of large column tables are not in the main memory, the first access to the table might be significantly slower. This is because all requested columns are loaded to the main memory before the query can be executed. This applies even if a single record is selected.

> Caution:
> Simply selecting all tables for preload in order to accelerate initial queries can slow down startup time considerably. The Preload checkbox is a tuning option and should be used carefully, depending on the individual scenario and requirements.

Memory Usage

The total amount of memory used by SAP HANA is referred to as used memory. It includes program code and stack, all data and system tables, and the memory required for temporary computations. In the Linux operating environment, memory is allocated for the program code (sometimes called the text), the program stack, and data. Most of the data memory, called the heap, is under program control.

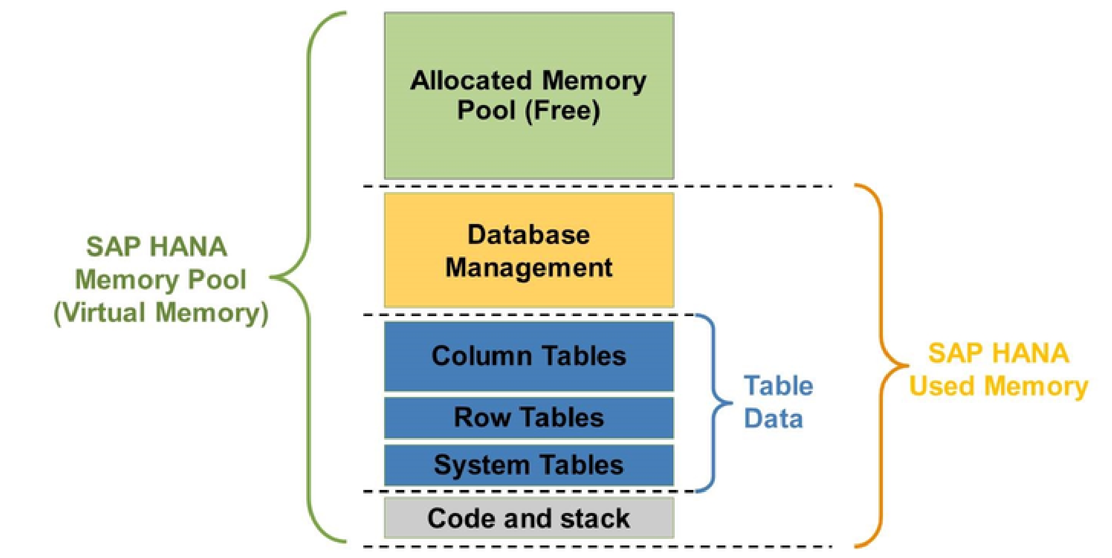

Figure 71: SAP HANA Memory Pool

As an in-memory database, it is crucial for SAP HANA to manage and track its own consumption of memory carefully. For this purpose, the SAP HANA database preallocates and manages its own data memory pool. The memory pool is used for storing in-memory tables, thread stacks, temporary computations, intermediate results, and other data structures. The use of memory by SAP HANA, therefore, includes its program code (exclusive and shared), the program stack, and the memory pool this includes all of the data tables (row and column), system tables, and created tables.

Parts of the pool are always in use for temporary computations. The total amount of memory in use is referred to as used memory. This is the most precise indicator of the amount of memory that the SAP HANA database uses.

Virtual, Physical, and Resident Memory

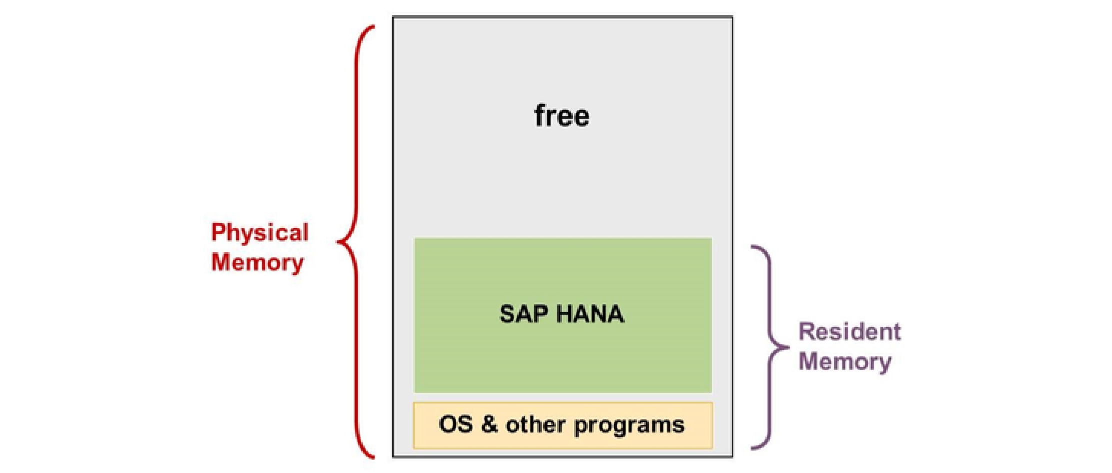

Figure 72: Virtual, Physical, and Resident Memory

When part of the virtually allocated memory needs to be used, it is loaded or mapped to the real, physical memory of the host and becomes "resident". Physical memory is the dynamic random-access memory (DRAM) installed on the host. On SAP HANA hosts, it ranges from 128 GB to 4 TB. It runs the Linux operating system, SAP HANA, and all other programs. Resident memory is the physical memory used by a process.

Memory Consumption

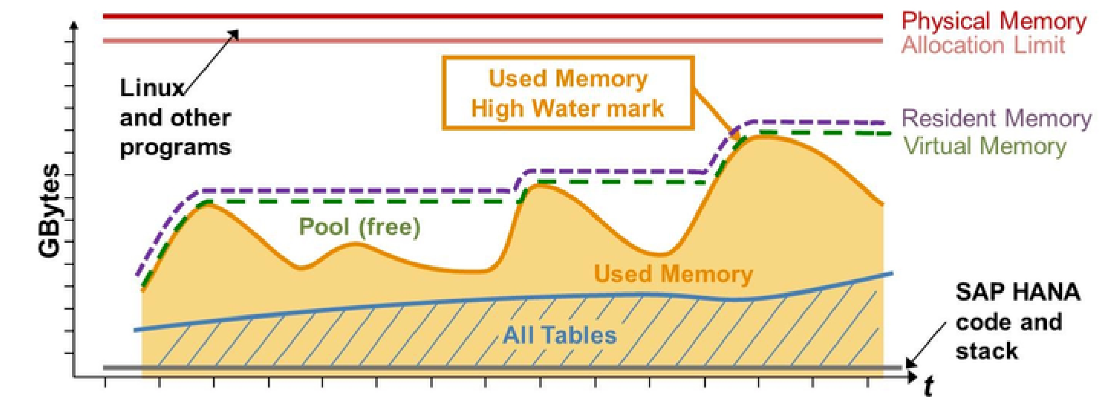

Figure 73: Memory Consumption

The SAP HANA database, across its different processes, reserves a pool of memory before actual use. This pool of allocated memory is preallocated from the operating system over time, up to a predefined global allocation limit. It is then efficiently used as needed by the SAP HANA database code.

When memory is required for table growth or for temporary computations, the SAP HANA code obtains it from the existing memory pool. When the pool cannot satisfy the request, the SAP HANA memory manager requests and reserves more memory from the operating system. At this point, the virtual memory size of the SAP HANA processes grows. Once a temporary computation completes or a table is dropped, the freed memory is returned to the memory manager. The manager recycles the memory to its pool, without informing Linux. Thus, from the perspective of SAP HANA, the amount of used memory shrinks. However, the virtual and resident sizes of the process are not affected. The used memory, therefore, can shrink to below the size of SAP HANA's resident memory, which is perfectly normal.

Note:
The database can also actively unload tables or individual columns from memory. For example, if a query or other processes in the database require more memory than is currently available. It does this based on a least recently used algorithm.

Caution:
Because of the preallocation of memory, Linux memory indicators such as top and meminfo do not accurately reflect the actual SAP HANA used memory size. Always base main memory monitoring on SAP HANA monitoring features.

Memory Management in the Column Store

The column store is the part of the SAP HANA database that manages data organized in columns in-memory. Tables created as column tables are stored here.

The column store is optimized for read operations, but it also provides good performance for write operations. This is achieved through two data structures: main storage and delta storage.

The main storage contains the main part of the data. Here, efficient data compression is applied to save memory and to speed up searches and calculations. Write operations on compressed data in the main storage are costly, however. Therefore, write operations do not directly modify compressed data in the main storage. Instead, all changes are written to a separate data structure called the delta storage. The delta storage uses only basic compression and is optimized for write access. Read operations are performed on both structures, while write operations only affect the delta.

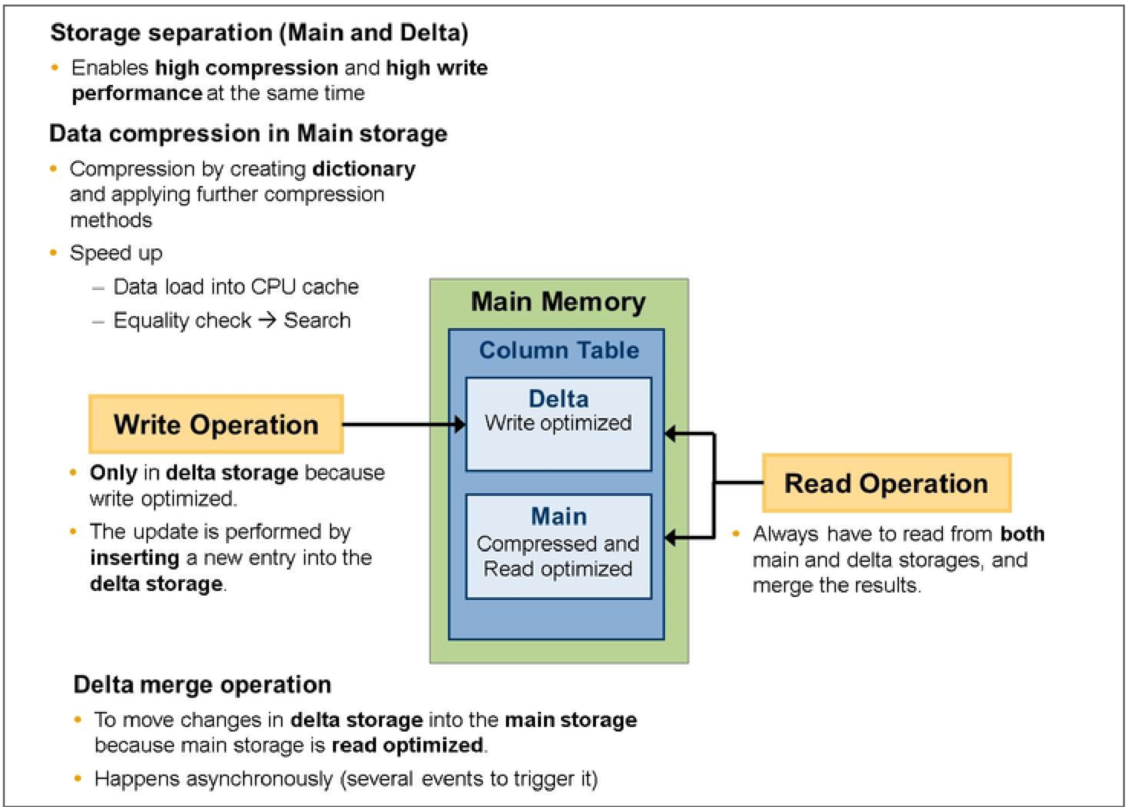

Figure 74: Column Store Memory Management

Delta Merge Operation

The purpose of the delta merge operation is to move changes collected in the delta storage to the read-optimized main storage. After the delta merge operation, the content of the main storage is persisted to disk. Its compression is recalculated and optimized, if necessary.

A further result of the delta merge operation is the truncation of the delta log. The delta storage structure itself exists only in-memory and is not persisted to disk. The column store creates its logical redo log entries for all operations executed on the delta storage. This log is called the delta log. If a system restart occurs, the delta log entries are replayed to rebuild the in-memory delta storages. After the changes in the delta storage are merged into the main storage, the delta log file is truncated by removing those entries that were written before the merge operation.

> Note:
> Data that is in the delta storage can only be fully loaded or unloaded. Partial loading is not possible. Therefore, if a delta merge has not been performed and the table's entire data is in the delta storage, the table is either fully loaded or unloaded.

Delta Merge Process

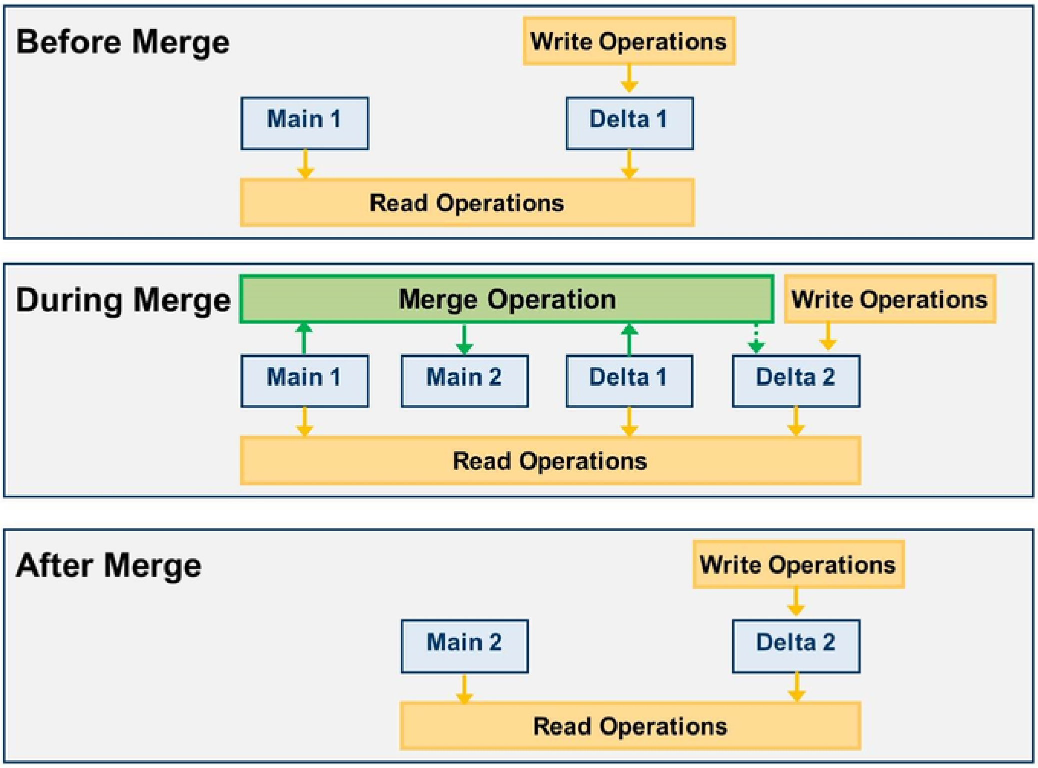

Figure 75: The Delta Merge Process

The following steps are performed in the merge process:

1. Before the merge operation, all write operations go to Delta 1 storage, and all read operations read from Main 1 and Delta 1 storages.
2. While the merge operation is running, the following events occur:
   a. All write operations go to the second delta storage, Delta 2.
   b. Read operations read from the original main storage, Main 1, and from both delta storages, Delta 1 and Delta 2.
   c. Uncommitted changes in Delta1 are copied to Delta2.
   d. The content of Main 1 and the committed entries in Delta 1 are merged into the new main storage, Main 2.
3. After the merge operation is completed, the following events occur:

a. Main1 and Delta1 storages are deleted.

b. The compression of the new main storage (Main 2) is reevaluated and optimized. If necessary, this operation reorders rows and adjust compression parameters. If compression has changed, columns are immediately reloaded into memory.

c. The content of the complete main storage is persisted to disk.

> Note:
> With this double buffer concept, the table only needs to be locked for a short time: at the beginning of the process when open transactions are moved to Delta2, and at the end of the process when the storages are "switched".

> Caution:
> The minimum memory requirement for the delta merge operation includes the current size of the main storage, the future size of main storage, the current size of delta storage, plus some additional memory. Even if a column store table is unloaded or partly loaded, the whole table is loaded into memory to perform the delta merge.

Reason for Expense of Delta Merge Operation

The delta merge operation can be expensive for the following reasons:

- The complete main storage of all columns of the table is rewritten in-memory. This uses some central processing unit resources, and temporarily duplicates the memory needed for the main storages (while Main 1 and Main 2 exist in parallel).
- The complete main storages are persisted to disk, even if only a relatively small number of records were changed. This creates disk input/output load.

This potentially negative impact on performance can be lessened by the following strategies:

- Executing memory-only merges

  A memory-only merge affects only the in-memory structures and does not persist any data.

- Splitting tables

  The performance of the delta merge depends on the size of the main storage. This size can be reduced by splitting the table into multiple partitions, each with its own main and delta storages. The delta merge operation is performed at partition level and only for partitions that actually require it. This means that less data is merged and persisted. Note that there are disadvantages to partitioning tables that should also be considered.

Merge Motivations

The request to merge the delta storage of a table into its main storage can be triggered in several ways. These are called merge motivations.

The following figure illustrates the different merge motivations and how they are triggered.

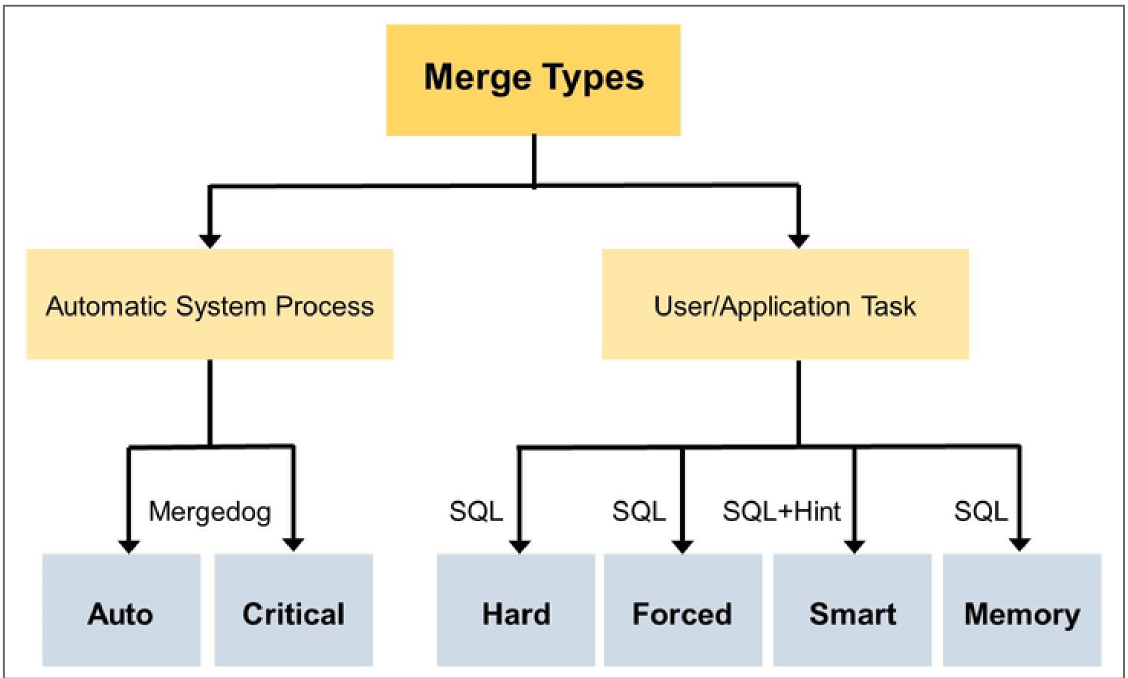

Figure 76: Merge Motivations

Auto Merge

The standard method for initiating a merge in SAP HANA is the auto merge. A system process called mergedog periodically checks the column store tables that are loaded locally. For each individual table or single partition of a split table, it determines whether or not a merge is necessary based on certain criteria (for example, size of delta storage, available memory, time since last merge, and others).

If the active parameter in the mergedog section of the `indexserver.ini` file is set to yes, auto merge is active.

> Note:
> You can activate or deactivate auto merge for an individual table in the TABLES (SYS) system view. Change the value in the AUTO_MERGE_ON column to **`TRUE`** or **`FALSE`**.

Smart Merge

If an application powered by SAP HANA requires more direct control over the merge process, SAP HANA provides a function that allows the system to check whether or not a delta merge makes sense. This function is called smart merge. For example, if an application starts loading relatively large data volumes, a delta merge during the load can have a negative impact both on the load performance and on other system users. Therefore, the application can disable the auto merge for those tables being loaded and send a hint to the database to merge once the load has completed.

When the application issues a smart merge hint to the database to trigger a merge, the database evaluates the criteria that determine whether or not a merge is necessary. If the criteria are met, the merge is executed. If the criteria are not met, the database takes no further action. Only a subsequent hint from the application triggers another evaluation of the criteria.

If the active parameter in the smart_merge_enabled section of the `indexserver.ini` file is set to yes , smart merge is active .

> Caution:
> For tables that you want to merge with the smart merge, disable the auto merge. Otherwise, the auto merge and smart merge may interfere with each other.

Hard and Forced Merges

You can trigger the delta merge operation for a table manually by executing the SQL statement `MERGE DELTA OF "<table_name>"`. This is called a hard merge. It causes the database to execute the delta merge for the table either immediately if sufficient system resources are available, or as soon as sufficient system resources become available. The hard merge is therefore subject to the merge token control mechanism.

If you want the merge to take place immediately regardless of system resource availability, you can pass an optional parameter. A forced merge may be useful in a situation where there is a heavy system load, but a small table needs to be merged. It is also useful if a missed merge of a certain table negatively impacts system performance. To execute a forced merge, execute the SQL statement `MERGE DELTA OF '<table_name>' WITH PARAMETERS ('FORCED_MERGE' = 'ON')`.

> Note:
> Unlike system-triggered delta merge operations, all of the manually-executed delta merge operations listed here do not later trigger an optimization of the compression of the table's new main storage. If the table was compressed before the delta merge operation, it remains compressed with the same compression strategy afterward. If it was not compressed before the delta merge operation, it remains uncompressed afterward. After a manual delta merge, you must therefore trigger compression optimization manually.

Critical Merge

To keep the system stable, the database can trigger a critical merge. For example, when auto merge is disabled and no smart merge hints are sent to the system, the size of the delta storage can grow too large for a successful delta merge to occur. The system initiates a critical merge automatically when a certain threshold is passed.

Parameters

Table 4: Threshold for Optimization Compression

The thresholds for optimization compression to occur are defined as parameter, as shown in the following table:

Parameter	Default	Description
Active	Yes	Compression optimization status
min_change_ratio	1.75	Minimum required change row count (ratio)
min_hours_since_last_merge_of_part	24	Minimum hours since the last merge of part

Parameter	Default	Description
min_rows	10240	Minimum required rows (which stored in the table)

Write operations on this compressed data are costly, as they require reorganizing the storage structure. Therefore, write operations in column store do not directly modify compressed data. All changes go into a separate area called the delta storage. The delta storage exists only in main memory. Only delta log entries are written to the persistence layer when delta entries are inserted.

Features of Delta Merge Operations

The features of the Delta merge operation are as follows:

- The delta merge operation is executed on table level.
- It moves changes collected in write-optimized delta storage into the compressed and read-optimized main storage.
- Read operations always read from both main storage and delta storage, and merge the results.
- The delta merge operation is decoupled from the execution of the transaction that performs the changes. It happens asynchronously at a later point in time.

Note:
For the delta merge operation, a double buffer concept is used. This has the advantage that the table only needs to be locked for a short time. For more information, see the Administration Guide.

Caution:
The minimum memory requirement for the delta merge operation includes the current size of main storage, the future size of main storage, the current size of delta storage, plus some additional memory. Even if a column store table is unloaded or partly loaded, the whole table is loaded into memory to perform the delta merge.

LESSON SUMMARY
You should now be able to:

- Explaining SAP HANA Memory Management and Data Persistence

Unit 4

Learning Assessment

1. From an architectural standpoint, which of the following components belong to the SAP HANA Indexserver ?

   Choose the correct answers.

   - [ ] **A** External Interfaces
   - [ ] **B** Preprocessor
   - [ ] **C** Relational Engines
   - [ ] **D** SAP Web Dispatcher

2. Which of the following events occur while the merge operation is running?

   Choose the correct answers.

   - [ ] **A** Uncommitted changes in the first delta storage are rolled back.
   - [ ] **B** The content of the complete main storage is persisted to disk.
   - [ ] **C** The content of the first main storage and the committed entries in the first delta storage are merged into the new second main storage.
   - [ ] **D** All write operations go to the second delta storage.

Unit 4

Learning Assessment - Answers

1. From an architectural standpoint, which of the following components belong to the SAP HANA Indexserver ?

   Choose the correct answers.

   [X] **A** External Interfaces

   [ ] **B** Preprocessor

   [X] **C** Relational Engines

   [ ] **D** SAP Web Dispatcher

   Correct! External interfaces, for example, Structured Query Language (SQL), Multidimensional Expressions (MDX), and Web interfaces allow clients to connect and communicate with the SAP HANA database. The table data in SAP HANA is kept in two different relational stores: Row Store and Column Store. Each of these stores shows significant differences with regard to the main memory management. The Preprocessor Server is a component of the entire SAP HANA system. It feeds unstructured data (for example text-documents) into SAP HANA. The SAP Web Dispatcher is the entry point for HTTP(s) request to SAP HANA. Read more on this in the lesson SAP HANA Architecture (Unit 4, Lesson 1) of the course HA200_14.

2. Which of the following events occur while the merge operation is running?

   Choose the correct answers.

   [ ] **A** Uncommitted changes in the first delta storage are rolled back.

   [ ] **B** The content of the complete main storage is persisted to disk.

   [X] **C** The content of the first main storage and the committed entries in the first delta storage are merged into the new second main storage.

   [X] **D** All write operations go to the second delta storage.

   Correct! The content of the first main storage and the committed entries in the first delta storage are merged into the new second main storage at the end of the merge operation. Read operations read from the original main storage and both delta-storages. The content of the complete main storage is persisted to disk after the merge operation is completed. Read more on this in the lesson SAP HANA Memory Management and Delta Persistence (Unit 4, Lesson 2) of the course HA200_14.

UNIT 5 SAP HANA Cockpit 2.0

UNIT OBJECTIVES

- Understand the SAP HANA Cockpit 2.0 architecture
- Install SAP HANA Cockpit 2.0
- Configure SAP HANA Cockpit 2.0
- Understand the SAP HANA Cockpit 2.0 revision strategy
- Update SAP HANA Cockpit 2.0 to a newer Support Package

Unit 5
Lesson 1

Introducing SAP HANA Cockpit 2.0

LESSON OBJECTIVES
After completing this lesson, you will be able to:

- Understand the SAP HANA Cockpit 2.0 architecture

Introduction to SAP HANA Cockpit 2.0

SAP HANA cockpit 2.0 provides a single point of access to a range of tools, which are used for the administration and detailed monitoring of multiple and individual SAP HANA databases. It also integrates the Structured Query Language (SQL) development features required by administrators. SAP HANA cockpit has a Web-based HTML5 user interface that you access through a browser. It runs on SAP HANA extended application services, advanced model. Use the cockpit to monitor and manage systems running SAP HANA 2.0 or SAP HANA 1.0 SPS 12.

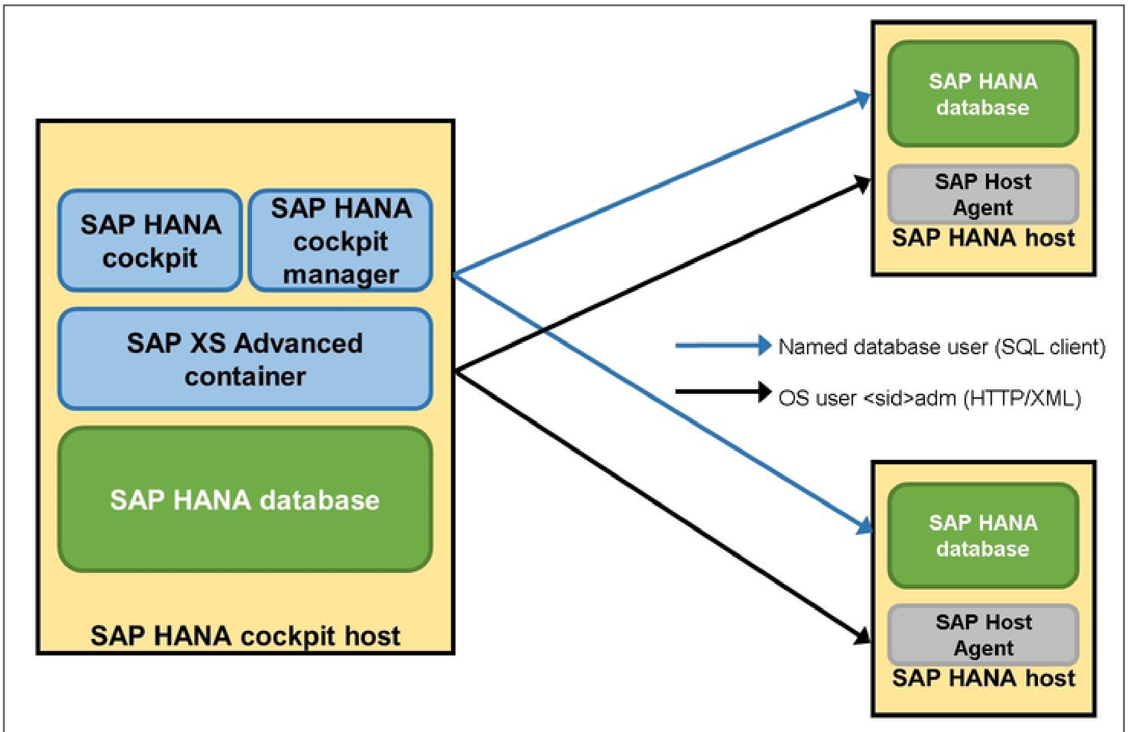

Figure 77: The SAP HANA Cockpit Architecture

Note:
While the cockpit was an integral part of earlier versions of SAP HANA, the new SAP HANA cockpit for SAP HANA 2.0 is installed separately on dedicated hardware. This provides more flexibility, because it allows you to manage more than one SAP HANA system in a single administration environment.

SAP HANA cockpit consists of two parts:

- The SAP HANA cockpit manager
- The SAP HANA cockpit

SAP HANA Cockpit Manager

With the SAP HANA cockpit manager, you can register resources and create groups of resources that other cockpit users can access within the SAP HANA cockpit. The resources are managed through the cockpit administrator user. This user is created during the installation of SAP HANA cockpit and is assigned the master password provided during the installation.

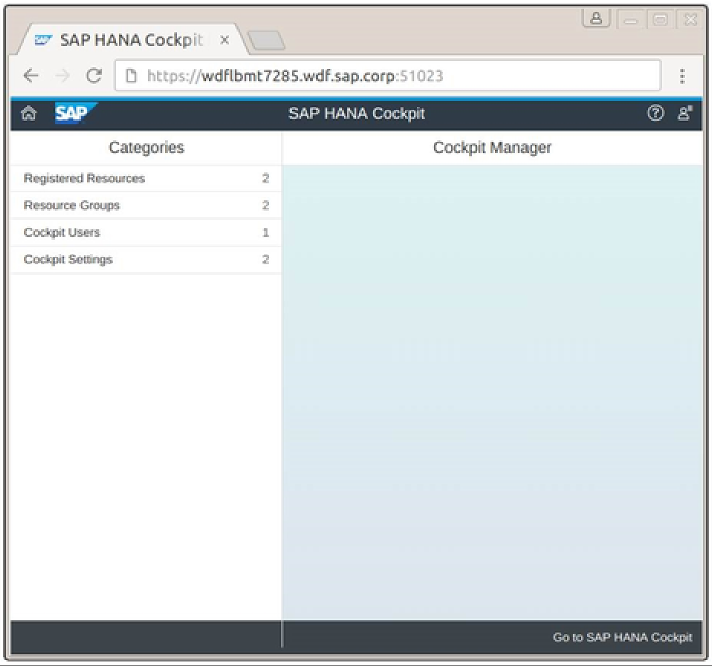

Figure 78: SAP HANA Cockpit Manager: Main Screen

A resource is an SAP HANA system, identified by a host and instance number. Suppose that a business unit has set up a new SAP HANA system and wants it to be managed through the cockpit. The first step is to register the SAP HANA system, or resource, in the cockpit.

SAP HANA Cockpit

The SAP HANA cockpit provides aggregate, system, and database administration features. For example, it provides database monitoring, user management, and data backup. Administrators can use the SAP HANA cockpit to start and stop services, to monitor the system, to configure system settings, and to manage users and authorizations. Cockpit pages that allow you to manage SAP HANA options and features (for example, SAP HANA dynamic tiering) are only available if the option or feature has been installed.

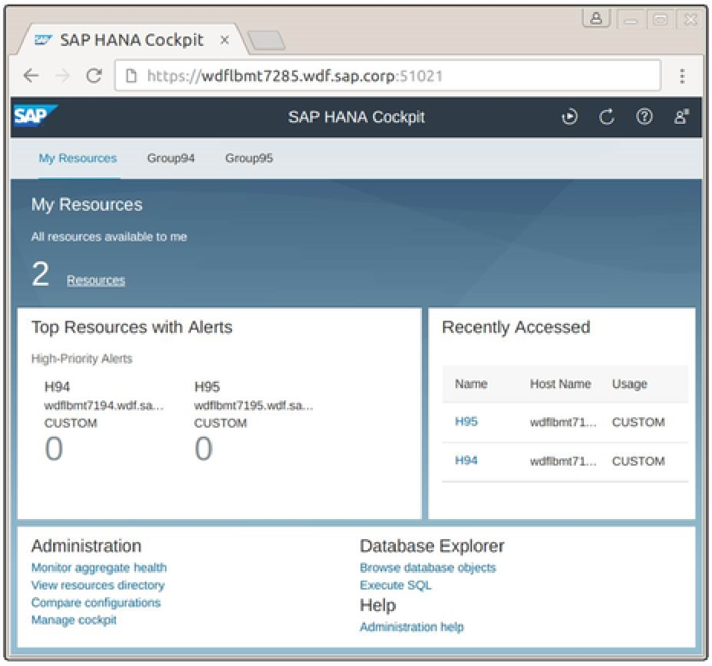

Figure 79: SAP HANA Cockpit: Main Screen

Initially, the SAP HANA cockpit displays data at a landscape or enterprise level. You can quickly drill down to an overview of an individual resource. For example, you can drill down to links, data, tiles, and different parts of a single tile. This provides access to more detailed information and functions.

The SAP HANA database explorer is integrated into the cockpit. It allows you to query information about the database using SQL and Multidimensional Expressions (MDX) statements. You can also view information about your database's catalog objects.

LESSON SUMMARY

You should now be able to:

- Understand the SAP HANA Cockpit 2.0 architecture

Unit 5
Lesson 2

Installing SAP HANA Cockpit 2.0

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Install SAP HANA Cockpit 2.0

Installation Overview

The SAP HANA database lifecycle manager (HDBLCM) installs the SAP HANA cockpit in a graphical user interface or the command-line interface.

During the installation process, you can create a fully authorized administration user named COCKPIT_ADMIN. The master password entered during the installation is assigned to this user.

Ports, through which the SAP HANA cockpit and the SAP HANA cockpit manager can be accessed, are assigned automatically by the installer. The local host name, ports, and the master password are required to set up the SAP HANA cockpit. Make sure that you pass this information along to the cockpit administrator.

> Note:
> At the end of the installation, the URLs of the SAP HANA cockpit and the SAP HANA cockpit manager are shown in the installation logs screen.

Hardware Requirements

SAP HANA cockpit can be installed on Intel-based hardware platforms or on IBM Power Systems. The minimum hardware requirements for a production environment are 16 GB of RAM on a dedicated server.

The supported operating systems are SUSE Linux Enterprise Server (SLES) and Red Hat Enterprise Linux (RHEL). For specific information about the supported operating systems for the SAP HANA cockpit, see the following SAP notes.

Minimum Hardware Requirements

- Intel-based hardware platforms
- IBM Power systems
- 16 GB of RAM on a dedicated server (for production)

Operating System Requirements

- SUSE Linux Enterprise Server (SLES)

- SAP Note 1944799 — SAP HANA Guidelines for SLES Operating System
- SAP Note 1984787 — SUSE Linux Enterprise Server 12.x for SAP Applications Config. Guide

- Red Hat Enterprise Linux (RHEL)
  - SAP Note 2009879 — SAP HANA Guidelines for Red Hat Enterprise Linux Operating System
  - SAP Note 2002167 — Red Hat Enterprise Linux 7.x Configuration Guide for SAP HANA

SAP HANA Cockpit 2.0 Installation

The SAP HANA cockpit 2.0 software can be downloaded from SAP Support Portal (https://support.sap.com/swdc) via the SAP Software Download Center.

Download the `.SAR` file to the SAP HANA cockpit server and unpack the using the command line:

```
SAPCAR -vxf SAPHANACOCKPIT##_$$.SAR -manifest SIGNATURE.SMF
```

> Caution:
> Unpacking the `.SAR` with the `-manifest SIGNATURE.SMF` creates the required signature file.
>
> `##` is the Support Package Stack number.
>
> `$$` is the Patch level.

The SAP HANA cockpit application cannot be deployed as an application for the advanced model of SAP HANA extended application services on an existing SAP HANA instance. Instead, you install or update it in combination with its own runtime environment for SAP HANA extended application services, advanced model. The runtime environment for the advanced model of SAP HANA extended application services that is supplied with the SAP HANA cockpit cannot be used to deploy applications for SAP HANA extended application services, advanced model.

Extract the SAP HANA Cockpit .SAR File

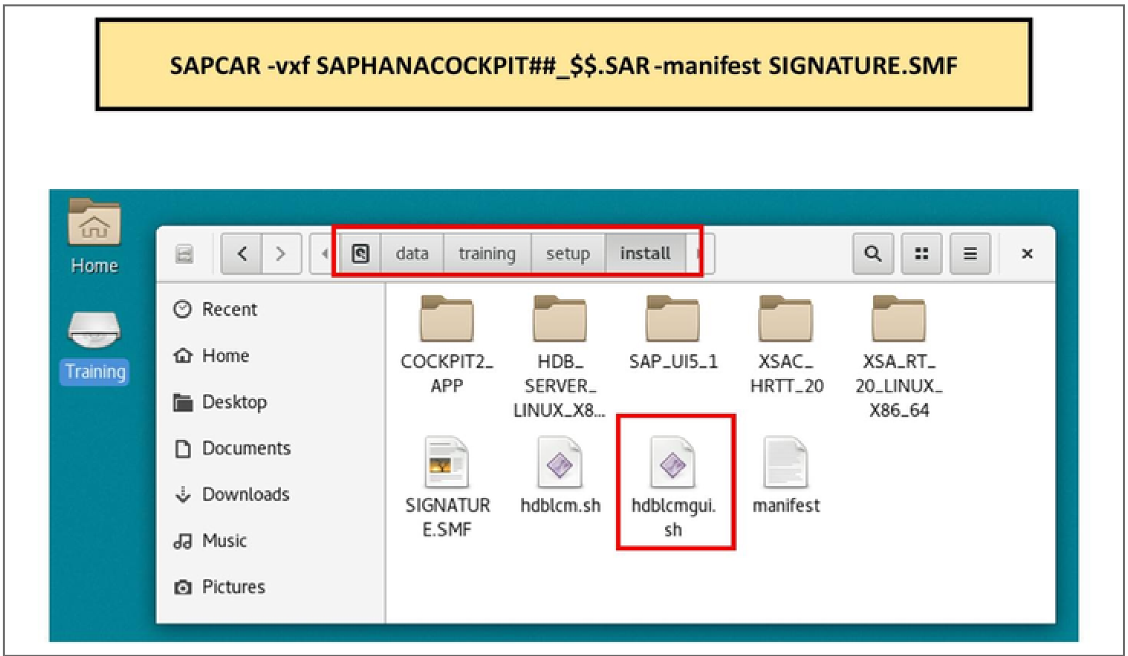

Figure 80: Extract the SAP HANA Cockpit .SAR File

After the .SAR file is unpacked, start the installation, as root user, by running the command hdblcmgui.sh .

The required input is preset when using the hdblcmgui, but it can be changed. You only need to provide the master password for the users created during the installation.

Predefined Installation Settings for SAP HANA Cockpit

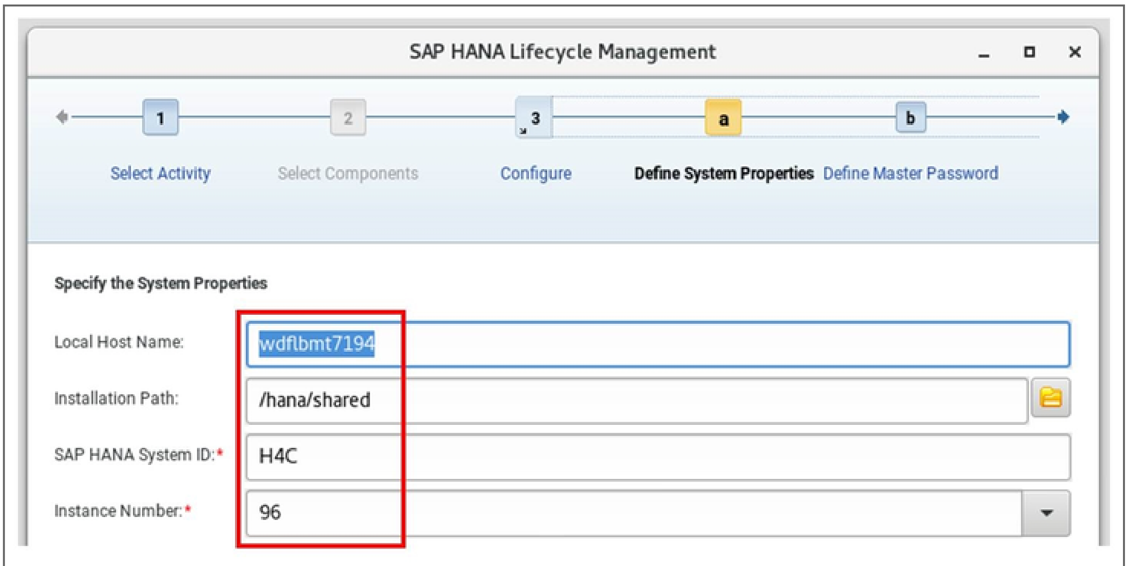

Figure 81: Predefined Installation Settings for SAP HANA Cockpit

As soon as the installation is finished, you can start configuring SAP HANA cockpit by creating Resources, Resource groups, and Cockpit users.

LESSON SUMMARY

You should now be able to:

- Install SAP HANA Cockpit 2.0

Unit 5
Lesson 3

Configuring SAP HANA Cockpit 2.0

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Configure SAP HANA Cockpit 2.0

Resources, Resource Groups, and Cockpit Users

With the Cockpit Manager, you can register resources and create resource groups and cockpit users. This allows you to manage all the SAP HANA systems in your landscape easily. With the created cockpit users, you can log on to the SAP HANA cockpit 2.0 and manage all the resources (SAP HANA systems) assigned to your resource group.

> Note:
> The tasks of registering resources, creating groups, and creating cockpit users do not have a set order.

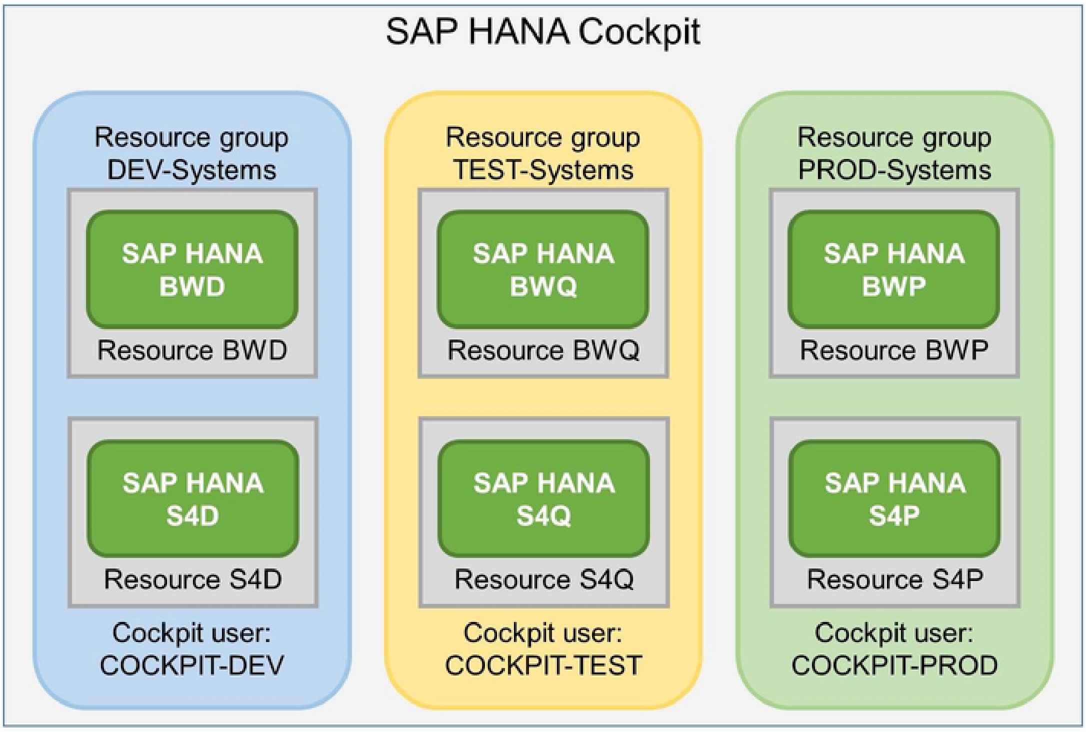

Figure 82: The Resources, Resource Groups, and Cockpit User Concept

You can register Cockpit Users, Resource Groups, and Resources with the Cockpit Manager. Start the Cockpit Manager with the following URL: https://<cockpit-host>:<port-number>. For example, in our training landscape the Cockpit Manager can be started using the URL: https://wdflbmt7285:51023.

Addition of Resources to SAP HANA Cockpit 2.0

To monitor and manage resources, register SAP HANA systems as a resource in SAP HANA cockpit 2.0.

> Hint:
> A resource is an SAP HANA system that you registered in SAP HANA cockpit 2.0.

You can register an SAP HANA system to the SAP HANA cockpit manager using the COCKPIT_ADMIN user account and its password. This account and password are created during the installation of SAP HANA cockpit 2.0.

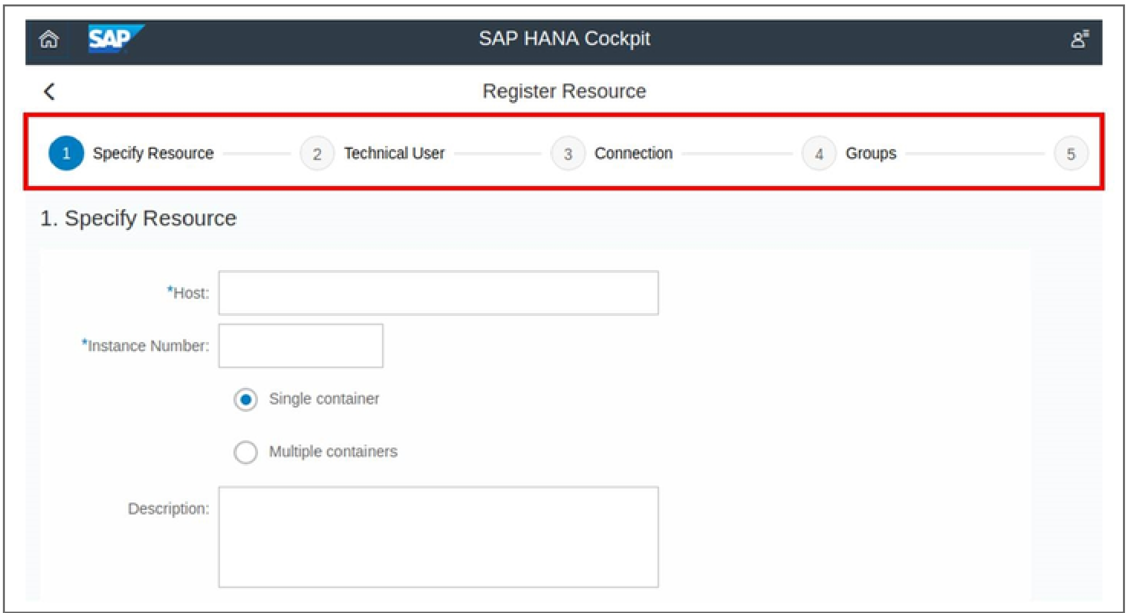

Figure 83: Register an SAP HANA System as a Resource

To register a resource, follow the guided procedure steps and provide the required information, like host, instance number, technical user, connection type, and group assignment.

Create Technical User

On the SAP HANA resource, create a technical user account that the cockpit will use to collect monitoring data (such as information on alerts and system performance). The technical user requires the CATALOG READ system privilege, and SELECT on the _SYS_STATISTICS schema.

> Note:
> Set up a dedicated account for the technical use. This technical user account should not be used by any other person.

Before you can register an SAP HANA system as a resource in SAP HANA cockpit 2.0, you need a technical user in that system. It's not possible to create the technical user by using the SAP HANA cockpit. Create this user manually, and grant the following minimum necessary authorization:

```
CREATE USER <username> PASSWORD < password> NO
FORCE_FIRST_PASSWORD_CHANGE;
GRANT CATALOG READ to <username>;
GRANT SELECT on SCHEMA _SYS_STATISTICS to <username>
```

If you plan to manage telemetry on the registered resource, grant additional authorization as follows:

```
GRANT SELECT, INSERT, UPDATE, DELETE, EXECUTE on SCHEMA _SYS_TELEMETRY
to <username>
```

Creation of Resource Groups

A resource group is a named set of one or more registered resources. Resource groups associate resources with cockpit users so that the cockpit users can manage and monitor specific resources through the cockpit. Use the Cockpit Manager to assign resources and cockpit users to resource groups.

Each registered resource belongs to a usage type resource group. These autocreated groups of resources (Production, Test, Development) are based on the system usage type of each resource. The system usage type is assigned during the installation.

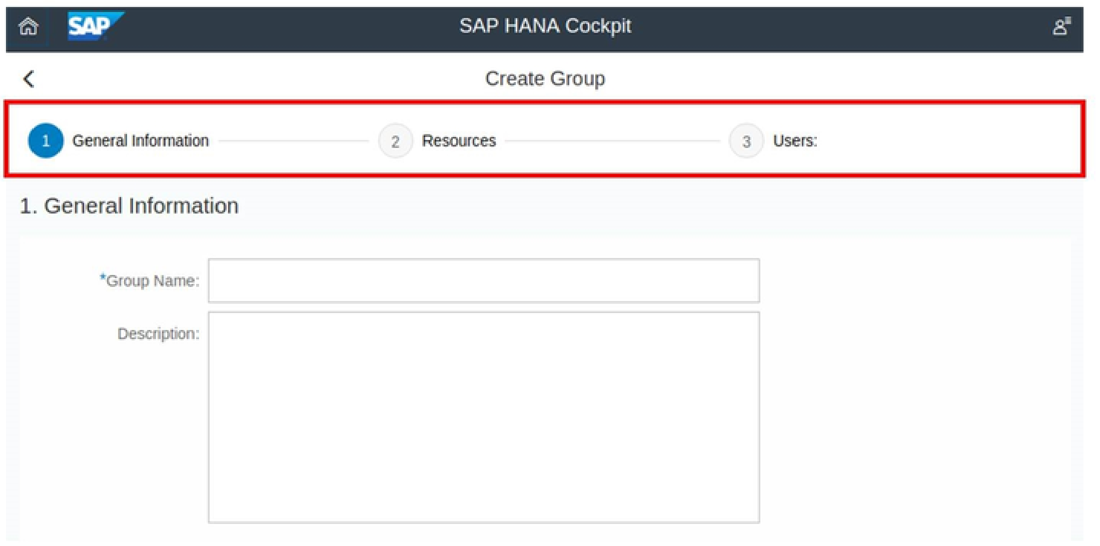

Figure 84: Create a Resource Group

To create a resource group, follow the guided procedure steps and provide the required information like group name, resource assignment, and cockpit user.

> Note:
> The system usage type can be changed in the global.ini file with the usage parameter in the system_information section. The resources can also belong to one or more groups that you create.

The cockpit administrator can use the resource groups to view and administrate similar SAP HANA databases (resources). The cockpit administrator can also control which other users have access to a resource.

In order to have access to a resource, a cockpit user must belong to one of the resource groups that you have created containing the resource.

> Note:
> You cannot assign users to an autocreated group.

After you have finished the setup of your landscape in Cockpit Manager, you can access the SAP HANA cockpit to monitor and administer the SAP HANA systems in your landscape.

Creation of Cockpit Users

With the Cockpit Manager, you can create cockpit users and assign them to groups of resources.

Cockpit users are SAP HANA cockpit application users. Therefore, they are separate from the database user credentials associated with the resources managed and monitored through the cockpit. Each cockpit user can be assigned access to groups of resources.

The COCKPIT_ADMIN, created during the installation, is the only administration user who can create other cockpit users (application users).

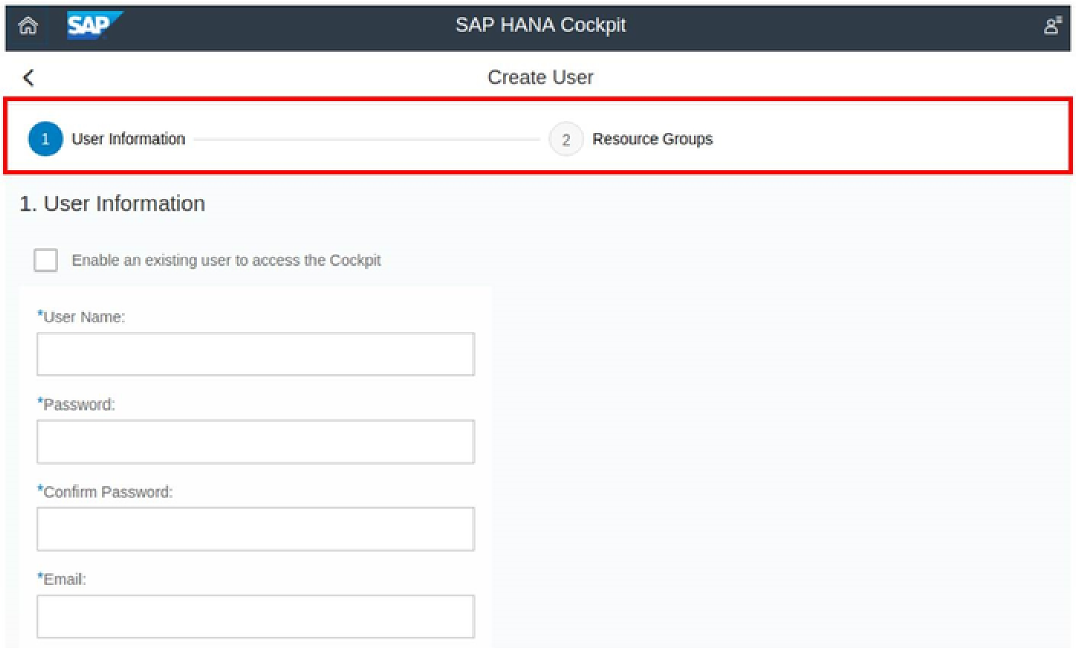

Figure 85: Create a Cockpit User

To create a cockpit user, follow the guided procedure steps and provide the required information, like user name, password, and group assignment.

Note:
In order to drill down into a specific system, each cockpit user also requires access as a database user. Database users are not managed through the Cockpit Manager, but rather through the Manage Users link on the Overview of a single resource.

When a cockpit user is assigned to a resource group, this user can monitor each of the resources within that group. The user also sees the aggregate data for the group. If no resource groups currently exist, you can add users to a new resource group after you have created resource groups.

A cockpit user that isn’t assigned to any groups, or is assigned to an empty group, can access the SAP HANA cockpit. However, the cockpit shows no data because the cockpit user has zero resources assigned.

SAP HANA Cockpit User Management

Each cockpit user must be assigned at least one cockpit role, which will dictate what portions of the cockpit or the cockpit manager they can acccess. (Cockpit roles are unrelated to the roles associated with database users. The latter govern which SAP HANA privileges are assigned to a database user).

The three available roles are Cockpit Administrator, Cockpit Resource Administrator and Cockpit User.

- Cockpit Administrator

  This role gives access to Manage Users and Cockpit setting , This role is automatically assigned to the user COCKPIT_ADMIN.

- Cockpit Resource Administrator

  This role gives access to the Registered Resource and Resource Groups sections of the Cockpit Manager,

- Cockpit User

  This role gives access to the SAP HANA Cockpit application where the user can view the resources in the assigned resource group.

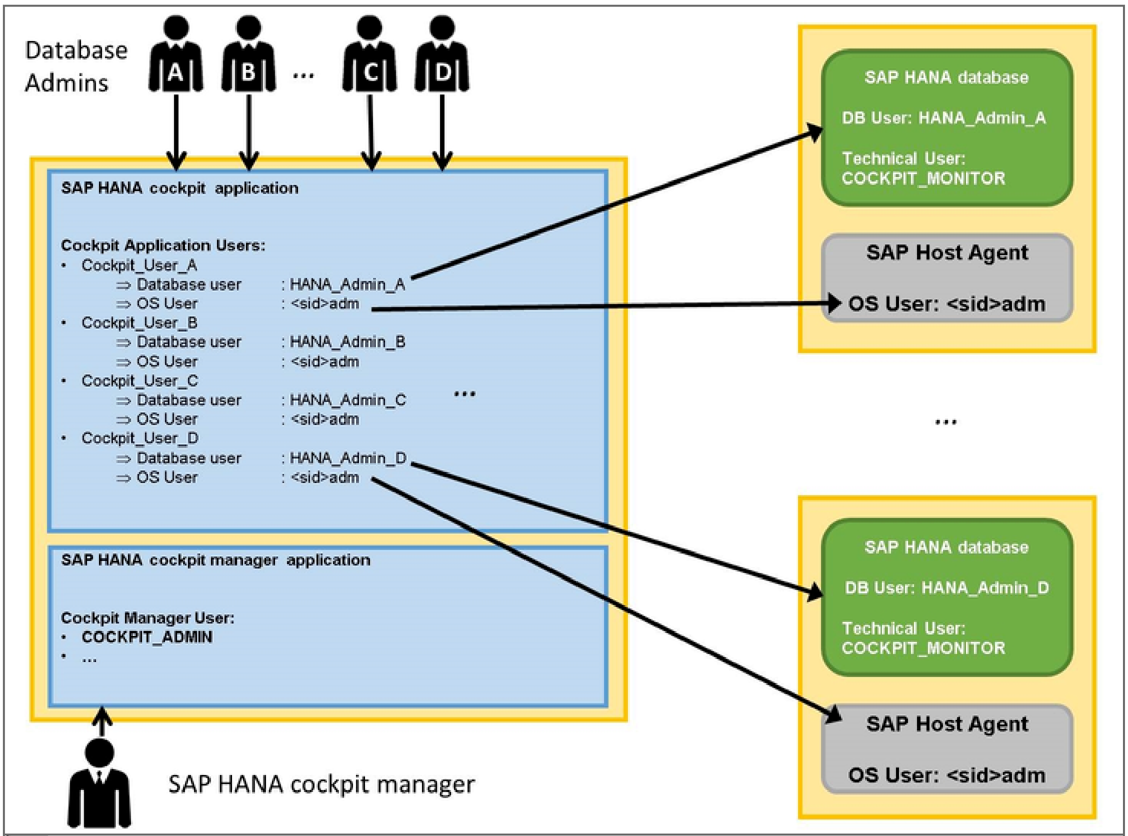

Figure 86: SAP HANA Cockpit User Management

During the setup and configuration there are several types of users created or assigned.

- The Cockpit Manager User (COCKPIT_ADMIN) is responsible for the creation and assignment of resources, groups, cockpit users.
- The Cockpit Application User (Cockpit_User_A) is a personal application user to grant the database administrator access to the SAP HANA Cockpit application.
- The DB User (HANA_ADMIN_A) is a personal user account in the resources for which the database administrator is responsible. The assignment from personal database user to the Cockpit User is done in the Resource Directory .
- The OS User is the Operating System User, usually <sid>adm, used to access the SAP Control process. This user is used for starting and stopping the database, and to control the restore process.
- The Technical User (COCKPIT_MONITOR) is a user account created per resource which is used for the monitor data collection in the Aggregate Health Monitor . The technical user needs to be created manually in each resource that needs to be monitored.

With this setup the segregation of duties can easily be implemented.

Opening of SAP HANA Cockpit

You can start SAP HANA cockpit 2.0 using the following URL: https://<cockpit-host>:<port-number> . For example, in our training landscape, the SAP HANA cockpit 2.0 can be started using the URL: https://wdflbmt7285:51021 . After the logon, you are presented with an overview of the resources assigned to your user account.

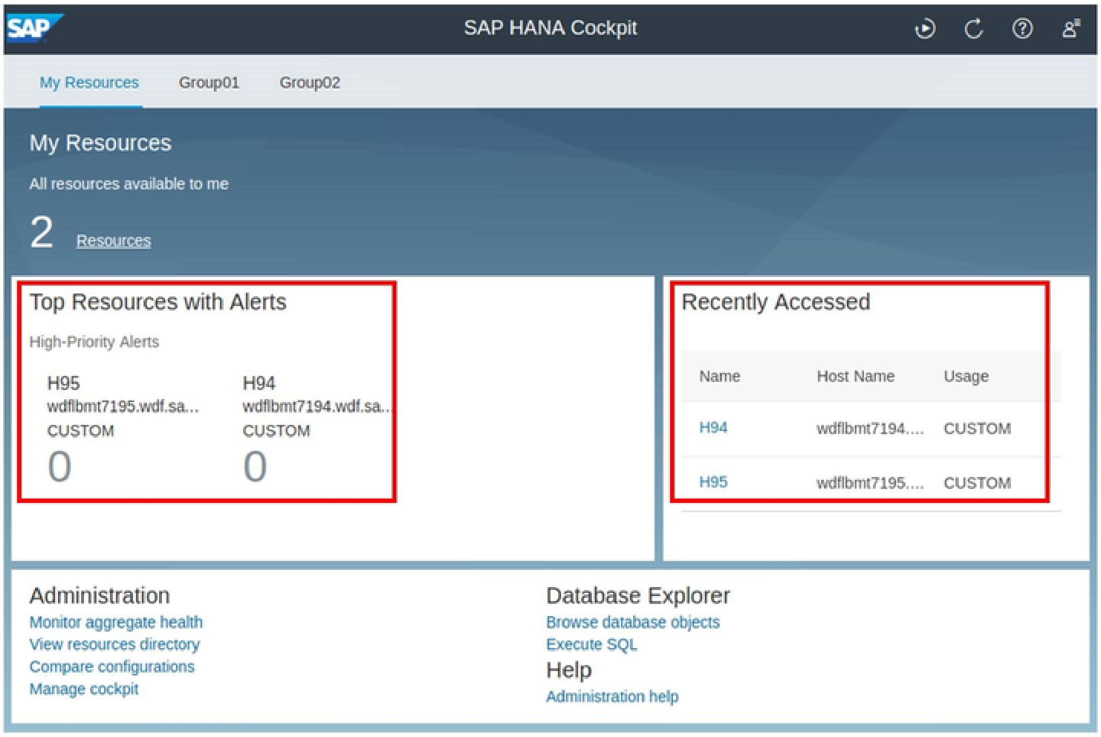

Figure 87: SAP HANA Cockpit Resource Overview

In the overview, to see the status of the SAP HANA systems quickly, select the Aggregated Health Monitor.

Detailed System Overview

From the Aggregated Health Monitor, navigate to the SAP HANA system overview page. This page displays the detailed status of the selected SAP HANA system.

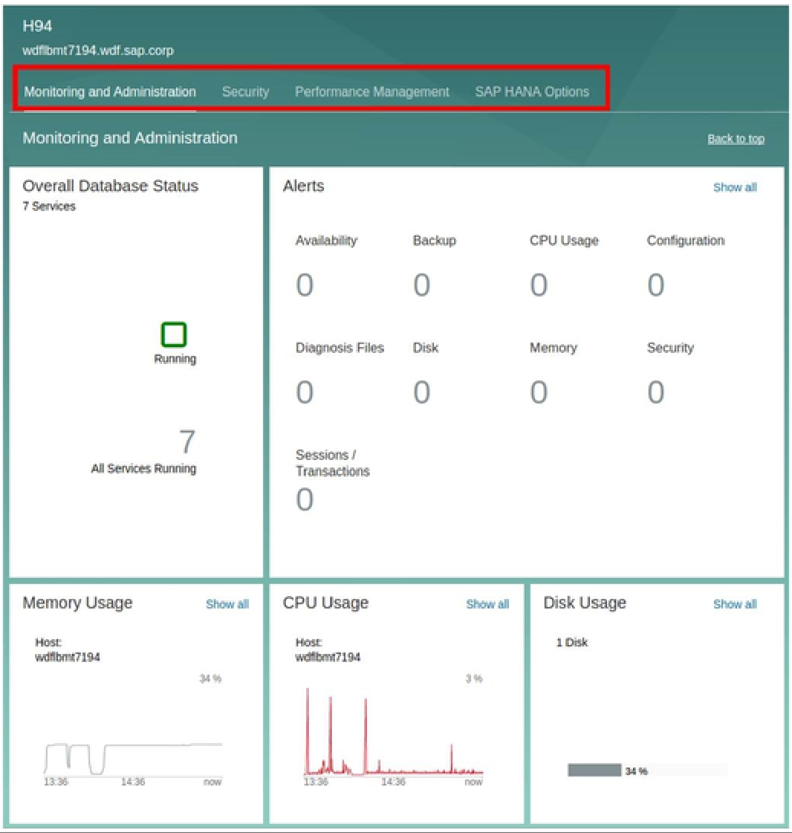

Figure 88: Detailed System Overview

The tiles in the detailed system overview provide information about the status of the SAP HANA database. The tiles refresh every few seconds and represent the current monitoring information.

LESSON SUMMARY

You should now be able to:

- Configure SAP HANA Cockpit 2.0

Unit 5
Lesson 4

Updating SAP HANA Cockpit 2.0

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Understand the SAP HANA Cockpit 2.0 revision strategy
- Update SAP HANA Cockpit 2.0 to a newer Support Package

SAP HANA Cockpit 2.0 Revision Strategy

SAP HANA cockpit 2.0 and all its subcomponents, including the Database Explorer and SQL Analyzer will receive new functionality, updates and corrections on a regularly basis. In this lesson you will learn about the SAP HANA cockpit 2.0 revision strategy.

SAP HANA Express, the XS Advanced (XSA) and a set of XSA multi-target applications are the main components of SAP HANA cockpit 2.0. An update to a newer SAP HANA cockpit 2.0 support package will update all the required SAP HANA cockpit 2.0 components.

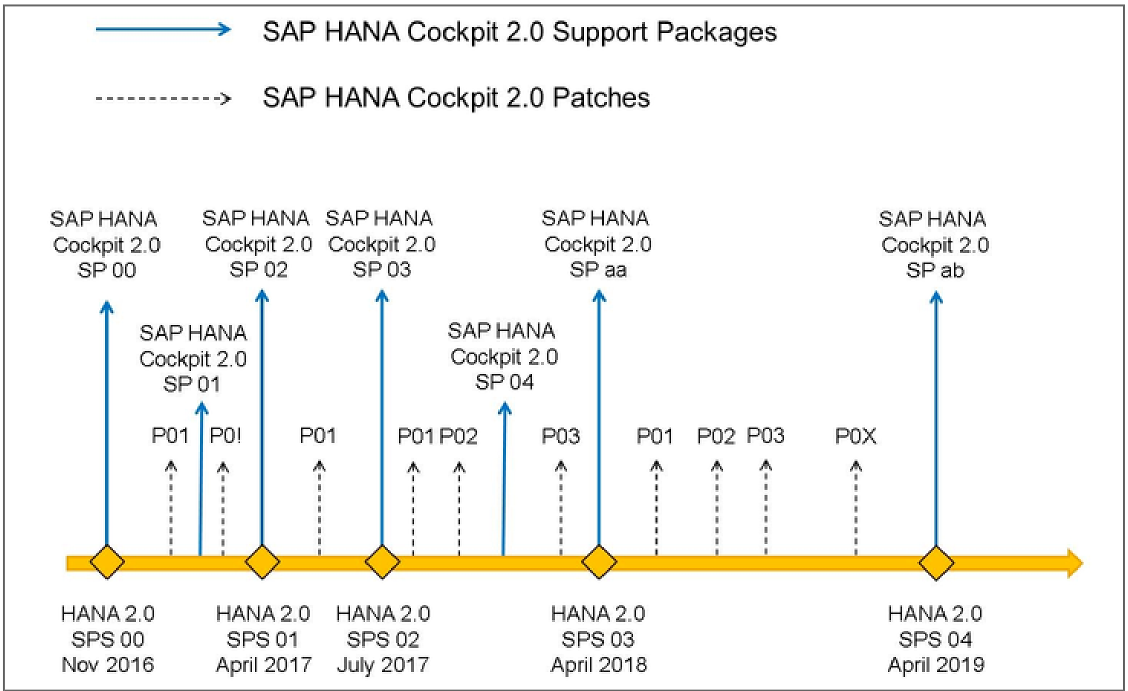

Figure 89: SAP HANA Cockpit Maintenance Strategy

Support Packages are periodically released for SAP HANA cockpit 2.0. These Support Packages include new functionality as well as fixes and security patches. Each new Support Package includes all the features and patches of the previous Support Packages.

Patches, which include only fixes and security patches, for the most recently released Support Package are provided until a new Support Package is provided. As soon as a new Support Package is released, SAP will stop providing patches for earlier support packages.

SAP HANA 2.0 cockpit support packages and patches are available from the SAP Software Downloads site (https://support.sap.com/en/my-support/software-downloads.html) in the Support Packages & Patches section, under SAP HANA Platform Edition 2.0. Each item is named "SP NN Patch MM for SAP HANA COCKPIT 2.0" where NN represents the Support Package number and MM represents the patch number.

SAP recommends that you consult the SAP Note for each Support Package release, accessible from the SAP HANA 2.0 Cockpit Central Release Note to define your upgrade path.

- SAP Note 2380291 - SAP HANA 2.0 Cockpit Central Release Note
- SAP Note 2433181 - SAP HANA 2.0 Cockpit Revision and Maintenance Strategy
- SAP Note 2513227 - SAP HANA cockpit 2.0 SP 04
- SAP Note 2378962 - SAP HANA 2.0 Revision and Maintenance Strategy

For the SAP HANA Platform, the Support Package Stacks are released on a yearly basis. These Support Package Stacks introduce new capabilities into SAP HANA. A new SAP HANA cockpit 2.0 is released as part of the SAP HANA Platform Support Package Stack. In between that yearly release cycle, SAP HANA 2.0 cockpit support packages can be released.

You do not have to align releases between SAP HANA cockpit 2.0 and the SAP HANA database systems to be able to administer and monitor these databases. SAP HANA cockpit 2.0 can administer and monitor SAP HANA databases from SAP HANA 1.0 SPS12 or later.

Update SAP HANA Cockpit 2.0

The SAP HANA cockpit can be updated using the SAP HANA database lifecycle manager (HDBLCM). This makes the update relatively easy as it is the same procedure as updating SAP HANA.

> Hint:
> Before starting the update it is recommend to perform a database backup or have a recent backup available.

To update an SAP HANA cockpit, you need to first download the installation files from Service Marketplace (SMP). This can be done manually using the SAP HANA database lifecycle manager (HDBLCM) Web user interface. Once the component packages have been prepared, the system update can be triggered from any of the three SAP HANA database lifecycle manager user interfaces.

> Caution:
> Do not update the SAP HANA cockpit components individually. Always use the SAP HANA database lifecycle manager (HDBLCM) to update the cockpit and all of its components in one step.

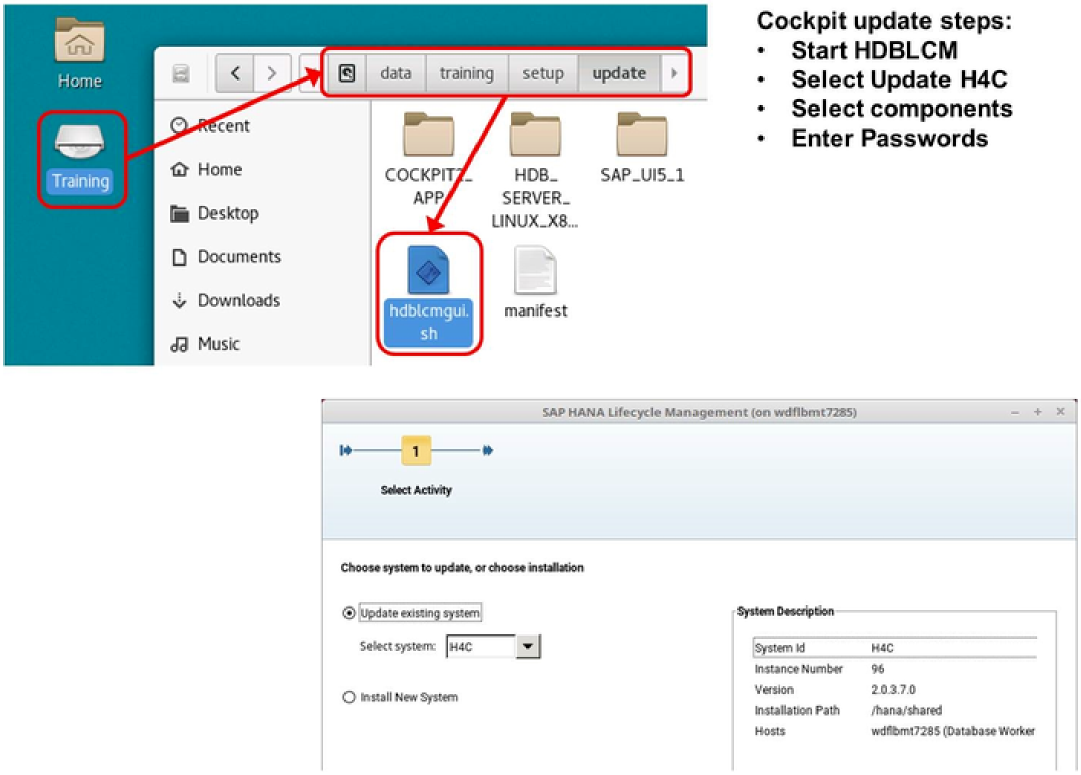

Figure 90: Start HDBLCM for HANA Cockpit 2.0 Update

The HDBLCM tool recognizes an installed SAP HANA cockpit, and will suggest updating the existing system. This update will update all the required components of the SAP HANA cockpit. HDBLCM will also request you to provide the password for the users <sid>adm, SYSTEM and the COCKPIT_ADMIN user.

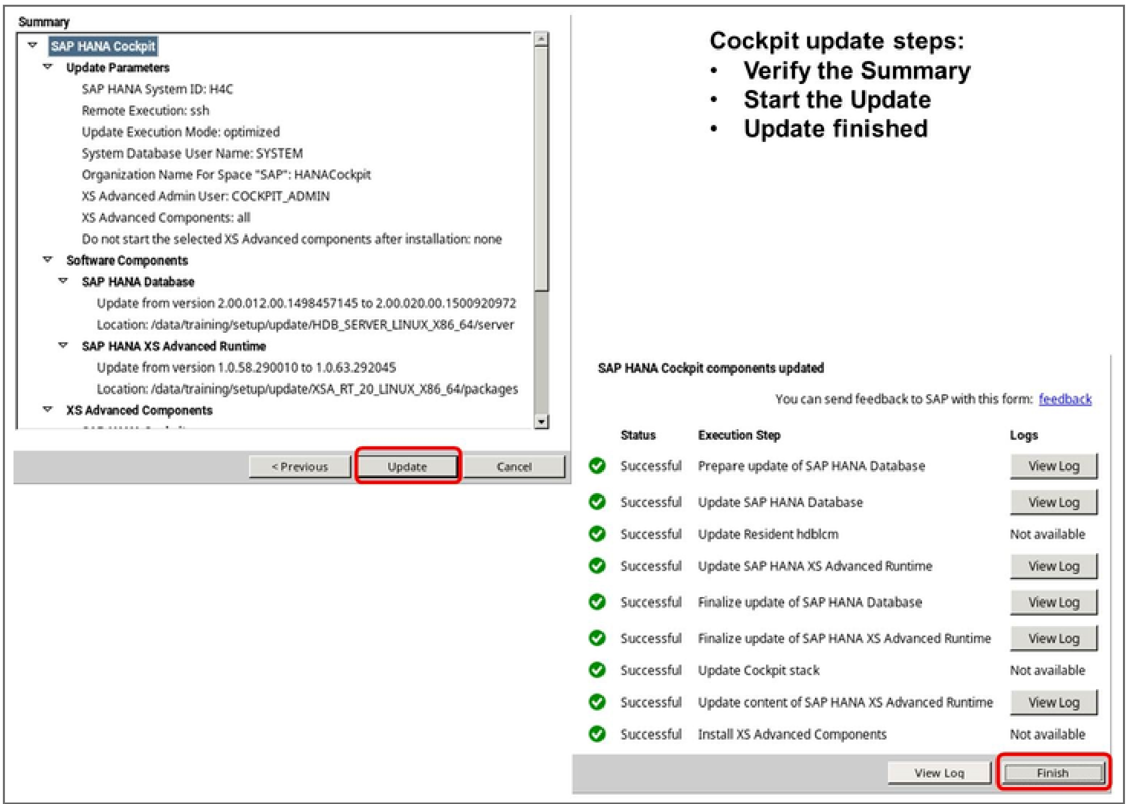

Figure 91: Start HANA Cockpit 2.0 Update

Note:
An update of SAP HANA cockpit takes around 30 minutes.

When all the required input is provided, a summary displays so that you can review your input. The database administrators should stop using the SAP HANA cockpit temporarily because the database and its components will be restarted during the update. The application data, registered resources, created groups, and users and their assignments are kept during the update.

LESSON SUMMARY
You should now be able to:

- Understand the SAP HANA Cockpit 2.0 revision strategy
- Update SAP HANA Cockpit 2.0 to a newer Support Package

Unit 5

Learning Assessment

1. The new SAP HANA Cockpit for SAP HANA 2.0 is installed separately on dedicated hardware.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

2. Which files are created after unpacking the SAP HANA Cockpit .SAR File?

   Choose the correct answers.

   - [ ] **A** Configuration File
   - [ ] **B** Signature File
   - [ ] **C** hdblcm.sh File

3. What is the order specified for registering resources and creating resource groups and cockpit users with the Cockpit Manager?

   Choose the correct answer.

   - [ ] **A** Resources, Resource Groups, Cockpit Users
   - [ ] **B** Resource Groups, Resources, Cockpit Users
   - [ ] **C** Cockpit Users, Resources, Resource Groups
   - [ ] **D** There is no order to follow

Unit 5

Learning Assessment - Answers

1. The new SAP HANA Cockpit for SAP HANA 2.0 is installed separately on dedicated hardware.

   Determine whether this statement is true or false.

   [X] True

   [ ] False

   Correct! This provides more flexibility because it allows you to manage more than one SAP HANA system in a single administration environment. In earlier versions of SAP HANA the SAP HANA Cockpit was an integral part of it. Read more on this in the lesson Introducing SAP HANA Cockpit 2.0 (Unit 5, Lesson 1) of the course HA200_14.

2. Which files are created after unpacking the SAP HANA Cockpit .SAR File?

   Choose the correct answers.

   [ ] **A** Configuration File

   [X] **B** Signature File

   [X] **C** hdblcm.sh File

   Correct! Unpacking the .SAR File with the parameter "-manifest SIGNATURE.SMF" creates the required signature file. Both the hdblcm.sh and the hdblcmgui.sh files will be unpacked. There is no configuration file created with the unpacking process. Read more on this in the lesson Installing SAP HANA Cockpit 2.0 (Unit 5, Lesson 2) of the course HA200_14.

3. What is the order specified for registering resources and creating resource groups and cockpit users with the Cockpit Manager?

   Choose the correct answer.

   - [ ] **A** Resources, Resource Groups, Cockpit Users
   - [ ] **B** Resource Groups, Resources, Cockpit Users
   - [ ] **C** Cockpit Users, Resources, Resource Groups
   - [x] **D** There is no order to follow

   Correct! There is no order to follow. The tasks of registering resources, creating groups and cockpit users do not have a set order. Read more on this in the lesson Configuring SAP HANA Cockpit 2.0 (Unit 5, Lesson 3) of the course HA200_14.

UNIT 6 SAP HANA Scenarios and Deployment Options

UNIT OBJECTIVES

- Talk about solution packages
- Explain the different deployment options for SAP HANA

Unit 6
Lesson 1

Describing SAP HANA Roadmap and Scenarios

LESSON OVERVIEW

This lesson focuses on SAP HANA use cases and scenario categories. These are discussed in conjunction with the SAP HANA roadmap and customer examples.

Business Example

While SAP HANA can be used as database management system in classic system setups for existing applications, it can also be the basis for a new generation of in-memory applications and use cases. For customers, it is important to understand the different use cases and scenario categories to be able to discuss potential roadmaps and migration paths for the system landscape.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

Overview of SAP HANA Roadmap

The figure, A Potential Roadmap for Using SAP HANA in Your System Landscape , shows a potential roadmap for the adoption of SAP HANA. The side-car scenarios allow you to start with a small SAP HANA system, implement clear scenarios, and solve existing issues. The use of SAP HANA as primary persistence for existing applications facilitates more comprehensive optimizations. The maximum improvement can be achieved by implementing tailor-made applications for SAP HANA.

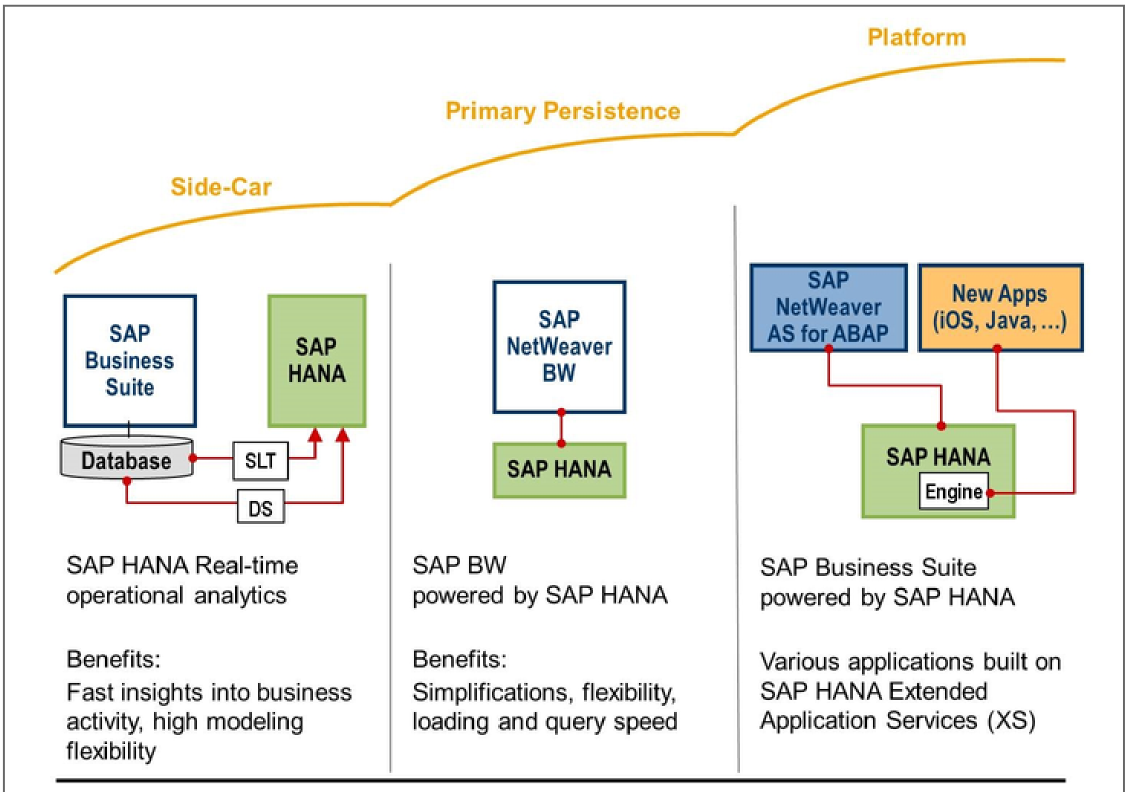

Figure 92: A Potential Roadmap for Using SAP HANA in Your System Landscape

> Note:
> This is an example for a potential roadmap with an increasing adoption of SAP HANA in the system landscape, and is not a standard recommendation. Depending on the customer requirements, other steps could be more reasonable, for example, using SAP HANA as the primary database in the first wave.

SAP HANA Scenarios

Depending on the system architecture, you can distinguish between side-by-side scenarios and integrated scenarios. In side-by-side scenarios, SAP HANA is added as an additional component to an existing landscape to facilitate analytical features or to accelerate processes. In integrated scenarios, SAP HANA is used as a primary database.

Furthermore, SAP HANA contains features that allow you to use it as a platform for a new generation of applications.

Examples of SAP HANA side-by-side scenarios are operational and agile data marts, and SAP HANA-based accelerators.

Data Mart Scenarios

A data mart is an industry term for a repository of data gathered from operational data originating in transactional systems or other sources. It is designed to serve a particular community of information workers by forming a basis for analytics, reporting, or a specific use in another type of application. The data mart aims to meet the specific needs of a particular group of users in terms of analysis, content, presentation, and ease-of-use.

With SAP HANA, operational data marts offer real-time analytics and reporting on data replicated from a transactional system's database. The raw tables are copied (structure and data) from the transactional system's database into SAP HANA. As new data is added into the relevant tables in the transactional system's database, copies of those records are transferred automatically into SAP HANA using replication technology. These replicated tables form the basis for specialized views that are created for analytics purposes. In some cases, the data modeling effort involved in developing these views can be significant, particularly when they convert raw transactional table data into a form that is best suited for analytics. Business Intelligence tools, such as the BI Tool Suite available from SAP BusinessObjects, are used for analysis and reporting.

Agile Data Marts

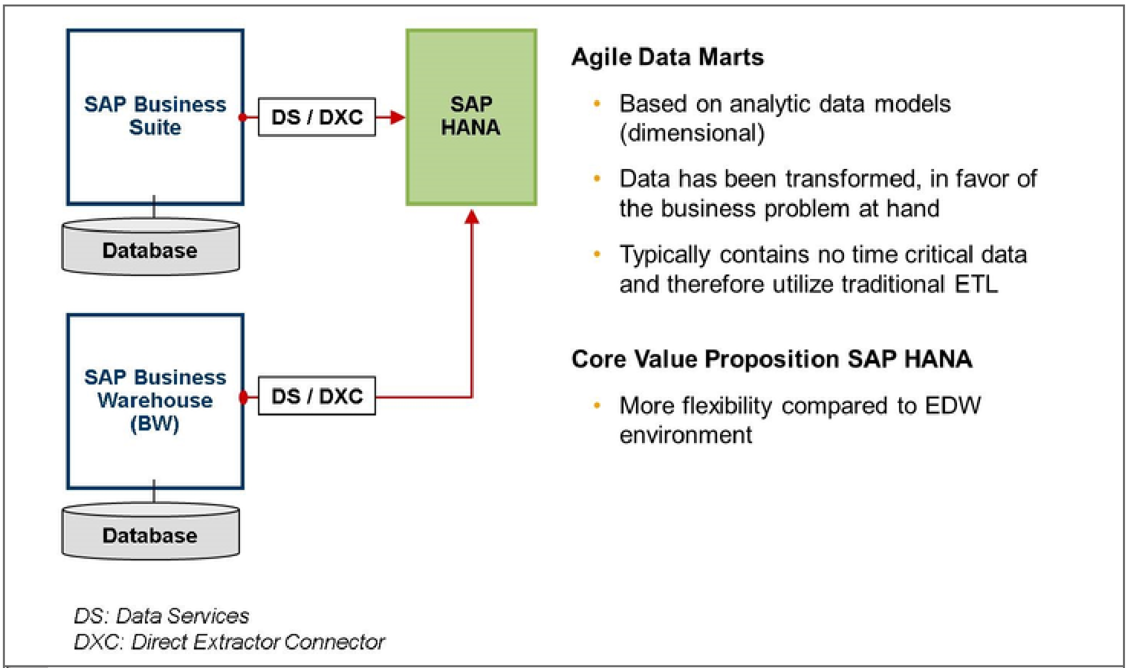

Figure 93: Side-by-Side Scenarios: Agile Data Marts

Agile data marts are a type of data mart that offer analytics and reporting on data acquired from a transactional system. When deployed in SAP HANA, they can offer greater flexibility compared to more comprehensive approaches to organizational information management; approaches such as deploying data marts within the context of an Enterprise Data Warehouse.

The goal of using SAP HANA as an agile data mart is to create more flexibility compared to an Enterprise Data Warehouse, because it is often realized using SAP Business Warehouse. Data is loaded by a traditional extract, transformation, and loading (ETL), for example, SAP BusinessObjects Data Services, and it has already been transformed. Furthermore, data models in SAP HANA can be implemented to connect data in different tables or apply additional logic. Agile data marts generally do not target the realization of real time reporting, but target the increase of modeling and reporting flexibility.

Operational Data Marts

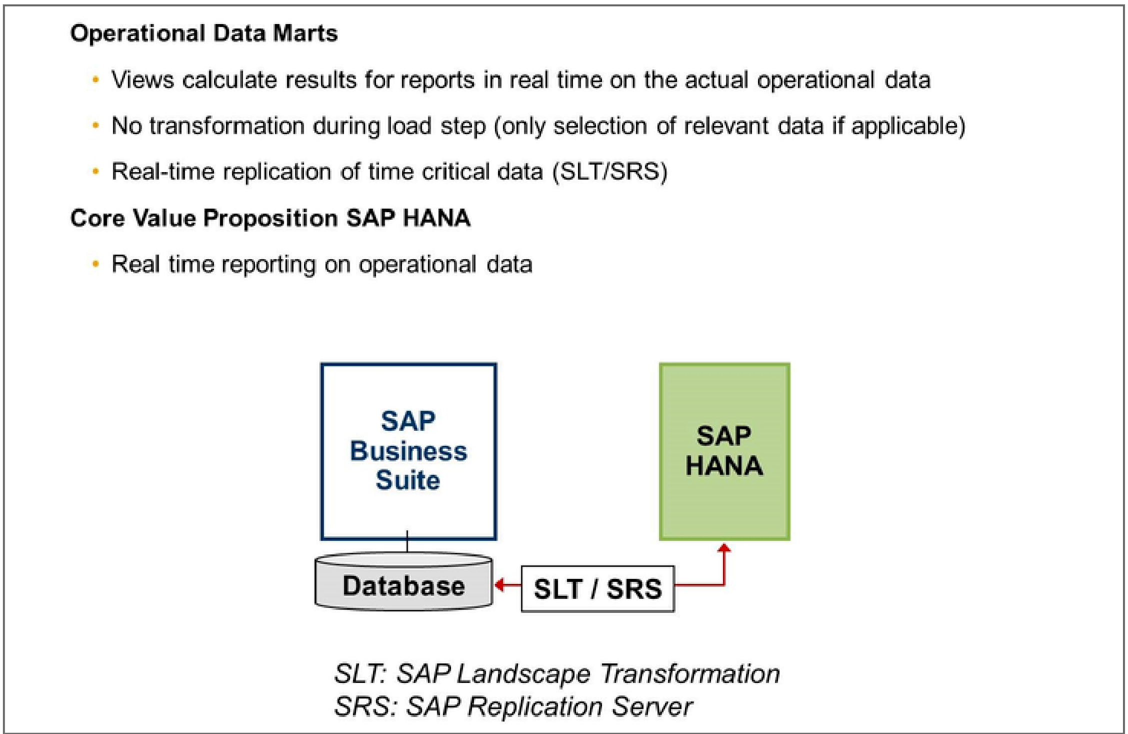

Figure 94: Side-by-Side Scenarios: Operational Data Marts

In contrast to agile data marts, operational data marts are focused on the requirements of operational reporting. Data can be acquired with low latency from SAP and non-SAP sources using SAP Landscape Transformation Replication Server for SAP HANA. The SAP Landscape Transformation Replication Server is an SAP NetWeaver ABAP-based tool that provides real time data replication. In addition, a log-based SAP Replication Server can also be used to provide real time data replication for an SAP Business Suite system.

Because data models implemented in SAP HANA do not require you to materialize aggregated data, the combination of using SAP HANA data models with real time data acquisition technologies allows you to implement reporting solutions that reflect data changes in the source systems immediately.

SAP HANA Accelerators

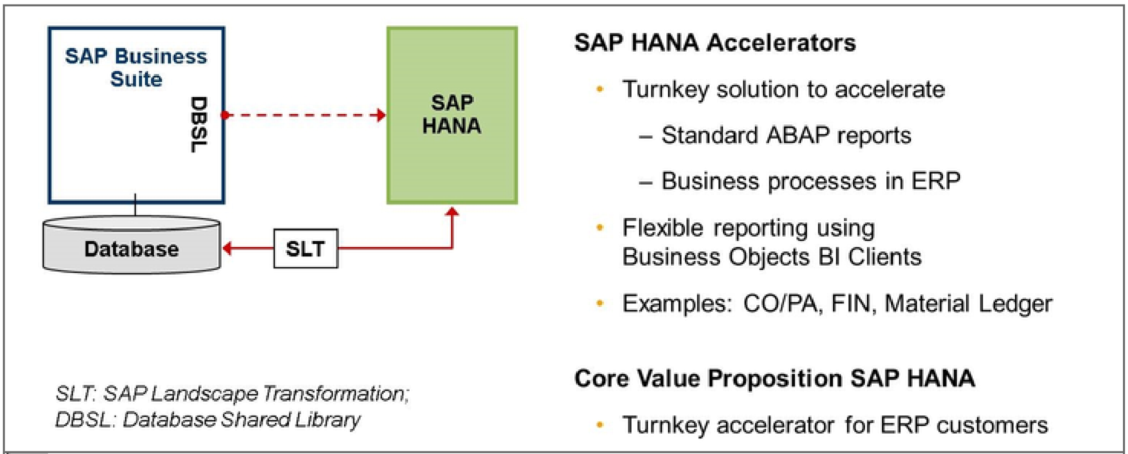

Figure 95: SAP HANA-Based Accelerators

SAP HANA accelerators enable the acceleration of standard ABAP reports, as well as selected business processes in SAP Business Suite systems. One example of this is a solution for SAP HANA accelerated finance and controlling that uses SAP HANA for financial accounting, controlling, material ledger, production cost analysis, and profitability analysis. It is also offered as RDS. For more information, see SAP Best Practices Explorer at https://rapid.sap.com/bp/.

Various other SAP HANA-based accelerators are offered by SAP. It is also possible to use SAP HANA as an accelerator for customer-individual implementations.

The typical approach for accelerators involves replicating data for data-intensive operations that are often bottlenecks for the given operation in an SAP HANAtable. A type of "switch" is then set in the SAP Business Suite application to indicate that whenever these specified tables are read, the read operation occurs in SAP HANA using a secondary database connection.

Architecturally, data is transferred with low latency to SAP HANA, which is used as secondary database. Using the appropriate Database Shared Library, the SAP Business Suite system accesses SAP HANA instead of the primary database for the specified reports or processes. This allows it to benefit from the acceleration or additional functionality implemented in SAP HANA.

SAP HANA as Primary Persistence for SAP NetWeaver-Based Applications

SAP Business Suite applications (SAP ERP, SAP CRM, SAP SCM, and so on), SAP Business Warehouse (SAP BW), and other SAP enterprise solutions are built on SAP's general platform, SAP NetWeaver.

SAP NetWeaver has two distinct aspects, ABAP and Java. Many applications built on SAP NetWeaver's ABAP or Java application servers are able to run on SAP HANA, where SAP HANA serves as the sole database in the architecture.

The technical interfaces are available for applications built on SAP NetWeaver Application Server for ABAP and SAP NetWeaver Application Server for Java to run on SAP HANA. However, specific development enablement is normally needed for each application to ensure that it runs optimally on SAP HANA. SAP Business Suite applications (SAP ERP, SAP CRM, SAP SCM, and so on), SAP Business Warehouse (SAP BW), and other SAP NetWeaver-based applications have been renovated to run on SAP HANA so that they can exploit its many advantages. Additionally, various components and complimentary applications that are built on SAP NetWeaver can also run on SAP HANA through the use of the provided SAP NetWeaver DB interfaces.

SAP HANA as Primary Persistence

In the scenario where SAP HANA is the primary persistence for SAP NetWeaver-based applications, note that SAP NetWeaver ABAP and Java application servers must run on separate hardware servers from the SAP HANA hardware.

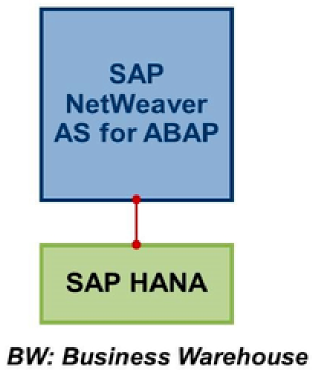

SAP BW on HANA

SAP Business Suite powered by SAP HANA

- SAP HANA Database becomes primary persistence of ABAP application server
- All Objects and BW loading procedures are accelerated by in memory technology
- High modeling flexibility

Core Value Proposition SAP HANA

- Speed and simplification for SAP BW and Business Suite

Figure 96: SAP HANA as Primary Persistence

In integrated scenarios, SAP HANA is used as the primary persistence for applications. This is achieved by migrating existing SAP Business Suite systems to SAP HANA, or by performing greenfield installations directly on SAP HANA. With becoming the primary persistence of the ABAP application server, all objects and processes can use the in-memory technology

Caution:
Although, architecturally, it looks as if the change solely affected the database layer, the application running on SAP HANA must be optimized explicitly in advance to use the capabilities and push down calculation intense logic to the database. Therefore minimum versions, respectively Enhancement Package levels, exist and contain SAP HANA support.

Deployment Options for SAP HANA and SAP NetWeaver AS for ABAP

SAP HANA and SAP NetWeaver AS for ABAP can be deployed on two different servers, or on one server.

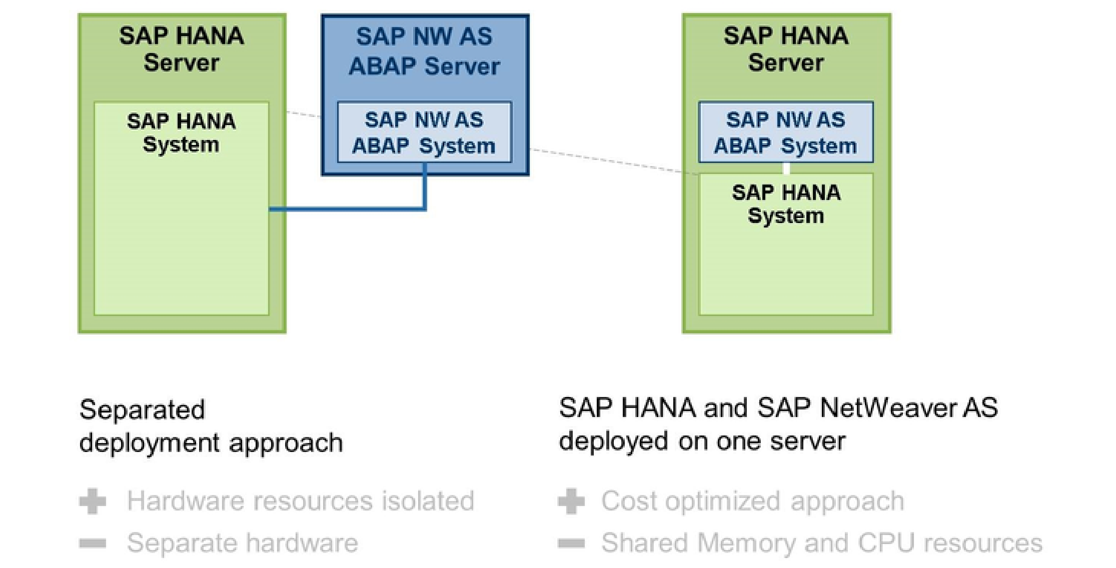

Figure 97: Deployment Options for SAP HANA and SAP NetWeaver AS for ABAP

Deployment of SAP HANA and SAP NetWeaver AS for ABAP on one hardware is available for all productive and nonproductive SAP HANA single node installations. All products based on SAP NetWeaver AS for ABAP 7.4 are supported.

The requirements are as follows:

- Additive sizing: additional memory resources for the SAP NetWeaver AS for ABAP system must be available on the SAP HANA server.
- Both systems require separate SIDs.

Connection of the ABAP Work Process to SAP HANA

The application data of the SAP NetWeaver-based applications are stored in the schema `SAP<Schema-ID>` of the SAP HANA system. The ABAP work processes connect to the SAP HANA system with the username `SAP<Schema-ID>` and the respective password of the user.

For automatic logon during the start of the SAP NetWeaver AS for ABAP system, the password is stored in the SAP HANA secure user store (`hdbuserstore`). This is a tool installed with the SAP HANA client software. It is used to store connection information to SAP HANA systems securely on the client so that client applications can connect to SAP HANA without users having to enter this information. It is typically used by the SAP NetWeaver AS for ABAP system or by scripts connecting to SAP HANA.

The connection information of the user `SAP<Schema-ID>` is stored in the “DEFAULT” key.

In an SAP NetWeaver AS for ABAP system, the `hdbuserstore` program is located in the `/usr/sap/<SID>/hdbclient` directory.

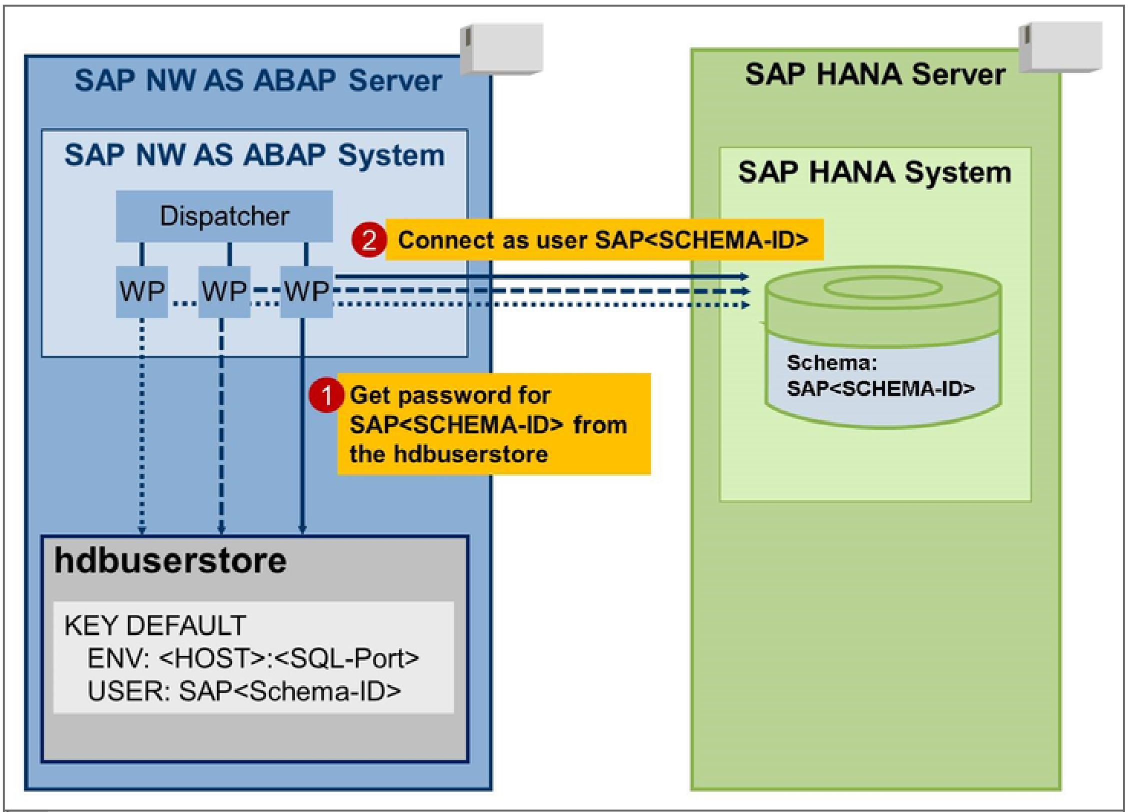

Figure 98: Connection of the ABAP Work Process to SAP HANA

> Note:
> As an alternative, you can use the "Secure Storage in File System (SAP NetWeaver AS for ABAP)" (SSFS) for the storage of the password for the ABAP database user. See SAP Note 1639578 : SSFS as password store for primary database connect .

If you change the `SAP<Schema-ID>` password of a NetWeaver based system, then you also must update the relevant DEFAULT connection entries in the `hdbuserstore`, or update the ABAP SSFS of the `<sapsid>adm` (Linux) or `Domain\SAPService<SAPSID>` (Windows) operating system user. Otherwise, the SAP NetWeaver AS for ABAP work processes continue to use the old password when establishing the connection.

Migration to SAP HANA

Technically, a migration to SAP HANA is only a change of the database and it does not affect most of the other components in the landscape. An SAP Business Suite system running on SAP HANA can still connect to, and be integrated with, other systems and hubs in the same way as a Business Suite system running on any other database. Furthermore, the same front ends and clients can be used to connect to the system. Even the application servers can be reused as they are, because they run on separate servers and not on the database host.

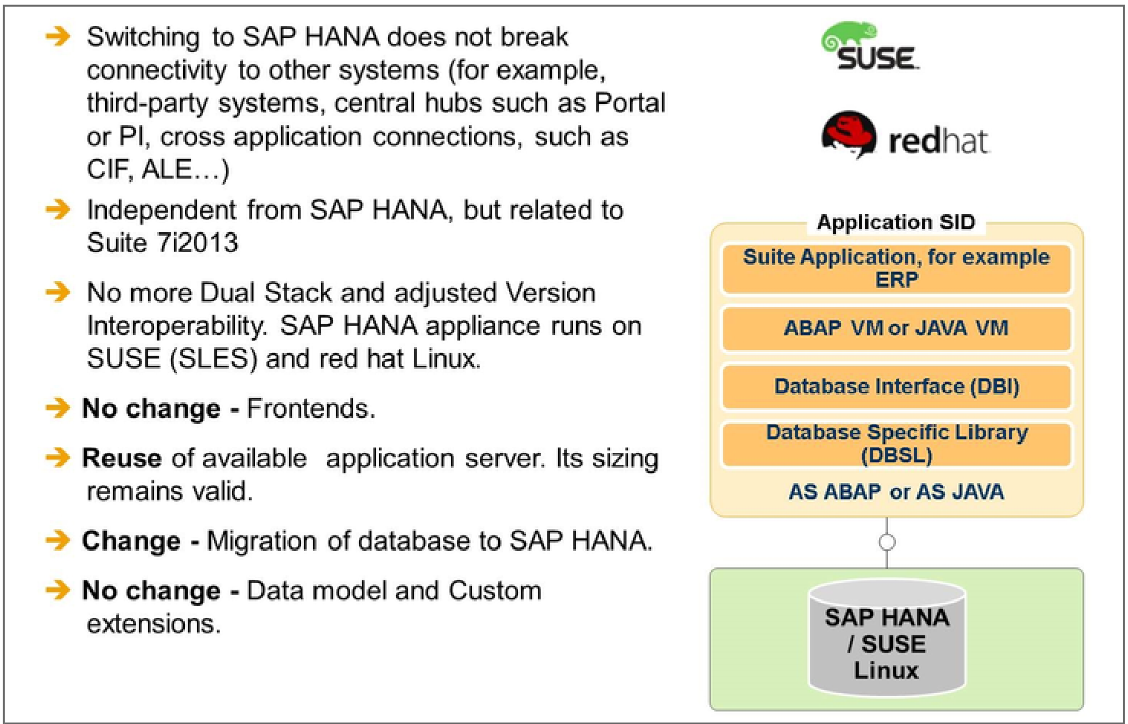

Figure 99: Overview: Migration of SAP Systems to SAP HANA

Migrating your existing SAP system to the SAP HANA database involves switching the SAP system to a new database that runs on a new host, because SAP HANA is an appliance.

Two Ways to Migrate to SAP HANA

A migration to SAP HANA can be performed in the following two ways:

- Heterogeneous system copy using Software Provisioning Manager (SWPM)
- Database migration option (DMO) of the Software Update Manager (SUM)

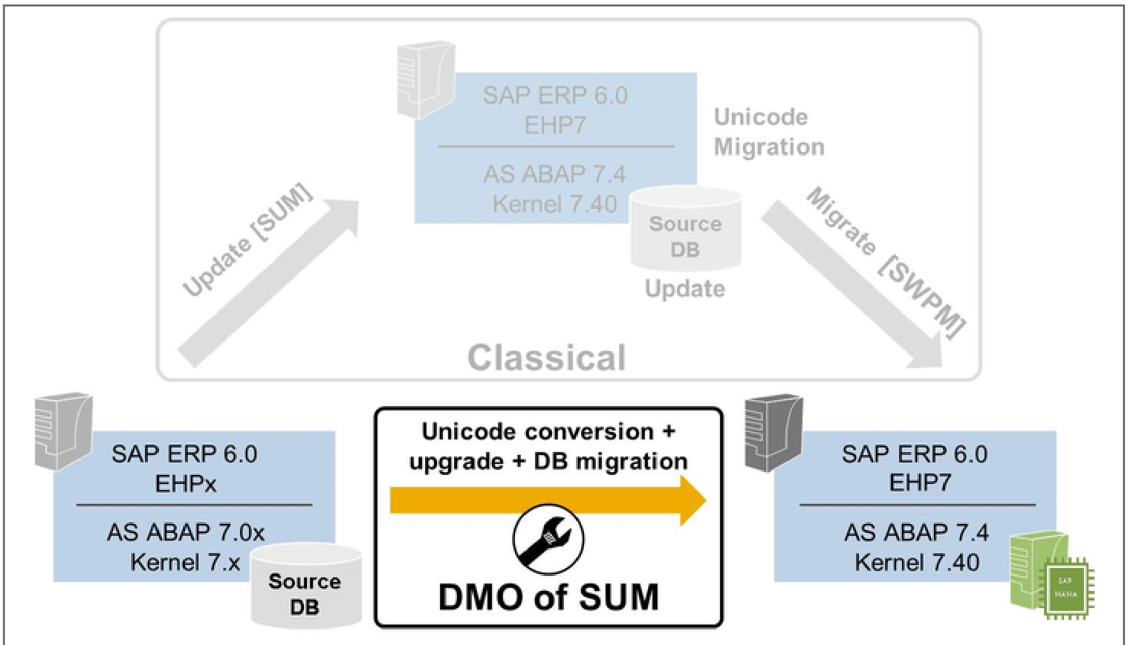

Figure 100: Two Ways to Migrate to SAP HANA

Initial Situation

The classical migration involves the sequence of SAP software update, which uses the Software Update Manager (SUM), and heterogeneous system copy, which uses the Software Provisioning Manager (SWPM). The database migration option simplifies the migration and is referred to as the one-step procedure to SAP HANA.

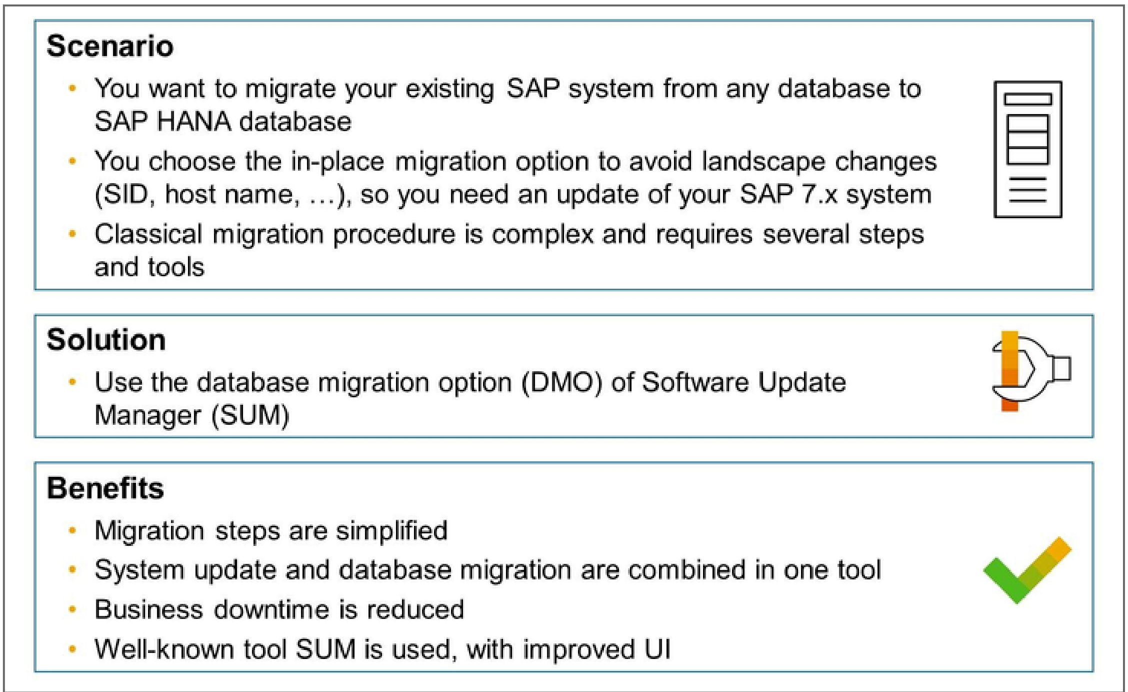

Figure 101: Initial Situation

The database migration option is not a new tool, it is just an option. A new option is an existing tool named Software Update Manager (SUM). SUM is the trusted tool for system maintenance, such as the following:

- Release upgrades
- Enhancement Package implementations
- Support package stacks for SAP NetWeaver-based systems

The Migration Process

For an inplace migration using the database migration option, the upgrade and migration are performed in a combined procedure, which reduces the total cost of ownership and risks.

Performing the migration in a combined procedure offers the following benefits:

- The combined procedure requires only one maintenance phase and not two.

  This reduces business downtime (total cost of ownership), and fewer regression tests are necessary.

- The original database is kept, and can be reactivated as a fallback.

  This reduces risk, no restore is required, and there is more time for testing before cutover.

- There are fewer prerequisites for SAP and DB start releases.

This reduces effort (total cost of ownership), and there are no additional licenses for traditional database updates.

- In-place migration keeps the application server and System ID stable.

  This has a low impact on system landscape because only the database server is new.

- For SAP BW, the database migration option can be applied when Post Copy Automation (PCA) is used,

SAP HANA as an Application Platform

SAP HANA provides the basis for an application development platform, where many different types of applications can be built on, and run on SAP HANA.

SAP HANA can be used as a sidecar to or primary database of existing applications. It can also be used as an entire application platform as follows:

- Any application can directly connect to SAP HANA using standard interfaces, such as Java Database Connectivity and Open Database Connectivity.
- Native SAP HANA applications can be implemented in SAP HANA, without requiring an additional application server from SAP HANA extended application services.

Note:
From SPS 11, SAP HANA includes an additional runtime environment for application development: SAP HANA extended application services, advanced model. SAP HANA extended application services, advanced model, represents an evolution of the application server architecture within SAP HANA. It builds on the strengths and expands the scope of SAP HANA extended application services, classic model. Customers and partners who want to develop new applications should use SAP HANA extended application services, advanced model.

If you want to migrate existing applications from SAP HANA extended application services, classic model, to the new SAP HANA extended application services, advanced model, runtime environment, first check the features available with the installed version of SAP HANA extended application services, advanced model. If the features match the requirements of the SAP HANA extended application services, classic model, that you want to migrate, then you can start the migration process.

SAP HANA Extended Application Services, Classic Model

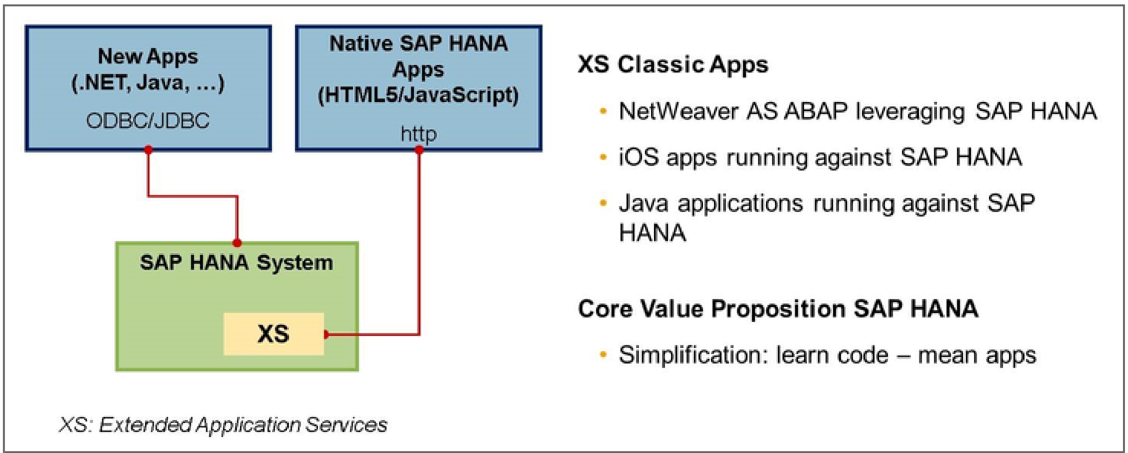

Figure 102: SAP HANA Extended Application Services, Classic Model

Within this category, there are two different types of applications that can be designed in this manner: Native SAP HANA applications, and applications with another application server that connects to SAP HANA. These applications can be described as follows::

- Native SAP HANA applications

  While SAP HANA is a database comprised of innovative technology, it is also much more than that. SAP HANA includes a small-footprint application server, a web server, and a repository for content, which provides lifecycle management functionality for development artifacts. Together with development tools, these components form an application development platform and runtime that can be used to build, deploy, and operate all kinds of SAP HANA-based software applications. These applications normally have a HTML or mobile app user interface that connects to SAP HANA using HTTP. The name for these capabilities is SAP HANA extended application services.

- SAP HANA-based applications with another type of application server (for example, .NET or Java)

  Various types of applications can be built on, and run on, SAP HANA, using the architecture of other well-known application servers and languages. Applications written using .NET are integrated with SAP HANA using Open Database Connectivity (ODBC), which is a standard, implementation-agnostic C-based API for accessing a database. Applications written using Java integrate using Java Database Connectivity (JDBC), which works similarly to ODBC. These interface types provide methods for creating and maintaining connections, transactions, and other mechanisms to create, read, update, and delete operations in SAP HANA. These methods map directly to the underlying SQL semantics, hiding the actual communication details. Essentially, any application that can use ODBC, ODBO, or JDBC can integrate with SAP HANA.

Hint:
The SAP HANA extended application services for development, modelling and tooling will be changing. They will be based on the SAP HANA extended application services, advanced model.

SAP plans to remove SAP HANA extended application services, classic model and the corresponding SAP HANA Repository with the next major product version of SAP HANA (see SAP note 2465027 - Deprecation of SAP HANA extended application services, classic model and SAP HANA Repository .

SAP HANA Extended Application Services

A combination of several technologies can be used for controlling the data processing and calculation logic, implementing the control flow, and creating the front end.

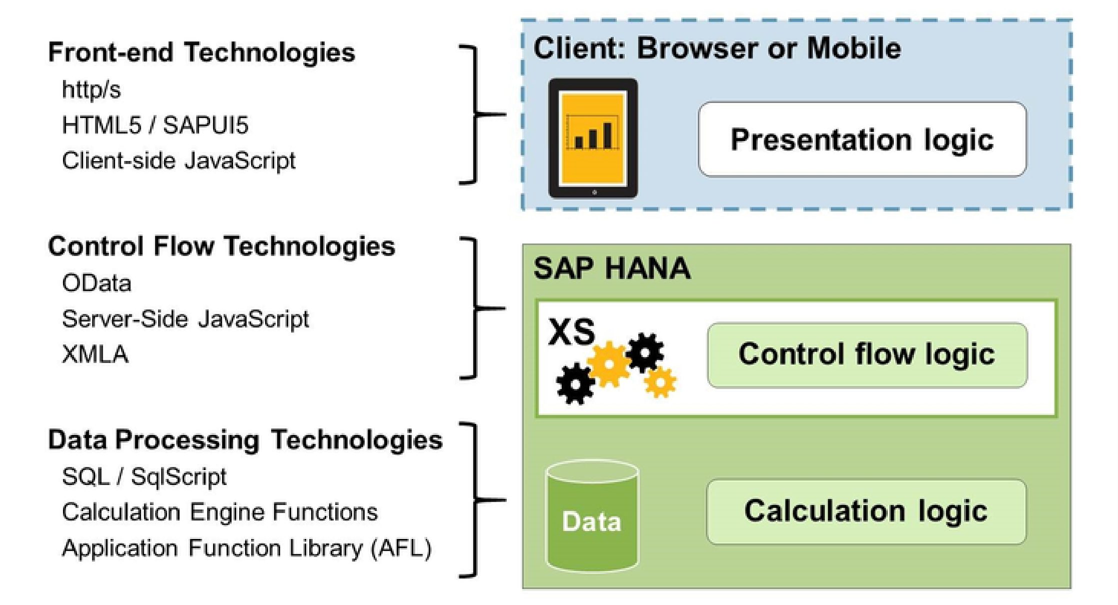

Figure 103: SAP HANA Extended Application Services

SAP HANA Extended Application Services, Advanced Model

SAP HANA extended application services, advanced model, provide a comprehensive platform for the development and execution of native data-intensive applications.

SAP HANA functions as a comprehensive platform for the development and execution of native data-intensive applications that run efficiently in SAP HANA, taking advantage of its in-memory architecture and parallel execution capabilities. Structured accordingly, applications can gain from the increased performance provided by SAP HANA because of the integration with the data source.

SAP HANA extended application services, advanced model, is a polyglot application platform that supports several programming languages and execution environments, for example, Java and Node.js. The JavaScript for SAP HANA extended application services, classic model, is supported by a framework running in the Node.js runtime.

In simple terms, SAP HANA extended application services, advanced model, is basically the Cloud Foundry open-source Platform-as-a-Service (PaaS) with a number of tweaks and extensions provided by SAP. These SAP enhancements include the following:

- Integration with the SAP HANA database
- OData support
- Compatibility with SAP HANA extended application services, classic model
- Additional features designed to improve application security

SAP HANA extended application services, advanced model, also provides support for business applications that are composed of multiple micro-services. These are implemented as separate Cloud Foundry applications, which combined are also known as Multitarget Applications (MTA). A multitarget application includes multiple modules, which are the equivalent of Cloud Foundry applications.

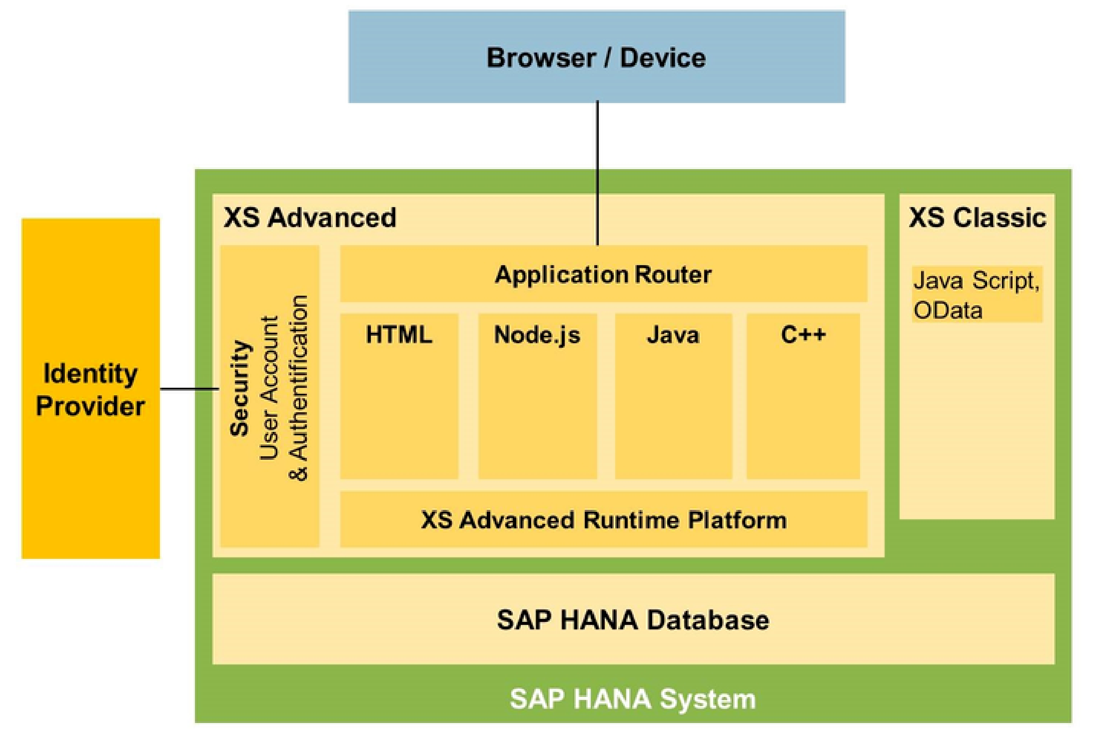

Figure 104: SAP HANA XS Advanced.pptx

Rules for Developers

For developers considering building applications that run on SAP HANA extended application services, advanced model, the following general rules apply:

- SAP HANA XSA applications

  If you want to develop new applications, for example, Java or JavaScript (including Node.js), which run solely on SAP HANA extended application services, advanced model, consider SAP HANA native development. These new advanced applications are multitarget applications (MTA), which comprise multiple modules (software components) that all share a common lifecycle for development and deployment.

- Integrated applications

  To enrich existing applications (for example, SAP HANA Live, SAP Fiori), combine SAP HANA native development with ABAP.

- SAP Business Suite or SAP BW

In the context of SAP Business Suite or SAP Business Warehouse (SAP BW), use ABAP if the scope is simply to optimize existing programs.

SAP HANA Platform

The SAP HANA platform combines database, data processing, and application platform capabilities. It provides libraries for predictive planning, text, spatial, and business analytics to enable business to operate in real time.

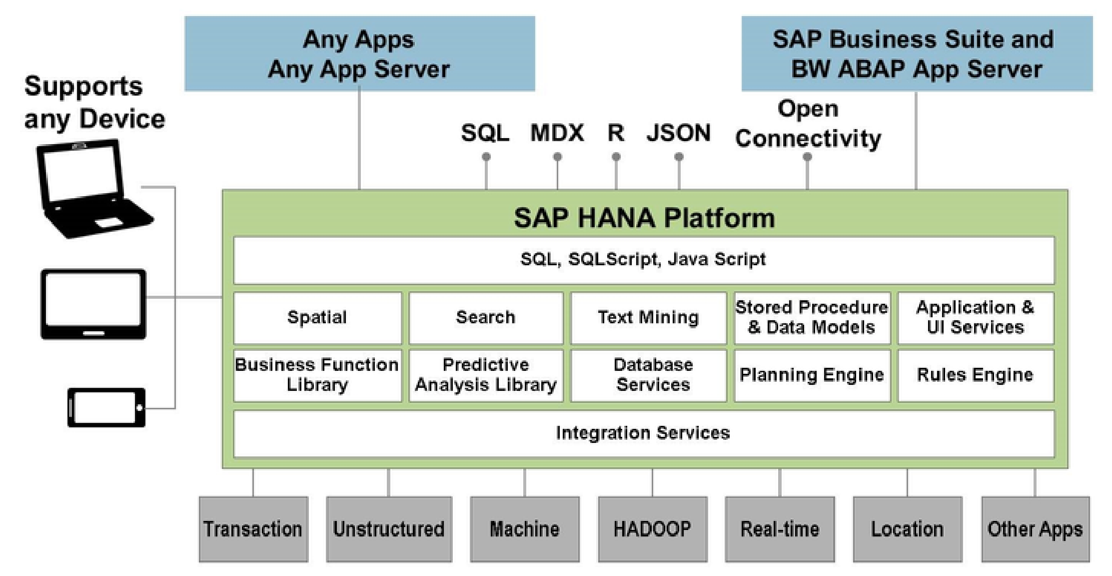

Figure 105: SAP HANA: A Platform to Innovate Enterprise Applications

SAP HANA is an in-memory database management system, but it also contains many additional features for specific use cases. Examples are spatial processing, search, and text mining, and integrated libraries. Some of these features can be used when running traditional applications on SAP HANA. Others are used in entirely new in-memory applications.

These features enable new scenarios and use cases. Running traditional applications on SAP HANA already provides significant advantages compared to traditional disk-based DBMS. By adapting applications to the innovative data model and architecture of SAP HANA, the advantages are even more comprehensive and enable entirely new business scenarios

Customers use SAP HANAin different scenarios. In addition to the optimization potential, the way in which SAP HANA is integrated into the system landscape also impacts aspects like system architecture, administration, operations, and security. Therefore, it is essential to include all stakeholders in the scenario discussion.

LESSON SUMMARY

You should now be able to:

Unit 6
Lesson 2

Identifying Deployment Options

LESSON OVERVIEW

Depending on the requirements for productive and nonproductive usage, various deployment options for SAP HANA exist. These are explained in detail in this lesson.

Business Example

At the time of its market introduction, SAP HANA only offered following an appliance model as a certified combination of hardware and software that could be deployed as an on-premise solution. Meanwhile, SAP is continuously working on increasing the flexibility and choice of deployment options for SAP HANA. For customers, it is essential to understand which deployment options exist, what their capabilities and limitations are, and which scenarios can be combined and run together on one SAP HANA server or database.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Talk about solution packages
- Explain the different deployment options for SAP HANA

SAP HANA Components

The SAP HANA platform is composed of several components.

SAP HANA, platform edition, is the technical foundation of the SAP HANA platform and various SAP HANA editions. The following figure shows the components of SAP HANA, platform edition.

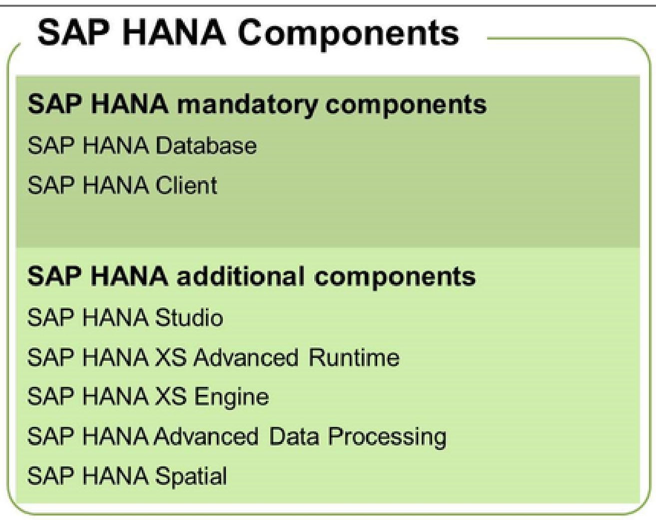

Figure 106: SAP HANA Components

SAP HANA system components can be installed, updated, or uninstalled using the SAP HANA database lifecycle manager (HDBLCM).

You can install the following additional SAP HANA system components:

- SAP HANA client
- SAP HANA Studio
- SAP HANA extended application services runtime
- SAP HANA extended application services engine
- SAP HANA Advanced Data Processing
- SAP HANA, spatial edition

.

The SAP HANA system is made up of the following mandatory components:

- SAP HANA server

  The SAP HANA database software is installed on the Linux operating system on certified hardware.

- SAP HANA client

  SAP HANA client software is required for connecting to the S AP HANA database. Versions exist for AIX, HP-UX, Linux, Microsoft Windows, and Solaris.

SAP HANA Additional Components

The SAP HANA system also has the following components:

- SAP HANA Studio

  The SAP HANA Studio is a collection of applications for the SAP HANA appliance software. It enables technical users to manage the SAP HANA database, to create and manage user authorizations, and to create new or modify existing models of data in the SAP HANA database.

- SAP HANA extended application services, advanced model, runtime

  SAP HANA includes an additional, new runtime environment for application development: SAP HANA extended application services, advanced model. SAP HANA extended application services, advanced model, represents an evolution of the application server architecture within SAP HANA by building upon the strengths (and expanding the scope) of SAP HANA extended application services, classic model.

- SAP HANA extended application services engine

  Native SAP HANA applications can be implemented in SAP HANA, without requiring an additional application server based on SAP HANA extended application services.

- SAP HANA Advanced Data Processing

  SAP HANA provides text mining functions for analyzing documents.

- SAP HANA , spatial edition

  SAP HANA supports the processing of spatial data for analyzing sales data.

SAP HANA Options

SAP HANA features, SAP HANA capabilities, and SAP HANA options provide additional functions. To use the SAP HANA options and SAP HANA capabilities, you need a dedicated license for the options or capabilities that you want to use (see the following disclaimer). The following figure shows some of the SAP HANA options and SAP HANA capabilities.

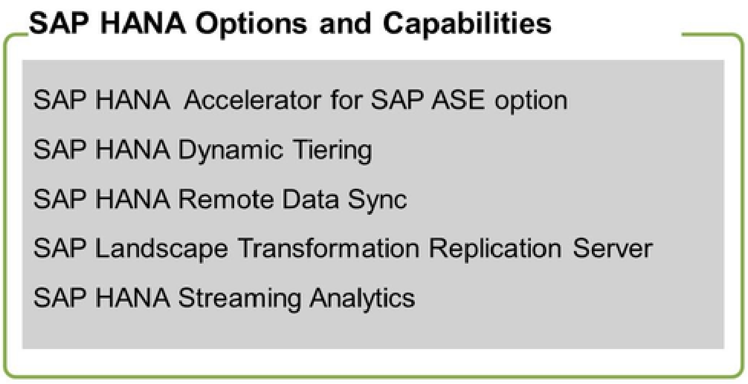

Figure 107: SAP HANA Options

Several extensions can be installed or added to the SAP HANA server, including the following:

- SAP HANA accelerator for SAP Adaptive Server Enterprise (SAP ASE) option
- SAP HANA dynamic tearing
- SAP HANA remote data sync
- SAP HANA landscape transformation replication server
- SAP HANA streaming analytics

> Note:
> For information about the availability of the SAP HANA features, SAP HANA capabilities, SAP HANA options on Intel-based hardware platforms or on IBM Power servers, see SAP HANA Hardware and Software Requirements in the SAP HANA Master Guide .
>
> The documentation for the SAP HANA optional components is available in SAP Help Portal at http://help.sap.com/hana_options .

> Caution:
> You need additional licenses for SAP HANA options. Contact your SAP sales representative for details.

SAP HANA Accelerator for SAP Adaptive Server Enterprise (SAP ASE) Option

The accelerator for SAP ASE adds analytics acceleration to the SAP ASE database engine using SAP HANA. SAP ASE users can run reports in SAP HANA using the data in SAP ASE for real time analytics. Users canb achieve this process by either replicating the data from SAP ASE to SAP HANA, or by creating virtual tables in SAP HANA that access SAP ASE data.

You can also use the SAP HANA accelerator for SAP ASE to accelerate SAP ASE stored procedures (not OTLP applications) by pushing down the stored procedure execution to SAP HANA. Minimal or no code changes to the existing stored procedures are needed. The stored procedures continue to execute against the SAP ASE reporting server with the execution being pushed to SAP HANA. The results are brought back to SAP ASE and then sent to the client SAP ASE application.

SAP HANA Dynamic Tearing and SAP HANA Remote Data Sync

SAP HANA dynamic tiering is a native big data solution for SAP HANA. Dynamic tiering adds smart, disk-based extended storage to your SAP HANA database. It enhances SAP HANA with large volume, warm data management capability.

By using dynamic tiering to place hot data in SAP HANAin-memory tables, and warm data in extended tables, the highest value data remains in-memory, and cooler less-valuable data is saved in the extended store. This can reduce the size of your in-memory database.

In remote data synchronization, many clients synchronize through the remote data sync server to central data sources.

SAP HANA Landscape Transformation Replication Server

The SAP Landscape Transformation Replication Server allows you to load and replicate data in real time from ABAP source systems and non-ABAP source systems to an SAP HANA environment.

The SAP Landscape Transformation Replication Server uses a trigger-based replication approach to pass data from the source system to the target system.

SAP HANA Streaming Analytics

The SAP HANA streaming analytics option processes high-velocity, high-volume event streams in real time. It allows you to filter, aggregate, and enrich raw data before committing it to your database.

With SAP HANA streaming analytics, you can accept data input from various sources, including data feeds, business applications, sensors, IT monitoring infrastructure, and so on. You can also apply business logic and analysis to the streaming data, and store your results directly in SAP HANA.

SAP HANA Deployment Options

SAP HANA can be deployed in different ways. Here is an overview of all available options offered by SAP today.

With SAP HANA deployment options, you can distinguish between a preconfigured on-premise appliance, a cloud deployment, and a hybrid model that combines cloud and on-premise instances. For each of the deployment options, various solutions exist, as follows.

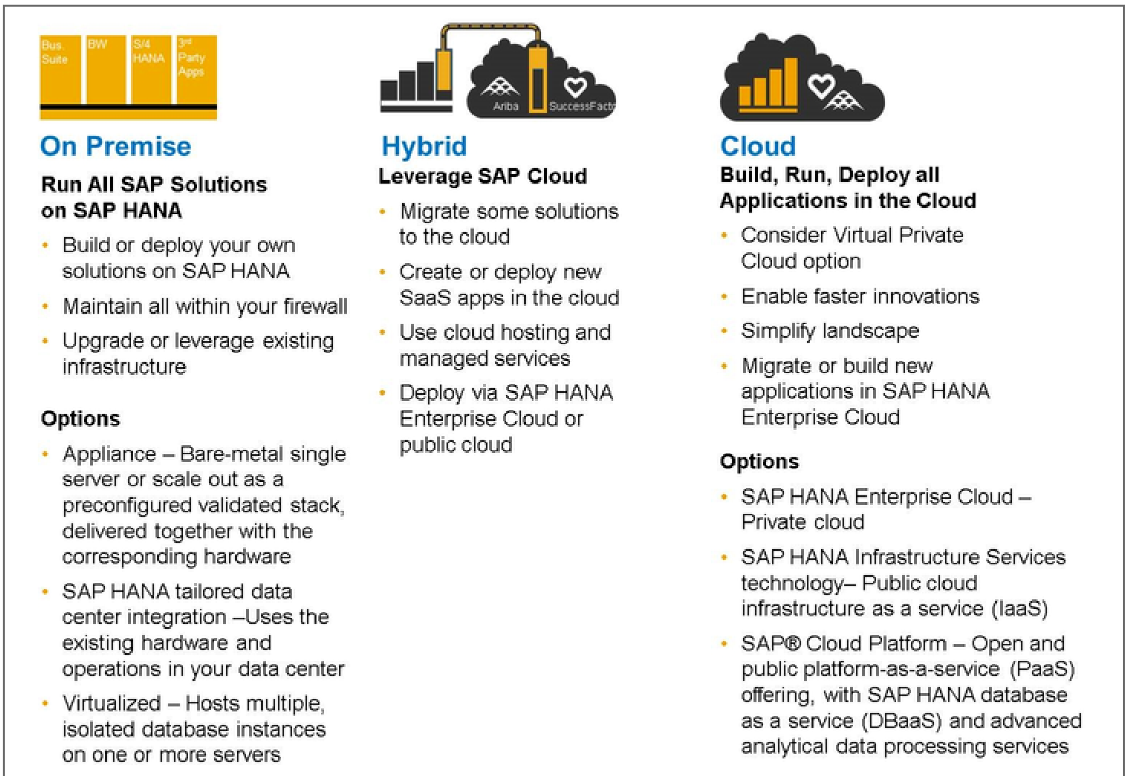

Figure 108: SAP HANA Deployment Options

SAP HANA On-Premise

A typical single-server deployment of SAP HANA can vary from a 2 CPU configuration with 128 GB of RAM as an entry-level system (see the certified and supported SAP HANA Hardware Directory for a complete list of all available entry-level configurations provided by SAP technology partners) to a high-end 8 CPU configuration with 8 TB of RAM (see Figure below). Special layouts for SAP S/4HANA support up to 20 TB per host and more.

If your data volume exceeds the volume applicable for single-server deployments, as detailed earlier, a scale-out configuration is the configuration of choice. A typical scale-out cluster consists of 2 to n servers per cluster, where in theory there is no limit for n. The largest certified configuration is 112 servers, while the largest tested configuration has more than 250 servers. Each server configuration is either 4 CPU/2 TB or 8 CPU/4 TB. Both architectures, scale-up and scale-out, provide full support for high availability and disaster recovery.

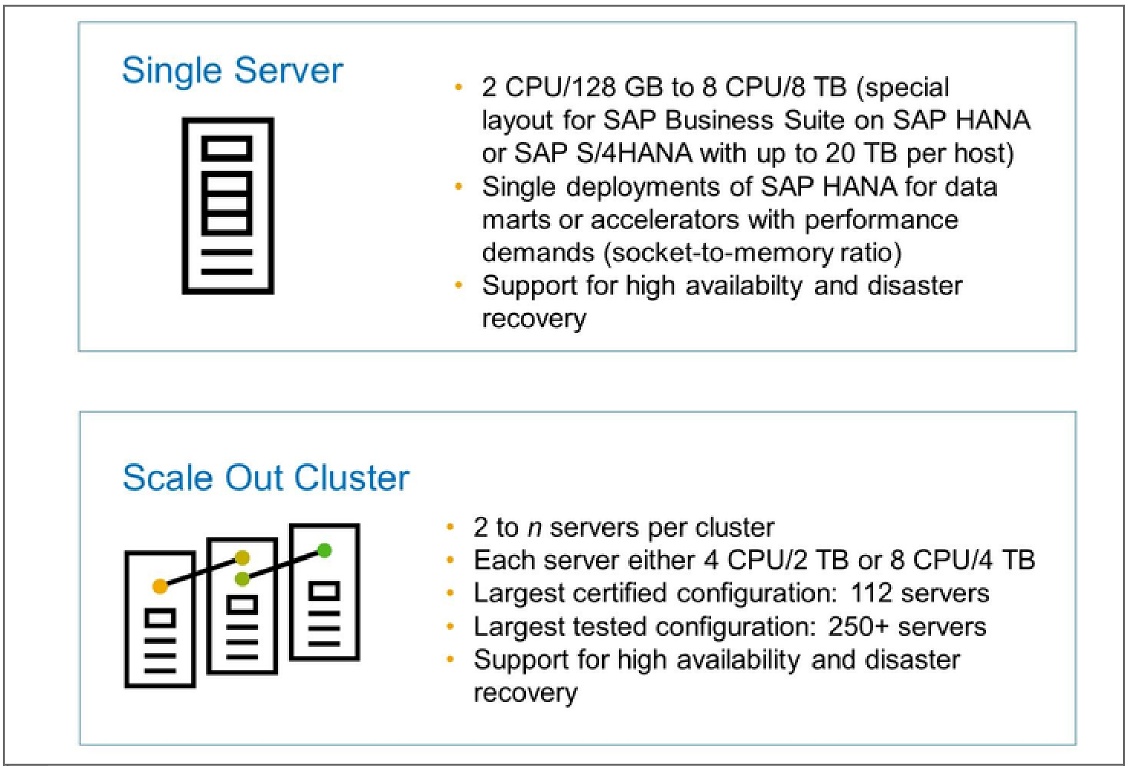

Figure 109: SAP HANA On-Premise

SAP HANA Tailored Data Center Integration Option: Overview

In an on-premise deployment, SAP HANA runs on dedicated hardware.

On-premise SAP HANA is deployed through the following offerings:

- As an appliance, SAP HANA combines software components from SAP optimized on proven hardware provided by hardware partners. This approach is valid for Intel-based hardware platforms only. While this approach is simple, it has limitations for hardware flexibility and compliance with existing IT operation processes. Therefore, SAP HANA tailored data center integration is offered as a new option to provide customers with greater flexibility.

- Compared with the appliance delivery approach, SAP HANA tailored data center integration is a more open and flexible approach for the integration of SAP HANA in the data center. The requirements for this deployment option are as follows:

  - The server is certified and belongs to the allowed hardware.

  - The storage solution has successfully passed SAP HANA hardware certification.

  - The components of SAP HANA can only be installed by certified hardware partners, or any person holding a certification, on validated hardware running an approved operating system.

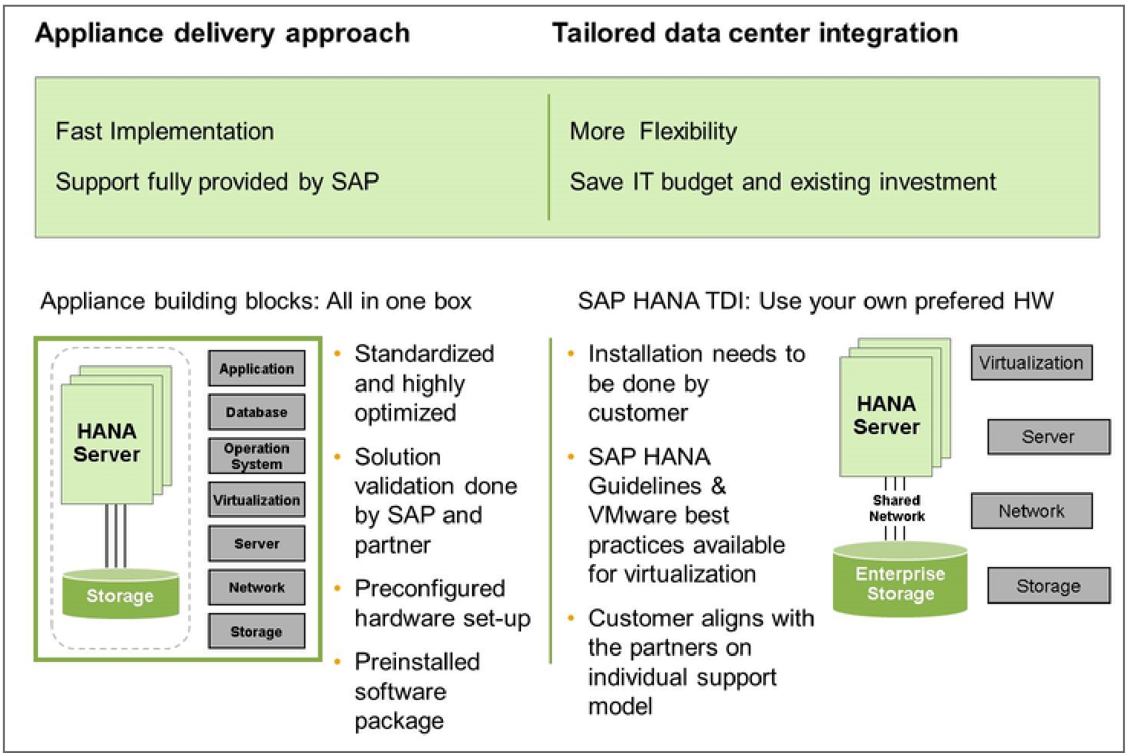

Figure 110: SAP HANA Tailored Data Center Integration Option: Overview

Tailored data center integration can reduce hardware and operations cost by reusing existing hardware components and processes.

Note:
Tailored data center integrations offers freedom and flexibility, which also leads to the customer's increased responsibility for the system, ranging from the installation up to running the landscape.

SAP HANA Tailored Data center Integration

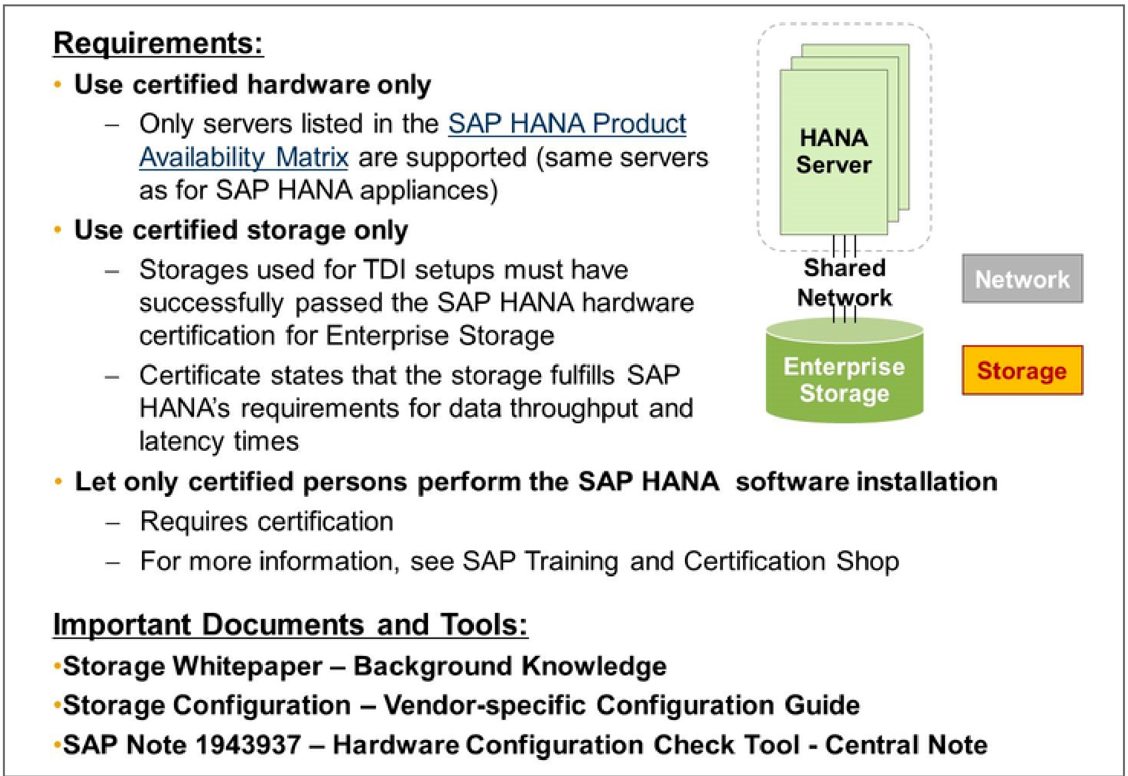

Figure 111: SAP HANA Tailored Data center Integration with Enterprise Storage

In the first wave, all storage that successfully passes the hardware certification can be used in combination with servers listed in the Product Availability Matrix. For details, see the Partner Information Center at http://www.sap.com/partners/overview.html

The following documents and tools provide information about SAP HANA tailored data center integration with enterprise storage:

- **Storage Whitepaper: Background Knowledge**
  - Conceptual storage layout, non-shared storage versus shared storage
  - Storage sizing formulas
  - Helps you to understand the impact of SAP data throughput KPIs in the daily operation of SAP HANA
  - Explains the high-availability support (Host Auto Failover) of SAP HANA
  - Storage requirements:: https://www.sap.com/documents/2015/03/74cdb554-5a7c-0010-82c7-eda71af511fa.html
- **Storage Configuration: Vendor-Specific Configuration Guide**
  - For certification, each storage vendor must file in a configuration guide
  - Explains how to configure the storage for optimal collaboration with SAP HANA
  - Get a copy directly from your storage vendor
- **SAP HANA HW Config Check Tool (HWCCT)**

- Command-line tool
- Used by storage vendors, SAP Support, and customers
- Measures the data throughput and latency times between the SAP HANA servers and the Enterprise Storage system
- Download it from SAP Support Portal
- Documented in PDF attachment of SAP Note [1943937](1943937)

Read the SAP Storage Whitepaper carefully, and get a copy of your storage vendor's configuration guide for SAP HANA. Make sure that you run the SAP HANA HW Config Check Tool to check your storage KPIs every time you change your storage configuration.

> Note:
> Customers should consider involving SAP Digital Business Services to perform a HANA Go-Live Check prior to going productive

SAP HANA Tailored Data Center Integration with Enterprise Network

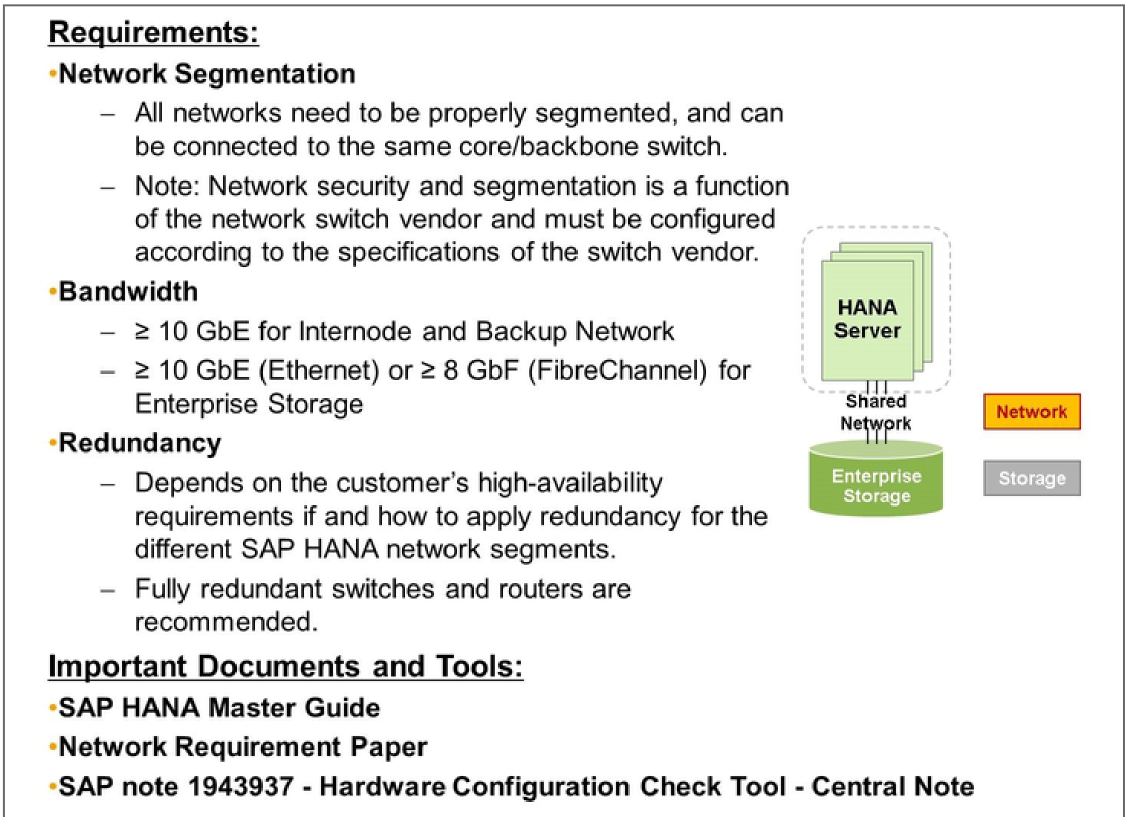

Figure 112: SAP HANA Tailored Data Center Integration with Enterprise Network

With the introduction of tailored data center integration with Enterprise Network, SAP supports hardware setups, which comply with the prerequisites mentioned in the figure, SAP HANA in the Cloud: Overview . Apart from that, no further approval by SAP is required. SAP does not introduce any certification of network components for tailored data center integration setups. Customers should consider involving SAP Active Global Support to perform an SAP HANA Go-Live Check prior to going productive.

SAP HANA in the Cloud

Several options are offered by SAP and partners to run SAP HANA in the cloud. SAP HANA is more than just an in-memory database, but rather a complete application development platform in its own right. Given this unique characteristic and the wide range of usage scenarios, three distinct packages that are logically based on each other are available.

Private Cloud, Managed Service	Public Cloud, Infrastructure-as-a-Service	Public Cloud, Pay-as-you-go
Deploy in a fully managed, secure cloud – with SAP HANA Enterprise Cloud. • **Reliable hosting environment** • **Managed services approach** • **Increased flexibility and low TCO**	Bring your own SAP HANA license to run on Amazon Web Services (AWS), Google Cloud Platform, Microsoft Azure, or IBM clouds – via SAP HANA Infrastructure Services. • **Fast time to value** • **No hardware investment** • **Flexible scalability**	Deploy SAP HANA One in third-party clouds: Amazon Web Services (AWS), or Alibaba. • **Hourly pay-as-you-go rate** • **Single database instance** • **Run custom applications only**

Figure 113: SAP HANA in the Cloud: Overview

Options for SAP HANA in the cloud

- **Private Cloud, Managed Service**

  SAP offers a secure, scalable, private cloud offering called SAP HANA Enterprise Cloud . It includes an SAP HANA software license, underlying cloud infrastructure, and SAP-managed services.

  SAP HANA Enterprise Cloud is a fully scalable and secure private cloud offering available only from SAP. It gives you the full power of SAP HANA in a private, managed cloud environment. The benefits include simplicity through rapid deployment, an integrated support model, and a comprehensive portfolio of innovative cloud solutions. Ideal for mission-critical applications – such as SAP Business Suite , SAP Business Warehouse , and custom SAP HANA applications – SAP HANA Enterprise Cloud is available at affordable, subscription-based pricing. It can be deployed for productive use or project acceleration.

- **Public Cloud, Infrastructure-as-a-Service (IaaS)**

  SAP HANA Infrastructure Services is a public cloud infrastructure that enables customers to quickly deploy and manage their prelicensed instance of SAP HANA without a hardware investment. Bring your own SAP HANA license to run on third-party public cloud providers: Amazon Web Services, Google Cloud Platform, IBM Bluemix Cloud Platform and Microsoft Azure. Supported scenarios include data marts, data warehouses, custom-built applications, and SAP applications –including SAP S/4HANA , SAP BW/4HANA and SAP Business One . Configurations ranging from 128 GB to 3 TB are available plus scale-out.

- **Public Cloud, Pay-as-you-go**

SAP HANA One is a public cloud pay-as-you-go offering. Ideal for department-scale projects, system integration, independent software vendors, and innovative startups.

SAP Cloud Platform

SAP Cloud Platform is an open and public PaaS offering that includes SAP HANA as a DBaaS, enabling developers to use advanced data stored in SAP HANA to build, extend, and integrate applications without investing in separate infrastructure or tools. SAP Cloud Platform provides a variety of innovative features and services to enrich your application with advanced analytics, the Internet of Things, and security capabilities and to integrate your cloud and on-premise applications.

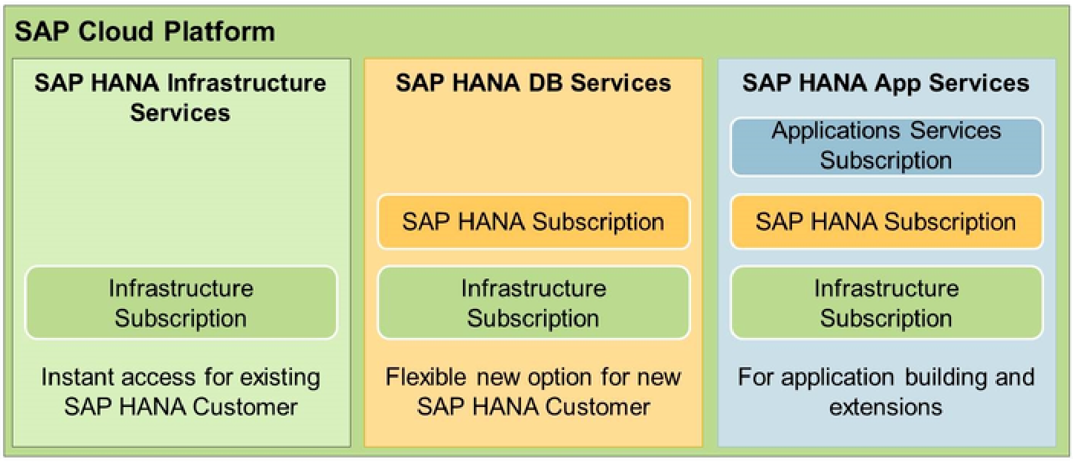

Figure 114: SAP Cloud Plattform

Business Cases for SAP HANA in the Cloud

When to choose which of these cloud offerings?

The following are reasons you may choose the **SAP HANA Enterprise Cloud** :

- You are re a net-new SAP HANA customer who does not want to set up a new or extend an existing data center
- You lack qualified resources for operations and maintenance of SAP HANA or don’t want to invest in them
- You want to lower software licensing costs or avoid high up-front investments
- You want to extend your on-premise landscape with innovative applications delivered through SAP HANA Enterprise Cloud

The following are reasons why you might choose **SAP HANA Infrastructure Services** and run SAP HANA in a public cloud:

- You have your own SAP HANA licenses but want to leverage cloud infrastructure for flexible scalability
- You want to reduce your hardware footprint

You may choose the **SAP HANA One** service if you pay for use to perform specific tasks that require one-time high-performance capabilities (for example, monthly or year-end close).

You may choose **SAP Cloud Platform** for the following reasons:

- You want to quickly build, extend, and integrate applications in a complete, cloud-based development environment based on SAP HANA
- You want to use a hybrid approach in which on premise and cloud deployments are combined. In such cases, specific requirements for data and process integration apply. The cloud integration service for SAP Cloud Platform supports these tasks by providing process and data integration technology as a service to connect your cloud and on-premise enterprise systems in real time

SAP HANA, express edition

SAP HANA, express edition is a streamlined version of SAP HANA that can run on laptops and other resource-constrained hosts, such as a cloud-hosted virtual machine. SAP HANA , express edition is free to use for in-memory databases up to 32GB of RAM.

Memory capacity increases beyond 32GB are available for purchase at the SAP Store .

SAP HANA , express edition is available for download from the SAP developer center free of charge: https://www.sap.com/developer/topics/sap-hana-express.html . This version of SAP HANA can be installed on laptop or desktop computers, Linux-based servers, virtual machines, or a private cloud.

Technical Deployment Options

The technical deployment options determine how SAP HANA systems, the hosts used for SAP HANA systems, and applications running on SAP HANA are deployed.

To run multiple scenarios on one system or database, you need to understand the availability and capabilities of the technical deployment options.

Technical Deployment Options

- Multitenant Database Container (MDC)
- Multiple Components on one Database (MCOD)
- Multiple Components on one System (MCOS)
- SAP HANA with Virtualization
- Technical Co-Deployment

Multitenant Database Containers

SAP HANA multitenant database containers establish a foundation for providing multitenancy in SAP HANA.

As of SAP HANA 2.0 SPS01 an SAP HANA system is installed in multiple-container mode. It can contain more than one multitenant database container.

A multiple-container system always has exactly one system database, which is used for central system administration. It also has any number of multitenant database containers (including zero), called tenant databases. An SAP HANA system installed in multiple-container mode is identified by a single system ID (SID). Databases are identified by a SID and a database name. From the administration perspective, there is a distinction between tasks performed at system level and those performed at database level. Database clients, such as the SAP HANA Studio , connect to specific databases.

All the databases in a multiple-container system share the same installation of database system software, the same computing resources, and the same system administration. However, each database is self-contained and fully isolated with its own, as follows:

- Set of database users
- Database catalog
- Repository
- Persistence
- Backups
- Traces and logs

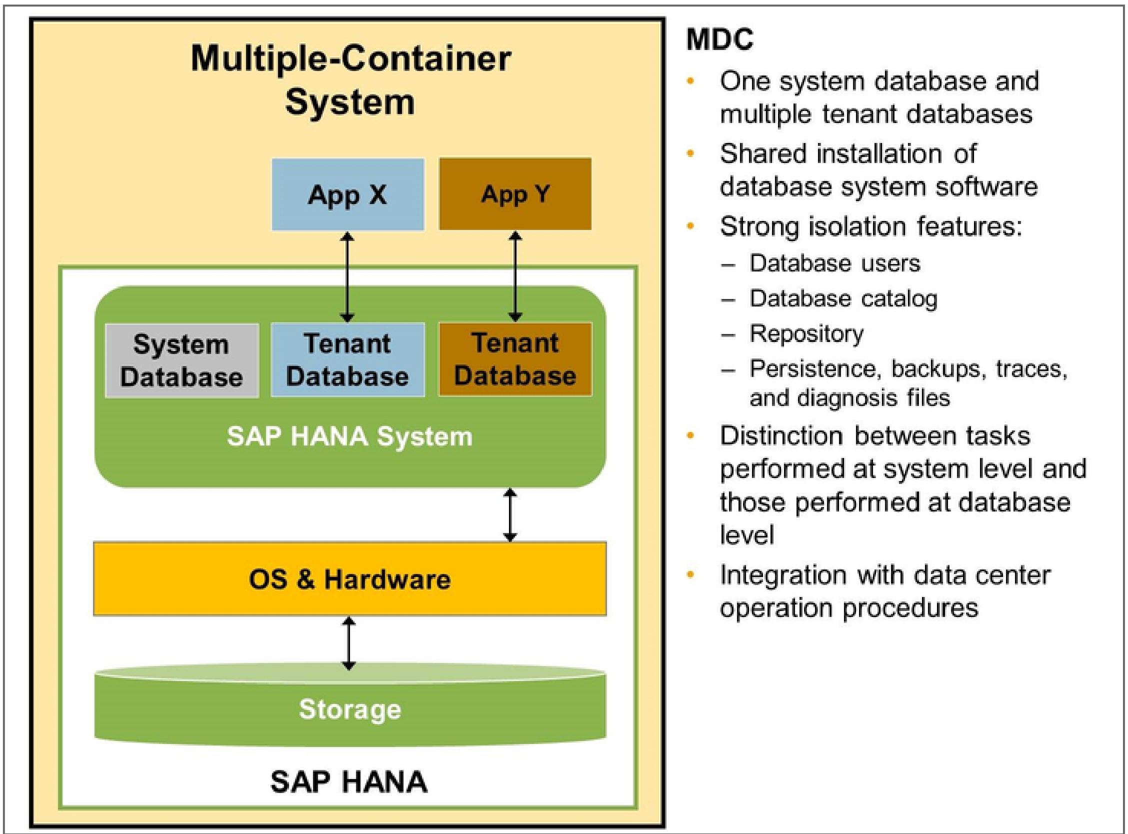

Figure 115: Multitenant Database Containers: Overview

Although database objects such as schemas, tables, views, procedures, and so on, are local to the database, cross-database SELECT queries are possible. This supports cross-application reporting, for example.

SAP HANA multitenant database containers allow you to manage several databases in one database management system. Note that a single database container is also called a tenant database.

Concept and Terminology

The concept and terminology of SAP HANA multitenant database containers are as follows:

Concept and Terminology of SAP HANA Multitenant Database Containers

- Run multiple tenant databases on one SAP HANA system
- Run and support multiple applications or scenarios on one SAP HANA system in production

- Strong separation of data and users
- Backup and restore available by tenant database
- Resource management by tenant (CPU, Memory)
- Move or copy tenant databases or applications to different hosts or systems
- Integration with existing data center operations procedures

Multiple Components on One Database (MCOD)

Multiple applications on one SAP HANA system are also known as Multiple Components on One Database (MCOD).

The technical deployment type MCOD refers to the scenario where more than one application, scenario, or component runs on one SAP HANA system. This deployment type is available, with restrictions, for production SAP HANA systems. For production systems, there are white lists that specify supported scenarios explicitly.

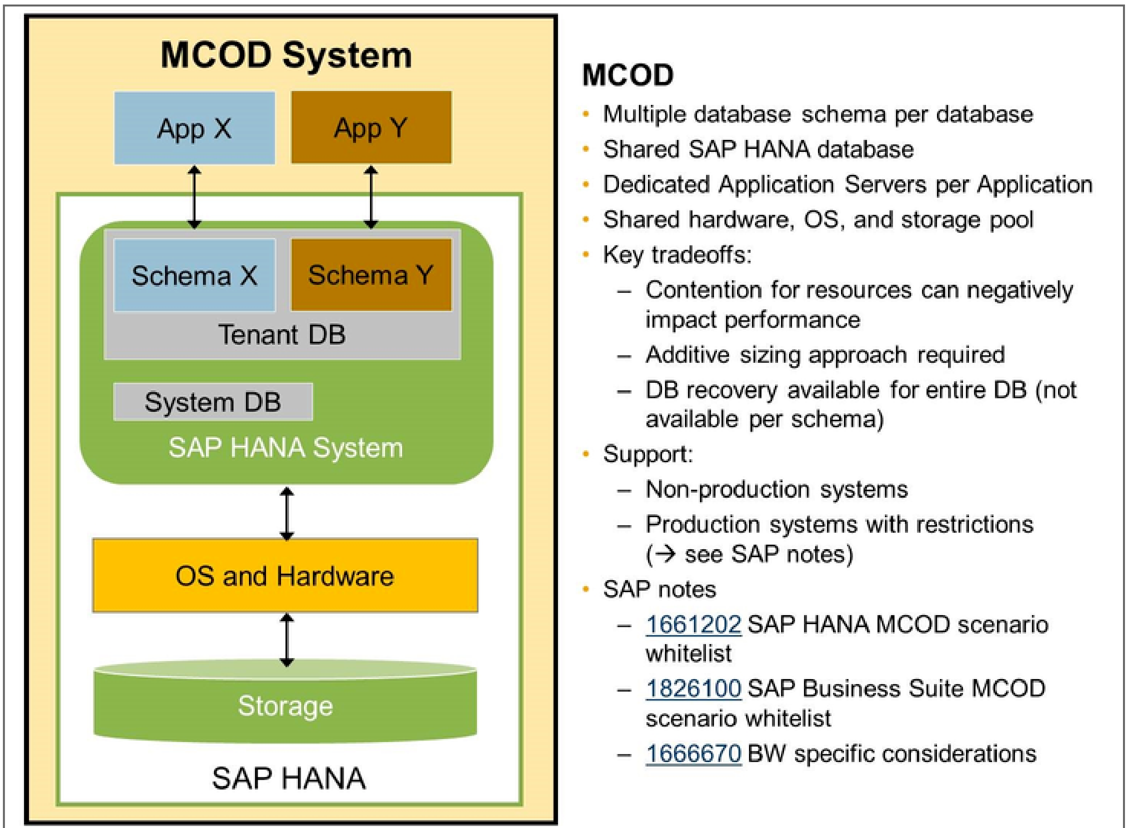

Figure 116: Multiple Components on One Database (MCOD)

Multiple Components on One System (MCOS)

Multiple SAP HANA systems on one host are also known as Multiple Components on One System (MCOS).

SAP does support running multiple SAP HANA systems (SIDs) on a single production SAP HANA host. This is restricted to single-host or scale-up scenarios only. Multiple components on one system requires significant attention to various detailed tasks related to system administration and performance management.

Production support is restricted to SAP HANA SPS 09 or higher because of the availability of some resource management parameters (for example affinity). Running multi-SID on one SAP HANA host can impact the performance of various types of operations, because competition for computing resources might occur (memory, cpu, i/o, and so on).

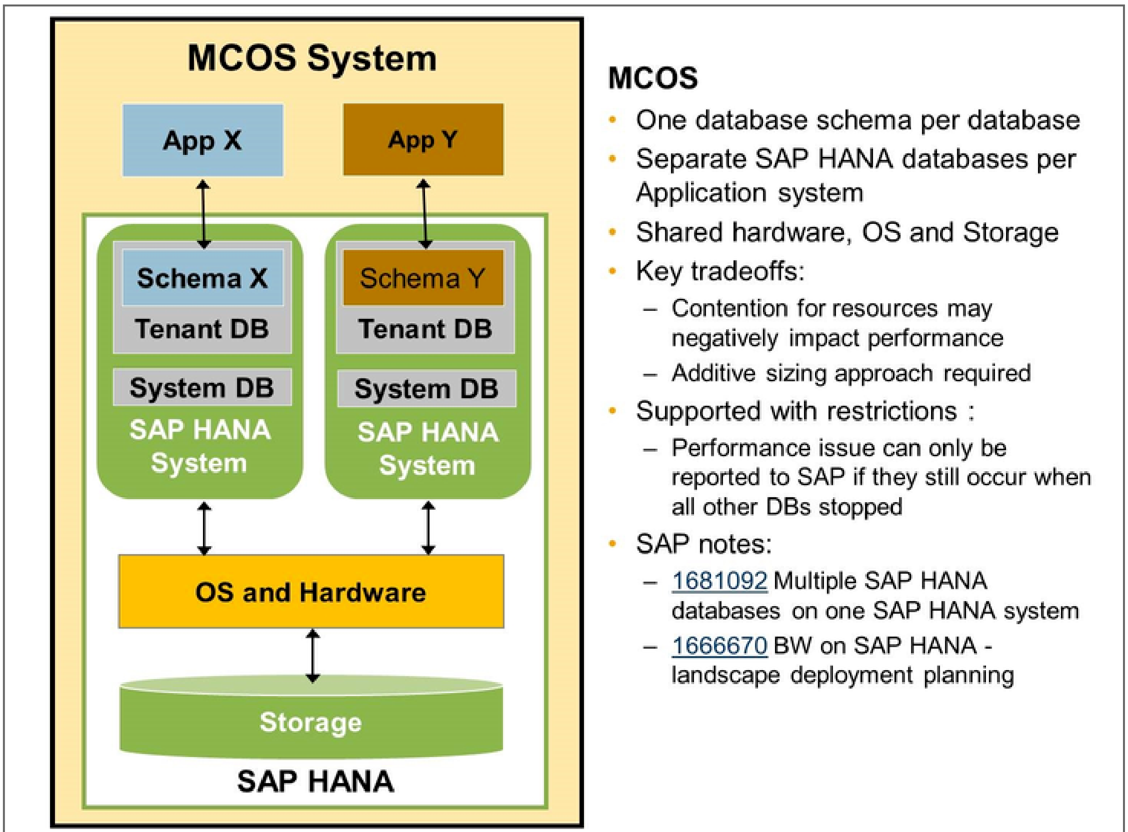

Figure 117: Multiple Components on One System (MCOS)

Perform the requisite testing in every project before going live. In general, stress or volume testing provides good indicators of the expected performance. When operating a system that features a multi-SID deployment, use the resource management features of SAP HANA (for example, parameters controlling memory limits, and influencing utilization of CPU cores, and so on) to optimize performance.

Virtualization

SAP HANA supports many popular virtualization supervisors, such as VMware vSphere 5.1 or newer, Hitachi LPAR, Huawei FusionSphere, IBM PowerVM, or others for nonproductive environments. You can choose virtualization for a wide variety of configurations, for single and multiple virtual machines, in single- or multinode configurations, for appliances of SAP HANA and SAP HANA tailored data center integration delivery methods.

The technical deployment type, SAP HANAw ith virtualization, refers to the scenario where one or more SAP HANA database SIDs are deployed on one or more virtual machines running on SAP HANA server hardware.

One benefit of virtualization is that you can assign dedicated CPU and memory resources to specific databases and, increase the flexibility of hardware usage.

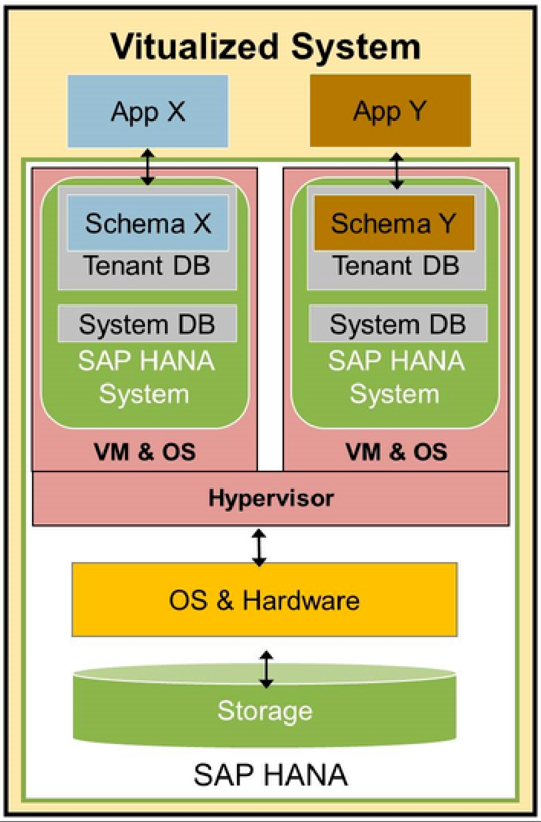

Virtualization

- One database schema per database
- Separate SAP HANA databases per SAP system
- Separate virtual machine and OS
- Shared hardware and storage
- Details about the supported scenarios:
  - SAP Note 1788665 – SAP HANA Support for virtualized / partitioned (multi-tenant) environments

Figure 118: Virtualization

For customers already standardizing on virtualization technology, SAP HANA offers the customer total cost of ownership reductions and additional options for planning and managing their systems landscapes, as follows:

- Ease of HW replacement and avoidance of recertification of OS and SAP installations
- Separation of IT Ownership (HW and SW layer)
- OS independent monitoring
- Low-cost HA capabilities in Dev and Test environments
- Positive impact on capital expenditures

Recommendations on Virtualizing SAP HANA Deployments

Depending on performance requirements, number of users, and technical parameters, among other criteria, there are best practices in which we recommend virtualizing a deployment of SAP HANA and others where we do not recommend this option. See Figure below for an overview.

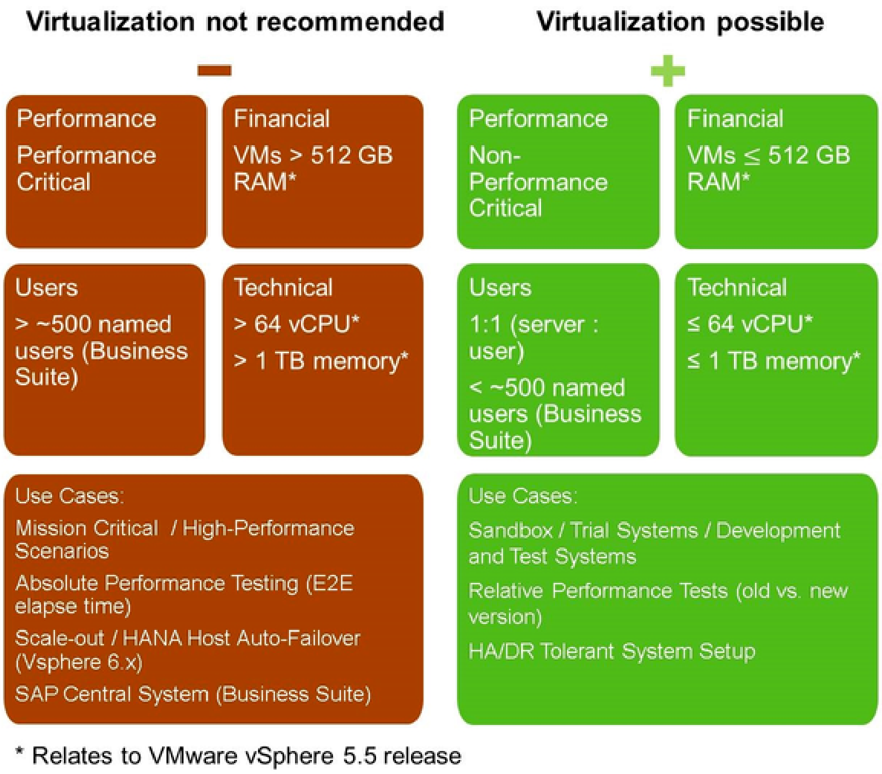

Figure 119: Use Cases for SAP HANA Virtualization

Technical Co-Deployment

Technical co-deployment is an additional alternative that can be used to combine several applications. This is available for Supplier Relationship Management (SRM) and Supply Chain Management (SCM). It is provided as an SAP ERP add-on and can be used productively.

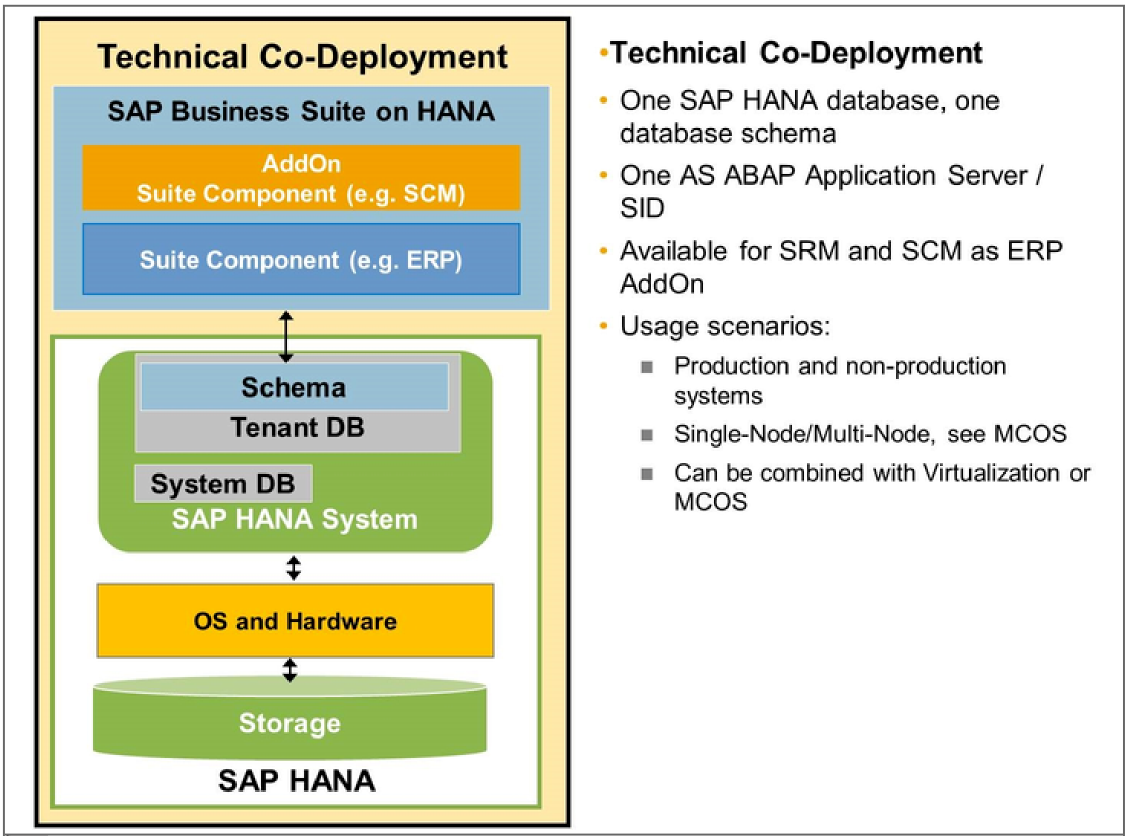

Figure 120: Technical Co-Deployment

Related Information

For additional information on SAP HANA deployment options, see the following SAP Notes:

- SAP Note 1666670 : Multiple SAP HANAdatabases on one appliance
- SAP Note 1661202 : Support for multiple applications on SAP HANA
- SAP Note 1826100 : Multiple applications SAP Business Suite powered by SAP HANA
- SAP Note 1681092 : SAP BW/4HANA - landscape deployment planning
- SAP Note 1788665 : SAP HANA running on VMware vSphere VMs

LESSON SUMMARY

You should now be able to:

- Talk about solution packages
- Explain the different deployment options for SAP HANA

Unit 6

Learning Assessment

1. Using the Database Migration Option (DMO) of the Software Update Manager (SUM) comes along with benefits of simplifying the migration steps by combining the system update and database migration into one tool and reducing the business downtime.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

2. Which of the following elements are SAP HANA system components?

   Choose the correct answers.

   - [ ] **A** SAP HANA Database
   - [ ] **B** SAP HANA Dynamic Tiering
   - [ ] **C** SAP HANA Advanced Data Processing
   - [ ] **D** SAP HANA Remote Data Sync

3. SAP HANA tailored data center integration is an open and flexible approach for the integration of SAP HANA in the data center. Tailored data center integration can reduce hardware and operations cost by reusing existing hardware components and processes.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

Unit 6

Learning Assessment - Answers

1. Using the Database Migration Option (DMO) of the Software Update Manager (SUM) comes along with benefits of simplifying the migration steps by combining the system update and database migration into one tool and reducing the business downtime.

   Determine whether this statement is true or false.

   [X] True

   [ ] False

   Correct! The database migration option simplifies the migration and is referred to as the one-step procedure to SAP HANA. The alternative way to migrate to SAP HANA is the classical way of a heterogeneous system copy using Software Provisioning Manager (SWPM). Read more on this in the lesson Describing SAP HANA Roadmap and Scenarios (Unit 6, Lesson 1) of the course HA200_14.

2. Which of the following elements are SAP HANA system components?

   Choose the correct answers.

   [X] **A** SAP HANA Database

   [ ] **B** SAP HANA Dynamic Tiering

   [X] **C** SAP HANA Advanced Data Processing

   [ ] **D** SAP HANA Remote Data Sync

   Correct! The SAP HANA database software is a mandatory component and installed on the Linux operating system on certified hardware. SAP HANA Advanced Data Processing is an additional component which provides text mining functions for analyzing documents. SAP HANA dynamic tiering is an SAP HANA Option, which provides a native big data solution for SAP HANA. Dynamic tiering adds smart, disk-based extended storage to your SAP HANA database. It enhances SAP HANA with large volume, warm data management capability. In remote data synchronization, which is a SAP HANA Option, many clients synchronize through the remote data sync server to central data sources. Read more on this in the lesson Identifying Deployment Options (Unit 6, Lesson 2) of the course HA200_14.

3. SAP HANA tailored data center integration is an open and flexible approach for the integration of SAP HANA in the data center. Tailored data center integration can reduce hardware and operations cost by reusing existing hardware components and processes.

   Determine whether this statement is true or false.

   - [x] True
   - [ ] False

   Correct! SAP HANA tailored data center integration is offered as a new option to provide customers a greater flexibility, but underlies some requirements to run this deployment option: The storage solution has successfully passed SAP HANA hardware certification. The server is certified and belongs to the allowed hardware. The components of SAP HANA can only be installed by certified hardware partners, or any person holding a certification, on validated hardware running an approved operating system. This approach is valid for Intel-based hardware platforms only. While this approach is simple, it has limitations for hardware flexibility and compliance with existing IT operation processes. Read more on this in the lesson Identifying Deployment Options (Unit 6, Lesson 2) of the course HA200_14.

UNIT 7 Post Installation Tasks

UNIT OBJECTIVES

- Perform the post-installation steps after a SAP HANA installation
- Update a SAP HANA database using the HDBLCM tool
- Explain the resident HDBLCM functions
- Explain the installation and use of SAP HANA InteractiveEducation (SHINE)
- Understand the SAP HANA maintenance strategy

Unit 7
Lesson 1

Performing Post-Installation Steps

LESSON OVERVIEW
In this lesson, you learn what to do after the SAP HANA installation.

Business Example

As part of the initial setup you have to establish SAP Solution Manager connectivity and configure a Remote Service Connection (via SAP Router).

LESSON OBJECTIVES
After completing this lesson, you will be able to:

- Perform the post-installation steps after a SAP HANA installation

Solution Manager Connectivity

In addition to running the on-site configuration tool, you can establish SAP Solution Manager connectivity and configure a Remote Service Connection (via SAP Router) as part of the initial setup.

Since Solution Manager 7.1 SP04, the SAP HANA databases can be integrated into SAP Solution Manager. These include the following:

- Performance Warehouse
- Alerting Infrastructure
- The transaction `DBACOCKPIT`

Remote service connection can be established through the SAP Router.

New connection type allows SAP support to access customer databases through local SAP HANA studio installation.

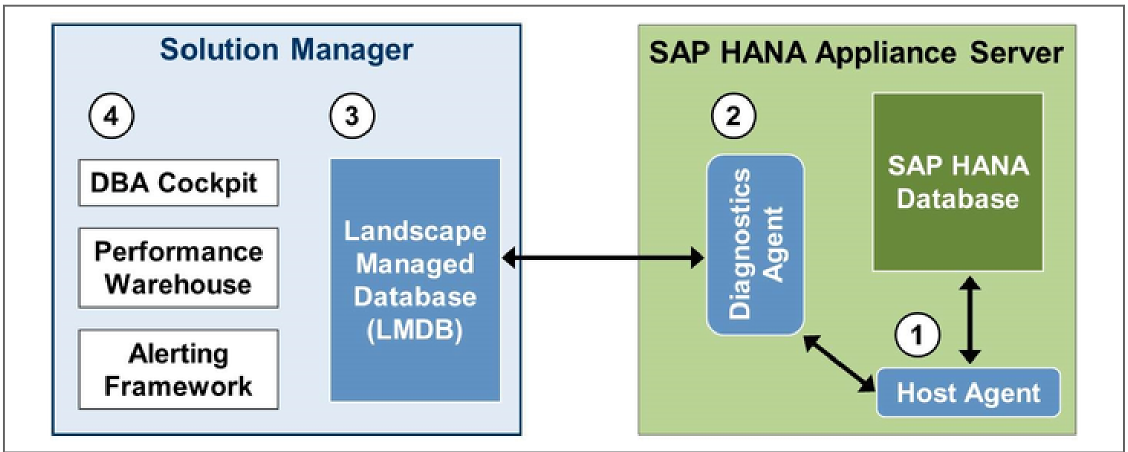

Figure 121: Solution Manager Connectivity: Technical System Overview

Involved Components

The Solution Manager connectivity involves the following components:

- Hostagent (Hostagent has to communicate with the SAP HANA Database)
- Diagnostics Agent (Diagnostics Agent has to communicate with the hostagent)
- Solution Manager (Diagnostics Agent has to be assigned to a Solution Manager)

Remote Connection to Solution Manager

As part of initial setup, establish the Solution Manager connectivity and the Remote Service Connection (via SAP Router).

SAP HANA can be connected to Solution Manager 7.1 greater than SP02 Level 3, but the recommendation is to use the newest one.

Establish standard SAP GUI and HTTP connection to Solution Manager (see SAP Note 962516).

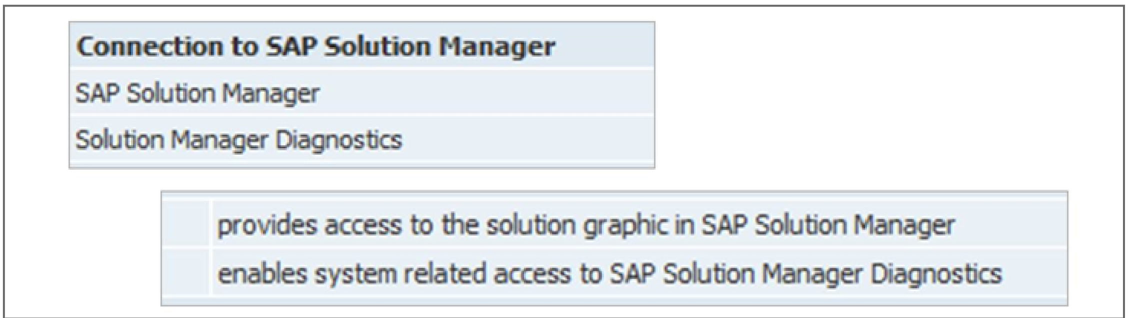

Figure 122: Remote Connection to Solution Manager

Configure Remote Support via SAP Router to SAP HANA Database Studio

To set up Root Cause Analysis, System Monitoring, and EarlyWatch Alert for SAP HANA with Solution Manager Version 7.10, see SAP Note 1747682 . The note has attachments that contain detailed instructions about the setup process.

Connection from the SAP Network to the customer network is through sapstartsrv on the SAP HANA Server to the Master Index Server.

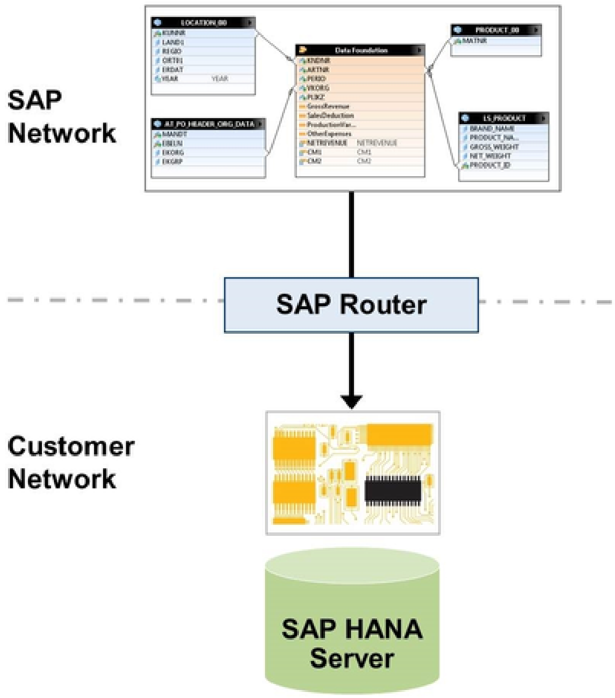

Figure 123: Configure Remote Support via SAP Router to SAP HANA Database Studio

SAP Router Information

The installation of saprouter is required at the customer site. SAP Support can provide support for customer systems through this saprouter connection. For more information on how to set up and configure the saprouter connection, see https://support.sap.com/remoteconnection.

The saprouter creates a secure SAP HANA studio connection from SAP Support to the customers SAP HANA database.

The customer opens the remote connection for a specific connection type, that is, SAP HANA database or ssh. The customer defines the different connection types in the system data.

For more information on the SAP remote connection for support, see the following SAP Notes:

- SAP Note 1592925 (SAP HANA database service connections)
- SAP Note 1635304 (Central note for SAP HANA support connections)

Remote Support via SAP Router to SAP HANA Database Studio

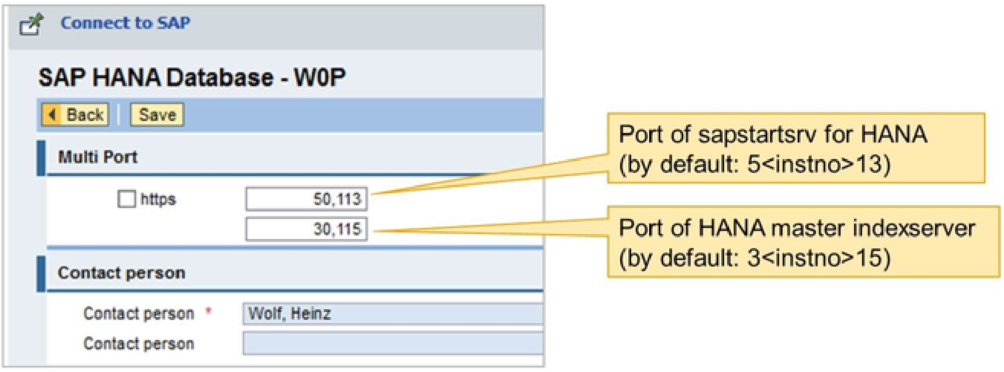

Figure 124: Remote Support via SAP Router to SAP HANA Database Studio

In the Connect to SAP configuration screen for the SAP HANA Database, you specify the port numbers of the sapstartsrv service and the master index server. The sapstartsrv is used to troubleshoot the problem when the SAP HANA database is down. The trace files can be reviewed by the sapstartsrv and the status of the SAP HANA services can be checked. When the database is running, the master index server can be used to troubleshoot the problem.

In some support cases, it might be necessary to provide OS-level access to SAP support, as follows:

1. For SAP HANA Linux systems, set up a SSH (SAP Note 1275351) or VNC (SAP Note 1327257) remote connection.
2. For Windows systems used for SAP BusinessObjects components, set up a Netviewer connection (SAP Note 1036616). A Netviewer connection requires the customer to actively Accept a connection request .
3. For unattended access, set up a Windows Terminal Server connection (see SAP Note 605795).

Perform the Post-Installation Steps

After the installation has finished successfully, the SAP HANA system is set up and running. You do not need to start the SAP HANA system, but you need to follow the post-installation instructions in the installation guide and the SAP installation notes.

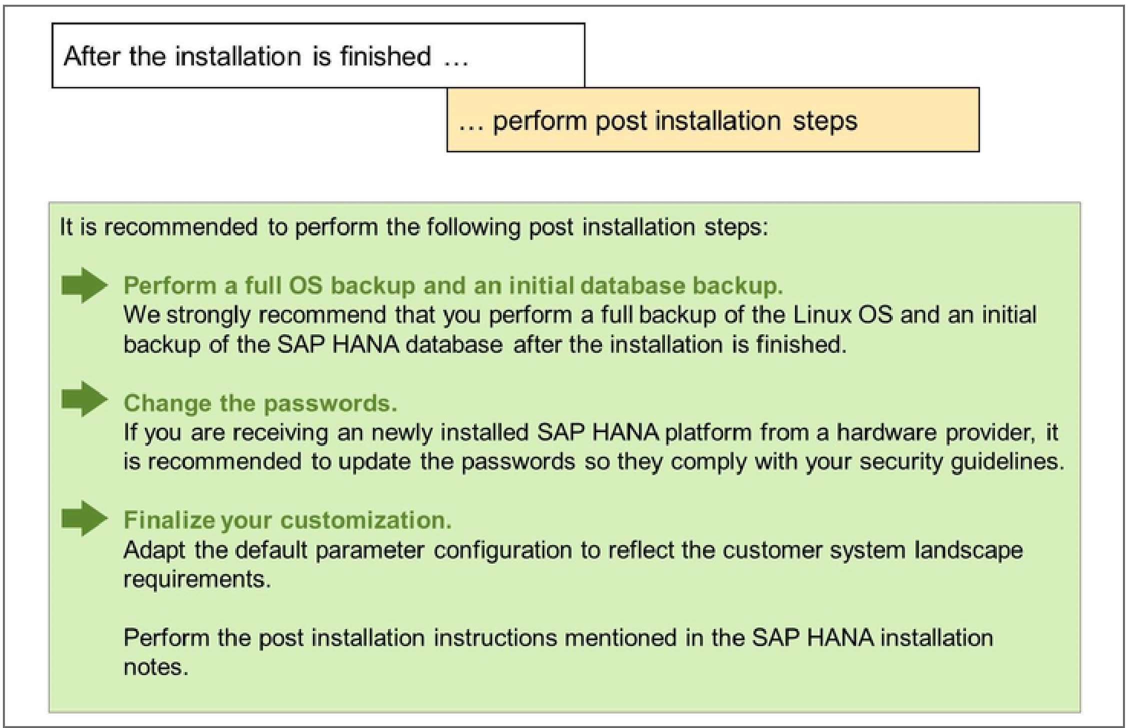

Figure 125: Post-Installation Steps

Implement the parameter recommendations for the SAP HANA database provided in the “Frequently Used Config Parameters in SAP HANA” document attached to SAP Note 2036111.

Other post installation tasks are as follows:

- Set the data backup location using parameter: `basepath_databackup`
- Set the log backup location using parameter: `basepath_logbackup`
- Activate log mirroring parameter: `basepath_logmirror`
- Check the database log mode using parameter: `log_mode`
- Assign the database memory allocation limit using parameter: `global_allocation_limit`
- Create personal user accounts for the database administrators.
- Deactivate the SYSTEM as described in the SAP HANA Administration guide (SAP HANA 2.0 SPS02) on page 674.
- Reserve Connections for Administrators as described in the SAP HANA Administration guide (SAP HANA 2.0 SPS02) on page 258.
- Install and configure SAP HANA Cockpit 2.0 to monitor and administrate the SAP HANA database.
- If required, update the SAP HANA database to the latest revision.

SAP HANA License

The General Licensing Process

As with all SAP products, you need a license from SAP to run SAP HANA. There are two kinds of license keys, as follows:

- Temporary license keys

  A temporary license key, which is valid for 90 days, is automatically installed in the system database of a new SAP HANA system and is effective for all tenant databases. During this period, you should request and install a permanent license key.

- Permanent license keys

  You can request a permanent license key on SAP Support Portal (https://support.sap.com/licensekey) under Launch the License Key Application. Permanent license keys are valid until the predefined expiration date. Furthermore, they specify the amount of memory licensed to the target SAP HANA database. Before a permanent license key expires, you should request and apply a new permanent license key. If a permanent license key expires in the system database, a temporary license key valid for 28 days is automatically installed. During this time, you can request and install a new permanent license key, for example, using the SAP HANA cockpit.

- License Keys for Tenant Databases

  You can install permanent license keys in individual tenant databases. The license key installed in a tenant database is valid for that database only and takes precedence over the license key installed in the system database. If a tenant-specific license key is not installed, the system database license key is effective in the tenant database.

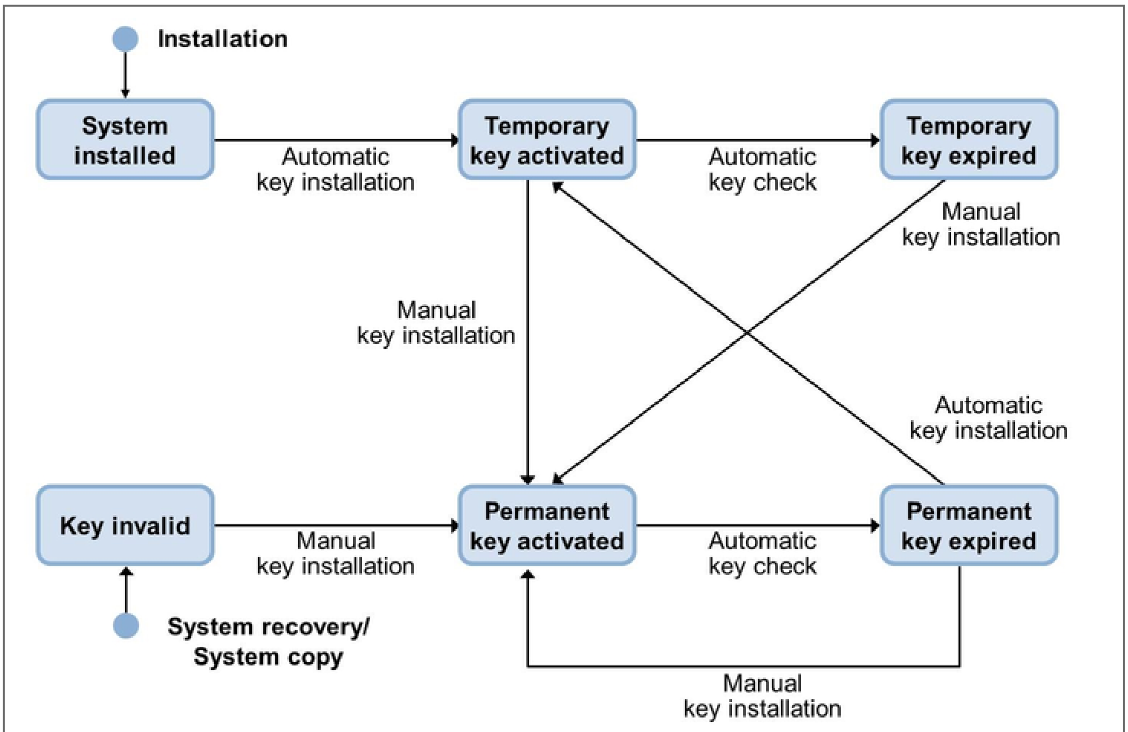

Figure 126: The General Licensing Process

Unenforced and Enforced License Keys

There are two types of permanent license key available for SAP HANA: unenforced and enforced.

- Unenforced license key

  The operation of SAP HANA is not affected if its memory consumption exceeds the licensed amount of memory.

- Enforced license key

  The database is locked down when the current memory consumption of SAP HANA exceeds the licensed amount of memory plus some tolerance. If this happens, either SAP HANA needs to be restarted, or a new license key that covers the amount of memory in use needs to be installed.

Database Lockdown

If the database goes into lockdown mode, it might be for one of the following reasons:

System database

The permanent license key has expired because:

- You did not renew the license within the 28 days.
- You changed the hardware and didn't request a new license key
- The amount of licensed memory was exceeded of the enforced license
- You deleted all the license keys

Tenant database

- The permanent license key has expired due to the same conditions as for the System database.
- The amount of licensed memory for the enforced license was exceeded.

  If the effective enforced license key is installed in the tenant database, it takes precedence over the license key installed in the system database. The tenant database remains in lock-down mode, even if there is a valid license key available in the system database.

In lockdown mode, it is not possible to query the database. Only a user with the system privilege LICENSE ADMIN can connect to the database and execute license-related queries, such as, obtain previous license data, install a new license key, and delete installed license keys. The database cannot be backed up in lockdown mode.

> Note:
> In a locked-down tenant database, deleting the locally installed license key will resolve the situation, assuming the system database has a valid license and the locally installed license key is not an enforced license key.

Additional licenses are required for certain applications running on the SAP HANA database, as well as certain SAP HANA options and capabilities. For more information, see SAP Note 1644792.

Install the SAP HANA License Key

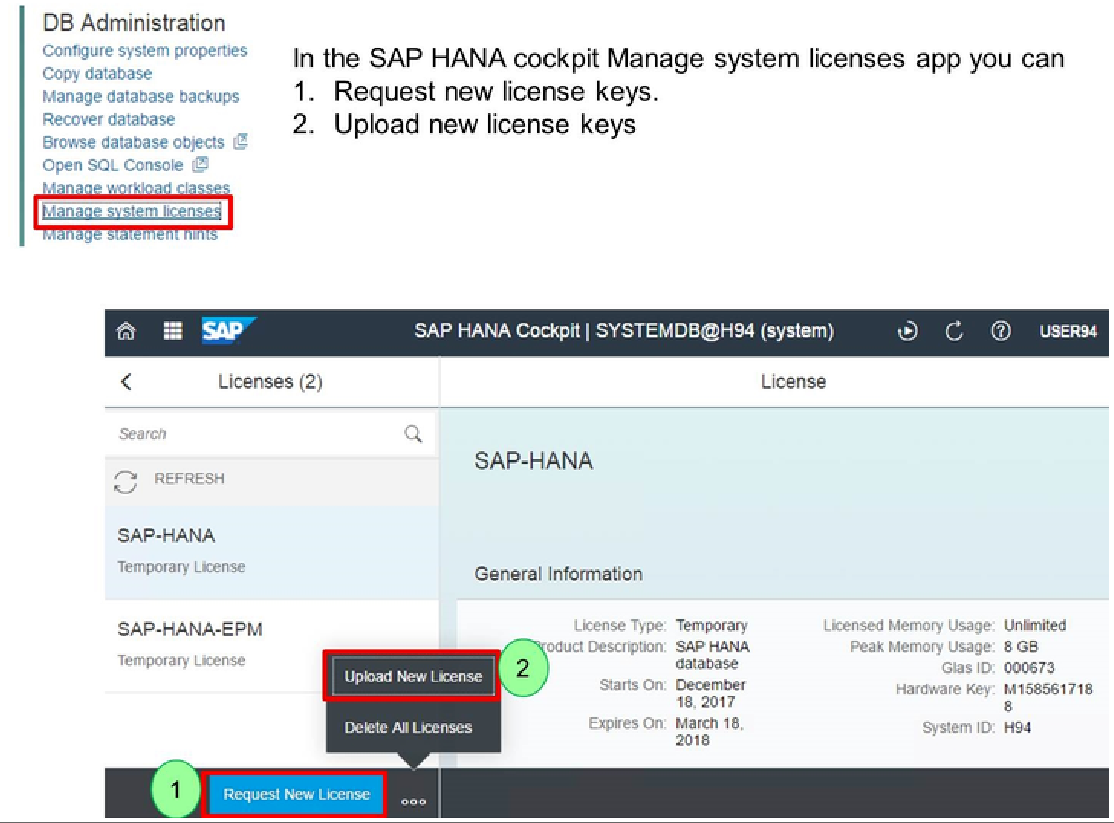

Figure 127: Install the SAP HANA License Key

Only a system with a valid license, that is not locked down, can be backed up. The license is also backed up and then restored with Recovery. When the Recovery of the backup is performed on the same system, there is no change in System ID and Hardware Key; the license key from the backup is recovered and used for license check. If the backup is too old, the license key from the backup might have expired. In this case, the database is locked after recovery and you need a new valid license key to unlock the database.

> Note:
> You are alerted 30 days before the license expires.

License Keys Further Information

For all tasks around the license management in the SAP HANA database, you need the system privilege LICENSE ADMIN.

The licensed memory is the amount of memory that a customer wants to assign to a particular SAP HANA instance. When a customer requests a license key from the SAP Service Marketplace, the customer must enter the amount of memory they require. The customer can decide how much they want to assign to the particular instance from the whole amount they bought. Then the specified number is put into the generated license key file. Once the license key is installed into the designated SAP HANA instance, the number is set in the SAP HANA instance and it shows in SAP HANA studio.

Memory allocation in SAP HANA Database uses a pool concept. The memory is pre-allocated from the operating system to gain performance on actual allocations done in SAP HANA Database code.

The global_allocation_limit parameter is used to limit the amount of memory that can be used by the database. The value is the maximum allocation limit in MB. By default, the global_allocation_limit parameter has the value 0. Depending on the amount of physical memory available in the SAP HANA server, the memory manager then allocates the memory according to the following rules:

- 90% of the first 64 GB of available physical memory
- 97% of each further GB

> Note:
> When the physical memory is less than 10 GB, then 1 GB is kept free for the operating system.

Licensing-Related SAP Notes

Table 5: Licensing-Related SAP Notes

The table contains licensing-related SAP Notes for further reference:

SAP Note	Description / Content
1644792	License key req./installation SAP HANA databases
1704499	System Measurement for License Audit
1817105	License-key / SMP only shows SAP HANA Platform Edition
2078724	Does SAP HANA Live Require a Separate License Key?
2366280	How to Create License Keys for SAP S/4HANA, on-premise edition
2375378	How to Create a License key for BW/4HANA 1.0

LESSON SUMMARY

You should now be able to:

- Perform the post-installation steps after a SAP HANA installation

Unit 7 Lesson 2

Updating SAP HANA

LESSON OVERVIEW

This lesson describes how you can update the SAP HANA database system using the SAP HANA lifecycle manager (HDBLCM).

Business Example

The project team has requested a SAP HANA 2.0 SPS02 sandbox system so that they can gain experience with the newest SAP HANA 2.0 SPS02 features like Multitenant Database Containers. Update the SAP HANA 1.0 SPS12 sandbox system to SAP HANA 2.0 SPS02.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Update a SAP HANA database using the HDBLCM tool

SAP HANA Database Update

The SAP HANA database lifecycle management tool (HDBLCM) can be used to update SAP HANA system components like the SAP HANA database and the SAP HANA client software. It can also update additional system components like Application Function Libraries (AFL and the product-specific AFLs POS, SAL, SCA, SOP, UDF), SAP liveCache applications (SAP LCA or LCAPPS-Plugin), SAP HANA extended application services advanced runtime applications, or SAP HANA smart data access (SDA)

Before you update the SAP HANA database, download the separate components from the SAP Support website (https://support.sap.com). There are several ways to download the required components, as follows:

- Manually via the SAP Software Download Center (https://support.sap.com/swdc).
- Using the Download Components option in SAP HANA cockpit.

  Using the Maintenance Planner from SAP Solution Manager.

Perform an SAP HANA system update from a local host. This minimizes the risk of a failed update because of network glitches.

Three Tools to Update the SAP HANA System

The SAP HANA lifecycle manager tools hdblcm (command line), hdblcmgui (X Window), and hdblcmweb (browser) can perform all actions offered in the SAP HANA platform lifecycle management portfolio. They are the only recommended tools for these actions.

You do not have to start hdblcmweb manually. The executable hdblcmweb is started automatically by the SAP Host Agent as soon as an action is triggered from the Web user interface.

> Caution:
> Since SPS09, the SAP HANA database lifecycle management tool (HDBLCM) is the only supported toolset for SAP HANA lifecycle management tasks. The SAP HANA lifecycle management tools replace all other tools, such as the SAP HANA unified installer, the on-site configuration tool, SUM for HANA, hdbinst, hdbupd, and the SAP HANA lifecycle manager tools.

Two Ways to Upgrade SAP HANA

There are two versions of the SAP HANA lifecycle manager tools: the hdblcm located on the installation media, and the resident hdblcm embedded in SAP HANA database. Both tools are required and are used to perform different administrative tasks.

The hdblcm located on the SAP HANA installation media installs and updates an SAP HANA database and its core components. No additional administrative tasks are possible.

The resident hdblcm performs administrative tasks for the database where it's embedded. The resident hdblcm installs or updates additional SAP HANA components, but it can't update the SAP HANA database executable files.

The figure, Task Overview of hdblcm and the resident hdblcm, outlines the differences between the two tools.

Activity	Installation media hdblcm	Resident hdblcm
Installation SAP HANA and it's components	+	–
Update SAP HANA and it's components	+	–
Add additional hosts to the SAP HANA system	–	+
Configure inter-service communication	–	+
Configure System Landscape Directory (SLD)	–	+
Rename the SAP HANA System	–	+
Uninstall SAP HANA components	–	+
Unregister the SAP HANA System	–	+
Install or update additional components (AFL, LCApps, SDA)	–	+

Figure 128: Task Overview of hdblcm and the resident hdblcm

Checks Before Update Process

Before updating the SAP HANA components, check that no read or write processes are running on the SAP HANA database. Perform the update process outside of business hours, because the SAP HANA system restarts during the update process. After the update finishes, the SAP HANA system is started and operational.

The business downtime is reduced by using the Phased System Update method or the Near Zero Downtime method. Both methods are explained in this course.

The Update Process

To perform the update, proceed as follows:

1. Uptime: Download all the required software components from the SAP Service Marketplace.
2. Uptime: Prepare the downloaded software archives for the update.
3. Uptime: Back up the SAP HANA database.
4. Start of downtime: Close all external access to the system.
5. Downtime: Perform the SAP HANA update.
6. Downtime: Update the depending components.
7. End of downtime: Open all external access to the system.
8. Uptime: Perform the post-update steps.

When you start the SAP HANA Lifecycle Management tool from SAP HANA cockpit, a user-friendly SAP Fiori interface displays.

How long will the upgrade take?

Time for upgrade = (Time for database shutdown) + (Time for database start) + 20 minutes.

The database restart time is influenced by the size of the ROW store and the input/output performance of the storage used for SAP HANA.

Start SAP HANA Lifecycle Management Tool

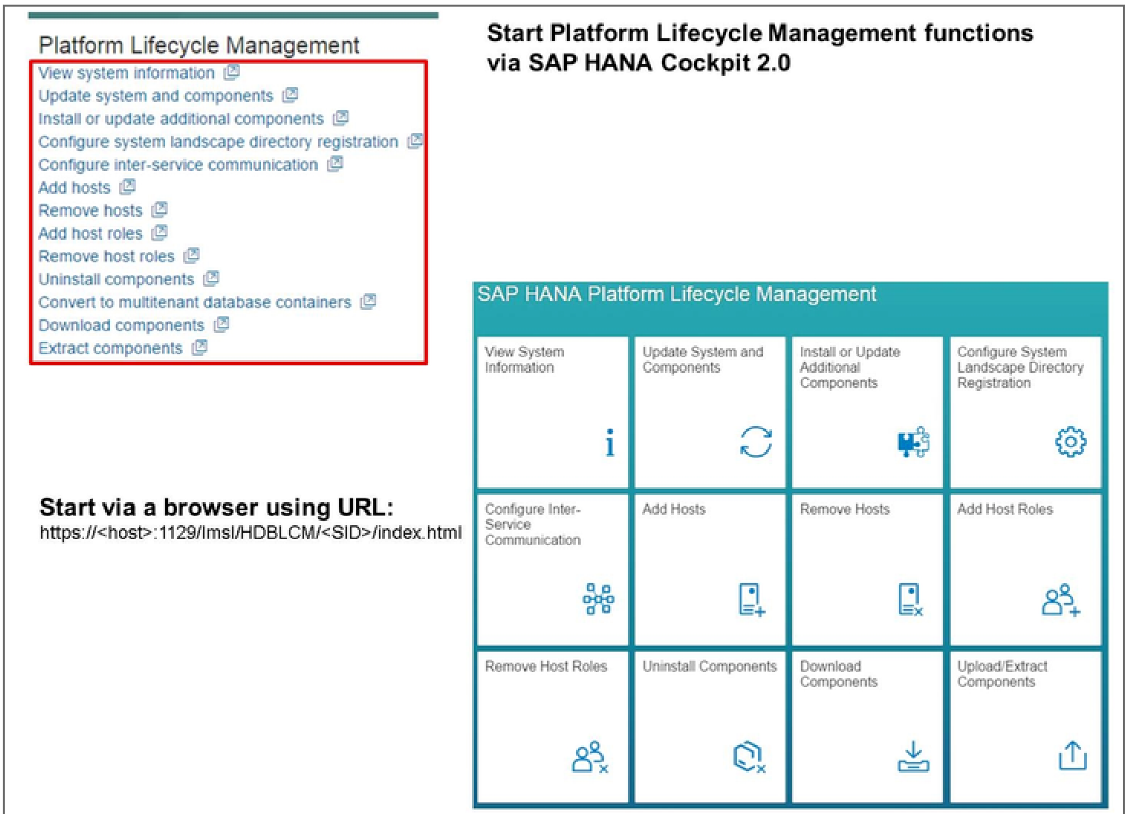

Figure 129: Start SAP HANA Lifecycle Management Tool

Troubleshooting with SAP HANA Lifecycle Manager

If the SAP HANA lifecycle manager does not work as expected, you can check the logs for the source of the problem, restart the lifecycle manager, or update to a more recent version.

Checking the Log Files

The SAP HANA lifecycle management tools hdblcm and hdblcmgui write log files during installation. The most recent log file is always available under `/var/tmp/hdblcm.log` or `/var/tmp/hdblcmgui.log`. Additionally, a copy of the log files is archived in the directory `hdb_<SID>_hdblcm_<action>_<date>`.

Because the SAP HANA lifecycle management tools hdblcm and hdblcmgui are wrappers for underlying component installers, you can also check the component logs. However, review and analyze the SAP HANA lifecycle management tools hdblcm and hdblcmgui logs first. Once the source of the problem is isolated to a specific component, you can analyze the component logs further.

The component log files are stored in the following path:

/var/tmp/hdb_<SID>_<action>_<time_stamp>

where <action>::= install | update | add host | uninstall | and so on

The following log files are written during performing the action:

- <hdbcommand>.log: can be read using a text editor
- <hdbcommand>.msg: XML format for the display in the installation tool with the GUI
- <hostname>_tracediff.tgz: provides a delta analysis of the original trace files that makes a detailed analysis more easy

Figure 130: Checking the Log Files

You can also view the last three log files in the SAP HANA studio using the Diagnosis Files administration function.

Phased Update

With a standard SAP HANA system update, the system goes offline once the update is triggered. This includes the preliminary checks and actual software switch.

Starting with SPS 10, to reduce system downtime, you can run an SAP HANA system update in two phases. The phased system update is performed in two steps, as follows:

1. Running the LCM update action with the prepare update checkbox selected. This phase is performed while the system is online.
2. Running the LCM update action a second time as usual, which resumes the updates and takes the system offline for the software switch.

You can perform the prepare update phase using either the SAP HANA database lifecycle manager graphical user interface, command-line interface, or Web user interface. The update resume phase can be performed from any of the three SAP HANA database lifecycle manager user interfaces.

Prerequisites

The prerequisites for a phase update are as follows:

- You are updating to a new Support Package Stack from an installation medium or you have prepared for update, either in the SAP HANA studio or manually.
- You have stopped the data replication.
- You have performed a system backup. Note that, during the update, there is a business downtime for your SAP HANA system.
- You know the <sid>adm, and database administrator passwords.
- You have applied a valid license key for the SAP HANA system.
- The SAP HANA system has been installed or updated with the SAP HANA database lifecycle manager (HDBLCM) Support Package Stack 10 or later.
- The SAP HANA database server is up and running. Otherwise, inconsistencies in the configuration occur.

Benefits of a Phased System Update

After downloading the SAP HANA software and preparing the downloaded archives for update execution, you can update your SAP HANA system in one step, or in a phased approach to minimize system downtime.

When you start the SAP HANA database lifecycle manager with the prepare_update checkbox selected, the SAP HANA database lifecycle manager extracts the packages (such as the SAP Host Agent and delivery units) from the new source, but does not perform the update. The software switch occurs when the SAP HANA database lifecycle manager runs a second time, which resumes the system update.

The phased update includes the following benefits:

- Decreased system downtime
- Reduced chance of a failed system update because the preliminary steps include archive preparation and dependency conflicts

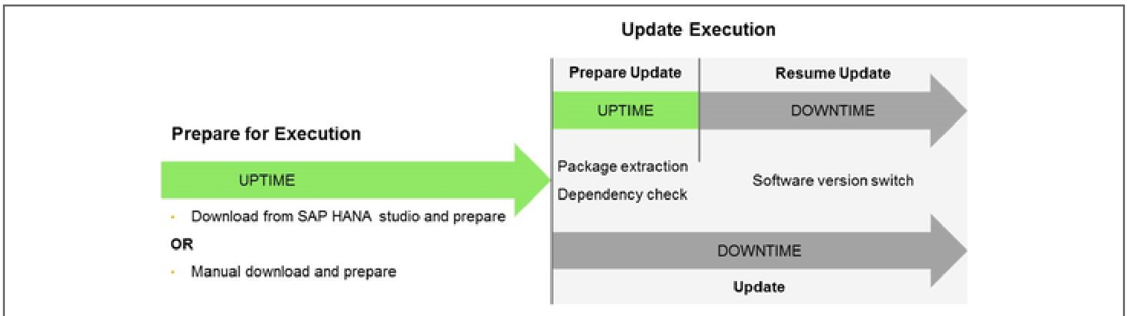

Figure 131: Phased Update

Perform a Phased System Update

Open the following directory on the installation medium: cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64

Perform the update preparation phase step with the SAP HANA database lifecycle manager using one of the following commands:

- `./hdblcmgui --action=update --prepare_update`
- `./hdblcm --action=update --prepare_update`

Resume the SAP HANA update. During the planned maintenance window, you can resume the prepared update using any of the standard update procedures.

LESSON SUMMARY
You should now be able to:

- Update a SAP HANA database using the HDBLCM tool

Unit 7
Lesson 3

Using the Resident SAP HANA Database Lifecycle Manager (HDBLCM) Tool

LESSON OBJECTIVES
After completing this lesson, you will be able to:

- Explain the resident HDBLCM functions

The Resident HDBLCM Tool

As an SAP HANA database administrator, you perform many complex tasks, like renaming an SAP HANA system or converting an SAP HANA database to a multitenant database. These complex system administration tasks can be streamlined with the Resident SAP HANA lifecycle manager (HDBLCM) tool. In SAP HANA 2.0, the Resident HDBLCM tool helps you to perform the following tasks using a simple procedure:

- Add or remove Host Roles.
- Add additional hosts to the SAP HANA database.
- Configure the Interservice communication in a scale-out environment.
- Configure the connection to the System Landscape Directory (SLD).
- Convert an SAP HANA database system to a Multitenant Database Containers (MDC) system.
- Uninstall the complete SAP HANA database or individual components.
- Install or update additional SAP HANA components.
- Remove an SAP HANA compute node from the SAP HANA storage to scale-up the compute node.
- Rename the System ID or Instance Number of an SAP HANA system.

When you start the resident hdblcmgui from the `/hana/shared/<SID>/hdblcm` directory, a graphical interface appears. You can also perform the same functions with the resident hdblcm .

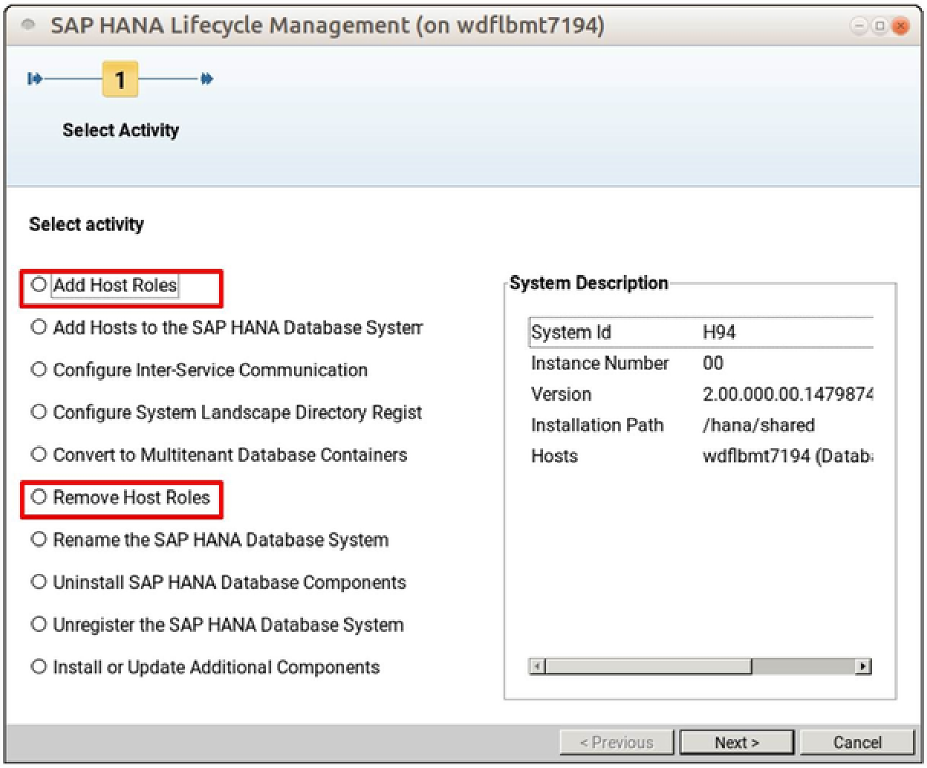

Figure 132: Adding or Removing Host Roles

Add or Remove Host Roles

The Add or Remove Host Role option allows you to specify additional roles for existing SAP HANA hosts during the SAP HANA option installation. Multiple host roles are not supported in production environments, but they can be useful in development and QA systems to keep the system landscape size under control. If an SAP HANA extended application services advanced runtime is installed, hosts can share multiple roles.

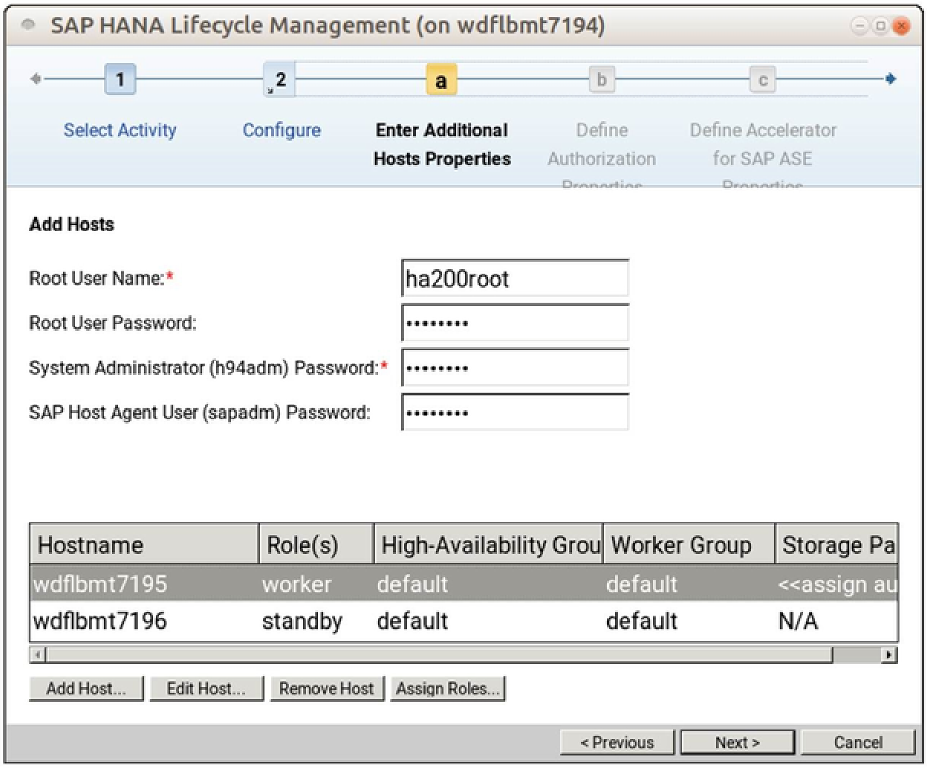

Figure 133: Add Hosts to an SAP HANA Database

Add Hosts to an SAP HANA Database System

If you want to increase the SAP HANA database size, you can add additional hosts to the system in a process called Scale-Out. This lets you increase the memory size and the number of CPUs in the SAP HANA system by adding hosts. You can specify the list of server and roles (worker or standby) to add with the resident HDBLCM tool. The installation and configuration is automatically performed by the resident HDBLCM tool.

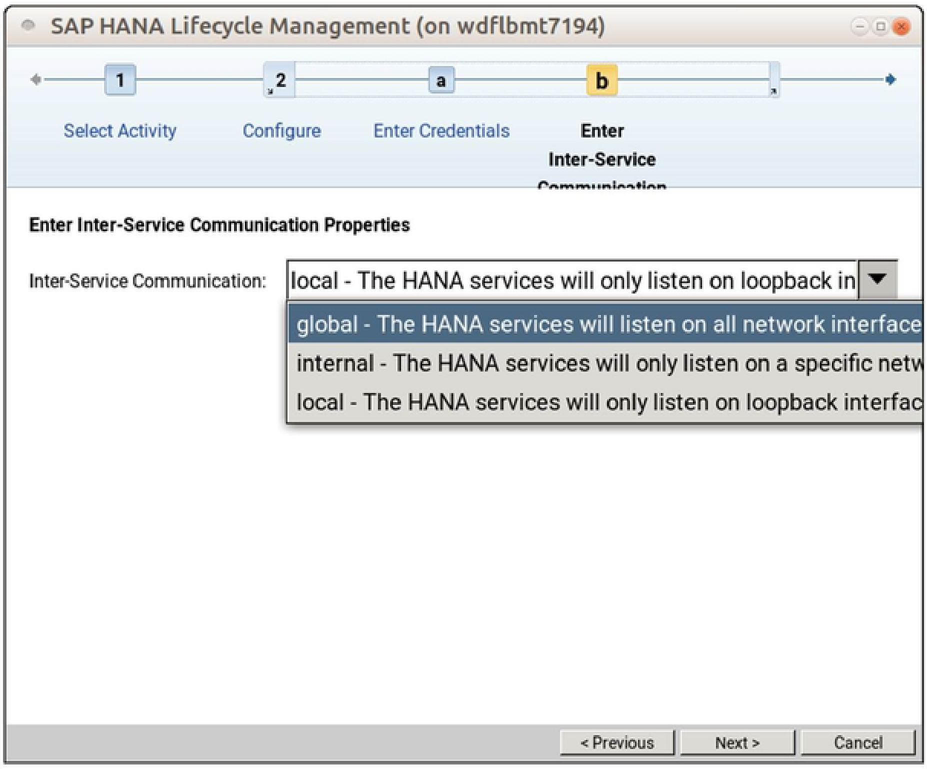

Figure 134: Inter-Service Communication Settings

Configure the Inter-Service Communication Settings

With this option, you can specify which network interface SAP HANA uses for internal network communication. The following options are available:

- Global: Binds the processes to all interfaces.
- Internal: Binds the processes to this address only and to all local host interfaces.
- Local: Opens the communication ports for internal usage on the local interfaces. This configuration is only an option for single installations because the server is not accessible from outside.

The global and local options do not require an internal network address entry. The internal option requires an internal network address entry.

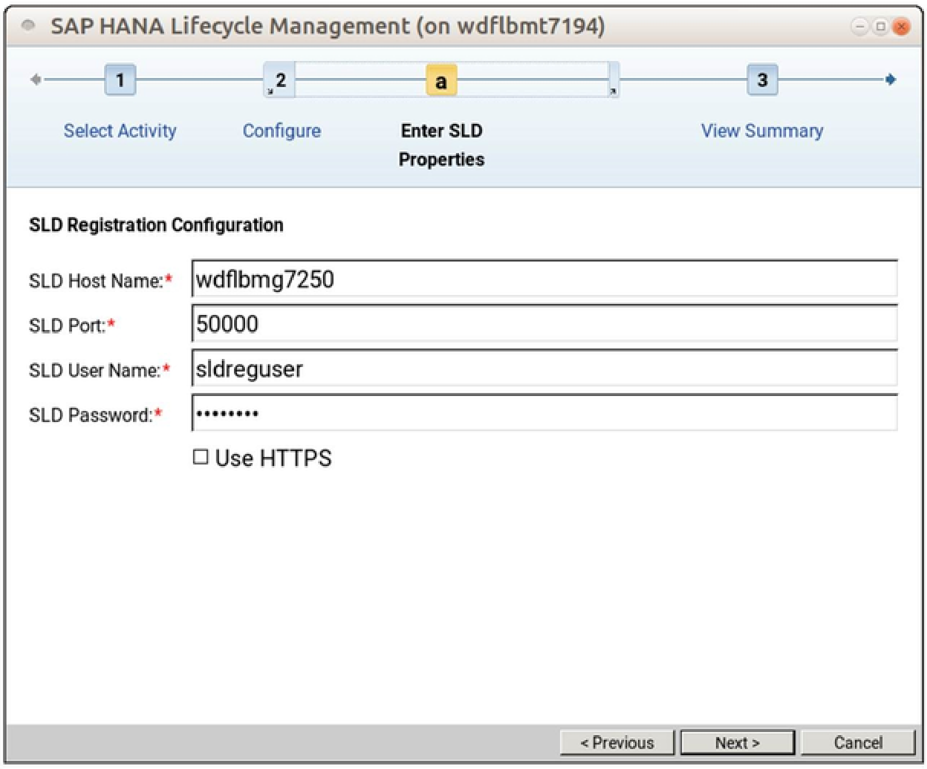

Figure 135: Configure System Landscape Directory (SLD)

System Landscape Directory Registration (SLD) Configuration

If you have SAP Solution Manager, you can use it to update your SAP HANA system like other SAP systems in your landscape. Register your SAP HANA system using the System Landscape Directory (SLD).

After you perform this configuration, the SAP HANA database automatically updates the SLD information on a regular basis to keep it current. This data can be used by the SAP Solution Manager to calculate update paths for your SAP HANA system.

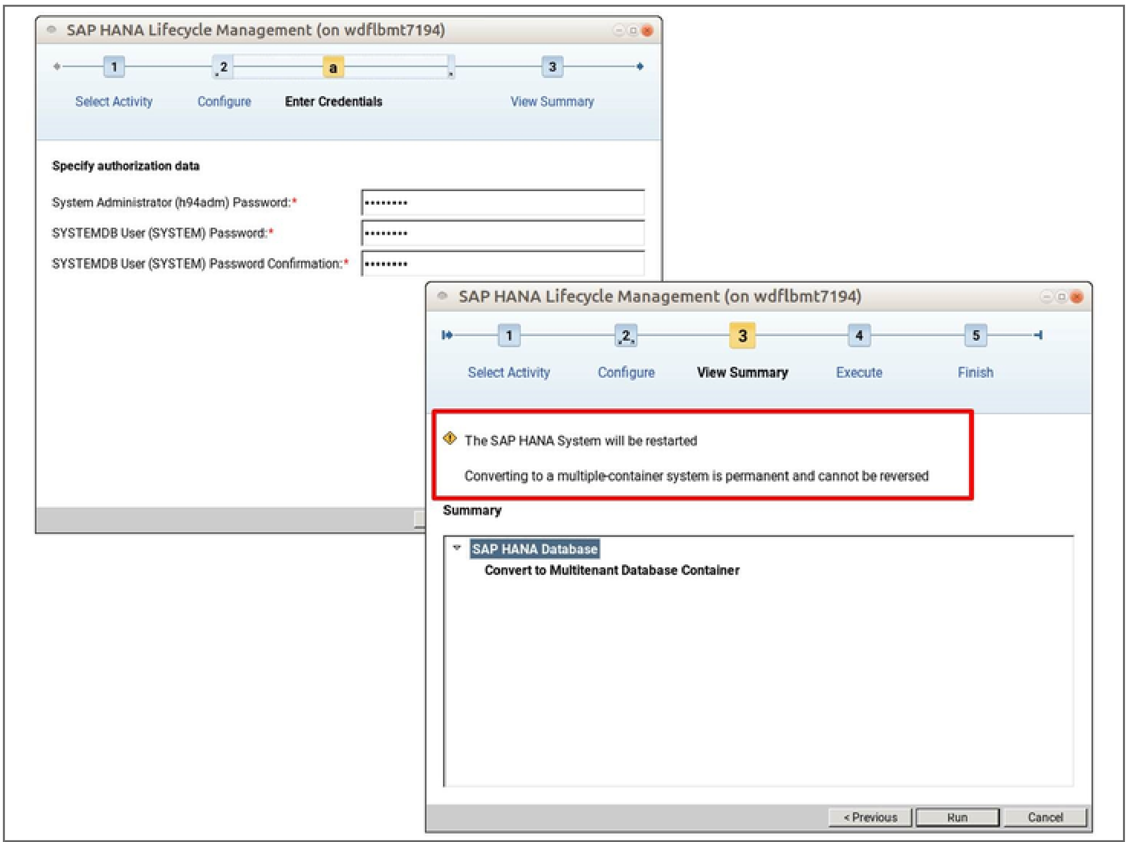

Figure 136: Multitenant Database Containers Conversion

Convert to Multitenant Database Containers (MDC)

With this option, you can convert an SAP HANA database to a multitenant configuration. A multitenant database container enabled system contains one system database and it can contain multiple tenant databases. During the conversion, the SAP HANA database restarts and the conversion cannot be reversed.

Since SAP HANA 2.0 SPS01, the default setup is Multitenant Database Containers. Thus, upgrading to SAP HANA 2.0 SPS01 initiates the MDC conversion automatically. New installations are installed as MDC database systems.

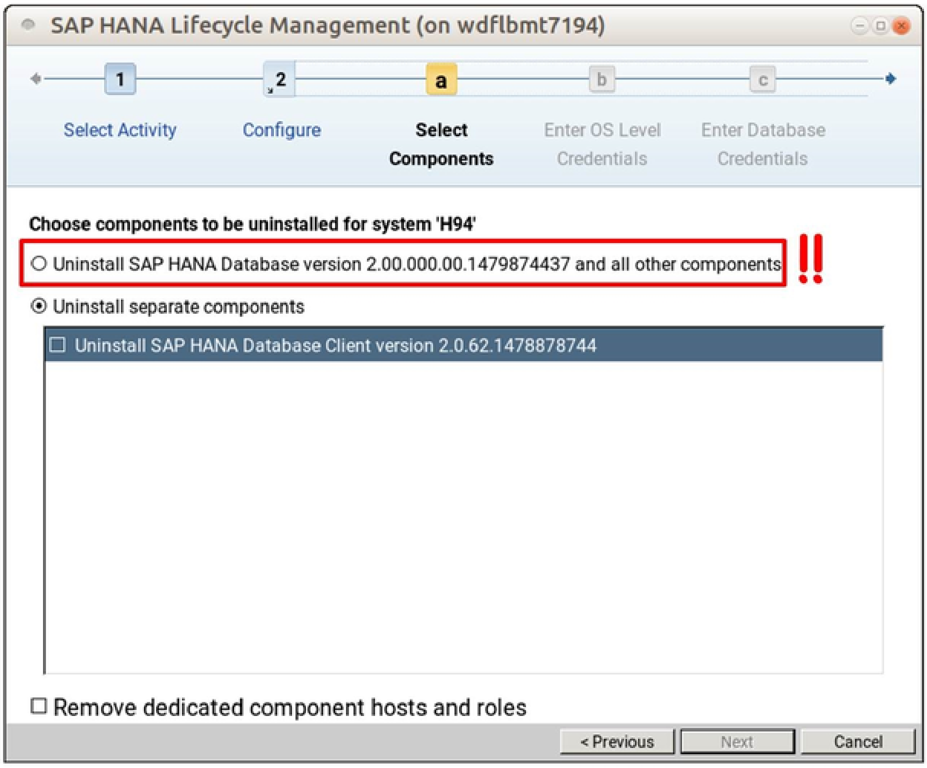

Figure 137: Uninstalling SAP HANA Database Components

Uninstall SAP HANA Database Components

With the Resident HDBLCM, you can uninstall the complete SAP HANA database system, or uninstall individual SAP HANA database system components. The following types of components can be managed:

- SAP HANA mandatory components (SAP HANA server and client)
- SAP HANA additional components (like Application Function Libraries, SAP liveCache applications, and SAP HANA smart data access)
- SAP HANA options (like SAP HANA dynamic tiering and SAP HANA smart data streaming)

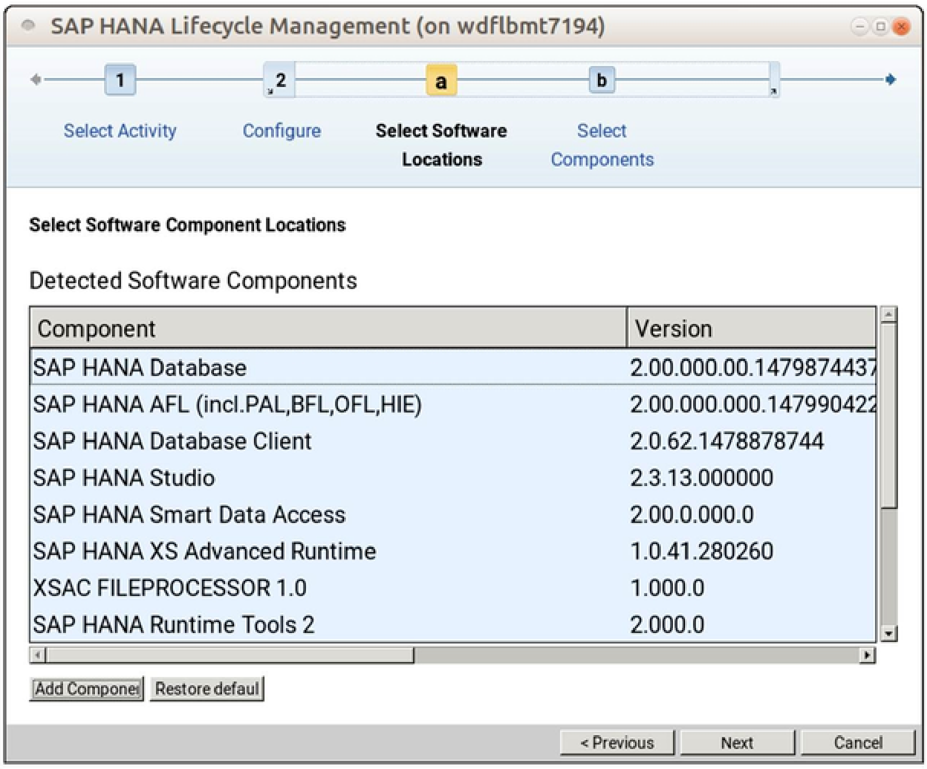

Figure 138: Install or Update Additional Components

Install or Update Additional Components

Once the SAP HANA database system is installed, you can easily add additional components to the system. This can even be done after the SAP HANA database has been upgraded to a higher Support Package Stack (SPS) or version.

To install or update SAP HANA components, first check that you have downloaded the components you want from the SAP Support Portal (https://support.sap.com). The component version must match the SAP HANA database version.

After you have unpacked the component, start the resident HDBLCM tool and point it to the download location. The directory is searched, and any found components are shown as options that can be installed or updated.

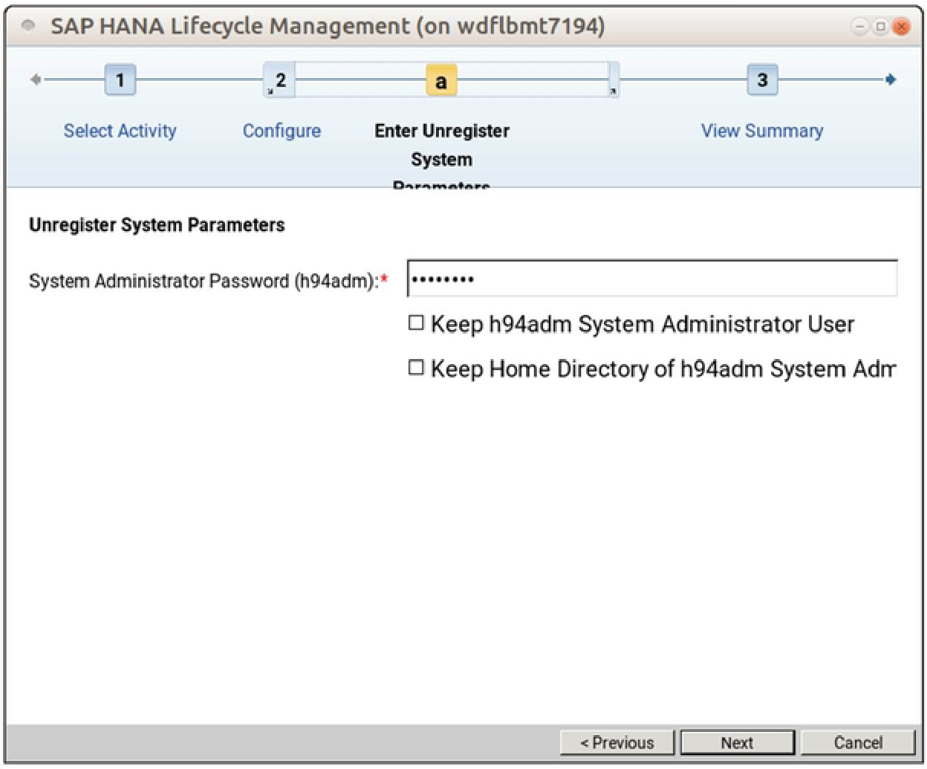

Figure 139: Unregistering an SAP HANA Server

Unregister the SAP HANA Database System

When an SAP HANA database runs out of memory storage, increase the memory of that SAP HANA system. This can be done by scale-out or scale-up. In the scale-up scenario, replace your current SAP HANA host with a new, larger host. To replace the smaller host with a larger host, unregister the smaller server from the SAP HANA database. The larger host can be registered to the SAP HANA database with the rename option in the resident HDBLCM tool.

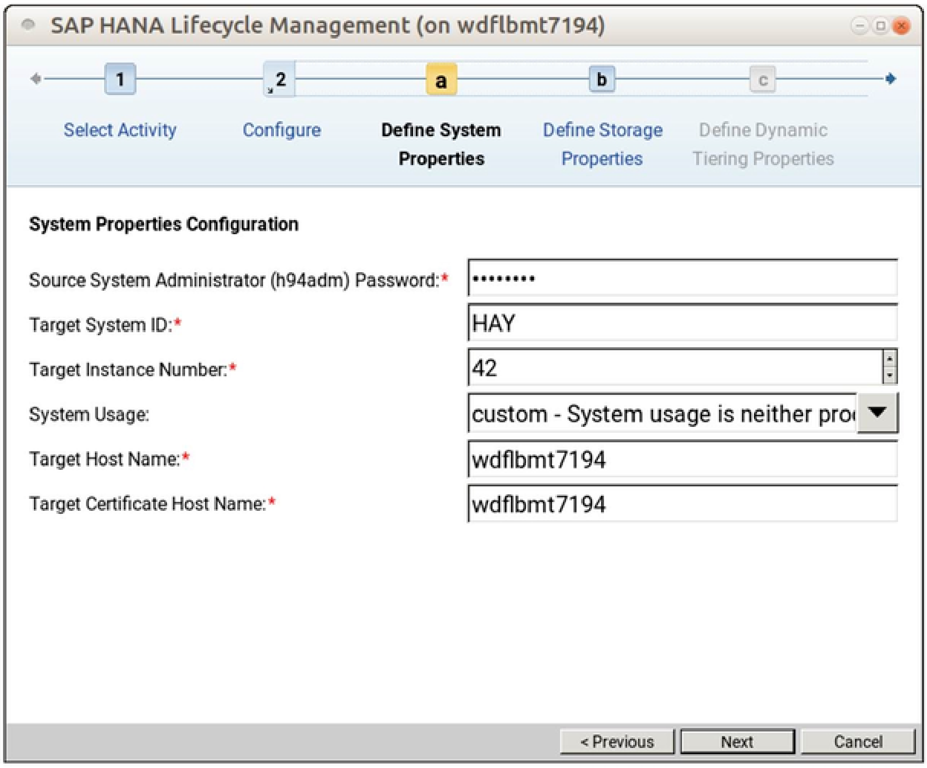

Figure 140: Renaming an SAP HANA database

Rename the SAP HANA Database

An SAP HANA system can be renamed by changing the system identifiers, such as host names, SID, and instance number. You can change system identifiers with the SAP HANA database lifecycle manager (HDBLCM).

Example Scenario: Unregister

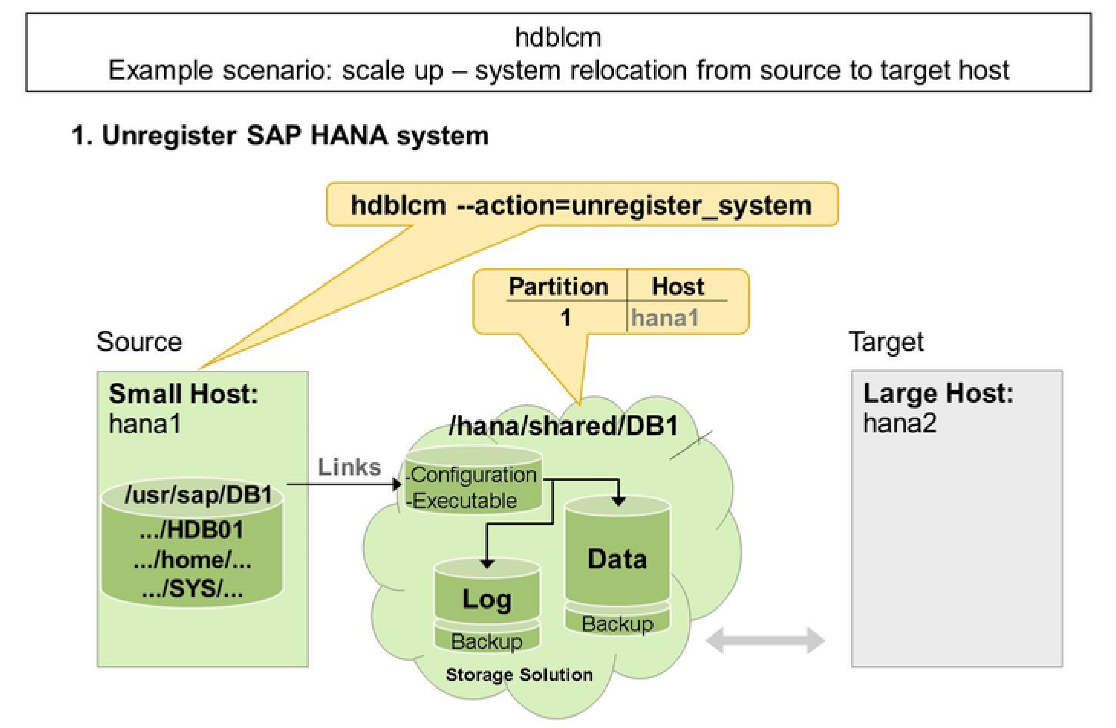

Figure 141: Example Scenario: Unregister

Relocation of SAP HANA

You might need to move the SAP HANA system to different, larger hardware. To do this, unregister the SAP HANA system on the old hardware, and reregister it on the new hardware. To complete this task, use the SAP HANA lifecycle management tool hdblcm(gui).

The SAP HANA platform lifecycle management tool hdblcm provides an easy way to relocate your current SAP HANA system onto new hosts, without the use of reinstall or backup and restore mechanisms.

The steps for this scenario are as follows:

- Logon to the source host.
- Navigate to the resident hdblcm.
- Execute `HDBLCM --ACTION=UNREGISTER_SYSTEM`.

Example Scenario: Register

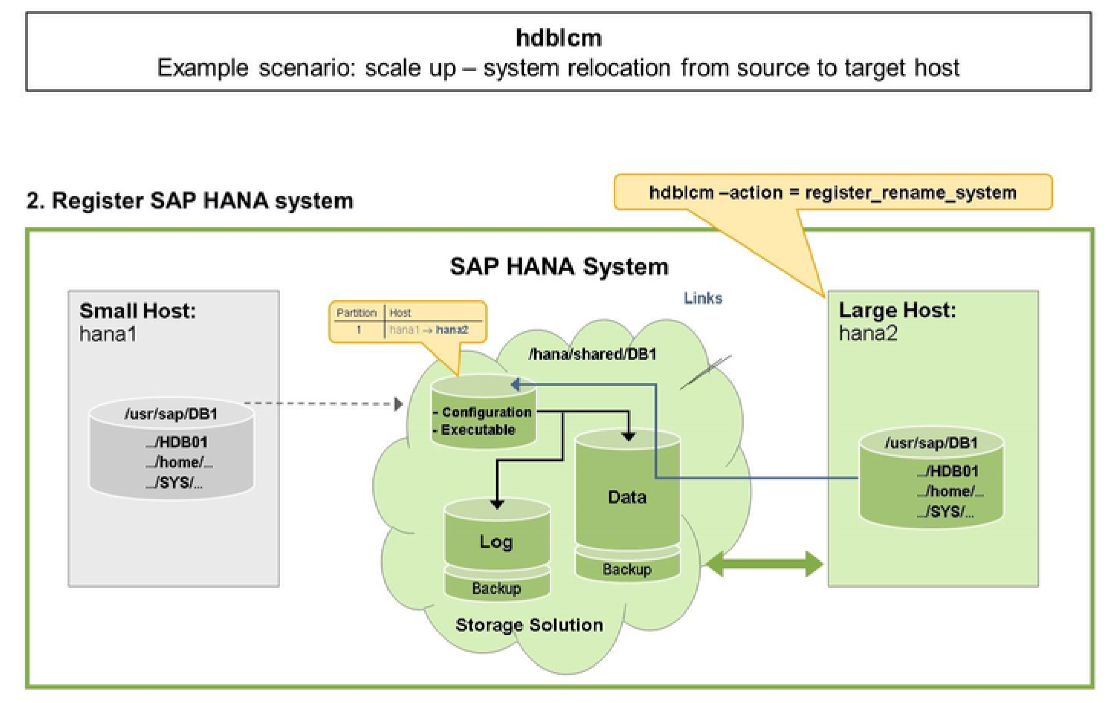

Figure 142: Example Scenario: Register

The steps to register the new host are as follows:

- Log on to the target host and mount the shared area.
- Execute `HDBLCM -ACTION=REGISTER_RENAME_SYSTEM`.
- Execute the host mapping.

Additional Features for the Resident HDBLCM Tool

With every new Support Package Stack, new or enhanced features are added to the resident HDBLCM tool. The command line version can have more features than the graphical version. Also, you can use the command line version in scripts to automate recurring system administration tasks. For more information on the use of the command line option for the resident HDBLCM, see the SAP HANA Installation and Update Guide.

LESSON SUMMARY
You should now be able to:

- Explain the resident HDBLCM functions

Unit 7
Lesson 4

Using SAP HANA Interactive Education (SHINE)

LESSON OVERVIEW

This lesson describes the SAP HANA Interactive Education (SHINE) application, which demonstrates how to build native SAP HANA applications.

Business Example

In your company, the SAP HANA Developers want a sandbox system with some content so that they can understand the features provided by SAP HANA SPS10.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explain the installation and use of SAP HANA InteractiveEducation (SHINE)

The SAP HANA Interactive Education (SHINE) Demo Application

SAP HANA Interactive Education, or SHINE, is a demo application that shows how to build native SAP HANA applications. The demo application, delivered with SAP HANA in a special delivery unit (DU), comes with sample data and design-time developer objects for the application's database tables, data views, stored procedures, OData, and user interface.

The delivery unit defines the following applications:

- Enterprise Procurement Model Admin Console

  You can generate large quantities of data for testing, and create synonyms for use in currency conversions.

- Enterprise Procurement Model Sample Application

  This is a sample Sales Order Dashboard and Purchase Order Worklist that shows how to construct similar native SAP HANA applications.

The delivery unit creates the SAP_HANA_DEMO schema. In this schema, the database objects, including the tables, are created. The views and procedures are created in the _SYS_BIC schema.

The delivery unit also comes with design-time objects for building the applications based on those database objects. These are located in the sap.hana.democontent.epm package.

Features Available in SHINE

After the SHINE demo application is installed, the developers and modelers can explore the features on the SAP HANA system.

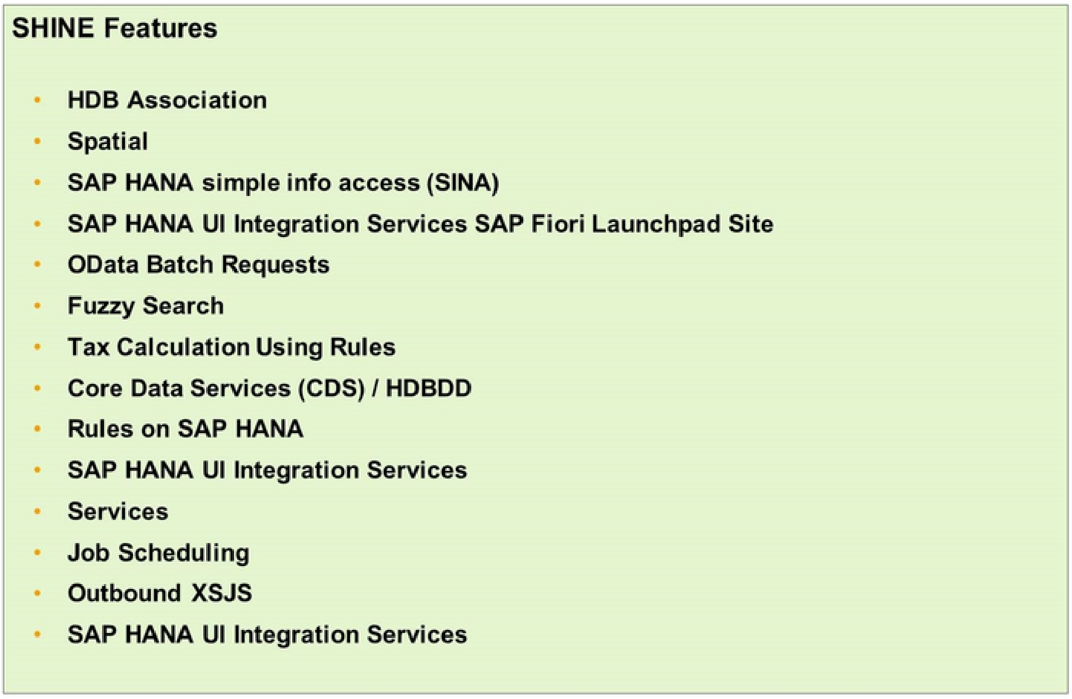

Figure 143: Available Features in SHINE

The following features are available:

- HDB Association

  SAP HANA extended application services allows you to use the core data services (CDS) syntax to create associations between entities. The associations are defined as part of the entity definition, which are design-time files in the repository.

- Spatial

  Spatial data describes the position, shape, and orientation of objects in a defined space. Spatial data is represented as 2-Dimensional (2D) geometries in the form of points, line strings, and polygons.

- SAP HANA simple info access (SINA)

  The SINA API is a client-side or front-end JavaScript API for developing browser-based search UIs.

- SAP HANA UI Integration Services SAP Fiori Launchpad Site

  The entry point to SAP Fiori apps on mobile or desktop devices.

- OData Batch Requests

  The OData standard collects multiple individual HTTP requests into one batched HTTP request.

- Fuzzy Search

  Fuzzy search is a fast and fault-tolerant search feature that can be used in SAP HANA.

- Tax Calculation Using Rules

  The Rules for Tax Calculation are used to determine the tax code based on the Company (Business Partner ID) and Product ID.

- Core Data Services (CDS) / HDBDD

  This is a new infrastructure for defining and using semantically rich data models in SAP HANA. CDS uses a data definition language (DDL), a query language (QL), and an expression language (EL). It includes write operations, transaction semantics, constraints, and more.

- Rules on SAP HANA

  This introduces business rules in the form of decision tables in SAP HANA database layer.

- SAP HANA UI Integration Services

  This provides the required services and UI patterns to create and design single applications, and sites based on SAP HANA extended application services native applications through efficient development tools, standardized services, and consistent UI experience.

- Services

  Many new services have been added.

- Job Scheduling

  Scheduled jobs define recurring tasks that run in the background.

- Outbound XSJS

  SAP HANA extended application services include a server-side JavaScript API (Outbound API) that allows access to a defined HTTP destination.

- SAP HANA UI Integration Services

  You want to allow the end user to personalize your application. You can use the personalization mechanism provided by the SAP HANA UI Integration Services (UIS).

SHINE Demo Application

To work with the demo application, a system administrator needs to perform the following tasks:

- Import the demo application delivery unit.
- Assign roles to developers who want to work with the demo application.

Afterwards, a developer with the proper role can perform the following tasks:

- Generate additional demo data, if necessary. The demo application comes with an initial set of data.
- View the demo application, and then explore the design-time objects for the demo applications to see how the applications were created.

Use of SHINE Demo Application

You can work with and explore the demo EPM application, and then view the code behind it to learn how it works. The application uses the purchase order data and sales order.

To launch and explore, you must have the sap.hana.democontent.epm.roles::User role assigned to your user. If you want to configure the SHINE demo application, you need the sap.hana.democontent.epm.roles::Admin role assigned to your user.

The Launchpad application is the entry point into the SHINE Demo Application. From the Launchpad, you can start the other applications that you want to explore.

SAP HANA Interactive Education Launchpad

Open the Launch Pad Application using the URL https://<hostname>.<xs port>/sap/hana/democontent/epm/ui/NewLaunchpad.html.

Replace <hostname> and <xs port> with the host name and port for your SAP HANA extended application services installation. The port is 43 plus the 2-digit instance number. For example, if the instance is 00, then the port is 4300.

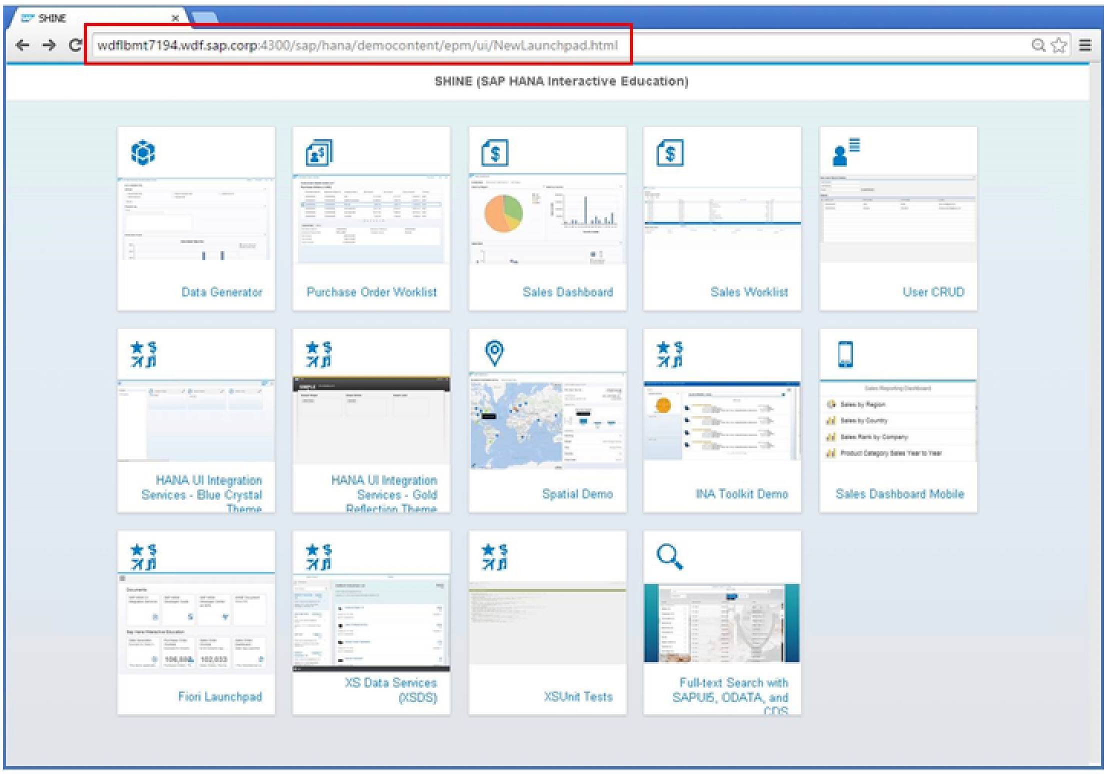

Figure 144: SAP HANA Interactive Education Launchpad

Some Examples from the SAP HANA SHINE Demo Apps

From the Launchpad application, you can explore all the SAP HANA Interactive Education applications. Some applications need extra configuration before being executed. This extra configuration is explained in the SHINE documentation, and it is also described in the introduction dialog boxes when you start an application for the first time.

> Note:
> To use the SAP HANA SHINE launchpad, the user needs the role sap.hana.democenter.epm.roles::admin .

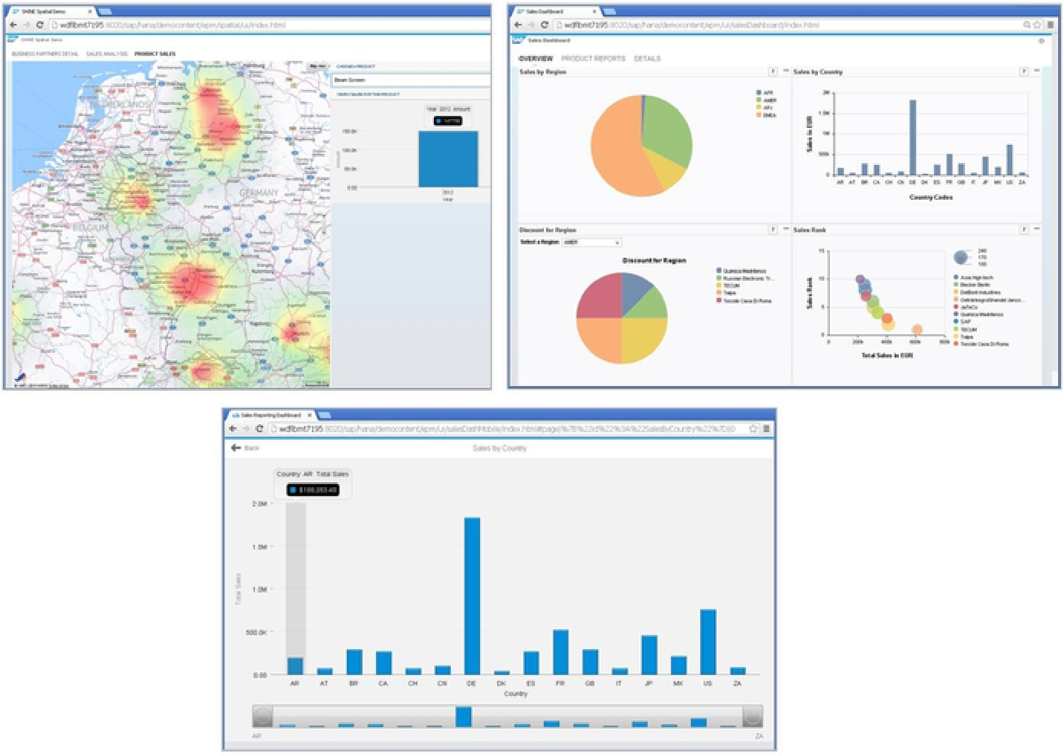

Figure 145: Some Examples from the SAP HANA SHINE Demo Apps

Import of SAP HANA SHINE Content

SAP HANA Application Lifecycle Management supports you in all phases of an SAP HANA application lifecycle, from modeling your product structure, to application development, transport, assembly, and installation.

To import delivery units, like SAP HANA SHINE, use the SAP HANA application lifecycle management tool. This is a browser-based tool that can be started by the URL: https://<host>::43$$/sap/hana/xs/lm/index.html, where $$= Instance Number.

> Hint:
> To use the SAP HANA application lifecycle management tool in a browser, you need the authorisation role: sap.hana.xs.lm.roles::Administrator .

SAP HANA application lifecycle management can deploy additional content to your SAP HANA system. You can use this to deploy the following SAP HANA application content:

- SAP HANA SHINE
- SAP HANA Live
- Customer Developments
- Third-party applications

Import SAP HANA Application Content

The figure, Import SAP HANA Application Content, shows how to start the SAP HANA application lifecycle management tool from the browser.

https://wdflbmt719x:4300/sap/hana/xs/lm/index.html

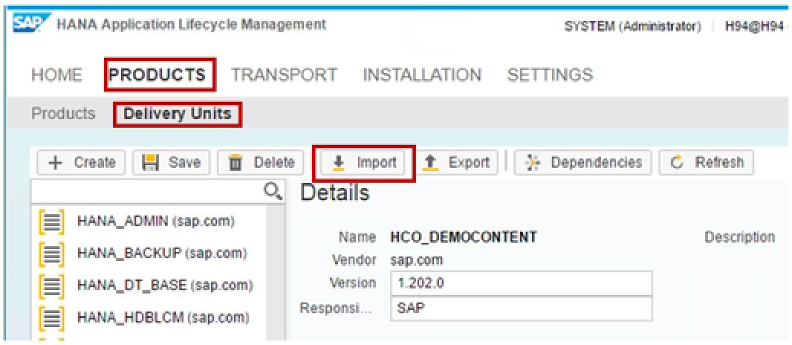

Figure 146: Import SAP HANA Application Content

Related Information

For additional information on the SAP HANA SHINE demo application, see the following:

- SAP HANA Interactive Education (SHINE) guide on http://help.sap.com/hana/SAP_HANA_Interactive_Education_SHINE_en.pdf
- SAP Note 1934114 : SAP HANA DEMO MODEL - SHINE Release & Information Note

LESSON SUMMARY

You should now be able to:

- Explain the installation and use of SAP HANA InteractiveEducation (SHINE)

Unit 7
Lesson 5

Explaining the Revision Strategy of SAP HANA

LESSON OVERVIEW

SAP HANA is a fast-changing and dynamic product. It is important to know the best way to keep your SAP HANA database up to date. This lesson provides information and recommendations about how to proceed.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Understand the SAP HANA maintenance strategy

SAP HANA Revision Strategy

As an IT architect, you need to set up an SAP HANA maintenance strategy that aligns with the customer's data center maintenance strategy.

SAP HANA Revision Strategy Terminology

In the context of the SAP HANA platform, there are several terms used to describe parts of the SAP HANA revision strategy, including the following:

- SAP HANA standard revisions
- SAP HANA maintenance revisions
- Support package stack (SPS)
- Support packages (SP)

Component	SPS 00	SPS 01	SPS 02
SAP HANA database	[Revision 000–00n]	[Revision 010–01n]	[Revision 020–02n]
SAP HANA client	[Revision 000–00n]	[Revision 010–01n]	[Revision 020–02n]
SAP HANA studio	[Revision 000–00n]	[Revision 010–01n]	[Revision 020–02n]
SAP HANA AFL	[Revision 000–00n]	[Revision 010–01n]	[Revision 020–02n]
SAP Host Agent	SPnn	SPnn	SPnn
SAP HANA Smart Data Access	SPnn	SPnn	SPnn
SAP HANA INA Toolkit	SPnn	SPnn	SPnn
...	...	...	...

Revisions

Support Packages (SP)

Support Package Stack (SPS)

This list had been simplified for illustration perspectives and does not show the entire bill of material of SAP HANA

Figure 147: SAP HANA Revision Terminology

SAP HANA Standard Revisions

The term "Standard Revision" or "Revision" refers to a maintenance package of an SAP HANA core component (SAP HANA Database, Studio, Clients, application function libraries, HWCC tool).

The SAP HANA standard revisions are shipments after release to customer (RTC) of the initial revision for a certain support package stack (SPS). For example 000, 001, or 011.

The standard revisions are shipped until the RTC of the next support package stack. The standard revisions provide incremental fixes based on the initial support package stack feature set.

SAP HANA Maintenance Revisions

The SAP HANA maintenance revisions are minor shipments after the last SAP HANA standard revision for certain support package stacks. For example 002.01, 012.01, or 022.01. Maintenance revisions are shipped after RTC the next version until the end of maintenance of the respective support package stacks.

Maintenance revisions provide fixes for the following bugs:

- Major bugs concerning critical functions in key SAP HANA scenarios (Business Suite on SAP HANA, SAP BW on SAP HANA, S/4 HANA, BW/4 HANA, SAP HANA Data Marts)
- Bugs without known workarounds

By upward compatibility, it is understood that when changing from software version n of a product to a product version n+1, existing functions of version n continue to be supported or replaced by equivalent functions in case a function was deprecated and is finally removed.

Data from version n can be transformed to and used with version n+1 without major changes. Interfaces of version n (APIs, Interfaces for other Systems/Products) remain unchanged.

> Note:
> For more information on possible update paths from SAP HANA maintenance to standard SAP HANA revisions, see SAP Note 1948334 .

Support Package Stack

The term support package stack (SPS) refers to the SAP HANA platform release, which bundles new features and functions through a set of component revisions and support packages (SPs). These are entirely validated and are provided as a single medium for consumption.

New functions are introduced once a year, every time a new SAP HANA support package stack (SPS) is released.

For easier handling, the numbering of SAP HANA SPS and revisions have been aligned. For example, revision 020 refers to the first SAP HANA revisions, which contains SAP HANA 2.0 SPS 02 capabilities.

Support Package

The term support package refers to all other parts of the SAP HANA platform that are non-core components for the SAP HANA database. That is, the SAP Host Agent or SAP HANA smart data access. These components are visible on SAP Support Portal support packages (SP).

Revisions and support packages have their own release cycle, and can be installed or upgraded separately.

SAP HANA Maintenance Strategy

The SAP HANA maintenance strategy is based on incremental, non-disruptive innovation updates.

New functions are introduced once a year, every time a new SAP HANA support package stack (SPS) is released. This happens at the beginning of April.

As of SAP HANA 2.0 SPS02, SAP is providing bug fixes and security patches for every SPS for two years after RTC. The last SPS of a major product version is in maintenance for five years after RTC.

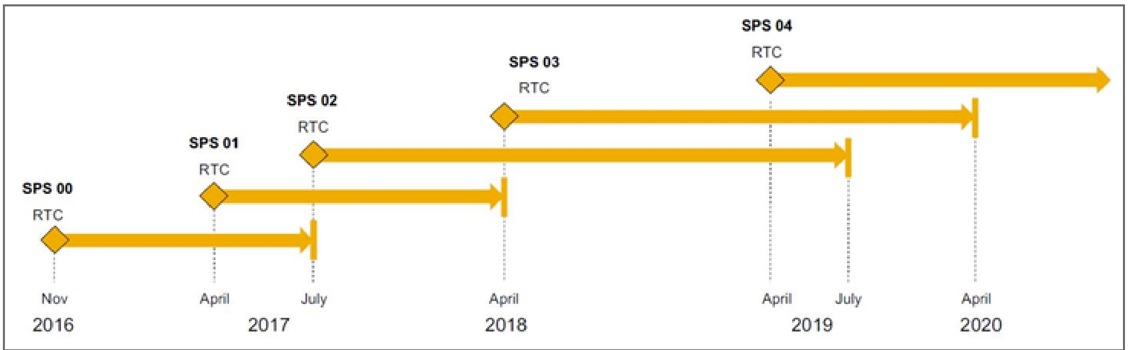

Figure 148: SAP HANA 2.0 Maintenance Strategy

SAP provides bug fixes and security patches for every SPS until the next SPS is released.

Note:
Adjust maintenance time lines and project go-live dates to this release schedule. See SAP Note 2378962 - SAP HANA 2.0 Revision and Maintenance Strategy for further details.

Because updates shipped for the SAP HANA platform are strictly downward compatible, earlier revisions can be removed from SAP Support Portal when a newer SAP HANA revision of the same SPS becomes available. Incompatible changes are considered for legal or security reasons, but are subject to a strict exception approval process.

The SAP HANA platform product remains in maintenance as long as any SAP business application releases built on top of SAP HANA are in mainstream maintenance, extended maintenance, or priority-one support.

Upgrade to SAP HANA 2.0

You can upgrade to SAP HANA 2.0 for systems running SAP HANA on SPS10 or newer. Systems running on SAP HANA SPS09 or lower first need to upgrade to an SAP HANA release after SPS10.

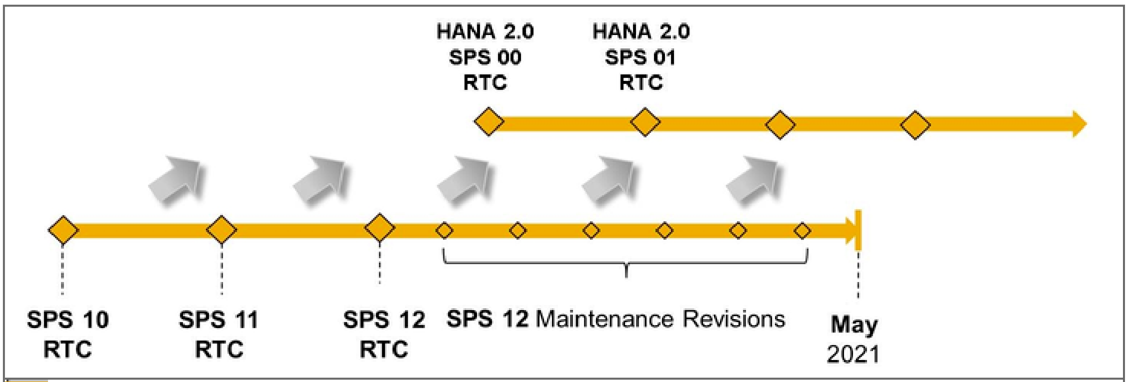

Figure 149: SAP HANA Upgrade Strategy

Note:
For details on possible update paths, see SAP Note 1948334 .

SAP recommends that you upgrade to the latest SAP HANA 1.0 SPS12 revision before upgrading to SAP HANA 2.0, so that you can use the capture and replay tool for regression tests.

Long-Term Maintenance Support for SAP HANA SPS12

Customers running mission-critical systems on SAP HANA 1.0 require a longer maintenance support window before upgrading to SAP HANA 2.0. To accommodate this, SAP will provide Maintenance Revisions for SAP HANA 1.0 SPS12 for a period of five years after RTC. This means that the end of maintenance for SPS12 is set for May 2021.

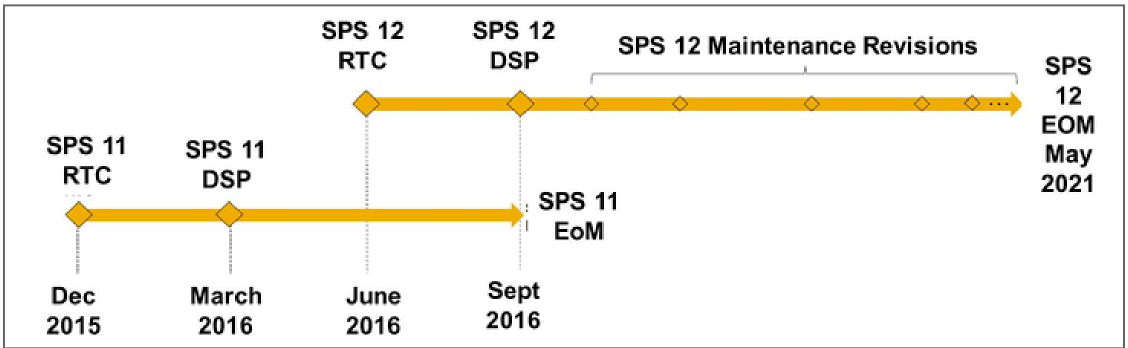

Figure 150: SAP HANA SPS12 Long-Term Support

The maintenance revisions for SAP HANA SPS12 are not scheduled. Rather, they are delivered on demand.

Note:
For the SAP HANA SPS12 long-term maintenance version, SAP provides regular upgrade paths from SPS12 to any newer SPS.

Regular Maintenance Using Standard Revisions

Customers can upgrade to any available revision of a new SPS at any time. Starting with SAP HANA 2.0 SPS00, SAP will no longer explicitly declare a DSP.

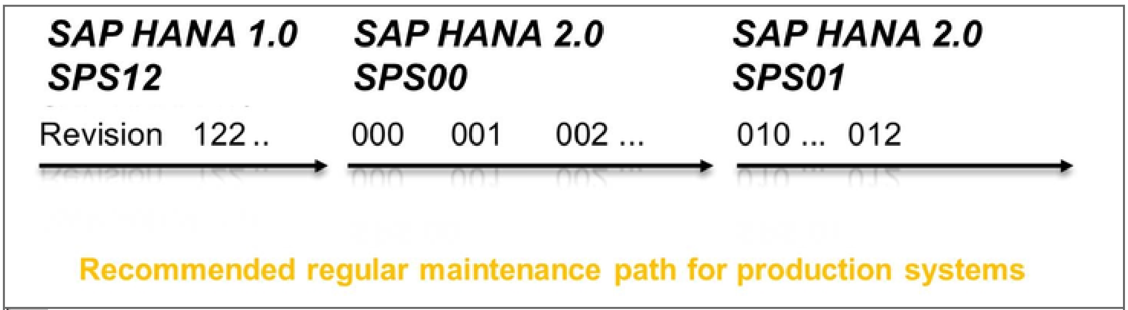

Figure 151: SAP HANA Regular Maintenance

> Note:
> DSP status is automatically given with release to customer (RTC) of a new SPS.

Extended Maintenance Using Maintenance Revisions

SAP provides SAP HANA standard revisions until the release to customer of the next SPS is declared. These standard revisions contain only incremental fixes.

SAP provides SAP HANA maintenance revisions between the RTC of the following and the over-next SPS. These maintenance revisions contain only major bug fixes, which are identified as relevant for production environments.

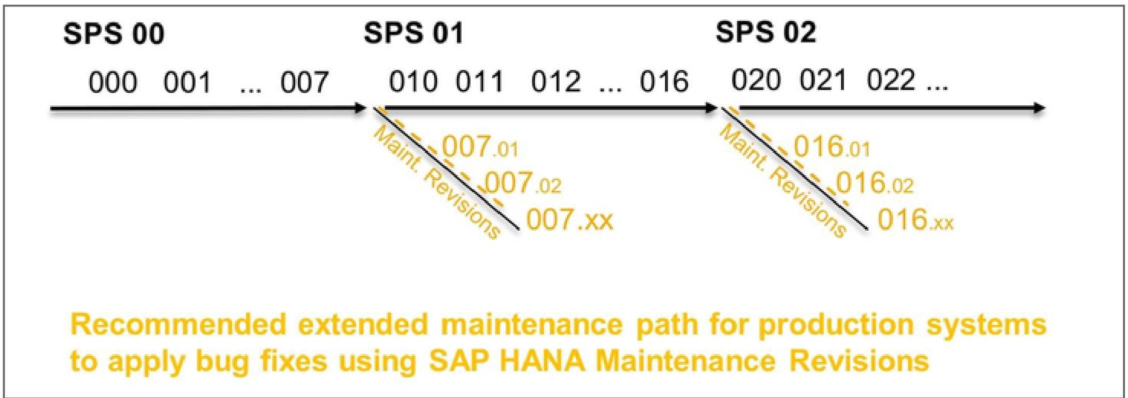

Figure 152: SAP HANA Extended Maintenance

The implementation of the highest SAP HANA revision available on SAP Support Portal benefits from incremental, but non-disruptive, improvements.

Capture and Replay Tool

SAP executes regression tests for applications shipped by SAP. Known regressions are documented in the SAP HANA release note of a revision. Customers should perform regression tests based on their actual setup and use of SAP applications.

When performing regression tests, it might be sufficient to perform only technical regression tests. Business acceptance tests are not required, as long as no new database features are used on application level.

Customers can contact SAP HANA product management to include their SAP HANA scenario in the regression tests.

Customer developments or third party applications can run full regression tests before updating to a newer SAP HANA revision.

Features

- Capture real system workload in productive environments
- Pre-process & replay captured workload on target systems
- Analyze runtimes and compare performance between system setups

Benefits

- Provides competitive replay, analysis and tuning capabilities for SAP HANA Database
- Helps ensure that landscape changes do not degrade system performance
- Enables real simulation of workload without the use of other 3rd party tools
- Reduces manual effort for testing changes in customer deployments

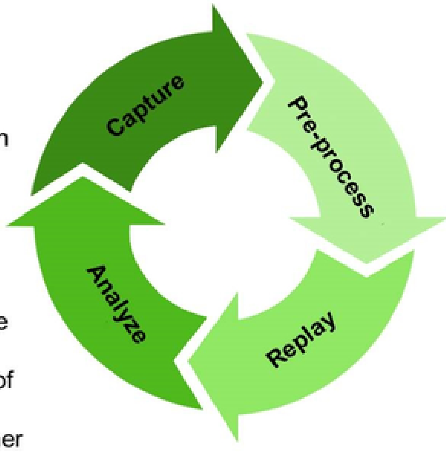

Figure 153: SAP HANA Capture and Replay

To support customers in performing regression tests, SAP provides the capture and replay tool. This tool allows you to capture real system workloads in productive environments, then prepare and replay the captured workload on the target system.

Important SAP Notes

For more information, see the following SAP Notes:

- SAP Note 1948334 : Recommended update paths for SAP HANA Maintenance Revisions
- SAP Note 2021789 : SAP HANA 1.0 Revision and Maintenance Strategy
- SAP Note 2378962 : SAP HANA 2.0 Revision and Maintenance Strategy
- SAP Note 2235581 : SAP HANA Supported Operating Systems

LESSON SUMMARY

You should now be able to:

- Understand the SAP HANA maintenance strategy

Unit 7

Learning Assessment

1. Which of the following SAP HANA license key types locks down the system when the current memory consumption of SAP HANA exceeds the licensed amount of memory?

   Choose the correct answer.

   - [ ] **A** Temporary license keys
   - [ ] **B** Unenforced license keys
   - [ ] **C** Enforced license keys

2. The resident hdblcm located on the SAP HANA installation media performs administrative tasks like installing and updating an SAP HANA database and its core components.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

3. Which of the following Inter-Service Communication properties binds the processes to entered address only and to all local host interfaces?

   Choose the correct answer.

   - [ ] **A** Global
   - [ ] **B** Internal
   - [ ] **C** Local

4. Which of the following features are part of SAP HANA SHINE?

   Choose the correct answers.

   - [ ] **A** SAP HANA simple info access
   - [ ] **B** SAP HANA SHINE application lifecycle management tool
   - [ ] **C** Outbound XSJS
   - [ ] **D** Core Data Services (CDS) / HDBDD

5. In the context of the SAP HANA maintenance strategy, which of the following terms is connected to bundle new features and functions to the SAP HANA Platform?

   Choose the correct answer.

   - [ ] **A** SAP HANA standard revisions
   - [ ] **B** SAP HANA maintenance revisions
   - [ ] **C** Support package stack (SPS)
   - [ ] **D** Support packages (SP)

Unit 7

Learning Assessment - Answers

1. Which of the following SAP HANA license key types locks down the system when the current memory consumption of SAP HANA exceeds the licensed amount of memory?

   Choose the correct answer.

   [ ] **A** Temporary license keys

   [ ] **B** Unenforced license keys

   [X] **C** Enforced license keys

   Correct! With enforced license keys, the system is locked down when the current memory consumption of SAP HANA exceeds the licensed amount of memory, plus some tolerance. If this happens, you can either restart SAP HANA, or install a new license key that covers the amount of memory in use. Temporary license keys are automatically installed by the SAP HANA system. This license is valid for 90 days. After 90 days, the license expires and the system is locked down. Once you have installed a valid permanent license, you can use your system until this license expires. With unenforced license keys, the operation of SAP HANA is not affected if its memory consumption exceeds the licensed amount of memory. Read more on this in the lesson Performing Post-Installation Steps (Unit 7, Lesson 1) of the course HA200_14.

2. The resident hdblcm located on the SAP HANA installation media performs administrative tasks like installing and updating an SAP HANA database and its core components.

   Determine whether this statement is true or false.

   [ ] True

   [X] False

   Correct! The resident hdblcm performs administrative tasks for the database where it's embedded. The resident hdblcm installs or updates additional SAP HANA components, but it can't update the SAP HANA database executable files. Read more on this in the lesson Updating SAP HANA (Unit 7, Lesson 2) of the course HA200_14.

3. Which of the following Inter-Service Communication properties binds the processes to entered address only and to all local host interfaces?

   Choose the correct answer.

   - [ ] **A** Global
   - [x] **B** Internal
   - [ ] **C** Local

   Correct! Internal: Binds the processes to this address only and to all local host interfaces. Global: Binds the processes to all interfaces. Local: Opens the communication ports for internal usage on the local interfaces. This configuration is only an option for single installations because the server is not accessible from outside. Read more on this in the lesson Using the Resident SAP HANA Database Lifecycle Manager (HDBLCM) Tool (Unit 7, Lesson 3) of the course HA200_14.

4. Which of the following features are part of SAP HANA SHINE?

   Choose the correct answers.

   - [x] **A** SAP HANA simple info access
   - [ ] **B** SAP HANA SHINE application lifecycle management tool
   - [x] **C** Outbound XSJS
   - [x] **D** Core Data Services (CDS) / HDBDD

   Correct! The SAP HANA simple info access (SINA) API is a client-side or front-end JavaScript API for developing browser-based search UIs. SAP HANA extended application services include a server-side JavaScript API (Outbound API) that allows access to a defined HTTP destination. Core Data Services (CDS) / HDBDD is a new infrastructure for defining and using semantically rich data models in SAP HANA. CDS uses a data definition language (DDL), a query language (QL), and an expression language (EL). It includes write operations, transaction semantics, constraints, and more. SHINE does not provide any application lifecycle management tool. The SAP HANA application lifecycle management can deploy additional content like SAP HANA SHINE to your SAP HANA system. Read more on this in the lesson Using SAP HANA Interactive Education (Shine) (Unit 7, Lesson 4) of the course HA200_14.

5. In the context of the SAP HANA maintenance strategy, which of the following terms is connected to bundle new features and functions to the SAP HANA Platform?

   Choose the correct answer.

   - [ ] **A** SAP HANA standard revisions
   - [ ] **B** SAP HANA maintenance revisions
   - [x] **C** Support package stack (SPS)
   - [ ] **D** Support packages (SP)

   Correct! Support package stack (SPS) refers to the SAP HANA platform release, which bundles new features and functions through a set of component revisions and support packages (SPs). New functions are only introduced twice a year, every time a new SAP HANA support package stack (SPS) is released. This happens at the beginning of April and the end of July. The SAP HANA standard revisions provide incremental fixes based on the initial support package stack feature set. SAP HANA maintenance revisions provide fixes for major bugs concerning critical functions in key SAP HANA scenarios (SAP S/4HANA, SAP BW/4HANA, SAP HANA data marts) and Bugs without known work-arounds. Support packages refer to all other parts of the SAP HANA platform that are noncore components for the SAP HANA database. That is, the SAP Host Agent or SAP HANA smart data access. Read more on this in the lesson Explaining the Revision Strategy of SAP HANA (Unit 7, Lesson 5) of the course HA200_14.

UNIT 8 Adminstration Tools

UNIT OBJECTIVES

- Explain the available Administration Tools
- Use the HDBSQL Command Line Tool
- Work with the DBA Cockpit
- Manage Complex SAP Landscapes
- Manage SAP Landscapes powered by SAP HANA
- Use SAP HANA Studio

Unit 8
Lesson 1

Explaining the Administration Tools

LESSON OVERVIEW

This lesson provides an overview of the tools that can be used for the administration of an SAP HANA database.

Business Example

Administrators of SAP HANA systems need to know about the tools for administration and monitoring, how they are integrated with SAP HANA, and what their capabilities are.

Previously, the most common tool was the SAP HANA studio, but, with SPS09, the SAP HANA Cockpit is used increasingly.

Because several tools are available for the administration of SAP HANA, it is important to know the differences between these tools.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explain the available Administration Tools

SAP HANA Administration Tools: Overview

Table 6: SAP HANA Administration Tools: Overview

The table provides an overview of the most important SAP HANA administration tools. A thorough overview is available in the SAP Administration guide.

Tool	Description
SAP HANA cockpit	SAP HANA cockpit provides a single point of access to a range of tools for the administration and detailed monitoring of multiple, individual, and tenant SAP HANA databases. It also integrates the SQL development functions required by administrators. SAP HANA cockpit, a Web-based HTML5 user interface that you access through a browser, runs on SAP HANA extended application services, advanced model. You can use the cockpit to monitor and manage systems running SAP HANA 2.0 or SAP HANA 1.0 SPS 12.
SAP HANA studio	The SAP HANA studio supports core administration and monitoring of SAP HANA databases, including system configuration, user management, and performance monitoring capabilities. It also supports SAP HANA development of content, including modeled views and stored procedures general system administration and monitoring tasks.
SAP HANA HDBSQL	SAP HANA HDBSQL is a command line tool for executing commands on SAP HANA databases.

Tool	Description
SAP Solution Manager	SAP Solution Manager is a central hub for the holistic lifecycle management of your SAP solution landscape. It contains central alerting infrastructure and business process analytics, end-to-end root cause analysis, monitoring and unified alert inbox for the entire solution landscape, guided issue resolution, and business process reporting.
SAP DBA cockpit	The DBA Cockpit is a platform-independent tool provided by ABAP systems, which you can use to monitor and administer your SAP HANA database.
SAP Landscape Management	SAP Landscape Management is a powerful automated solution designed to simplify, centralize, and orchestrate the management and operations of the entire application landscape powered by SAP HANA. With the enterprise edition of SAP Landscape Management, you can perform various operations on SAP HANA instances, such as start, stop, monitor, or system replication operations.

SAP HANA Administration Tools: Overview

The SAP HANA Studio is both the central development environment and the main administration tool for the SAP HANA database. Additionally, with SPS09, a first version of the web-based SAP HANA cockpit was introduced for monitoring SAP HANA. The SAP HANA cockpit can also be used on mobile devices. Furthermore, SAP HANA is fully integrated into SAP Solution Manager .

Following the cloud strategy of SAP, SAP HANA offers Web-based tools for monitoring and administration.

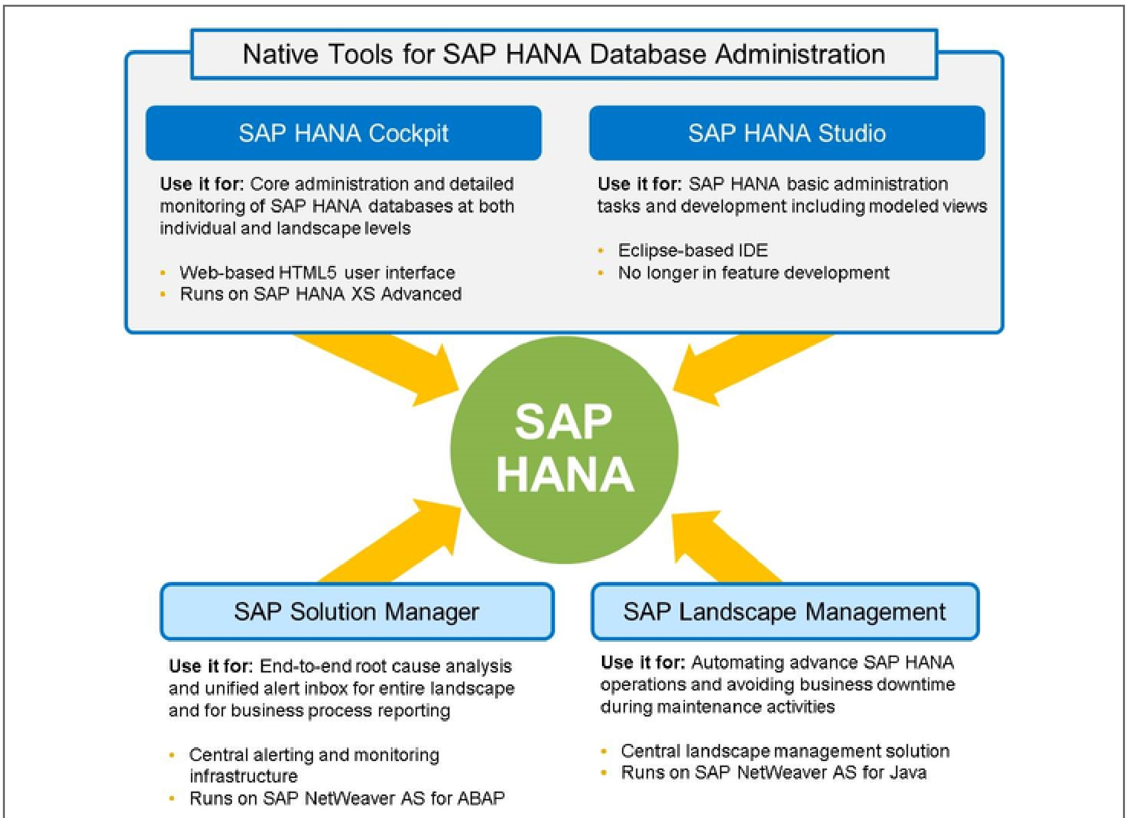

Figure 154: SAP HANA Administration Tools: Overview

Note:
In the future, the administration perspective of SAP HANA studio will be replaced with SAP Database Control Center and SAP HANA cockpit.

LESSON SUMMARY
You should now be able to:

- Explain the available Administration Tools

Unit 8 Lesson 2

Using the HDBSQL Command Line Tool

LESSON OVERVIEW

This lesson describes the use of HDBSQL and some of the most important commands.

Business Example

SAP HANA HDBSQL is a command line tool for entering and executing SQL statements, executing database procedures, and querying information about SAP HANA databases. Administrators can execute statements from a command line or schedule scripts that access the SAP HANA database.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Use the HDBSQL Command Line Tool

SAP HANA Connection with HDBSQL

SAP HANA HDBSQL is a command line tool for executing commands on SAP HANA databases.

With SAP HANA HDBSQL , you can execute SQL statements and database procedures, and query information about the database and database objects. SAP HANA HDBSQL is installed with the SAP HANA software. It accesses databases both on your local computer and on remote computers.

Call SAP HANA HDBSQL with the command `hdbsql [options]` from the following location: `/usr/sap/<SIDY/HDB<instance>/exe`. You can execute individual commands interactively or non-interactively. You can also import commands from a file and execute them in the background.

Features of HDBSQL

HDBSQL has the following features:

- Execute SQL statements
- Execute database procedures
- Request information about the database catalog
- Execute shell commands
- Execute commands (command syntax and options)
- Overview of all HDBSQL call options
- Overview of all HDBSQL commands

To use HDBSQL interactively and to execute some commands, log on to the database as a database user.

> Note:
> The user logging on must be a database user. If you do not specify a username and password of a database user, you can log on using Kerberos authentication.

Options to Connect to an SAP HANA System

The following options are available:

- One-step logon with username and password
  - Specify credentials in the start command of hdbsql
- Two-step logon with username and password
  - Start hdbsql first
  - Connect to the system

Logon Option	How
One-step logon with username and password	**Enter the following command:** hdbsql [<options>] -n <database_host> -i <instance_id> -u <database_user> -p <database_user_password>
Two-step logon with username and password	**1. Start HDBSQL by entering the following command:** hdbsql [<options>] **2. Log on to the database by entering the following command:** \c [<options>] -n <database_host> -i <instance_id> -u <database_user> -p <database_user_password>

Figure 155: Options to Connect to an SAP HANA System

> Note:
> In a multiple-container system, the option –d specifies the name of the multitenant database container.

Logon with HDBSQL: Example

In this example, HDBSQL is used to connect to a SAP HANA system with instance number 01 on the localhost. Database user MONA is specified with password RED.

To start HDBSQL and log on to the SAP HANA database as user MONA with the password RED, enter the following command:

```
hdbsql -n localhost -i 1 -u MONA -p RED
```

To display general information about the database, enter the following command:

```
\s
host: wdfd00245293a:30015
database: ORG
user: SYSTEM
kernel version: 1.50.00.000000
SQLDBC version: libSQLDBCHDB 1.50.00 Build 0000000-0120
autocommit: on
```

Figure 156: Logon with HDBSQL: Example

Execution of Commands

HDBSQL commands can be executed in interactive and noninteractive mode. To execute commands, log on to the database.

Mode	Steps
Interactive (session) mode	1. **Call HDBSQL by entering the following command:** `hdbsql [<options>]` 2. **Enter the following command and press** `Enter` HDBSQL executes the command 3. **Exit HDBSQL by entering the following command:** `exit | quit | \q`

Non-interactive (command) mode	**Run a command:**
	`hdbsql [<options>] <command>`
	HDBSQL executes the command and exits
	
	Run multiple commands from a batch file:
	`hdbsql [<options>] -I <file>`
	HDBSQL imports the commands from the specified file and processes them in the background.

Note
To execute an SQL statement or a database procedure as a command, place the statement or procedure in quotation marks.

Figure 157: HDBSQL: Executing Commands

In addition to executing commands individually, you can execute multiple commands from a batch file. HDBSQL imports the commands from the specified file and processes them in the background.

> Note:
> When you execute from a batch file, the AUTOCOMMIT mode is activated by default. If you deactivate it, the batch file must contain an explicit COMMIT statement to ensure that HDBSQL executes the SQL statements immediately after the batch file is imported.

HDBSQL Commands

The figure, HDBSQL Commands, shows a subset of HDBSQL commands.

> Note:
> For a detailed description of the features, see the SAP HANA Administration Guide.

Command	Description
`\?`	Displays all HDBSQL commands
`\h[elp]`	
`\a[utocommit] [ON/OFF]`	Switches AUTOCOMMIT mode on or off
`\al[ign] [ON/OFF]`	Switches formatted output of the results of SQL statement on or off
`\es[cape] [ON/OFF]`	Switches the escape output format on or off
`\c[onnect]`	Logs a user onto the database
`\dc [PATTERN]`	Lists all table columns that correspond to the `PATTERN` HDBSQL lists only those tables to which the current user has access
`\de [PATTERN]`	Lists all the indexes of database objects that correspond to the `PATTERN`
`\di[sconnect]`	Logs the user off from the database
`\dp [PATTERN]`	Lists all database procedures that correspond to the `PATTERN`

Figure 158: HDBSQL Commands

> Note:
> Instead of an HDBSQL command, you can also enter an SQL statement or a database procedure.

Secure User Store (hdbuserstore)

The secure user store (`hdbuserstore`) is a tool installed with the SAP HANA client. Use it to store connection information to SAP HANA systems securely on the client. This process ensures that client applications can connect to SAP HANA without users having to enter this information. It is used by scripts connecting to SAP HANA.

The secure user store allows you to store SAP HANA connection information, including user passwords, securely on clients. In this way, client applications can connect to SAP HANA without the user having to enter host name or logon credentials. You can also use the secure store to configure failover support for application servers in a 3-tier scenario (for example, SAP Business Warehouse) by storing a list of all the hosts that the application server can connect to.

> Note:
> The secure user store can only be used for SQLDBC and JDBC-based connections. The SAP HANA studio does not use the SAP HANA secure user store. Instead, it uses the Eclipse secure storage. For more information, see the Eclipse documentation.

hdbuserstore Program

The secure user store is installed with the SAP HANA client package. After you install the SAP HANA client, the hdbuserstore program is located in one of the following directories:

- /usr/sap/hdbclient (Linux or UNIX)
- %SystemDrive%\Program Files\sap\hdbclient (Microsoft Windows)

The secure user store is user-specific, so only the operating system user who owns the corresponding secure store file can access the secure store. However, you can, with the appropriate operating system privileges, manage another user's secure store. This option is needed, for example, to manage the connection details for ABAP on Microsoft Windows because the application server runs under a different user (SAPService<SAPSID> instead of <sapsid>adm).

hdbuserstore Commands

Use the hdbuserstore program to store and manage connection information in the secure store. For more information about the available commands, see SAP HANA Security Guide .

Example:
Sample hdbuserstore Commands

Action	Command	Examples
Create a user key in the user store and store the password under this user key	hdbuserstore SET <user_key> <env> <user_name> <password>	hdbuserstore SET millerj localhost:30115 MILLERJ Student00
List all available user keys (passwords are not displayed)	hdbuserstore LIST <user_key>	hdbuserstore LIST millerj **The following information is displayed:** • KEY: millerj • ENV: localhost:30115 • USER: JohnMiller
Call hdbsql **with the user key**	hdbsql -U <user_key>	hdbsql -U millerj
Configure failover support for application servers by using hdbuserstore **to specify a list of host names that the server can connect to**	hdbuserstore SET default "<hostname_node1>:3<inst>15, ….. <hostname_node(n)>:3<inst>15" <sapsid> <Password>	hdbuserstore SET default "Id9490:33315;Id9491:33315;Id9492:33315;Id9493:33315" <sapsid> <Password>

Figure 159: Using hdbuserstore for Connecting to SAP HANA: Example

Note:
The tool did not check if the user really exists.

You can also use the `hdbuserstore` to configure failover support for application servers in a 3-tier scenario (for example, SAP NetWeaver Business Warehouse) by storing a list of all the hosts that the application server can connect to.

LESSON SUMMARY
You should now be able to:

- Use the HDBSQL Command Line Tool

Unit 8
Lesson 3

Working with the DBA Cockpit

LESSON OVERVIEW

This lesson describes the functions of the DBA Cockpit.

Business Example

The DBA Cockpit in SAP Solution Manager provides a detailed insight into the status of the database. This uses the same data that you can see in the SAP HANA studio for your in-memory database, but the DBA Cockpit also supports other databases. If you have heterogeneous databases in your environment because your business applications still run on traditional databases, the DBA Cockpit enables you to use the same tool for the different databases.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Work with the DBA Cockpit

DBA Cockpit: Overview

The DBA Cockpit is a platform-independent tool that monitors and administers databases from an SAP NetWeaver AS for ABAP environment. It allows you to monitor SAP HANA databases remotely using the SAP Solution Manager. For SAP HANA databases, the DBA cockpit offers many of the same functions as the SAP HANA Studio. However, you can also use the DBA Cockpit to schedule database backups.

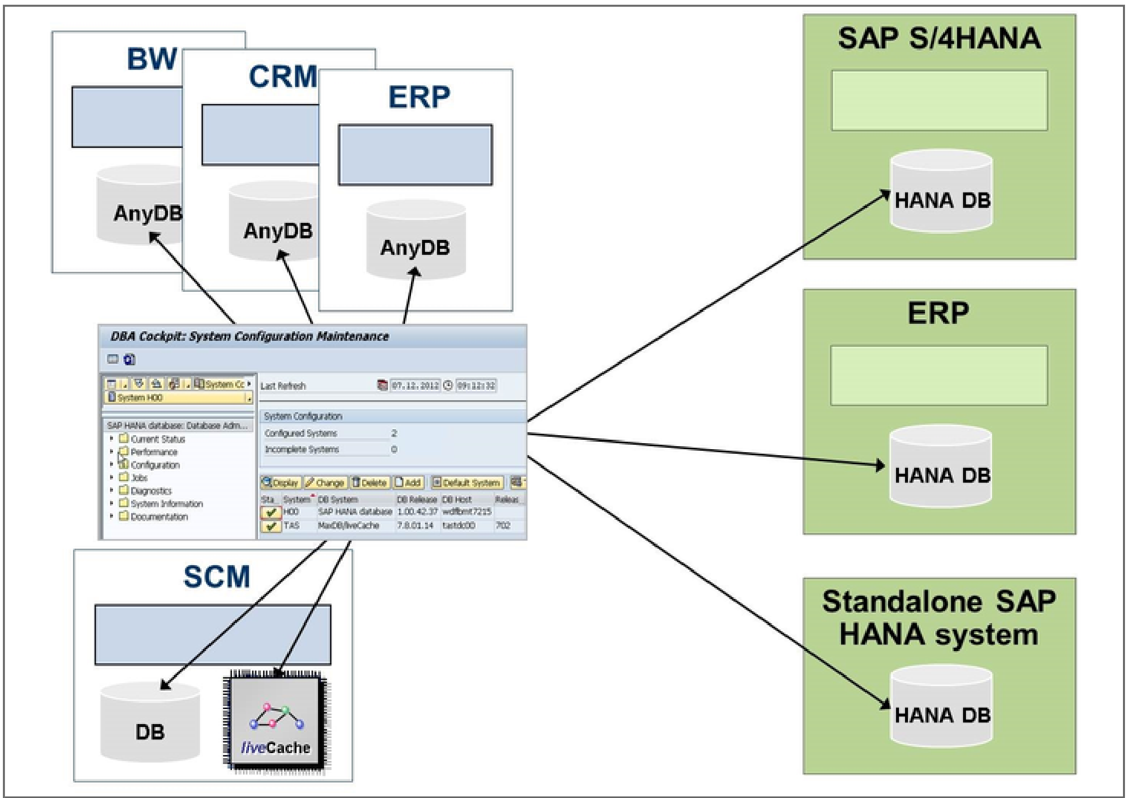

Figure 160: DBA Cockpit: Remote Monitoring

To start the DBA Cockpit, use transaction code DBACOCKPIT. Alternatively, you can use the transaction codes for specific SAP monitoring tools to open the corresponding application within the DBA Cockpit.

Screen Layout in the DBA Cockpit

The initial screen of the DBA Cockpit is divided into the different areas as indicated in the figure, Screen Layout in the DBA Cockpit .

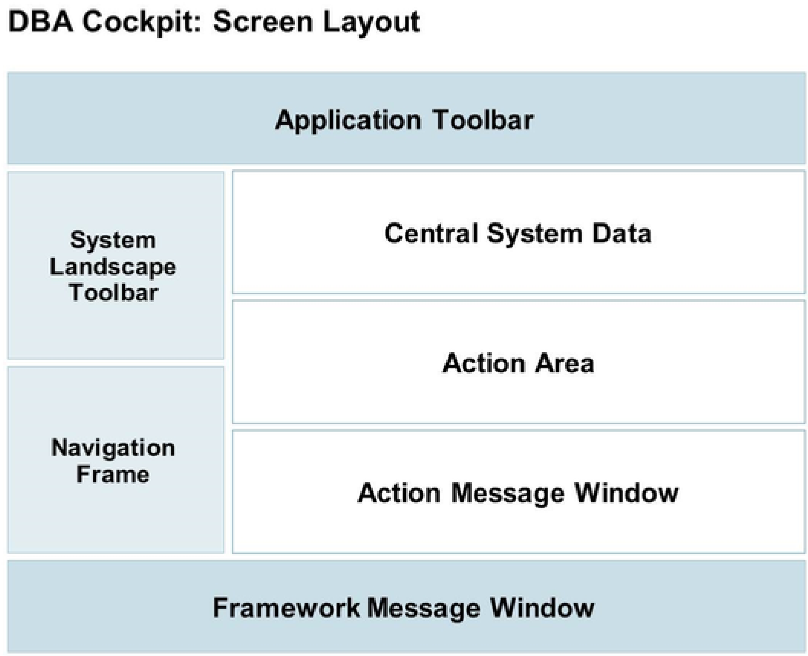

Figure 161: Screen Layout in the DBA Cockpit

The initial screen of the DBA Cockpit is divided into the following areas:

- Application toolbar

  Provides basic functions. For example, to display or hide the System Landscape toolbar and the navigation frame.

- System landscape toolbar

  Provides central functions to manage the system landscape. For example, to manage database connections and to choose the system to monitor.

- Navigation frame

  Provides quick access to a range of analysis information. For example, performance monitoring, space management, and job scheduling.

- Framework message window

  The framework message window contains a complete history of the messages sent during the session.

  The navigation frame on the left shows the available functions. For example, Overview and Alerts under the `Current Status` folder, INI files under the `Configuration` folder, Performance, Jobs, Diagnostics, System Information, and so on.

- Central system data

  Provides information, such as the time of last refresh, database startup time, and the name of the database.

- Action area

  Displays the details of the currently selected action.

- Action message window

Displays additional information for the selected action.

Some functions are only available in particular tools and not in others. For example, the DBA Planning Calendar is only available in the DBACOCKPIT, and is not yet available in the SAP HANA studio.

DBA Cockpit for SAP HANA: Functionality Overview

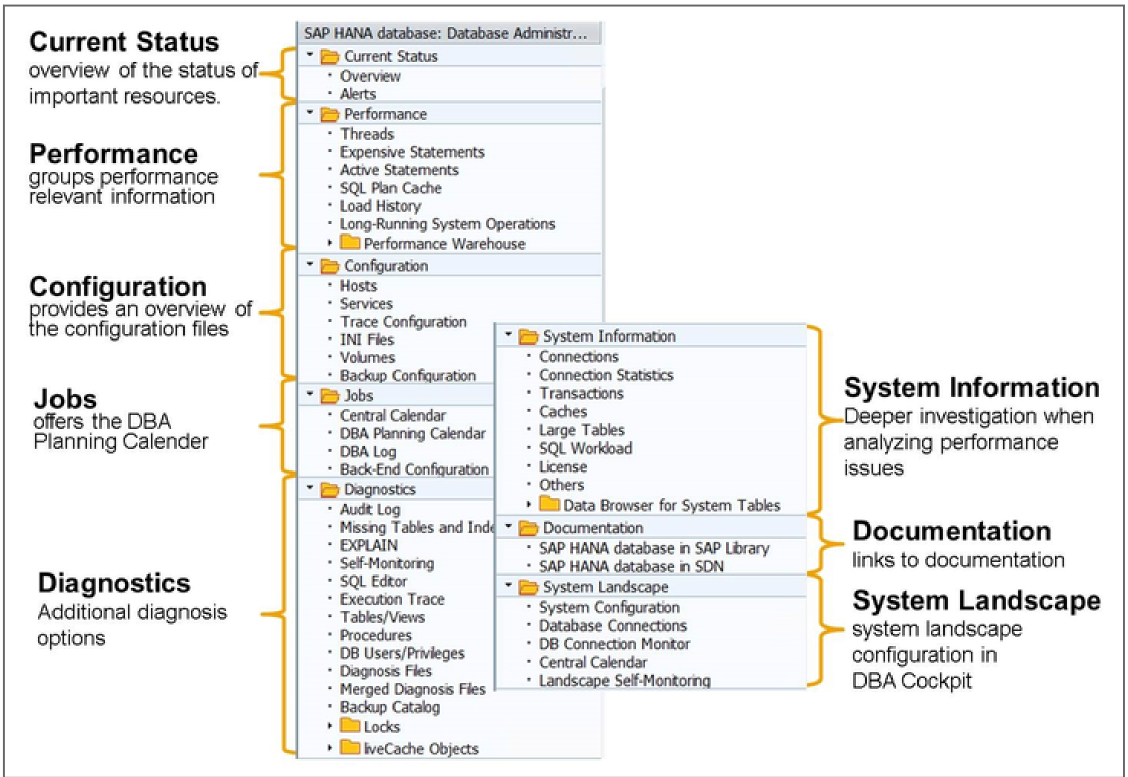

Figure 162: DBA Cockpit for SAP HANA Database

The figure shows DBA Cockpit for SAP HANA Database .

Monitor: Current Status

The Current Status monitor provides an overview of the statuses of the most important database resources.

The section overview provides information about the following:

- The status of the available disk space and physical memory
- The status of the services
- The time at which the database was started
- Current alerts
- Memory and CPU consumption from the SAP HANA database
- Disk consumption from the SAP HANA database
- Memory and CPU consumption from the operating system
- Disk space used on a particular host, from the operating system

DBA Cockpit Compared to SAP HANA studio: Overview

To use the DBA Cockpit, choose Current Status → Overview . You can display the overall SAP HANA system status, for example, CPU, Disk, and Memory allocation. The same values are shown under the Overview tab in the SAP HANA Studio Administration Console .

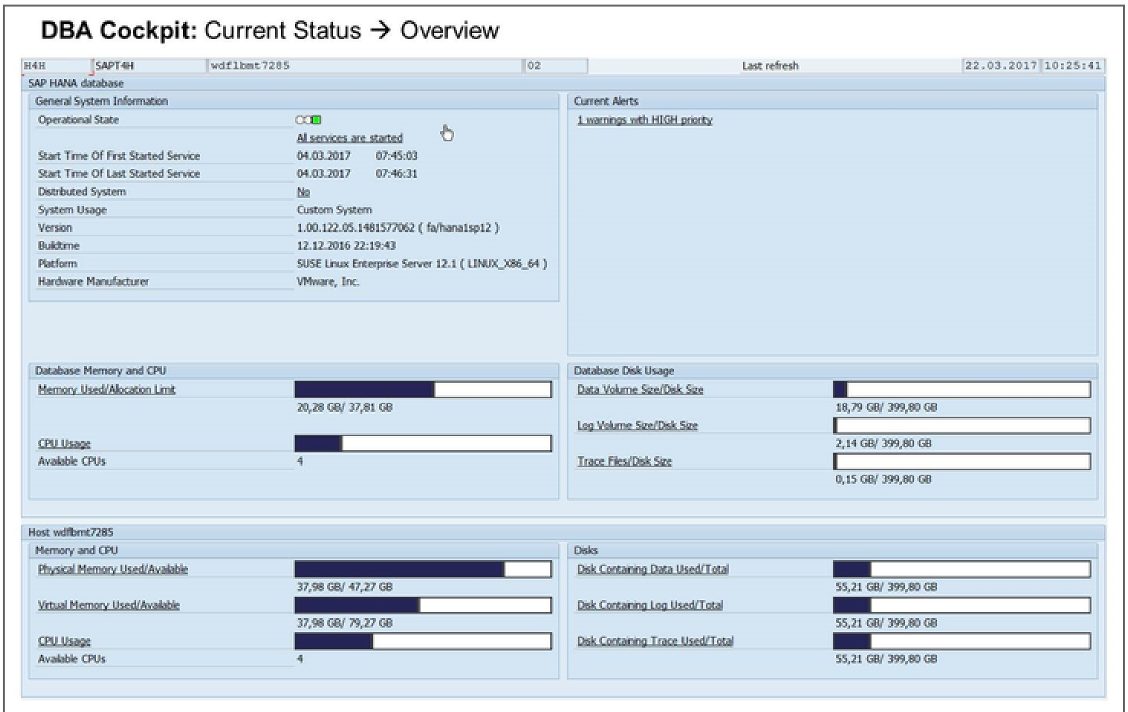

Figure 163: DBA Cockpit Compared to SAP HANA studio: Overview

> Note:
> Even if the database is unavailable, the Overview section is always available, and jobs can always be scheduled.

The other sections in this monitor provide more information on the status of the system services, currently active threads, and the use of disks and volumes.

Monitor: Performance

You can analyze performance data of your database system using the Performance Warehouse. As a prerequisite, it requires an SAP Solution Manager system with the Solution Manager Diagnostics (SMD) enabled.

In the Performance Warehouse, all of the relevant performance indicators that are collected by the DBA Cockpit are stored in an SAP Business Intelligence (BI) system. This SAP BIsystem is used by the Solution Manager Diagnostics (SMD) back end of an SAP Solution Manager system. SMD already uses this SAP BI to store the workload data of SAP applications. To configure the extraction of data into the SMD BI, use the SMD Setup Wizard .

Based on this architecture, the DBA Cockpit uses SAP BI technology to provide reports for performance analyses, which you can customize according to your needs. All collected data has a time dimension, so you can analyze the database performance for any point in time or over a specified period.

Almost all reports are displayed as a chart to visualize the key performance indicators (KPIs). In addition, there is a detailed table view. To navigate within these reports, use the SAP BI drilldown feature. To make you immediately aware of performance issues, violations of performance thresholds are highlighted based on predefined SAP BI exceptions.

The Performance Warehouse is shipped with predefined content that you can use to create your own reports, according to your needs.

Monitor: Diagnostics

The Diagnostics node contains the following sections:

- Audit Log

  The DBA audit log records all actions that change the database. For example, starting, stopping, and reconfiguring services, changes to parameters in configuration files, deletion of trace files, and table imports

- Missing Tables and Indexes

  Missing Tables and Indexes shows the differences between the database in the SAP system and the ABAP dictionary

> Note:
> The Missing Tables and Indexes function is only available for local systems or for ABAP systems, for which an additional RFC destination has been assigned. It is not available for remote systems.

- EXPLAIN

  EXPLAIN shows the execution plan for SELECT, INSERT, UPDATE, or DELETE statements

- SQL Editor

  Use the SQL Editor to execute SQL statements

- Tables or Views

  You can display a table view, a view, or a monitoring view

- Diagnosis Files

  Used for SAP HANA databases that are offline and that cannot be reached by SQL

- SQLDBC Trace

  Activating, deactivating, and analyzing the SQLDBC Trace

- Database Trace

  Activating, deactivating, and analyzing the SQLDBC Trace

Monitor: System Information

The information displayed in the sections of this monitor can be helpful for analyzing performance issues.

- **Connections**

  Provides detailed information about open connections

- **Transactions**

Displays open transactions

- **Connection Statistics**

  Provides information about open connections, such as network input/output statistics

- **Caches**

  Provides information about caches created by the SAP HANA database. The Total Size column shows the size of the available caches.

- **Query Cache**

  Provides information about the query cache, which is where executed SQL statements are cached

- **Large Tables**

  Provides information about the largest tables in the SAP HANA system. This information is helpful for analyzing performance and system dimensions. You can see the table sizes in the main memory, the delta sizes, and the fastest growing tables.

- **SQL Workload**

  Provides an overview of statements that are executed

- **Data Browser for System Tables**

  Provides an overview of the tables in the SYS schema and the _SYS_STATISTICS schema. These tables contain data that is useful for analyzing system performance. To display the content of a table, select the table and choose Display Table/View Content .

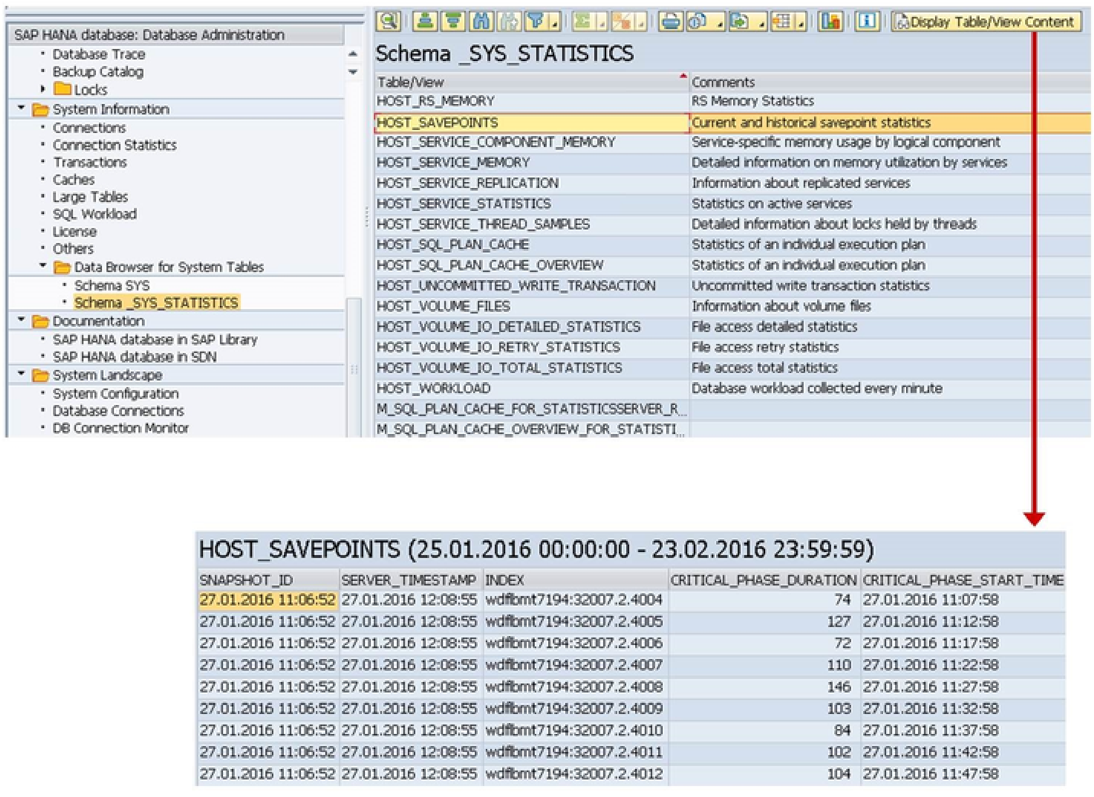

Figure 164: System Information: Data Browser for System Tables

Integrating SAP HANA as a Remote Database

With SAP Solution Manager Version 7.10 Support Package 4 or higher, SAP HANA can be integrated into monitoring as a remote database and included in the end-to-end database analysis.

The prerequisites for the Solution Manager integration are as follows:

- Installation of the SAP HANA client software
- Supported kernel version (at least 7.20 Patch 100)
- SAP HANA DBSL (minimum 7.20 Patch 110)
- SAP Host Agent (at least 7.20 Patch 84)
- SAP Solution Manager Diagnostics Agent

You can also refer to the following SAP notes:

- SAP Note 1664432 : DBA Cockpit: SAP HANAdatabase as remote database
- SAP Note 1612172: Additional corrections for setting up the DBA Cockpit using the SAP Solution Manager
- SAP Note 1672429 : Corrections with regard to the technical system HANA DATABASE for the setup in the SAP Solution Manager
- SAP Note 1721598 : Corrections regarding the technical system HANA DATABASE; the system also saves required attributes in the Landscape Management Database (LMDB)

Adding an SAP HANA System as Remote Database

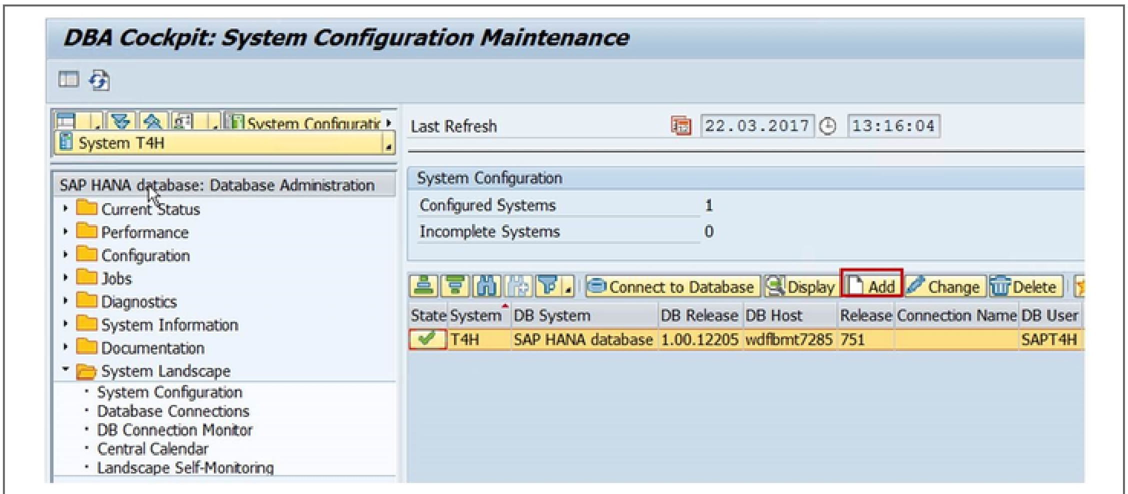

Figure 165: Adding an SAP HANA System as Remote Database

If the prerequisites are met, an SAP HANA system can be added to DBA Cockpit by choosing the Add button.

Database Connections: Add Connection Entry

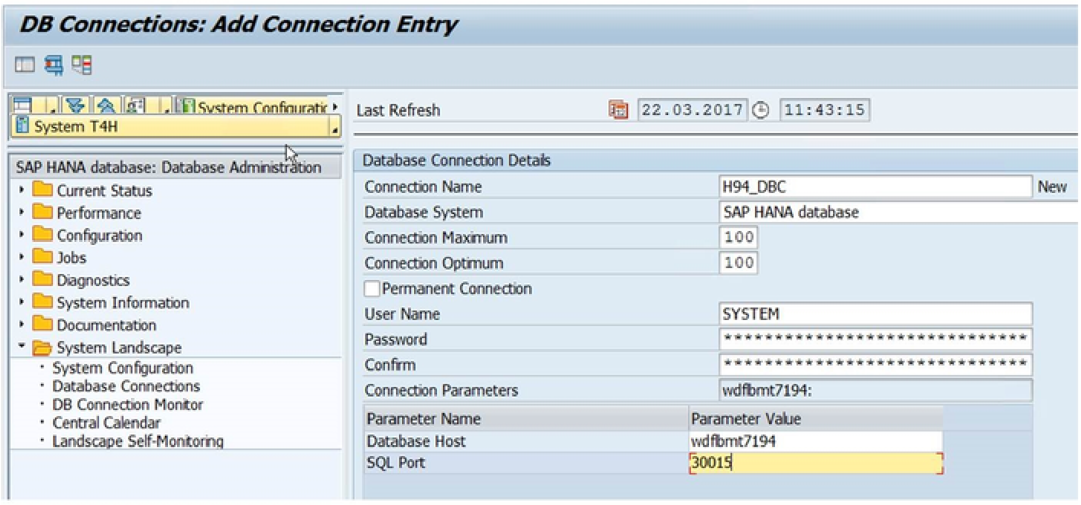

Figure 166: Database Connections: Add Connection Entry

To connect to a remote SAP HANA database, first add a respective secondary database connection.

Specify the following parameters:

- Connection Name
- Database System (SAP HANA database)
- User Name (SAP HANA database user with at least monitoring privileges)
- Password (SAP HANA database user password)
- Database Host (hostname of SAP HANA database)
- SQL Port (3<instancenumber>15)

Add System Entry

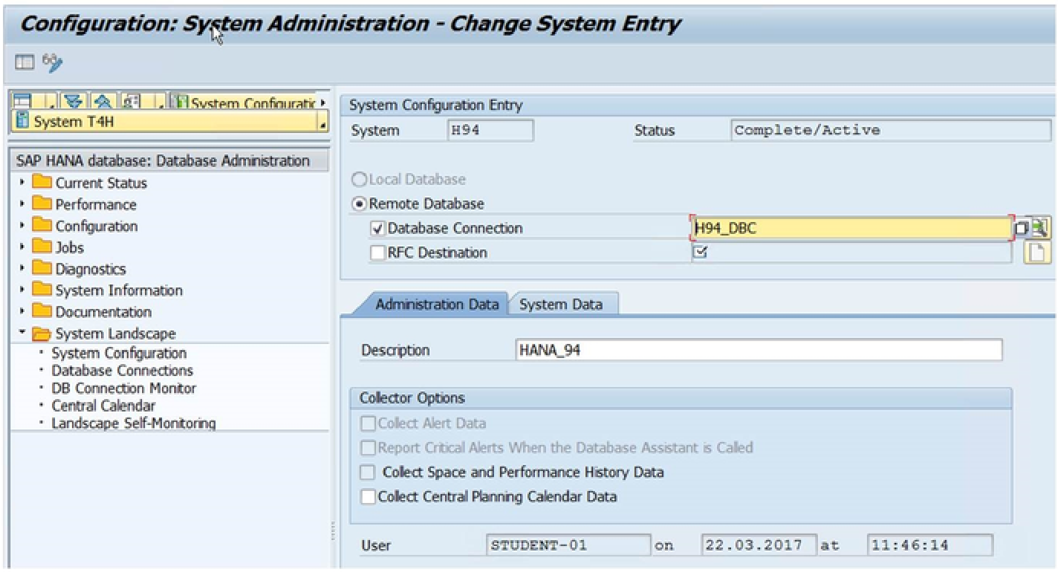

Figure 167: Add System Entry

A new system entry for the SAP HANA database can then be added. This entry refers to the database connection created in the previous step.

Result of the Configuration

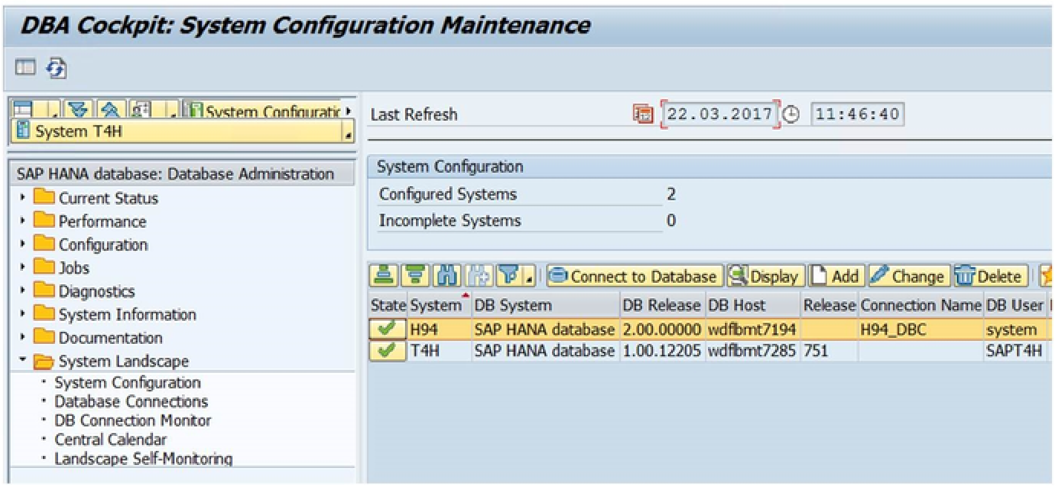

Figure 168: Result of the Configuration

After choosing Save , the DBA Cockpit stores the information and tries to connect to the newly added system. The SAP HANA system appears in the System Landscape Toolbar (H00 in the figure). To display the available functionality, choose the SAP HANA system.

LESSON SUMMARY

You should now be able to:

- Work with the DBA Cockpit

Unit 8
Lesson 4

Managing SAP Landscapes

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Manage Complex SAP Landscapes
- Manage SAP Landscapes powered by SAP HANA

SAP Landscape Management

SAP Landscape Management (LaMa) automates and simplifies SAP system and landscape management. SAP Landscape Management allows central management and provisioning of an SAP landscape including SAP systems powered by SAP HANA running in physical, virtual, and cloud infrastructures. It is possible to streamline and automate critical business processes and maximize landscape visibility and control.

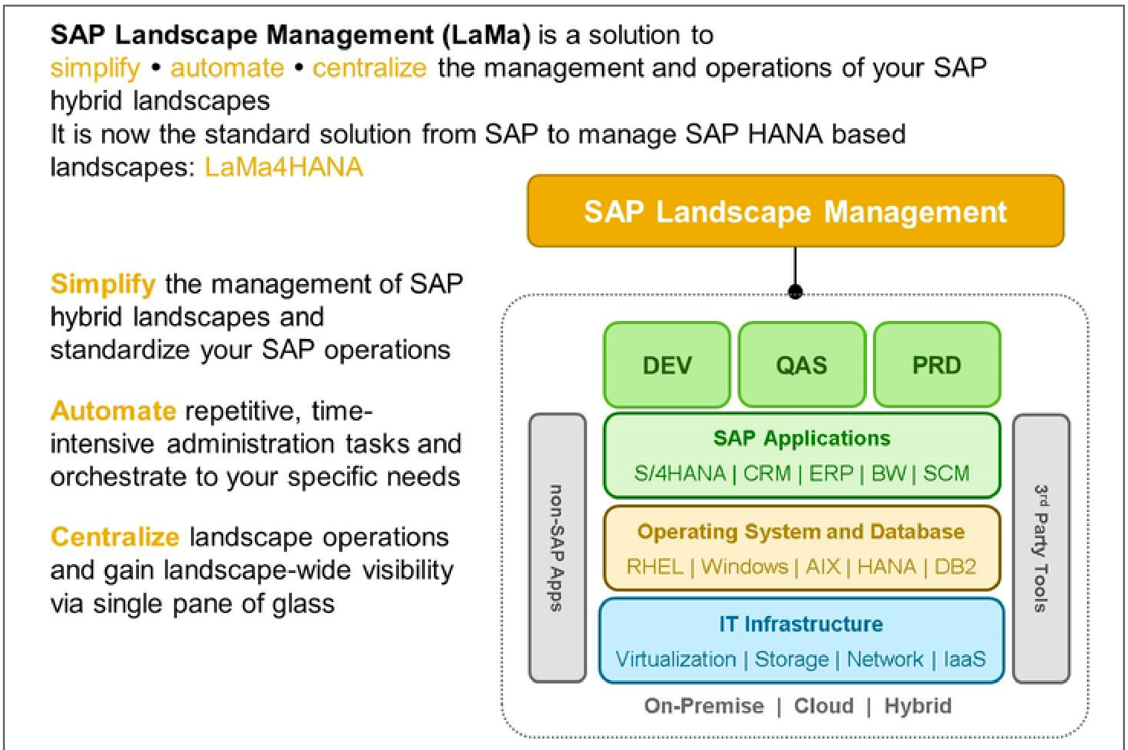

Figure 169: Overview SAP Landscape Management

SAP Landscape Management is an add-on to SAP NetWeaver 7.5, installed as an application with the SAP NetWeaver Application Server for Java (SAP NetWeaver AS for Java). It uses the underlying SAP NetWeaver AS for Java technology. Therefore, to monitor and manage the SAP Landscape Management system, you use the monitoring and managing mechanisms of SAP NetWeaver AS for Java.

- SAP Landscape Management software is a licensed product and the product is available in two editions: Standard edition and enterprise edition. Standard edition is available to all SAP customers with at least one active SAP application license and has no license cost associated with it. Enterprise edition is available to all SAP customers as a separately licensed product. Please contact your SAP Account Executive or Sales Representative for further details.

The following figure gives an overview about the operations that can be performed with SAP LaMa .

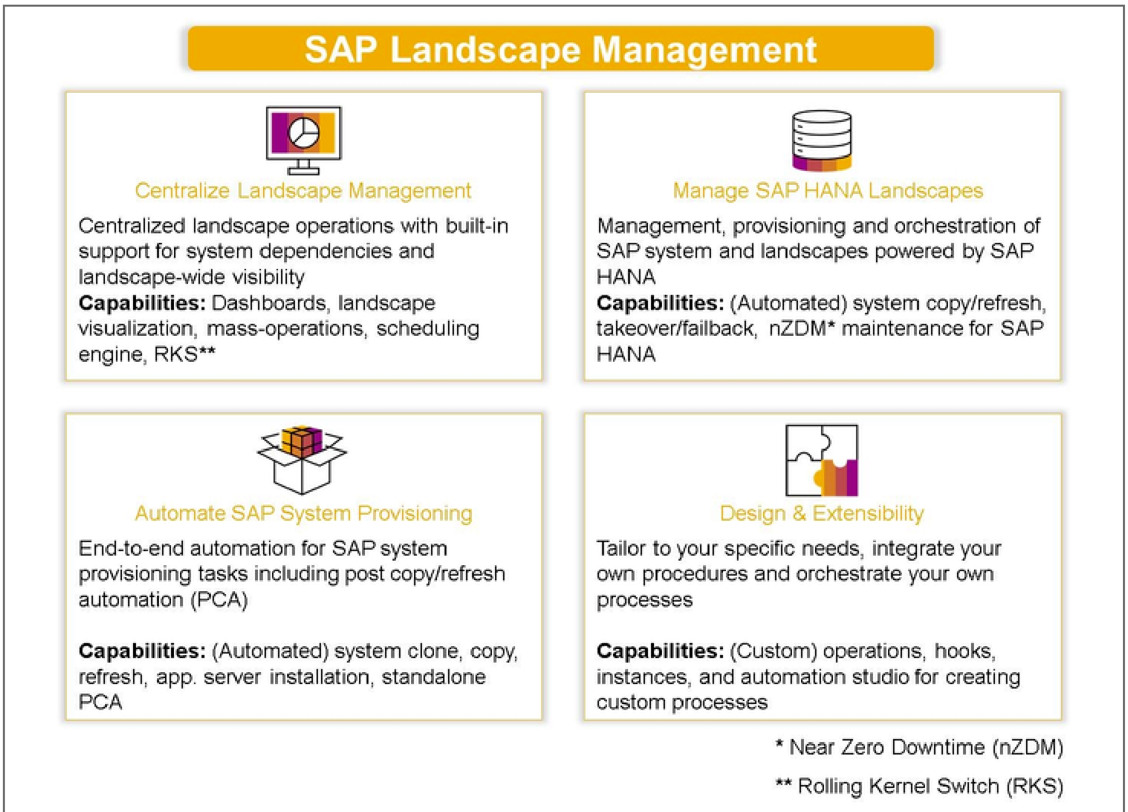

Figure 170: Managing Complex SAP Landscapes

Managing SAP Landscapes powerded by SAP HANA

SAP Landscape Management is the standard solution from SAP to manage SAP HANA based landscapes: SAP LaMa4HANA . It performs the following operations:

- Simplifies the management of SAP hybrid landscapes and standardize your SAP operations
- Automates repetitive, time-intensive administration tasks and orchestrate to your specific needs
- Centralizes landscape operations and gain landscape-wide visibility via single pane of glass

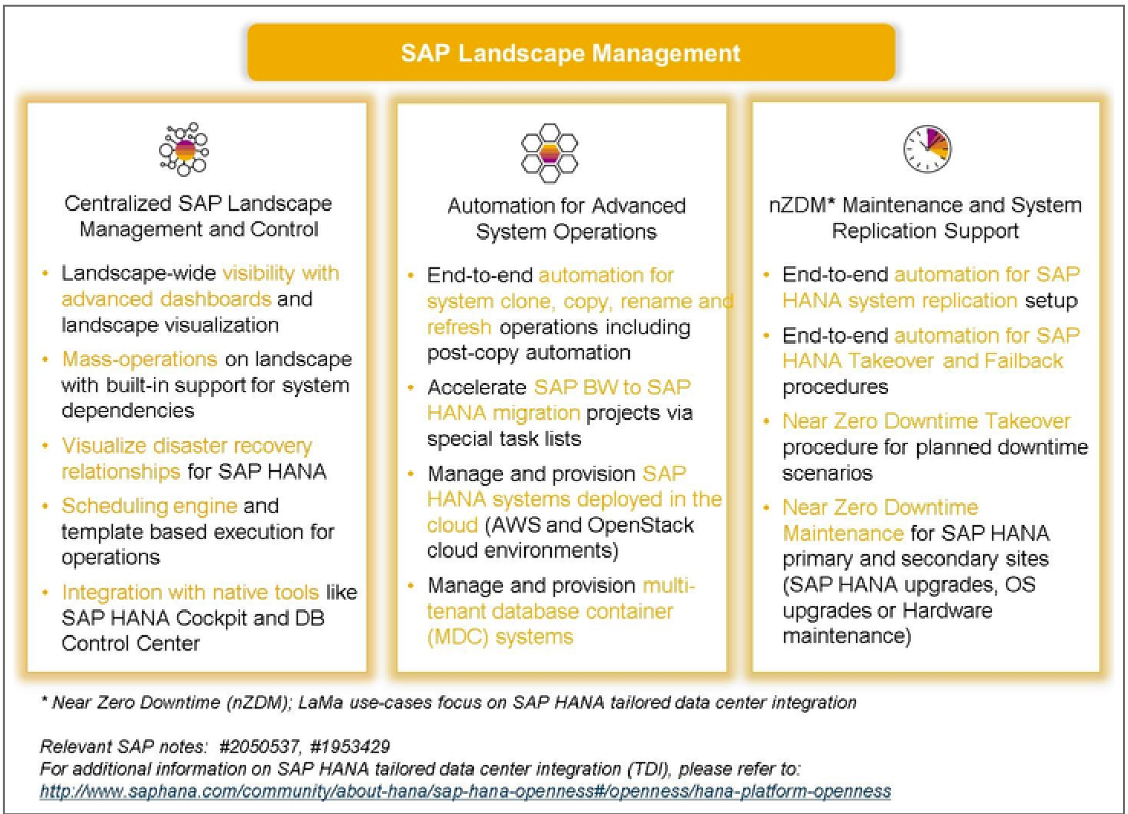

Figure 171: Managing SAP Landscapes powered by SAP HANA

In the following figures two examples are shown relating to how SAP HANA system administration can benefit from automation.

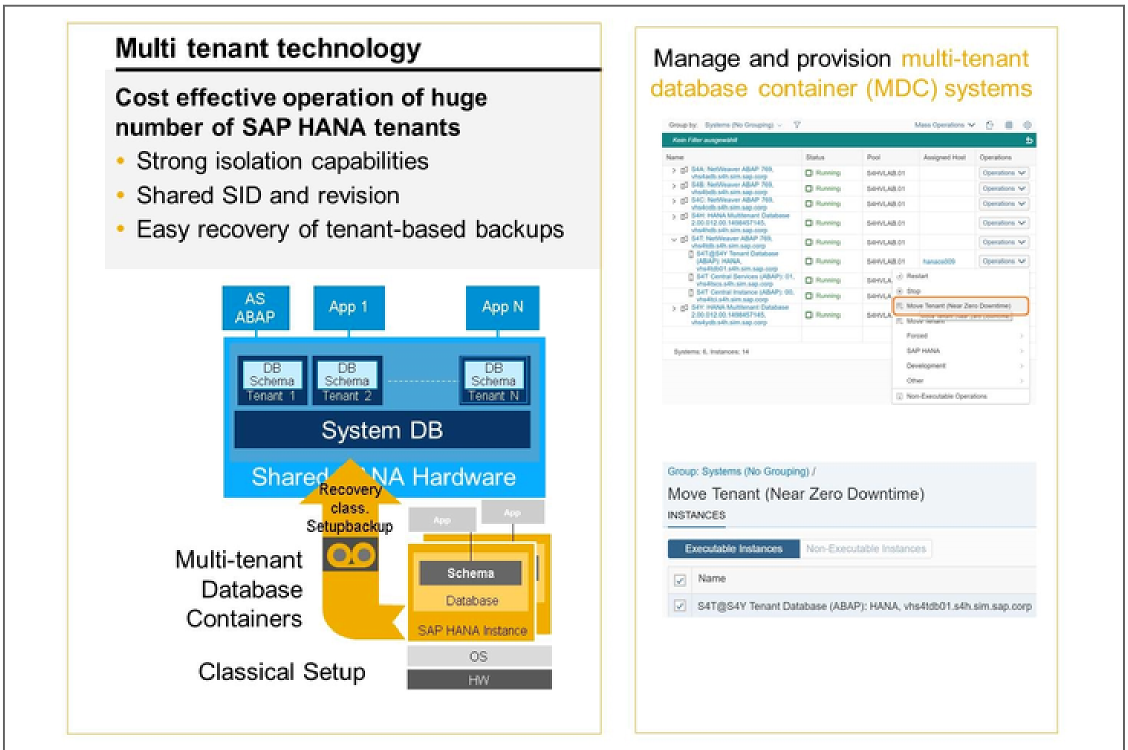

Figure 172: Management and Provisioning for MDC Systems

Features to benefit from Automation

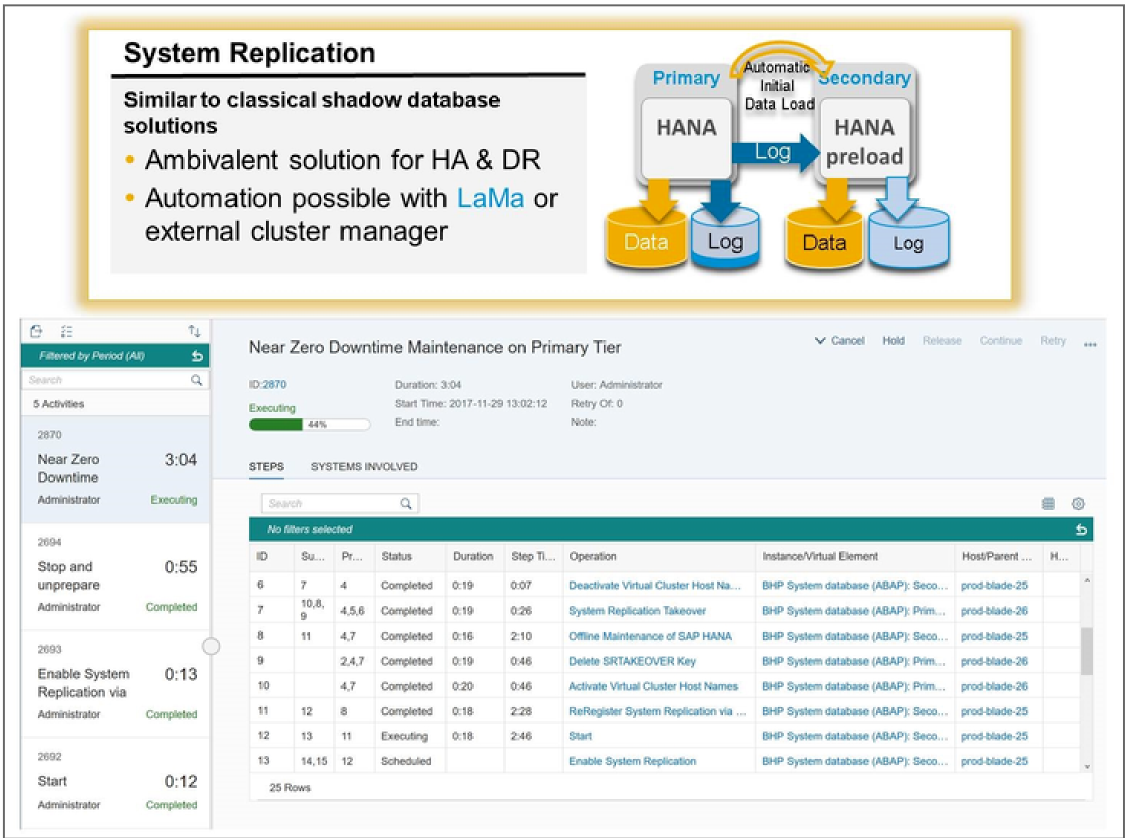

Figure 173: Features to benefit from Automation for System Replication

LESSON SUMMARY

You should now be able to:

- Manage Complex SAP Landscapes
- Manage SAP Landscapes powered by SAP HANA

Unit 8
Lesson 5

Using SAP HANA Studio

LESSON OVERVIEW

This lesson gives an introduction to the SAP HANA Studio and SAP HANA Cockpit and explains some basic features of both tools. Further details are discussed in the respective topics in other lessons of this course.

Business Example

With SPS09 a new Administration tool is available, the SAP HANA cockpit.

The SAP HANA cockpit is a SAP Fiori launchpad based tool. It is currently not a replacement of SAP HANA Studio, but it is a long-term goal to have only one administration tool. The current version is not yet complete, but more features will be included in the next version.

The well-known SAP HANA studio runs on the Eclipse platform, and is both the central development environment and the main administration tool for SAP HANA.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Use SAP HANA Studio

SAP HANA Studio

Consists of several applications / perspectives

- Administration console
- Information modeler
- Lifecycle management

Perspectives can be switched easily

- Quick link buttons
- Welcome screen

Several SAP HANA databases can be integrated

- Navigator tree
- Folder structure

Context menu provides easy access to all functions

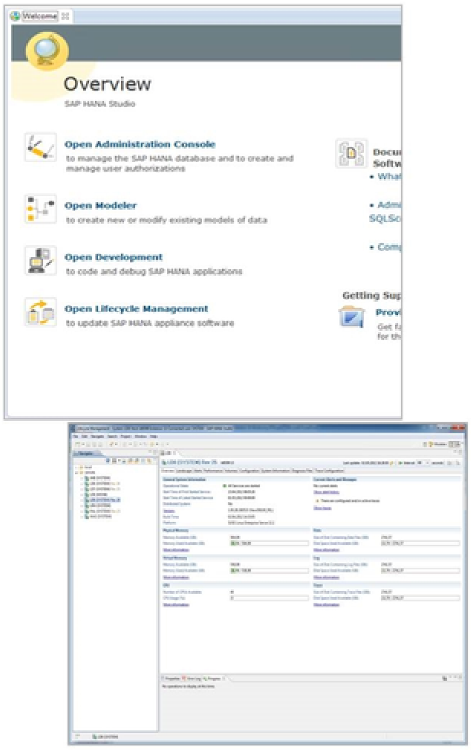

Figure 174: SAP HANA Studio: Overview

Administrators use the SAP HANA studio, for example, to start and stop services, to monitor the system, to configure system settings, and to manage users and authorizations. The SAP HANA studio accesses the servers of the SAP HANA database by SQL.

Developers can use the SAP HANA studio to create content such as modeled views and stored procedures. These development artifacts are stored in the repository, which is part of the SAP HANA database.

SAP HANA Studio: Available Perspectives

The SAP HANA studio is developed in Java and is based on the Eclipse platform. The SAP HANA studio presents its various tools in the form of perspectives. Database administration and monitoring features are contained within the SAP HANA Administration Console perspective. Additional perspectives include the SAP HANA Modeler perspective and the SAP HANA Development perspective.

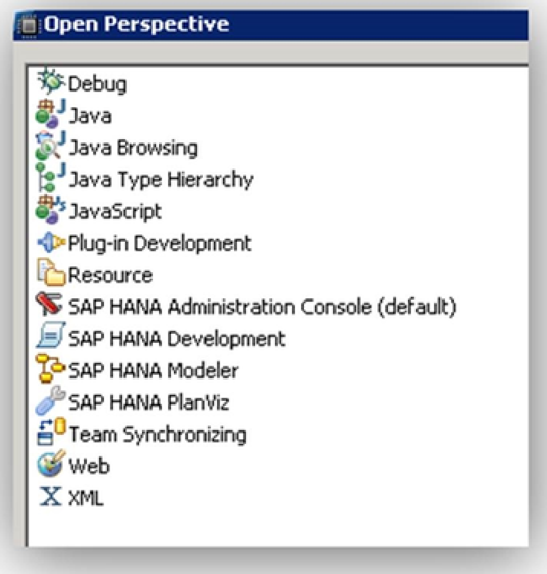

Figure 175: SAP HANA Studio: Available Perspectives

SAP HANA Studio: Adding a System

After the installation, SAP HANA Studio does not contain any system. To add SAP HANA systems, right click the system window. The following two options are available:

- Add System...
- Add System Archive Link...

Figure 176: SAP HANA Studio: Adding a System

Add System

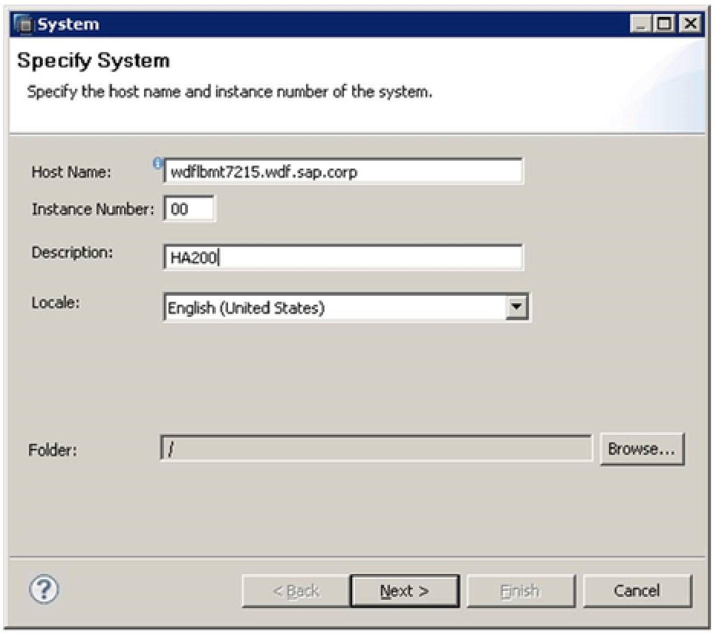

Figure 177: Add System

The first option adds an individual SAP HANA system following a guided procedure, in which you can provide the following information:

- Host name
- Instance number
- Description (for identification in SAP HANA Studio only)
- Locale
- Database User
- Database Password

Add System Archive Link

The Add System Archive Link is a new feature that allows users to connect to multiple SAP HANA systems.

A centrally-stored archive of SAP HANA systems is an efficient way of deploying system information to all users of the SAP HANA studio. Users no longer have to obtain the connection details of all systems individually or add them individually.

One user can manage the list of all systems in a centrally-accessible archive (through File → Export → SAP HANA → Landscape. Other users can then link to this archive in read-only mode.

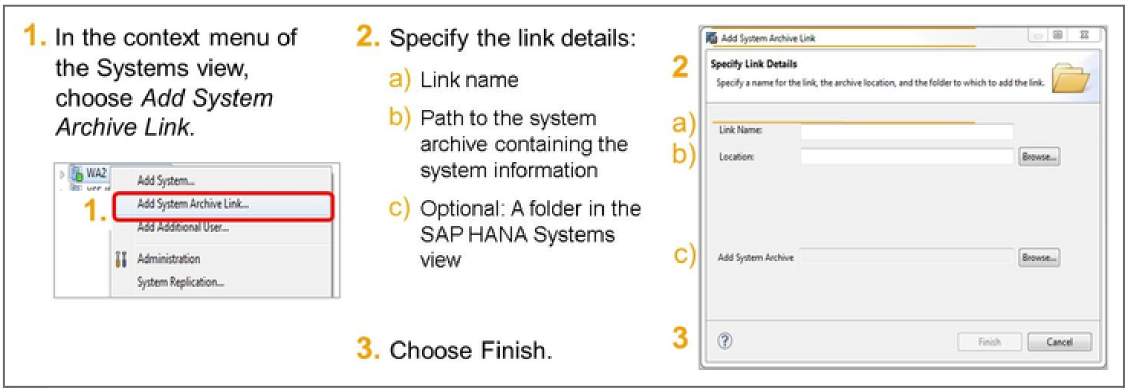

Figure 178: Add System Archive Link

The second option, Add System Archive Link, allows you to insert a link to a centrally-stored archive of SAP HANA systems.

To allow users who work in the SAP HANA studio to connect efficiently to multiple SAP HANA systems, you can manage a list of all systems in a centrally-accessible archive. Users can then link to this archive.

A centrally-stored archive of SAP HANA systems is an efficient way to deploy system information to all users of the SAP HANA studio, for example, developers, content modelers, and other administrators. Users no longer have to obtain the connection details of all systems individually or add them individually. In addition, if you change the central file, for example to add new systems or change the host of an existing system, you can ensure that users always have up-to-date system access.

SAP HANA Studio: System Navigator

Backup	Backup configuration, catalog, snapshot
Catalog	Schemas with tables, functions, procedures, and so on
Content	With packages and views
Provisioning	Smart data access, remote data sources, and so on
Security	Users and roles

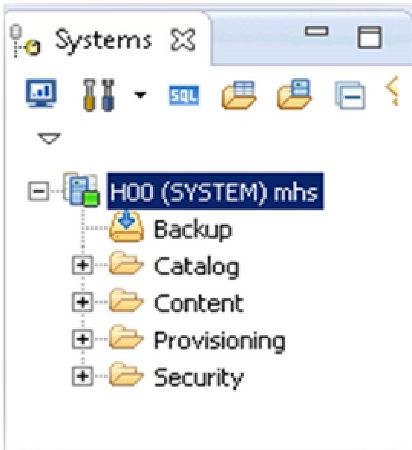

Figure 179: SAP HANA Studio: System Navigator

Once the system is added, it appears in the system navigator screen on the left of the SAP HANA Studio window. It contains the following elements:

- Backup

  You can configure the backup (Destination, File size), and view the backup catalog.

- Catalog

The catalog contains all schemas that have the respective column and row tables. While some schemas exist by default for internal SAP HANA use, others can be created by users, respectively administrators.

- Content

  The content folder holds packages that store development and modeling artifacts.

- Provisioning

  Provisioning relates to the functionality of "Smart Data Access". It contains remote sources and proxy tables.

- Security

  In the security folder, you can maintain users, roles, and other security settings.

Administration Console Perspective: Screen Areas

The figure, Administration Console Perspective: Screen Areas, provides an overview of the administration and monitoring activities of SAP HANA using the administration console of the SAP HANA studio (the studio). The administration console of the studio allows system administrators to manage the database, including creating and managing user authorizations. The studio also contains perspectives for other tasks, such as the information modeler that allows modeling users to create new or modify existing models of data, and the lifecycle management that allows you to update the SAP HANA system.

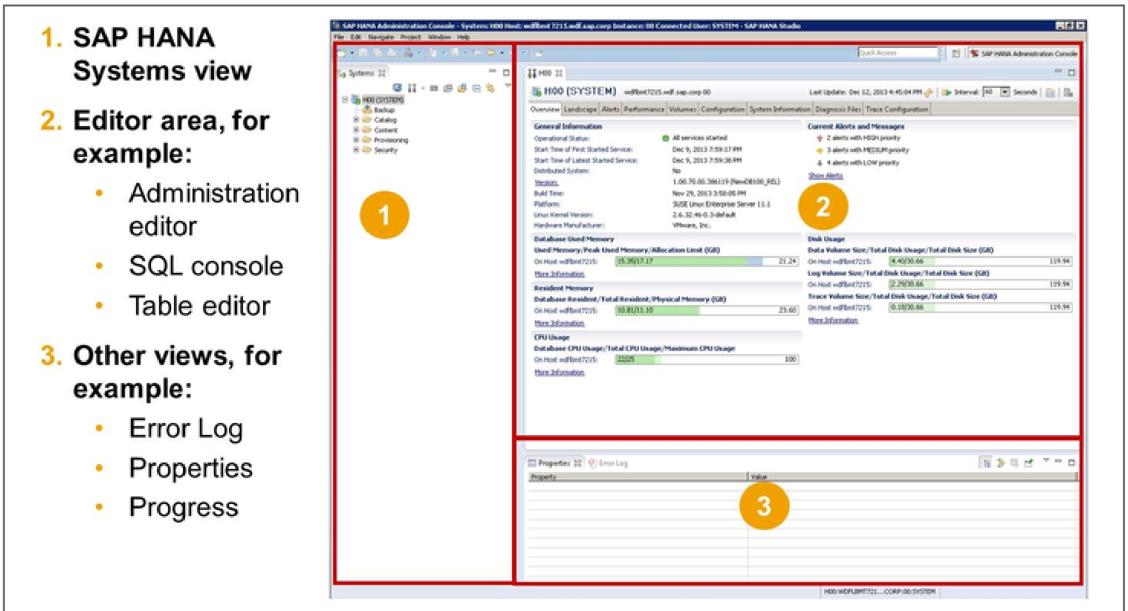

Figure 180: Administration Console Perspective: Screen Areas

The administration console is predelivered by SAP.

You can access the administration console in one of the following ways:

- Choose the Administration button in the upper right corner.
- Double-click the system in the system monitor.
- Double-click the system in the Navigator view.

View Context Menu of SAP HANA Studio

The SAP HANA Systems view provides a hierarchical view of all the SAP HANA systems managed in the SAP HANA studio and their contents (database catalog, users, roles). This view allows you to see the status of your systems at a glance. It is also the central access point for performing system-specific administration and monitoring activities.

You can access a range of administrative functions from the context menu of the systems view.

Administration Console: Overview Tab

The administration console contains the following tabs:

- Overview
- Landscape
- Alerts
- Performance
- Volumes
- Configuration
- System Information
- Diagnosis Files
- Trace Configuration

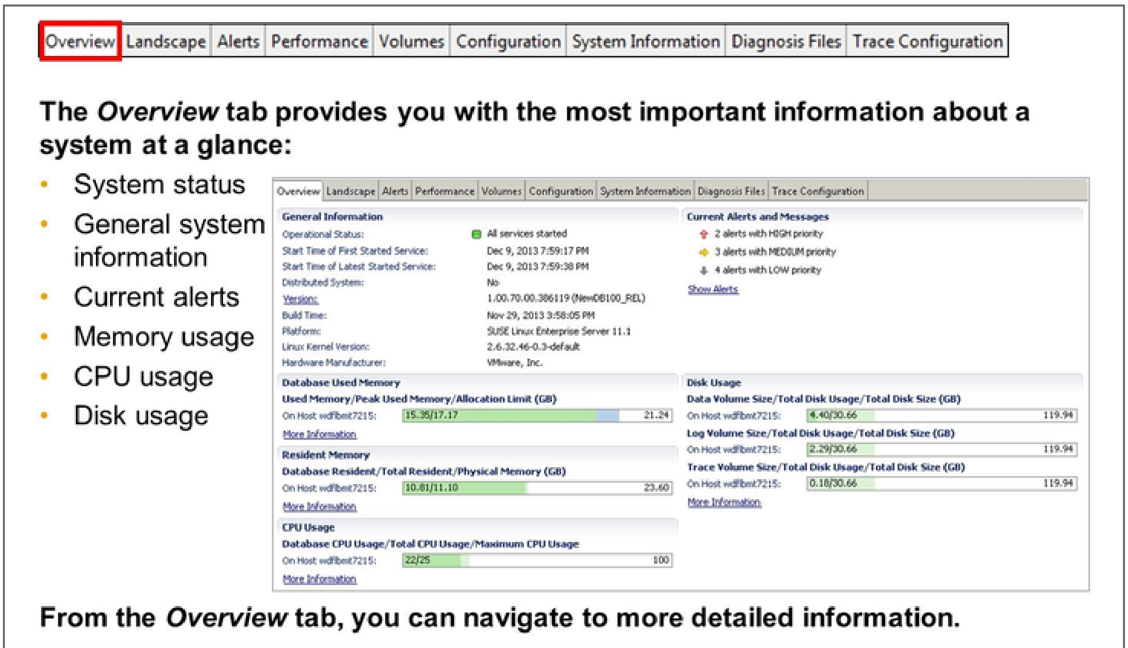

Figure 181: Administration Console: Overview Tab

Check the database status on the Overview tab page of the Administration editor regularly. To open the Administration editor, choose Administration in the context menu, or double-click the database entry. The most important database information is displayed here.

In the upper part of the screen, the overall database state and general database information (software versions, and so on) are displayed. The warning section shows the latest warnings generated by the statistics server. The bar views provide an overview of important database resources: the amount of memory, CPUs, and storage space available on the server, as well as the amount used by these resources (used by all processes, not only by the SAP HANA database).

In a distributed landscape, the amount of available resources is aggregated over all servers. In addition, the resource information of the server with the highest resource consumption is displayed. Links in each section guide you to more detailed information about the specific topic. For example, it shows a database version history, a detailed alert list, or detailed storage information.

The Administration Editor: Diagnosis Mode

Administration editor in diagnosis mode

- The system has not yet started or is down
- No SQL connection available
- SAP HANA studio collects information using the SAP start service (sapstartsrv)
- Analyze any problems that may occur during startup or while the system is stopped
- You can also access diagnosis files

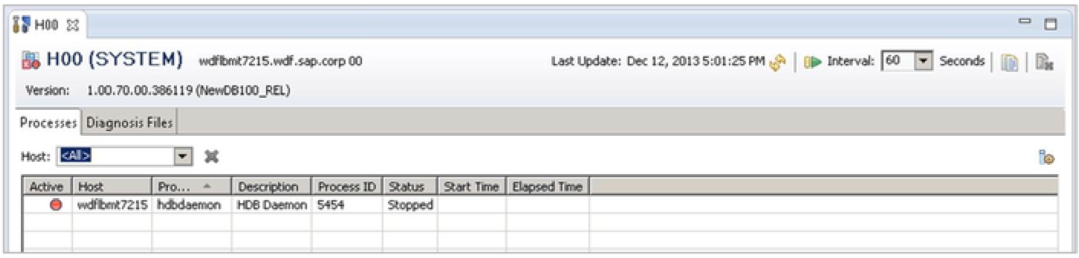

Note: You can open the Administration editor in diagnosis mode only as the operating system user, sidadm.

Figure 182: The Administration Editor: Diagnosis Mode

The SAP HANA studio collects information about the system using SQL. However, when the system has not yet started or is down, no SQL connection is available. In this situation, the SAP HANA studio collects information about the database using the connection of the SAP start service (sapstartsrv).

You can view this information in the Administration editor in diagnosis mode. In this way, you can analyze any problems that occur during startup or while the system is stopped. You can also access diagnosis files.

You can only open the Administration editor in diagnosis mode as the operating system user, <sid>adm.

User-Defined SQL Statements for System Monitoring

When you save the Administration editor, all statements, together with the defined folder structure, are saved to a single XML file. They are available on the System Information tab of the Administration editor for all systems registered in the SAP HANA Studio .

LESSON SUMMARY

You should now be able to:

- Use SAP HANA Studio

Unit 8

Learning Assessment

1. When was the first release of SAP HANA Cockpit?

   Choose the correct answer.

   - [ ] **A** SAP HANA 1.0 SPS09
   - [ ] **B** SAP HANA 1.0 SPS11
   - [ ] **C** SAP HANA 1.0 SPS12
   - [ ] **D** SAP HANA 2.0

2. When connected to the SAP HANA database, which hdbsql command displays general information about the database?

   Choose the correct answer.

   - [ ] **A** \di
   - [ ] **B** \s
   - [ ] **C** \c
   - [ ] **D** \a

3. Which of the following database resources can be monitored at the DBA Cockpit Current Status Overview monitor?

   Choose the correct answers.

   - [ ] **A** The time at which the database was started.
   - [ ] **B** Disk consumption from the SAP HANA database.
   - [ ] **C** Connection statistics about open connections.
   - [ ] **D** Memory and CPU consumption from the SAP HANA database.

4. SAP Landscape Management is a powerful automated solution de-signed to simplify, centralize, and orchestrate the management and operations of the entire application landscape powered by SAP HANA. With the enterprise edition of SAP Landscape Management, you can perform various operations on SAP HANA instances, such as start, stop, monitor, or system replication operations.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

5. Which of the following elements are contained by the SAP HANA Studio System Navigator?

   Choose the correct answers.

   - [ ] **A** Backup
   - [ ] **B** Content
   - [ ] **C** Configuration
   - [ ] **D** Security

Unit 8

Learning Assessment - Answers

1. When was the first release of SAP HANA Cockpit?

   Choose the correct answer.

   [X] **A** SAP HANA 1.0 SPS09

   [ ] **B** SAP HANA 1.0 SPS11

   [ ] **C** SAP HANA 1.0 SPS12

   [ ] **D** SAP HANA 2.0

   Correct! With SAP HANA 1.0 SPS09 a new Administration tool is available, the SAP HANA cockpit. With SAP HANA 1.0 SPS09, a first version of the web-based SAP HANA cockpit was introduced for monitoring SAP HANA. Read more on this in the lesson Explaining the Administration Tools (Unit 8, Lesson 1) of the course HA200_14.

2. When connected to the SAP HANA database, which hdbsql command displays general information about the database?

   Choose the correct answer.

   [ ] **A** \di

   [X] **B** \s

   [ ] **C** \c

   [ ] **D** \a

   Correct! \s displays general information about the database. \di logs the user off from the database. \c logs a user onto the database. \a switches AUTOCOMMIT mode on or off. Read more on this in the lesson Using the HDBSQL Command Line Tool (Unit 8, Lesson 2) of the course HA200_14.

3. Which of the following database resources can be monitored at the DBA Cockpit Current Status Overview monitor?

   Choose the correct answers.

   [X] **A** The time at which the database was started.

   [X] **B** Disk consumption from the SAP HANA database.

   [ ] **C** Connection statistics about open connections.

   [X] **D** Memory and CPU consumption from the SAP HANA database.

   Correct! The database start time can be found in the General System Information section. The Database Disk Usage section provides disk consumption details. Detailed information about open connections, such as network input/output statistics is provided its own monitoring area: System Information — Connections and Connection Statistics. Memory and CPU consumption from the database can be seen in the Memory and CPU area. Read more on this in the lesson Working with the DBA Cockpit (Unit 8, Lesson 3) of the course HA200_14.

4. SAP Landscape Management is a powerful automated solution de-signed to simplify, centralize, and orchestrate the management and operations of the entire application landscape powered by SAP HANA. With the enterprise edition of SAP Landscape Management, you can perform various operations on SAP HANA instances, such as start, stop, monitor, or system replication operations.

   Determine whether this statement is true or false.

   [X] True

   [ ] False

   Correct! SAP Landscape Management Enterprise Edition software helps users reduce the total cost of ownership (TCO) of their SAP systems and improve their business agility by simplifying and automating the efforts required to configure, provision, deploy, monitor, and manage their systems in both physical and virtualized infrastructures. Read more on this in the lesson Managing SAP Landscapes (unit 8, Lesson 4) of the course HA200_14.

5. Which of the following elements are contained by the SAP HANA Studio System Navigator?

   Choose the correct answers.

   [X] **A** Backup

   [X] **B** Content

   [ ] **C** Configuration

   [X] **D** Security

   Correct! You can configure the backup (Destination, File size), and view the backup catalog. The content folder holds packages that store development and modeling artifacts. In the security folder, you can maintain users, roles, and other security settings. The Administration Console contains the tab: Configuration. Not the System Navigator pane. It provides Configuration Information where you can check and change system configuration parameters. Read more on this in the lesson Using SAP HANA Studio (Unit 8, Lesson 5) of the course HA200_14.

UNIT 9 Database Administration Tasks

UNIT OBJECTIVES

- Start and stop SAP HANA
- Configure the SAP HANA database

- Perform regular database administration tasks
- Configure SAP HANA traces
- Work with diagnosis information
- Use the SQL console
- Perform SAP HANA table administration
- Transport changes
- Understand the administration tasks that are still in SAP HANA studio

Unit 9
Lesson 1

Starting and Stopping SAP HANA

LESSON OVERVIEW

The goal of this lesson is to learn about the different ways to start and stop SAP HANA. The following tools are outlined: SAP HANA cockpit, HDB, and sapcontrol.

Business Example

As the administrator of the SAP HANA database, you need to be able to start, stop, and restart the SAP HANA database system for maintenance purposes. To determine the most suitable tools, you investigate the following tools: SAP HANA cockpit, HDB, and sapcontrol.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Start and stop SAP HANA

Tools for Starting and Stopping the SAP HANA Database

To start and stop a SAP HANA database, you require the credentials of the <sid>adm operating system user. This operating system user is created during the SAP HANA database installation. Alternatively, you can also use the root user to start and stop SAP HANA database using the sapcontrol operating system command.

The SAP start service (sapstartsrv) is the standard SAP mechanism for starting and stopping systems. It starts all necessary database services, such as the name server, index server, and for example the preprocessor services.

You can start or stop the SAP HANA database using the SAP HANA cockpit or by using the operating system commands HDB or sapcontrol.

> Note:
> SAP HANA studio can also be used to start and stop the SAP HANA database system. However, it is not shown in this lesson because SAP HANA studio is depreciated.

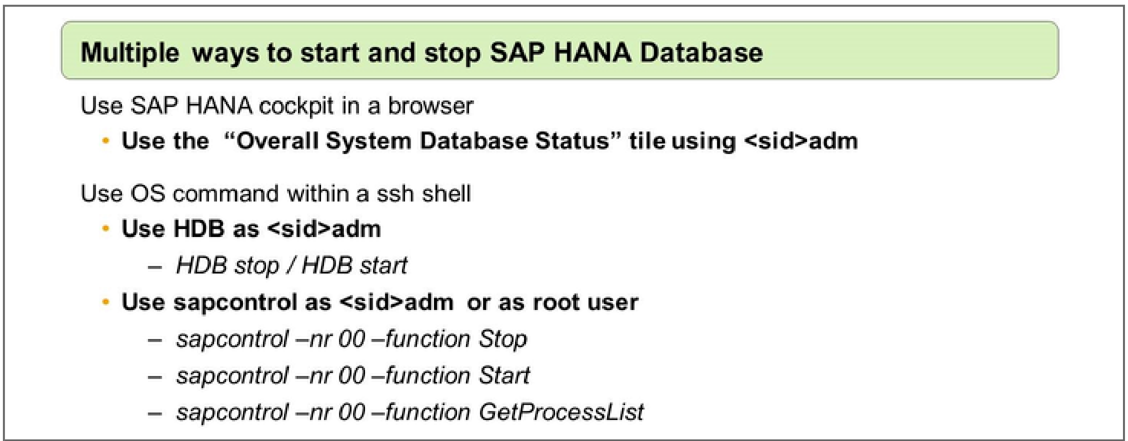

Figure 183: Tools for Starting and Stopping the SAP HANA Database

Starting and Stopping the SAP HANA Database with SAP HANA Cockpit

SAP HANA is designed to deliver continuous system availability, but there are some reasons when the SAP HANA database system needs to be restarted. In the following situations a restart might be necessary:

- During planned maintenance, when the SAP HANA database is updated, or to activate configuration parameter changes that can only be activated by a restart of the system.
- During unplanned maintenance, due to a hardware or software failure. The SAP HANA database is unresponsive or down.

Planned Maintenance: Start or Stop SAP HANA

During planned maintenance, you can start and stop the SAP HANA database system using the Overall System Database Status tile in the SYSTEMDB overview screen of the SAP HANA cockpit. The Overall System Database Status tile shows the status of the services used by the SAP HANA SYSTEMDB. Depending on the current SAP HANA database system state choose the Start System or the Stop System button to start or stop the database.

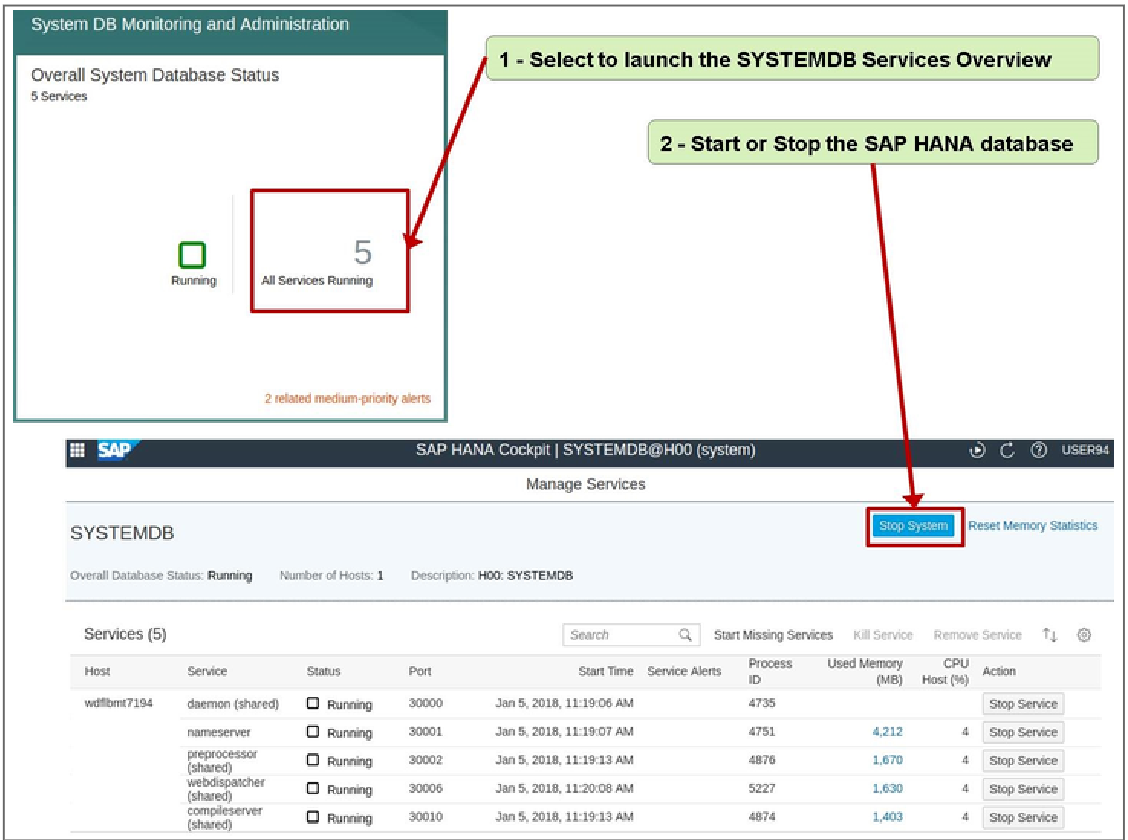

Figure 184: Planned Maintenance: Start or Stop SAP HANA

To start and stop the SAP HANA database using the SAP HANA cockpit in a browser, make sure that you provided the database administrator credentials (personal database administration user or the SYSTEM user) and the SAP Control user credentials (a personal operating system user or the <sid>adm) in the SAP HANA cockpit — Resource Directory.

SAP recommends that you do not use the SYSTEM user for day-to-day activities in production environments. Instead, use the SYSTEM user to create personal database users with the minimum privilege set required for their daily duties.

> Hint:
> If your corporate security policies recommend the use of personal operating system users, instead of the <sid>adm user, to perform database administration tasks, read SAP Note 1881622 - “SAP HANA DB: Access to trace files via sapstartsrv”. This note explains how to setup personal operating system users for sapstartsrv.

Unplanned Maintenance

During the unplanned maintenance, the SAP HANA database system is unresponsive or down. The root cause of this problem might be software or hardware related, and needs further investigation. In the SAP HANA cockpit for the SYSTEMDB in the Alerting & Diagnostics area, you can troubleshoot an unresponsive system and view the trace and diagnosis files.

Before you restart the SAP HANA database, use the Manage full system information dumps application to save all the important log and trace files by generating a full system dump. After this investigation, you can use the Overall System Database Status tile to start the SAP HANA database system.

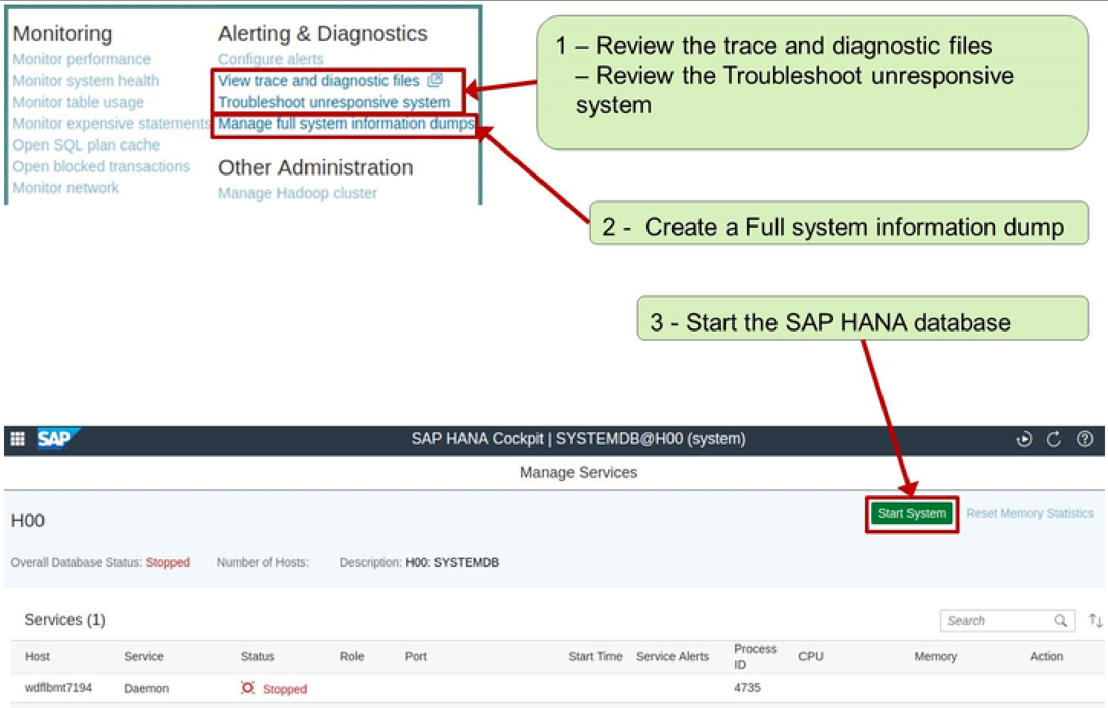

Figure 185: Unplanned Maintenance: Start or Stop SAP HANA

> Note:
> To start and stop SAP HANA database using the Overall System Database Status application, you need to provide the <sid>adm credentials. These credentials are needed because the starting and stopping is performed by the sapstartsrv.

The Troubleshoot unresponsive system application runs via the SAP Start service (sapstartsrv). Therefore it can investigate a unresponsive SAP HANA database system from the operating system level.

Stopping the SAP HANA Database Using SAP HANA Cockpit

When stopping the SAP HANA database system using the Overall System Database Status application, you can specify an immediate shutdown or a soft shutdown with a timeout. The soft shutdown waits for all the running statements to finish. If the running statements aren't finished within the specified timeout, an immediate shutdown is performed.

When stopping the SAP HANA database, you can define how you want to stop the system, as outlined in the following table.

Table 7: Stopping the SAP HANA Database Using SAP HANA Cockpit

Option	Description
Immediately	The system is stopped immediately. Open transactions are aborted and rolled back.

Option	Description
Softly — timeout (sec)	The system is stopped after all running statements have finished. If the system doesn't stop before the specified timeout, it is stopped immediately. The default timeout is 5 minutes.

System Start Activities

When the SAP HANA database system is started, the following activities are executed:

1. The data volume of each service is accessed in order to read and load the restart record.
2. The list of open transactions is read into memory.
3. Row tables are loaded into memory.
4. Open transactions are processed using the redo log as follows:
   - Write transactions that were open when the database was stopped are rolled back.
   - Changes of committed transactions that were not written to the data area are rolled forward.
   - The first column tables start being reloaded into memory because they are accessed for roll forward.

> Note:
> After this step, the database is technically available and logon is possible.
>
> Because a regular or soft shutdown writes a savepoint, there are no replay log entries to be processed in this case.

5. Aborted transactions are determined and rolled back.
6. A savepoint is performed with the restored consistent state of the database.
7. Column tables that are marked for preload and their attributes are asynchronously loaded in the background (if they have not already been loaded as part of log replay).
8. Column tables that were loaded before restart and their attributes start reloading asynchronously in the background (if they have not already been loaded as part of log replay or because they are marked for preload).

Reloading column tables as described in steps 7 and 8 restores the database to a fully operational state more quickly. However, it does create performance overhead and might not be necessary in nonproduction systems.

You can deactivate the reload feature in the **`indexserver.ini`** file by setting the **`reload_tables`** parameter in the sql section to false. In addition, you can configure the number of tables whose attributes are loaded in parallel using the **`tables_preloaded_in_parallel`** parameter in the parallel section of **`indexserver.ini`**. This parameter also determines the number of tables that are preloaded in parallel.

Starting and Stopping the SAP HANA Database Using OS Commands

On operating system level, the SAP HANA database can be started or stopped using the commands **`sapcontrol`** or **`HDB`**.

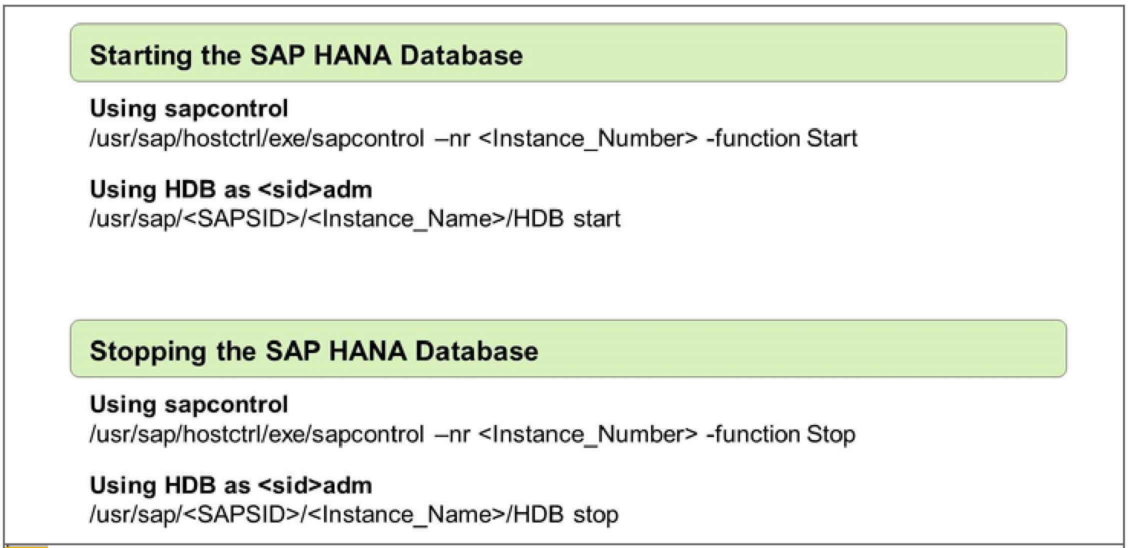

Figure 186: StartStop_OSCommands.ppt

Displaying the Process List

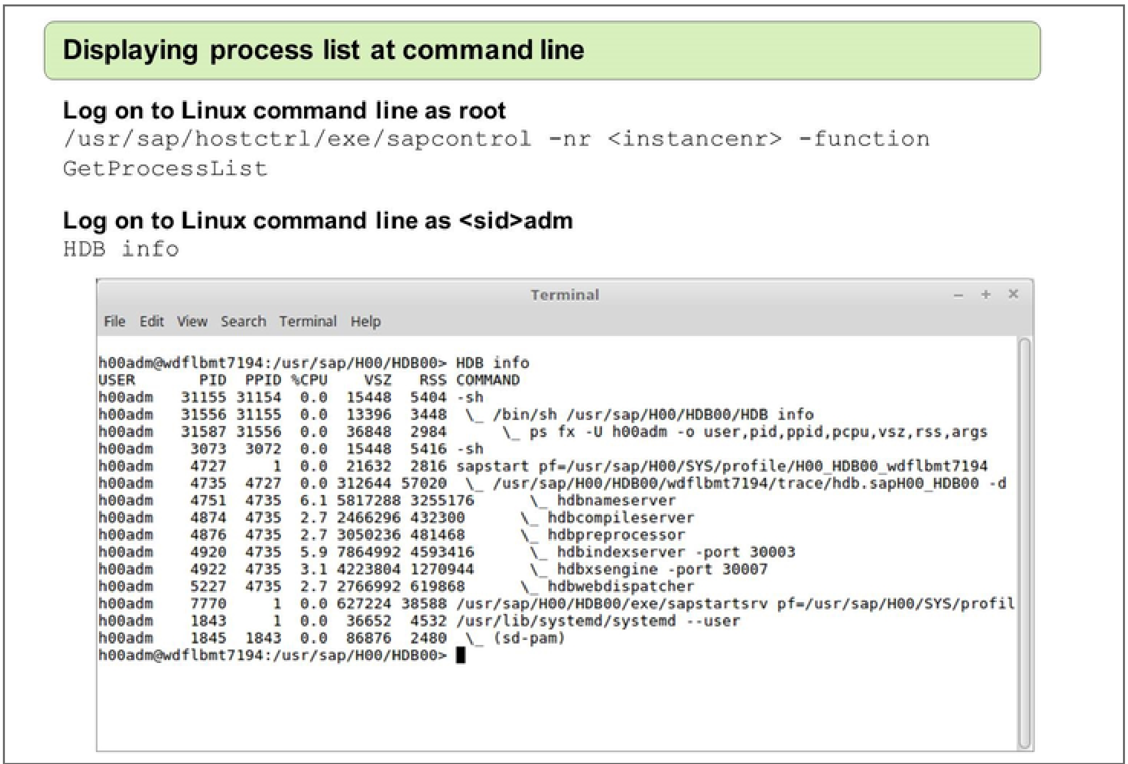

Figure 187: Displaying the Process List

Starting and Stopping Individual Tenants

With the introduction of Multitenant Database Containers (MDC) in the SAP HANA database, the database administrator(s) should not always stop and start the whole SAP HANA database system. Any SAP HANA database systems running version SAP HANA 2.0 SPS 01, or later, is set in multiple-container mode, so it is more likely that only an individual tenant needs to be stopped or started. SAP HANA cockpit supports stopping and starting of tenant databases.

> Note:
> SAP HANA studio doesn’t support starting and stopping of individual tenants.

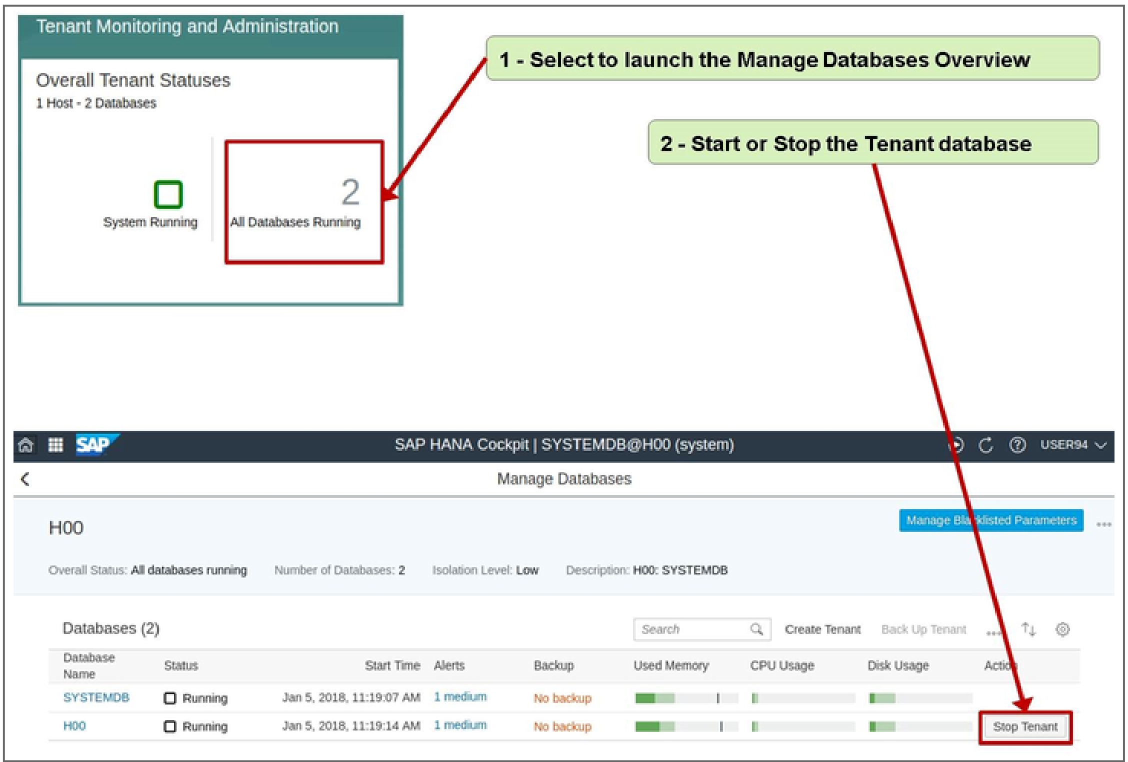

Figure 188: Start and Stop Individual Tenants

As the administrator of a tenant (database) system, you are responsible for creating and configuring new tenant databases, subsequently monitoring the availability and performance of those databases, and performing certain database administration tasks. You can perform many of these basic administration tasks on tenant databases using the SAP HANA cockpit, but for some advanced operations you might require the SQL command-line tools.

As a system administrator, you can start or stop tenant databases individually using the following SQL commands:

- `ALTER SYSTEM START DATABASE <tenant name>`
- `ALTER SYSTEM STOP DATABASE <tenant name>`

If you stop a tenant database individually, you can only start it again individually. You cannot start it with a full system restart.

Starting and stopping an SAP HANA system with multitenant database containers affects the system database and all the tenant databases.

> Note:
> If you stopped the database, it is a hard stop. The database is stopped immediately, even if users are connected. Open transactions are aborted and rolled back; no savepoint operation is forced. You cannot back up a stopped database.

Starting and Stopping of a Distributed SAP HANA Database System

HDB start or HDB stop only starts and stops the local host. It cannot be used to start or stop the complete SAP HANA system.

Use `SAPCONTROL` to start or stop all the hosts in a scaled-out SAP HANA system from the command line.

Action	Command
Start the system	/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StartSystem
Stop the system	/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function StopSystem
Query current host status	/usr/sap/hostctrl/exe/sapcontrol -nr <instance_number> -function GetSystemInstanceList

Figure 189: Starting and Stopping Distributed SAP HANA Systems Using Sapcontrol

> Note:
> You need to be logged on to the SAP system host as user <sid>adm, or as a user with root permissions.

Starting and Stopping of Individual SAP HANA Database Services

You can stop and start the individual database services (nameserver, indexserver, xsengine, and so on) running on hosts.

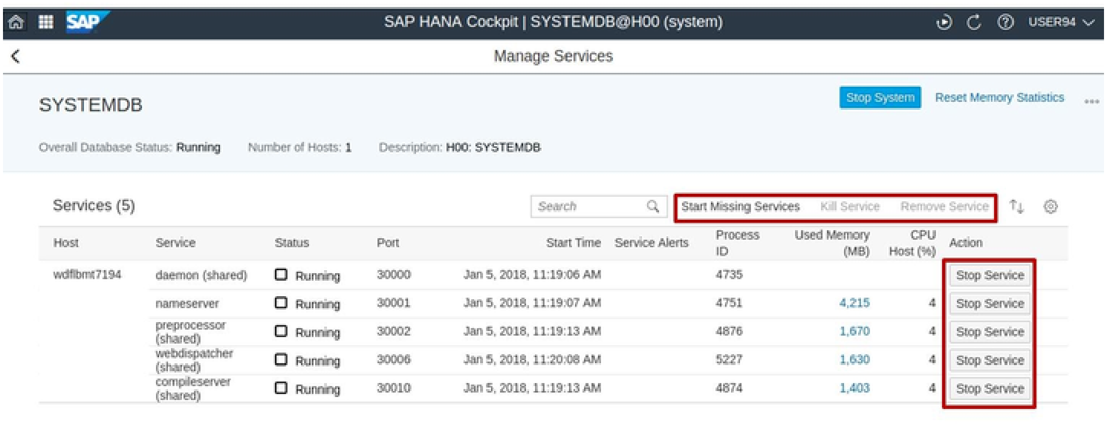

Figure 190: Starting and Stopping Individual Database Services

To stop, start, or restart database services, you must have the system privilege SERVICE ADMIN.

The following are examples of situations where you have to restart an individual database service:

- A host in a distributed system failed and a standby host took over. However, the services of the failed host remain inactive even after the host is available again. In this case, you need to restart the services manually.
- After an update of SAP HANA extended application services, you need to restart the xsengine service.

Options for Stopping and Starting Database Services

Table 8: Options for Stopping and Starting Database Services

Option	Description
Stop Service	The service stops normally and then typically restarts.
Kill Service	The service stops immediately and then typically restarts.
Start Missing Services	Any inactive services start.

Note:
The SAP HANA database provides several features in support of high availability, one of which is service auto-restart. If a failure occurs, or if an intentional intervention by an administrator disables one of the SAP HANA services, the SAP HANA service auto-restart function automatically detects the failure. It then restarts the stopped service process.

LESSON SUMMARY
You should now be able to:

- Start and stop SAP HANA

Unit 9
Lesson 2

Configuring the SAP HANA Database

LESSON OVERVIEW

This lesson explains how to configure the SAP HANA Studio and the SAP HANA Database.

Business Example

You are an administrator and want to adjust the configuration of the SAP HANA Studio, and to change database parameters according to your requirements.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Configure the SAP HANA database

SAP HANA Database Configuration Parameters

The properties of an SAP HANA database system are defined by the parameters in its configuration files. The configuration files are separated into sections; these sections bundle parameters of the same category.

Opening the Configure System Properties Tool

You can display and change parameters using the SAP HANA cockpit 2.0 through the Configure system properties option. To change the parameters of configuration files, you require the system privilege INIFILE ADMIN.

> Note:
> Do not change parameters directly in the configuration files on the operating system level.

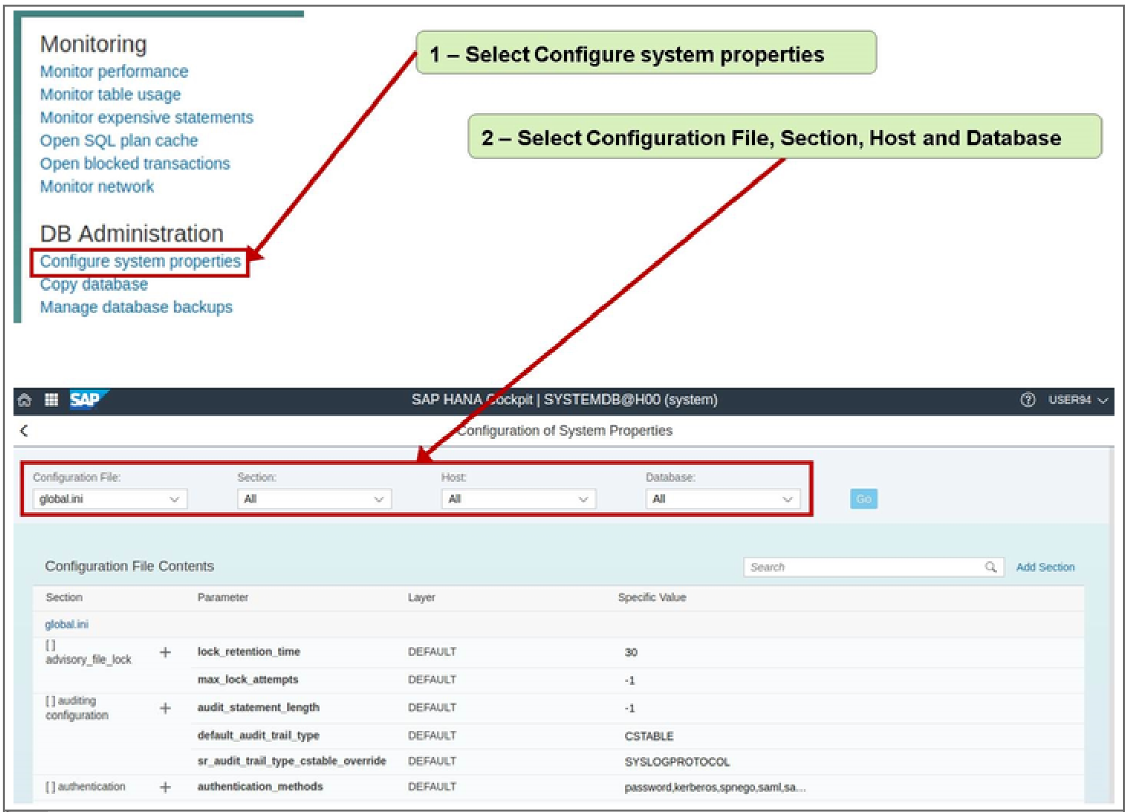

Figure 191: Opening the Configure System Properties Tool

Finding Configuration Parameters Using Search

SAP HANA has many configuration files and parameters. Therefore, to find a parameter easily in the parameter structure, use the Search function.

In the Search field, enter the name of a parameter, or few characters of a parameter, and press Enter on your keyboard.

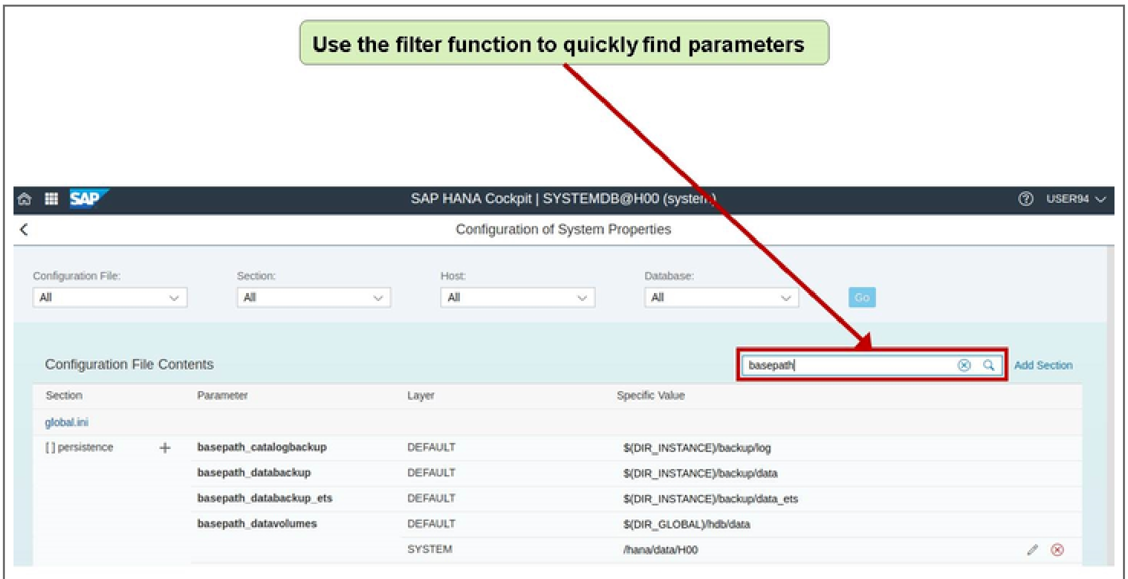

Figure 192: Using the Search Function to Find Configuration Parameters

If you upgrade to a new revision of SAP HANA, the newest parameter settings based on the newest experiences are included automatically, but your own changes remain unchanged from this update.

You can easily distinguish between parameters that still have their default value or are changed. The parameters that are changed have the Pencil icon behind the name.

Changing of Configuration Parameter Values

In the list, you can find the parameters that you searched for. Now you can change the parameter by choosing the plus (+) button in front of the parameter that you want to change.

> Caution:
> Only change the configuration parameters of the SAP HANA database if these changes are recommended in SAP documentation, SAP Notes, or by SAP employees (for example, consulting, development, and support).

To guarantee optimal performance and the highest stability, SAP appliance hardware partners can deliver SAP HANA systems with settings that deviate from the standard. For more information see, SAP Note 1730999 : Configuration changes in SAP HANA appliance.

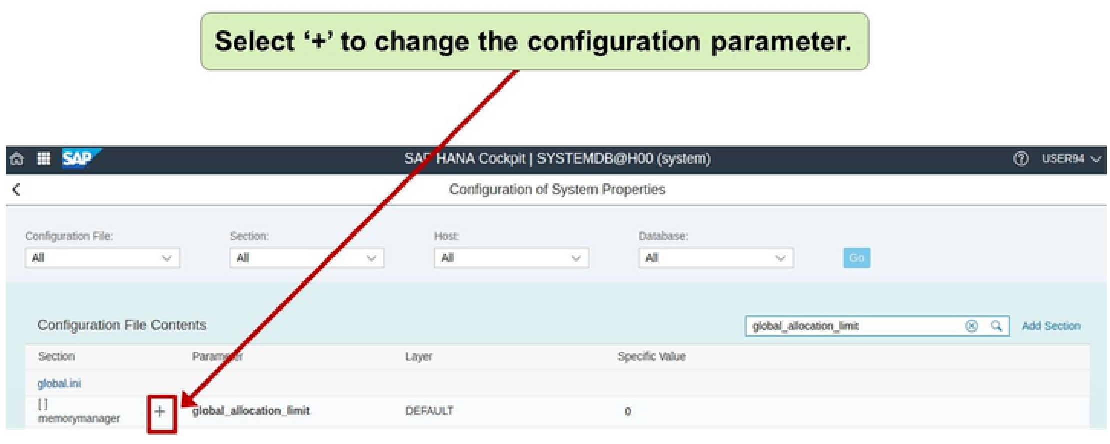

Figure 193: Changing Parameter Values in SAP HANA Cockpit

Changing Parameter Values in SAP HANA Cockpit

After you choose the plus (+) button, a dialog box window appears where you can maintain the parameter values.

> Hint:
> Before choosing the plus (+) button, copy the parameter name to the clipboard. In this way, you can paste the parameter name into the Key: input field.

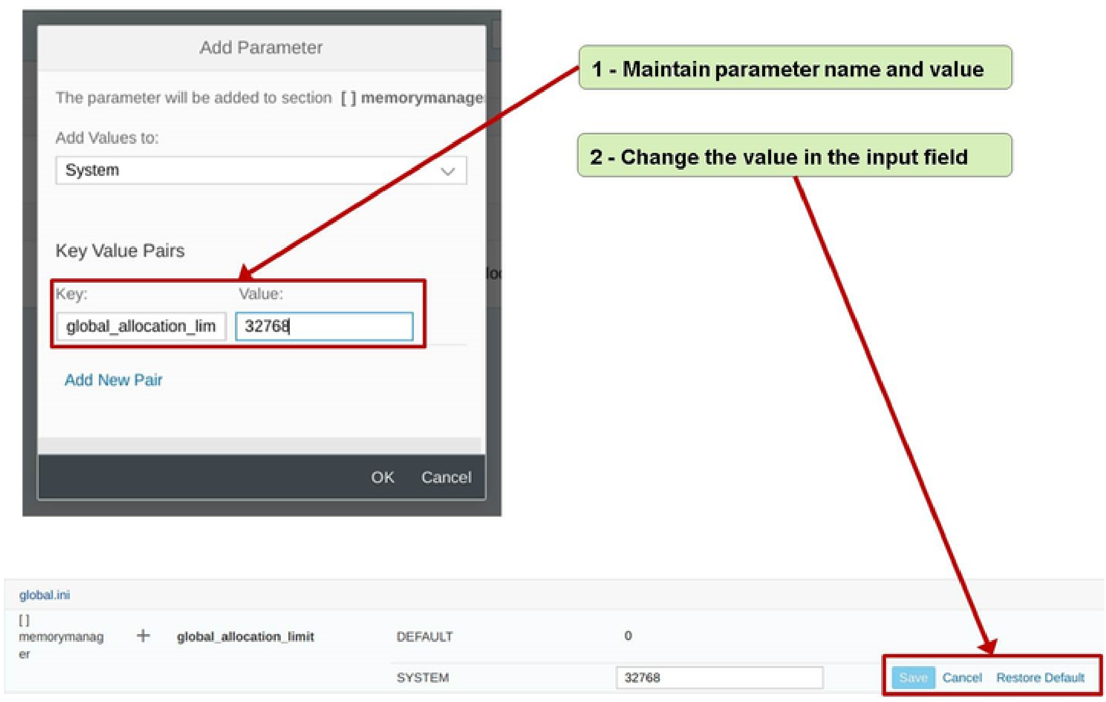

Figure 194: Changing Parameter Values in SAP HANA Cockpit

To change or restore a changed configuration parameter value to its default value, choose the Pencil icon behind the value. If you select the Restore Default option, the change is reverted.

Important Configuration Parameters

The following section introduces several important SAP HANA configuration parameters. Check these parameters and, if needed, change them during the post-installation process.

Parameter: global_allocation_limit

Parameter: global_allocation_limit

Type	Range	Unit	Changeable
integer	0 – max. memory	MB	online

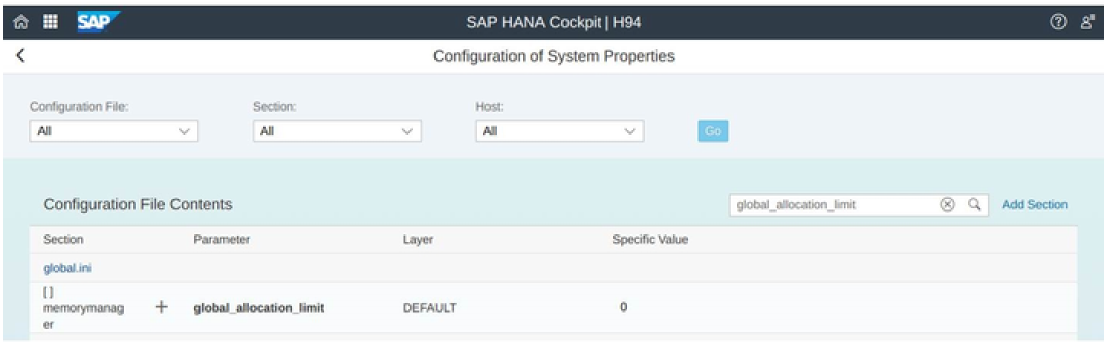

Figure 195: Parameter: global_allocation_limit

The global_allocation_limit parameter limits the amount of memory used by the database. The value is the maximum allocation limit in MB.

> Note:
> If there is a missing entry or a value of 0, the system uses the default settings.

The global allocation limit is calculated as 90% of the first 64 GB of available physical memory on the host, plus 97% of each further GB. Or, for small physical memory, it is calculated as the physical memory minus 1 GB.

If you only enter a value for the system, it is used for all hosts. For example, if you have five hosts, and you set the limit to 5 GB, the database can use up to 5 GB on each host (25 GB in total). If you enter a value for a specific host, then the specific value is used for that host, and the system value is used for all other hosts. This is relevant only for distributed systems.

> Hint:
> For information on the memory allocation of SAP HANA, see the lesson on memory management and data persistence.

Parameter: savepoint_interval_s

The savepoint_interval_s parameter controls how often the internal buffers are flushed to the disk, and when a restart record is written. After a power failure or crash, the log since the last savepoint needs to be replayed. Thus, this parameter indirectly controls the restart time.

If you set the savepoint to a lower value, the startup shortens, but the CPU load increases slightly. If you set the savepoint to a higher value, the startup time increases, but the CPU load decreases a little.

If you set the savepoint to 0, the writing of savepoints is disabled. With no savepoints, you cannot recover the database. Only use this setting for testing purposes, like log I/O performance tests.

> Note:
> Never use savepoint_interval_s = 0 for a productive SAP HANA database system.

Parameter: savepoint_interval_s

Type	Range	Unit	Changeable
integer	0,10 – 7200	second	online

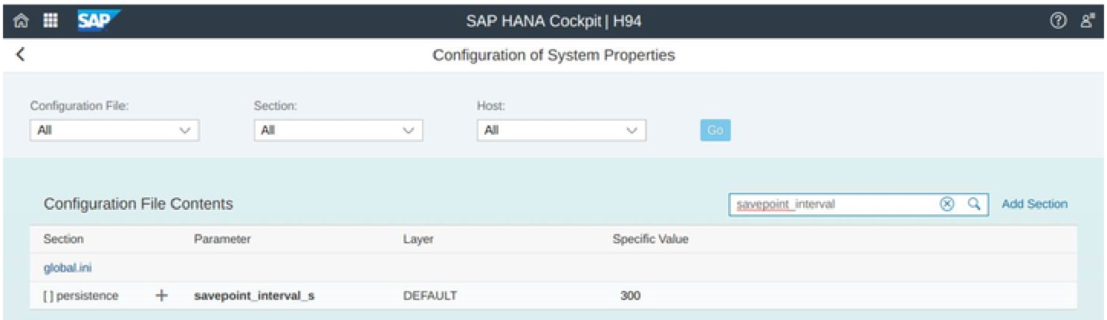

Figure 196: Parameter: savepoint_interval_s

> Note:
> Because changes to data are persisted to the log area synchronously, they are not lost if a power failure or crash occurs.

Parameter: enable_auto_log_backup

Automatic log backup is always backing up closed log segments of the database. Generated backups are stored in the location set by the basepath_logbackup parameter.

You can enable or disable automatic log backup with the enable_auto_log_backup parameter. The default setting is **`enable_auto_log_backup = yes`**.

Parameter: enable_auto_log_backup

Type	Range	Unit	Changeable
Boolean	yes or no		online

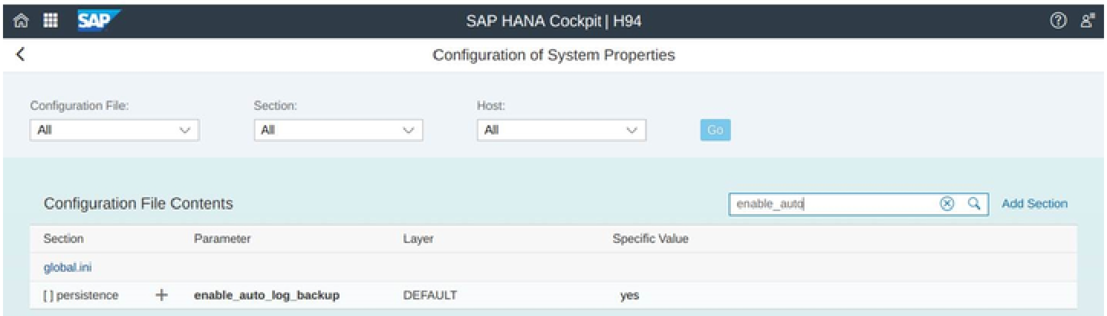

Figure 197: Parameter: enable_auto_log_backup

During normal system operation (log mode normal), keep the automatic log backup activated.

Caution:
If the automatic log backup is disabled and log mode normal is used, the log area increases until the file system is full. If the file system is full, the database freezes.

Parameter: log_mode

When the log_mode is set to **`normal`** and the parameter enable_auto_log_backup is enabled, then the log segments are backed up automatically. The normal setting for log_mode provides support for point-in-time recovery.

After the system backs up the full log segment, it can reuse the space that the full log segment occupied in the log area to overwrite it with new log entries. If the log area becomes full and new log segments cannot be created on disk, a log-full situation arises and the database freezes. When the log area is full, the system is unable to write more log entries until a log backup is completed.

Parameter: log_mode

Type	Range	Unit	Changeable
enum	normal or overwrite		offline

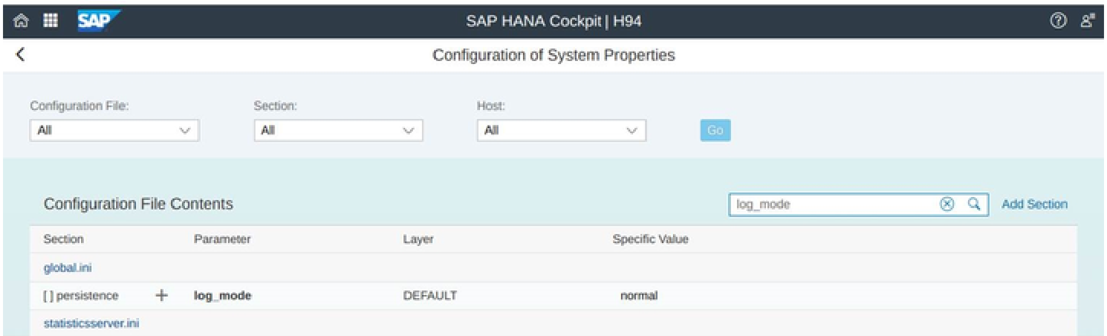

Figure 198: Parameter: log_mode

Another mode is also available. If the log_mode is set to **`overwrite`**, log segments are freed by savepoints and no log backup is performed. For example, this is useful for test installations that do not require back up or recovery. Automatic log backups can prevent log-full situations from arising.

Note:
Do not use the **`overwrite`** setting for production systems. When log_mode is set to **`overwrite`**, point-in-time recovery is not possible. For recovery, only data backups are used; the logs are not used. You can only select the **`Recover the database to a specific data backup`** recovery option.

Caution:
When you change the log mode, restart the database system to activate the changes, and create a full data backup of the database.

Parameter: log_buffer_size_kb

The log_buffer_size_kb parameter sets the size of one in-memory log buffer in kilobytes. If you set a higher buffer size, throughput increases, but at the cost of COMMIT latency.

Parameter: log_buffer_size_kb

Type	Range	Unit	Changeable
integer	128 – 16384	KB	online

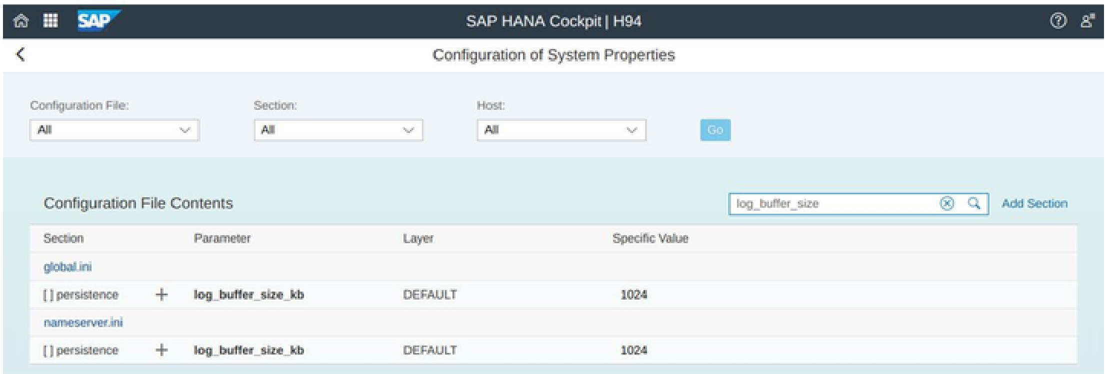

Figure 199: Parameter: log_buffer_size_kb

A higher buffer size increases the throughput at the cost of COMMIT latency. During COMMIT of a transaction, if all preceding buffers are already flushed, this data must be flushed to the I/O subsystem.

Parameter: content_vendor

A delivery unit is a collection of packages that are transported together. To ensure that all the packages belonging to your application are transported consistently together within your system landscape, assign them to the same delivery unit. Each delivery unit has a unique identity.

The identity of a delivery unit consists of two parts: a vendor name, and a delivery-unit name. The combined ID ensures that delivery units from different vendors are easily distinguished and that they follow a pattern that SAP uses for various software components.

To create and manage delivery units, first maintain the identity of the vendor with whom the delivery units are associated, and in whose namespace the packages that make up the delivery unit are stored.

So, before creating a delivery unit, define the content_vendor parameter in the indexserver.ini file.

Note:
If the parameter value for content_vendor is UNDEFINED, you cannot create delivery unit.

Parameter: content_vendor

Type	Range	Unit	Changeable
string	n/a	n/a	online

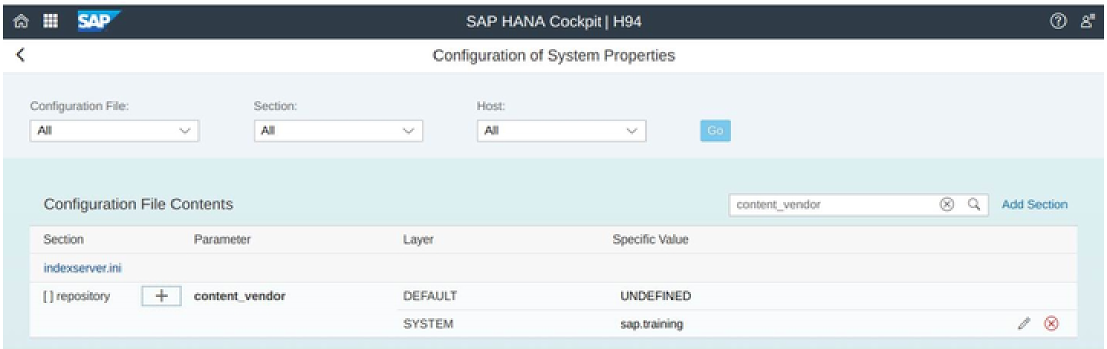

Figure 200: Parameter: content_vendor

Configuration Files

SAP HANA stores the configuration parameters in configuration files on operating system level. During the start of the SAP HANA database, these files are read and the changes are activated.

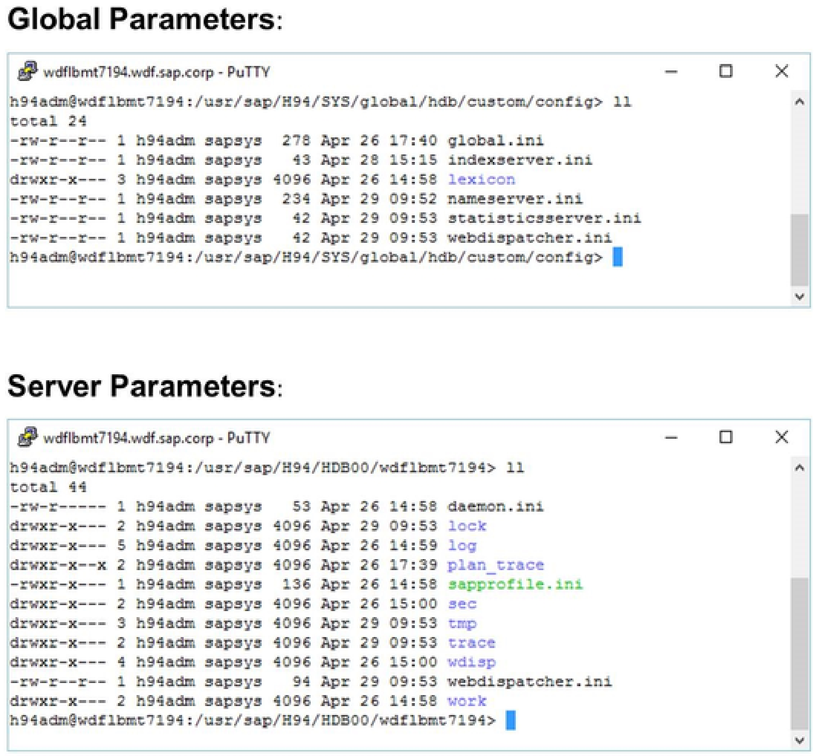

Figure 201: Configuration File Locations

Configuration files (.ini files) are only created in the file system if customer-specific changes are made to them after installation. If no customer-specific changes are made, these directories remain empty.

The configuration files are located in the following directories:

- `/usr/sap/<SID>/SYS/global/hdb/custom/config`
- `/usr/sap/<SID>/HDB<instance number>/<host name>`

During installation of SAP HANA database, the following customer-specific configuration files are created:

- `sapprofile.ini`

  This contains system identification information, such as the system name (SID) or the instance number.

- `daemon.ini`

  This contains information about which database services to start.

- `nameserver.ini`

  This contains global information for each installation. The landscape section contains the system-specific landscape ID and assignments of hosts to roles MASTER, WORKER, and STANDBY.

LESSON SUMMARY

You should now be able to:

- Configure the SAP HANA database

Unit 9
Lesson 3

Performing Regular Database Administration Tasks

LESSON OVERVIEW

This lesson describes typical tasks of an administrator and how SAP HANA cockpit can be used to support their execution.

Business Example

After installation, you require an overview of which tasks you need to perform as administrator and how to complete these tasks using SAP HANA cockpit. You want to ensure good performance for the processing of your SAP HANA database. Therefore, perform regular checks and take preventive action, if required.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Perform regular database administration tasks

Overview of Administrative Tasks

You can perform administrative tasks for the SAP HANA database using the SAP HANA cockpit. The detailed overview of the SAP HANA cockpit allows the database administrator to manage the SAP HANA database. It also allows you to create and manage users and roles.

The administrative tasks are divided into three categories: initial, regular, and on-demand tasks.

Overview of Database Administrative Tasks

The administrative tasks are as follows:

- Initial tasks
  - Perform a full data and a file system backup
- Regular tasks
  - Check the system status
  - Check the status of the services
  - Perform data backups
  - Check the alerts and error logs
  - Check the performance
  - Check the volume configuration

    - Maintain configuration
    - Check the system information
- On-demand tasks
    - Check the diagnosis files
    - Activate and analyze additional traces
    - Avoid LOG FULL situations
    - Avoid log backup area becoming full
    - Monitor disk space that is used for diagnosis files

Initial Tasks

Performing an Initial Backup

After the initial setup and initial load, you must perform a full data and file system backup, including a configuration backup. This to safeguard the changes that you made to the database data and configuration.

As of SAP HANA 2.0 SPS01, the SAP HANA database is set to Multitenant Database Containers by default. This means that the initial full data backup needs to be performed on the SystemDB and the Tenant database to safeguard the recoverability of the SAP HANA system.

Note:
The topic of backup and recovery is covered in detail in the following unit, Backup and Recovery.

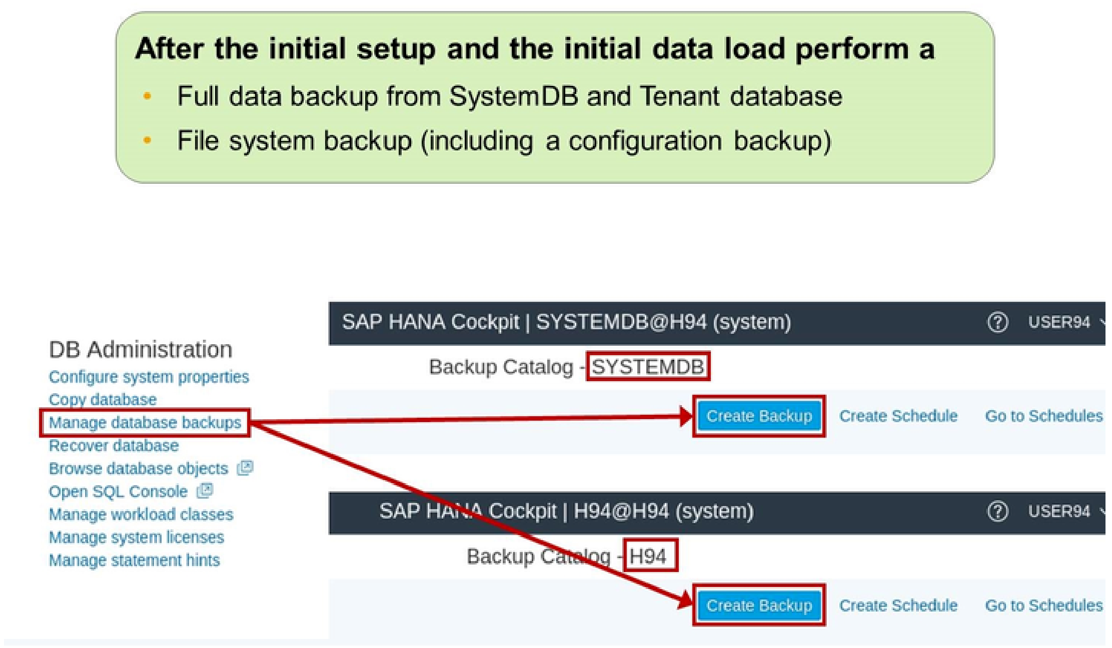

Figure 202: Backup from SAP HANA Cockpit

Installing a Valid License for the SAP HANA Database

At least one license key is required to use the SAP HANA system. This license key must be installed in the system database. There are two kinds of license key: temporary license keys and permanent license keys.

License Keys for Tenant Databases

You can install permanent license keys in individual tenant databases. The license key installed in a tenant database is valid for that database only and takes precedence over the license key installed in the system database. If a tenant-specific license key is not installed, the system database license key is effective in the tenant database. The earlier unit, Post Installation Tasks, explains how to install a license into the SAP HANA database.

Regular Administration Tasks

Checking the Database System Health

From the My Resources page, select the Aggregate Health Monitor application to regularly check the high-level status of all the SAP HANA databases under your responsibility. All the resources, SystemDB and Tenant databases, assigned to your SAP HANA cockpit user account will be displayed in the Aggregate Health Monitor application.

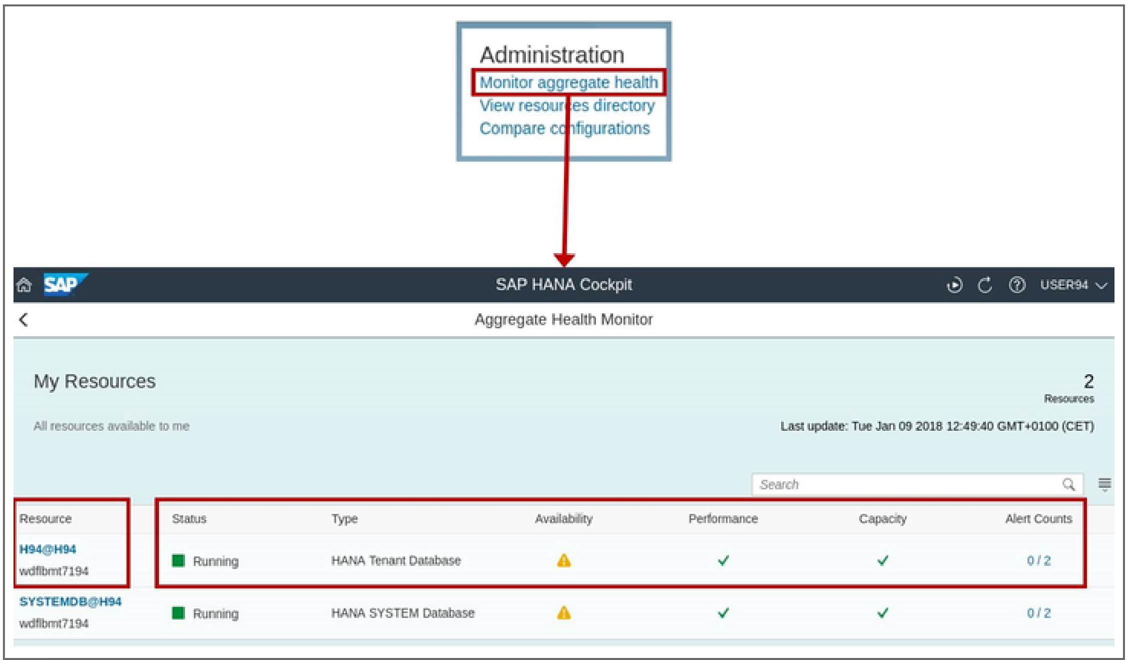

Figure 203: Monitor Aggregated Health Application

The Aggregate Health Monitor displays the five high-level status indicators:

- Status

  Are managed resources running? The possible statuses are Running and Stopped.

- Availability

  Are the managed resources reachable on the network? Are they able to serve the business needs of their users, including humans and applications? Performance and capacity issues can affect availability.

- Performance

  Are the managed resources meeting the response time expectations of database users, including humans and applications?

- Capacity

  Do the managed resources have the system resources to support their applications?

- Alerts

  The monitor displays the number of high and medium priority alerts.

Detailed Database Overview

When you select a database in the Aggregate Health Monitor application, you will navigate to the detailed database overview page of the SAP HANA cockpit 2.0. This overview page displays the most important system KPIs, such as the following:

- The Overall Database Status tile shows the status of all database services.
- The Alerts tile shows a summary of the most important system alerts generated by the statistics service, which is the monitoring tool for the database. It collects statistical and performance information using SQL statements.
- The Memory Usage tile shows memory usage of the host over the last 30 minutes.
- The CPU Usage tile shows the CPU usage of the host over the last 30 minutes.
- The Disk Usage tile shows the disk usage of the host over the last 30 minutes.

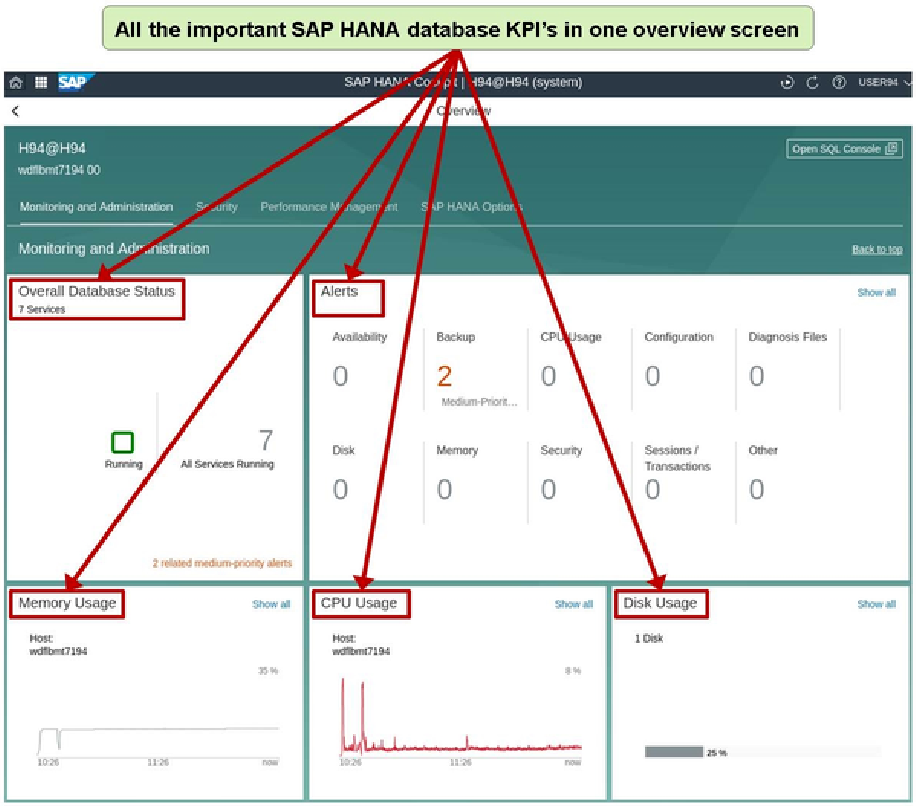

Figure 204: Checking the Overall System Status Overview in SAP HANA Cockpit

> Note:
> For the SystemDB, the Overview page also shows the Overall Tenant Status and Top Tenant Databases with Alerts tiles. With these two additional tiles, you can easily identify tenants that need your attention.

Overall Database Status Tile

To monitor the health of your SAP HANA database in more detail, for example, to troubleshoot performance bottlenecks, you can analyze the status and resource usage of individual database services.

The overall database status is displayed on the Overall Database Status tile. If high priority alerts exist, these are shown on the tile as well. Open the Overall Database Status app by choosing the tile.

The status of all the services in the system is displayed. For each service, detailed information about its memory, CPU consumption, and status is available. You can customize the view and add further columns.

As an administrator, you may need to perform certain operations on all or selected services (for example, start missing services, or stop or kill a service).

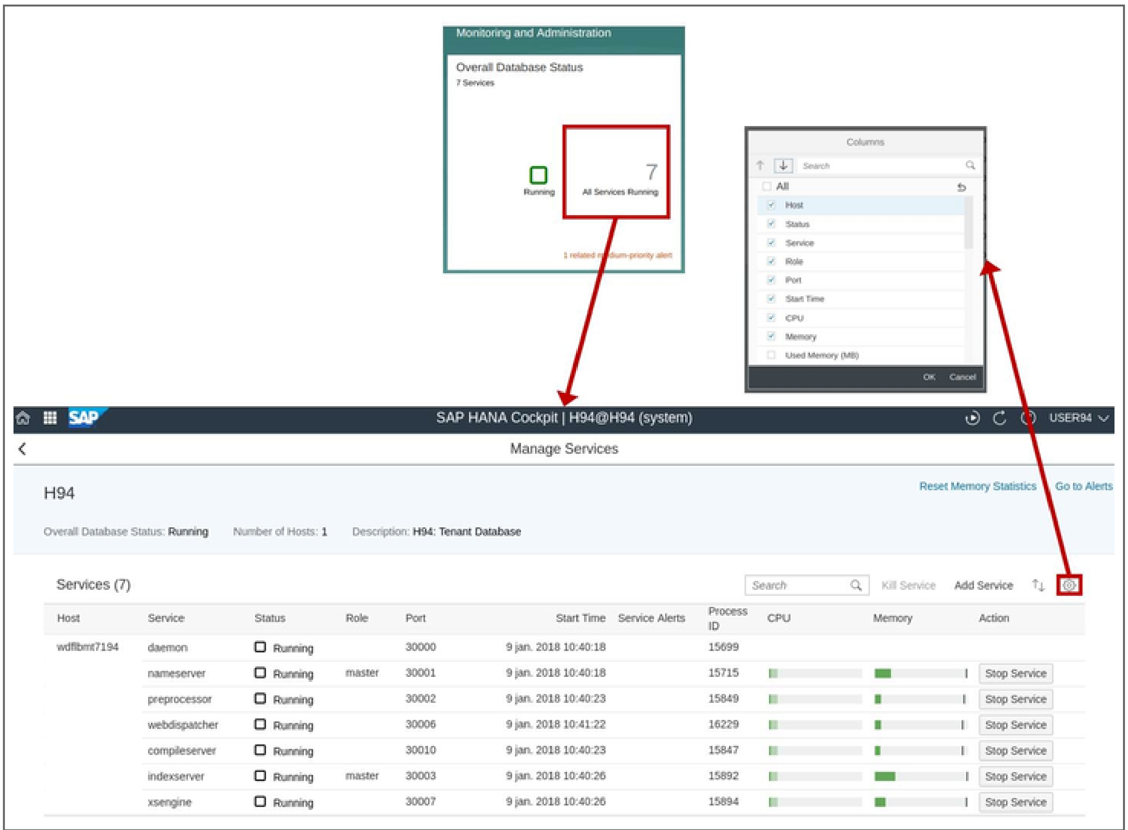

Figure 205: Overall Database Status Tile

Checking the Status of the Services

On the Manage Services page, check that all services that belong to your database are running. Verify that the name server, preprocessor, index server, webdispatcher, xsegine, and compileserver are running on the database system. The Manage Services page also contains information about the status of all database services. Correct running services are indicated with a green icon and have the status **Running**.

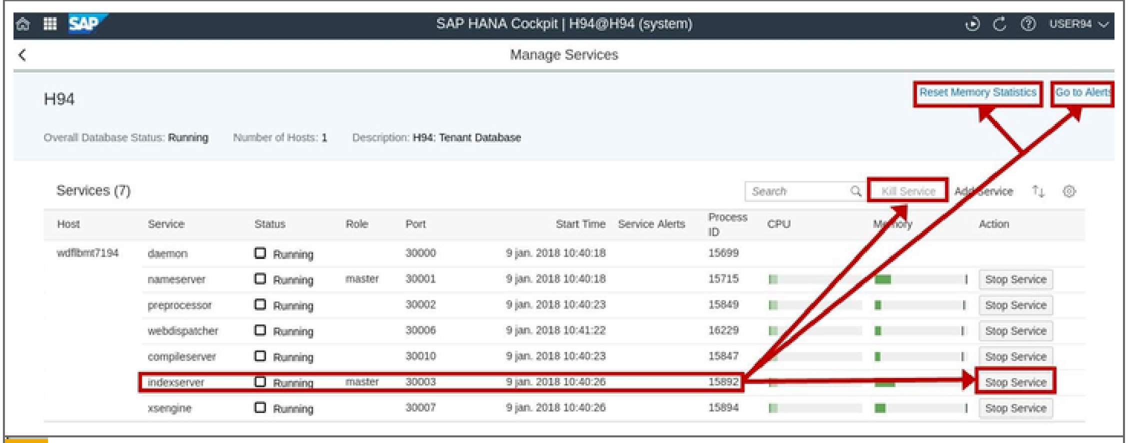

Figure 206: The Manage Services Screen

In addition to the information about resource usage, on the bottom task bar you can choose Start Missing Services or Stop a selected service. Use the More (◦◦◦) button to open the menu where you can choose to Kill or Remove services. When using the Stop or Kill command, the selected service is stopped or killed, and then starts automatically again. Because all services are restarted automatically when they are stopped, there is no need to start single services manually.

You can reset the memory statistics from the More (◦◦◦) menu.

Memory and CPU Usage Details

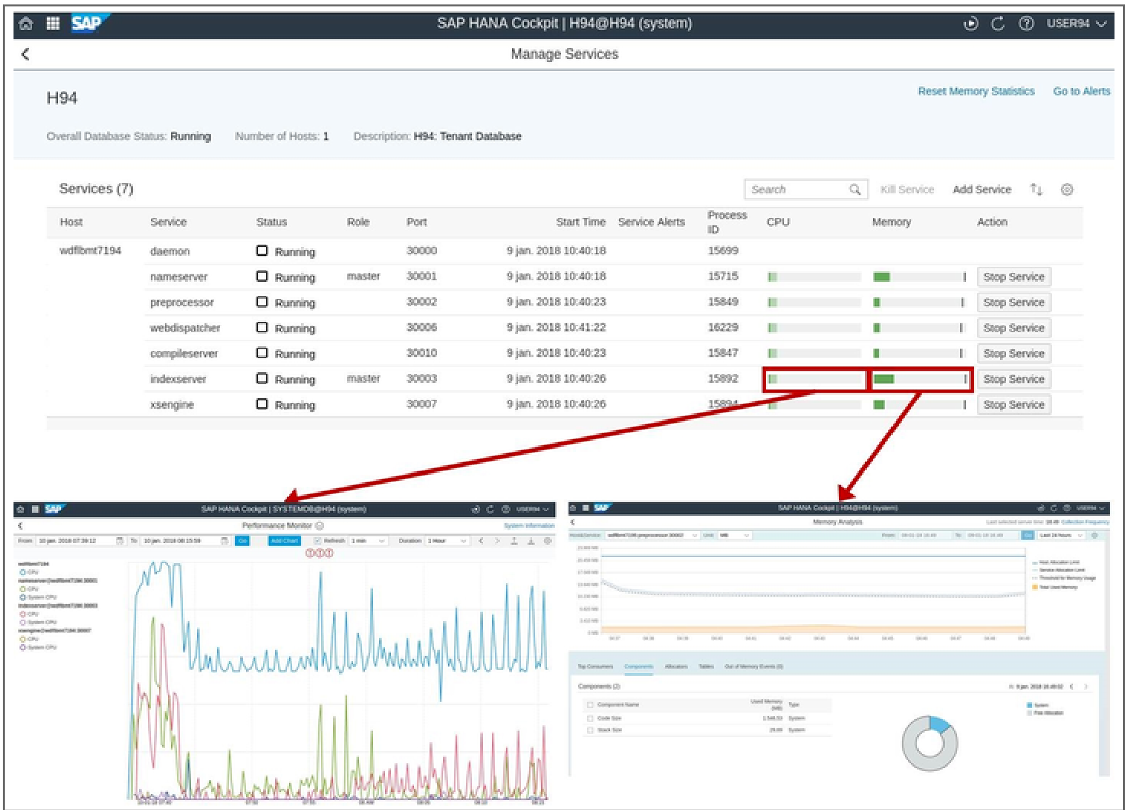

Figure 207: Memory and CPU Usage Details

The Used CPU bar shows the overall CPU usage per service, but this CPU usage is used by several services. Choose the CPU bar to open the CPU Performance Monitor screen.

The Used Memory bar shows the overall memory usage per service, but this memory is used by several components. Choose the Memory bar to open the Detailed Memory Allocation Statistics screen.

Use the Memory Allocation Statistics application to visualize and explore the memory allocation history of the components for each service in the SAP HANA database collected by the statistics service.

Memory Allocation Statistics

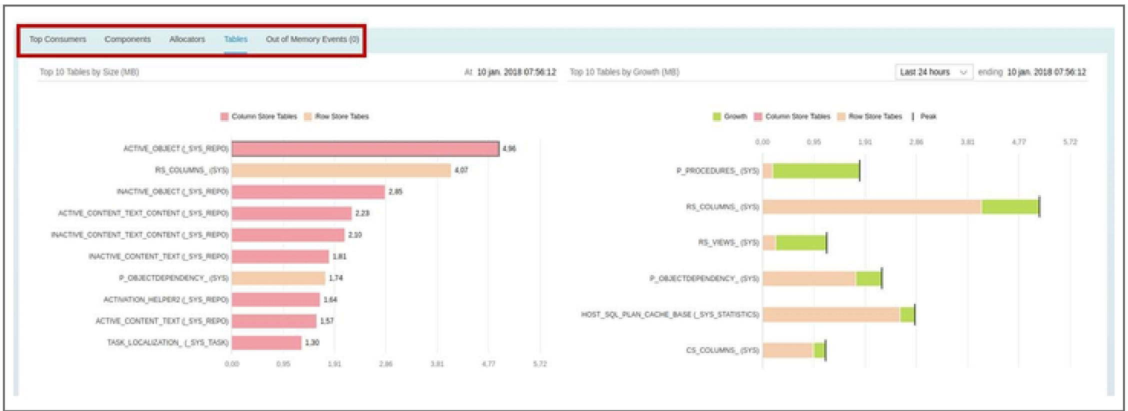

Figure 208: Memory Allocation Statistics

The SAP HANA Cockpit provides a graphical breakdown of the following main categories of memory usage: physical memory, SAP HANA database, table data, and database management. You can view other information regarding the current size of used resources on the Overview tab of the Administration editor.

The following information displays in the screen areas that are identified in the figure, Memory Allocation Statistics:

1. The components of the selected service listed in descending order of current used memory, by default.
2. The current breakdown of SAP HANA used memory is displayed as a pie chart.
3. Allocations of the selected component are listed in descending order of current used inclusive memory, by default.
4. The current breakdown of memory usage of the 10 highest consuming allocations is displayed as a pie chart.

CPU Performance Monitor

Figure 209: CPU Performance Monitor

The CPU Performance Monitor enables you to visualize and explore the history of the CPU usage. Additional KPIs can be added for a greater insight into the used system resources.

CPU Performance Monitor Application

The features of the CPU Performance Monitor are as follows:

- Displays the CPU usage history broken down into several components as a graph.
- More KPIs can be added to the graph.
- The time period shown in the graph can be changed to the required time frame.

Disk Performance Monitor

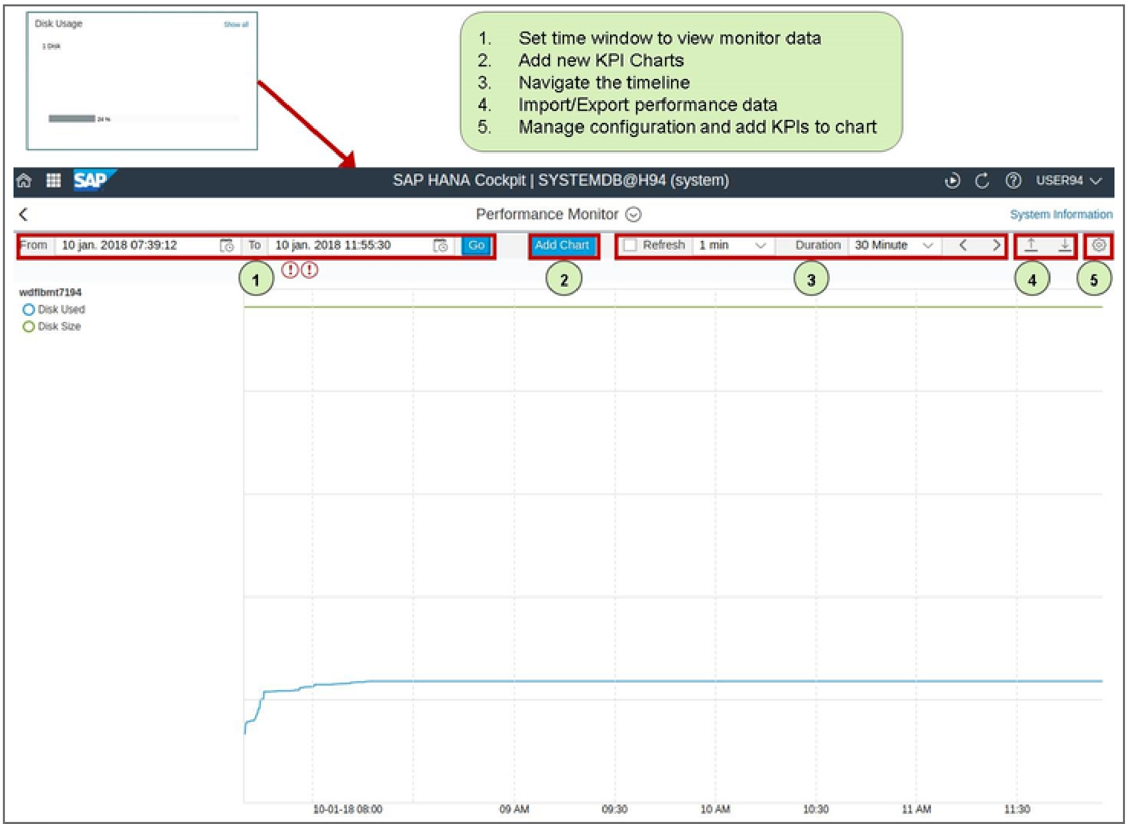

Figure 210: Disk Performance Monitor

The Disk Usage tile displays the fill level per disk and indicates space problems. Open the Disk Performance Monitor application to investigate possible problems over time.

The Disk Performance Monitor tile indicates disk usage on disks belonging to the SAP HANA database. Values are displayed for all disks. The disk with the highest or most critical disk usage is also shown.

This tile provides access to the Resource Utilization application where you can visualize and explore the history of disk usage. Other key system resources (such as CPU, memory, and disk) can be added if required.

The tile analyzes bottlenecks, identifies patterns, and forecasts requirements. Open it through the context-menu of the specific SAP HANA system.

Alerts Tile

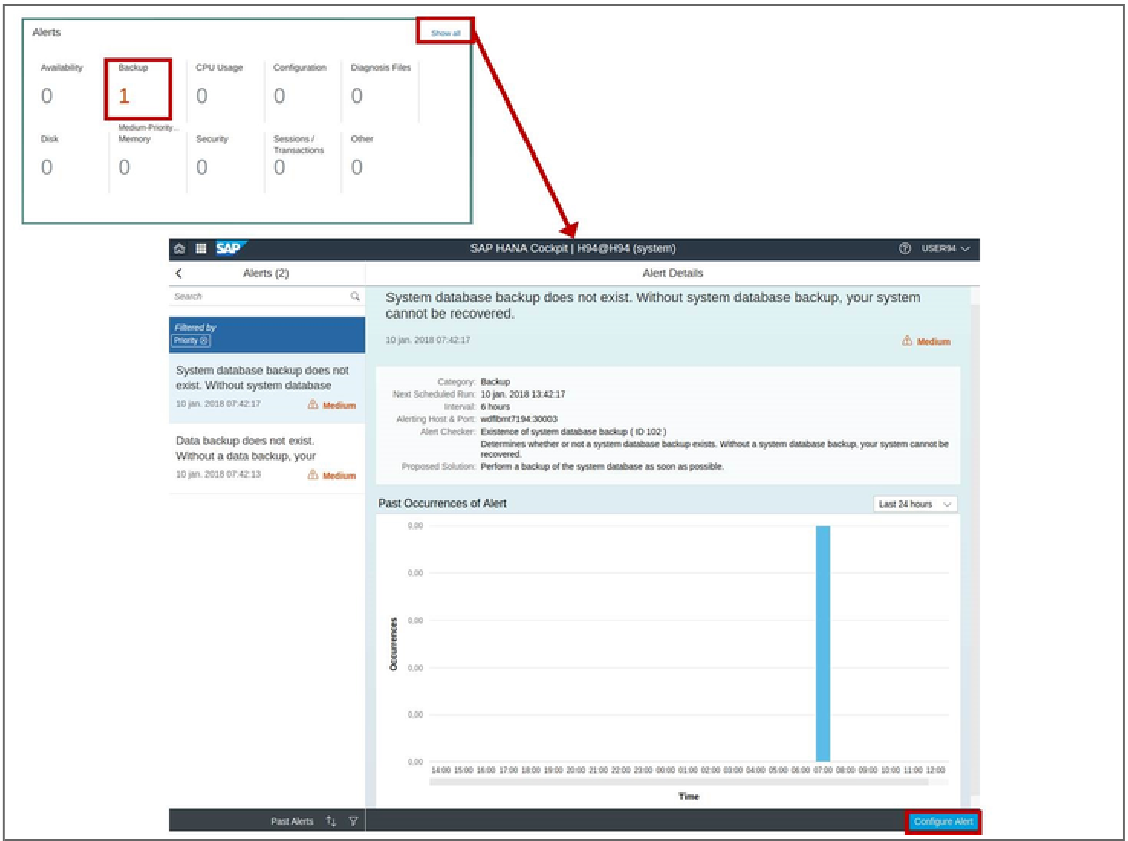

Figure 211: Alerts Tile: Show All

The statistics service is one of the main components of the monitoring infrastructure of the SAP HANA database. It performs regular checks and issues an alert when an alert condition is fulfilled. The alerts are aggregated and the most important alerts are shown on the Alerts tile.

You can use the filter function to filter for specific alerts or priorities.

The current alerts are listed. However, when you select Past Alerts, you can view a summary of all the alerts that occurred in the system over the last 30 days.

Alerts Tile: Alert Configuration

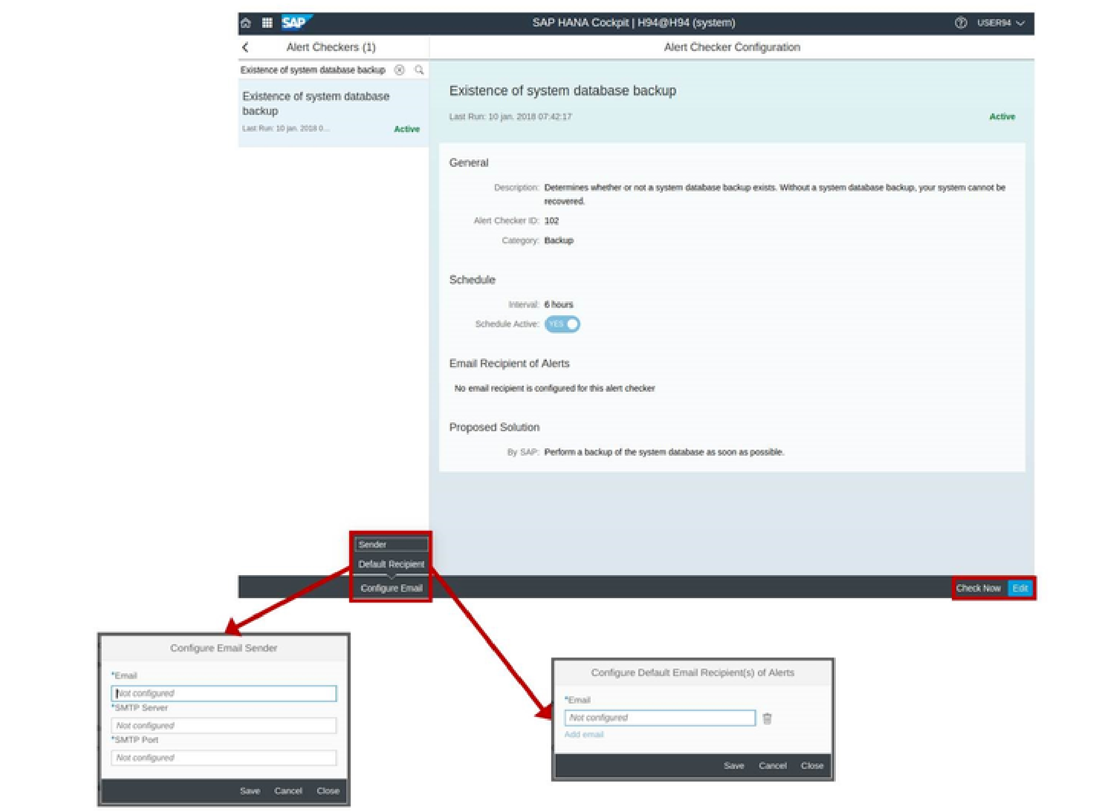

Figure 212: Alerts Tile: Alert Configuration

When you select an alert, you can set up e-mail notifications and edit the alerts properties. You can also run the Check Now feature, to check if there are new alerts.

With the e-mail notification feature, you can specify the e-mail sender, SMTP Server, SMTP Port, and Recipients. After the e-mail notification is set up, the SAP HANA database system sends notification e-mails to the recipients whenever this alert occurs.

Configure Alerts Application

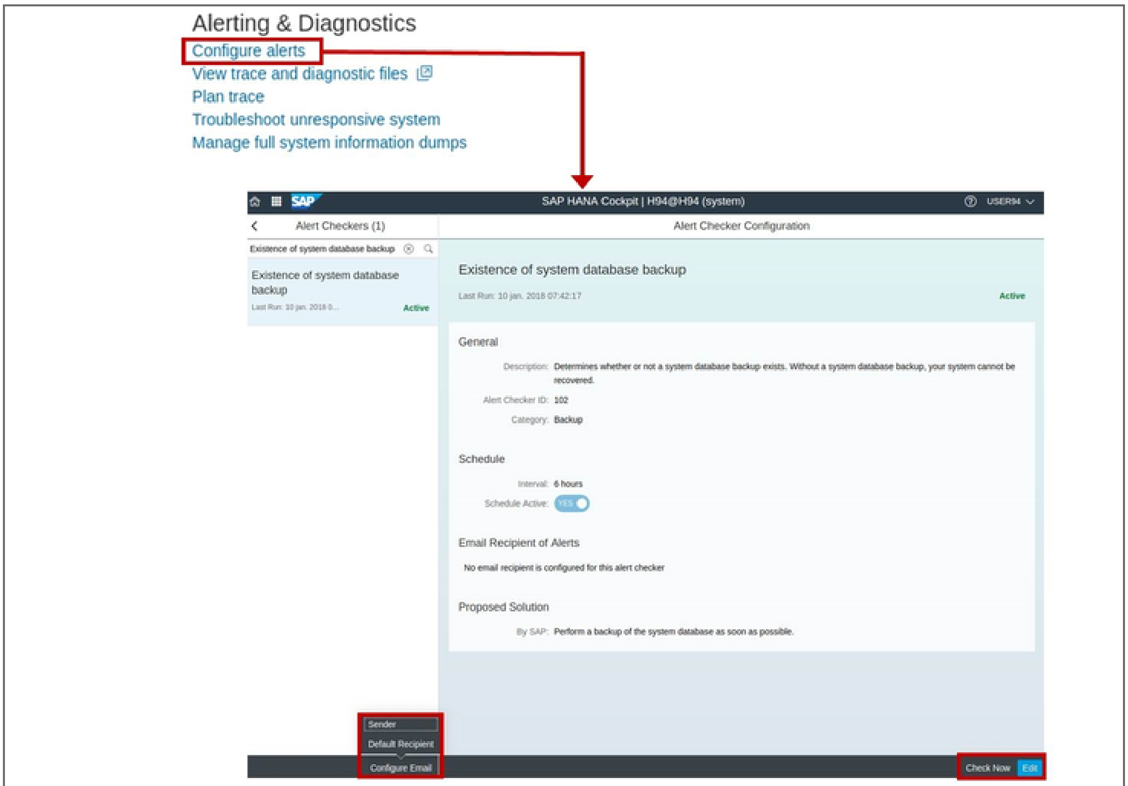

Figure 213: Configure Alerts Application

With the Configure Alerts application, you can view of all the SAP HANA database system alerts that are available. You can also edit the properties and set up e-mail notifications for selected alerts.

The Configure Alerts application also provides a search function so that you can easily find a specific alert.

All checks are displayed in list format on the left. For detailed information about a specific check on the right, simply select it.

You can change the threshold values that trigger alerts of different priorities. In addition, you can switch checks off and on.

In the edit mode of an alert, you can also set the priority. The priority of the alert indicates the severity of the problem. It depends on the nature of the check and the configured threshold values. For example, by default, a low priority alert is issued if 90% of available disk space is used. If 98% is used, a high priority alert is issued.

For each check, you can configure the following:

- A specific e-mail alert
- A threshold
- The schedule and the activation of this concrete check

In addition, you can choose the Run the check now button. This is useful for when the check is running every six hours (backup), for example, and you want to know if the process is working now.

Threads Tile

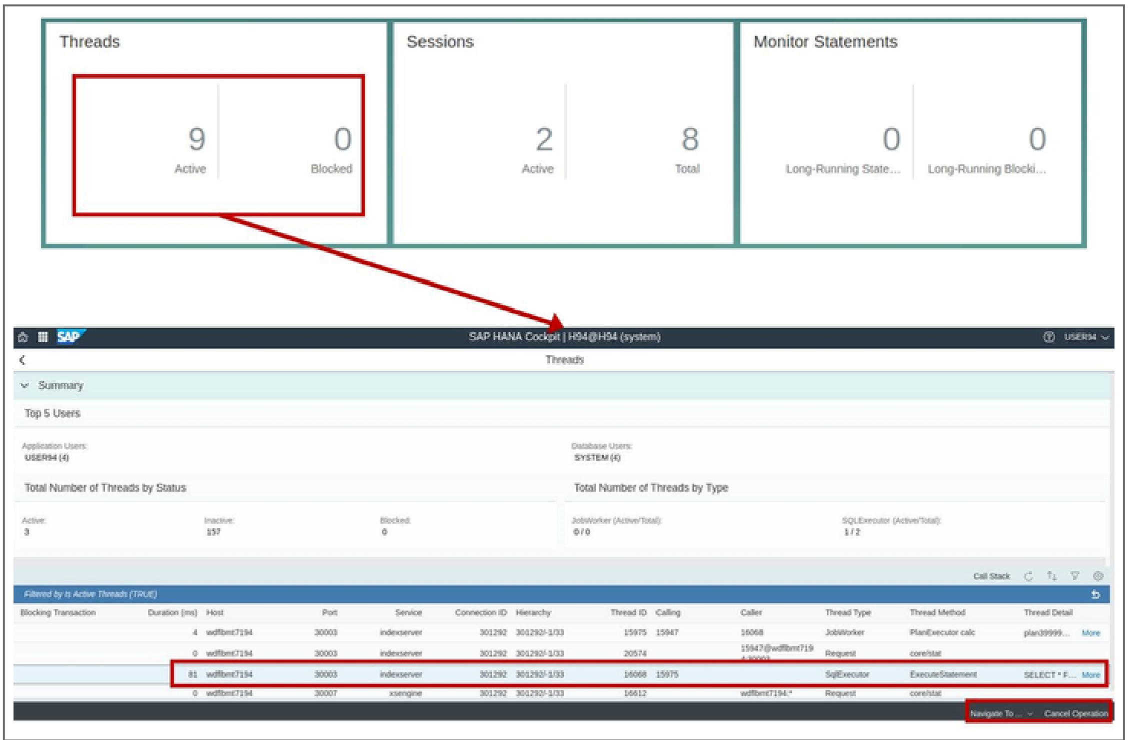

Figure 214: Threads Tile

The Threads tile provides an overview of the active and blocked number of threads in the SAP HANA database system. This information is useful when investigating performance problems because it indicates how many threads are active and or blocked in the systems.

When you select the Threads tile, you get detailed information on the current threads running in the SAP HANA database system.

An overview of the top five users is available for quick analysis. At the bottom, all the current threads are displayed, with detailed information. You can use the Filter option to search for a specific thread.

When a thread causes problems in the system, you can cancel this thread. Select the thread and choose the Cancel Operations button to cancel the running thread. Only do this when you are sure that this thread is causing the problems.

Sessions Tile

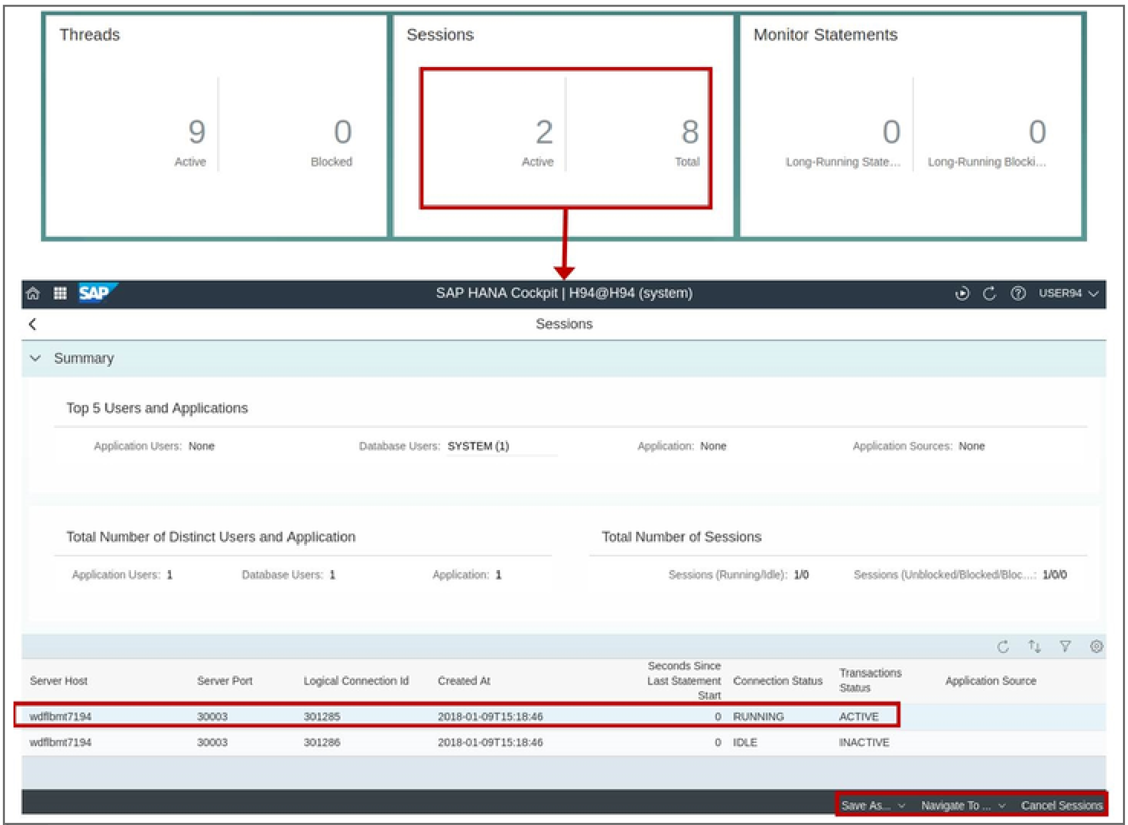

Figure 215: Sessions Tile

The Sessions tile provides an overview of the active and total number of sessions connected to the SAP HANA database system. This information is useful when investigating performance problems because it indicates how many sessions are connected to the system.

When you select the Sessions tile, you can view detailed information on the current sessions that are connected to the SAP HANA database system.

A top five of users and applications is available for quick analysis. At the bottom, all the current sessions are displayed, with detailed information. You can use the Filter option to search for a specific session.

When a session causes problems in the system, you can also cancel this session. Select the session and choose the Cancel Sessions button to cancel the running session. Only do this when you are sure that this session is causing the problems.

Monitor Statements Tile

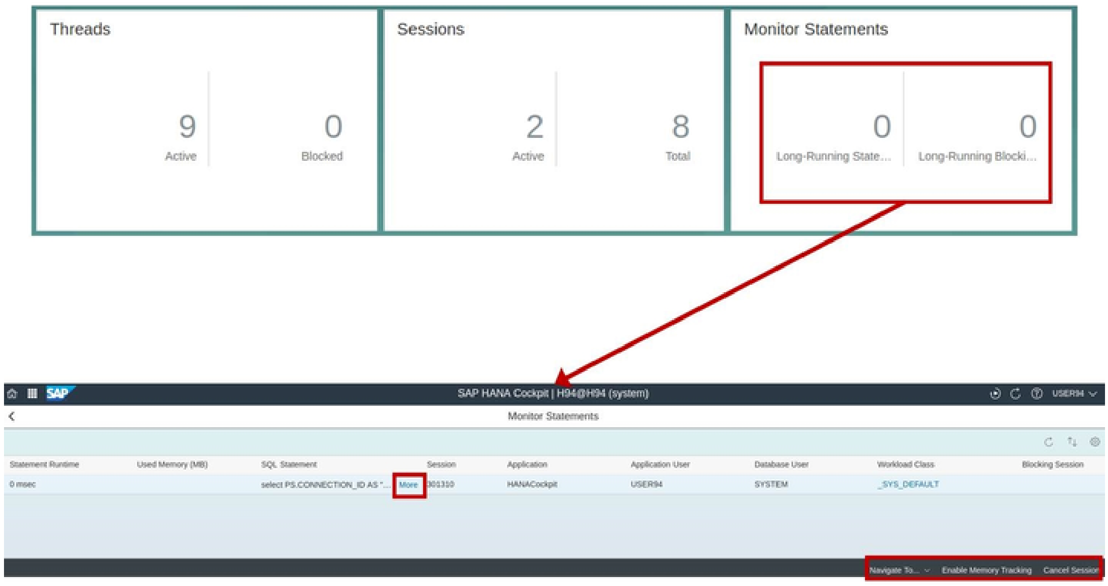

Figure 216: Monitor Statements Tile

The Monitor Statements tile provides an overview of Long Running Statements and indicates if there are Blocked Statements. This information is useful when investigating performance problems because it shows how many long running statements are blocked in the systems.

When you select the Monitor Statements tile, you get detailed information on the current statements running in the SAP HANA database system.

The list shows all the Long Running Statements in the system. In the details, you can see the Session ID and the user information, the database user, and the application user. With this user information, you can quickly identify which user session is causing the problem. Also, the Blocking Session is shown.

When a running statement is the cause of the problem in the system, this statement can be stopped by canceling its corresponding session. Select the statement and choose the Cancel Session button to cancel the running statement. Only do this when you are sure that this statement is causing the problems.

To get more detailed memory usage information on running statements, you can enable or disable memory tracking.

System Replication Tile

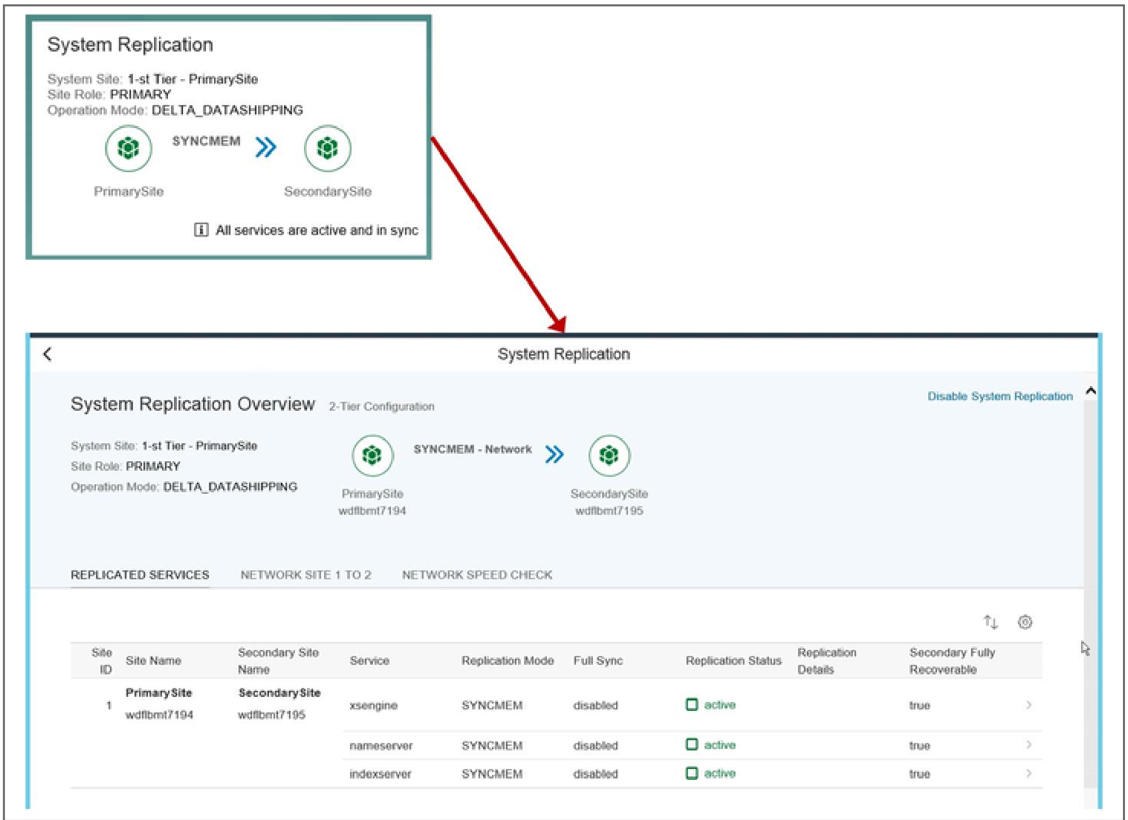

Site ID	Site Name	Secondary Site Name	Service	Replication Mode	Full Sync	Replication Status	Replication Details	Secondary Fully Recoverable
1	PrimarySite wdflbmt7194	SecondarySite wdflbmt7195	xsengine	SYNCMEM	disabled	active		true
			nameserver	SYNCMEM	disabled	active		true
			indexserver	SYNCMEM	disabled	active		true

Figure 217: System Replication Tile

The System Replication tile provides an overview of the System Replication status. This information is useful when system replication is set up and active because it provides a quick overview on of the System Replication status.

When you select the System Replication tile, you get detailed information on the current status of the system replication running in the SAP HANA database system.

The list shows the information on the primary and secondary site status. It shows the chosen replication mode and replication status. With this information, you can quickly identify if the secondary site is ready for a system takeover.

In this course, the unit on high availability and disaster tolerance provides more information on the possible system replication options and shows how to set up system replication.

General Information Tile

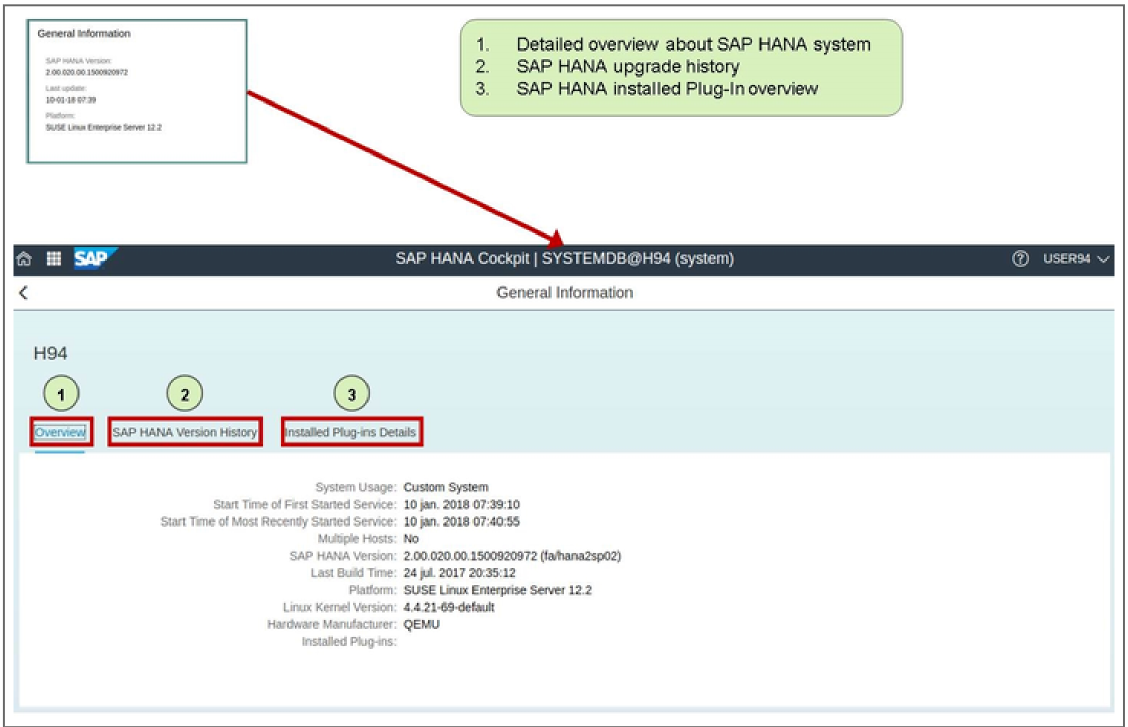

Figure 218: General Information Tile

The General Information tile provides an overview of the installed SAP HANA and Linux versions. This information is useful when determining if a certain patch is installed or not.

When you select the General Information tile, you get detailed information like System Usage Type, SAP HANA Version, Linux Kernel Version, and Hardware Manufacturer.

On the SAP HANA Version History tab, a list shows the upgrade history of the SAP HANA database system.

The Installed Plug-ins Details tab shows detailed information on the installed plug-ins.

On-Demand Tasks

- In case of problems with the SAP HANA database, you can check log and trace files for errors from the View Trace and Diagnostics application under the Alerting and Diagnostics area of the SAP HANA cockpit.
- You can turn on and configure several traces in the SAP HANA cockpit Alerting and Diagnostics area.
- In certain situations, you have to restart the system (for example, after a power failure). This can be done using the Overall Database Status tile in the SAP HANA cockpit.
- Avoid **LOG FULL** (file system full) situations by having the Auto Log Backup activated.
- Avoid log backup area becoming full by deleting obsolete log backup tiles.
- Monitor disk space using the Disk Usage tile in the SAP HANA cockpit.

Avoiding Log Full Situations

When the log is backed up, the backed up log segments remain on disk until they are released automatically after a savepoint. After the log is released, the oldest unused log segment is overwritten with new log entries. If there are no unused log segments, new log segments are created. If the disk becomes full and no more log segments can be created, a log full situation arises. When the log is full, no more logging is possible until the log backup is complete. Automatic log backup prevents log full situations from arising.

Prevent the log backup area from becoming full. Regularly archive old log backups to a different location by using the operating system commands.

If there are problems with the SAP HANA database, you can check log and trace files for errors. These log files are available in the View trace and Diagnostic files link in the Alerting and Diagnostics area of the SAP HANA cockpit.

LESSON SUMMARY

You should now be able to:

- Perform regular database administration tasks

Unit 9
Lesson 4

Configuring Traces

LESSON OVERVIEW

This lesson describes how to activate the traces within SAP HANA studio.

Business Example

You are an administrator and you want to activate a trace to analyze an issue.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Configure SAP HANA traces

Trace Configuration

This lesson explains how to activate and configure the different trace tools available in SAP HANA 2.0.

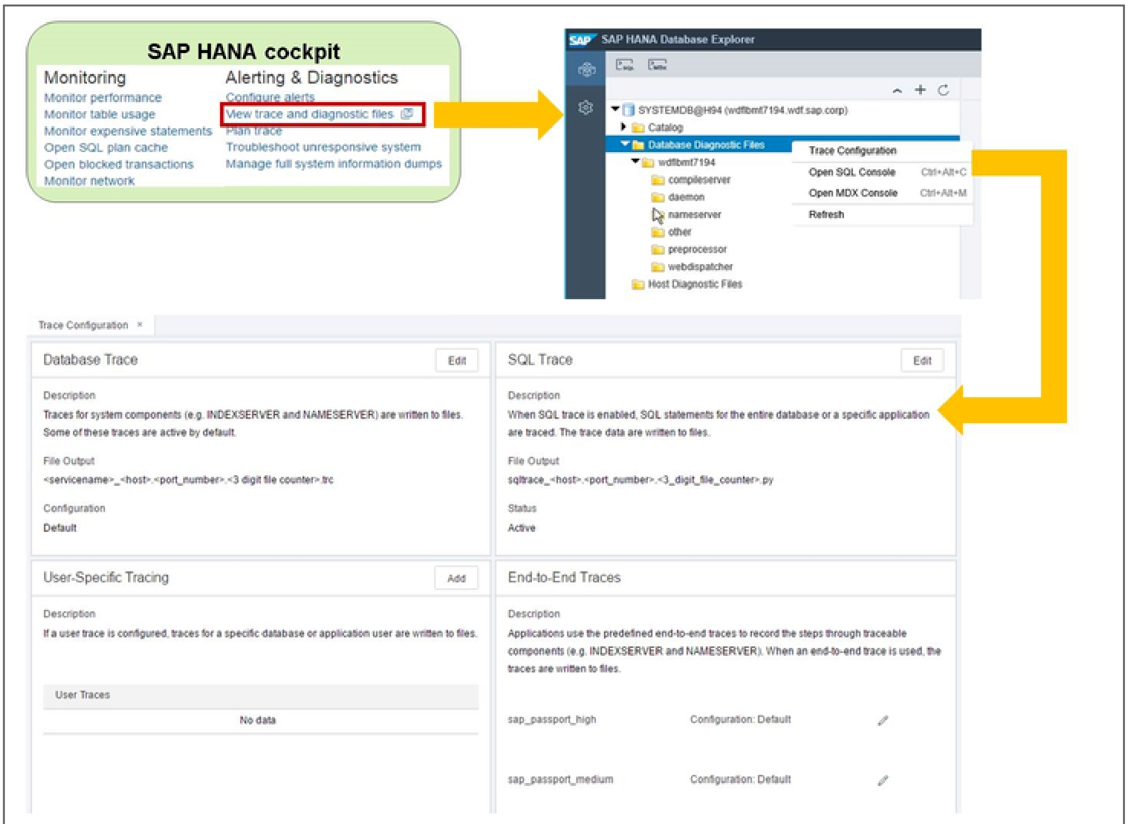

Figure 219: Trace Configuration

The various traces can produce detailed information about the actions of the database system. You can activate and configure traces in the Administration area of SAP HANA cockpit or via the SAP Web IDE for SAP HANA. Different configuration options are available for each trace.

> Note:
> To configure traces, you must have the system privilege **TRACE ADMIN**. To configure the kernel profiler, you must have the **SAP_INTERNAL_HANA_SUPPORT** standard role.

Useful Traces

You can use the following traces:

- **Database trace** (including user-specific and end-to-end database traces)

  The database trace records information about activity in the components of the SAP HANA database. Use this information to analyze performance and to diagnose and debug errors. Each service of the SAP HANA database writes to its own trace file. By default, the database trace is active with the default trace level ERROR.

- **SQL trace**

  The SQL trace collects information about all executed SQL statements and saves it as an executable python program. This is useful for recording a scenario. By default, the SQL trace is inactive.

- **Expensive statements trace**

  Expensive statements are individual SQL queries that have an execution time above a configured threshold. The expensive statements trace records information about these statements for further analysis. By default, the expensive statements trace is inactive.

- **Performance trace**

  The performance trace is a performance tracing tool built into the SAP HANA database. It records performance indicators for individual query processing steps in the database kernel. By default, the performance trace is inactive.

- **Plan trace**

  With the plan trace, you can visualize and analyze the execution plans for every query that has been executed in the specified application.

- **Kernel profiler**

  The kernel profiler is a sampling profiler built into the SAP HANA database. It collects, for example, information about frequent and expensive paths during query processing. By default, the kernel profiler is inactive.

> Note:
> Only SAP development support has the technical expertise required to interpret the information collected by the performance trace and the kernel profiler.

Expensive Statements Trace Configuration

Trace with Default Configuration Status

Table 9: Trace with Default Configuration Status

The table lists the traces and their default configuration status.

Trace	Default Configuration or Status
Database trace	Active with default trace level ERROR
SQL trace	Inactive
Performance trace	Inactive
Kernel profiler	Inactive
Expensive statements trace	Inactive
Plan trace	Inactive

To Activate the SQL Trace in the SAP Web IDE

1. Open the following URL: https://<hostname>:43##/sap/hana/ide/trace

   Note that ## is the instance number.

2. If requested, enter the following credentials:

Field	Value
User name	**SYSTEM**
Password	**Welcome1**

3. In the SAP Web IDE for SAP HANA, in the Trace window, choose the Configuration button.

4. In the SQL Trace tab, choose the Edit Configuration button (pencil icon) for the Index Server.

5. In the SQL Trace Configuration for Index Server dialog box, select the Active radio button.

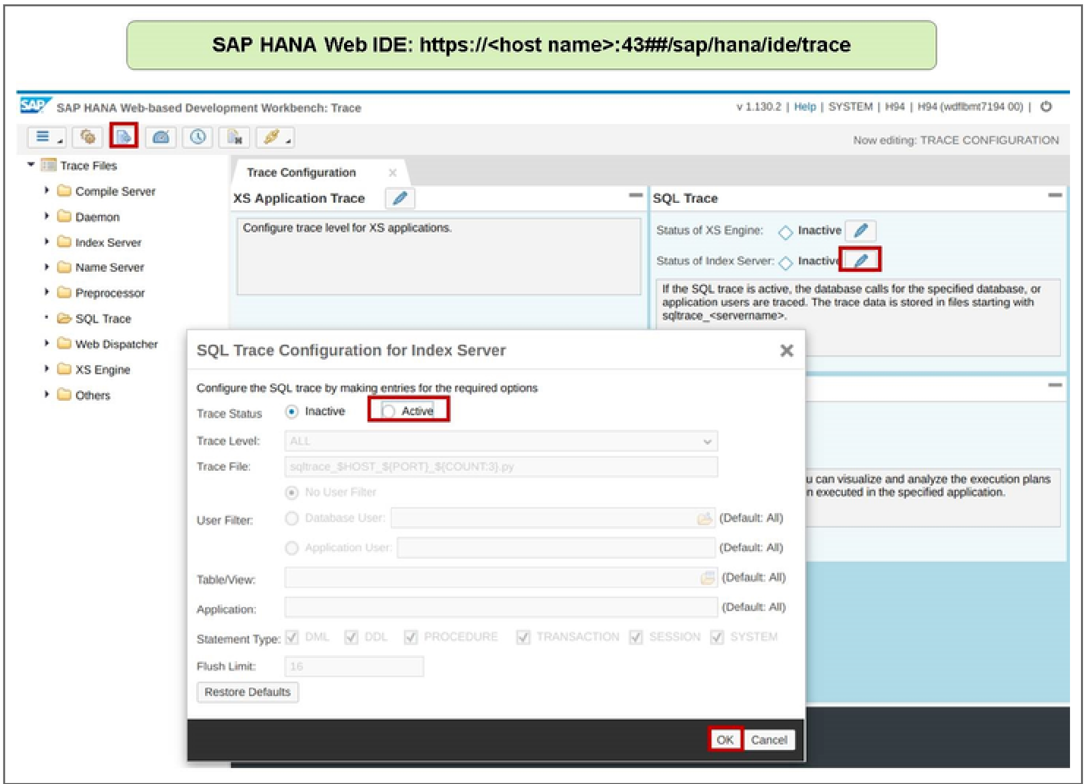

Figure 220: Activation of SQL Trace

6. Select the User Filter and the Table/View that you require.

7. To activate the SQL trace, choose OK.

> Note:
> Remember to disable the trace after sufficient trace information is gathered.

LESSON SUMMARY

You should now be able to:

- Configure SAP HANA traces

Unit 9
Lesson 5

Working with Diagnosis Information and Diagnosis Files

LESSON OVERVIEW

This lesson explains how to deal with SAP HANA diagnosis files.

Business Example

When there is an issue in the system, as an SAP HANA system administrator, you need to analyze diagnosis files for issue resolution.

As an SAP HANA administrator, when receiving support from SAP, you can send the diagnosis files to SAP. The configuration files must be backed up periodically with the database backup.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Work with diagnosis information

Working with Diagnosis Files

Working with Diagnosis Files in the Database Explorer

The SAP HANA database explorer allows you to diagnose and analyze errors in an SAP HANA database by viewing the relevant diagnostic files

In the database browser, diagnostic files for online databases are grouped by host and then by service. In a multi-host system, check each host folder to view all diagnostic files associated with a particular service.

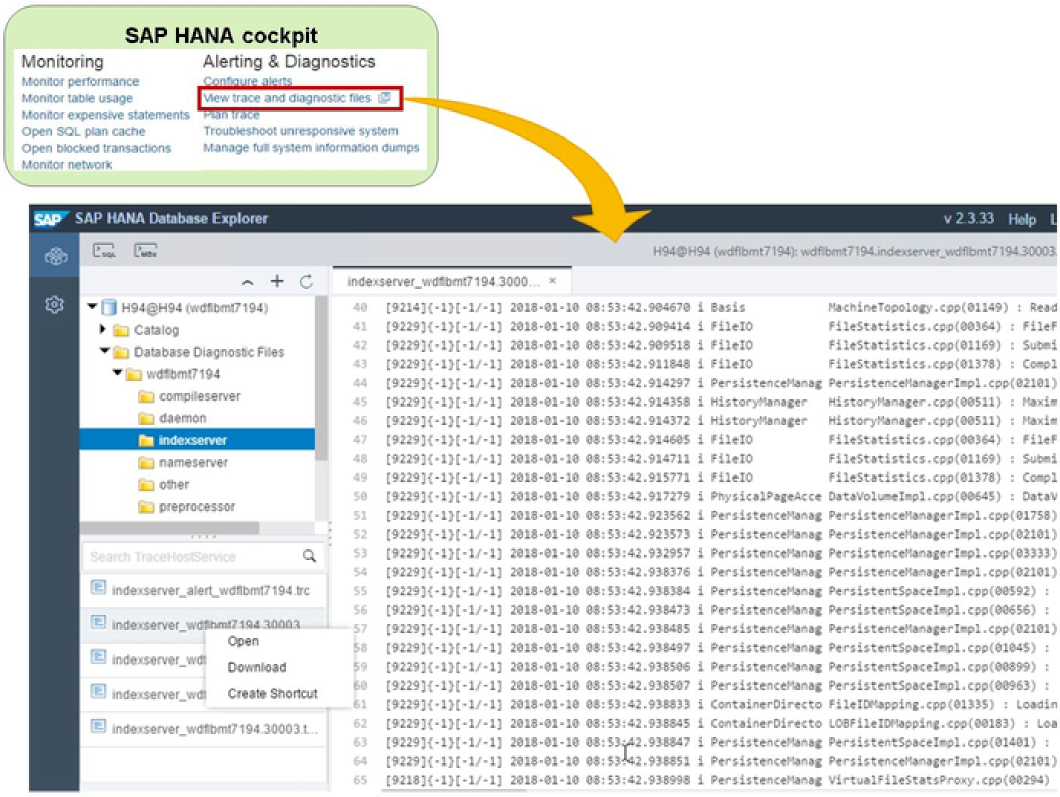

Figure 221: Diagnosis Files in the Database Explorer

Working with Diagnosis Files in SAP Web IDE for SAP HANA: Traces

Diagnosis files include log and trace files, as well as a mixture of other diagnosis, error, and information files. If there are problems with the SAP HANA database, you can check these diagnosis files for errors. You can also filter, delete, and download diagnosis files.

To access the trace and diagnostic files, open the Working with Diagnosis Files in SAP HANA WebIDE - Traces tool in a browser using the following URL: https://<host name>:43##/sap/hana/ide/trace. Note that ## is the instance number.

You can also view the trace and diagnostic files from the operating system level. By default, the trace and diagnostic files are stored in the following location: `/usr/sap/<SID>/HDB<instance>/<host>/trace`.

In this location, you can monitor the disk space that is used for diagnosis files, and delete files that you no longer need.

Note:
To view diagnosis files and delete trace files, you must have the **TRACE ADMIN** system privilege.

Diagnosis Files in SAP Web IDE: Overview

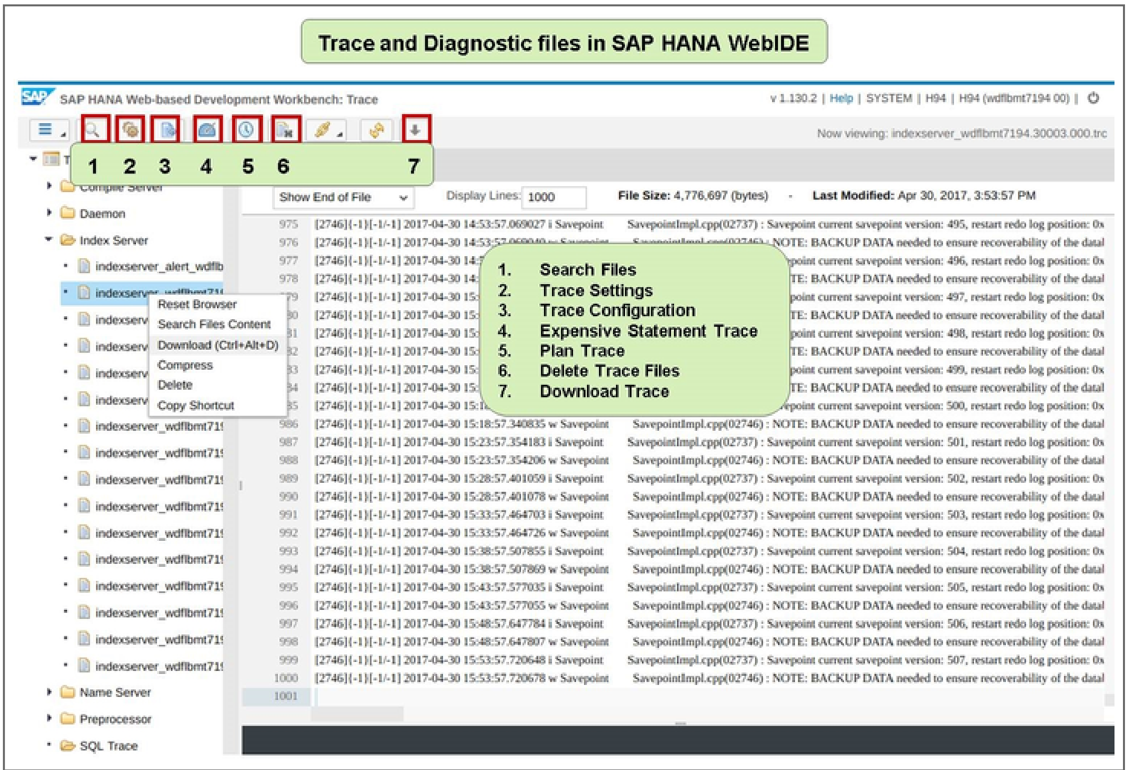

Figure 222: Diagnosis Files in SAP Web IDE: Overview

In the WebIDE: Trace files view, you can view the trace files and diagnostic files. The different files are displayed in separate tabs.

For large files, you might want to use a different tool than the browser. For this reason, a download option is provided in the context menu when you right-click the file.

Check the Diagnosis Files

If there are problems with the database, you can check the log and trace files for errors. These diagnosis files are available in the SAP Web IDE for SAP HANA on the Traces tile. When you select a file in the Trace Files list on the left, it opens in the editor. Every file opens in a new tab.

Trace file rotation prevents trace files from growing indefinitely by limiting the size and number of trace files. You can configure trace file rotation globally for all services in the system and for individual services.

> Note:
> The parameters maxfiles and maxfilesize , which are found in the global.ini, control the log rotation.

Filtering the Diagnostic Files List

To shorten the long list of diagnostic files, you can filter the list on a specific filename pattern.

The filter function is located in the context menu when you right-click Trace Files or a folder of an SAP HANA service. For example, to view everything related to nameserver, right-click Trace Files , select the Filters menu, and enter the filter pattern **`nameserver`**.

Display File

When you choose a file in the file list, it opens automatically in a new tab on the right side of the screen. Use the Show End of File, Show Start of File, or Show Entire File menu options to navigate large files more easily. When you filter the file in this way, you can specify how many lines you want to view.

> Note:
> Depending on the type of data in the diagnosis file, the number of lines actually displayed might be greater than or less than specified. This is because the data in some diagnosis files is fetched in bytes, and the number of bytes per line varies.

Display Diagnosis Files Ending with .gz (zipped) File. The features of .gz files are as follows:

- The .gz (zipped) file is automatically extracted and shown in the SAP Web IDE for SAP HANA.
- The last 1,000 lines are displayed by default.
- The Download button in the SAP Web IDE downloads the file to your local download folder. This location depends on your browser settings.

Compress Files

If you need to download a diagnosis file (for example, to send it to SAP Support), you can compress it first on the server. This is useful for large diagnosis files and for slow connections. To compress a file, right-click it and choose Compress. After compression, the file has the *.zip file format. You can select multiple files to compress.

Delete Files

You can delete files using the following options:

- Delete individual log files, trace (*.trc) files and other nontrace files shown in the file list.

  To delete one or more individual files from the list, select the file or files, and, in the context menu, choose Delete.
- Mass delete trace files (*.trc)

  You can mass delete trace files, for example delete all the trace files of a specific service, by choosing Delete Trace Files..., and selecting the required files.

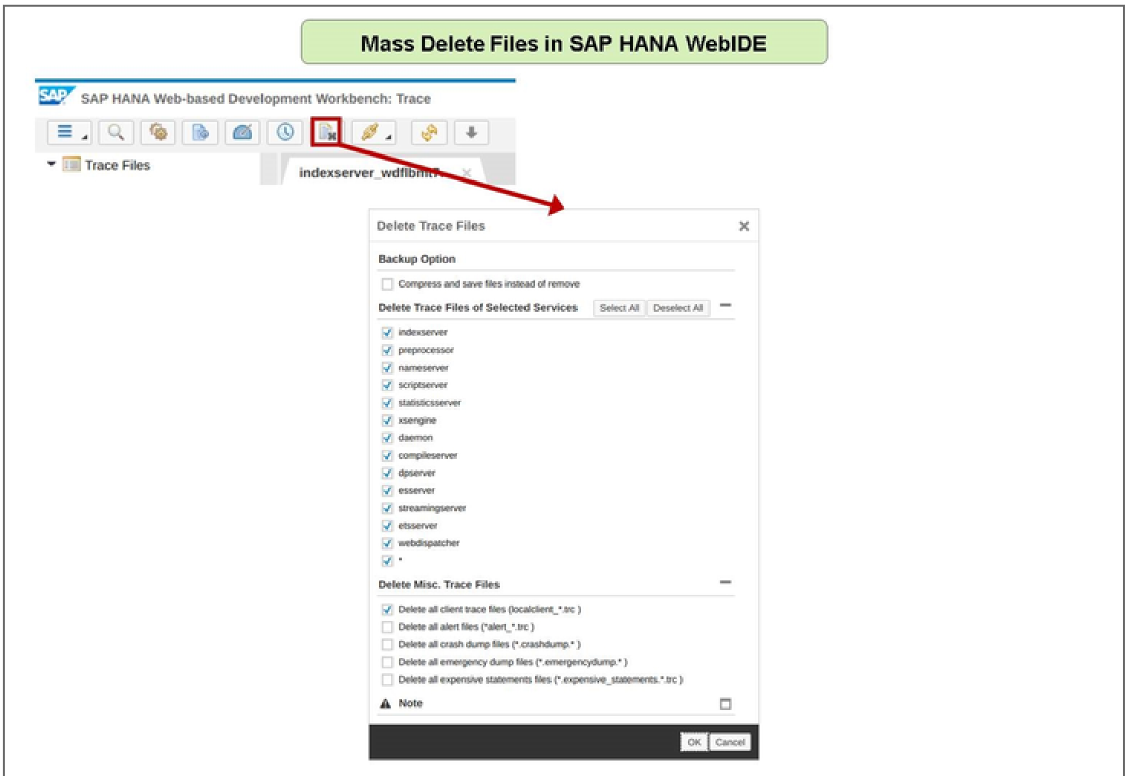

Figure 223: Diagnosis Files in SAP Web IDE for SAP HANA: Mass Delete

The deleted file might actually not be deleted. If a running service is currently writing to the file, it cannot be deleted. If so, the file disappears from the list in the SAP HANA studio and is hidden in the file system at the operating system level. As long as a service is still writing to the file, it still exists and uses disk space. Once the file reaches its maximum size, the system stops writing to it and creates a new trace file. The physical deletion of the file depends on how trace file rotation is configured.

If the trace files are open, you cannot delete the trace files. In this case, the contents of the files are cleared, but the file still exists and its size is reduced.

Download Files

To download a diagnosis file for offline analysis, right-click, and choose Download . You can select multiple files to download.

Collection and Download of Diagnosis Information

To help SAP Support analyze and diagnose problems with your system, you can collect a range of diagnosis information from your system into a zip file. You can trigger the collection of diagnosis information from the SAP HANA cockpit, and the command line.

In SAP HANA cockpit you can use Alerting & Diagnostics → Manage full system information dumps to analyze and diagnose problems with the SAP HANA database system. It collects all the important diagnosis information into a zip file, which you can download to a local PC. You can attach this downloaded zip file to an SAP Support Message.

The Manage full system information dumps includes the following features:

- Collect important diagnosis information
- Collect RTE (Runtime Environment) dump files

- Display an overview of the collected diagnosis information zip files
- Download and delete collected diagnosis information

When you start collecting diagnosis information, the system collects the relevant information by executing the Python script `fullSystemInfoDump.py`. You can execute this script within the Manage full system information dumps in SAP HANA cockpit or directly from the command line on the SAP HANA server as <sid>adm user.

Collecting Diagnosis Information Process

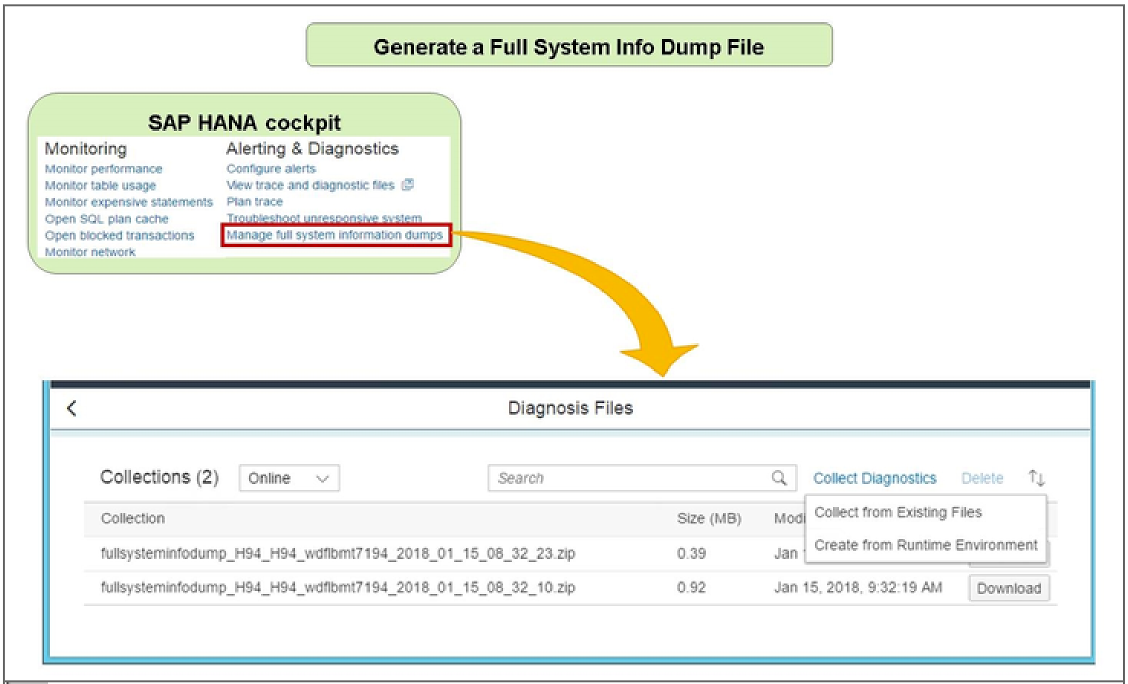

Figure 224: Offline Administration: Collecting Diagnosis Information

To collect diagnosis information using SAP HANA cockpit Alerting & Diagnostics → Manage full system information dumps , proceed as follows:

1. In the SAP HANA Cockpit, scroll down to the Alerting & Diagnostics area, and choose the Manage full system information dumps link.
2. Choose the Collect Diagnostics button, and choose Collect from Existing Files or Create from Runtime Environment .
3. In the dialog box, specify the required information, and choose Start .
4. When the system has finished collecting the relevant information, the zip file appears in the diagnosis files list.

Diagnosis Information

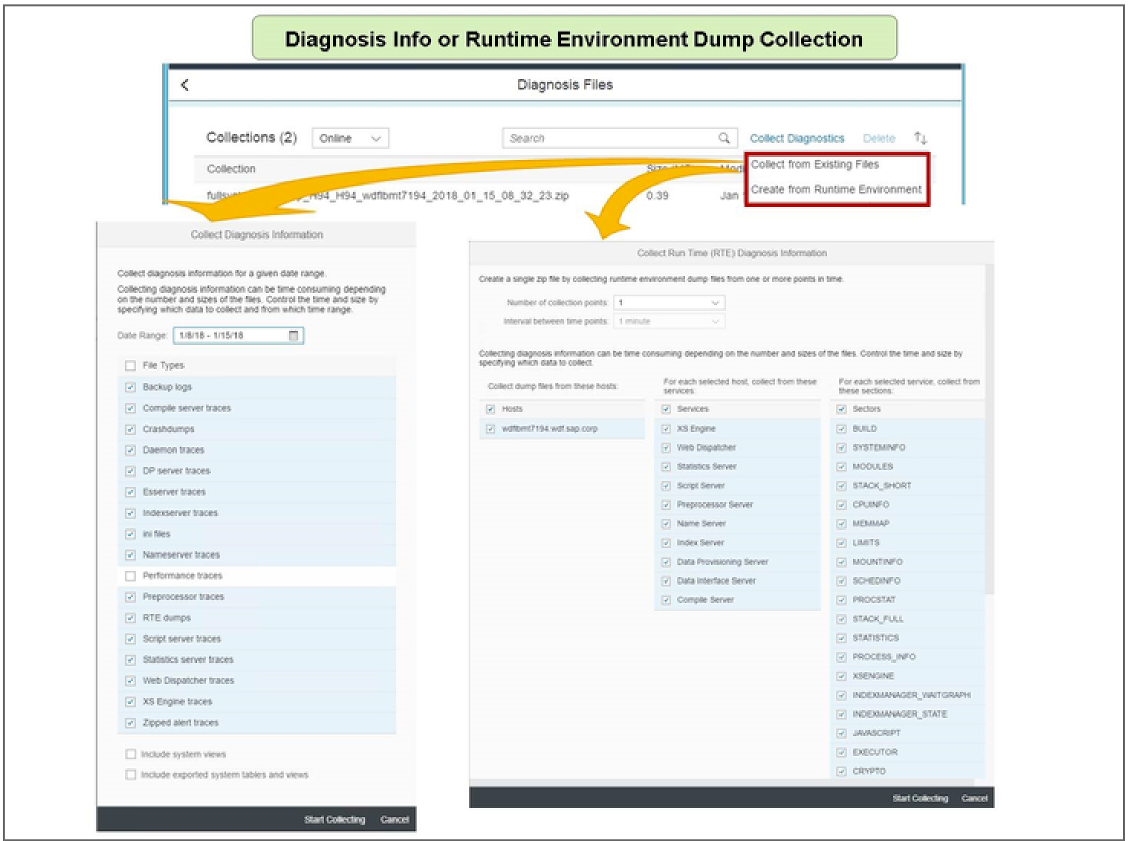

Figure 225: Offline Administration: Collect Diagnosis Information

When collecting Diagnosis Information, you can specify the date range. Make sure that you specify the date range so that it includes the time when the problem occurred.

For the Runtime Environment Dump File collection, specify the hosts, number of sets and time interval to start the collection. Be aware that collecting Runtime Environment information is a time-consuming operation.

The Python script can collect more information when the SAP HANA database is online because it can connect to the database and retrieve monitoring information from the SAP HANA monitor view. These views are not available when the database is offline.

RTE Dump File Collection

For each service, the RTE dump file contains information about System, CPU, Memory, Threads, Loaded Modules, and so on. A file named `<service name>_<host name>_<port>_runtimedump_<date>.trc` is added to the full system dump zip. These files are stored unabridged in the full System Dump file.

Collecting Diagnosis Information Using the Command Line

The `fullSystemInfoDump.py` script is part of the server installation and can be run from the command line. It is located in the directory `$DIR_INSTANCE/exe/python_support`.

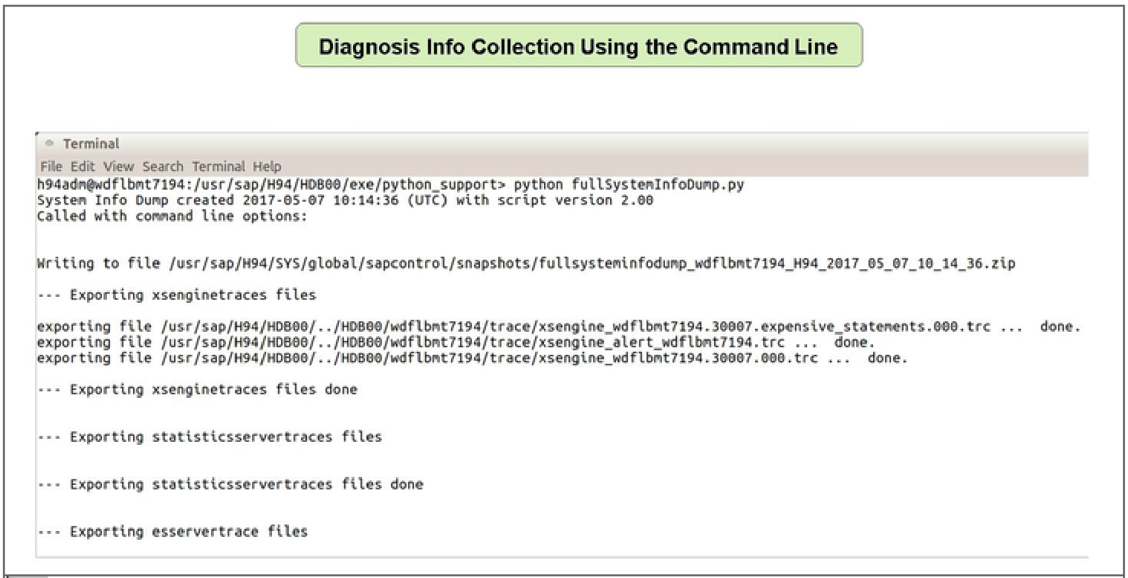

Figure 226: Collecting Diagnosis Information Using the Command Line

Start the script from its location with the command: `python fullSystemInfoDump.py`. Make sure you run the script as <sid>adm user.

The Python script collects the important trace file information and adds a copy to the dump file under the same name as the trace file. For storage reasons, only the trace files from the last seven days are collected unabridged. From older trace files, only the most recent 10,000 lines are collected.

Configuration Files Collection

All configuration .ini files are collected and are stored in the full system dump file using the original .ini file name.

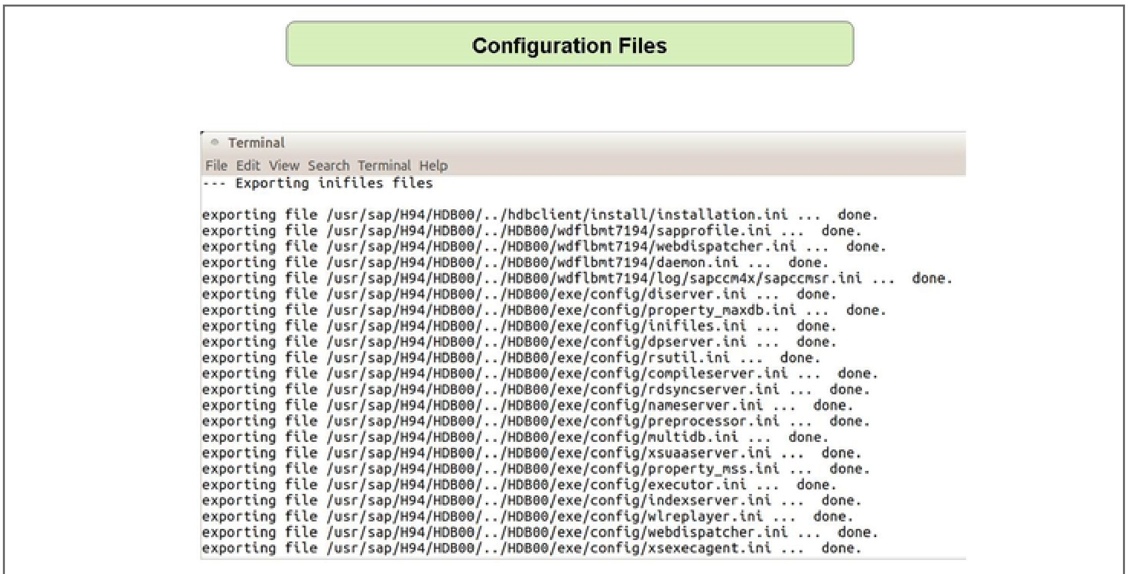

Figure 227: Configuration Files Collection

The following backup log files are collected unabridged:

- `$DIR_INSTANCE/<SAPLOCALHOST/trace/backup.log`

- `$DIR_INSTANCE/<SAPLOCALHOST/trace/backint.log`

Additional Information Collected if SQL Connection Is Available

All rows of the following system tables and monitoring views are exported to a CSV file with the name of the table:

- SYS.M_INIFILE_CONNECTIONS with CONNECTION_ID > 0
- SYS.M_DATABASE_HISTORY
- SYS.M_INIFILE_CONTENTS
- SYS.M_LANDSCAPE_HOST_CONFIGURATION
- SYS.M_SERVICE_STATISTICS
- SYS.M_SERVICE_THREADS
- SYS.M_SYSTEM_OVERVIEW
- SYS.M_TABLE_LOCATIONS
- SYS.M_TABLE_LOCKS
- SYS.M_TABLE_TRANSACTIONS
- _SYS_STATISTICS.STATISTICS_ALERT_INFORMATION
- _SYS_STATISTICS.STATISTICS_ALERT_LAST_CHECK_INFORMATION
- _SYS_STATISTICS.STATISTICS_ALERTS
- _SYS_STATISTICS.STATISTICS_INTERVAL_INFORMATION
- _SYS_STATISTICS.STATISTICS_LASTVALUES
- _SYS_STATISTICS.STATISTICS_STATE
- _SYS_STATISTICS.STATISTICS_VERSION

Note:
The first 2,000 rows of all remaining tables in the _SYS_STATISTICS schema are exported as ordered by column SNAPSHOT_ID.

Additional Information Collected if SQL Connection is Unavailable

All available topology information is exported to a file named `topology.txt`. It contains information about the host topology in a tree-like structure. The keys are grouped using brackets, while the corresponding values are referenced by the ==> symbol.

The following figure shows an example of the content of the topology.txt file.

```
[]
 ['host']
  ['host', 'ld8521']
   ['host', 'ld8521', 'role']
     ==> worker
   ['host', 'ld8521', 'group']
     ==> default
   ['host', 'ld8521', 'nameserver']
    ['host', 'ld8521', 'nameserver', '30501']
     ['host', 'ld8521', 'nameserver', '30501', 'activated_at']
       ==> 2011-08-09 16:44:02.684
     ['host', 'ld8521', 'nameserver', '30501', 'active']
       ==> no
     ['host', 'ld8521', 'nameserver', '30501', 'info']
       ['host', 'ld8521', 'nameserver', '30501', 'info', 'cpu_manufacturer']
       ==> GenuineIntel
       ['host', 'ld8521', 'nameserver', '30501', 'info', 'topology_mem_type']
       ==> shared
       ['host', 'ld8521', 'nameserver', '30501', 'info', 'sap_retrieval_path_devid']
       ==> 29
       ['host', 'ld8521', 'nameserver', '30501', 'info', 'build_time']
       ==> 2011-07-26 17:15:05
       ['host', 'ld8521', 'nameserver', '30501', 'info', 'net_realhostname']
       ==> -
       ['host', 'ld8521', 'nameserver', '30501', 'info', 'build_branch']
       ==> orange_COR
       ['host', 'ld8521', 'nameserver', '30501', 'info', 'mem_swap']
       ==> 34359730176
       ['host', 'ld8521', 'nameserver', '30501', 'info', 'mem_phys']
```

Figure 228: Contents of the Topology File

Troubleshooting an Unresponsive System

With the SAP HANA cockpit for Offline Administration, you can also analyze an unresponsive SAP HANA system. The information is collected using a Python script that is run by the sapstartsrv. You must provide the credentials of the operating system user <sid>adm.

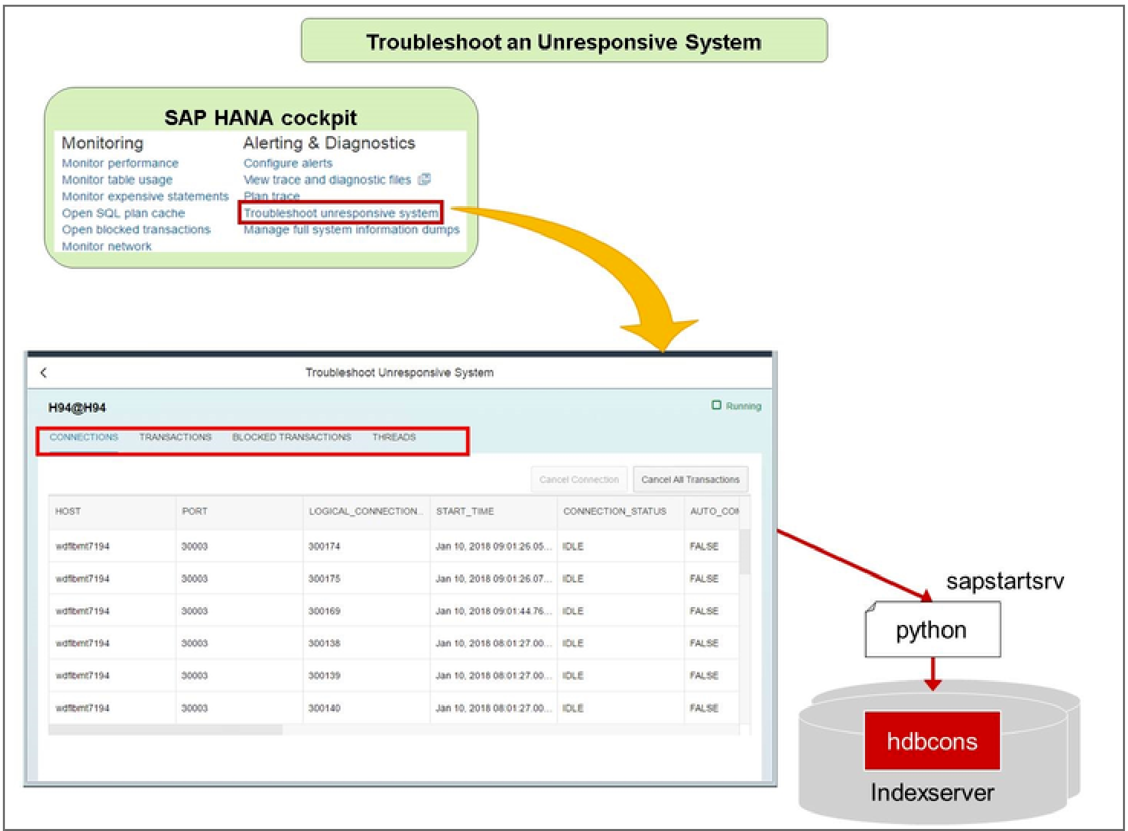

Figure 229: Troubleshoot an Unresponsive System

When a system cannot be reached by SQL or if it is experiencing major performance issues because of high load or blocking situations, you can instead collect information using the connection of the SAP start service (sapstartsrv).

The operational status of all connections, transactions, blocked transactions, and threads in the system is collected. If necessary, you can cancel individual connections and transactions, or even cancel all transactions.

LESSON SUMMARY

You should now be able to:

- Work with diagnosis information

Unit 9
Lesson 6

Using the SQL Console

LESSON OVERVIEW

The lesson briefly describes the following topics:

- Executing SQL statements in the SAP HANA studio
- Query analysis features
- Plan Visualizer, which creates a graphical representation of the query

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Use the SQL console

Execution of SQL Statements in SAP HANA Cockpit

Some tasks might require you to work with SQL statements. For example, certain administration tasks can only be performed using SQL. In the Database Explorer, you can enter, execute, and analyze SQL statements.

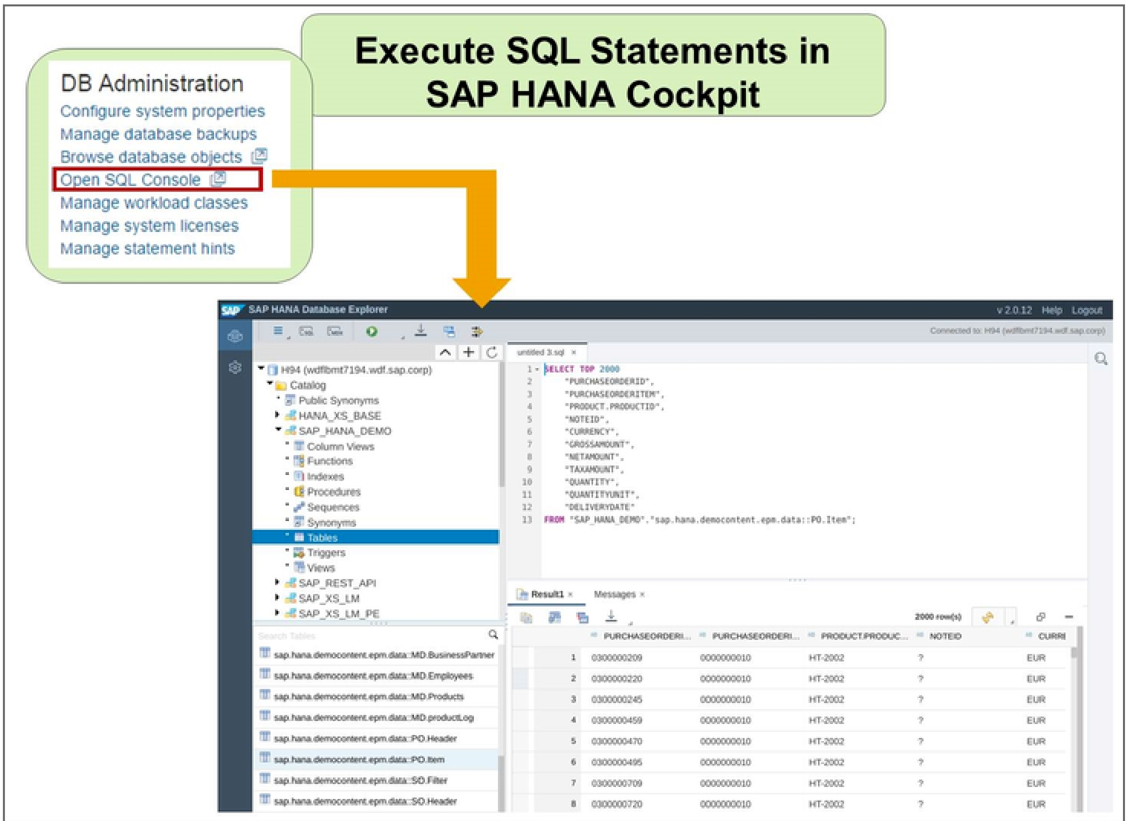

Figure 230: Executing SQL Statements in the SAP HANA Cockpit

To open the SQL Console, in SAP HANA cockpit, scroll down to the DB Adminstration area and choose the Open SQL Console link. This opens the Database Explorer where you can add your database.

Add a Database to the Database Explorer

Initially, the Database Explorer view is empty. However, you can add a database by clicking on the plus sign (+). In the next dialog screen, select the database you want to add and choose OK.

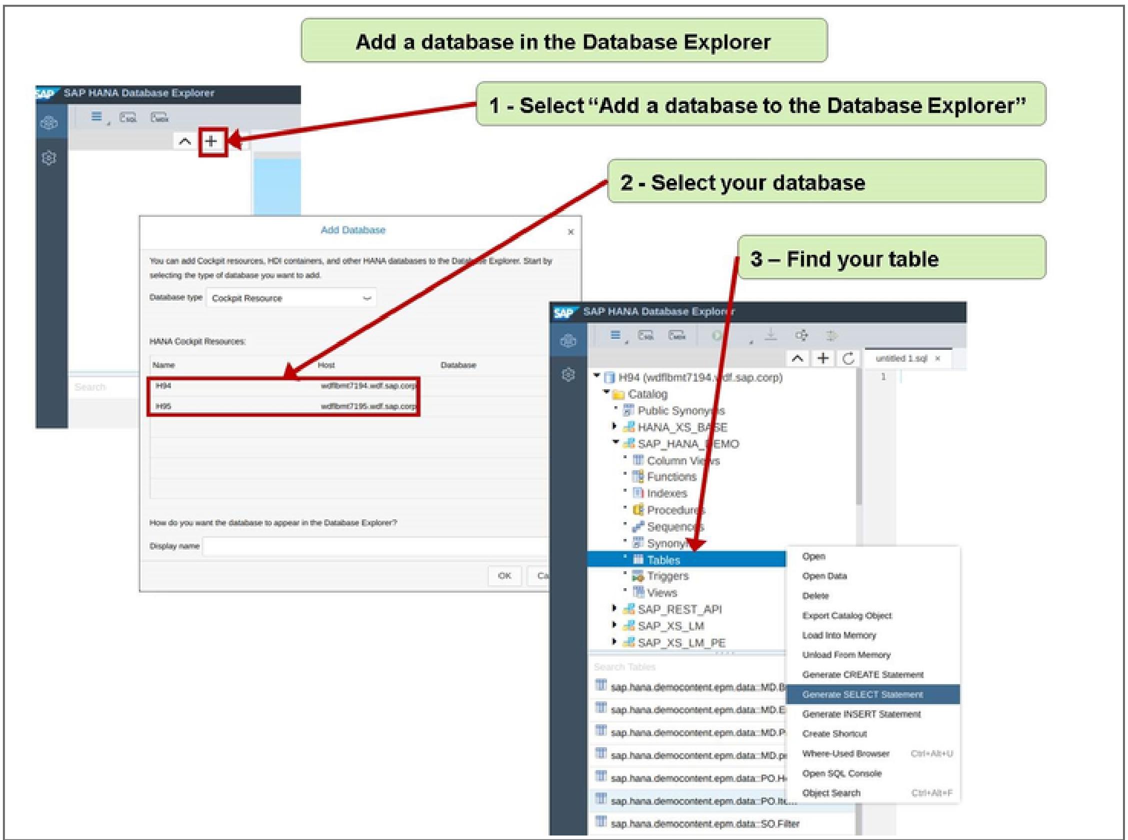

Figure 231: Add a Database to the Database Explorer

Now you can explore the database by opening the Catalog folder and navigating to the schema where the table is located. When you right-click a table, the context menu opens. This menu contains many options for analyzing the selected table.

Monitor Expensive Statements Application

With the Monitor Expensive Statements application, you can activate and view the Expensive Statement Trace. This information is useful for investigating which specific SQL statement is causing the high system load.

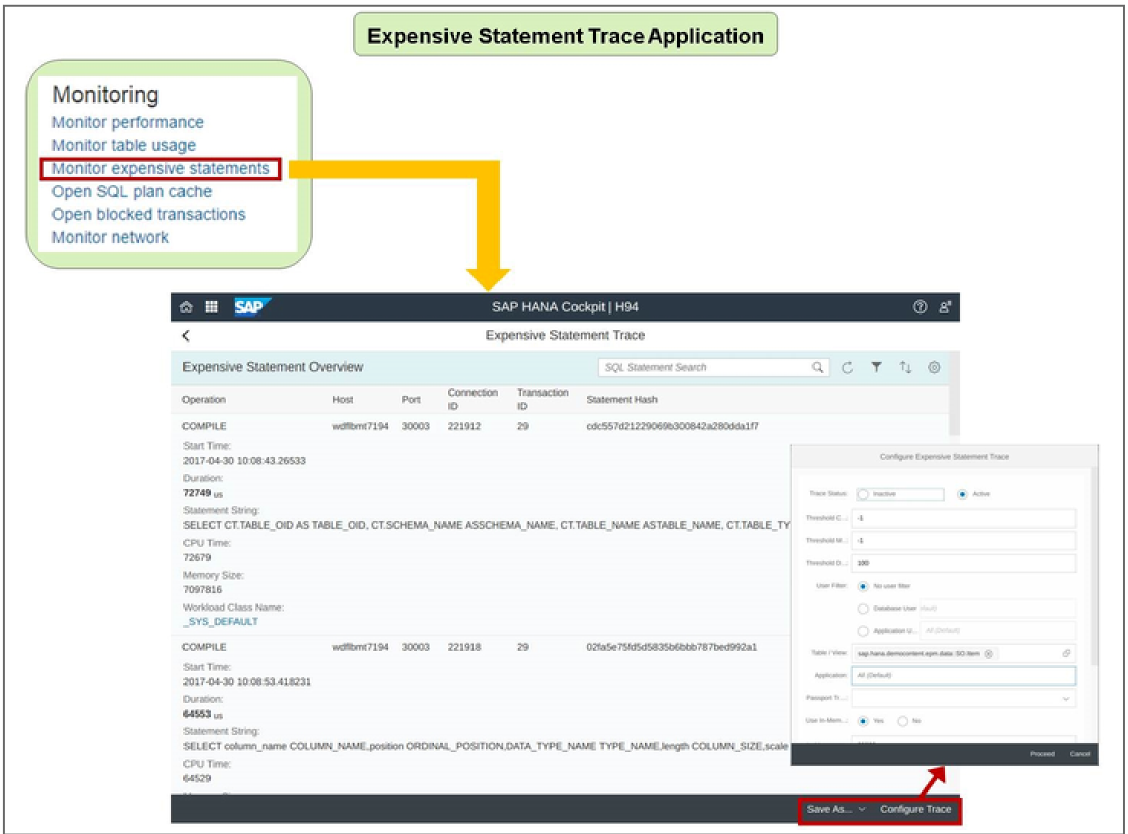

Figure 232: Monitor Expensive Statements Application

When you choose the Monitor expensive statements link in the Monitoring area, the Expensive Statement Trace is active, and you get information on the traced expensive statements. If no trace information is shown, you can activate the Expensive Statement Trace using the Configure Trace button.

Additional Trace Attributes

On the Configure Expensive Statement Trace dialog screen, you can activate or deactivate the trace. To create a more accurate trace, you can specify several additional trace attributes such as the following:

- Threshold values on CPU time, memory, and duration
- A user filter on the database or application user
- A filter on a specific table, view, or application
- Define a passport trace level
- The target location of the trace to be in-memory or file system
- The number of trace records kept in-memory and the trace flush interval

Use the Expensive Statement Trace to analyze SQL statements, but in normal production operations, check that it is disabled by default.

SQL Plan Cache Application

With the SQL Plan Cache application, you can activate and view the SQL Plan Cache kept by the SAP HANA database system.

When a SQL statement is executed for the first time, an execution plan is generated. The process of generating an execution plan is an expensive procedure. Therefore, to improve the performance of generating execution plans, the SAP HANA database stores generated execution for later reuse.

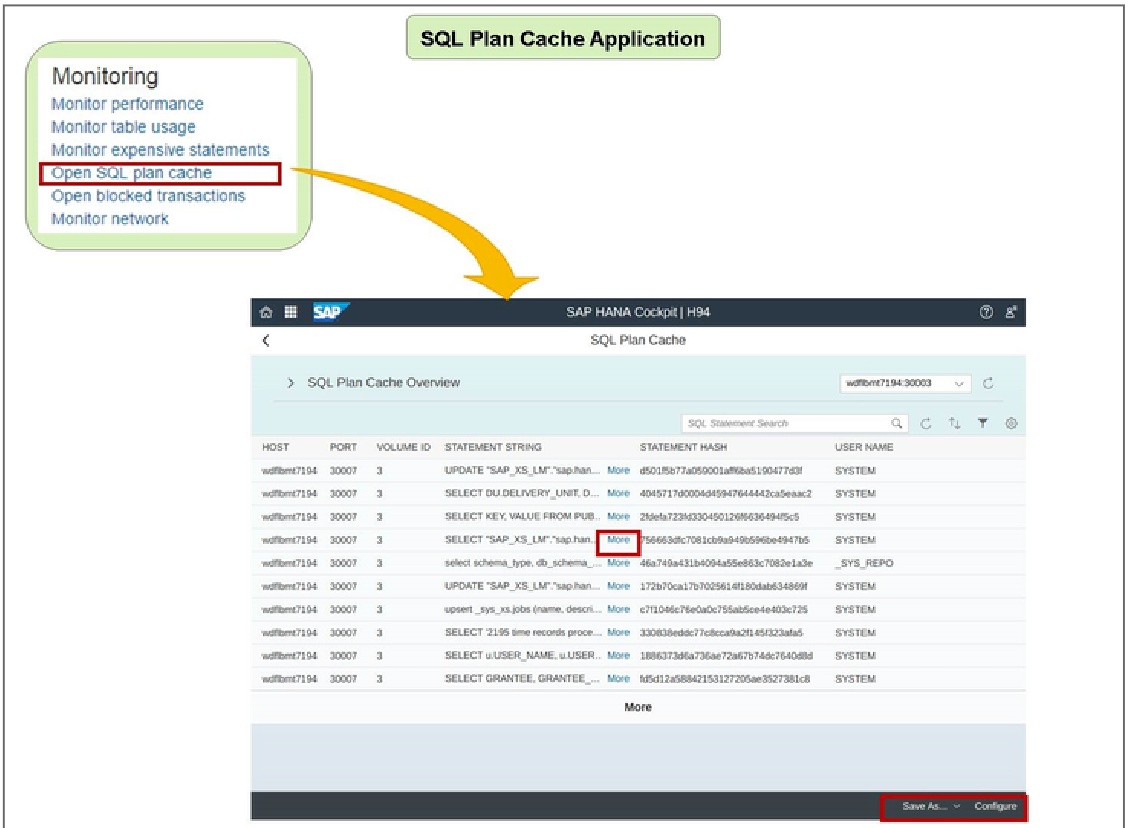

Figure 233: SQL Plan Cache Application

When you select the Open SQL plan cache link in the Monitoring area, you get detailed information on the cached SQL Plans. To see the full SQL statement that is cached, choose the More link. You can add additional columns with measured KPIs by choosing Configure Columns .

To get an summarized overview of the SQL plan cache, choose the SQL Plan Cache Overview link. The overview gives information on KPIs like cache size, use count, and total execution time.

With the Configure button, you can enable the collection of SQL Plan Cache Statistics. This allows you to analyze SQL Plan Cache performance problems.

LESSON SUMMARY

You should now be able to:

- Use the SQL console

Unit 9
Lesson 7

Performing SAP HANA Table Administration

LESSON OVERVIEW

Table administration is an important task for SAP HANA administrators.

This lesson covers details of table definition and partitioning, and explains various administrative tasks in this area.

Business Example

You are an administrator, and you need to create tables, optimize partitioning, and perform administrative tasks in this context.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Perform SAP HANA table administration

Column-Based and Row-Based Storage

The SAP HANA database supports both row-based and column-based storage. However, it is optimized for column storage. When you create a table, choose in advance whether it is stored by row or column.

When to Use Column Store

- Calculations on a small number of columns
- Table is searched based on the values of a few columns
- Table has large number of columns
- Table has large number of rows and columnar operations are required
- High compression rates shall be achieved

Tables that are organized in columns are read-optimized. They have better compression rates than tables organized in rows. Furthermore, some features of the SAP HANA database, such as partitioning, are available only for column tables. Column-based storage is used for large tables with bulk updates. However, update and insert performance is better on row tables. Row-based storage is used for small tables with frequent, single updates.

> Note:
> You can join row tables with column tables in the SAP HANA database. However, it is more efficient to join tables of the same storage type.

When to Use Row Store

When to Use Row Store

- Processing single records at one time / many selects and updates
- Accessing complete records
- Columns contain mainly distinct values
- No aggregations or fast search required
- Small number of rows

> Hint:
> You can change an existing table from one storage type to the other (ALTER TABLE ALTER TYPE).

Table Creation

To load data into the SAP HANA database, you need to create tables. Tables can be kept in row storage or column storage, depending on the use case.

> Note:
> To create a table, you must be authorized to create objects in the selected schema.

Tables can be created using SQL or the SAP HANA Studio interface.

Sample SQL command for creating a column table:

```
CREATE COLUMN TABLE "TRAINING"."CUSTOMER"
("CUSTOMER_ID" NVARCHAR(10) DEFAULT '' NOT NULL ,
 "CUSTOMER_TEXT" VARCHAR(30) DEFAULT '',
 "COUNTRY" VARCHAR(5) DEFAULT '',
 "REGION" VARCHAR(5) DEFAULT '',
 "LOCATION" VARCHAR(30) DEFAULT '',
 "DISCOUNT_GROUP_ID" VARCHAR(4),
 PRIMARY KEY ("CUSTOMER_ID"))
```

Details and Options: See also SAP HANA SQL Reference

Figure 234: Sample SQL Command for Creating a Column Table

The figure, Sample SQL Command for Creating a Column Table, shows a sample SQL command for creating a column table. The column table CUSTOMER is created within database schema TRAINING. It contains five different columns, of which CUSTOMER_ID is the primary key.

Hint:
For details and options, see SAP HANA SQL Reference.

Using SAP HANA Studio

Alternatively, you can create a table directly within SAP HANA Studio, as shown in the figure, Using SAP HANA Studio.

From the schemas context menu, select *New Table*:

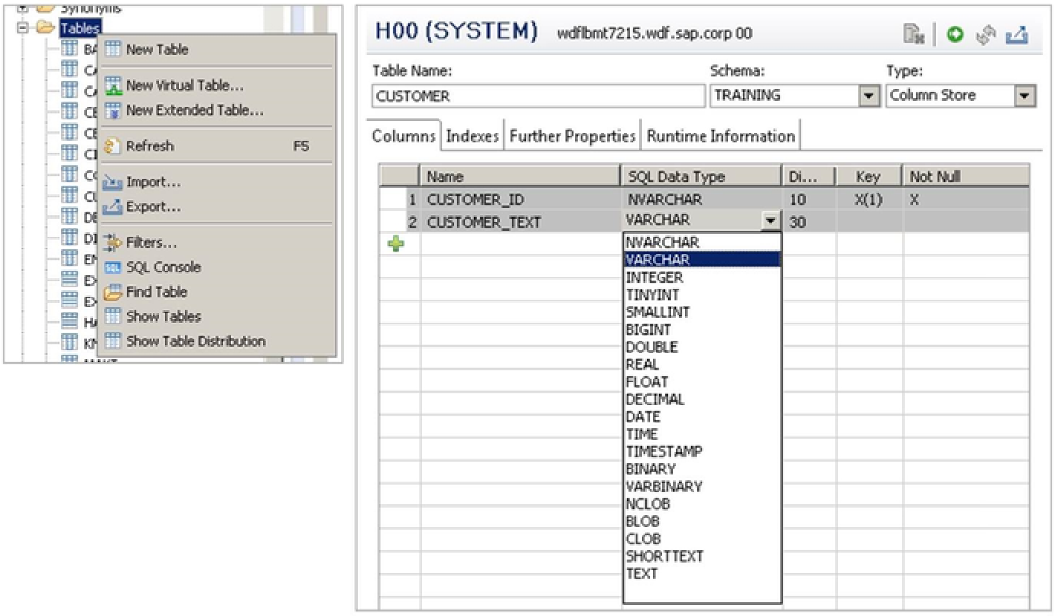

Figure 235: Using SAP HANA Studio

Note:
Note: You do not have to define an index with the table; you can create it at any time.

To Create Tables Using SAP HANA Studio

1. In the Systems view, open the catalog schema in which you want to create the new table.
2. In the context menu of the schema in which you want to create the table, choose New Table .
3. Enter the table name and table type (column store or row store).
4. Define the columns of your table (name and properties).
5. If required, you can add indexes.
6. Choose Create Table .

Display of Table Definition and Content

The Database Explorer in SAP HANA cockpit and SAP HANA Studio offer multiple options for displaying table definition and content.

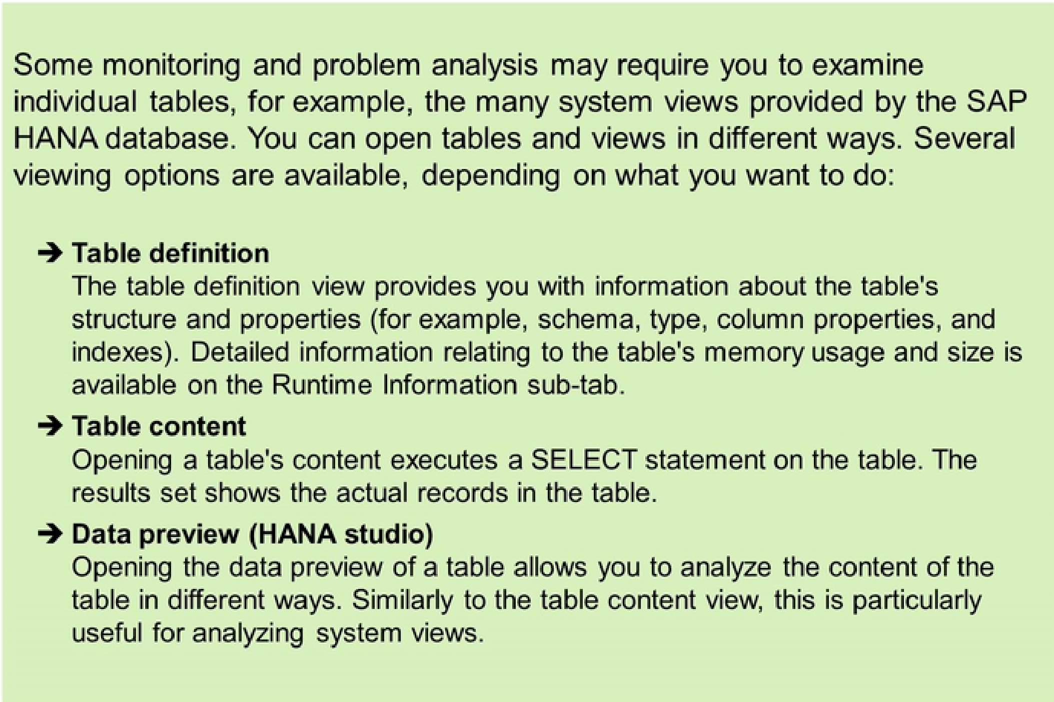

Figure 236: Displaying Table Definition and Content

Displaying Table Definition

To display catalog object definitions and change existing catalog objects, you require specific privileges. If these are not granted to you, the error **`Insufficient privilege`** displays.

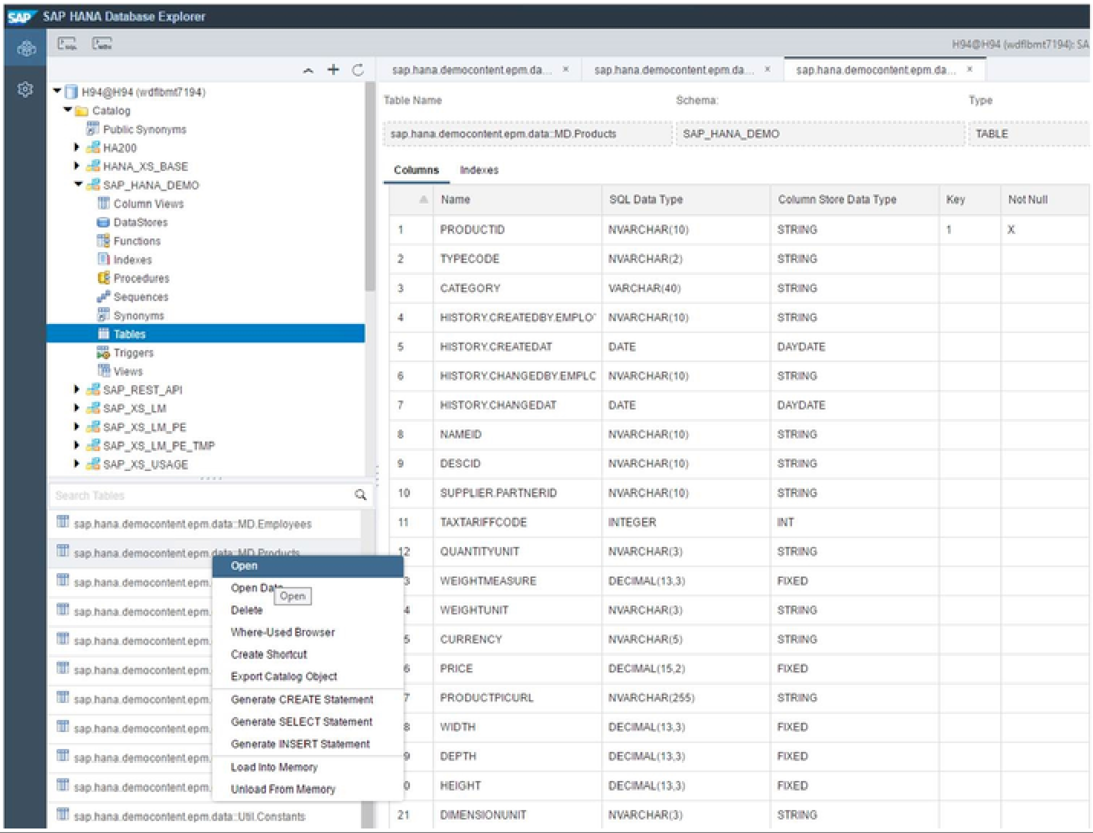

Figure 237: Displaying Table Definition in the Database Explorer

To open the table editor in SAP HANA studio, choose Open Definition in the context menu of a specific table.

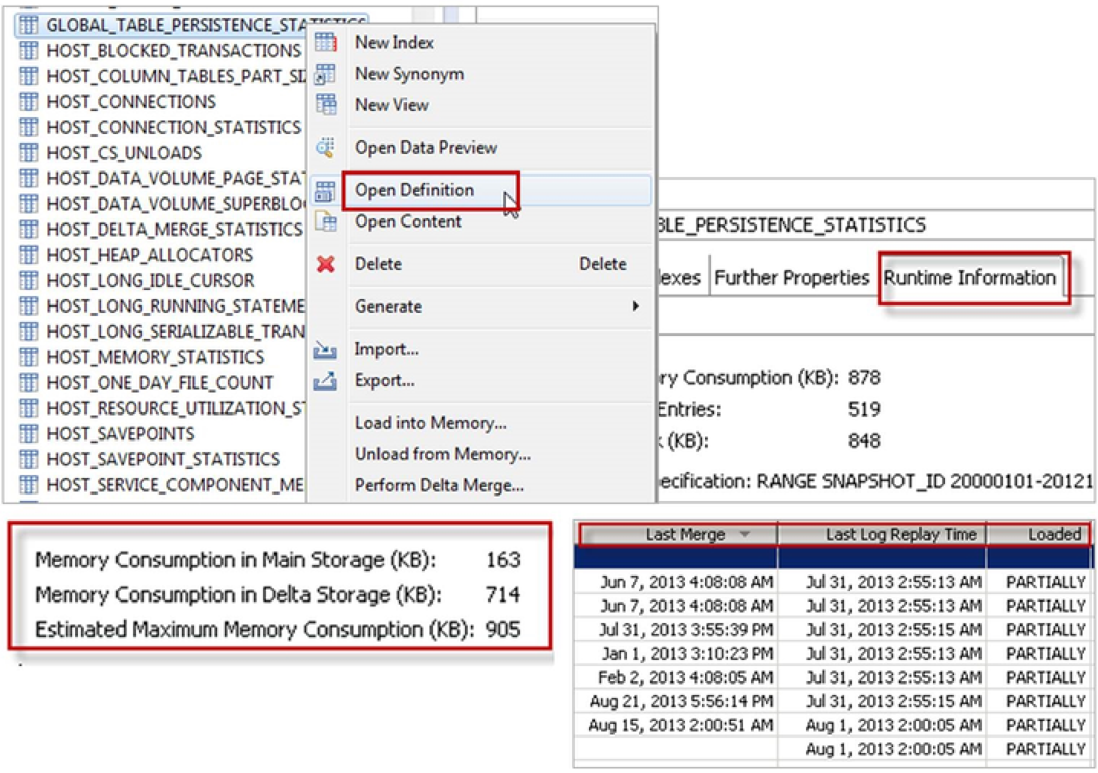

Figure 238: Displaying Table Definition in SAP HANA Studio

The table definition displays columns and indexes in the Runtime Information column, details about the memory and disk consumption, and information about the compression of individual columns.

Note:
By default, to open the definition of a table, double-click the table in the Systems view. You can configure this setting in the preferences of the SAP HANA studio.

Displaying Table Content

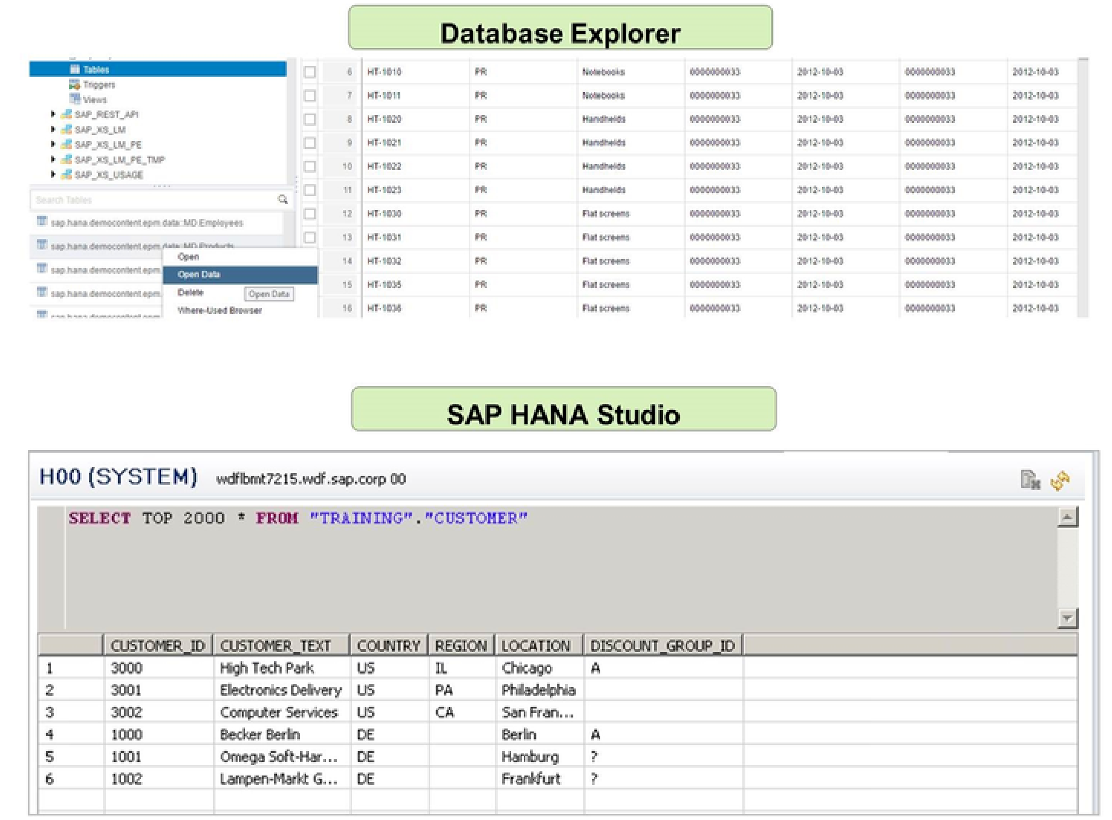

Figure 239: Displaying Table Content

Displaying the table content can be useful, for example, if you want to view the content of a system view to help you understand what is happening in the database.

Note:
By default, only the first 1,000 rows are displayed. To change this setting in the preferences of the SAP HANA studio, choose SAP HANA→ Runtime → Catalog .

Data Preview

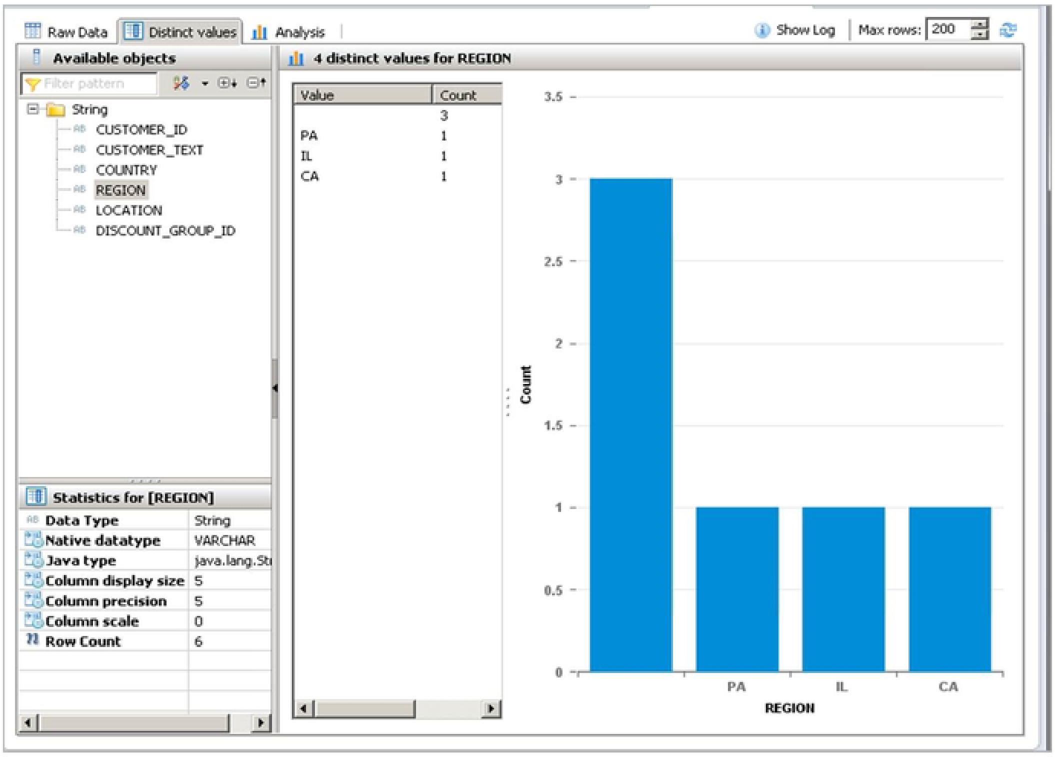

Figure 240: Data Preview in SAP HANA Studio

Table Partitioning and Distribution

With the partitioning feature of the SAP HANA database, you can split column-store tables horizontally into disjunctive subtables or partitions. Thus, large tables can be broken down into smaller, more manageable parts.

> Hint:
> Partitioning is typically used in distributed systems, but it may also be useful for single-host systems.

Additional DDL Statements for Partitioning in the SAP HANA Database

- Create table partitions
- Repartition tables
- Merge partitions to one table
- Add/delete partitions
- Move partitions to other hosts

When a table is partitioned, the split ensures that each partition contains a different set of rows of the table. There are several different ways of assigning the rows to the partitions of a table, for example, hash partitioning, partitioning by range, or value.

Advantages of Partitioning

- Load balancing
- Overcoming size-limitation of column store tables
- Parallelization
- Partition pruning
- Improved performance of delta merge operation
- Explicit partition handling

Partitioning has the following advantages:

- **Load balancing in a distributed system**

  Individual partitions can be distributed across multiple hosts. Therefore, a query on a table is not processed by a single server but by all the servers that host partitions.

- **Overcoming the size limitation of column-store tables**

  A nonpartitioned table cannot store more than 2 billion rows. You can overcome this limit by distributing the rows across several partitions. Each partition must not contain more than 2 billion rows.

- **Parallelization**

  Operations can be parallelized by using several execution threads for each table.

- **Partition pruning**

  Queries are analyzed to determine whether or not they match the given partitioning specification of a table. If a match is found, you can determine the actual partitions that hold the data being queried. This method reduces the overall load on the system and improves the response time.

- **Improved performance of the delta merge operation**

  The performance of the delta merge operation depends on the size of the main index. If data is only modified on some partitions, fewer partitions need to be delta-merged, and therefore performance is better.

- **Explicit partition handling**

  Applications can actively control partitions, for example, by adding partitions to store the data for an upcoming month.

Single-Level Partitioning: Supported Specifications

When a table is partitioned, its rows are distributed to partitions according to different criteria, known as partitioning specifications. The SAP HANA database supports several single-level partitioning specifications.

- Hash
- Range
- Round Robin

The SAP HANA database supports the following single-level partitioning specifications:

- **Hash Partitioning**

  Hash partitioning distributes rows to partitions equally for load balancing. It expands the 2 billion row limitation. The number of the assigned partition is computed by applying a hash function to the value of a specified column. Hash partitioning does not require an in-depth knowledge of the actual content of the table.

- **Range Partitioning**

  Range partitioning creates dedicated partitions for certain values or certain value ranges in a table. This requires an in-depth knowledge of the values that are used or that are valid for the chosen partitioning column. For example, a range partitioning scheme creates one partition for each calendar month.

- **Round Robin Partitioning**

  Round-robin partitioning produces an equal distribution of rows to partitions. However, unlike hash partitioning, you do not have to specify partitioning columns. With round-robin partitioning, new rows are assigned to partitions on a rotation basis. The table must not have primary keys.

Hint:
For more information, see SAP HANA Administration Guide.

Note:
In addition to single-level partitioning, SAP HANA has various options for multilevel partitioning. Details are described in the SAP HANA Administration Guide.

Time Selection Partitioning: Aging

SAP HANA SPS7 includes a new feature called "Time Selection Partitioning (Aging)" .

The SAP HANA database offers a special time selection partitioning scheme, also called aging. Time selection or aging allows application data to be horizontally partitioned into different temperatures, like hot and cold.

Applications can use aging to separate hot (current) data from cold (old) data by using time selection partitioning to:

- Create partitions and re-partition
- Add partitions
- Allocate rows to partitions
- Set the scope of Data Manipulation Language (DML) and Data Query Language (DQL) statements

Figure 241: Time Selection Partitioning: Aging

Example for Creating a Hash-Partitioned Table Using SQL

Example for Creating a Hash-Partitioned Table Using SQL:

```
CREATE COLUMN TABLE MY_TABLE
(       a INT,
        b INT,
        c INT,
        PRIMARY KEY (a,b))
PARTITION BY HASH (a, b) PARTITIONS 3
```

Figure 242: Example for Creating a Hash-Partitioned Table Using SQL

In the figure, Example for Creating a Hash-Partitioned Table Using SQL, three partitions are created on columns a and b of the table MY_TABLE.

Table Distribution Editor

Alternatively, you can use the Table Distribution Editor in SAP HANA Studio.

To support the analysis and monitoring of performance issues in a distributed SAP HANA system, a table distribution editor is available in which you can see how tables and table partitions are distributed across the hosts. Detailed information about tables/partitions (for example, memory usage and size) is also available.

You access the Table Distribution editor from the SAP HANA Systems view.

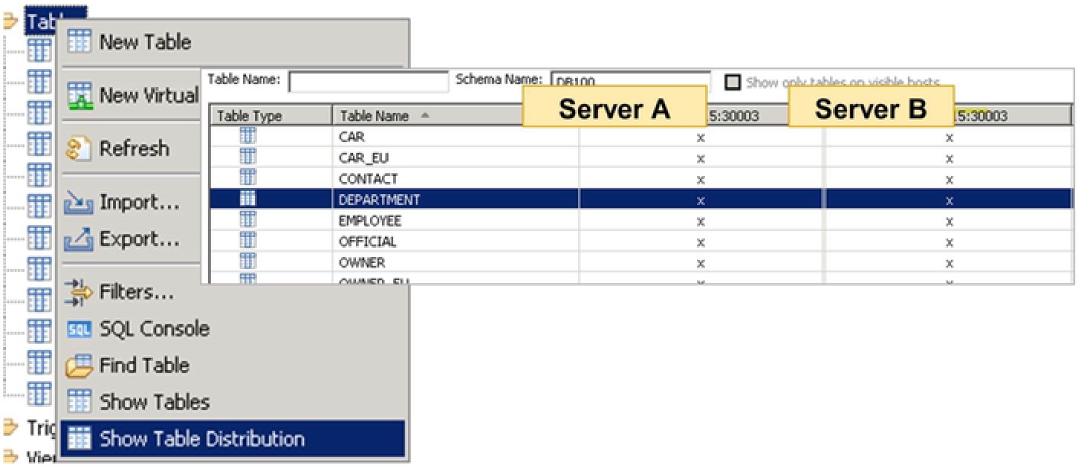

Figure 243: Table Distribution Editor in SAP HANA Studio

The Table Distribution editor outlines the distribution of tables in a distributed system. To open it, choose the context menu on the Console folder, or any schema or tables folder in the Navigator. Because of performance reasons, it only displays 1000 tables of the selected schema. To configure this number, choose Preferences → Administration Console → Common → Table Distribution Editor . If more tables exist in the selected schema, a message displays.

Table Distribution Editor: Additional Actions

If a table is distributed to several partitions, it displays the host that stores each of these partitions. Existing partitions can be moved to different hosts. You can move tables that are not partitioned to other hosts as well. However, you cannot split a table or change the partitioning using this view.

In the Table Distribution editor, you can perform the following additional actions:

- Move tables and partitions to other hosts in the system
- Partition non-partitioned tables using the supported partitioning specifications (hash, range, and round-robin)
- Change a partitioned table into a non-partitioned table by merging its partitions

➔ **Note:**

- Before moving tables or partitions, the system checks that the host has sufficient memory.

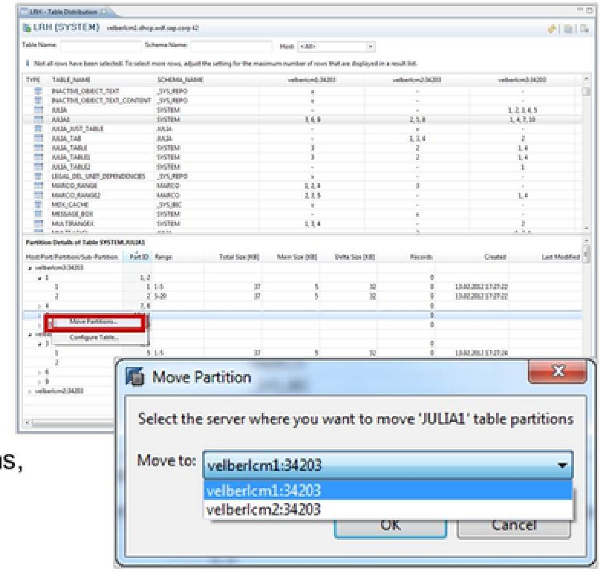

Figure 244: Table Distribution Editor: Additional Actions

SAP HANA Table Administration - Administrative Tasks

To ensure consistency for partitioned tables, execute checks and repair statements, if required.

You can call general and data consistency checks for partitioned tables to check, for example, that the partition specification, metadata, and topology are correct.

Two types exist:

➔ **General check: Consistency check**

```
CALL CHECK_TABLE_CONSISTENCY('CHECK_PARTITIONING',
'<schema>', '<table>')
```

➔ **Data check: General check plus check whether all rows are located in correct parts**

```
CALL CHECK_TABLE_CONSISTENCY('CHECK_PARTITIONING_DATA',
'<schema>', '<table>')
```

Repairing rows that are located in incorrect parts:

```
CALL CHECK_TABLE_CONSISTENCY('REPAIR_PARTITIONING_DATA',
'<schema>', '<table>')
```

Figure 245: Administrative Tasks: Checking Partitioning Consistency

Note:
The data checks can take a long time to run, depending on the data volume.

Administrative Tasks: Table Replication (Tuning Option)

Starting with SAP HANA SPS7, a new tuning option called is available. The option Table Replication allows you to replicate tables to multiple hosts.

In a scale-out system, tables may be replicated to multiple hosts. This is useful when slowly changing master data often has to be joined with tables or partitions of other tables that are located on multiple hosts, and you want to reduce network traffic.

Example for creating column store tables with replicas on all hosts:

```
CREATE COLUMN TABLE MY_TABLE (I INT PRIMARY KEY
        REPLICA AT ALL LOCATIONS
```

- For prerequisites and additional details, see also SAP HANA Administration Guide.

Figure 246: Administrative Tasks: Table Replication (Tuning Option)

Note:
Several aspects need to be considered. See the SAP HANA Administration Guide for details.

Administrative Tasks: Loading and Unloading Column Tables

Under normal circumstances, the SAP HANA database manages the loading and unloading of tables into and from memory independently, the aim being to keep all relevant data in memory. However, you can manually load and unload individual tables and table columns if necessary.

Options:

Loading and unloading tables using the Database Explorer

- Right click the table in *Tables* view
- In the context menu, choose *Load Into Memory* or *Unload from Memory*

Loading and unloading tables using SQL

- `LOAD <table_name>`
- `UNLOAD <table_name>`

Loading and unloading individual columns

- `LOAD <table_name> (<column_name>, …)`
- `UNLOAD <table_name> (<column_name>, …)`

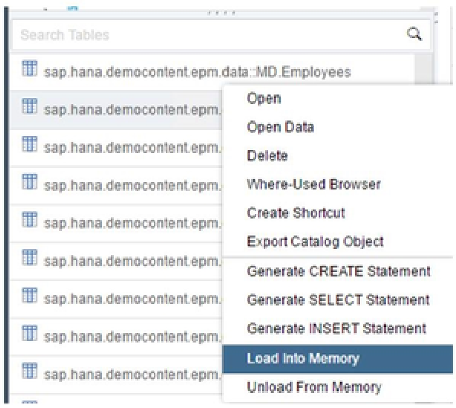

Figure 247: Administrative Tasks: Loading and Unloading Column Tables

Because the SAP HANA database manages the loading and unloading of tables automatically, you do not have to interfere with this process. However, if necessary, you can load and unload individual tables and table columns manually. For example, this occurs in the following cases:

- To measure, precisely, the total memory, or the amount of memory used by a particular table (load)
- To actively free up memory (unload)

> Hint:
> For more information about a table's current memory usage and load status, view its table definition.

Administrative Tasks: Performing Manual Delta Merge Operations

By default, SAP HANA controls the delta merge process automatically. However, it may be necessary or useful to trigger a merge operation manually in some situations. For example, this occurs in the following cases:

- An alert has been issued because a table exceeds the threshold for the maximum size of delta storage.
- You need to free up memory.

Delta merges can be triggered manually using SAP HANA Studio or SQL.

You can trigger the delta merge operation for a column table manually, for example, if you need to free up memory.

Options:

Menu command in SAP HANA Studio

- Right-click the table in *Systems* view
- In the context menu, choose *Perform Delta Merge*

SQL

- `MERGE DELTA OF '<table_name>'`

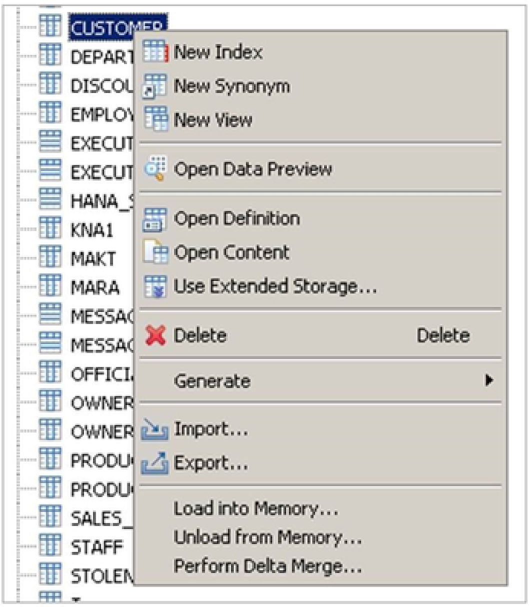

Figure 248: Administrative Tasks: Performing Manual Delta Merge Operations

Note:
Additional options exist. For more information, see the SAP HANA Administration Guide.

Hint:
Even though the delta merge operation moves data from the delta storage to the main storage, the size of the delta storage is not zero. This is because records written by open transactions are moved to the new delta storage while the delta merge operation takes place. Furthermore, even if the data containers of the delta storage are empty, they still need some in-memory space.

Load, unload, and merge are available in the context menu of a specific column store table. You can select multiple tables at once. That operation is then executed for all selected tables.

Administrative Tasks: Importing and Exporting Tables

You can easily export and import tables, as other catalog objects, back into another database, as shown in the figure Administrative Tasks: Importing and Exporting Tables.

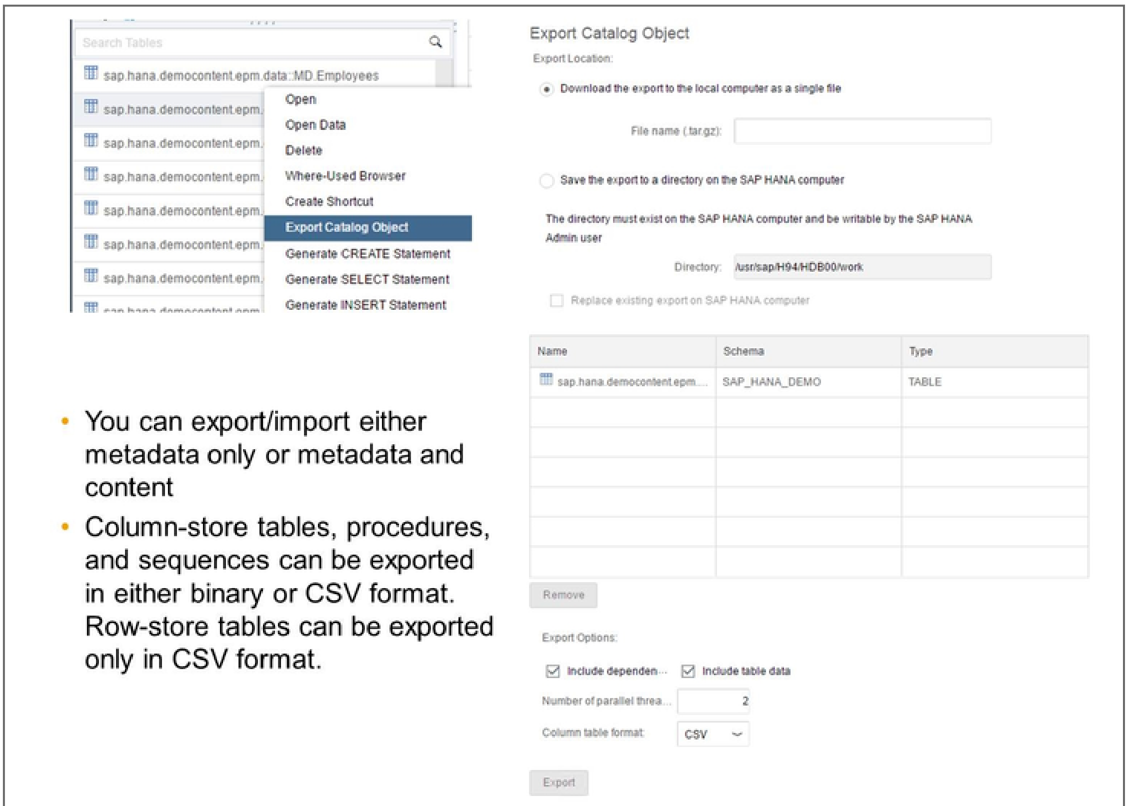

Figure 249: Administrative Tasks: Importing and Exporting Tables

Note:
The size of a .CSV format file can be large compared to the Binary file size.

By default, the exported data is stored on the database server. However, you can also export the data to the local client machine. Importing data creates the tables in the same schema as in the source system. If the table already exists, select the checkbox so that it can be overwritten. Otherwise, the import aborts with an error message.

LESSON SUMMARY

You should now be able to:

- Perform SAP HANA table administration

Unit 9
Lesson 8

Transporting Changes

LESSON OVERVIEW

This lesson provides an overview of the transporting options.

Business Example

SAP HANA lifecycle management covers the following two aspects:

- Platform lifecycle management for customizing and updating your SAP HANA platform
- Application lifecycle management for managing SAP HANA content products and transports

In this lesson, you will learn how to manage SAP HANA content, including modelling, change recording, transports, and installation.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Transport changes

SAP HANA Application Lifecycle Management: Overview

Application lifecycle management includes all the activities that you need to ensure that the software components you develop for SAP HANA are not only produced and shipped in a regulated way, but also meet the requirements laid out for the SAP HANA platform.

In SAP HANA, several objects can be developed to build standalone applications or to integrate with other products such as SAP systems.

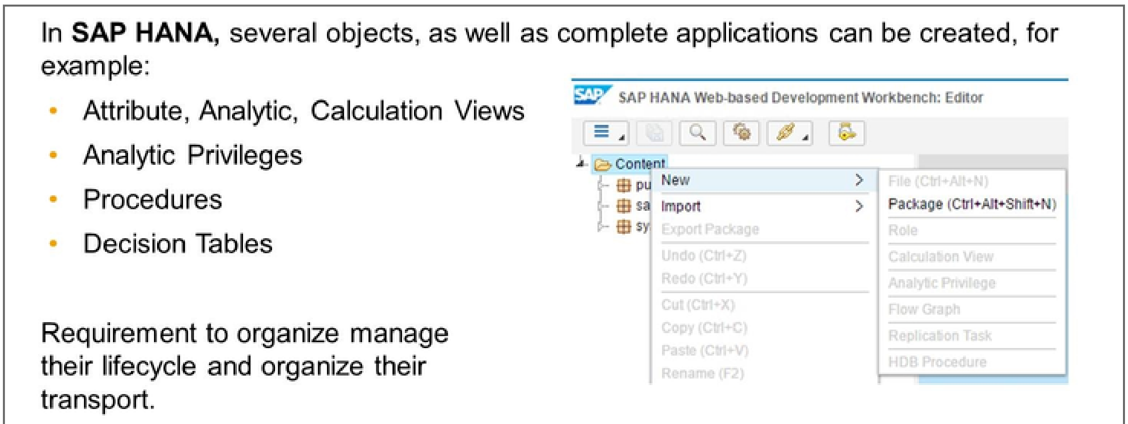

Figure 250: SAP HANA Transports: What Do We Need Them For?

SAP HANA Content

These objects are regarded as transportable content. What is defined as SAP HANA content is shown in the figure, What is Defined as SAP HANA Content?

SAP HANA content defined:

- **Not** part of the core SAP HANA database installation itself
- Is delivered by SAP as part of SAP HANA optimized solutions
- Is created in SAP HANA-based development projects (partner, customer)
- Sometimes called “objects” or “artifacts”

Content comprises all kinds of objects, for example:

- Schemas and table definitions
- Attribute views, analytic views, and calculation views
- Procedures and privileges
- SQLScript, JavaScript and HTML
- Roles and permissions

Figure 251: What is Defined as SAP HANA Content?

Design Time and Runtime Objects: SAP HANA Repository

It is important to distinguish between design time and run time objects. Design time objects are regarded as SAP HANA content.

The repository and lifecycle management of objects

- Native feature of SAP HANA providing “backend” functionality for content lifecycle management
- Used to manage various types of design time objects (Content)
- During deployment/activation, the design time objects become runtime objects (Catalog)

Key functions provided by the repository:

- Object versioning
- Namespace concept
- Support for server-based development

SAP HANA Studio Modeling Perspective

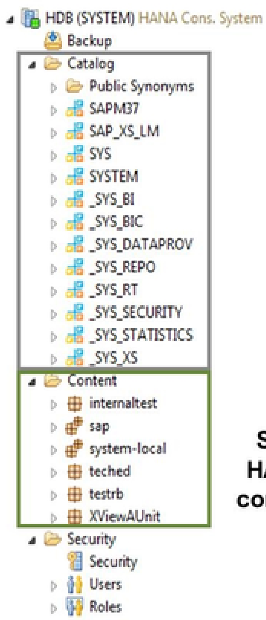

SAP HANA content

Figure 252: Design Time and Runtime Objects: SAP HANA Repository

To ensure consistency when transporting objects, ship objects that belong together at the same time.

Context: Packages

In the repository, SAP HANA objects that belong together are made up of packages, and packages can be assigned to a delivery unit.

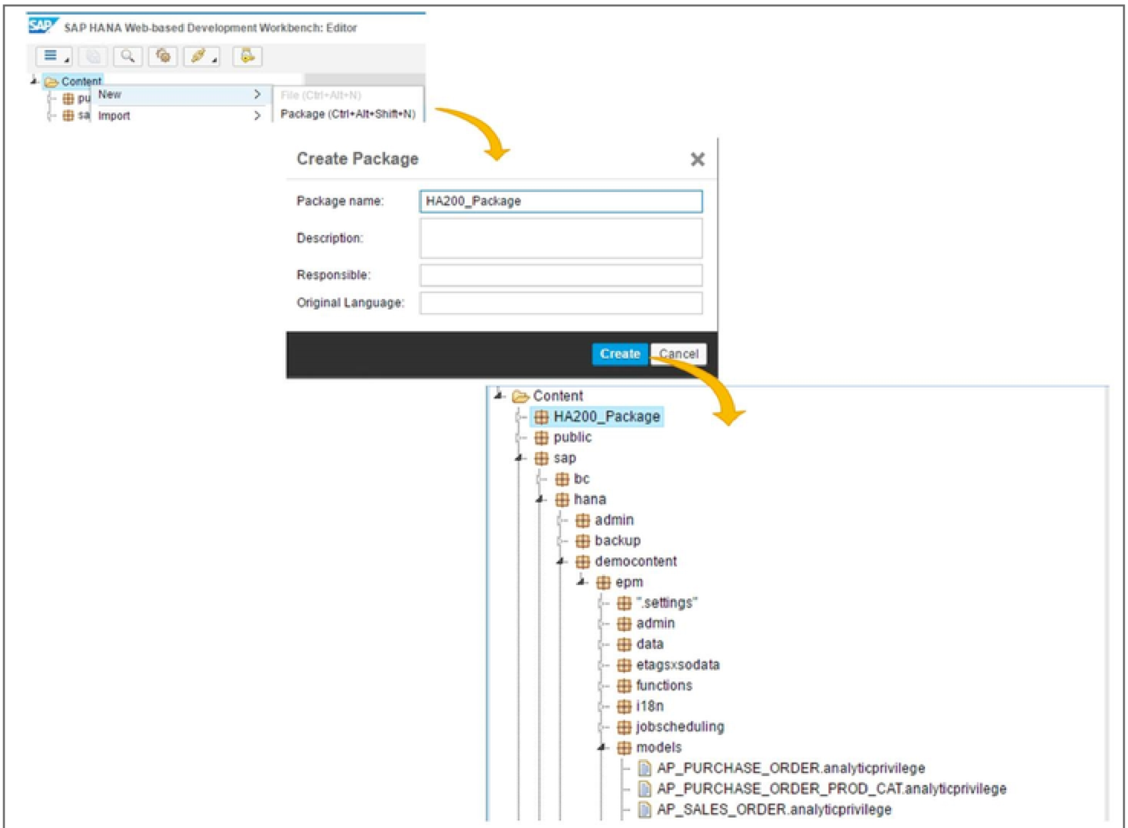

Figure 253: Context: Packages

All content delivered as part of the application that you develop for SAP HANA is stored in packages in the SAP HANA repository. The packages are arranged in a hierarchy that makes the process of maintaining the packages transparent and logical.

Packages enable you to group together the artifacts that you create and maintain for your applications. Be aware of the privileges that the application developers require to access (and perform operations on) the packages.

Context: Delivery Units

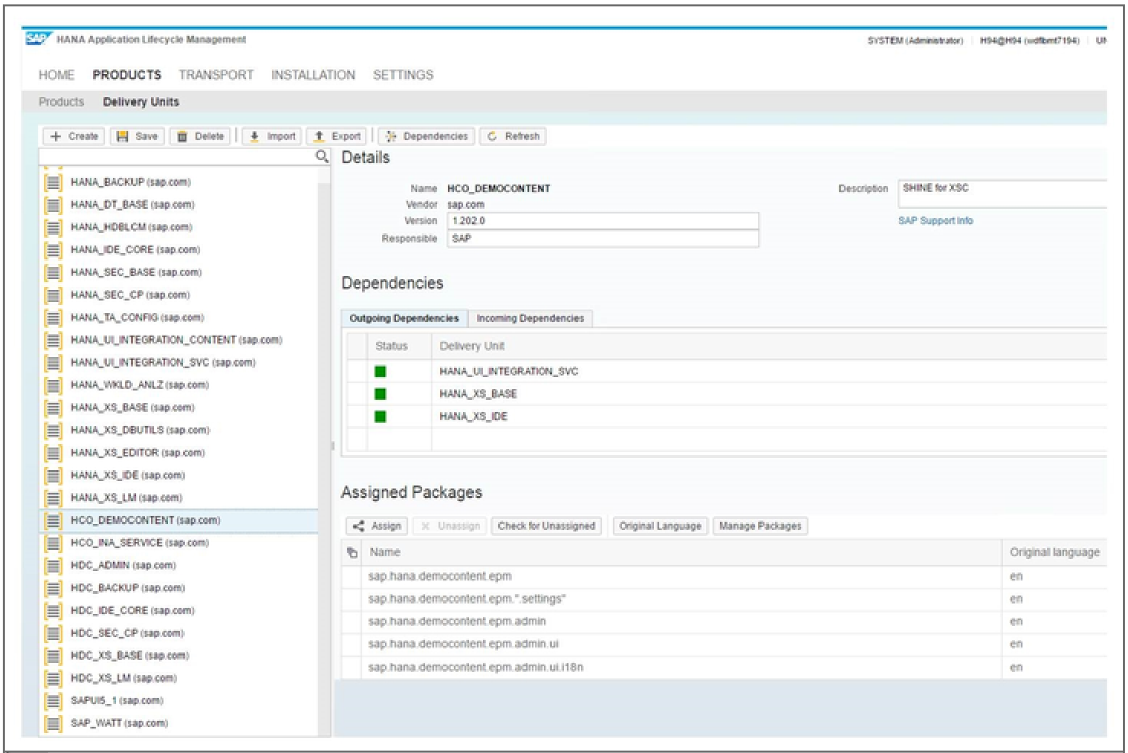

Figure 254: Context: Delivery Units

A delivery unit is a collection of packages that are transported together. Assign all the packages belonging to your application to the same delivery unit. This ensures that they are transported consistently together within your system landscape. Each delivery unit has a unique identity.

The identity of a delivery unit consists of two parts: a vendor name, and a delivery-unit name. The combined ID ensures that delivery units from different vendors are distinguished easily. It also ensures that they follow a pattern that SAP uses for various software components.

To create and manage delivery units, you first need to maintain the identity of the vendor. The delivery units are associated with the vendor, and the packages that make up the delivery unit are stored in the vendor’s namespace.

Delivery Unit

The delivery unit has the following features:

- Collection of packages to be transported together
- Helps to ensure consistent transports of all packages of one application
- Unique identity
  - Vendor name (compare content_vendor attribute)
  - Delivery-unit name

Overview: Products, Delivery Units, Packages, and Objects

Delivery Units are associated with a product instance. A product corresponds to an application, which could be an SAP-delivered application, a partner application, or customer application developed on a project basis.

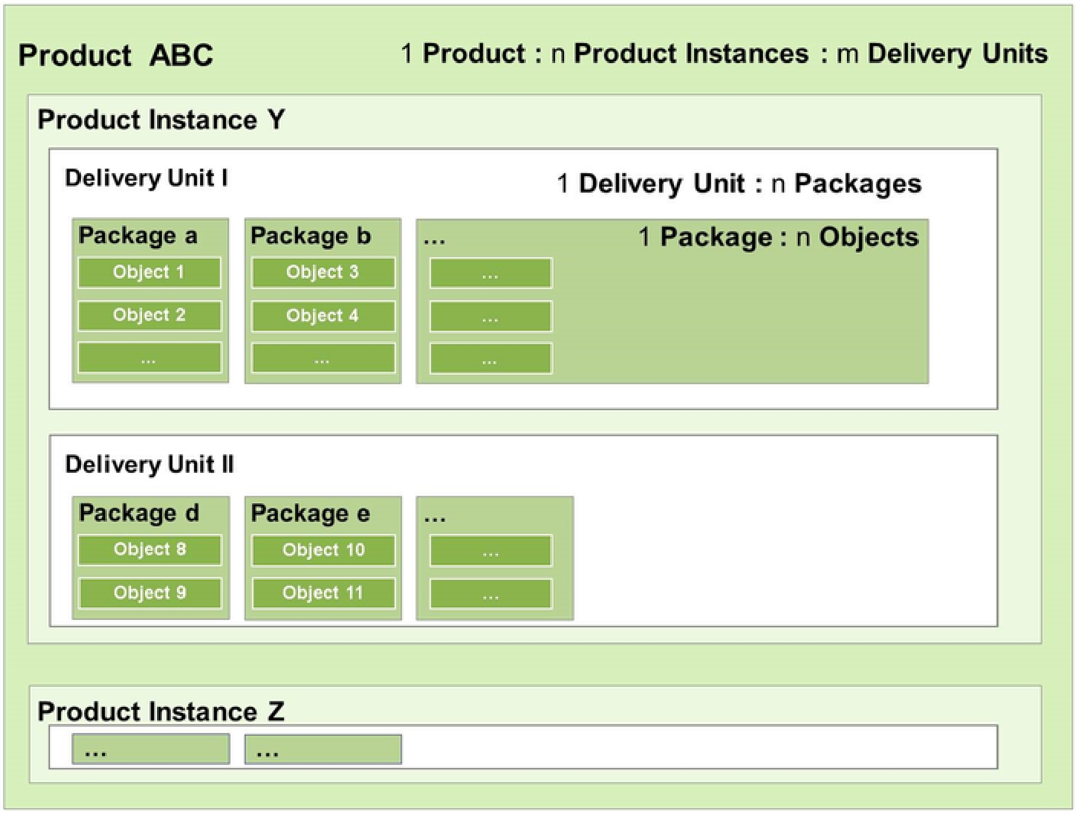

Figure 255: Overview: Products, Delivery Units, Packages, and Objects

Typical Basic Transport Landscape

A typical basic transport landscape for SAP HANA consists of a development system, a test system, and a productive system.

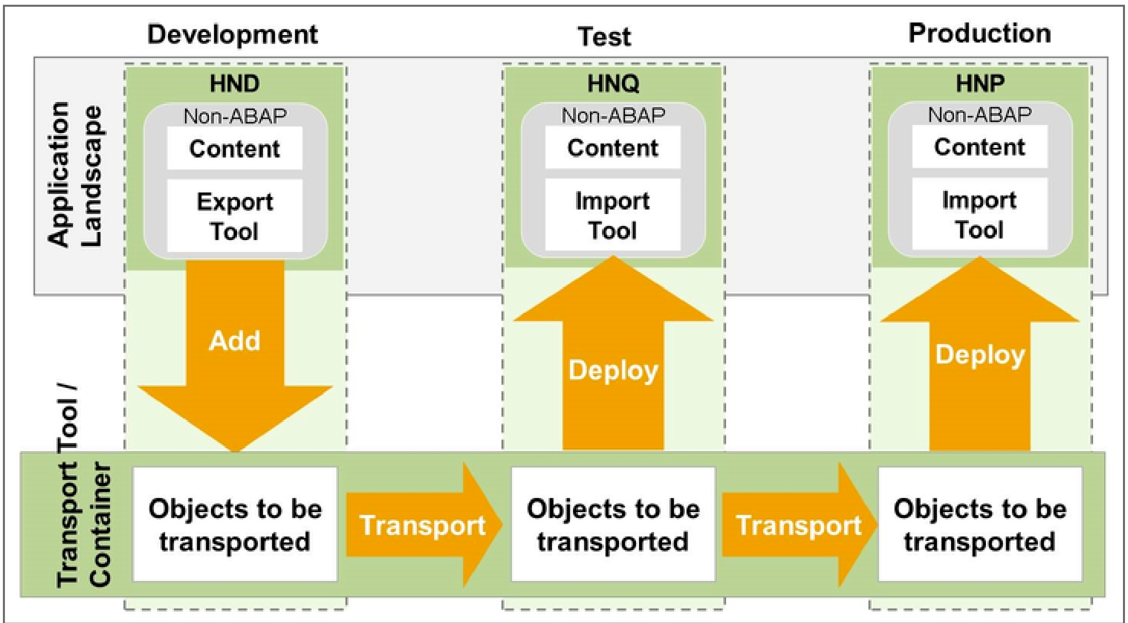

Figure 256: TypicalBasicTransportLandscape_Image.ppt

Transporting SAP HANA Content: Available Options

There are multiple options for transporting SAP HANA content. Which one is suitable depends on the use case and integration scenario, as follows:

- **Native SAP HANA Content**

  SAP HANA Application Lifecycle Manager can be used to transport native SAP HANA content. Because this is an SAP HANA standalone transport management tool, it is suitable for customers without an ABAP footprint. It is a lightweight and easy-to-use transport tool.

- **Native SAP HANA Content or as Part of a Solution**

  With the Enhanced Change and Transport System (CTS+), SAP HANA content can be transported like any other non-ABAP content. This facilitates integration in the existing CTS transport landscape and integration in SAP process tools (ChaRM, QGM).

- **SAP HANA Content Exclusively Used by ABAP**

  An alternative for transporting SAP HANA content are the SAP HANA Transport Containers, which are used exclusively by ABAP (ABAP for SAP HANA). With that, SAP HANA artifacts can be transported with standard ABAP transport. This also ensures integration in the existing CTS transport landscape and in SAP process tools.

- **Content That Needs to be Transferred Quickly Without Transport Management System**

  Content can be quickly transferred from one SAP HANA system to another using the export and import functionality. This facilitates moving objects with little effort. However, in many cases using a transport management solution is a better option than this manual approach.

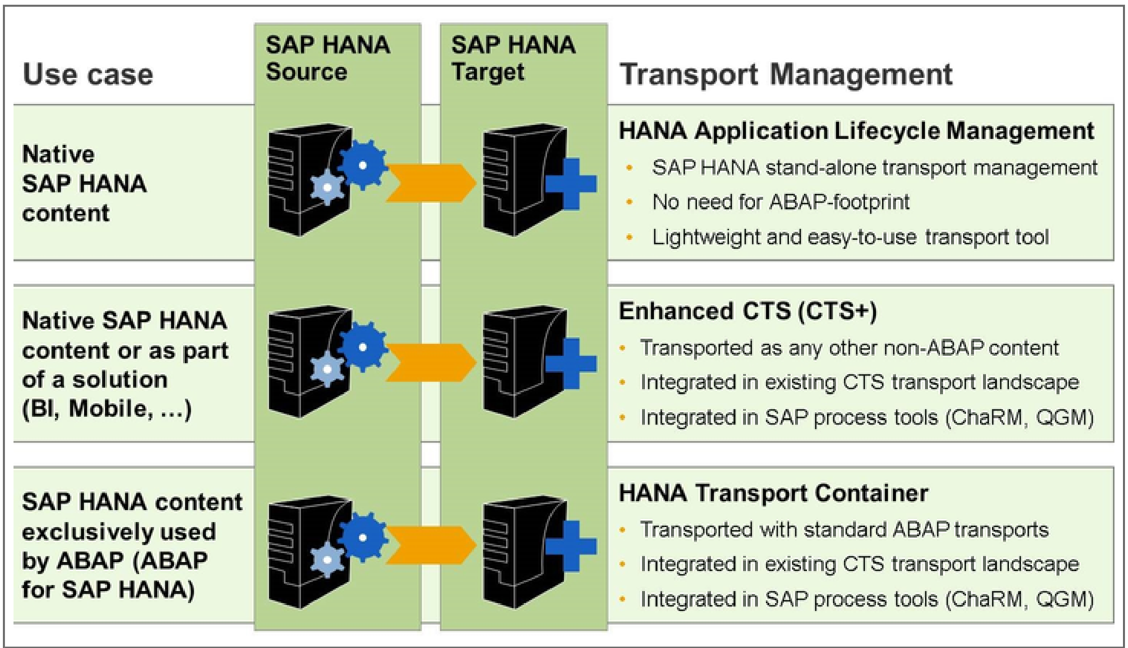

Figure 257: Transporting SAP HANA Content: Available Options

Transport of Native SAP HANA Content with SAP HANA Application Lifecycle Manager (HALM)

The SAP HANA Application Lifecycle Manager enables you to create your product, delivery unit, package, and basic application components. Additionally, the SAP HANA Application Lifecycle Manager enables administrators to set up the transport of delivery units and changes, start and monitor transports, and upload or download delivery unit archives.

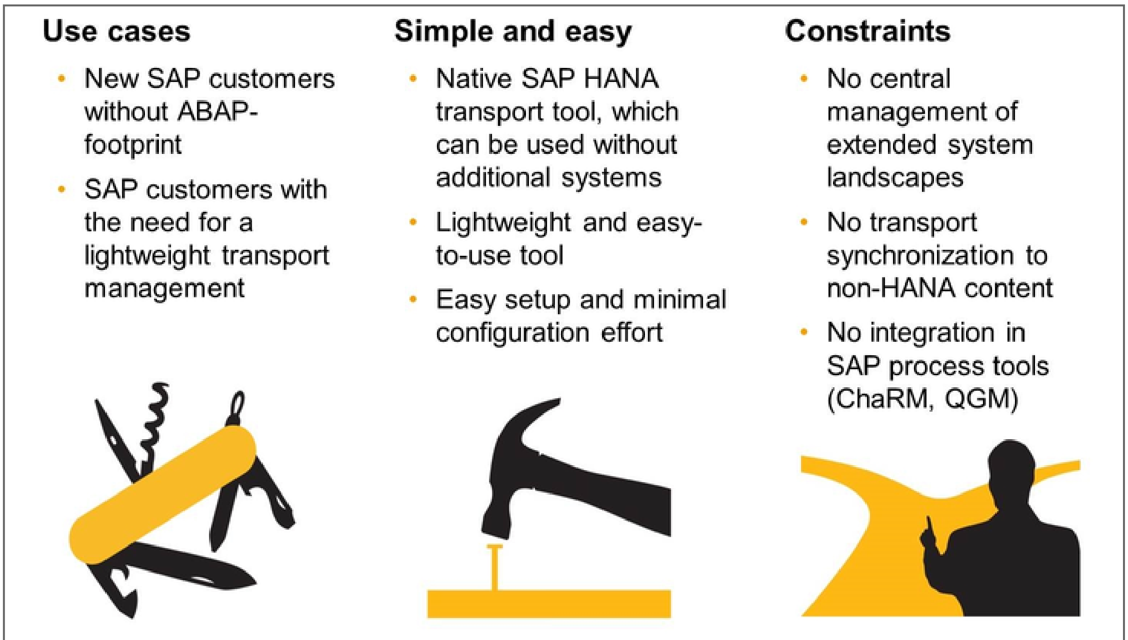

Figure 258: SAP HANA Application Lifecycle Manager: Use Cases and Constraints

SAP HANA Application Lifecycle Manager: Capabilities

- **Product Management Capabilities**

It provides product management capabilities to define and maintain products, delivery units, and packages for metadata setup for SAP HANA content development.

- **Transport Management Capabilities**

  It provides transport management capabilities to manage content changes between two native SAP HANA systems.

SAP HANA Application Lifecycle Manager: Transport Process

As an administrator, you can use the SAP HANA Application Lifecycle Manager as a single point of access to perform the following tasks:

- Assign the appropriate delivery units or changes to the transport route
- Execute exports and imports (uploads and downloads)
- Monitor the transport processes

> Note:
> The SAP HANA Application Lifecycle Manager tool is available on the SAP HANA XS Web server.

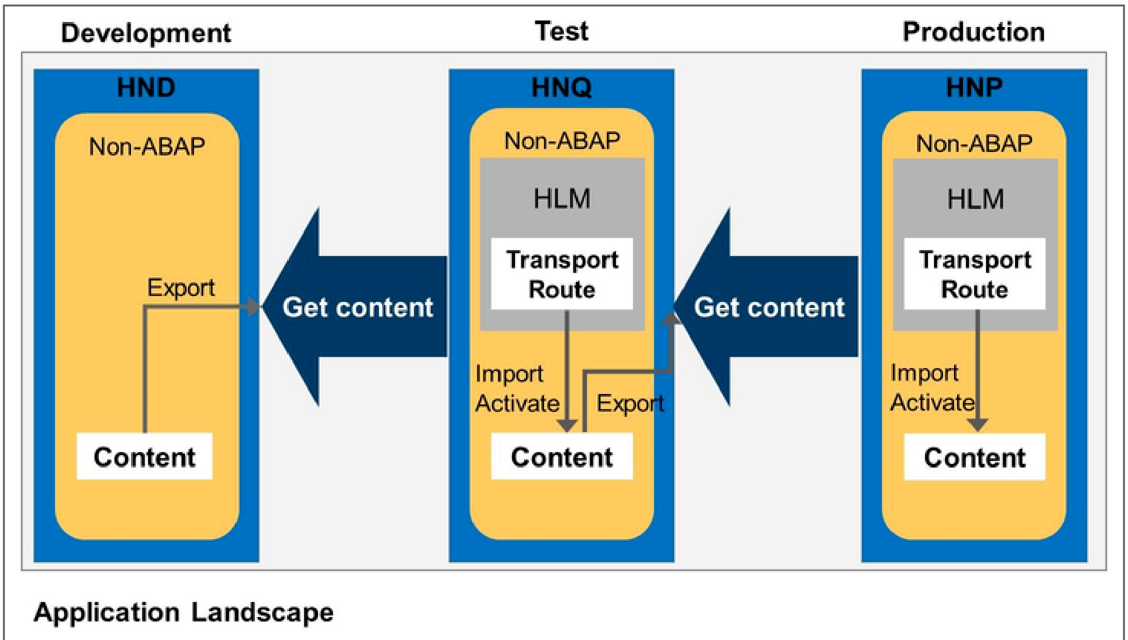

Figure 259: SAP HANA Application Lifecycle Manager: Transport Process

SAP HANA Lifecycle Manager: Web Application

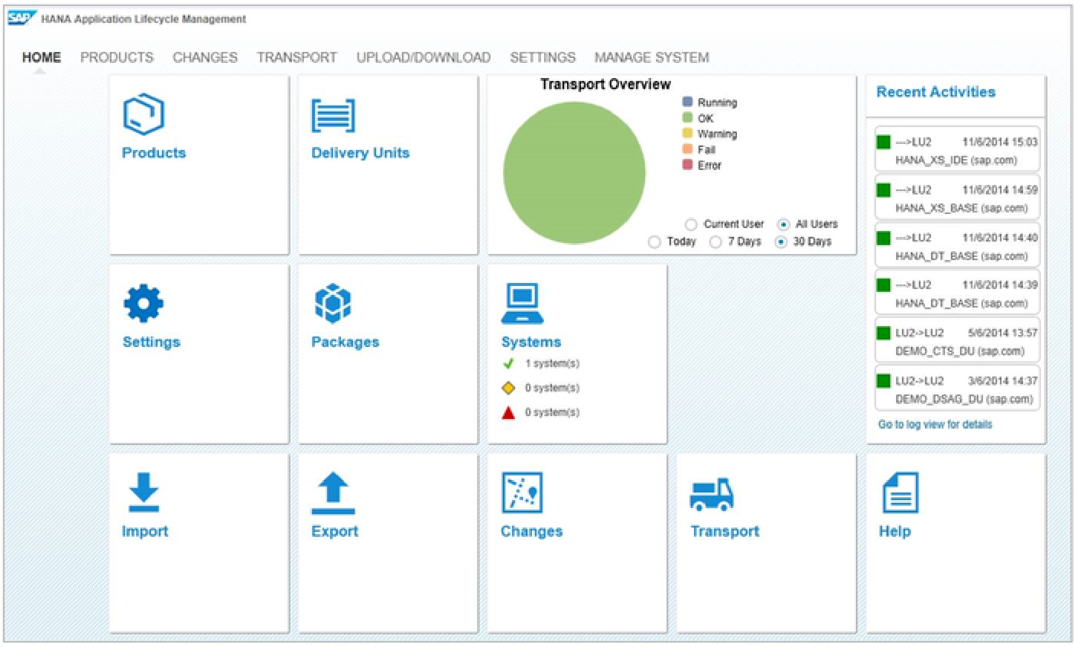

Figure 260: SAP HANA Lifecycle Manager: Web Application

The responsibility for common application-lifecycle management performed with the SAP HANA Application Lifecycle Manager is shared between the various lifecycle management roles. These roles must be assigned to the SAP HANA users who start the SAP HANA Application Lifecycle Manager.

For example, the Administrator role provides access to all options and tools in the SAP HANA Application Lifecycle Manager. To start a transport operation based on a defined route, you only need the privileges assigned with the ExecuteTransport user role. The Display role allows a user to view details of the delivery units, routes, and transports, but they cannot make any changes.

Granularity of Transports

The granularity of transports in SAP HANA Lifecycle Manager is as follows:

- Full Deliver Unit or Product (without Change Recording)
- Full Released Delivery Unit or Product (with Change Recording enabled)
- Change (with Change Recording enabled)

Change Recording in SAP HANA

Change recording provides the infrastructure to record changes during development.

Change recording provides the following:

- Automatic recording and grouping of object changes
- Decoupling of activation and transport
- Predecessor calculation of changes

Change Recording can be enabled as global system setting in your development environment.

Transporting without change recording has the following features:

- Delivery Unit transport contains all active objects in the packages of that particular Delivery Unit.
- If an object is ready to be transported, its Delivery Unit must be activated.

Change Tracking

Transporting with change recording has the following features:

- Automatic recording of object changes to a change list when an object is activated
- Team Development

  Allows a developer (or team) to work on a development artifact and release the “change” only when the artifact is ready to promote to the test system. For developers not contributing to this change, the objects are locked.
- Release in two steps

  Contributors have to approve first before a change can be released.
- Transport

  Delivery Unit transport contains only objects where their change is released.

- Optional mode for managing content in a promote-to-production system landscape (activate in HALM)
- Automatic recording of changes to a change list
- Allows a developer (or team) to work on a development artifact and release the “change” only when the artifact is ready to promote to the test system
- Provides more precise control over which objects get transported from the development system

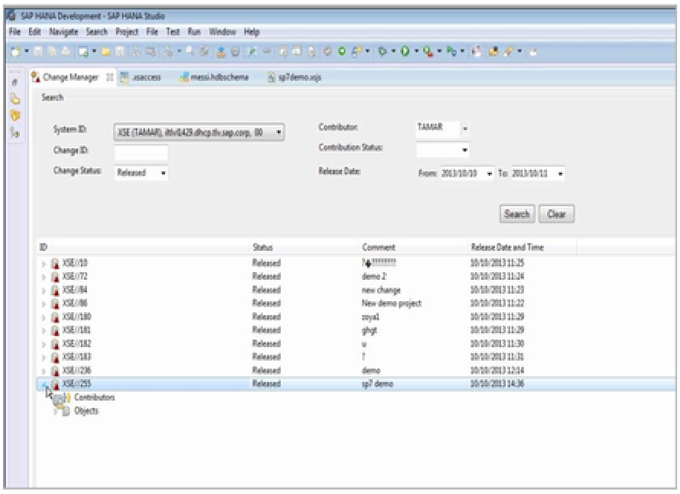

Figure 261: Change Tracking

Dependency Viewer

SAP HANA Lifecycle Manager includes a graphical tool to display dependencies between delivery units.

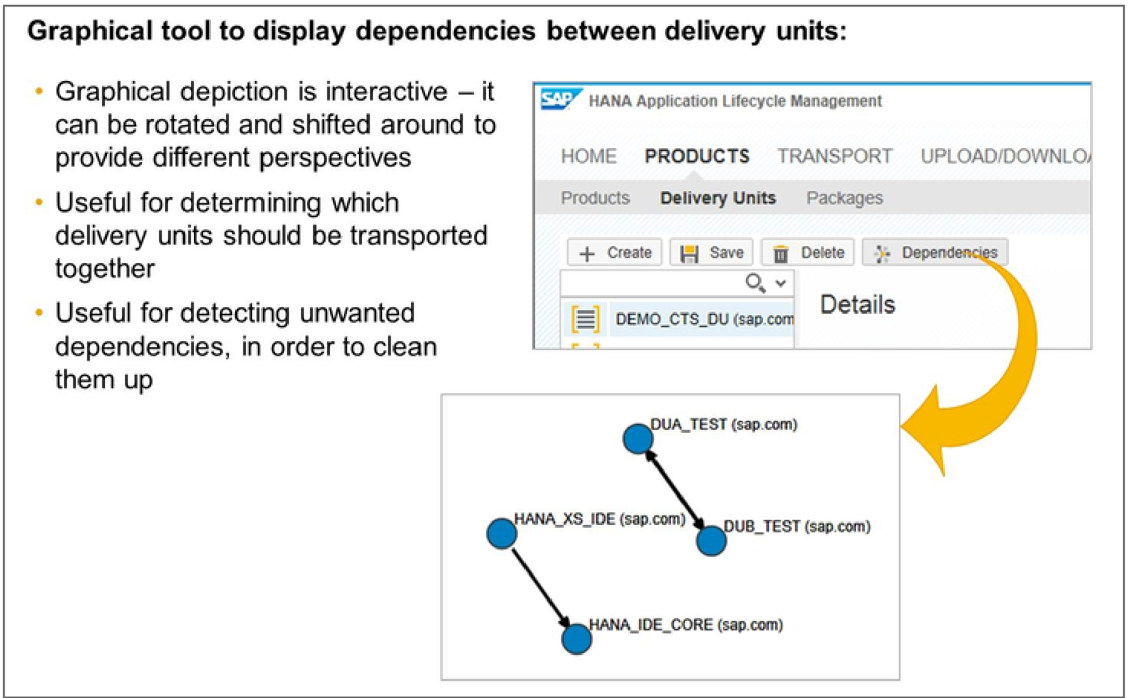

Figure 262: Dependency Viewer

Using SAP HANA Application Lifecycle Manager

> Hint:
> For more information on setting up and using SAP HANA Application Lifecycle Manager, see the SAP HANA Developer Guide.

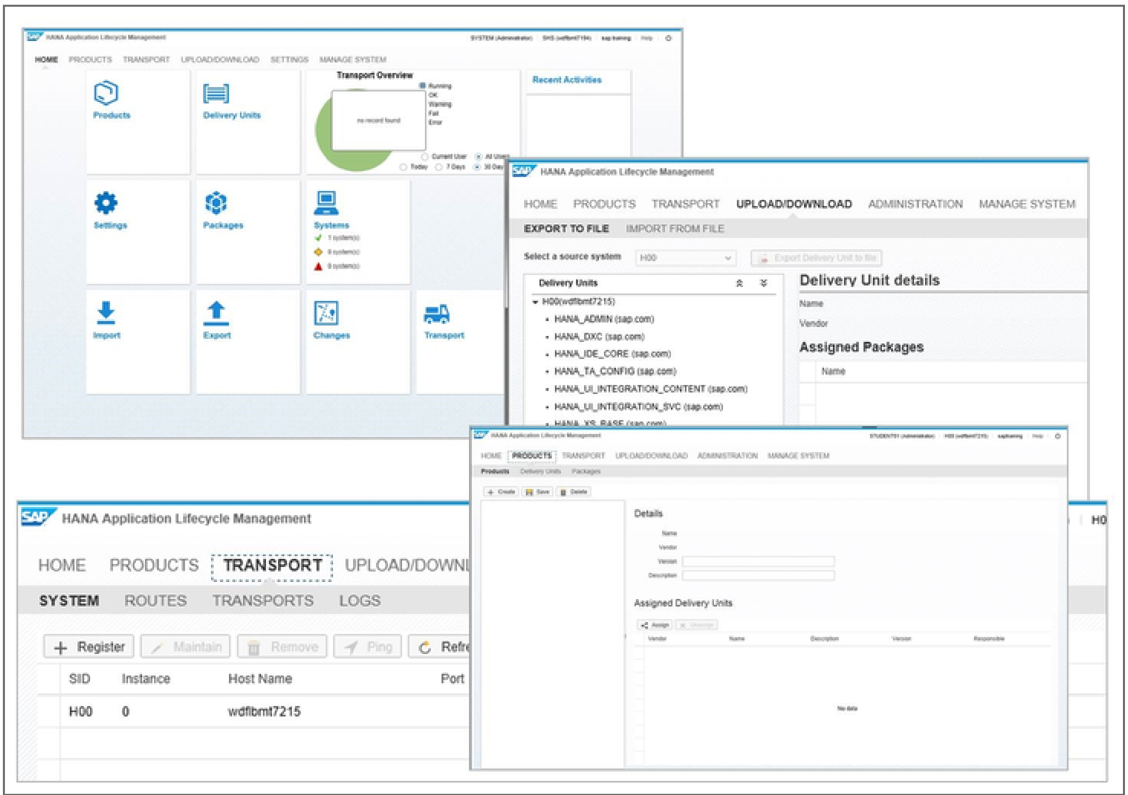

Figure 263: Using SAP HANA Application Lifecycle Manager

Transport of SAP HANA Content with the Enhanced Change and Transport System (CTS+)

The Change and Transport System (CTS) of ABAP has been enhanced so that it can also be used for transporting non-ABAP objects, known as CTS+ or enhanced CTS.

You might already use CTS, for example to manage non-ABAP transports for applications like the SAP Enterprise Portal or to transport your SAP BW ABAP objects. If so, you might want to use the same tool to transport the SAP HANA objects as well. With the integration of SAP HANA into CTS, this is now possible. You can model your landscape for your SAP HANA systems in Transport Management System (TMS) in the same way as with any other non-ABAP application supported by CTS.

To use SAP HANA with CTS as described in this lesson, your systems have to fulfill certain prerequisites.

Using CTS+ with SAP HANA: Prerequisites

- CTS+ System
  - SAP Solution Manager 7.1 SPS05, SAP NetWeaver 7.3 including enhancement package 1, or SAP NetWeaver 7.4 (AS ABAP and AS JAVA)
  - CTS plug-in installed on the CTS system (taken from SL Toolset 1.0 SP04 at least — always use the newest CTS plug–in available)
  - SAP Note 1731044 or SAP Note 1730989 must be implemented on the host of the CTS Deploy Web Service
- SAP HANA

- SAP HANA Platform

Using CTS+ with SAP HANA: Landscape

The figure, Using CTS+ with SAP HANA: Landscape, shows the systems that are involved in the scenario. The figure shows, as an example, a three system landscape consisting of a development, a test, and a production system. This is a basic example. You can set up larger or simpler landscapes in CTS. All of the options that you might know from TMS are available for SAP HANA systems as well. You can, for example, have several systems in a row, or more than one target system at once.

In addition, you need a system where CTS is configured. For the setup, use an SAP Solution Manager or SAP NetWeaver where the CTS Plug-In contained in Software Logistics (SL) Toolset is installed. The set-up is described in the How-To Guide on: http://scn.sap.com/docs/DOC-8576. In this lesson, this system is referred to as "CTS system".

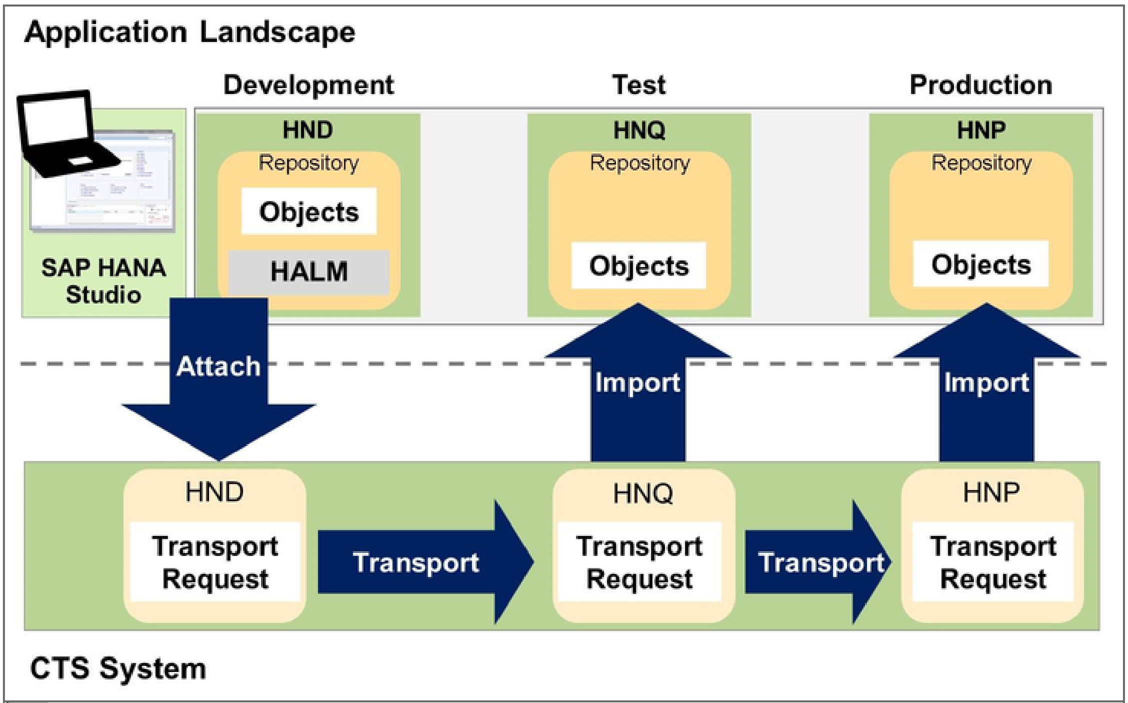

Figure 264: Using CTS+ with SAP HANA: Landscape

The figure, Using CTS+ with HANA: Landscape, also illustrates the process of exporting and importing objects with SAP HANA. The front end is the SAP HANA Studio, or (starting with SPS08) the SAP HANA Application Lifecycle Management (HALM). You can start the export from the SAP HANA Developer Studio or SAP HANA Application Lifecycle Management. You no longer need to use the option of exporting content to a file system and attaching it manually to a transport request.

The next step is to release the transport request. Depending on your configuration, this is either done automatically, or by the Transport Organizer Web UI. You can then start the import. This is done on the CTS system.

> Note:
> Since SAP HANA studio SP05, you are no longer required to export the SAP HANA content to the file system and attach it manually to a CTS transport request. You can now export SAP HANA content and attach it to a transport request in one step (referred to as "Close Coupling"). This is now the preferred way of exporting SAP HANA content to a transport request.

Using CTS+ with SAP HANA: Export Process in SAP HANA Studio

Before you use CTS with SAP HANA, configure your CTS system and the SAP HANA development system (remember that you have to install the CTS plug-in).

On the CTS system, the following elements require configuration:

- The Deploy Web Service is needed to start the deployment on the target systems.
- The Transport Organizer is used to manage transport requests for non-ABAP applications.

After performing these two steps, the systems and the transport route in CTS are ready. As a last configuration step, you have to configure the connection from your SAP HANA development (source) system to the CTS system. This configuration is done in SAP HANA Application Lifecycle Management (HALM).

If CTS is enabled, you have the following options for transports:

- Transport full Delivery Units (DU) based on the active state of the contained objects.
- Transport only the changed objects per Delivery Unit based on released changes (as of SAP HANA SPS08, if Change Recording is enabled).

For exporting, you can either use SAP HANA Application Lifecycle Management (more details are provided in How-To Guide on: http://scn.sap.com/docs/DOC-8576) or the SAP HANA studio (more details are provided in How-To Guide on: http://scn.sap.com/docs/DOC-8576).

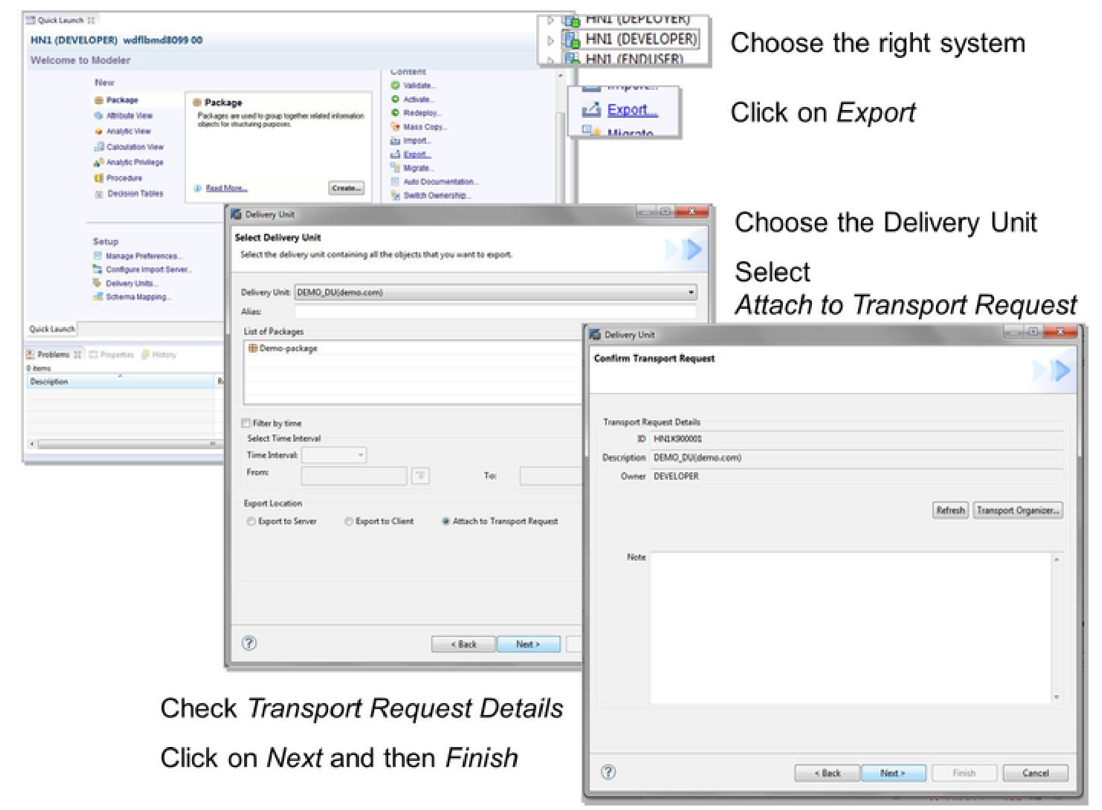

Figure 265: Using CTS+ with SAP HANA: Export Process in SAP HANA Studio

Using CTS+ with SAP HANA: Import Process in TMS

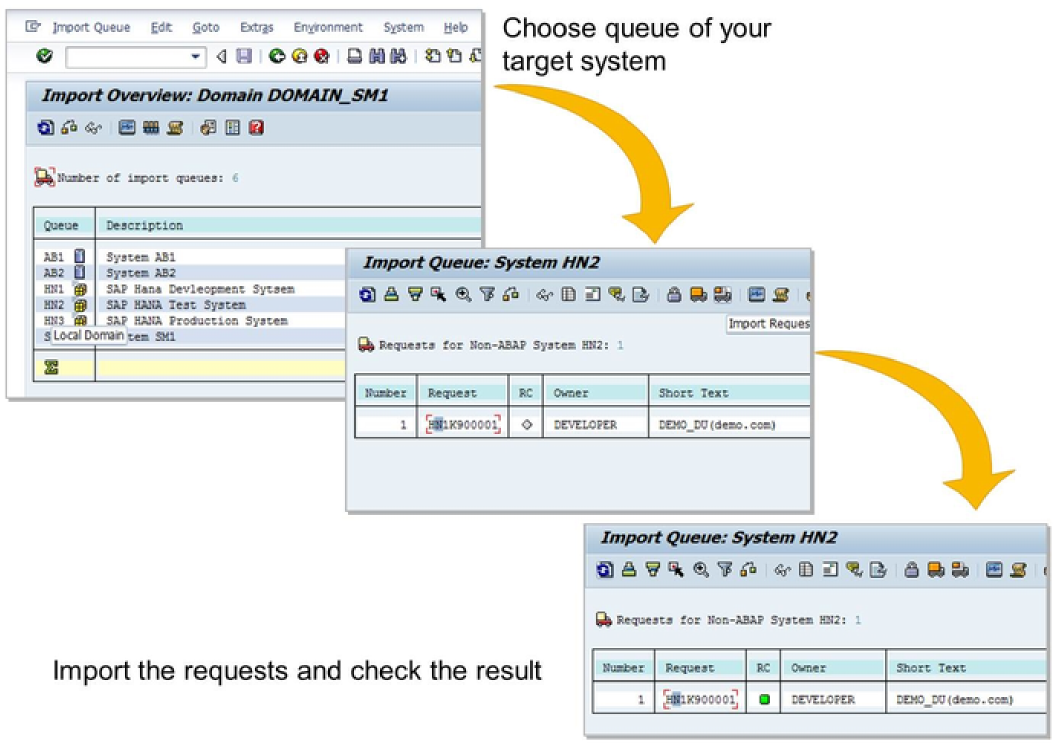

Figure 266: Using CTS+ with SAP HANA: Import Process in TMS

Hint:
For more information, see: http://scn.sap.com/docs/DOC-8576 and https://scn.sap.com/docs/DOC-45659 .

Use of SAP HANA Transport Containers for ABAP for SAP HANA Content

Since SAP NetWeaver 7.4, there are numerous SAP HANA-related optimizations that help to develop ABAP applications for SAP HANA.

The development of ABAP coding and SAP HANA artifacts that belong together means that they also need to be transported together consistently through the system landscape. The SAP HANA Transport Container (HTC) can be used for this.

SAP HANA Transport Container (HTC): Overview

- Targeted at ABAP for SAP HANA applications:
  - Transporting ABAP coding and SAP HANA artifacts together consistently
- SAP HANA Transport Container for ABAP:
  - ABAP development object
  - Allows integration of SAP HANA content into standard Change and Transport System (CTS)

SAP HANA Transport Container (HTC): Landscape

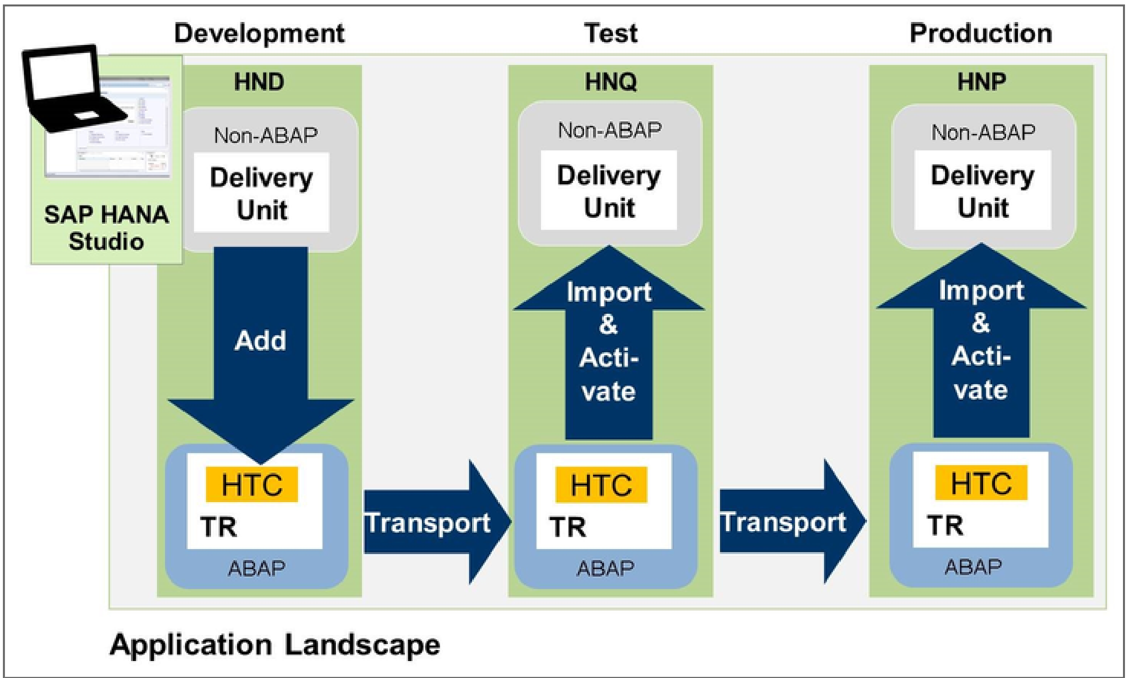

Figure 267: SAP HANA Transport Container (HTC): Landscape

The SAP HANA Transport Container is an ABAP development object, which integrates SAP HANA repository content into the standard Change and Transport System (CTS). Since SAP NetWeaver AS for ABAP 7.4, the SAP HANA Transport Container has been integrated into the Transport Organizer of SAP NetWeaver AS for ABAP. In this way, the SAP HANA repository content is integrated into the Change and Transport System. It ensures an efficient delivery of applications built out of ABAP and SAP HANA content through the ABAP transport mechanism. SAP HANA Transport Container (HTC) transports full DUs based on the active state of the contained objects.

Note:
This means that ABAP for SAP HANA applications is transported as normal, as with any classic ABAP-based application.

SAP HANA Transport Container (HTC): Procedure Overview

- Source System
  - Create Delivery Unit and Assign Packages
  - Create SAP HANA Transport Container
  - Release Transport Request
- Target System
  - Import Transport Request
  - Activate SAP HANA Content

Hint:
For more information, see: http://scn.sap.com/docs/DOC-43035

Options for Exporting and Importing SAP HANA Content Manually

As an alternative to using a transport management solution for a quick test transfer, you can use the export and import functionality of SAP HANA. Exporting and importing is possible as client-side and server-side.

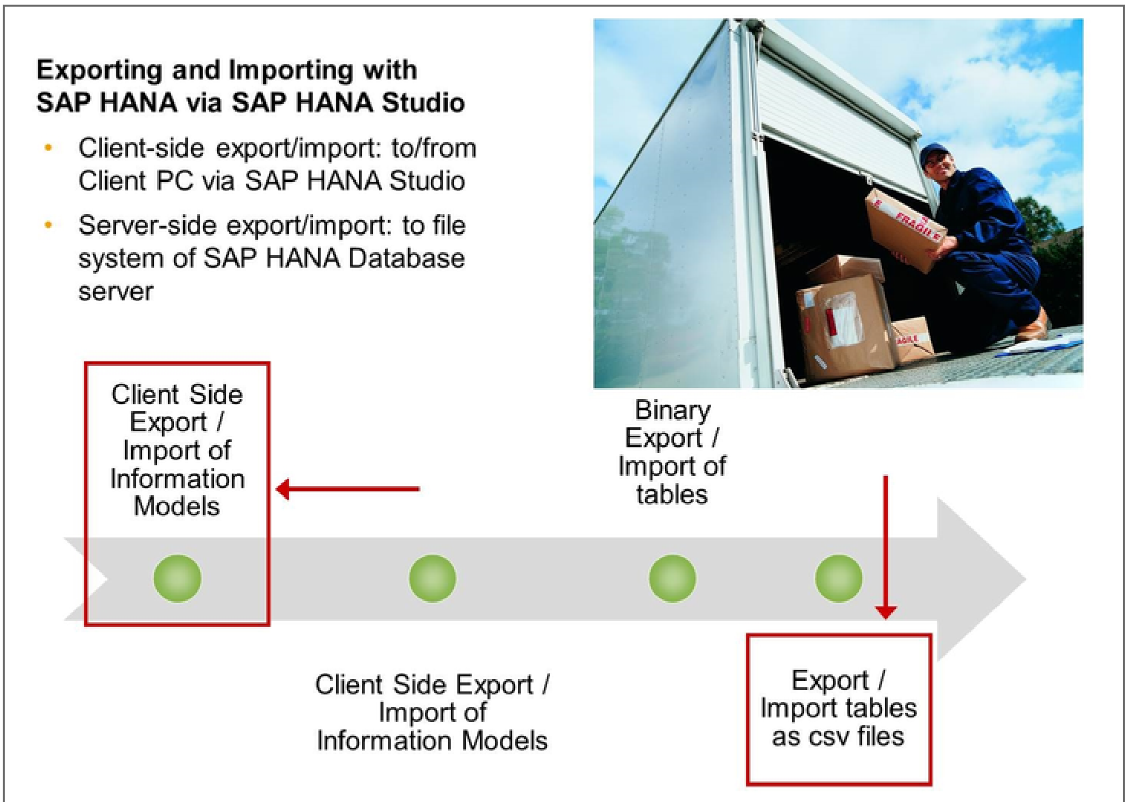

Figure 268: Using Client Export and Import for Models

Using Client Export and Import for Models

You can export all catalog objects to a file system and then import them back into another database. For example, if you want to move data from a test system to a production system, clone your system, or provide the data to SAP Support so that they can replicate a scenario.

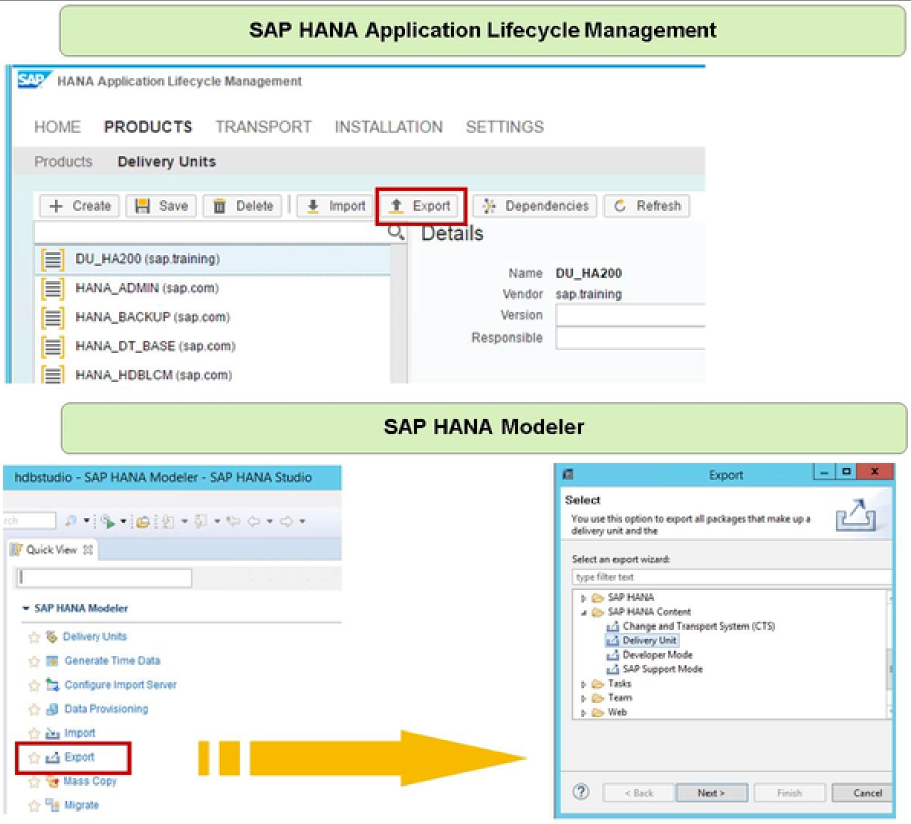

Figure 269: Using Client Export and Import for Models

Note:
If you want to specify a different directory in the server's file system, it must already exist and the database must have authorization to access it.

Export and Import of Tables: Considerations

Difference between binary export and csv export

CSV export does not contain DDL statement

- You can only re-import into existing table (or have to create table “somehow”)
- Binary export contains DDL as well → can also create the table

CSV export is human readable

- May be a security issue

CSV export is not compressed

- Export about factor 10 larger than binary export
- Example: Table MARA exported as binary and CSV
  binary is 1.7 MB
  csv is 14 MB

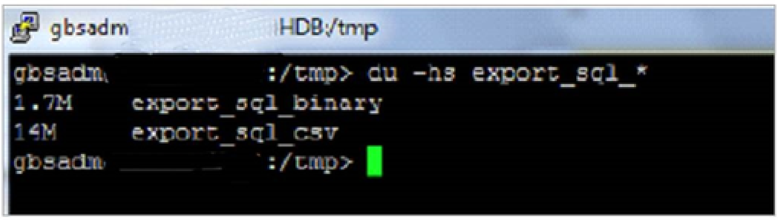

Figure 270: Export and Import of Tables: Considerations

Note:
You can use your authorization to prevent the export of content. For more information, see the developer guide.

Hint:
For the export of small tables or catalog-only exports, a CSV export to the client file system is appropriate. However, consider the maximum file size of your operating system. A binary export on the server is recommended for large exports (for example, exports over 2 GB).

LESSON SUMMARY
You should now be able to:

- Transport changes

Appendix: Administration Tasks in SAP HANA Studio

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Understand the administration tasks that are still in SAP HANA studio

Administration Tasks in SAP HANA Studio

Monitoring Hosts in a Distributed System

In a distributed system, you can monitor hosts from the *Hosts* subtab in SAP HANA Studio.

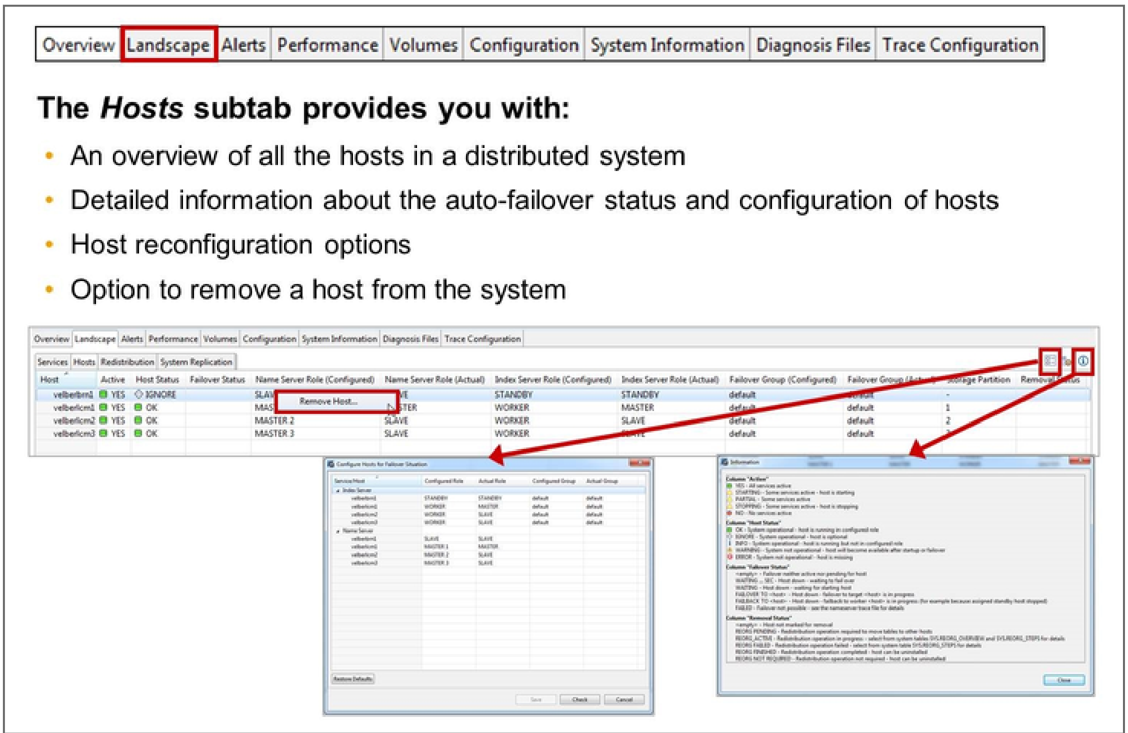

Figure 271: The Administration Editor: Hosts Subtab

Redistributing Data in a Scale-Out System

In a distributed system, tables and table partitions are assigned to an index server on a particular host at the time of their creation. However, you can change this assignment, and in certain situations, changing assignments is necessary. SAP HANA supports several “redistribution operations” that use complex algorithms to evaluate the current distribution and determine a better distribution, depending on the situation.

Overview | Landscape | Alerts | Performance | Volumes | Configuration | System Information | Diagnosis Files | Trace Configuration

The ***Redistribution*** subtab allows you to execute redistribution operations if you need to:

- Redistribute data before removing a host from the system
- Redistribute data after adding a new host to the system
- Optimize current table distribution
- Optimize table partitioning

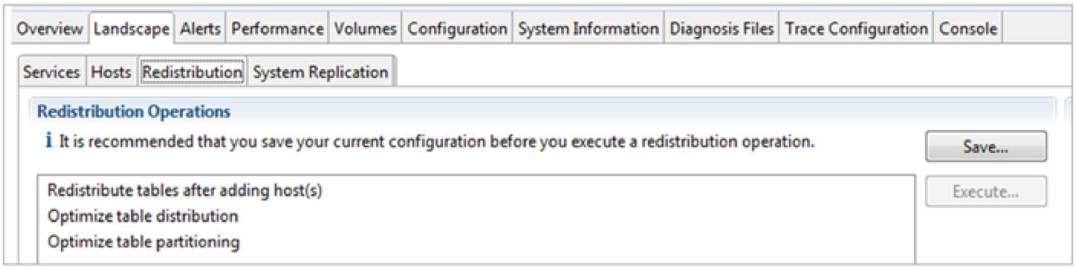

Note: It is always recommended that you save the current distribution before executing a table redistribution operation in case you need to restore it.

Figure 272: The Administration Editor: Redistribution Subtab

Setting Up and Monitoring System Replication

System replication ensures the high availability of an SAP HANA system. Through the continuous replication of data from a primary to a secondary system, including in-memory loading, system replication facilitates a rapid failover in the event of a disaster. Productive operations can resume with minimal downtime.

You can set up and monitor SAP HANA system replication from within the administration console.

Overview | Landscape | Alerts | Performance | Volumes | Configuration | System Information | Diagnosis Files | Trace Configuration

On the ***System Replication*** subtab, you can:

- Perform the initial setup, that is, enable system replication and establish the connection between two identical systems.
- Monitor the status of system replication to ensure that both systems are in sync.
- Trigger failover to the secondary system in the event of a disaster and failback once the primary system is available again.

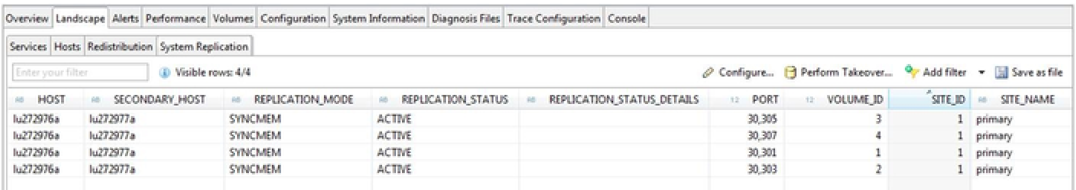

HOST	SECONDARY_HOST	REPLICATION_MODE	REPLICATION_STATUS	REPLICATION_STATUS_DETAILS	PORT	VOLUME_ID	SITE_ID	SITE_NAME
lu272976a	lu272977a	SYNCMEM	ACTIVE		30,305	3	1	primary
lu272976a	lu272977a	SYNCMEM	ACTIVE		30,307	4	1	primary
lu272976a	lu272977a	SYNCMEM	ACTIVE		30,301	1	1	primary
lu272976a	lu272977a	SYNCMEM	ACTIVE		30,303	2	1	primary

Figure 273: The Administration Editor: System Replication Subtab

Extended System Replication Configuration

- To offer higher levels of availability, you can **link together multiple systems** in a multitier system replication landscape.
- After setting up a basic system replication scenario, you add a third system to provide another level of redundancy.
- In a multitier setup, the primary system is always on tier 1, a tier 2 secondary has a primary system as its replication source, and a tier 3 secondary has the tier 2 secondary as its replication source.
- The tier 3 secondary can now also be added using SAP HANA studio.
- On each node, only those actions are provided, which are currently possible, depending on the system status.

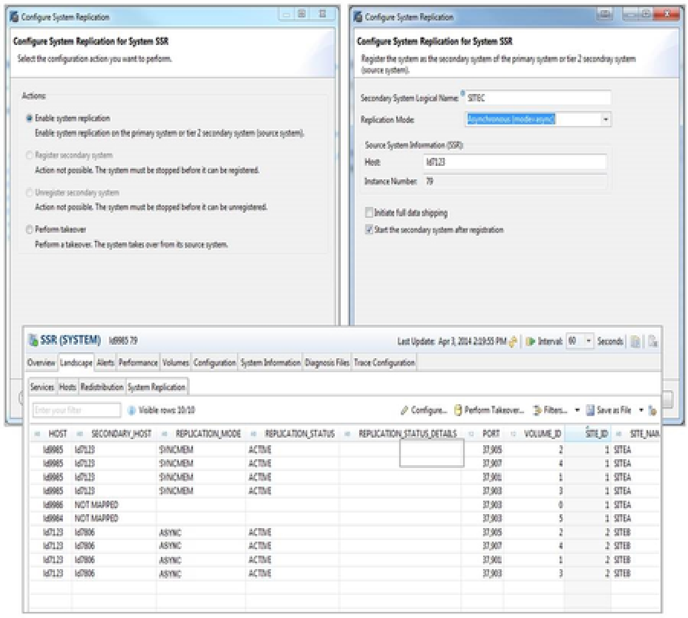

Figure 274: Extended System Replication Configuration

Assessing Performance Information

Gathering and analyzing data about the performance of your SAP HANA systems is important for root-cause analysis and the prevention of future performance issues.

General information about overall system performance is available in the System Monitor and on the Overview tab of the Administration editor. You can monitor more detailed aspects of system performance on the Performance tab.

You can monitor the following detailed aspects of system performance on the ***Performance*** tab:

- Threads
- Sessions*
- Blocked transactions
- Expensive statements*
- SQL plan cache*
- Job progress*
- Load

*The information displayed on these sub-tabs of the Performance tab is detailed and highly customizable. To support administrators performing complex analyses, user-specific column and filter settings are saved when the Administration editor is closed. These settings are restored the next time the tab is opened, independent of system.

Figure 275: The Administration Editor: Performance Tab

The Administration Editor: Threads Subtab

You can monitor running threads in the Threads subtab.

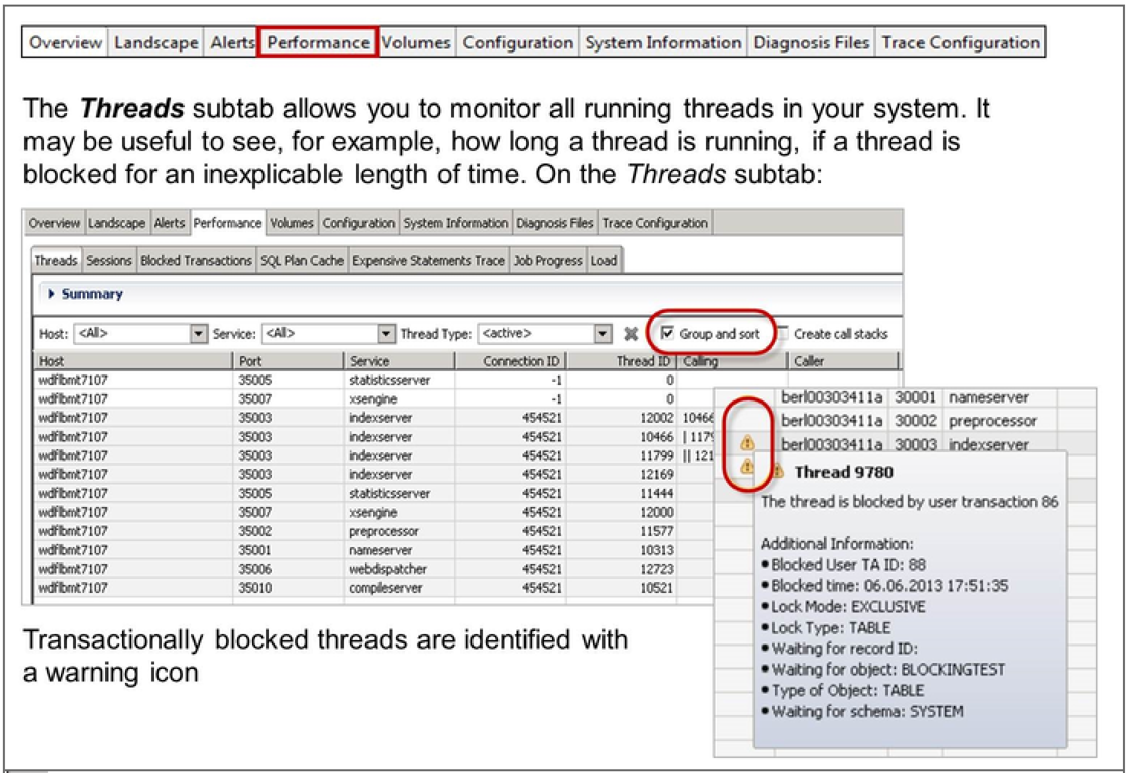

Figure 276: The Administration Editor: Threads Subtab

The Administration Editor: Threads Subtab Actions

The Group and sort filter provides a meaningful and clear structure for thread analysis, as follows:

- Threads with the same connection ID are grouped.
- Within each group, the call hierarchy is depicted.
- Groups are displayed in order of descending duration.

Additional actions can be performed here.

Overview | Landscape | Alerts | Performance | Volumes | Configuration | System Information | Diagnosis Files | Trace Configuration

On the ***Threads*** sub-tab, you can also perform the following actions:

- End the operations associated with a thread
- See the full details of a thread
- Jump to the related calling and called threads of a thread by right clicking the thread
- View the call stack for a specific thread

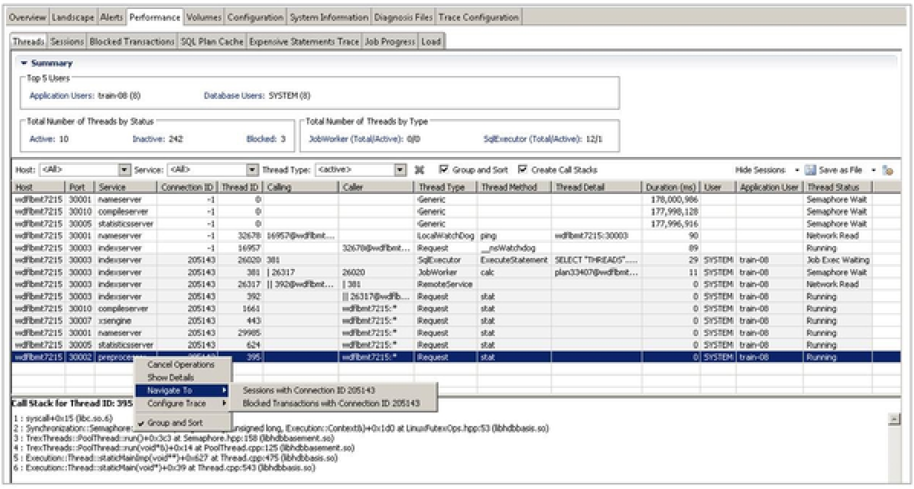

Figure 277: The Administration Editor: Threads Subtab Actions

The Administration Editor: Sessions Subtab

The Sessions subtab allows you to monitor all sessions in the current landscape.

Overview | Landscape | Alerts | Performance | Volumes | Configuration | System Information | Diagnosis Files | Trace Configuration

The ***Sessions*** subtab allows you to monitor all sessions in the current landscape:

- Active/inactive sessions and their relation to applications
- Whether a session is blocked and if so, which session is blocking
- The number of transactions that are blocked by a blocking session
- Statistics like average query runtime and the number of DML and DDL statements in a session

You can also cancel sessions.

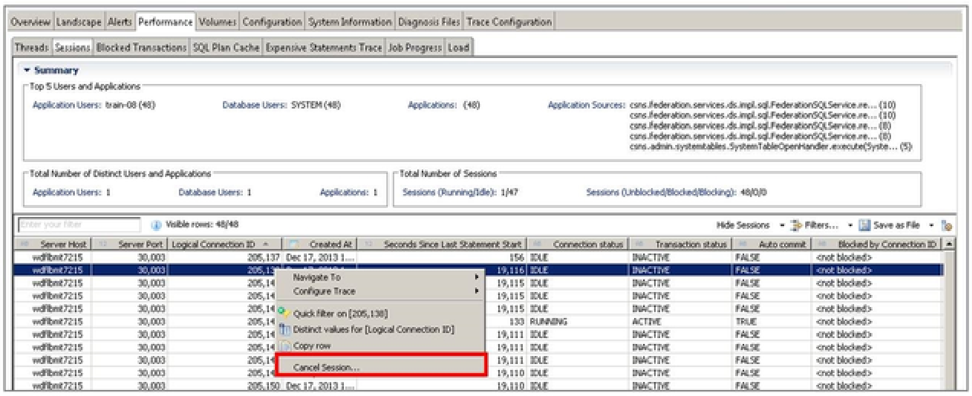

Figure 278: The Administration Editor: Sessions Subtab

The monitor identifies active and inactive sessions, and their relationship to applications. It also shows if a session is blocked and if so, by which other session. It shows if a session is blocking other sessions and how many transactions are inside. Statistics, such as average query runtime and the number of DML and DDL statements in a session, are included.

The table shows the result from the system information statement sessions. To cancel a session, right click the session, and choose Cancel Session .

The Administration Editor: Blocked Transactions Tab

Blocked transactions, or transactionally blocked threads, impact the responsiveness of applications.

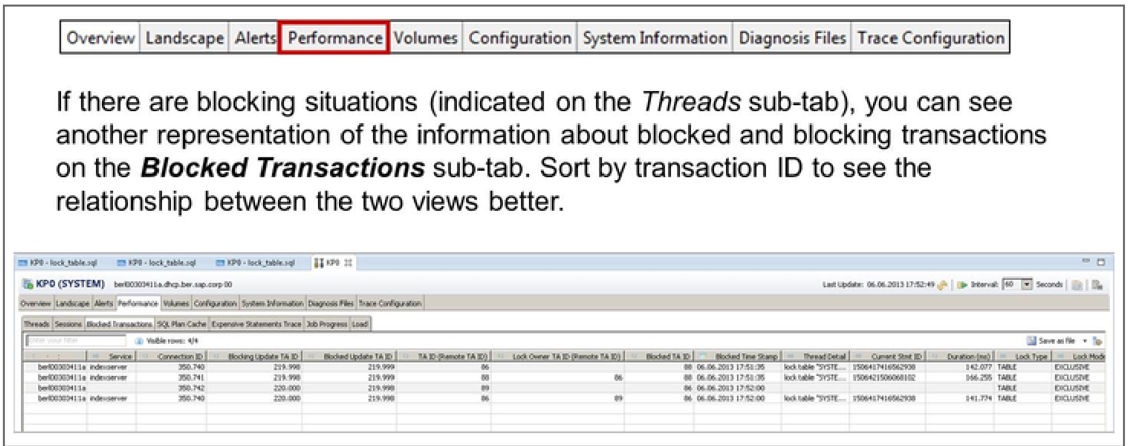

Figure 279: The Administration Editor: Blocked Transactions Tab

Blocked transactions are transactions that cannot be processed further because they need to acquire transactional locks (record or table locks) that are currently held by another transaction. Transactions can also be blocked while waiting for other resources, such as network or disk (database or metadata locks).

The Administration Editor: SQL Plan Cache Tab

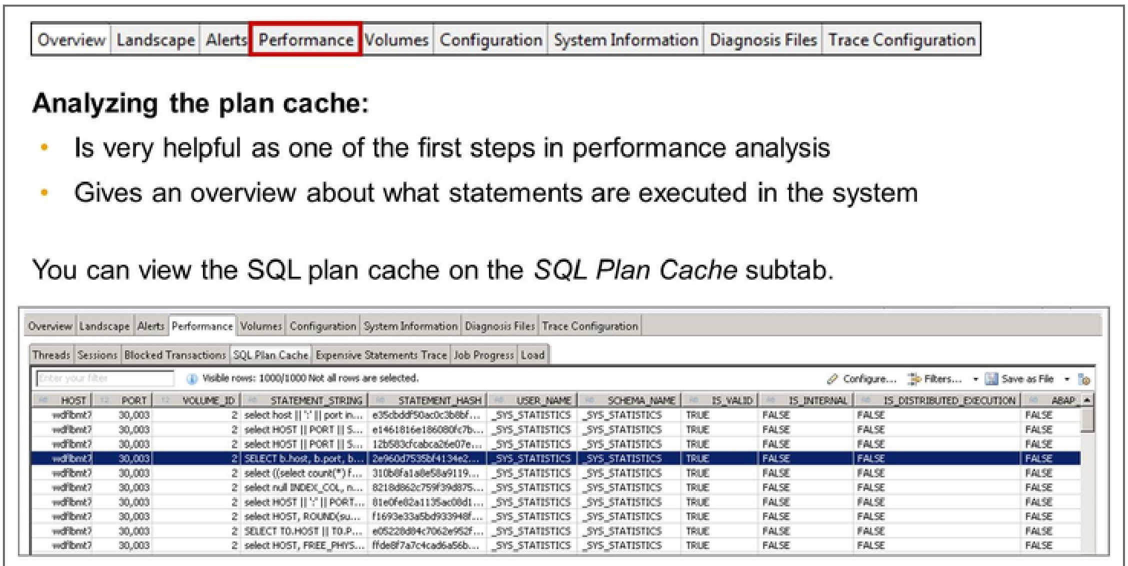

Figure 280: The Administration Editor: SQL Plan Cache Tab

The plan cache stores compiled execution plans of SQL statements for reuse, which improves recompilation at each request. For monitoring reasons, the plan cache keeps statistics about each plan. For example, it stores the number of executions, the minimum, maximum, total, and average runtime, as well as lock and wait statistics. Analyzing the plan cache gives an overview of the statements that are executed in the system, which is a useful first step in performance analysis.

> Note:
> The setup of a cache means that seldom-used entries are removed from the plan cache.

Because the SQL plan lists frequently executed queries, it provides an outline of the workload in the system.

The Administration Editor: Expensive Statements Tab

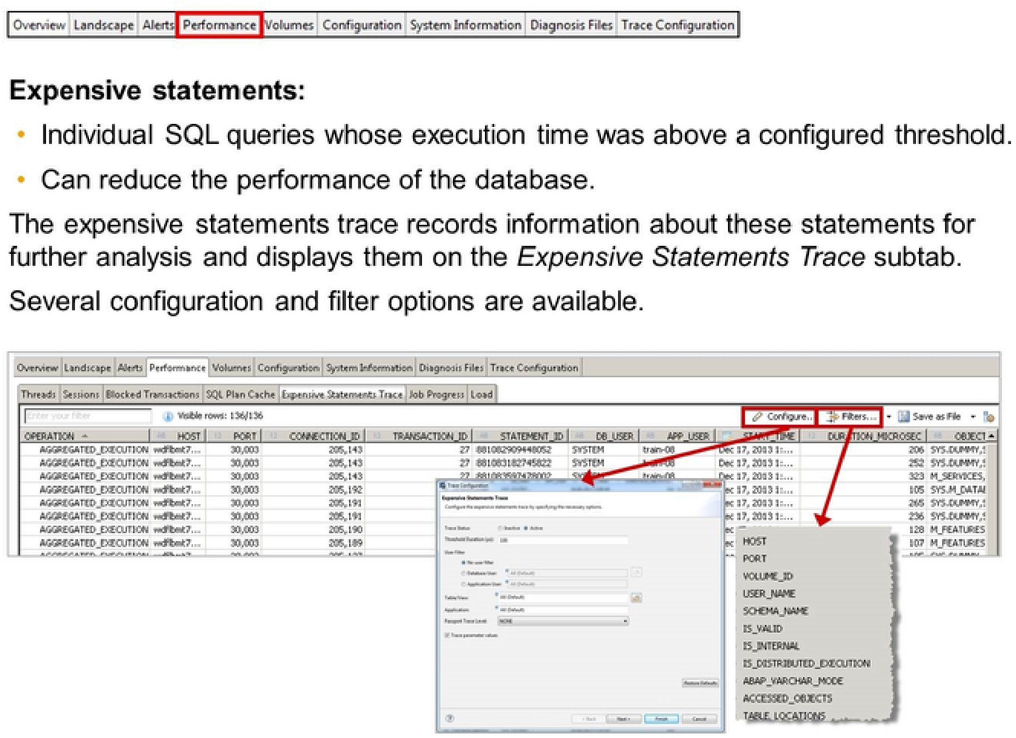

Figure 281: The Administration Editor: Expensive Statements Tab

Expensive statements are individual SQL queries that have an execution time above a configured threshold. The expensive statements trace records information about these statements for further analysis and displays them in the Administration editor.

> Note:
> The expensive statements trace is deactivated by default.

Personalized Administrator View

The individual steps of statement execution are displayed in a hierarchical tree structure underneath the aggregated statement execution information.

Hint:
Some administrator views in SAP HANA Studio are personalized.

The settings are restored the next time the view is opened. The procedure is independent of the system.

This function applies to the following tabs:

- Sessions
- SQL Plan Cache
- Expensive Statements Trace
- Job Progress
- System Replication

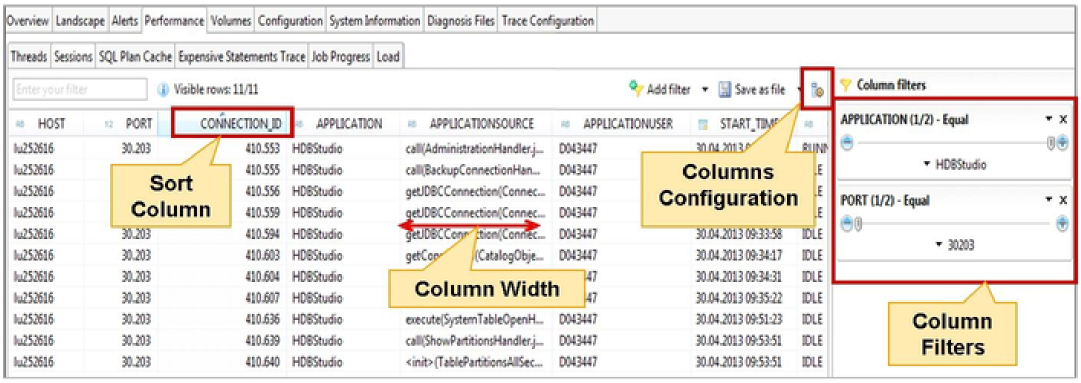

Figure 282: Personalized Administrator View

The Administration Editor: Job Progress Tab

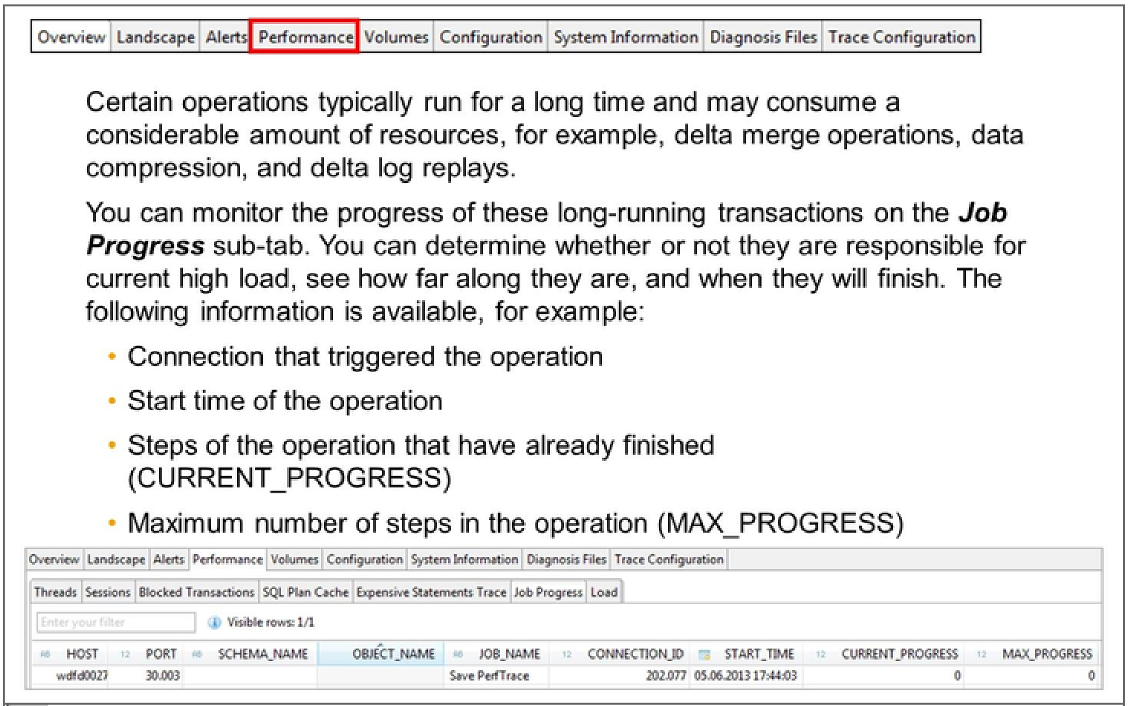

Figure 283: The Administration Editor: Job Progress Tab

The Administration Editor: Load Subtab

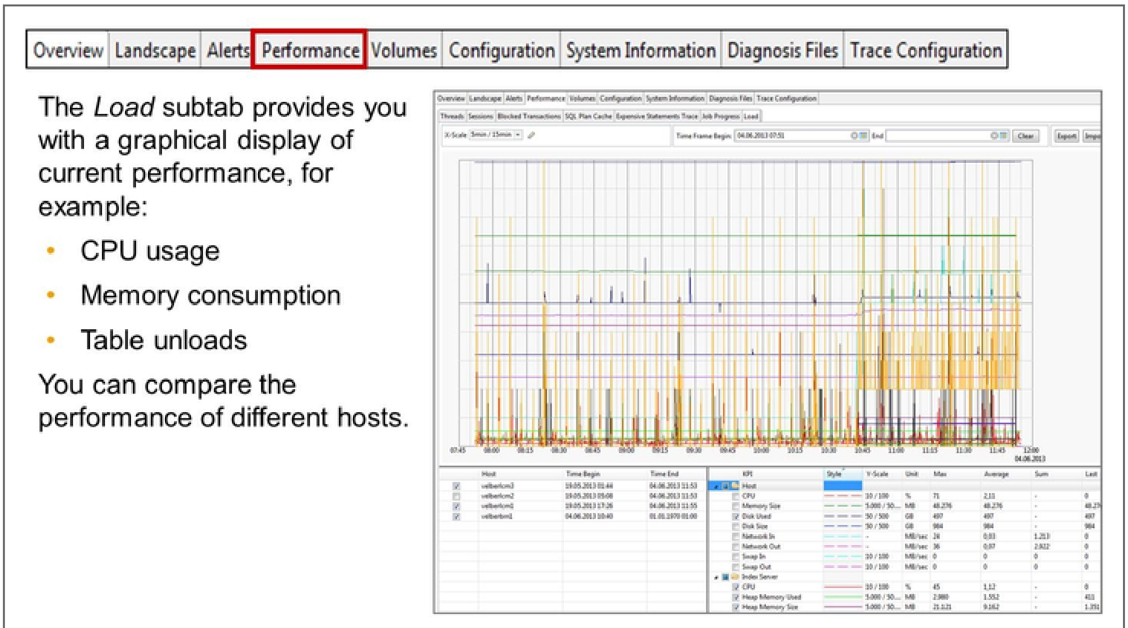

Figure 284: The Administration Editor: Load Subtab

Use the load graph for performance monitoring and analysis. For example, use it to identify the number of blocked transactions that exist now and in the past, or to troubleshoot the root cause of slow statement performance.

Monitoring Disk Usage and Volumes

Overview | Landscape | Alerts | Performance | Volumes | Configuration | System Information | Diagnosis Files | Trace Configuration

To ensure that the database can always be restored to its most recent committed state, you must ensure that there is enough space on disk for data and log volumes. On the ***Volumes*** tab, you can monitor:

- Disk usage
- Volume size
- Other disk activity statistics

There are two views available for monitoring the size of volumes on disk: service, and storage type. You can also filter by host.

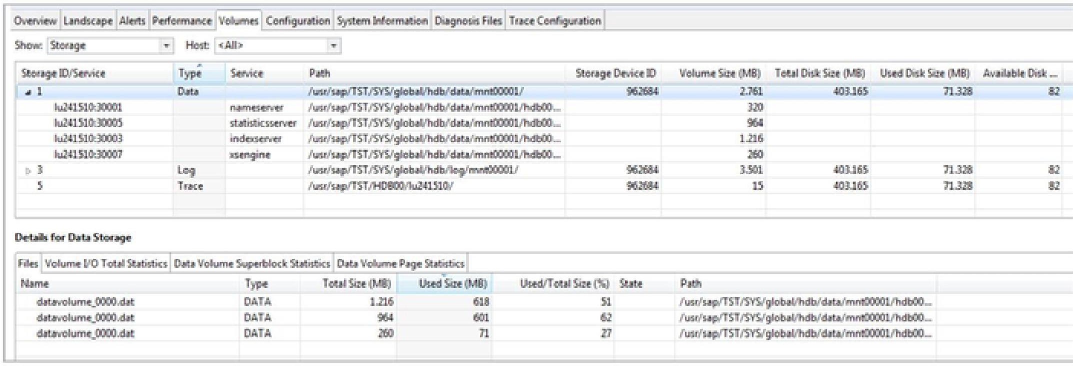

Overview | Landscape | Alerts | Performance | Volumes | Configuration | System Information | Diagnosis Files | Trace Configuration

Show: Storage   Host: <All>

Storage ID/Service	Type	Service	Path	Storage Device ID	Volume Size (MB)	Total Disk Size (MB)	Used Disk Size (MB)	Available Disk ...
1	Data		/usr/sap/TST/SYS/global/hdb/data/mnt00001/	962684	2.761	403.165	71.328	82
lu241510:30001		nameserver	/usr/sap/TST/SYS/global/hdb/data/mnt00001/hdb00...		320			
lu241510:30005		statisticsserver	/usr/sap/TST/SYS/global/hdb/data/mnt00001/hdb00...		964			
lu241510:30003		indexserver	/usr/sap/TST/SYS/global/hdb/data/mnt00001/hdb00...		1.216			
lu241510:30007		xsengine	/usr/sap/TST/SYS/global/hdb/data/mnt00001/hdb00...		260			
3	Log		/usr/sap/TST/SYS/global/hdb/log/mnt00001/	962684	3.501	403.165	71.328	82
5	Trace		/usr/sap/TST/HDB00/lu241510/	962684	15	403.165	71.328	82

Details for Data Storage

Files | Volume I/O Total Statistics | Data Volume Superblock Statistics | Data Volume Page Statistics

Name	Type	Total Size (MB)	Used Size (MB)	Used/Total Size (%)	State	Path
datavolume_0000.dat	DATA	1.216	618	51		/usr/sap/TST/SYS/global/hdb/data/mnt00001/hdb00...
datavolume_0000.dat	DATA	964	601	62		/usr/sap/TST/SYS/global/hdb/data/mnt00001/hdb00...
datavolume_0000.dat	DATA	260	71	27		/usr/sap/TST/SYS/global/hdb/data/mnt00001/hdb00...

Figure 285: The Administration Editor: Volumes Tab

To ensure that the database can always be restored to its most recent committed state, check that there is enough space on disk for data and log volumes. Monitor disk usage, volume size, and other disk activity statistics on the Volumes tab of the Administration editor.

The Volumes tab has two views for monitoring the size of volumes on disk: service, and storage type (that is data, log, and trace).

Hint:
Although trace files are not stored in volumes, they are displayed on the Volumes tab in the Storage view. This is because they use disk space, and therefore, need to be monitored.

Retrieving System Information

Overview | Landscape | Alerts | Performance | Volumes | Configuration | System Information | Diagnosis Files | Trace Configuration

The *System Information* tab lists several predefined SQL SELECT statements on system views. These statements provide you with easy access to important system information.

If you have compiled your own SQL statements for monitoring purposes, you can save these statements on the *System Information* tab for convenient repeated execution.

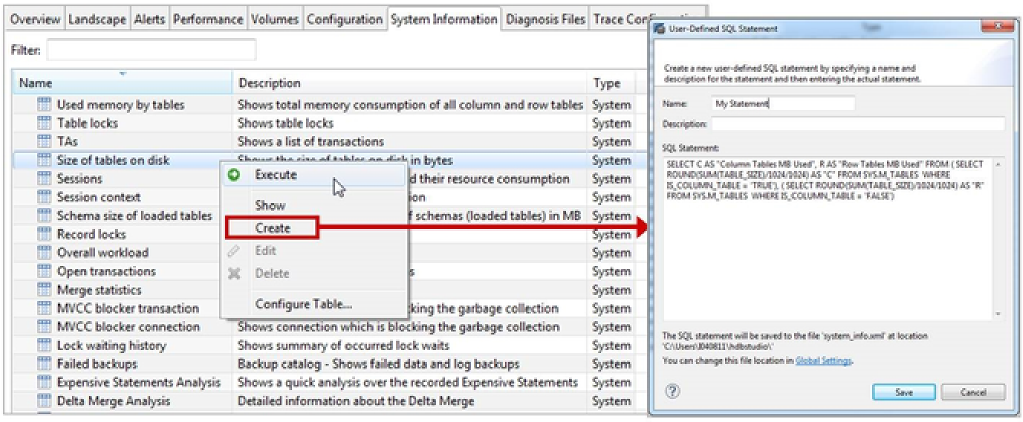

Figure 286: The Administration Editor: System Information Tab

If you double-click an entry in this list, you execute the underlying statement. To see the actual statement, from the context menu, choose Show .

SAP HANA Mini Checks

During the analysis of complex problems, you might want to determine special database information that is not entirely available in standard functions. Because of this, SAP provides a collection of useful SQL statements for SAP HANA database analysis. For information on how to download the SQL statements, see SAP Note 1969700 .

Note:
For the mini checks attachment, see SAP Note 1969700 : SQL statement collection for SAP HANA

For the mini checks documentation, see SAP Note 1999993 : How-To: Interpreting SAP HANA Mini Check Results

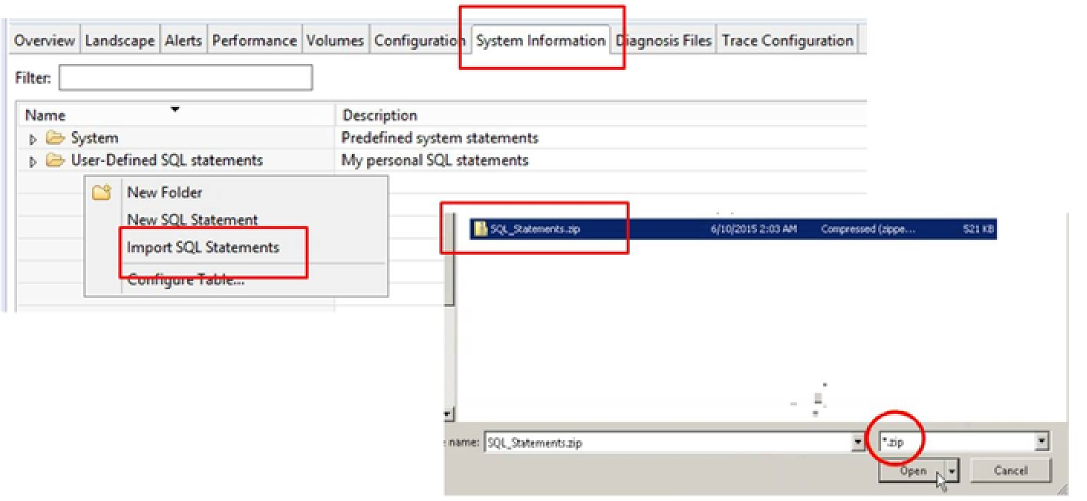

Figure 287: Implement SAP HANA Mini Checks

After you import the `SQL_Statements.zip` file, you can execute these checks in the System Information tab to help with daily monitoring and SAP HANA system analysis. For structured storage, create a separate folder before you import the Mini checks.

SAP HANA Mini Checks Implemented

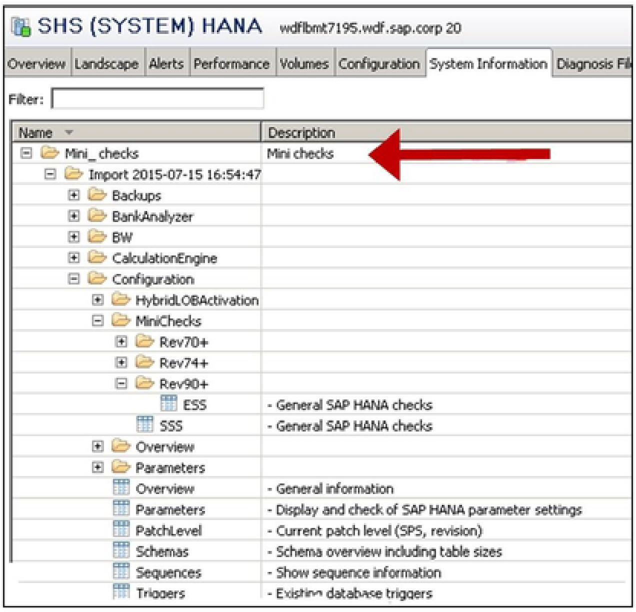

Figure 288: Mini Checks Implemented

Mini Checks Usage

You can use each statement separately, but you can use Mini Checks to execute the most important statements with one call. Use the version that best fits your system environment, so that the most comprehensive set of checks is executed. It is important to know your SAP HANA revision number and if you are using a standalone or embedded statistics server. The statistics server assists you with monitoring the SAP HANA system, collects historical performance data, and warns you of system alerts (such as resource exhaustion). The historical data is stored in the _SYS_STATISTICS schema.

Note:
The new Statistics Server is also called the embedded Statistics Server or Statistics Service. Before SP7, the Statistics Server was a separate server process - like an extra Index Server with monitoring services on top of it. The new Statistics Server is now embedded in the Index Server. This simplifies the SAP HANA architecture and helps avoid out-of-memory issues in the Statistics Server. By default, the Statistics Server is set to use only 5% of the total memory. SP7 and SP8 still use the old server, but you can migrate to the new service by implementing SAP Note 1917938 .

A drag option is available to help you organize your folder structure. You can also delete queries and folders if they are not of use.

Start the mini check each day, so that you know what is going on in your system. If you plan to move your system to a newer version, you can precheck your system with the corresponding version of your target revision.

SAP HANA Mini Checks Results

CHID	DESCRIPTION	HOST	VALUE	EXPECTED_VALUE	C	SAP_NOTE
****	GENERAL					
10	Analysis date		2015/07/15 16:55:50 (CEST)			
11	Database name		SHS			
12	Revision level		100.00	>= 100.00		2021789
110	Everything started		yes	yes		2177064
111	Host startup time variation (s)		0	<= 600		2177064
115	Service startup time variation (s)	wdflbmt7195	98	<= 600		2177064
****	OPERATING SYSTEM					
208	Supported operating system	wdflbmt7195	yes	yes		
210	Minimum CPU rate (MHz)	wdflbmt7195	2493	>= 2000		1890444
211	Hosts with varying CPU rates		no	no		1890444
220	Current CPU utilization (%)	wdflbmt7195	100	<= 80	X	2100040
221	Peak CPU utilization (%, last day)	wdflbmt7195	100	<= 90	X	2100040
222	Time since CPU utilization > 95 % (h)	wdflbmt7195	0.03	>= 12.00	X	2100040
230	Current memory utilization (%)	wdflbmt7195	32	<= 90		1999997
[illegible]1	Time since memory utilization > 95 % (h)		never	>= 12.00		1999997

Figure 289: SAP HANA Mini Checks Results

To analyze the results further, you can export your results to a flat file and import the results to Microsoft Excel.

Note:
The values in the Expected_Value column are updated regularly, so it is important to import the newest version of the SQL collection occasionally.

Filtering on areas that deviate from their expected values (filter on X in column C) shows you which areas to focus on. The example in the figure shows that there is a problem with the CPU and that we should refer to SAP Note 2100040 : FAQ: SAP HANA CPU to understand the CPU consumption of SAP HANA and learn how to resolve the issue.

To Activate the JDBC Trace

1. Right click the system.
2. Choose Properties .
3. In the JDBC Trace , select the Enable trace checkbox.

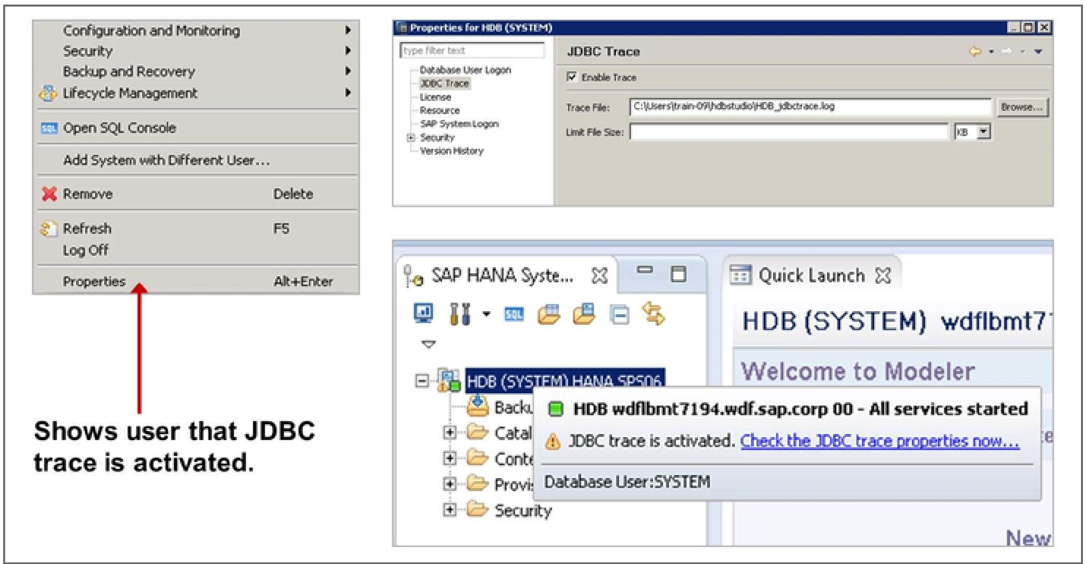

Figure 290: JDBC Trace Activation

Note:
A warning decorator and tooltip appears when the JDBC trace is activated. A message also displays on the administration overview screen.

Analyze SQL in SAP HANA Studio

Query Analysis Features

The following functions for analyzing query execution are available in the SQL console:

- Explain Plan
- Visualize Plan

The Visualize Plan is also available from the context menu of the SQL Plan Cache and *Expensive Statement* sub-tabs of the *Performance* tab.

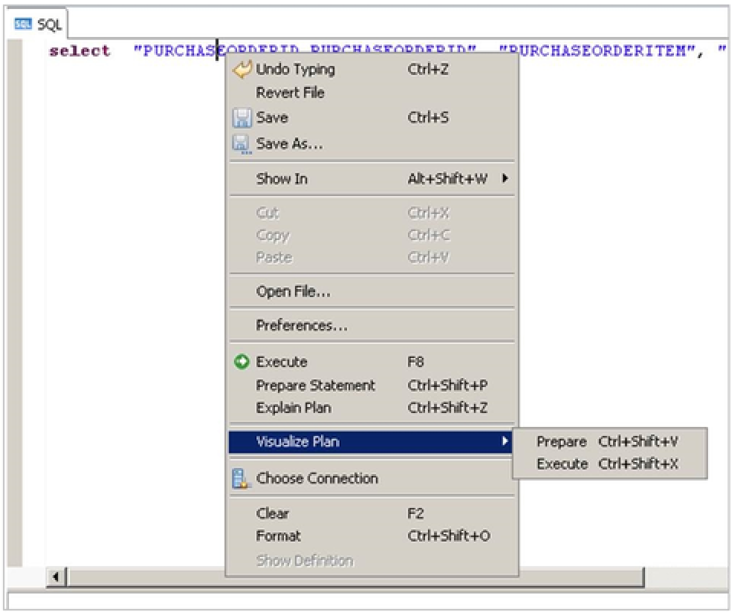

Figure 291: Query Analysis Features

Any SQL statement can be executed in the SQL editor. For SELECT statements, you can generate the explain plan. This option is available in the context menu. You can enter multiple SQL statements, each separated by the configured separator character. These are then executed one after the other.

You can change the connection of the SQL editor to a different system or user, which means that you can run the same statements on a different database. The used tables must exist in that database as well.

To help you to understand and analyze the execution plan of a SQL statement, you can generate a graphical view of the execution plan.

You can visualize the explain plan of the SQL statement in one of the following ways:

- Enter the statement in the SQL console, and choose Visualize Plan from the context menu.
- On the SQL Plan Cache tab, or on the Expensive Statements Trace tab of the Performance tab, right click the statement, and choose Visualize Plan .

Plan Visualizer

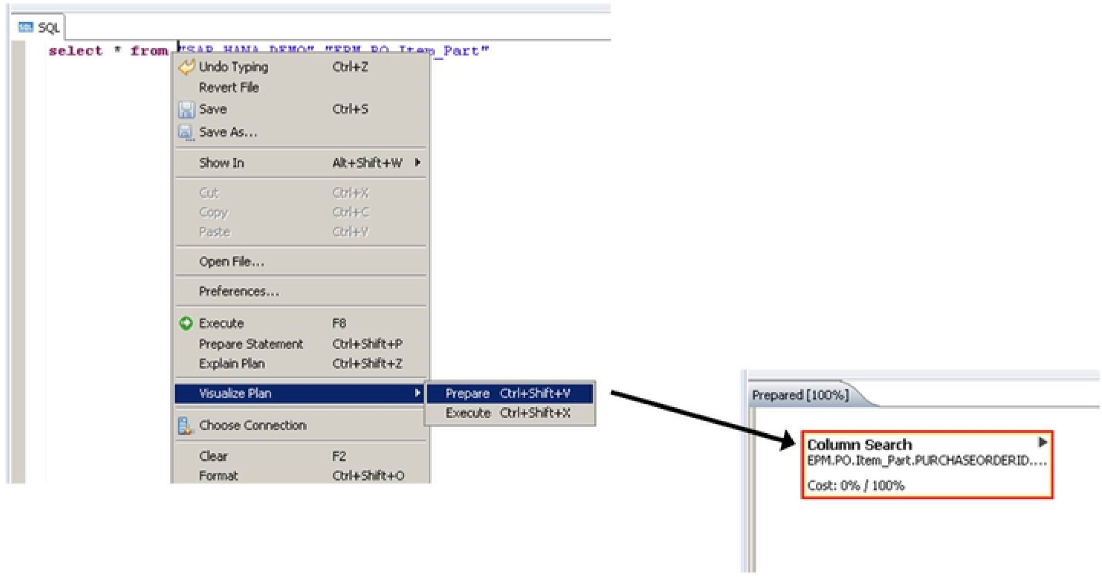

Figure 292: Visualize the Plan

Note:
Execution time is given as the following pair of values:

- Self: the execution time of the node
- Inclusive: the execution time, including the descendent nodes

If the query used the SAP HANA Column Engine, you can view the details of the various database operations by choosing Visualize Column Plan from the context menu. A detailed graphic is displayed.

This graphic is a powerful tool for studying the performance of queries on SAP HANA databases. You can explore the graphic further. For example, you can expand, collapse, or rearrange nodes on the screen. You can also save the graphic as an image or XML file, for example, so that you can submit it as part of a support query.

Execute the Plan for Validation

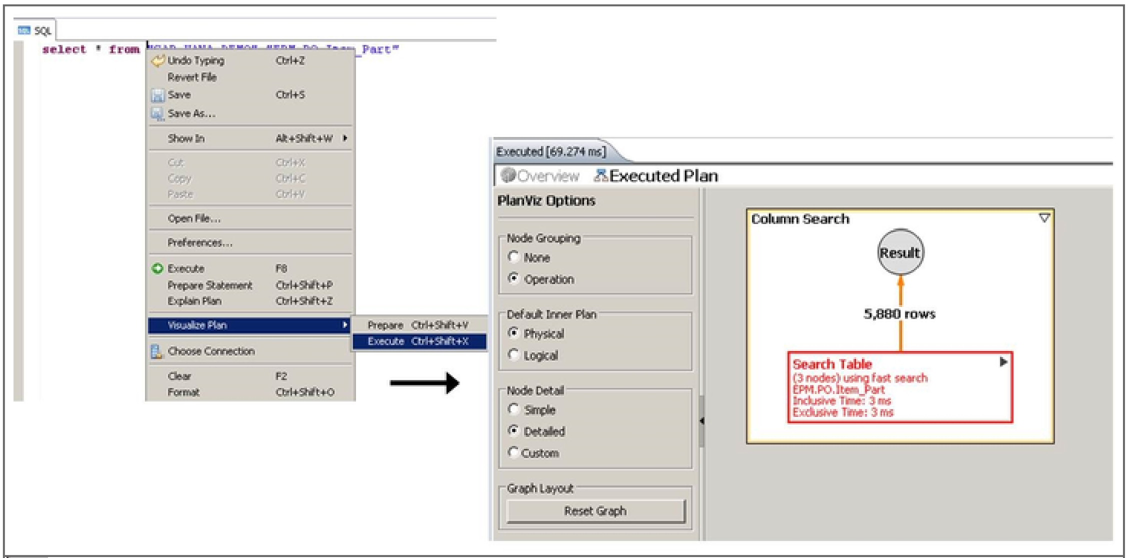

Figure 293: Execute the Plan for Validation

The figure shows how to execute the plan for validation.

Timeline View

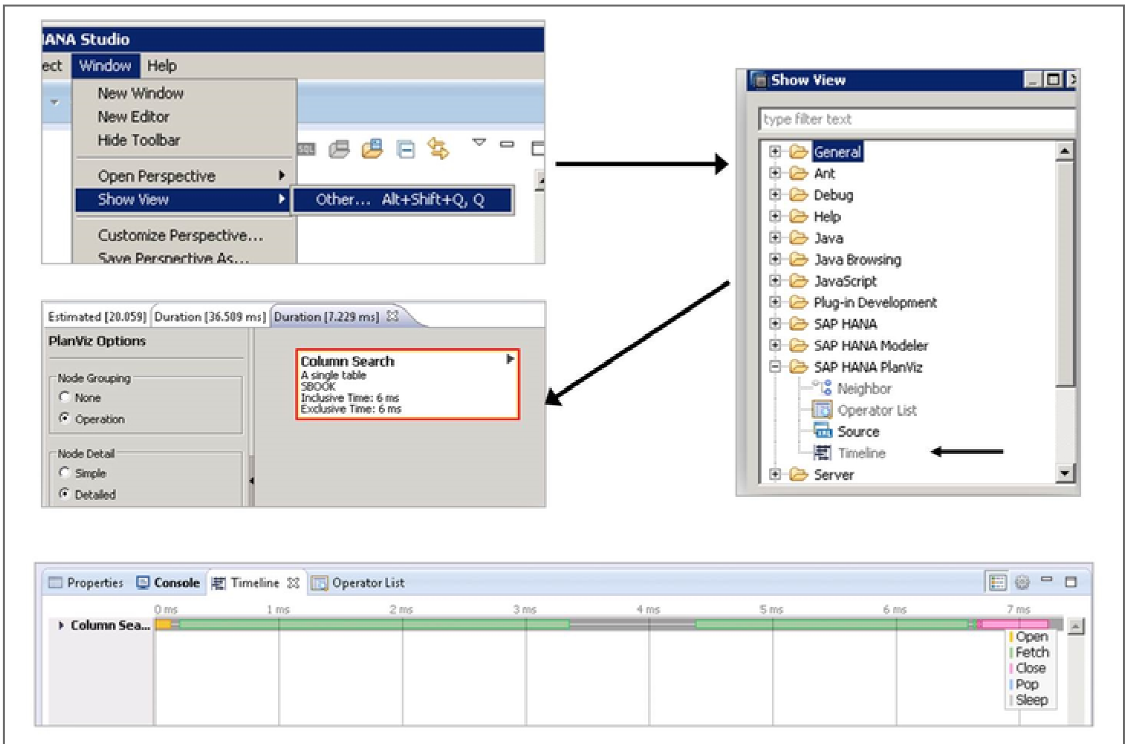

Figure 294: Timeline View

To see a temporal breakdown of the individual operations processed in the execution of the query, open the Timeline view.

Direct Access to Plan Visualizer

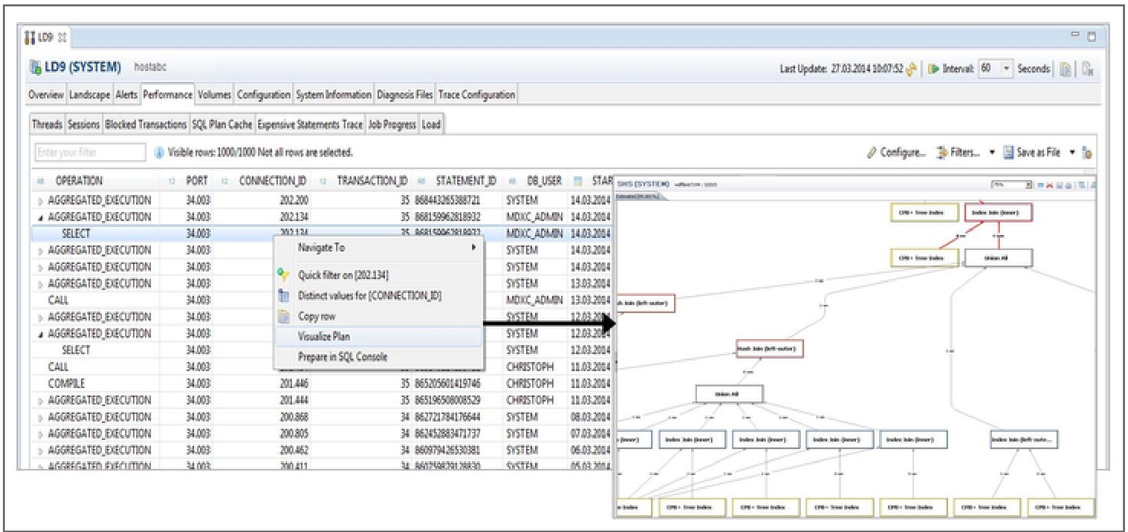

Figure 295: Direct access to Plan Visualizer

You can display the visualized plan for a specific statement in the Expensive Statements Trace or SQL Plan Cache view.

Analyzing SQL Execution with the SQL Plan Cache

The SQL plan cache collects statistics on the preparation and execution of SQL statements. Thus, it is an important tool for understanding and analyzing SQL processing. You can access the SQL plan cache in the Administration editor on the Performance tab. The two monitoring views associated with the SQL plan cache are M_SQL_PLAN_CACHE and M_SQL_PLAN_CACHE_OVERVIEW in the _SYS_STATISTICS schema.

Table 10: Useful Filtering Columns

Column	Description
TOTAL_EXECUTION_TIME	The total time spent for all executions of a plan. This helps to identify which statements are dominant in terms of time.
AVG_EXECUTION_TIME	The average time it takes to execute a plan execution. This can help you identify long-running SQL statements.
EXECUTION_COUNT	The number of times a plan has been executed. This can help you identify SQL statements that are executed more frequently than expected.
TOTAL_LOCK_WAIT_COUNT	The total number of waiting locks. This can help you identify SQL statements with high lock contention.

Column	Description
USER_NAME	The name of the user who prepared the plan and, therefore, where the SQL originated (ABAP/index-server/statistics server)

Statement Analysis

Operator List for Plan Visualizer

- The Operator List lists detailed characteristics of all operators within a current plan (visualized + executed).
- It can be used to dynamically explore the operator set along user defined filters to pinpoint specific operators of interest.
- For example you might:

1. Filter all operators consuming a certain minimal CPU time
2. Order those operators along the number of input rows
3. Further restrict the filter to a specific operator type (for example "Column Search")
4. Double-click an operator of interest to check its positioning within a visualized plan

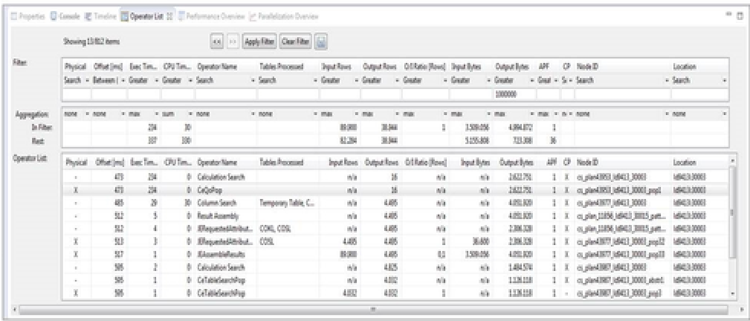

Figure 296: Statement Analysis

The view supports the display of various KPIs, including the following:

- Execution time
- CPU time
- Setting of filters along all the columns
- KPIs display the number of operators within the filtered set (top left corner) with immediate aggregated information (max, min, sum, ...)

Parameter Set

Plan Cache: Keep Set of Parameters

- The parameter set of a prepared statement is stored in monitoring view M_SQL_PLAN_CACHE_PARAMETERS to allow for the easy re-execution of a specific query:
  - First parameter set is stored when the total execution time of the statement is greater than parameter plan_cache_parameter_sum_threshold (default 1000ms).
  - Afterwards, the parameter values will be updated if the statement's single execution time is greater than parameter plan_cache_parameter_threshold (default 100ms) and its last captured execution time.
- Parameters to enable this feature:
  - plan_cache_parameter_enabled: Activate/deactivate capture of parameter values of prepared statements
  - plan_cache_parameter_for_lob_enabled: Activate/deactivate capture of BLOB /CLOB/NCLOB parameter values.

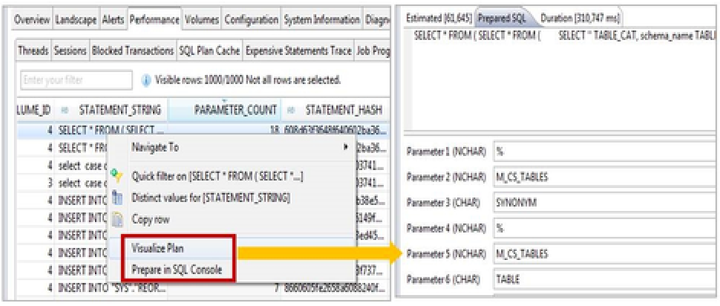

Figure 297: Parameter Set

In the Administration editor of SAP HANA Studio, on the Performance → SQL Plan Cache tab, the stored parameter set is used when you choose Visualize Plan or Prepare in SQL Console .

If a statement is removed from the SQL plan cache, its parameter information is also removed from the M_SQL_PLAN_CACHE_PARAMETERS view.

Additionally, you can use the M_SQL_PLAN_CACHE_PARAMETERS_FOR_STATISTICSSERVER_RESET monitoring view to reset the parameter list and to view, for example, hourly statistics. Use this in combination with M_SQL_PLAN_CACHE_STATISTICSERVER_RESET.

For **plan_cache_parameter_for_batch_enabled** , currently, plan cache captures the first parameter set of batch execution to reduce performance drop.

To capture all parameter sets of batch execution, turn this configuration on.

Link Between SQL Plan Cache and Expensive Statements Trace

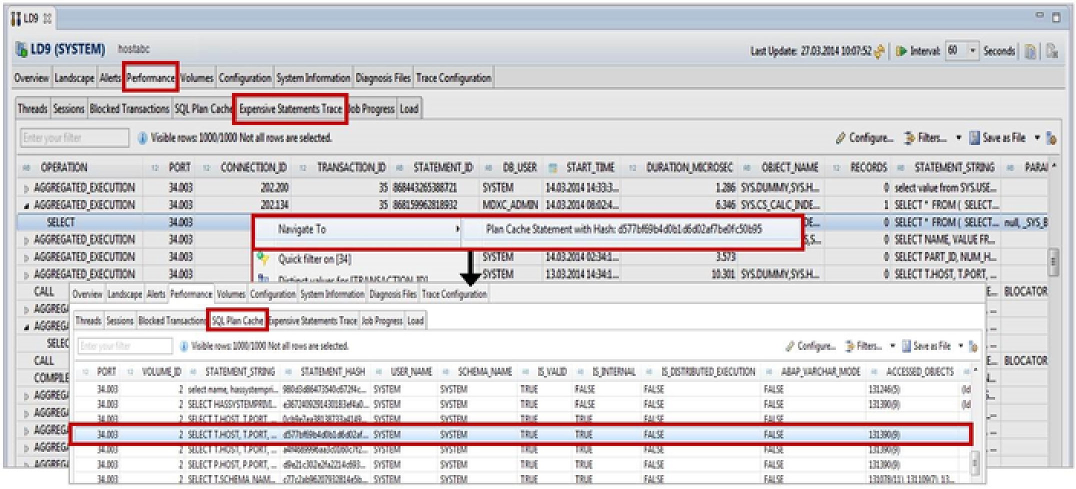

Figure 298: Link Between SQL Plan Cache and Expensive Statements Trace

Navigation between the Expensive Statements Trace and the SQL Plan Cache view has been improved.

You can display a specific statement in the other view using the context menu option Navigate To .

Execution in a Distributed System

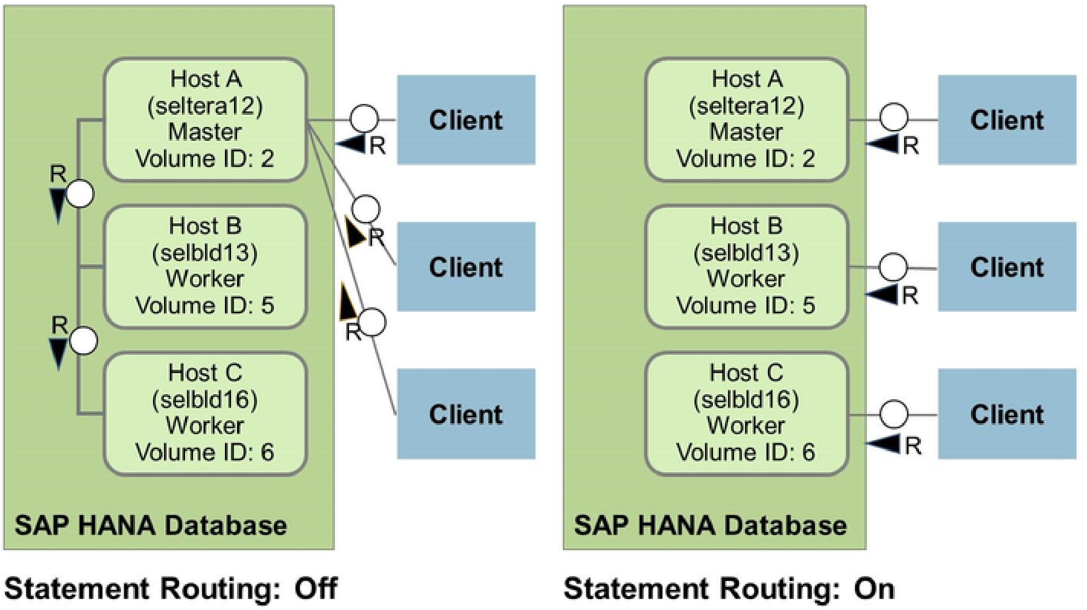

Figure 299: Execution in a Distributed System

In distributed SAP HANA systems, tables and table partitions are located on multiple hosts. You might have to execute requests received from database clients on multiple hosts, depending on where the requested data is located, as follows:

- Statement routing is not enabled

Requests from the database client are executed on the contacted index server (in this case, the master index server). The required data is fetched from the index server on the relevant host or hosts.

- Statement routing is enabled

  Request execution is routed directly to the host on which the required data is located after initial query compilation.

> Note:
> The execution time improves when statement routing is enabled. Statement routing is controlled by the client_distribution_mode parameter in the indexserver.ini file. It is enabled by default (value = statement).

LESSON SUMMARY

You should now be able to:

- Understand the administration tasks that are still in SAP HANA studio

Unit 9

Learning Assessment

1. Which of the following users are able to start and stop a distributed SAP HANA Database System

   Choose the correct answers.

   - [ ] **A** SYSTEM user
   - [ ] **B** <SID>ADM user
   - [ ] **C** ROOT user

2. Which of the following configuration files are created if customer-specific changes are made after or during installation of SAP HANA database?

   Choose the correct answers.

   - [ ] **A** sapprofile.ini
   - [ ] **B** hanaconfig.ini
   - [ ] **C** daemon.ini
   - [ ] **D** nameserver.ini

3. After the initial setup and after the initial load, you must perform a full data and file system backup, including a configuration backup.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

4. Which of the following traces has the default configuration status of “Active”?

   Choose the correct answer.

   - [ ] **A** Performance trace
   - [ ] **B** Database trace
   - [ ] **C** SQL trace
   - [ ] **D** Kernel profiler

5. The parameters “maxfiles” and “maxfilesize”, which are found in the global.ini, control the log rotation.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

6. In the Monitoring Expensive Statements trace, which additional trace option needs to be set to “false” to write the trace data directly to file.

   Choose the correct answer.

   - [ ] **A** Threshold memory
   - [ ] **B** Use in-memory tracing
   - [ ] **C** Trace flush interval
   - [ ] **D** In-memory tracing records

7. In which circumstances do you create a column store table?

   Choose the correct answers.

   - [ ] **A** When the calculations are based on a small number of columns
   - [ ] **B** When accessing complete records
   - [ ] **C** When aggregations or fast search is required
   - [ ] **D** When the table has a large number of columns

8. A delivery unit is a collection of packages that are transported together. Assign all the packages belonging to your application to the same delivery unit. This ensures that they are transported consistently together within your system landscape. Each delivery unit has a unique identity.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

Learning Assessment - Answers

1. Which of the following users are able to start and stop a distributed SAP HANA Database System

   Choose the correct answers.

   [ ] **A** SYSTEM user

   [X] **B** <SID>ADM user

   [X] **C** ROOT user

   Correct! To start and stop an SAP HANA system, you require the credentials of the operating system user (<sid>adm) that were created when the system was installed. Alternatively, root users can start and stop SAP HANA. The system user is the database superuser. Read more on this in the lesson Starting and Stopping SAP HANA (Unit 9, Lesson 1) of the course HA200_14.

2. Which of the following configuration files are created if customer-specific changes are made after or during installation of SAP HANA database?

   Choose the correct answers.

   [X] **A** sapprofile.ini

   [ ] **B** hanaconfig.ini

   [X] **C** daemon.ini

   [X] **D** nameserver.ini

   Correct! sapprofile.ini contains system identification information, such as the system name (SID) or the instance number. daemon.ini contains information about which database services to start. nameserver.ini contains global information for each installation. The landscape section contains the system-specific landscape ID and assignments of hosts to roles MASTER, WORKER, and STANDBY. There is no configuration file named hanaconfig.ini. Read more on this in the lesson Configuring the SAP HANA Database (Unit 9, Lesson 2) of the course HA200_14.

3. After the initial setup and after the initial load, you must perform a full data and file system backup, including a configuration backup.

   Determine whether this statement is true or false.

   [X] True

   [ ] False

   Correct! This to safeguard the changes that you made to the database data and configuration. It is recommended that you perform a data backup after the initial load, at regular intervals, before the database software is upgraded to a new version, and after any situation that causes log writing to be interrupted. Read more on this in the lesson Performing Regular Database Administration Tasks (Unit 9, Lesson 3) of the course HA200_14.

4. Which of the following traces has the default configuration status of “Active”?

   Choose the correct answer.

   [ ] **A** Performance trace

   [X] **B** Database trace

   [ ] **C** SQL trace

   [ ] **D** Kernel profiler

   Correct! The database trace records information about activity in the components of the SAP HANA database. Use this information to analyze performance and to diagnose and debug errors. Each service of the SAP HANA database writes to its own trace file. By default, the database trace is active with the default trace level ERROR. The performance trace is a performance tracing tool built into the SAP HANA database. It records performance indicators for individual query processing steps in the database kernel. By default, the performance trace is inactive. The SQL trace collects information about all executed SQL statements and saves it as an executable python program. This is useful for recording a scenario. By default, the SQL trace is inactive. The kernel profiler is a sampling profiler built into the SAP HANA database. It collects, for example, information about frequent and expensive paths during query processing. By default, the kernel profiler is inactive. Read more on this in the lesson Configuring Traces (Unit 9, Lesson 4) of the course HA200_14.

5. The parameters “maxfiles” and “maxfilesize”, which are found in the global.ini, control the log rotation.

   Determine whether this statement is true or false.

   [X] True

   [ ] False

   Correct! Trace file rotation prevents trace files from growing indefinitely by limiting the size and number of trace files. You can configure trace file rotation globally for all services in the database and for individual services. For this you need the system privilege INIFILE ADMIN. Configure“maxfiles” by specifying the maximum number of trace files that may exist and “maxfilesize” by specifying in bytes the maximum size an individual trace file may reach. Read more on this in the lesson Working with Diagnosis Information and Diagnosis Files (Unit 9, Lesson 5) of the course HA200_14.

6. In the Monitoring Expensive Statements trace, which additional trace option needs to be set to “false” to write the trace data directly to file.

   Choose the correct answer.

   [ ] **A** Threshold memory

   [X] **B** Use in-memory tracing

   [ ] **C** Trace flush interval

   [ ] **D** In-memory tracing records

   Correct! In-memory tracing is active information cached in memory. Otherwise, data is written directly to file. Threshold memory usage of statement is executed in bytes. When set to 0, all SQL statements are traced. Trace flush interval is the number of records after which trace file is flushed. In-memory tracing records is the maximum number of trace records (per service) stored in memory This setting only takes effect when in-memory tracing is active. Read more on this in the lesson Using the SQL Console (Unit 9, Lesson 6) of the course HA200_14.

7. In which circumstances do you create a column store table?

   Choose the correct answers.

   [X] **A** When the calculations are based on a small number of columns

   [ ] **B** When accessing complete records

   [X] **C** When aggregations or fast search is required

   [X] **D** When the table has a large number of columns

   Correct! When each column acts as an individual table, each of these individual mini-tables can be indexed (=sorted) and compressed (=process of removing duplicates). This makes sure that each of these tables only contains a unique entry. In case of analytic applications where aggregations are used and fast search and processing is required, row-based storage is inefficient. In row based tables all data in a row has to be read even though the requirement may be to access data from a few columns. Hence these queries on huge amounts of data take a lot of time. When multiple columns need to be searched or aggregated, each of these operations can be assigned to a different processor core. When the application needs to only process a single record at one time (many selects and/or updates of single records) row store is more efficient. Read more on this in the lesson Performing SAP HANA Table Administration (Unit 9, Lesson 7) of the course HA200_14.

8. A delivery unit is a collection of packages that are transported together. Assign all the packages belonging to your application to the same delivery unit. This ensures that they are transported consistently together within your system landscape. Each delivery unit has a unique identity.

   Determine whether this statement is true or false.

   [X] True

   [ ] False

   Correct! The identity of a delivery unit consists of two parts: a vendor name, and a delivery-unit name. The combined ID ensures that delivery units from different vendors are distinguished easily. It also ensures that they follow a pattern that SAP uses for various software components. To create and manage delivery units, you first need to maintain the identity of the vendor. The delivery units are associated with the vendor, and the packages that make up the delivery unit are stored in the vendor's namespace. The delivery unit is a collection of packages to be transported together. This helps to ensure consistent transports of all packages of one application. It is uniquely identified by its name and the vendor name. Read more on this in the lesson Transporting Changes (Unit 9, Lesson 8) of the course HA200_14.

UNIT 10 Backup and Recovery

UNIT OBJECTIVES

- Explain backup and recovery
- Perform data area backup
- Configure a log area backup
- Describe additional backup topics
- Perform database recovery
- Explain backup and recovery using data snapshots
- Explain the scenarios for a database copy

Unit 10
Lesson 1

Explaining Backup and Recovery

LESSON OVERVIEW

Business Example

You have to perform backups for the SAP HANA database. Therefore, you need to know the backup and recovery concept of the SAP HANA database.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explain backup and recovery

SAP HANA Persistence

To ensure optimal performance, the SAP HANA database holds most of its data in-memory. However, it still uses persistent storage to provide a fallback in case of failure.

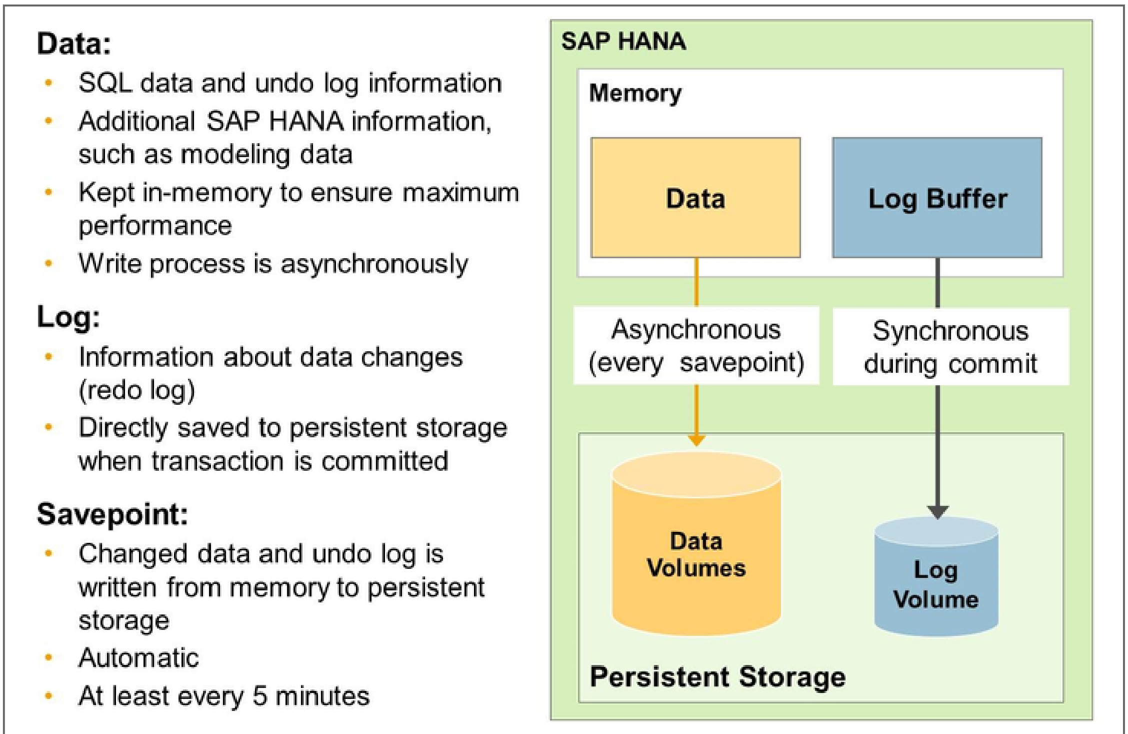

Figure 300: SAP HANA Persistence

During normal database operation, data is automatically saved from memory to disk at regular savepoints. Additionally, all data changes are recorded in the redo log. The redo log is saved from memory to disk with each committed database transaction. After a power failure, you can restart the database as you would with any disk-based database. It then returns to its last consistent state by replaying the redo log since the last savepoint.

Although savepoints and log writing protect your data against power failures, savepoints do not help if the persistent storage itself is damaged. To protect against data loss because of disk failures, backups are required. Backups save the payload (the actual data) of the data area and log area to different locations. Unused space in the database is not backed up.

The data backup includes all the data structures that are required to restore the database. This includes user data, information models, topology information, and the secure storage file system (SSFS). A data backup does not include customer-specific configuration.

Overview of Backup and Recovery

Backups are performed while the database is running. The impact of backups on system performance is negligible, and users can continue to work while the backup is running.

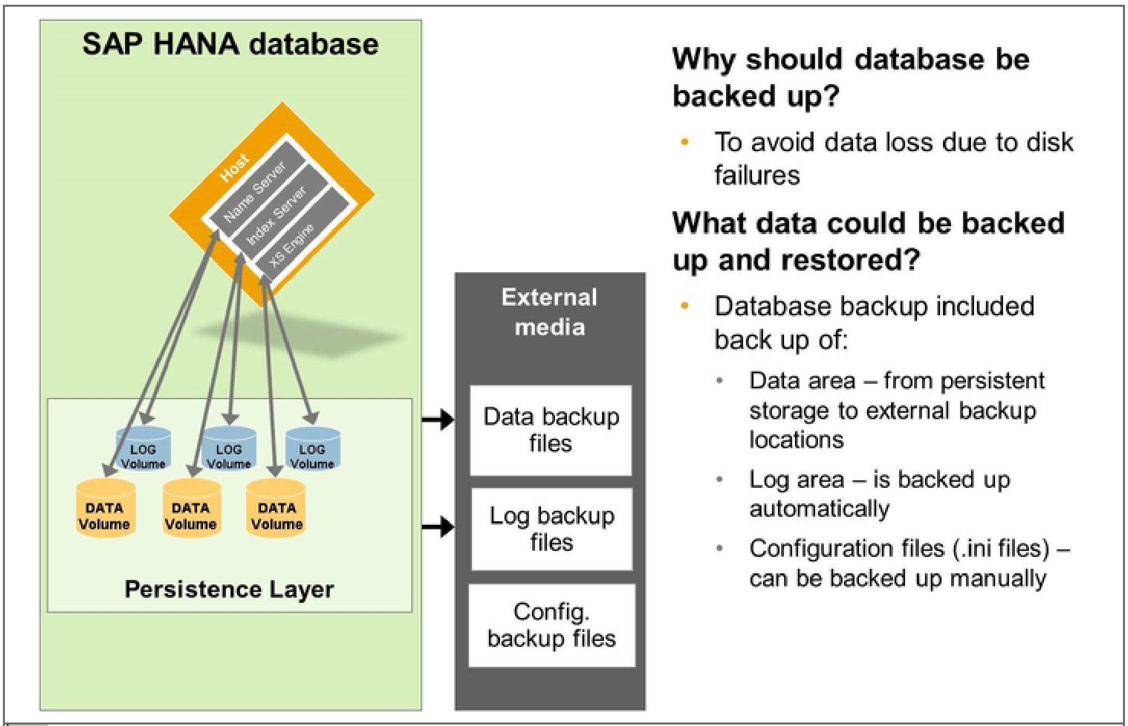

Figure 301: Overview of Backup and Recovery

The data area is backed up in parallel for each of the SAP HANA services. If SAP HANA is running on multiple hosts, a data backup includes all the service-specific backup parts for all the hosts.

While a data backup is running, some data integrity checks are performed. If these checks are successful, the data is written to the backup destination.

Data backups save the content of the data area to a different location in the file system. Depending on the scenario, this includes the replicated business data from SAP ERP and all the modeling data.

Overview

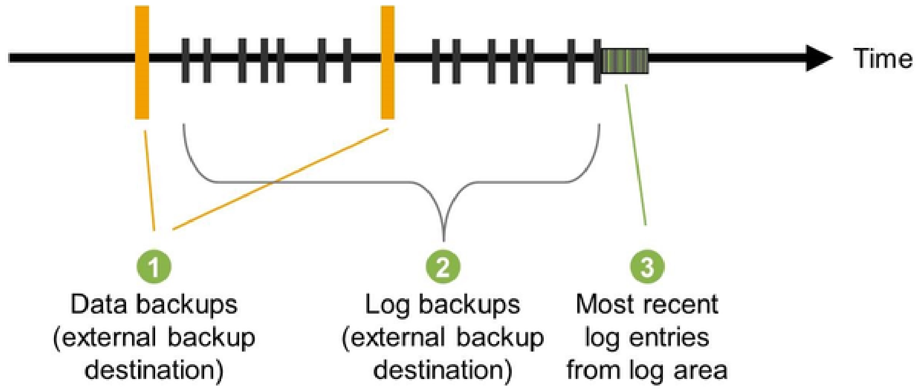

Figure 302: Overview

The properties of an SAP HANA system are defined in the parameters of its configuration files. These files are not backed up as part of the database backup. If you want to back up configuration files that contain customer-specific changes, you can do so manually.

In a recovery situation, configuration files can help to identify and restore the customer-specific changes. The configuration files are not essential to perform a recovery. If you want to use a customer-specific configuration, reconfigure the recovered system using the SAP HANA studio.

Overview of Backup and Recovery Capabilities

SAP HANA supports the following backup and recovery capabilities:

- Full backup
- Delta backup
- Redo log backups
- Backup and recovery using third-party tools
- Integrity checks for backups
- Backup lifecycle management
- Recovery to the most recent state
- Recovery to a specific point-in-time
- Recovery to a specific data backup
- Database copy using backup and recovery

Performing Backups

Backups can be created using the following tools:

- Creating data backups manually:
  - SAP HANA cockpit

  - SAP HANA studio
  - DBA Cockpit
  - SQL commands (command line)
- Scheduling data backups:
  - Backup scheduler in SAP HANA cockpit
  - Planning calendar in DBA Cockpit
  - Using scripts (via the SQL interface)

Destinations for Backups

You can specify whether data and log backups are written to the file system (see SAP Note 1820529), or using third-party backup tools (see SAP Note 1730932). The BACKINT software development kit (SDK) for the SAP HANA interface performs all the actions needed to write the backup data to external storage. The backup tools communicate directly with the SAP HANA database through the BACKINT SDK for the SAP HANA interface.

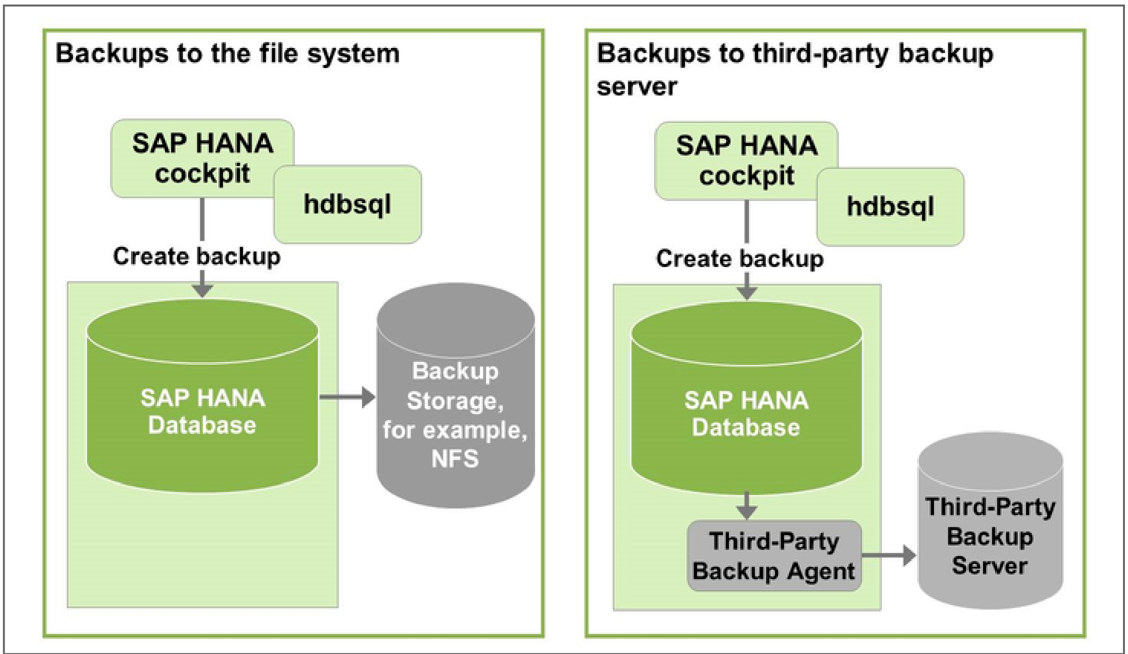

Figure 303: Destinations for Backups

BACKINT SDK for SAP HANA

BACKINT SDK for SAP HANA is an application programming interface (API) that can be implemented by a third-party backup agent. It has the following features:

- It provides functions for backup, recovery, query, and delete.
- The third-party backup agent runs on the SAP HANA server and communicates with the third-party backup server.
- Backups are transferred through pipes.
- It has full integration with SAP HANA studio (configuration and execution of backups to BACKINT).

- It can be configured for data backups and for log backups.

> Note:
> SAP certification is required for BACKINT SDK for SAP HANA implementations by third-party vendors.

The default configuration is defined when a third-party backup tool is installed. After a backup tool has been installed, you can back up and recover the SAP HANA database without making any further changes.

Backup of Multitenant Database Containers

The usual SAP HANA backup and recovery principles apply for multitenant database containers. The system database and all tenant databases create their own backup.

Overview: Backup of Multitenant Database Containers

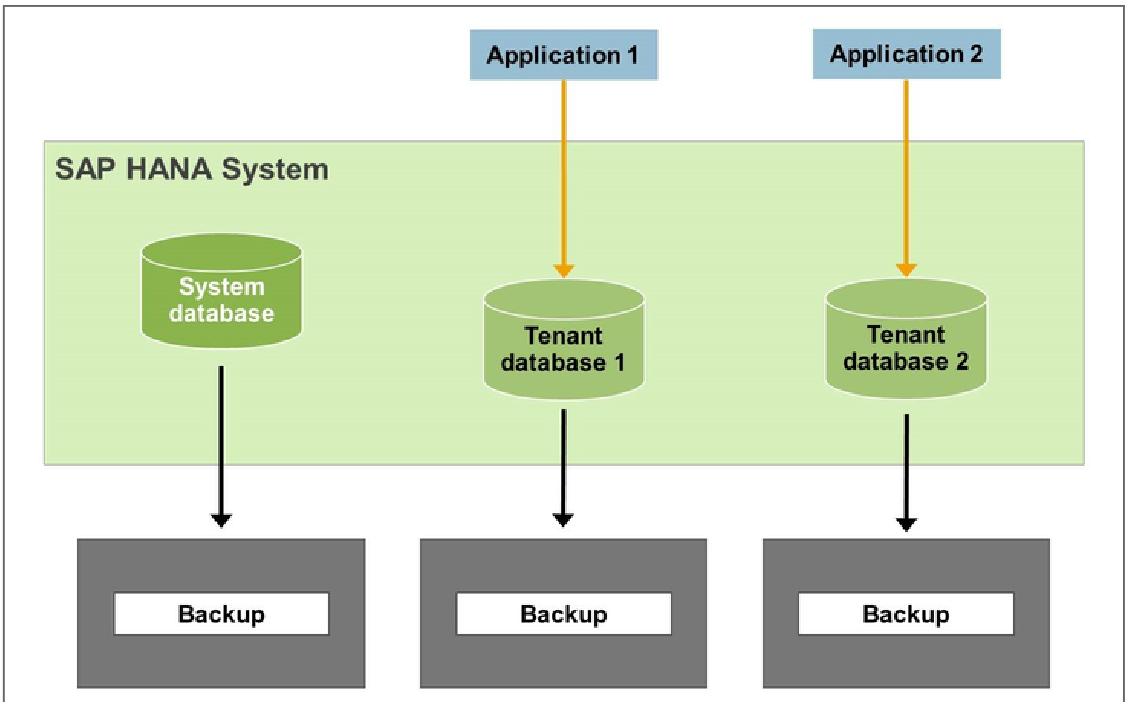

Figure 304: Overview: Backup of Multitenant Database Containers

Backup and Recovery Strategy

This is an overview of information to consider when planning your backup and recovery strategy with SAP HANA database. You can find more information on the individual points in the following sections.

Points to Note

- Backups can only be created when SAP HANA is online. All the configured SAP HANA services must be running.
- While backups are being created, the impact on system performance is negligible, and users can continue to work normally.

- With a data backup, only the actual data is backed up; unused space in the database is not backed up. A full data backup includes all the data that is required to recover the database to a consistent state. This includes both business data and administrative data.
- The system database plays a central role. It can initiate backups of both the system database itself and of individual tenant databases.
- A tenant database can create its own backups without the need to connect through the system database.
- System database and tenant databases have their own backup catalogs.
- The data backup reflects the consistent database state from the time at which the data backup was started.
- Backup and recovery always applies to the whole database. You cannot back up and recover individual database objects.

Release Compatibility of SAP HANA Backups

In some situations, backups from earlier SAP HANA releases can be used for a recovery.

- SAP HANA backups created with release 1.0 SPS10 or newer can be used to recover to SAP HANA 2.0.
- A backup of an SAP HANA single-container system can only be recovered to a tenant database.
- A backup of an SAP HANA single-container system cannot be recovered to a system database.

Privileges for Backup and Recovery

Backup and recovery operations can only be performed by users that have the appropriate authorizations. In SAP HANA multitenant database containers, the required authorization depends on whether administrative tasks are performed at system level or at database level.

Table 11: Required Authorizations

To perform operations related to backup and recovery, the following authorizations are required:

Task	**Required authorizations**
Back up SAP HANA using SAP HANA cockpit or SAP HANA studio	BACKUP ADMIN or BACKUP OPERATOR
	CATALOG READ This privilege is required to collect the information needed by the backup wizard
Back up the database without a user interface	BACKUP ADMIN or BACKUP OPERATOR (recommended for batch users only)
Recover the database without a user interface	This is supported for an SAP HANA single container system or the system database in an SAP HANA multitenant database container. The recovery is executed as the operating system user (<sid>adm). You therefore require the logon credentials of this user.

Task	Required authorizations
Physically delete data, and log backups and obsolete versions of the backup catalog from the backup location	BACKUP ADMIN
Administration tasks executed on a tenant database through the system database	DATABASE ADMIN

Difference Between the BACKUP ADMIN and BACKUP OPERATOR

What is the difference between the BACKUP ADMIN and BACKUP OPERATOR?

The BACKUP ADMIN and BACKUP OPERATOR system privileges exist so that you can implement a clearer separation of duties, if this is necessary in your organization.

A user with the BACKUP ADMIN system privilege can perform all of the backup-related operations, including backup deletion and configuration.

A user with the BACKUP OPERATOR system privilege can only perform backups.

For example, if you have automated the regular performance of backups using Cron, it is more secure to use a user with the BACKUP OPERATOR privilege to avoid the malicious deletion of backups.

Related Information

For more information, see the following:

- SAP HANA documentation
  - SAP Help Portal: http://help.sap.com/hana_appliance
  - SAP HANA Administration Guide, chapter: "Backing Up and Recovering the SAP HANA Database"
  - SAP HANA Technical Operations Manual
- SAP Notes
  - SAP Note 1642148 : FAQ: SAP HANA database backup and recovery
  - SAP Note 1730932 : Using backup tools with BACKINT
  - SAP Note 1812980 : Changes to the backup catalog as of revision 45
  - For further notes on backup and recovery, see BC-DB-HDB-BAC
- BACKINT for SAP HANA certification
  - Certification announcement and description

LESSON SUMMARY

You should now be able to:

- Explain backup and recovery

Unit 10
Lesson 2

Performing Data Area Backup

LESSON OVERVIEW
The goal of this lesson is to understand how the data backup works.

Business Example

You have to define a backup strategy for your SAP HANA database. Therefore, you need to know how to perform data area backups.

LESSON OBJECTIVES
After completing this lesson, you will be able to:

- Perform data area backup

Overview of Data Area Backup

The following principles for data backups apply for an individual database in a multitenant database container system.

A data backup includes all the data that is required to recover the database to a consistent state.

With a data backup, only the actual data is backed up; unused space in the database is not backed up.

> Note:
> A data backup does not include the log area or customer-specific configuration settings.

The data area is backed up in parallel for each of the SAP HANA services. If SAP HANA is running on multiple hosts, a data backup includes all the service-specific backup parts for all the hosts.

While a data backup is running, some data integrity checks are performed. If these checks are successful, the data is written to the backup destination.

The payload of the data area can be backed up by performing a complete data backup or a delta backup. Delta backups contain data that changed since the last complete data backup.

Delta backups allow you to reduce the amount of data that is backed up, compared to full data backups. In turn, this means that delta backups are faster to create than full data backups.

Types of Delta Backups

Two types of delta backups are available, as follows:

Differential

It stores all the data that changed since the last full data backup.

The amount of data to be saved with each differential backup increases.

Incremental

It stores the data that changed since the last full data backup or the last delta backup (incremental or differential).

If data remains unchanged, it is not saved to more than one backup. For this reason, incremental backups are the smallest of the backup types.

Overview Data Area Backup

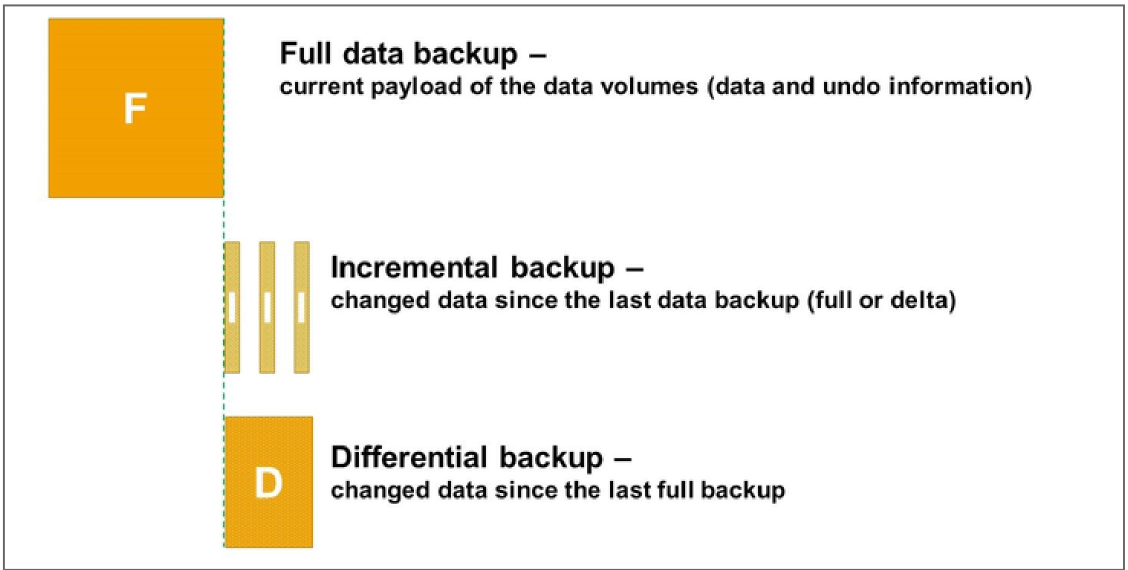

Figure 305: Overview Data Area Backup

Note:
Delta backups are data backups. In contrast to delta backups, log backups contain the redo log entries of a closed log segment.

The type of delta backup to use depends on your specific backup and recovery requirements. You can also mix incremental and differential backups.

Note that, in terms of backup and recovery, “changed data” is related to the physical layout of the data in the SAP HANA persistent storage. This does not always correlate to the amount of data actually changed.

For example, a delta merge of a column store partition does not change the content of an SAP HANA database at all. However, it recreates the whole partition for optimized read access of the data. In this case, the whole partition is backed-up in a delta data backup, even if no data is changed.

Note:
SAP HANA supports seamless integration of SAP-certified third-party backup tools. Normally, delta backups work with the default configuration settings. However, in some situations, additional steps might be required to create delta backups with a third-party backup tool.

Overview Complete Data Area Backup

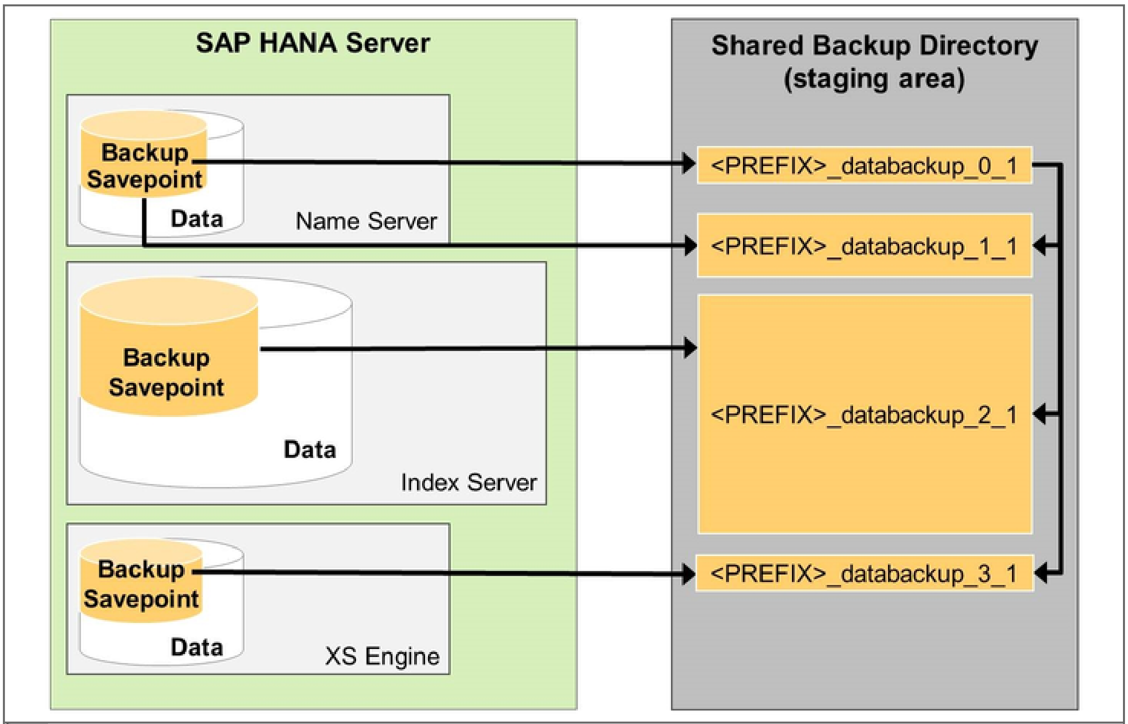

Figure 306: Overview Complete Data Area Backup (Tenant Database)

When the data area is backed up, all the payload data from all the servers is backed up. This happens in both single-host and multihost environments.

File-Based Backups

Default location: Parameter `basepath_databackup`

global.ini

persistence	
basepath_databackup	$(DIR_INSTANCE)/backup/data
basepath_datavolumes	$(DIR_GLOBAL)/hdb/data
basepath_logbackup	$(DIR_INSTANCE)/backup/log
basepath_logvolumes	$(DIR_GLOBAL)/hdb/log

Structure of file names for Backups:

Full Backup: <prefix>_databackup_<suffix>

Differential: <prefix>_databackup_differential_<suffix>

Incremental: <prefix>_databackup_incremental_<suffix>

Rules for the location of the backup files:

- Should never be in the same directory and on the same file system as the data
- The folder must already exist before the backup process is started
- The Volume ID is automatically added to a specified file name
- The backup ID is unique for each persistence only

Figure 307: Location of Backup Files

The data backup files are written to the location specified by the basepath_databackup parameter in the persistence section of the global.ini configuration file. By default, the location for data backup files is `$(DIR_INSTANCE)/backup/data`.

To use a different location, specify a different path when you perform the backup. If you need to, you can specify a different path for each backup. Alternatively, you can change the value of the basepath_databackup parameter in global.ini → persistence.

If you change the backup location in basepath_databackup, the change occurs immediately.

For improved data safety, specify a path to an external backup location. The backup location should never be on the same file system as the data or log areas.

Note that all the files for a particular data backup are written to the same location. The files belonging to the same data backup cannot be written to multiple locations. Different data backups can be written to different locations, but all the files belonging to one particular data backup are written to the same location. Create the directory structures before the backup is started.

Note the following information for file-based backups:

- The configured destination for data and log backups must be valid throughout the whole system, not only for specific hosts.
- Use shared backup storage to make the backup area available to all the nodes in a database.

  Shared backup storage allows the master name server to perform availability checks for file-based backups at the beginning of the recovery.

In addition, shared storage offers support for database copy.

- The backup location in the file system is specified system-wide. Backups of tenant databases are always created in subdirectories of this location.

> Note:
> The default backup destination can only be changed for file-based backups. Backups made using third-party tools always use the destination `/usr/sap/<SID>/SYS/global/hdb/backint`. Because of this, you cannot change the backup destination for third-party tools.

Elements of Backup Files

Each backup file name contains the following elements: `<<path>/<prefix>_<string>_<suffix>>`. These elements are described as follows:

- The `<path>` is optional. If no complete path is specified, the default backup location is used.
- You can specify a `<prefix>` for the backup file name, or you can use the prefix proposed by the system.
- The `<string>` defines the type of the data backup.
- The system adds a unique `<suffix>` to each backup file name that indicates the volume ID and the partition ID. Because this is done for each service that is included in the backup, you only need to specify one file name prefix for all the backups on the different hosts.

  For delta backups the suffix contains also the backup ID of the backup that the delta backup is based on.

  The suffix that is appended to a file name prefix is only unique for each service. So, the next time you back up a service, the system assigns the same backup suffix to the backup file for that service. If you do not change the file name, the existing backup file for that service is overwritten by the new backup.

During the backup process, a backup file for each service is created in the backup location.

> Note:
> For file-based backups, use a unique prefix for each data backup name. For example, a timestamp.
>
> If you use the same prefixes, then replicate a data backup to a new destination as soon as the backup is created. Otherwise, an existing complete data backup with the same name will be overwritten by the next data backup.

The configuration of backup settings (for example, third-party backup tool integration, backup destination paths, log backup settings) is available in the Backup Editor of SAP HANA Studio.

Overview of Configuration Options

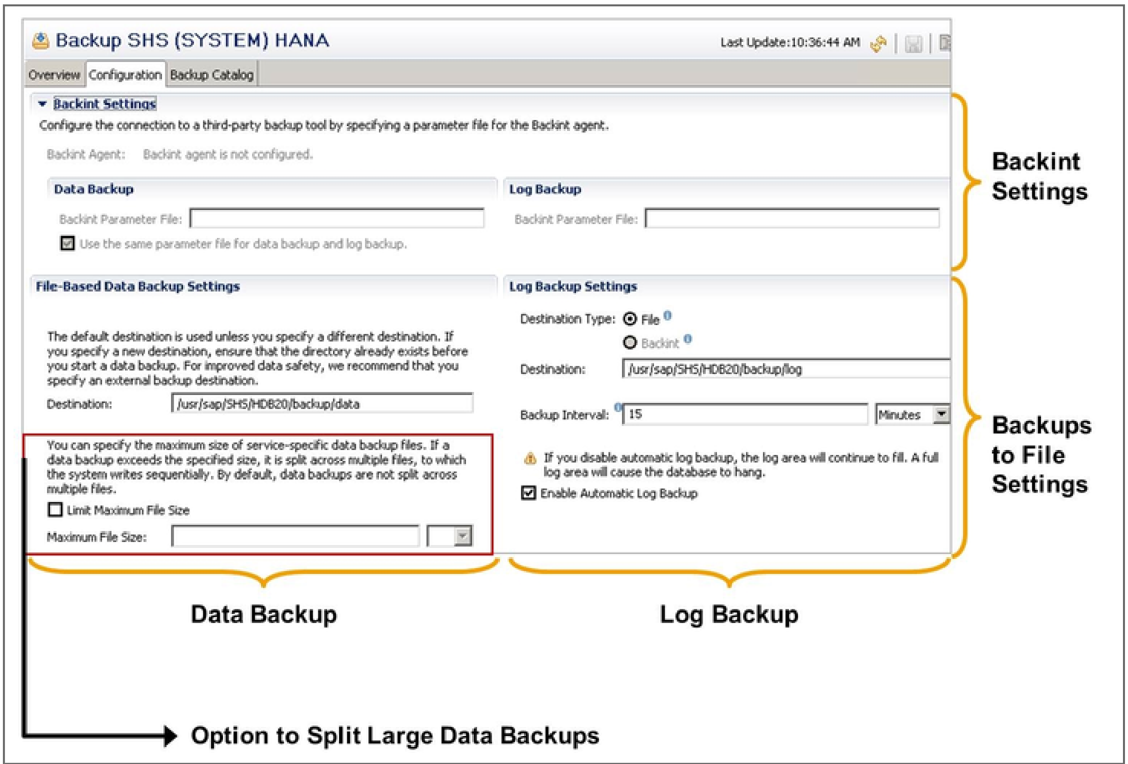

Figure 308: Overview of Configuration Options in SAP HANA Studio

In large SAP HANA systems, data backup files might be larger than the maximum file size that can be stored on the respective file system. The configuration options allow you to specify the maximum file size for backup files. If a backup exceeds this size, it is split into several files.

Estimation of the Size of a Backup

Ensure that sufficient free space is available in the file system for backups. If there is not enough space, the backup fails. For this reason, before you back up the database, estimate the amount of space that is needed in the backup destination.

Estimation of the backup size

- Estimate the size of the backup

  `select * from M_BACKUP_SIZE_ESTIMATIONS`

SQL | Result

```
select * from M_BACKUP_SIZE_ESTIMATIONS
```

	HOST	PORT	SERVICE_NAME	ENTRY_TYPE_NAME	ESTIMATED_SIZE
1	wdflbmt7194	32,003	indexserver	complete data backup	2,862,178,304
2	wdflbmt7194	32,003	indexserver	differential data backup	107,175,936
3	wdflbmt7194	32,003	indexserver	incremental data bac...	107,175,936
4	wdflbmt7194	32,007	xsengine	complete data backup	68,423,680
5	wdflbmt7194	32,007	xsengine	differential data backup	262,144
6	wdflbmt7194	32,007	xsengine	incremental data bac...	262,144

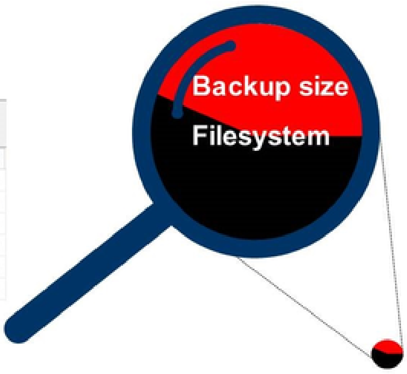

- Estimate the space required for a backup
- BACKUP CHECK USING FILE <data_file_definition> SIZE <size>

SQL

```
BACKUP CHECK USING FILE ('/usr/sap/SHS/HDB20/backup/data') SIZE 2862370816
```

```
Statement 'BACKUP CHECK USING FILE ('/usr/sap/SHS/HDB20/backup/data') SIZE 2862370816'
successfully executed in 0 ms 689 µs  (server processing time: 0 ms 192 µs) - Rows Affected: 0
```

Figure 309: Estimation of the Size of a Backup

The Administrator ensures that sufficient free space for the backup files is available. They calculate the amount of free space that is needed in the backup directory.

To estimate the size of a backup, you can use the system table M_BACKUP_SIZE_ESTIMATIONS in the SQL Editor in the SAP HANA studio. This system table contains information about the used blocks.

> Note:
> The actual size of a data backup can be larger or smaller than the estimated size. For example, if data is changed in the database after the size has been estimated and before the backup is performed, the actual backup size may be different from the estimated size.
>
> It is therefore recommended to keep some additional free space in reserve.

> Hint:
> The more difficult part is the sizing for log backups, because this depends on the amount of data changes that occur in the database. This in turn is a unique quantity for each system and timeframe.
>
> When loading data, the experiences shows that the disk size of log entries is typically at least twice the size of the loaded data after compression in SAP HANA.

Performing Backups Using SAP HANA Cockpit

Backup of a Multitenant Database Container System

Using SAP HANA cockpit, you can create data backups and delta backups (differential backups and incremental backups). You need to create a backup for the system database and all tenant databases.

Data backups of the system database are needed on a regularly basis.

The system database contains information about the system as a whole and all tenant databases. It is used for central system administration.

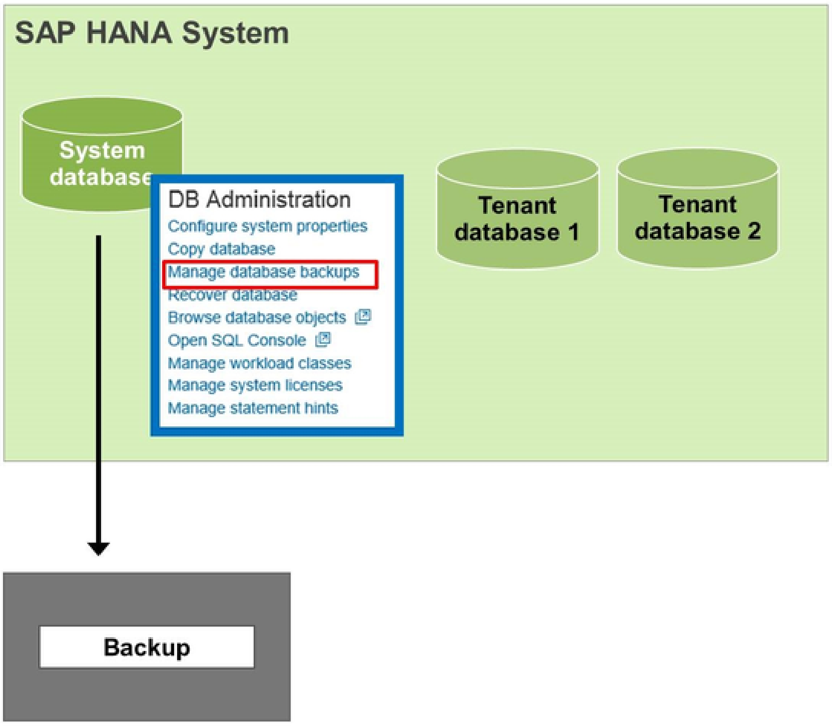

Figure 310: Multitenant Database Containers: Backing Up the System Database

You can perform a data backup of the system database with SAP HANA cockpit. Select the system database in the resources list of SAP HANA cockpit and choose DB Administration → Manage database backups. Then, specify your backup settings and start the backup.

Because data backups of the system database only contain information about the system as a whole, data backups of the tenant databases are also needed on a regularly basis.

The tenant databases contain the business data. They have their own index servers.

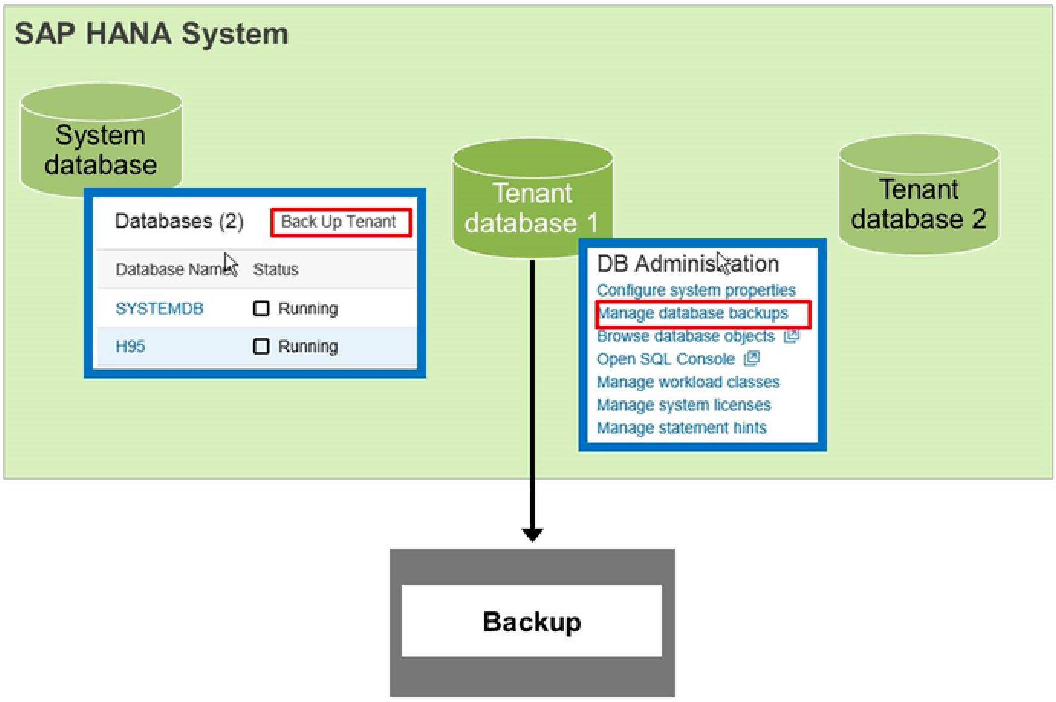

Figure 311: Multitenant Database Containers: Backing Up a Tenant Database

You can perform a data backup of a tenant database with SAP HANA cockpit. Select the system database in the resources list of SAP HANA cockpit and choose Overall Tenant Statuses. Then select the tenant database on the Manage Databases screen and choose Back Up Tenant. Then, specify your backup settings and start the backup.

Depending on the system configuration, you can also initiate a data backup directly from a tenant database. Select the tenant database in the resources list of SAP HANA cockpit and choose DB Administration → Manage database backups. Then, specify your backup settings and start the backup.

Performing Backups Using SAP HANA Cockpit

If a backup is started, the backup wizard also shows the estimated backup size. See the figure, Performing Backups Using SAP HANA Cockpit.

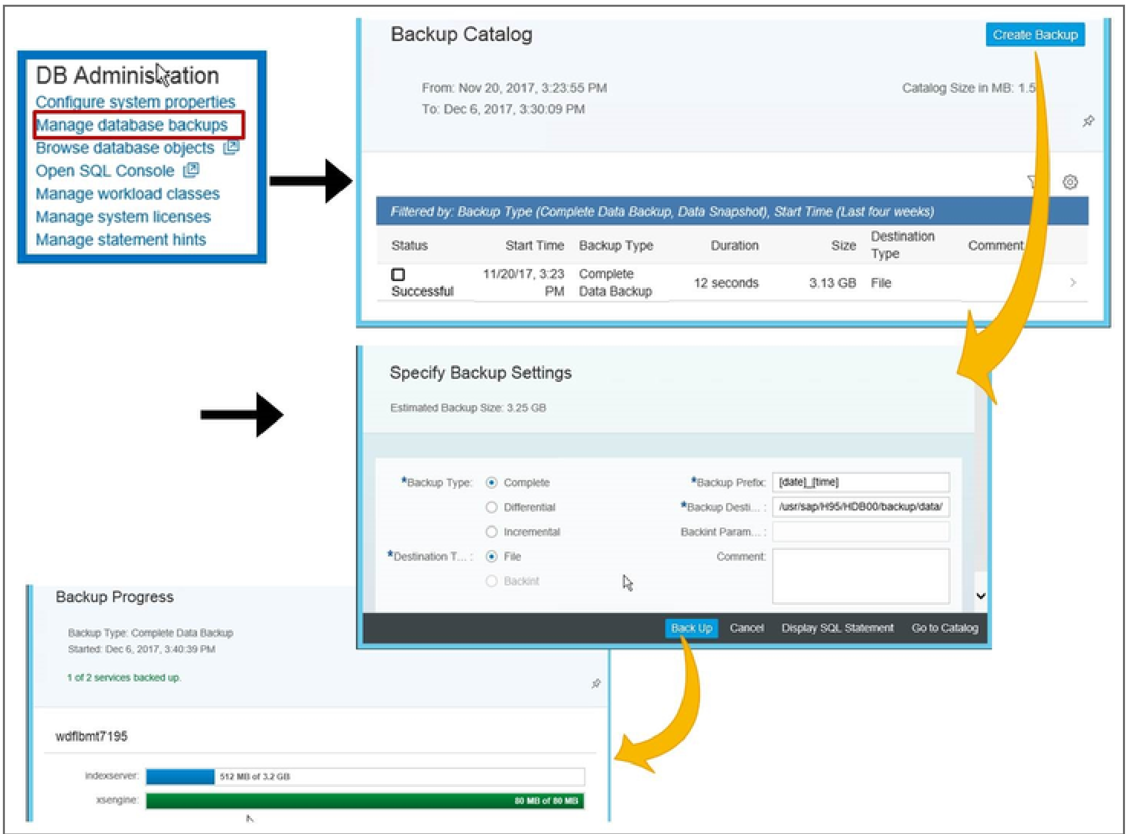

Figure 312: Performing Backups Using SAP HANA Cockpit

Performing a Data Backup Using SAP HANA Cockpit

To create a data backup using SAP HANA Cockpit, perform the following steps:

1. In the SAP HANA cockpit, choose Manage database backups .
2. To open the backup settings page, choose the Create Backup button at the bottom of the backup catalog.
3. Select the type of the data backup from one of the following:
   - Complete Data Backup
   - Differential Data Backup
   - Incremental Data Backup
4. Specify the location (directory) and the backup file prefix to use, and choose Back Up .

   The SAP HANA Cockpit uses the time stamp for the backup file prefix by default.

   The default location shows the path specified in global.ini under the backup parameter basepath_databackup .
5. Once you have started the backup, the progress is displayed.

Overview of Backup Operations

Once you have started the backup, the progress is displayed. When the backup is finished, the backup details are shown.

You can cancel a running data backup from the progress details screen.

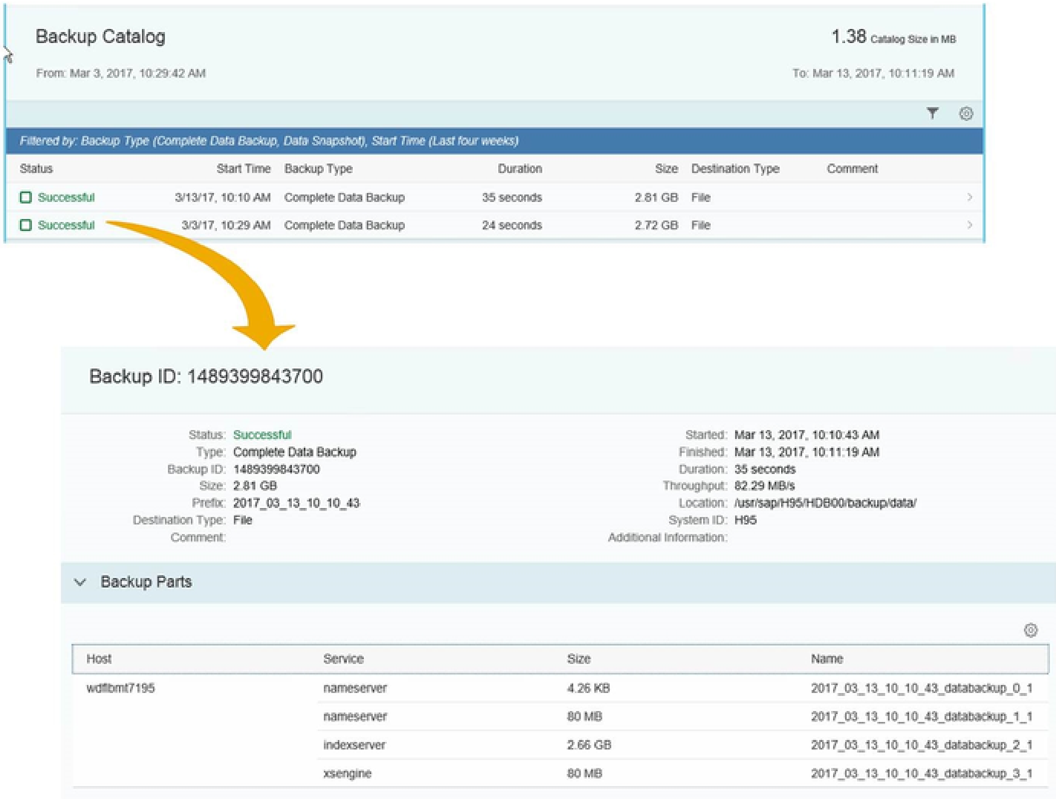

Figure 313: Overview of Backup Operations

Performing a Data Backup Using SAP HANA Studio

To create a data backup using SAP HANA Studio, perform the following steps:

1. In the Navigator view, select the system that you want to back up.
2. From the context menu, choose Back Up .
3. Select the type of data backup from one of the following:
   - Complete Data Backup
   - Differential Data Backup
   - Incremental Data Backup
4. Specify the location (directory) and the backup file prefix to use, and choose Next .

   The default location shows the path specified in global.ini under the basepath_databackup backup parameter.
5. When all the settings are correct, choose Finish . The backup then starts. The progress of the backup is shown for all types of services (for example, the name server, and index servers).When all the volumes have been backed up, a confirmation message displays.

> Note:
> A data backup performed with the SAP HANA studio only saves the payload of the data volumes of the database. The database configuration files (and .ini files) are not backed up. Configuration files (.ini files) that contain customer-specific changes can be backed up manually so that you can easily identify and restore customer-specific changes in a recovery situation.

Performing a Data Backup Using SQL Commands

You can enter SQL commands either by using the SQL editor in SAP HANA studio, or by using the hdbsql program on the command line.

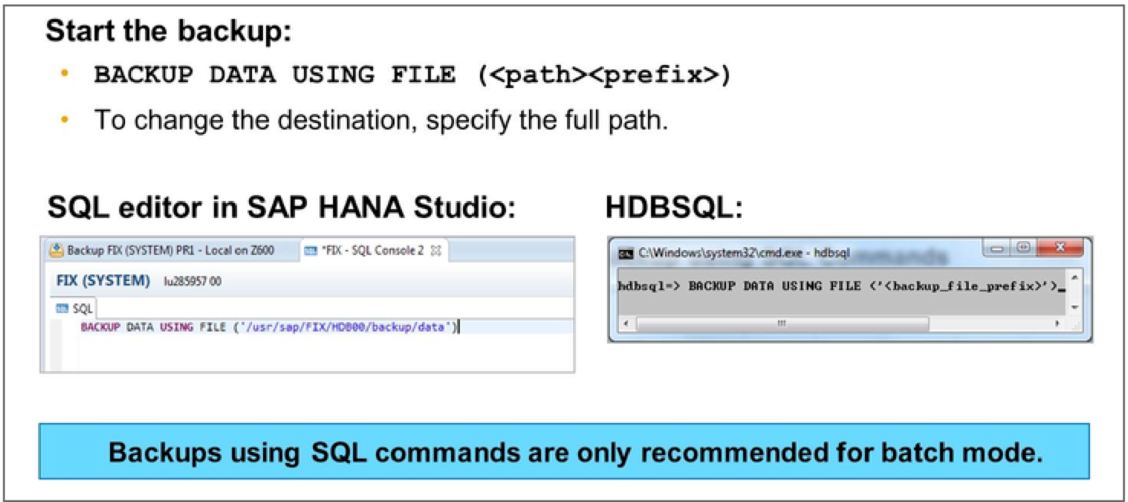

Figure 314: Database Backup Using SQL Commands

> Note:
> Only use backups with SQL commands for batch mode (see the section on backup and recovery in the administration guide).

Scheduling Backups Using SAP HANA Cockpit

Scheduling Backups using SAP HANA Cockpit

Data backups can be scheduled with the SAP HANA cockpit. You can schedule complete, incremental, or differential backups, as well as backups to file or third-party (BACKINT) tools.

To open the backup scheduling page, from within the Data Backup application, choose Create Schedule at the bottom of the backup catalog. To view existing data backup schedules, choose Go to Schedules .

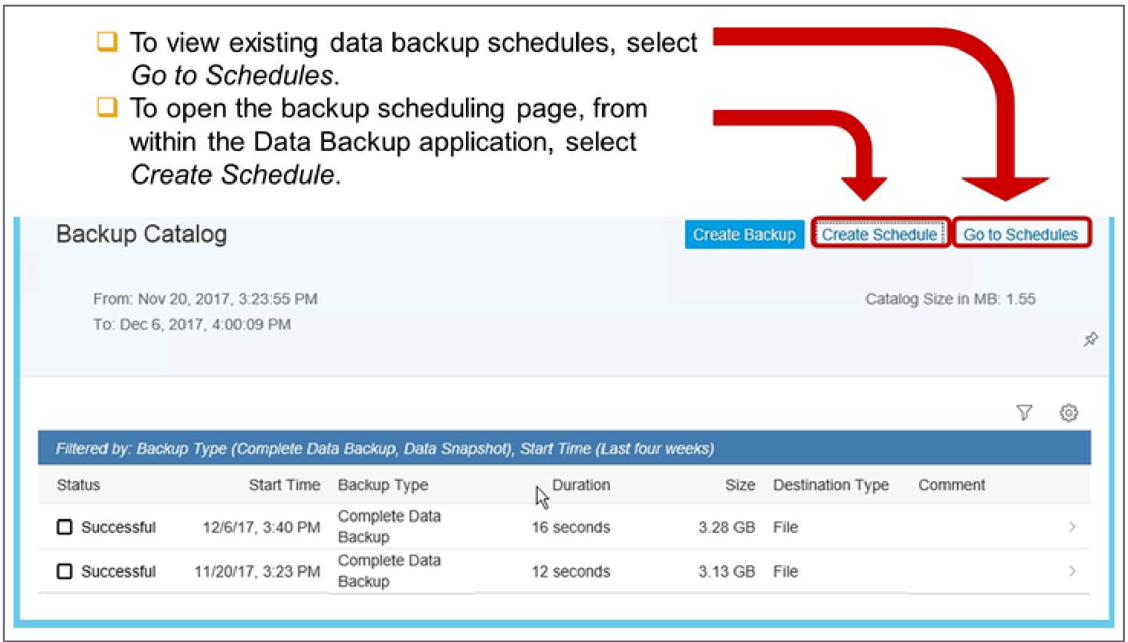

Figure 315: Schedule Backups in SAP HANA Cockpit

Note:
SAP HANA cockpit 2.0 cannot schedule backups for SAP HANA 1.0 databases.

Schedule Data Settings

Specify the schedule data settings as follows:

- Select the general data backup settings:
  - Backup Type
  - Destination Type
  - Backup Prefix
  - Backup Destination
- Specify the proper schedule data settings:
  - Schedule Name
  - Start of Schedule
  - Recurrence pattern
  - Execution time (in UTC time)

Note:
Currently snapshots cannot be scheduled.

Caution:
To transition from SAP HANA 1.0 to SAP HANA 2.0, note the following:

- Backup schedules created with SAP HANA cockpit 1.0 are not compatible with SAP HANA cockpit 2.0.
- Before you upgrade from SAP HANA 1.0 to SAP HANA 2.0, use the SAP HANA cockpit 1.0 to delete all the backup schedules created with SAP HANA 1.0.
- After you upgrade to SAP HANA 2.0, create new backup schedules.

Prerequisites for Scheduling Backups

Backup schedules must be activated globally.

From the backup catalog overview in SAP HANA cockpit, choose Go to Schedules, then set Backup Schedules to `On`.

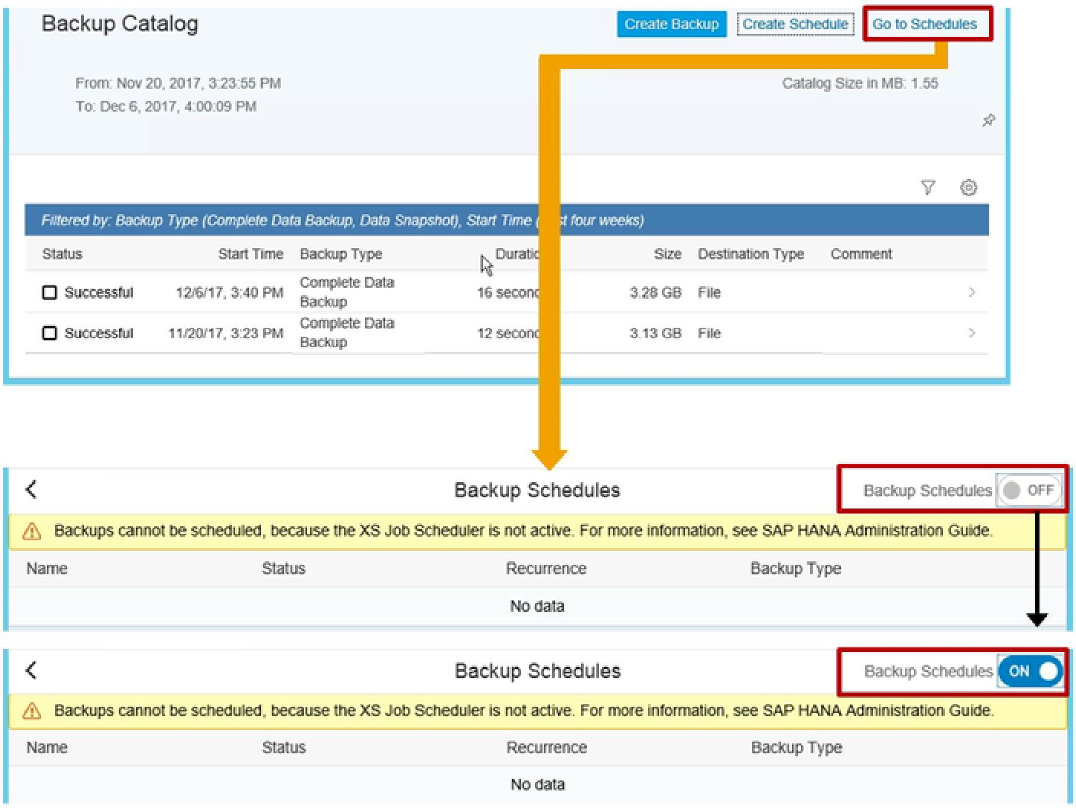

Figure 316: Activate the XS Job backupjob

Enable the Job Scheduler

The XS job scheduler has to be activated for the system database and each tenant database. A backup of a tenant database must be scheduled through the tenant database itself. A backup of a tenant database cannot be scheduled through the system database.

For the system database, the XS scheduler must be enabled in the nameserver.ini file. To enable the XS scheduler, you can use the following SQL statement:.

```
ALTER SYSTEM ALTER configuration ('nameserver.ini','SYSTEM') SET
('scheduler','enabled')= 'true' WITH reconfigure
```

For each tenant database, the XS scheduler must be enabled in the xsengine.ini file. To enable the XS scheduler, you can use the following SQL statement:

```
ALTER SYSTEM ALTER configuration ('xsengine.ini','SYSTEM') SET
('scheduler','enabled')= 'true' WITH reconfigure
```

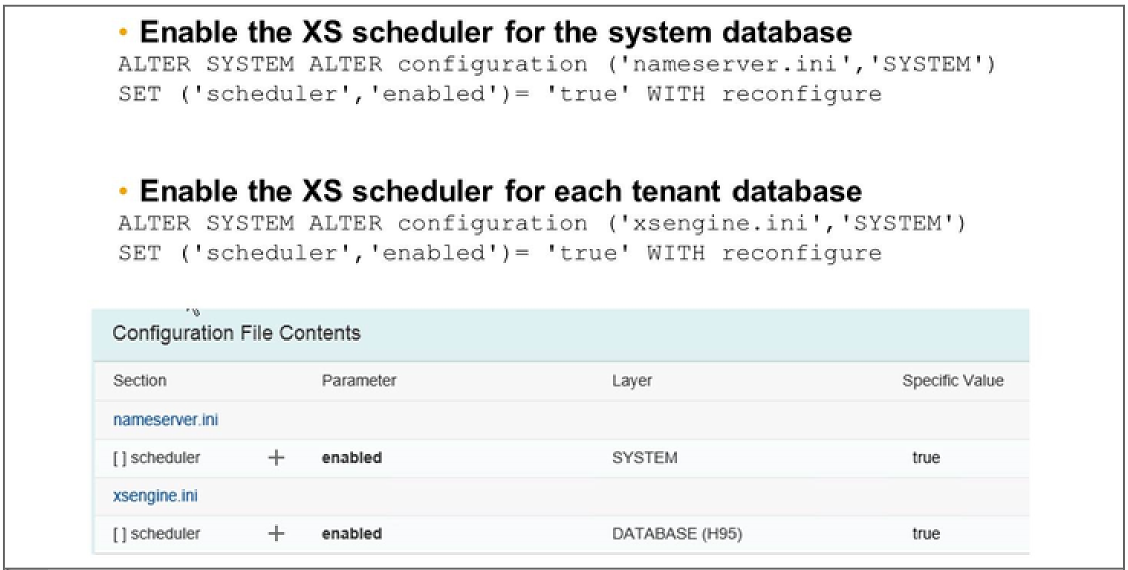

Figure 317: Enable the Job Scheduler

To schedule backups, you require the BACKUP ADMIN system privilege and read authorization for the _SYS_XS.JOB_SCHEDULES and _SYS_XS.JOBS tables.

Overview of Scheduled Data Backups

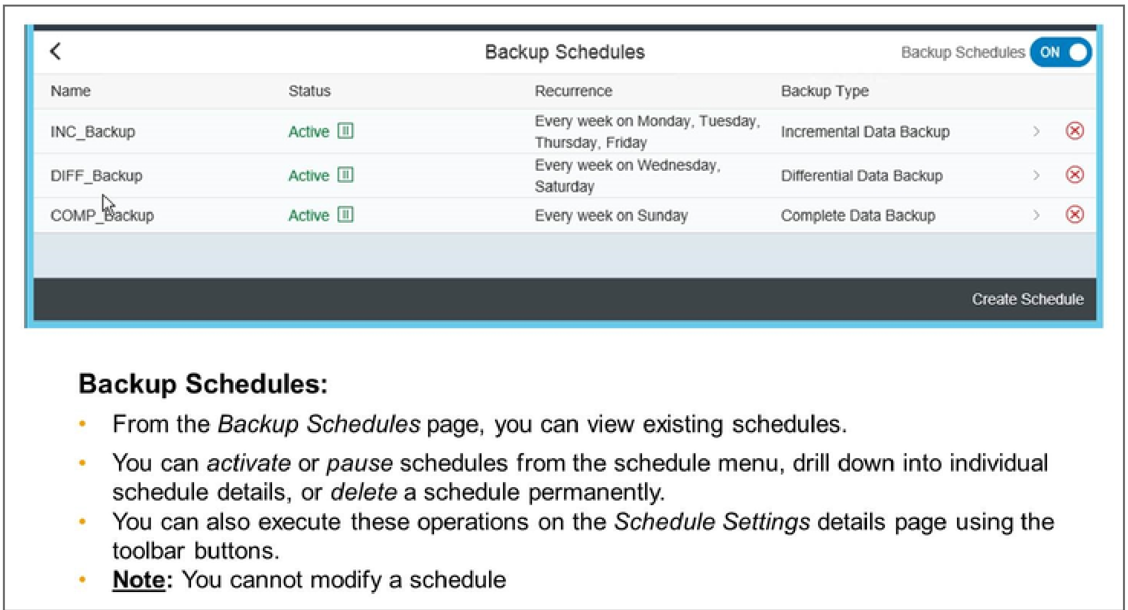

Figure 318: Overview of Scheduled Data Backups

> Note:
> To view data backup prefixes in the Backup Overview page, configure the prefix column by choosing the gear icon.

DBA Planning Calendar

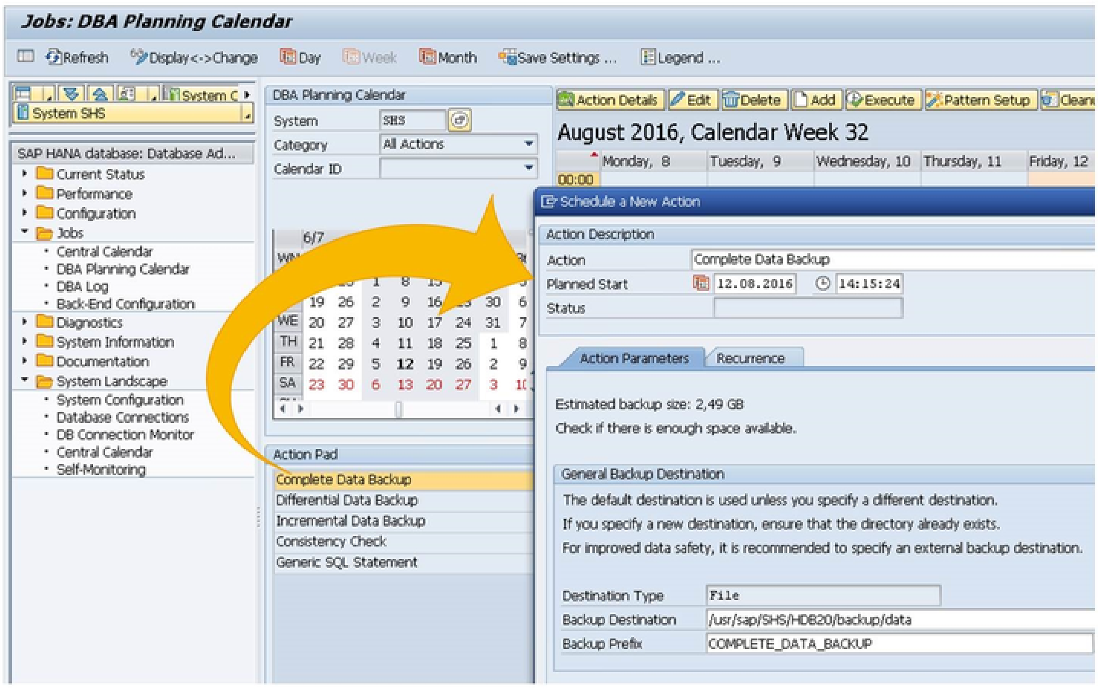

Figure 319: Scheduling Backups Using the DBA Cockpit

The DBA Planning Calendar can be used to schedule, execute, and check almost all regular database administration actions, including data backup, and consistency checks. The DBA Planning Calendar does the following tasks:

- Executes scheduled actions automatically
- Displays actions that are scheduled to run in the background

To start the DBA Cockpit, use transaction code `DBACOCKPIT`.

Performing Scheduling

To schedule an action, proceed as follows:

1. To open the DBA Planning Calendar, in the DBA Cockpit, choose Jobs → DBA Planning Calendar .
2. To create a new action, perform one of the following:
   - Double-click a calendar row.
   - Select a calendar cell and choose Add.

- Move an action from the Action Pad to a calendar cell in the future. You can also move actions to reschedule them.
- To copy an action, hold down the Ctrl key while dragging.

3. Specify the following action details:

- Planned Start: Specify the start date and time of the action.
- Action Parameters: If different from the default, specify the location and prefix for the file.
- Recurrence: Specify when the action will be repeated or if it will be executed only once.

Multistreaming Data Backups with Third-Party Backup Tools

When creating a data backup, a third-party backup tool can use multiple channels to write the backup data for each service.

For example, this capability allows you to distribute backup data in parallel to multiple devices.

By default, SAP HANA uses one channel for data backups. If required, you can configure SAP HANA to use additional channels. When multiple channels are used, SAP HANA distributes the data equally across the available channels. All the parts of a multistreamed backup are approximately the same size.

In the example below, three parallel streams have been configured.

- The index server backup is distributed across three streams.
- Because both name server and XS engine backups are smaller than 128 GB, they are not distributed across several streams

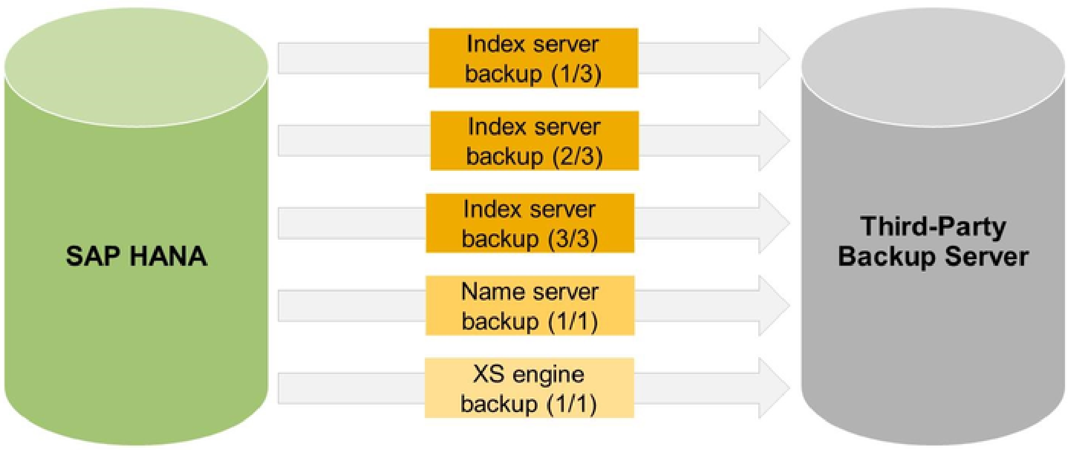

Figure 320: Parallel Backup Streaming via BACKINT

Change the Number of Channels for Multistreaming

If parallel streams have been configured, the individual service backups are distributed across all available streams. Note that the different services always use dedicated backup streams. Backups are only distributed if they are larger than 128 GB. Both full and delta backups are supported.

To configure the number of parallel streams, use the parallel_data_backup_backint_channels.ini file parameter (default: 1, max: 32).

During recovery, the number of streams used is the same as during backup. This is independent of the current setting of the parameter.

> Note:
> To create multistreamed data backups, configure the third-party backup tool to use multiple channels with good performance.
>
> For more information about the configuration of the backup tool, consult the vendor documentation.

Backup of Scale-Out-Systems

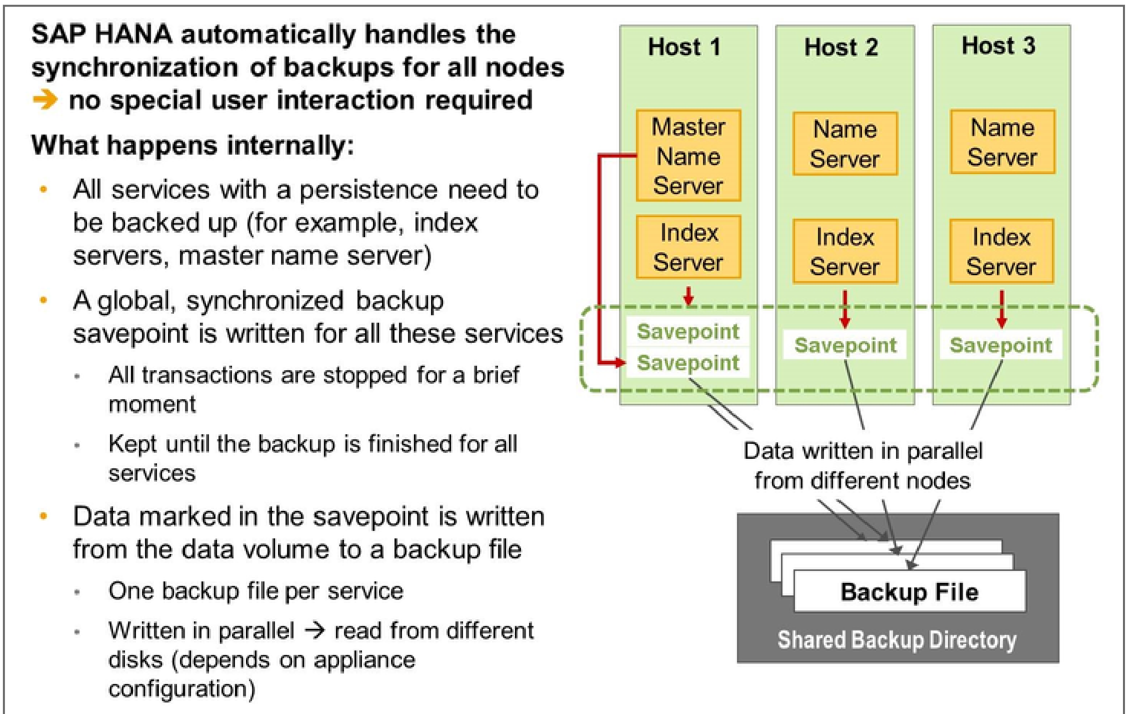

Figure 321: Single-Node and Scale-Out-Systems

SAP HANA automatically handles the synchronization of backups for all nodes. The figure, Single-Node and Scale-Out-Systems, lists the internal process that happens while performing backups for all nodes.

LESSON SUMMARY
You should now be able to:

- Perform data area backup

Unit 10
Lesson 3

Configuring a Log Area Backup

LESSON OVERVIEW

This lesson gives you an overview of the configuration and the different log modes.

Business Example

You have to define a backup strategy for your SAP HANA database. In addition to performing data area backups, you have to configure a log area backup.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Configure a log area backup

Log Management

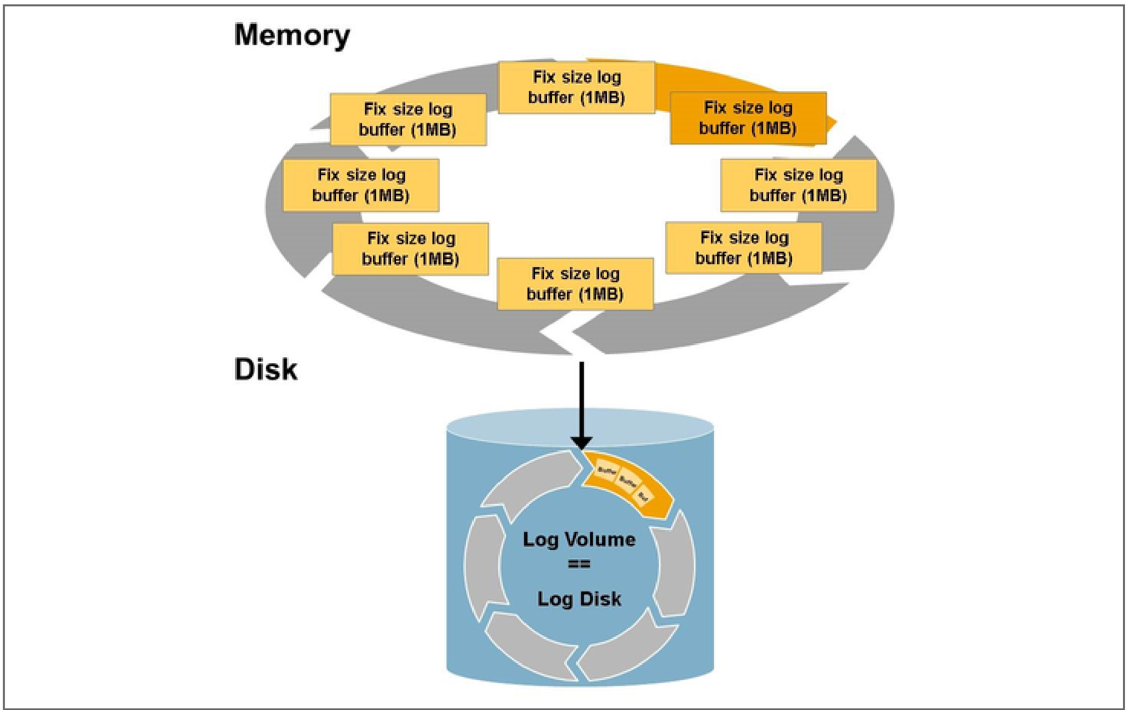

Figure 322: Log Writing

A log is written to Log-Buffers in-memory.

If a Log-Buffer becomes full or a commit entry is written, the Log-Buffer is written to the assigned log volume.

The log is finally written into log segments, where multiple Log-Buffers can be combined.

You can find many log files as log segments (1 GB) on the log volume.

Log Mode

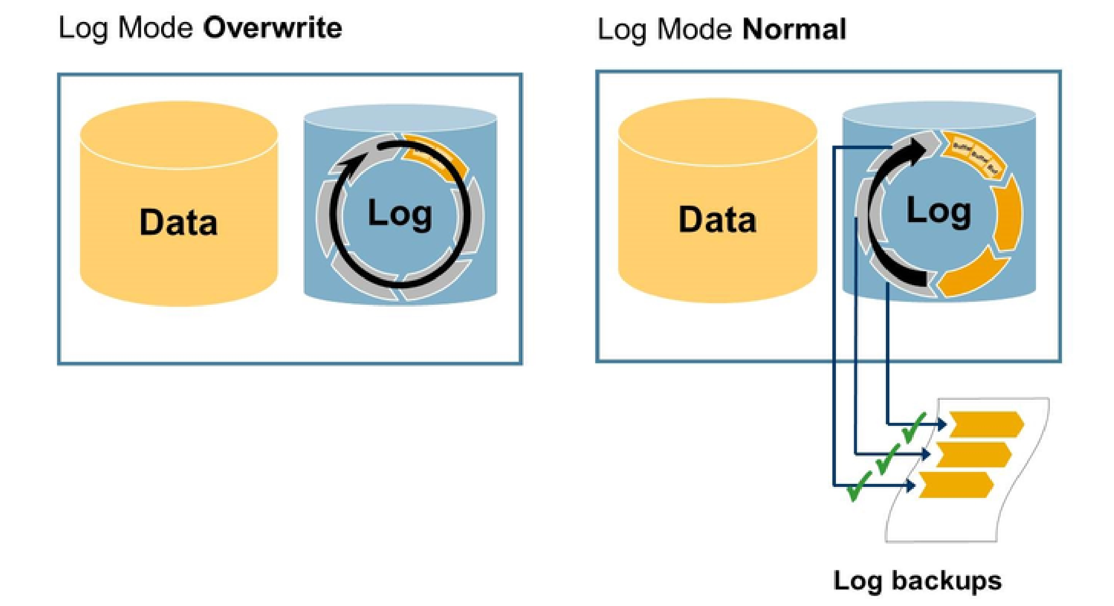

Figure 323: Log Mode

The system can reuse the space that is occupied in the log volume by the log segments. The log_mode parameter controls how the log segments are reconsumed.

Overwrite mode is as follows: `log_mode = overwrite`.

Log segments are freed by savepoints and no log backup is performed. This can be useful, for example, for test installations that you do not need to back up or recover.

Caution:
The overwrite mode, `log_mode = overwrite`, is not recommended for production systems.

With `log_mode = overwrite`, no point-in-time recovery is possible. For recovery, only data backups are used; the logs are not used. The Recover the database to a specific data backup recovery option is the only option that can be selected.

Normal mode is as follows: `log_mode = normal` (default). The features of normal mode are as follows:

- Keeps log segments until backup
- Automatic log backup available (time-based or when segment is full)
- Log backup directory configured with parameter basepath_logbackup
- Backup catalog maintenance
- Restoring of any available data backup with log replay to the last committed state
- Restoring of any available backup without log replay

> Note:
> After installation, SAP HANA temporarily runs in overwrite log mode. After you create the first full data backup, SAP HANA automatically switches to the default normal log mode.

Overview of Log Area Backup

For productive systems, use normal log mode because it provides the highest security for the restoration of data during a recovery of the SAP HANA database. In normal log mode, the system automatically creates log backups that can be used for a recovery in addition to the data backups. However, more backup space is required in this log mode because of the log backups. Therefore, an operational concept for administrating data and log backups is a prerequisite for using normal log mode.

After changing the log mode parameters, restart the database system to activate the changes. Also, create a full data backup of the database.

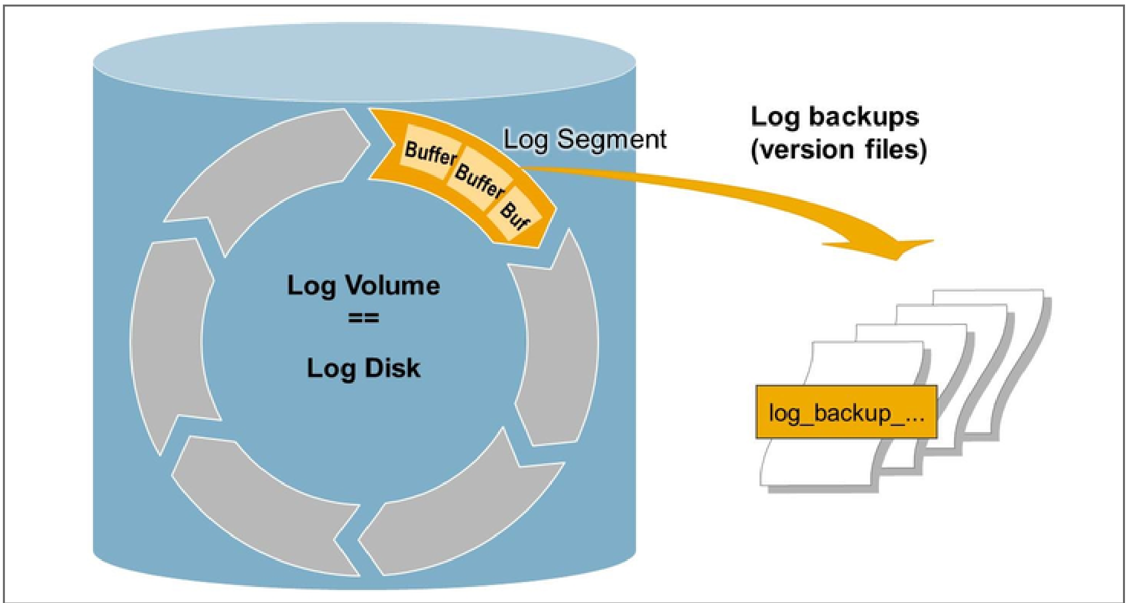

Figure 324: Overview of Log Area Backup

Log Segment Backups

To allow the reuse of log segments, the system can perform regular log backups. During a log backup, the payload of the log segments is copied from the log area to service-specific log backup files.

A log segment is backed up in the following situations:

- The log segment is full.
- The log segment is closed after exceeding the configured time threshold.
- The database is started.

If you do not regularly move the log backup files to an external destination, you risk the file system becoming full. Log segments can only be overwritten by the system after they have been backed up.

Caution:
Do not delete log segments on the operating system level, because the log area will become unusable and the database might stop working immediately.

Note:
If backups go to the file system, you must also regularly archive the BACKUPS log to avoid the BACKUP DESTINATION log from becoming full.

Configuration of the Log Backup

Log backup behavior: Configured by parameters in the *global.ini* configuration file.

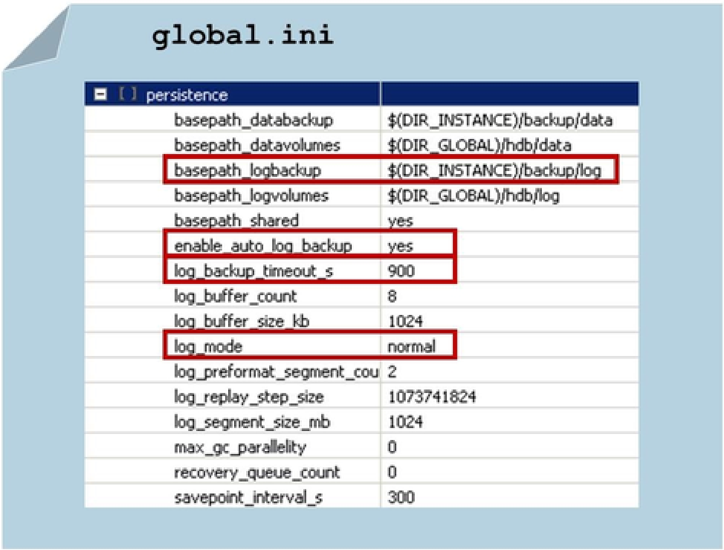

Figure 325: Backup Configuration for the Log Area

Location of the Log Backup Files Using Destination Type FILE

The log backup files are written to the location specified by the basepath_logbackup parameter in the persistence section of the global.ini configuration file. By default, the location for log backup files is `$(DIR_INSTANCE)/backup/log`.

The default backup destination can only be changed for file-based backups. If you change the backup location in basepath_logbackup, the change occurs immediately.

Note:
Backups made using third-party tools always use the destination: `/usr/sap/<SID>/SYS/global/hdb/backint`. For this reason, it is not possible to change the backup destination for third-party tools. For a destination for third-party tools, only named pipes are created in the file system. Named pipes occupy no space in the file system.

Automatic Log Backup

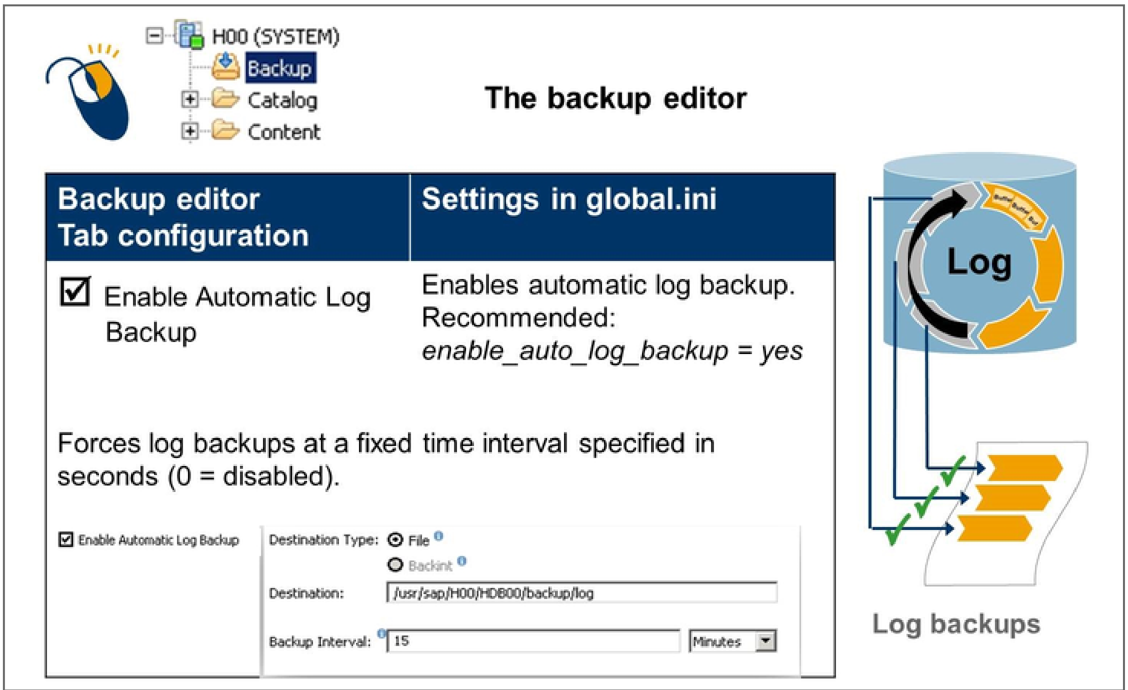

Figure 326: Automatic Log Backup

You can enable or disable automatic log backup with the enable_auto_log_backup parameter. The default setting is `ENABLE_AUTO_LOG_BACKUP = YES`.

> Note:
> In the default log_mode normal, if automatic log backup is disabled, the log area grows until the file system is full. At that stage, the database freezes.

Log Backup Timeout Parameter

The parameter log_backup_timeout_s in the global.ini configuration file defines the interval at which log backups are created. By default, the log backup interval is 15 minutes (900s).

Specifying an appropriate interval for log backups enables you to recover an SAP HANA database with a good Recovery Point Objective (RPO). In the event of database failure, the RPO is the maximum time span of data that will be lost if the log area cannot be used for recovery, and if only data backups, delta backups, and log backups are available.

> Note:
> If the log segments become full before the log backup interval, the logs are backed up automatically. A time interval of 0 means that log backups are created only when a log segment is full and when services are restarted.

> Note:
> The log_backup_timeout_s parameter only takes effect if enable_auto_log_backup is set.
>
> For `LOG_MODE = NORMAL`, these parameters must have the following values:
>
> - `ENABLE_AUTO_LOG_BACKUP = YES`
> - `LOG_BACKUP_TIMEOUT_S > 0`

To provide full point-in-time recoverability, enable automatic log backups in production systems. An alert notifies administrators when automatic log backups are disabled.

> Note:
> Improvements for log backups when using BACKINT include the following:
>
> - In some cases, third-party backup tools have encountered deadlocks when two SAP HANA database services requested log backups from the same tape (there is no concurrent access to the tape). The internal recovery handling of SAP HANA has been adapted to avoid deadlock situations when retrieving log backups from a third-party backup tool that uses tapes.
> - In some scenarios, the start of a third-party backup agent for a log backup can take longer than the actual log backup itself. During times of high load, this can cause many pending log backups and, in the worst case, log full situations (log segments are only released for overwrite after a successful log backup). SAP HANA now uses a single backup call to the third-party agent for all log segments of a service that are ready for backup.

Options to Optimize the Workload of the Log Backup

With a well-operated SAP HANA system, many log backups are created. It is common to have over 2,000 backup files per day. External backups tools can have problems with digesting so many files.

SAP HANA uses several options to reduce the number of backup files that are created.

SAP HANA uses several options to reduce the number of created backup files

Service consolidation

- By checking if several services have to write backups at the same time:
  - Collect these files into one package
  - Results in one backup activity instead of several before
- Alternatively, for older SAP HANA releases, the log segment size can be increased to 4 gigabytes (max.) per segment

Log backup interval mode

- Instead of event-driven frequency (log segment finished), interval can be reduced to 15 min
- Parameter: LOG_BACKUP_INTERVAL_MODE
  - `immediate` (default) – existing event driven behavior
  - `service` – wait for the next service (15 minutes interval)

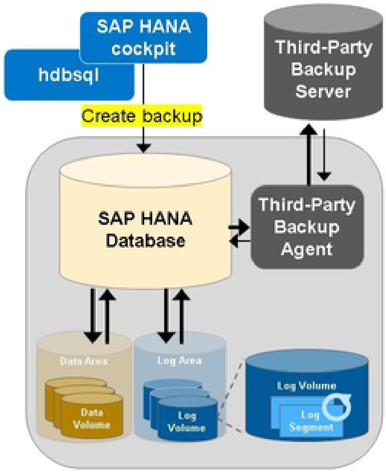

Figure 327: Options to Optimize the Workload of the Log Backup

Writing Multiple Log Segments to One Log Backup

To improve the performance of log backups, SAP HANA can write all the log segments of a service that are ready to be backed up at a particular time to a single log backup. You can define the maximum size of this single log backup. This option is supported for both file-based backups and third-party tools.

If a single log backup operation takes a long time, several other log segments can be queued for backup during that time. During periods of high load, log segments are closed and queued for backup faster than a single backup operation is completed.

Log segments in the log area are only released for overwrite after a successful log backup. In some situations, there can be a delay in releasing log segments because they are waiting to be backed up. As a result of this delay, the log area can grow. If the log segments cannot be backed up and released faster than the log area is growing, the log area can become even more full.

To remedy this issue, SAP HANA can write all the log segments of a service that are ready to be backed up at a particular time to a single log backup.

The maximum size of the log segments to be processed by a single backup operation is defined by the parameter max_log_backup_size in the backup section of the global.ini file. The default value is 16. This means that one backup operation creates log backups with a maximum size of 16GB.

LESSON SUMMARY

You should now be able to:

- Configure a log area backup

Unit 10
Lesson 4

Describing Additional Backup Topics

LESSON OVERVIEW

This lesson explains how the backup catalog provides information about the backups you have performed.

Business Example

You have to define a backup strategy for your SAP HANA database. Therefore, you have to define a strategy to backup the configuration files of your database. After you have defined a strategy for the data area and the log area backup, you need information about the execution of backups and their history.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Describe additional backup topics

The Backup Log

The `backup.log` and `backint.log` files record information about backups. This information can be used to diagnose errors.

> Note:
> As more data is written to `backup.log` and `backint.log`, the files grow, but their increased size does not impact database performance. If `backup.log` or `backint.log` become too large for the available disk space, you can safely delete or rename either file as required.

The `backup.log` records information about data backups, log backups, the progress of a backup, and the backup catalog. It also records information about recovery operations.

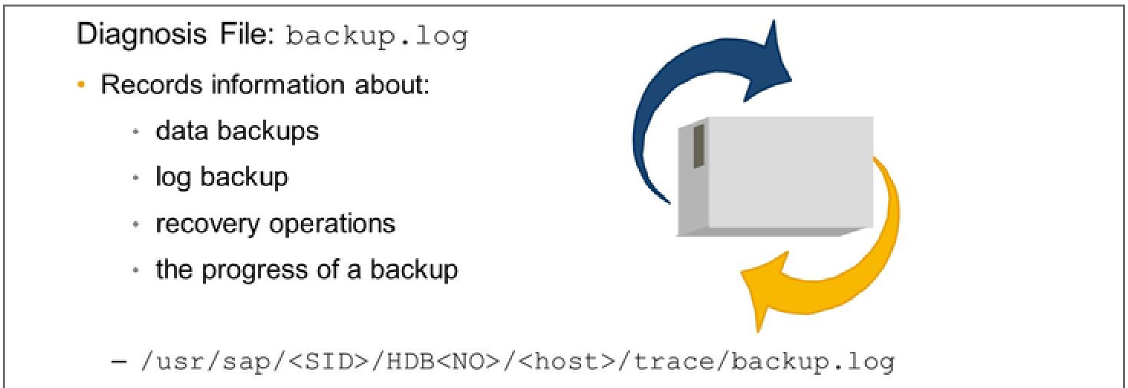

Figure 328: Monitoring Backups

The `backup.log` file records information about the data and log backups.

The `backint.log` file contains information about the activities of the BACKINT agent. The BACKINT agent is part of a third-party backup tool.

Backup Catalog

The backup catalog provides information about the execution of backups and their history. It enables the system to do the following:

- Determine whether a recovery is possible
- Choose which data and log backup to use to recover the database
- Determine which backup files are no longer needed for a recovery

Information in the Backup Catalog

The backup catalog contains information on the following:

- Backups created for a database
- The start and completion times of the backups
- Whether a backup is still running
- Whether a backup was successful or not
- Volumes that were backed up
- Log backups and what part of the log they contain
- Backup destinations and their sizes
- The destination type
- The backup ID
- An external backup ID when using a third-party backup tool

Backup Catalog in the SAP HANA Cockpit

The Backup Catalog displays a list of past backups. This list allows you to see the status of each catalog entry, as well as its key information, at a glance. To see the full details of a particular entry, select it from the list. More information appears in the Backup Details area. This includes, for example, backup start and completion times, duration, size, throughput time, and a breakdown for each service.

By default, only full data backups are displayed. To see delta backups, select the Differential Data Backup or the Incremental Data Backup checkboxes.

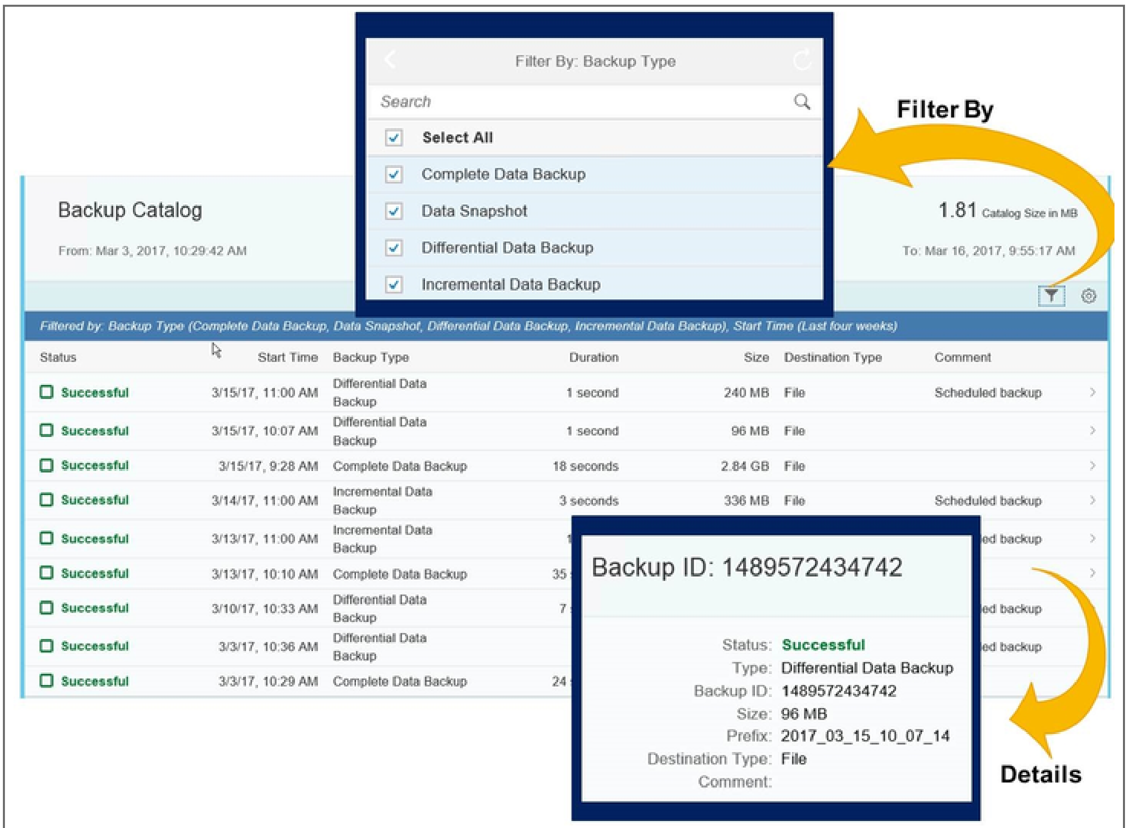

Figure 329: Backup Catalog in the SAP HANA Cockpit

Note:
You can also access the backup catalog using the Backup Editor in SAP HANA Studio.

Backing Up the Backup Catalog

The backup catalog is backed up and versioned after every completed backup operation. This allows the backup catalog to be accessed during a recovery. Even when `log_mode = overwrite` is set, where logs are not created, the backup catalog is still backed up.

If the backup catalog is backed up using a third-party tool, the versioning of the backup catalog is handled by the backup tool.

SAP HANA writes one backup of the backup catalog for concurrent log backups of different services. This means that the backup of the backup catalog covers all log backups that were written since the previous backup of the backup catalog. This action is enabled by default. To disable it, choose global.ini → backup, and set the enable_accumulated_catalog_backup database configuration parameter to **`False`**.

The backup catalog is written as a separate backup. By default, the backup catalog is backed up to the same destination as the log backups.

Caution:
If you change the default destination for the log backups, the backup catalog is not automatically backed up to the same location.

The parameters to configure the destination of the backup catalog are shown in the following figure.

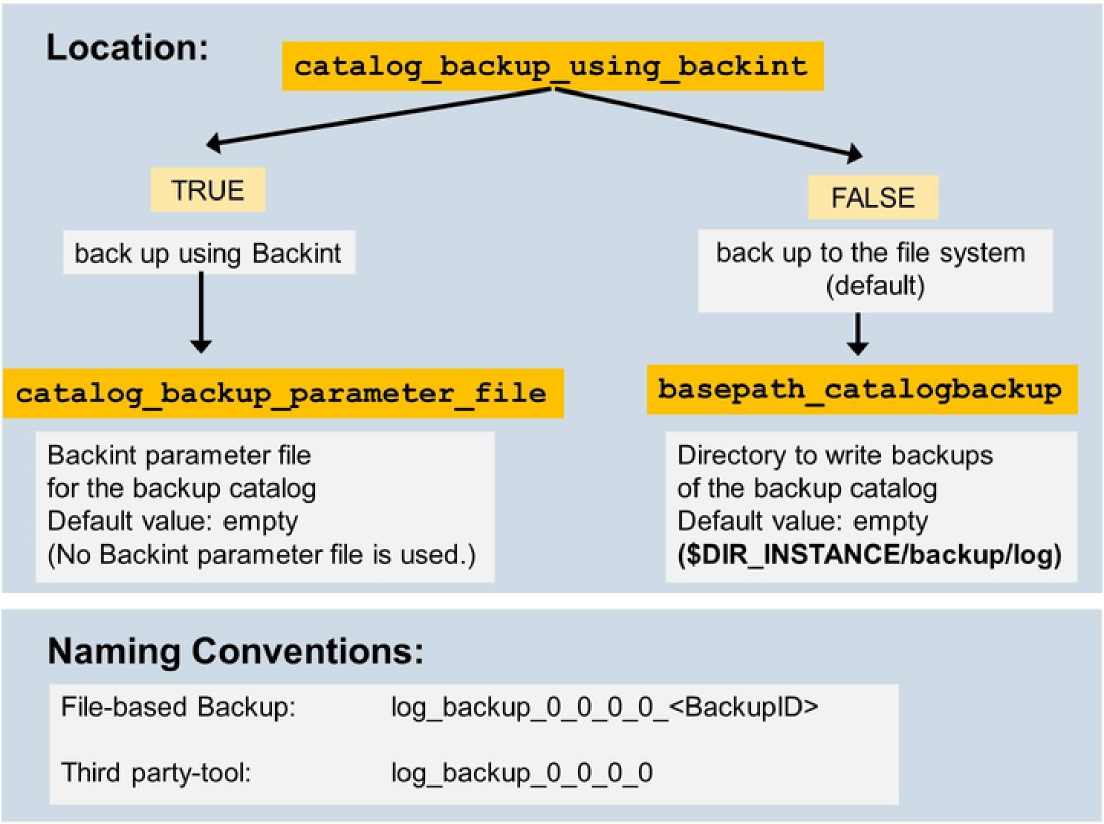

Figure 330: Backing up the Backup Catalog

Monitoring Views of the Backup Catalog

You can access the backup catalog using monitoring views. Monitoring views are located in the SYS schema.

The monitoring views M_BACKUP_CATALOG, M_BACKUP_CATALOG_FILES, and M_BACKUP_PROGRESS provide different overviews of information from the backup catalog.

- M_BACKUP_PROGRESS provides detailed information about the most recent data backup.
- M_BACKUP_CATALOG_FILES provides information about the backups created, and the backup destinations for data and log backups.
- M_BACKUP_PROGRESS provides information about currently running and last finished backups only. It is cleared at database restart.

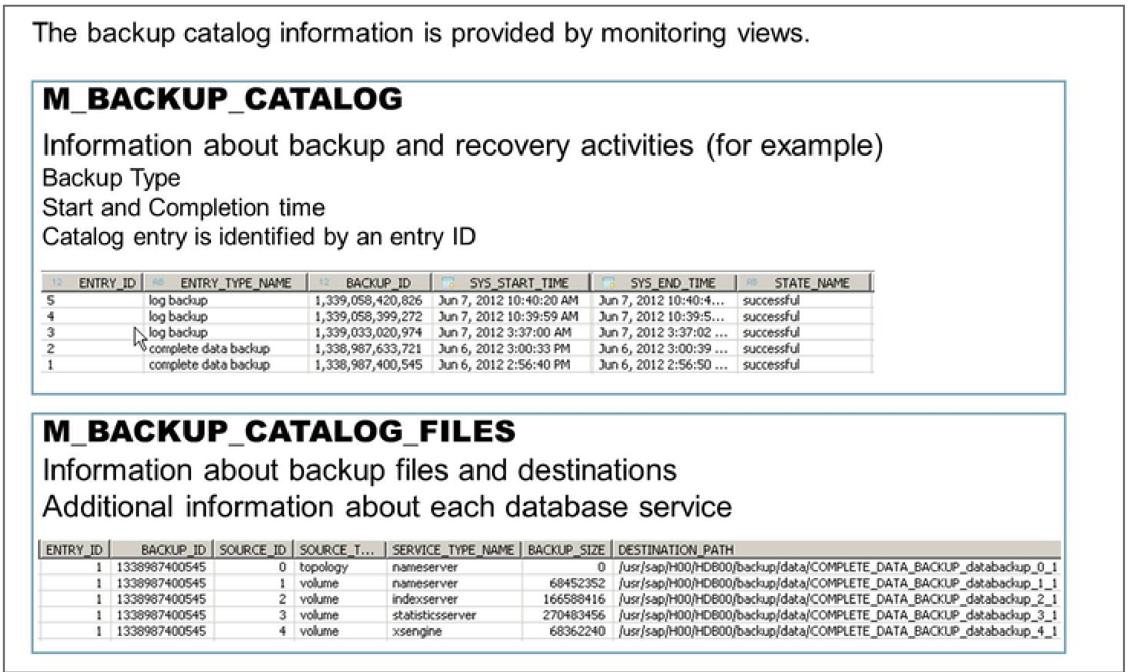

ENTRY_ID	ENTRY_TYPE_NAME	BACKUP_ID	SYS_START_TIME	SYS_END_TIME	STATE_NAME
5	log backup	1,339,058,420,826	Jun 7, 2012 10:40:20 AM	Jun 7, 2012 10:40:4...	successful
4	log backup	1,339,058,399,272	Jun 7, 2012 10:39:59 AM	Jun 7, 2012 10:39:5...	successful
3	log backup	1,339,033,020,974	Jun 7, 2012 3:37:00 AM	Jun 7, 2012 3:37:02 ...	successful
2	complete data backup	1,338,987,633,721	Jun 6, 2012 3:00:33 PM	Jun 6, 2012 3:00:39 ...	successful
1	complete data backup	1,338,987,400,545	Jun 6, 2012 2:56:40 PM	Jun 6, 2012 2:56:50 ...	successful

ENTRY_ID	BACKUP_ID	SOURCE_ID	SOURCE_T...	SERVICE_TYPE_NAME	BACKUP_SIZE	DESTINATION_PATH
1	1338987400545	0	topology	nameserver	0	/usr/sap/H00/HDB00/backup/data/COMPLETE_DATA_BACKUP_databackup_0_1
1	1338987400545	1	volume	nameserver	68452352	/usr/sap/H00/HDB00/backup/data/COMPLETE_DATA_BACKUP_databackup_1_1
1	1338987400545	2	volume	indexserver	166588416	/usr/sap/H00/HDB00/backup/data/COMPLETE_DATA_BACKUP_databackup_2_1
1	1338987400545	3	volume	statisticsserver	270483456	/usr/sap/H00/HDB00/backup/data/COMPLETE_DATA_BACKUP_databackup_3_1
1	1338987400545	4	volume	xsengine	68362240	/usr/sap/H00/HDB00/backup/data/COMPLETE_DATA_BACKUP_databackup_4_1

Figure 331: Backup Catalog

M_BACKUP_PROGRESS provides detailed information about the most recent data backup.

Backup Lifecycle Management

To keep your backup storage space at an optimum level, delete backups that are no longer needed for a recovery regularly.

To free backup storage, you can physically delete data backups and log backups, and delete their associated entries in the backup catalog. To reduce the size of the backup catalog, delete the records of individual data backups from the backup catalog. However, retain the physical backups, for example, to comply with legal requirements for data retention.

Backup lifecycle management provides a framework to delete old data and log backups from the backup catalog only, or from the backup catalog and physically from the backup location. You can delete backups from the file system or from a connected third-party backup server via the BACKINT interface. This allows you to manage your backup storage space or to fulfill regulatory deletion requirements.

The deletion functionality is available in SAP HANA studio and on the command line using SQL commands.

Delete old backups:

- From the Backup Catalog
- From the Backup Catalog and Backup Location
- Available in SAP HANA Studio

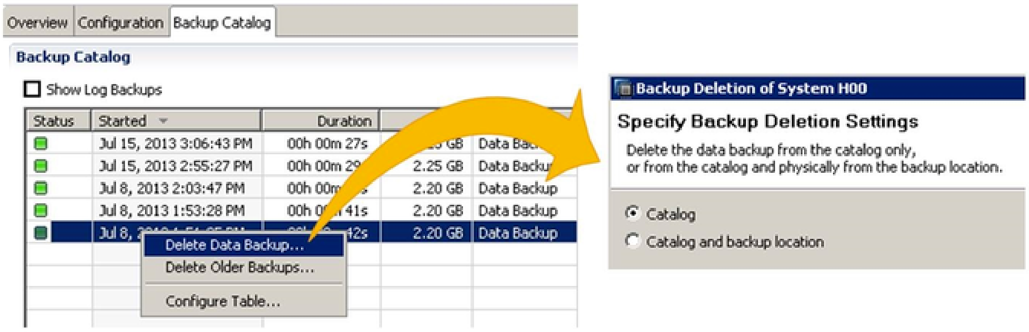

- Available using SQL commands:

```
BACKUP CATALOG DELETE BACKUP_ID <backup_id> [COMPLETE]
BACKUP CATALOG DELETE ALL BEFORE BACKUP_ID <backup_id>
[WITH FILE] | [WITH BACKINT] | [COMPLETE]
```

Figure 332: Backup Lifecycle Management

There is an audit event that you can enable that creates an entry in the audit trail whenever a backup is deleted using this function.

> Note:
> If data snapshots are deleted from the backup catalog, they are not deleted physically. Data snapshots that are no longer needed must be deleted manually.

Performing Backups with Scripts

In addition to performing backup and recovery operations using the SAP HANA Cockpit and the SAP HANA studio, you can also use SQL statements. The syntax for these statements is described in the SAP HANA Administration Guide.

You can use these SQL statements to define scripts that trigger a database backup with SAP HANA backup functionality. For an example of such a backup script, see SAP Note 1651055 .

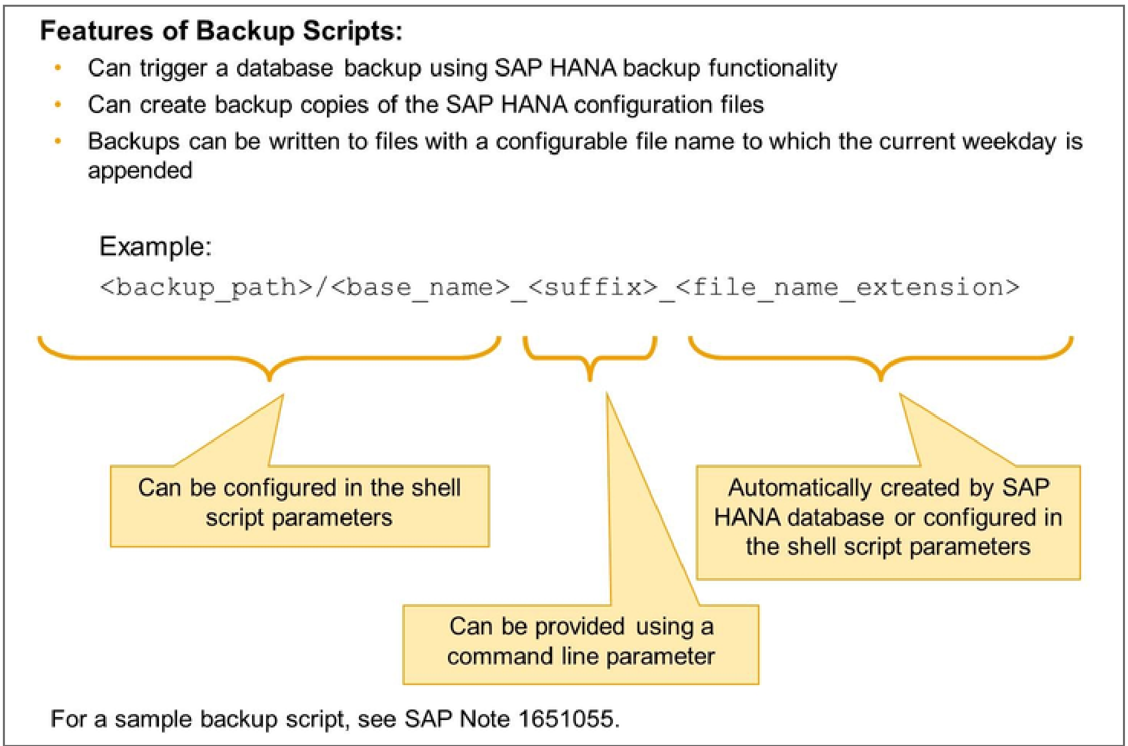

Figure 333: Scheduling a Backup Using Scripts

Example of Database-Specific Parameters

Table 12: Example of Database-Specific Parameters

Name	Default	Description
SID *	- - -	SYSTEM ID of the SAP HANA database system
INSTANCE *	- - -	Instance number of the SAP HANA database system
HOSTNAME *	- - -	Hostname (the local host name of the database server. Do not use “localhost” . Do not use the fully qualified (<host-name>.<domain>) name.
SIDPATH **	`/usr/sap/${SID}`	The directory into which the binaries of the SAP HANA database system are installed
INSTPATH **	`${SIDPATH}/HDB${INSTANCE}`	The directory containing the instance data of the SAP HANA database

Note: (*) means that the parameter must be adjusted to your particular installation; (**) means that the parameter refers to a default setting of the SAP HANA database that is unlikely to be changed in any database installation.

Specify parameters as follows:

- `<name>=<value>`

- No space on either side of the "=" operator
- The name of parameters must not be changed
- Case-sensitive
- Use `${<name>}` for parameters that reference other parameters

Command Line Options

Table 13: Command Line Options

The backup script offers the following command line options:

Name	Description
`-h`	Display usage information and exit (regardless of any other command line parameters given).
`-t`	Test mode: Do not create or delete backup files, that is, do not create data backup, do not create configuration file backup. Writes log messages into file ${SCRIPT_LOG}.
`-q`	Suppress wait time and information output (recommended in batch mode).
`-d`	Only create a data backup. Do not back up configuration files.
`-c`	Only back up configuration files. Do not run a database backup.
`-p`	Add script parameterization and command line switches to the script log file.
`--suffix=<value>`	Create backup files that do not contain the weekday as part of the name, but <value> instead. Note: There must not be any white space on either side of the "=" sign.

Backup of Configuration Files

Customer-specific configuration settings (*.ini files) are not backed up automatically as part of a full backup. The configuration settings are not essential to perform a database recovery. If you want to back up configuration files that contain customer-specific changes, you can do so manually.

In a recovery situation, a backup of the configuration settings can be helpful to more easily identify and restore customer-specific changes to the default settings. If you want to use customer-specific configuration settings after a recovery, you need to reconfigure the recovered system using SAP HANA cockpit or SAP HANA studio.

The properties of an SAP HANA system are defined in the parameters of its configuration files.

The configuration files (.ini files)

- Contain details of the configuration of the database
- Customer specific modification in .ini files could be backed up manually
- Should be copied to an external backup destination

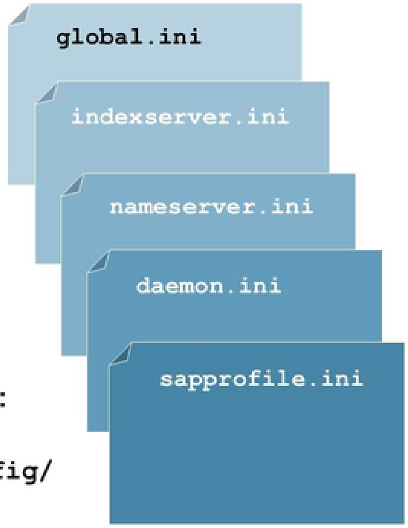

The .ini files are located in the following directories:

Global configuration settings
/usr/sap/<SID>/SYS/global/hdb/custom/config/

Configuration settings for a tenant database
/usr/sap/<SID>/SYS/global/hdb/custom/config/DB_<tenant_name>

For host-specific configuration settings:
/usr/sap/<SID>/HDB<instance-no.>/<hostname>

Figure 334: Backup of Configuration Files

Configuration files are only created in these directories if customer-specific changes are made to them after installation. If no customer-specific changes have been made, these directories are empty.

The nameserver.ini file contains global information for each installation. The landscape section contains the system-specific landscape ID and assignments of hosts to the MASTER, WORKER, and STANDBY roles. If the system landscape changes, for example, hosts are added or removed, the landscape section of the nameserver.ini also changes.

> Caution:
> The sapprofile.ini contains information that is specific to each host. For this reason, in a recovery situation, do not copy the sapprofile.ini file manually to a different host, because it will not be compatible with a new landscape.

Binary Configuration File

In addition to the configuration files, all customer-specific changes are also saved in one separate (binary) configuration file. This file is created when SAP HANA is installed and it is stored in the same directory as the configuration files.

The binary configuration file is versioned. When the file is changed, a new version is created and the previous version is renamed sequentially. All the file versions are stored in the same directory.

If you want to back up customer-specific configuration changes, back up all the versions of the binary configuration file manually together with the other configuration files.

In a recovery scenario, if you want to restore customer-specific settings, use both the configuration files (.ini files) and the binary configuration file.

To restore customer-specific configuration settings from the binary file, use the command line tool `hdbparam`. If you do not want to restore the most recent version of the binary file, use `hdbparam` to check the individual parameter values and decide which version of the binary file to restore.

LESSON SUMMARY

You should now be able to:

- Describe additional backup topics

Unit 10
Lesson 5

Performing Database Recovery

LESSON OVERVIEW

This lesson explains when it is necessary to recover SAP HANA and how you can do this.

Business Example

Because of a hardware error, the database can no longer be started. After solving the hardware problem, you perform a recovery of the database.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Perform database recovery

Recovery of Multitenant Database Containers

An SAP HANA database can be recovered using data backups and log backups . For SAP HANA multitenant database containers, you can recover the system database. You can also recover a tenant database via the relevant system database. System database and tenant database can be recovered one by one in the same system (recovery) or in a different system type (system copy) of the SAP HANA multitenant database containers. For a recovery, the source database and target database must have identical configurations.

A recovery of the system database may be needed, for example, if there are physical errors in the system database's volumes.

Recovery of an SAP HANA Database

When is it Necessary to Recover the SAP HANA Database?

- Disk crash of the data area
- Disk crash of the log area
- Reset system to a certain point in time for special recovery purpose
- Create a copy of the database

What Kind of Recovery Procedures are Available for the SAP HANA Database?

- Recovery to status before failure
- Point-In-Time recovery
- Recovery to a specific data backup

Steps to Recover Database

The steps to recover the database depend on the recovery scenario and the reason for the recovery. This lesson describes some recovery scenarios.

Data Area is Unusable

If the data area is unusable, and all the data changes after the last complete data backup are still available in the log backups and log area, you can recover the data from committed transactions that was in-memory at the time of failure.

No committed data is lost.

For recovery, the data backups, the log backups, and the log area are used. When you restore the data backup successfully, the log entries from the log backups and the log area are replayed automatically.

You can also recover the database using an older data backup and log backups. All relevant log backups made after the data backup are needed for the recovery.

For more information, see SAP Note 1705945 : Determining the files needed for a recovery.

Log Area is Unusable

If the log area is unusable, you can only replay the log backups. Therefore, any changes that are made after the most recent log backup are lost. In addition, all the transactions that were open during the log backup are rolled back.

You can still recover the database to a point in time within the existing log backups.

For recovery, the data backups and the log backups are used. When the data backup is successfully restored, the log entries from the log backups are automatically replayed. To prevent the recovery of entries from the unusable log area, in the Recovery Wizard , specify the Initialize log area option.

Logical Error: Point in Time Recovery

To reset the database to a particular point in time, you need a data backup from before the point in time to recover to, the subsequent log backups, and the log area.

All changes made after the recovery time are lost. If you perform this recovery, consider recovering the database to a different system.

Recovery Types

The figure, Backup History, shows an overview of the possible backup types during normal operation.

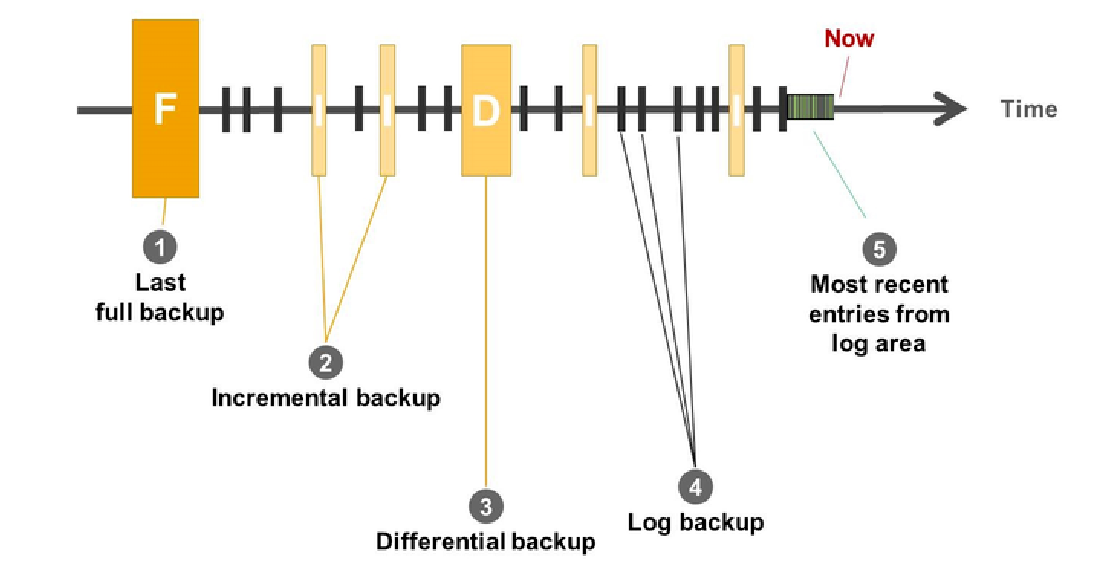

Figure 335: Backup History

Recovery Types

The following recovery types are available:

(A) Recover the database to its most recent state

This option recovers the database to as close as possible to the current time.

This recovery option uses the following data:

- The most recent data backup
- Log backups made since the most recent data backup
- Log area

(B) Recover the database to a specific point in time

This recovery option uses the following data:

- The last data backup available before the specified point in time
- Log backups made since the data backup to be used
- Log area

(C) Recover the database to a specific data backup

This recovery option uses the following data:

- The specified data backup

Note:
Option (C) is not supported for delta backups. Log entries are not replayed, neither from the log backups nor from the log area. All log entries that still exist in the log area are deleted.

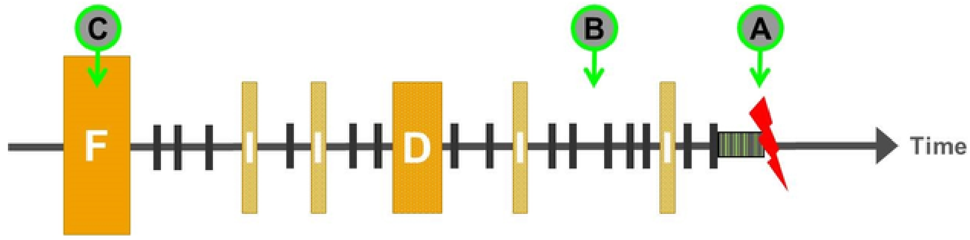

(A) Until the latest point in time

- Uses the latest available data backup + all log backups since the latest data backup + all log entries that are still available in the log area (if it was not destroyed by the crash)

(B) Until a specified point in time (in the past)

- Uses the last data backup before the specified point in time + the log entries from the log backups between that data backup and the specified point in time

(C) Using a specified data backup (with clear log)

- Uses the specified data backup. All log entries that still exist in the log area are deleted, no log will be replayed

Figure 336: Recovery Types

Requirements

Requirements for Performing a Recovery

- SAP HANA database must be shut down
- User must have system privilege **BACKUP ADMIN**
- The <sid>adm operating system user is required
- At least one data backup is required
- Before a recovery can start, all data and log backups must be either accessible in the file system or available through the third-party backup tool
- The number and type of services in source and target systems must be identical
- To recover customer-specific configuration settings, configure the customer-specific settings before starting the recovery

Constraints

- Recovery to a lower software version is not possible

Performing SAP HANA Database Recovery

To perform an SAP HANA database recovery, the following requirements must be met:

- The SAP HANA database software must be installed, so that an initial database exists. In a recovery situation, you can use the SAP HANA cockpit or the SAP HANA studio to restore customer-specific changes to this initial database.

  Note that if you want to restore customer-specific configuration settings, you can do this either before you restore the database and the log backups, or at the end of the recovery.

- Ensure that the target system and the source system have identical configurations. The number and types of services (for example, index server) on each host must be identical for both system landscapes.
- At the beginning of a recovery, all the data and log backups to be used must be either accessible in the file system or available through the third-party backup tool.
- At least one data backup must be available before the recovery is started.
- To restore the database to a particular point in time, you need a data backup and all the log backups up to the point in time for recovery (including the log backups made after the desired point in time of the recovery).

The following constraints apply to SAP HANA database recovery:

- An SAP HANA database cannot be recovered to an SAP HANA database with a lower software version.

  The SAP HANA software version used for the recovery must always be the same version or higher than that of the SAP HANA database used to create the data backup,
- SAP HANA backups created with release 1.0 SPS10 or newer can be used to recover to SAP HANA 2.0. This is true for both SAP HANA single-container systems and tenant databases. For SAP HANA running on IBM Power systems, different release compatibilities apply.

> Note:
> To recover the SAP HANA database, the database needs to be shut down. For this reason, during recovery, the database cannot be accessed by end users or applications.

Checking of Backups

The success of a database recovery can only be ensured if the required backups are available and have not been changed since they were created. For this reason, manually check backups periodically, or if you suspect that they have been changed in some way since they were created.

When SAP HANA data or log backups are created, the integrity of the data to be backed up is automatically checked while the backups are being written. The data is written to the backup destination only if the integrity check is successful.

When a recovery is started, the integrity of the backups to be used is checked automatically. If an error is detected, the recovery is stopped, and must be repeated.

In addition to the automatic backup checks performed by SAP HANA, you can manually check data backups and log backups without performing a recovery. You can check if all the backups needed for a recovery are available and can be accessed, and if backups have changed or moved since they were first written.

Tools to Perform Manual Backup Checks

Perform manual backup checks using the following tools:

- `hdbbackupcheck`

  Checks whether individual data backups and log backups have changed since they were created.

- `hdbbackupdiag`

  Determines which data backups and log backups are required to complete a recovery. It also checks if these backups are available and if they can be accessed.

> Note:
> Use `hdbbackupcheck` periodically to check the consistency of the metadata of a backup.
>
> `hdbbackupdiag` can be used before the start of a recovery.

> Caution:
> To maintain good recovery performance, and to allow the check to be completed quickly, `hdbbackupdiag` only checks the metadata of a backup. It does not check the integrity of the backup content on the block level.
>
> In some situations, although a backup appears to be consistent and its metadata correct, it might have internal errors. In such cases, use `hdbbackupcheck` to check for corruption in individual data or log backups.

Checking Individual Backups

Use the `hdbbackupcheck` tool to check the integrity of individual data backups and log backups manually.

A data backup of an SAP HANA instance consists of multiple parts, each with the same prefix. A part of a backup is a backup file in the system storage or a backup object that has been transferred to an external backup tool. To check a data backup, start `hdbbackupcheck` for each individual part of the data backup.

The `hdbbackupcheck` tool notifies you if any errors were detected in the checked part of the backup. The notifications are as follows:

- 0: if no errors were detected
- 1: if an error was detected

hdbbackupcheck [parameters] <backup> [-i <backupid>] [-e <ext-backupis>]

```
> hdbbackupcheck 2017_12_06_14_40_39_databackup_3_1
Backup '<directory>/2017_12_06_14_40_39_databackup_3_1' successfully checked.
> hdbbackupcheck -v 2017_12_06_14_40_39_databackup_3_1
Check backup '<directory>/2017_12_06_14_40_39_databackup_3_1'.
Destination of Type: file, Version: 9
Destination header information:
        DatabaseId: 5a12ff7d-a095-3f91-e100-00000a1671e4
        InternalStartTime: 1512571239998 / 2017-12-06T15:40:39+01:00
        CurrDestInformation: [FILE][<directory>/2017_12_06_14_40_39_databackup_3_1]
        backupId: 1512571239938
        ServiceName: indexserver
        NumberOfVolumes: 2
        HostName: wdflbmt7195
        VolumeId: 3
        DestId: 1
        NumberOfDest: 1
        SID: H95
        DatabaseName: H95
        HanaVersion: 2.00.020.00.1500920972
        HanaWeekstone: 0000.00.0
        Architecture: little endian
        WorkerGroups: default
        SrcPoolInformation[0]: [DATABASE_SNAPSHOT]@node[3] BackupId: 1512571239938
        DstPoolInformation[0]: [FILE][<directory>/2017_12_06_14_40_39_databackup_3_1]
    Source header information:
        SrcType: 1
        SourceInformation: [DATABASE_SNAPSHOT]@node[3]
        srcVersion: 6
        sourceSize: 3439333376
        encryption: NOT ENCRYPTED
  Check backup content '[DATABASE_SNAPSHOT]@node[3]'.
  Backup content '[DATABASE_SNAPSHOT]@node[3]' successfully checked.
Backup '<directory>/2017 12 06 14 40 39 databackup 3 1' successfully checked.
```

Figure 337: Example of Output From hdbbackupcheck

Checking the Backups Required for a Recovery

The `hdbbackupdiag` tool determines which backups are required to complete a recovery in a specified point in time. It also checks if these backups are available and if they can be accessed.

hdbbackupdiag

The `hdbbackupdiag` tool is used in the following situations:

- For file-based backups:
  - The backup is available in the file system, either at the location to which it was written or at a location specified by a search path.
  - The current operating system user has read authorization for the file.
  - The actual size of the backup file is the same as the size recorded in the backup file header.
  - The backup ID is identical to the backup ID specified in the backup catalog.
- For backups created using a third-party tool:
  - The backup is available in the third-party tool.

Display all the backups required to recover the database

- In this example the metadata of the backups is not checked. (The `--check` option displays also checks of the metadata.)
- From this list, the backup names can be easily included in shell scripts.
- The time specified is UTC time, not local time.

```
> hdbbackupdiag -f -d /usr/sap/H95/HDB00/backup/log/DB_H95 -u "2017-12-08 10:00:00"
found backup catalog 1512723940142 from file <directory>log_backup_0_0_0_0.1512723940142
using backup catalog 1512723940142 from file <directory>log_backup_0_0_0_0.1512723940142
2017_12_06_14_40_39_databackup_0_1
2017_12_06_14_40_39_databackup_2_1
2017_12_06_14_40_39_databackup_3_1
log_backup_2_0_9022848_9028800.1512572078843
log_backup_2_0_9028800_9033728.1512572978847
log_backup_2_0_9033728_9040000.1512573819340
log_backup_2_0_9040000_9045376.1512574718828
log_backup_2_0_9045376_9051648.1512575560814
log_backup_2_0_9051648_9057152.1512576460728
log_backup_2_0_9057152_9063232.1512577360449
log_backup_2_0_9063232_9068736.1512578260433
 ...
log_backup_3_0_289517440_289618048.1512720912865
log_backup_3_0_289618048_289745216.1512721812875
log_backup_3_0_289745216_289846592.1512722712889
log_backup_3_0_289846592_289947776.1512723612897
```

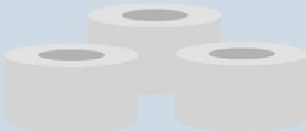

Figure 338: Example of Output From hdbbackupdiag

Database Recovery

Recovery with SAP HANA Multitenant Database Containers

For a recovery with SAP HANA multitenant database containers, note the following information:

- To perform a recovery, an SAP HANA database needs to be shut down.
- To recover a complete SAP HANA system, the system database needs to be recovered first, and then all the tenant databases are recovered individually.
- The recovery of a tenant database is always initiated from the system database. If tenant databases need to be recovered, they are recovered individually, and not all together in one single operation.
- If a tenant database is shut down for recovery the system database and any other tenant databases remain online.
- The system database only needs to be recovered if it is corrupted. If only a tenant database is corrupted, the system database does not need to be recovered.
- If the system database is shut down for recovery all its tenant databases are automatically shut down as well.
- To recover a database, at least one full backup (data backup or data snapshot) must be available.
- When an SAP HANA multitenant database container is recovered, the services needed are generated automatically in the tenant databases.

- With SAP HANA 2.0, an SAP HANA single-container system can be recovered to a tenant database.
- An SAP HANA database cannot be recovered to an SAP HANA database with a lower software version.

Recovering a Tenant Database
To recover a tenant database, proceed as follows:

- Recovery of tenant databases can only be initiated from the system database.
- The system database and other tenant databases are not affected.
- Select the tenant database to be recovered.
- Specify your recovery type and further recovery settings, and start the recovery.

> Note:
> While a tenant database is being recovered, the system database and any other tenant databases remain online.

> Note:
> The system does not support tenant database copy using backup and recovery with third-party tools.
>
> A backup of a tenant database can be recovered to a different SAP HANA multitenant database container using file system-based backups.

Recovering the System Database
To recover the system database, proceed as follows:

- The whole system is shut down, including all tenant databases.
- Specify your recovery type and further recovery settings, and start the recovery.
- The system database is recovered and restarted.
- Restart the tenant databases.
- The content of the tenant databases is not affected by the system database recovery.

> Note:
> If the system database is shut down for recovery, its tenant databases are shut down automatically as well. This means that, until the recovery of the system database is completed, all its tenant databases are unavailable.

A recovery of a tenant database may be required if, for example, a logical error occurred in the tenant database.

If tenant databases need to be recovered, they are recovered individually, and not all together in one single operation.

Recovery of a Database with SAP HANA Cockpit

The options to recover a database using SAP HANA Cockpit are as follows:

- To its most recent consistent state
- To a specific point in time

> Note:
> Using SAP HANA cockpit, only a tenant database can be recovered to a point in time. To recover a system database to a point in time, use SQL.

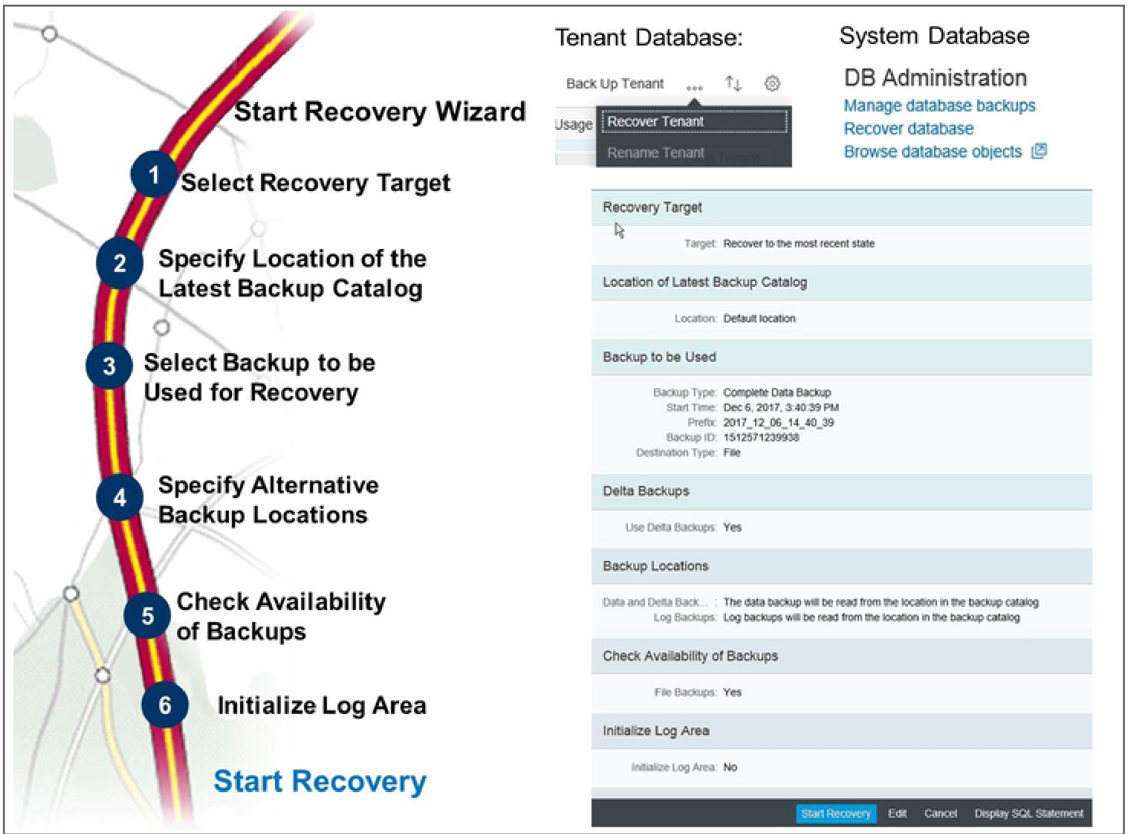

Figure 339: Perform Recovery

To recover an SAP HANA database, perform the following steps:

1. In the SAP HANA cockpit, choose the following:
   - Overall Tenant Statuses , select the tenant database from the overview and choose Recover Tenant .
   - Recover Database to recover the system database.
2. Specify the recovery type from one of the following:
   - Recover the database to the most recent state

- Recover the database to a specified point in time

3. Specify the location of the most recent backup catalog.

   An overview of the available full backups is displayed.

4. Select the complete data backup to be used for the recovery.

5. Specify whether to use delta backups.

6. If you are using backups that differ from those recorded in the backup catalog, specify their locations.

7. Check whether the backups are available.

   Decide if SAP HANA checks before the recovery starts if all the backups are available and if they can be accessed.

8. Specify whether to initialize the log area.

   If you initialize the log area, the content of the log area is lost. No log entries from the log area can be replayed during the recovery. The log entries from the log backups are replayed if they are needed.

   Initialize the log area when the log area is unusable or when you are recovering the database to a different system.

9. Choose Review .

   An overview of the recovery settings displays.

   To change any settings, choose Edit. All the settings that you specified are retained until you change them.

10. To display the SQL statement that is used for the recovery, choose Display SQL Statement .

11. Choose Start Recovery .

    The progress of the recovery for each SAP HANA service displays.

When the recovery is complete, a message confirms this, and shows the timestamp to which the recovery was completed.

The database is restarted automatically after the recovery.

Recover a Database using SAP HANA Studio
The options to recover a database using SAP HANA Studio are as follows:

- To its most recent consistent state
- To a specific point in time
- To a specified full data backup

Before you start the recovery, shut down the SAP HANA database.

To recover an SAP HANA database, perform the following steps:

1. Open the Recovery Wizard.

2. Confirm that the database can be shut down.

3. Specify the recovery type from one of the following:
   - Recover the database to its most recent state
   - Recover the database to a specified point in time
   - To a specified full data backup
4. If the log backups are not in the original location, specify a new location.

   An overview of data backups displays.
5. From the backup catalog, select a complete data backup.
6. Finalize the recovery settings.
7. Start the recovery

The progress of the recovery for each SAP HANA service is displayed in the dialog box. When the recovery is complete, a message confirms this, and it shows the timestamp for when the recovery was completed.

The database restarts automatically after the recovery.

Recovery Features

The recovery features are as follows:

Automatic checks for file system backups at the start of a recovery

In addition to checking for missing backups at the start of a recovery, SAP HANA also checks file system backups for corruption automatically.

Example of a corruption might be if the size or backup ID do not match with the information recorded in the backup catalog. If SAP HANA detects a corruption, the recovery is not started. The details are displayed in the recovery wizard and are written to the backup log file.

> Note:
> The extended checks are executed for file system backups only. If a third-party backup tool is used, only the existence of the backups on the third-party backup server is verified.

Progress reporting for a recovery shows the recovery process in detail

After the initial collection of system information for the recovery, the recovery wizard shows the following phases (progress per service):

- Phase 1: Data recovery

  Using data backup or snapshot
- Phase 2: Log recovery

  Using log backups or log that is still available in the log area
- Phase 3: Restart

Recovery with the Command Line Tool

To recover an SAP HANA database, it is strongly recommended that you use SAP HANA cockpit. You can also recover the system database of a SAP HANA multitenant database container system or an SAP HANA single-container system using the Python script `recoverSys.py`.

SQL statements for recovery cannot be executed using the normal SQL clients such as `hdbsql` and they cannot be executed when the database is online. For this reason, the Python script `recoverSys.py` is used to pass SQL statements to SAP HANA.

> Caution:
> Tenant databases cannot be recovered using the command line tool. To recover a tenant database, use SAP HANA cockpit or the `recoverSys.py` tool to execute SQL statements on the online system database to recover an offline tenant database..

Recovery Procedure Using Command Line Tool

The following is the procedure for performing Recovery using the Command Line Tool:

1. The administrator calls the script with the required parameters, specifying recovery target time, recovery type, and further options.

   ```
   HDBSettings.sh recoverSys.py [<parameters>]
   ```

2. The script stops the SAP HANA database, prepares, and executes the recovery.

3. After the master name server of the SAP HANA database starts successfully, the script terminates.

   Note that, at this point, the recovery is not complete yet. Call the script using the --wait option, because this ensures that the script waits until the recovery finishes.

```
HDBSettings.sh recoverSys.py --wait
[140737354004224, 0.002] >> starting recoverSys (at Tue Mar 14
14:21:39 2017)
[140737354004224, 0.002] args: ()
[140737354004224, 0.002] keys: {'wait': True}
own pid: 26697
recoverSys started: 2017-03-14 14:21:39
testing master: wdflbmt7195
wdflbmt7195 is master
shutdown database, timeout is 120
stop system
stop system: wdflbmt7195
stopping system: 2017-03-14 14:21:39
stopped system: 2017-03-14 14:22:04
creating file recoverInstance.sql
restart database
restart master nameserver: 2017-03-14 14:22:09
start system: wdflbmt7195
2017-03-14T14:22:16+01:00 P026831 15accfbc475 INFO RECOVERY
state of service: nameserver, wdflbmt7195:30001, volume: 1,
RecoveryExecuteTopologyAndSSFSRecoveryFinished

2017-03-14T14:23:16+01:00 P026831 15accfbc475 INFO RECOVERY
RECOVER DATA finished successfully, reached timestamp
2017-03-14T14:21:30+01:00, reached log position 188651136


recoverSys finished successfully: 2017-03-14 14:23:17
[140737354004224, 97.835] 0
[140737354004224, 97.835] << ending recoverSys, rc = 0 (RC_TEST_OK),
after 97.833 secs
```

To check that the recovery is successful, see the backup.log.

Resumption of Recovery

You can resume an interrupted recovery, instead of repeating the entire recovery. It is normally only necessary to resume a recovery in exceptional circumstances.

If a recovery is canceled or interrupted, an SAP HANA database cannot start. Before work can continue in the database, the recovery must be completed. In many situations, a recovery can be repeated from the beginning reasonably quickly, and the database can run again with only minimal loss of work time.

However, having to repeat an interrupted recovery from the beginning may sometimes cause a significant loss of work time. Because of this, the option to resume an interrupted recovery can save a significant amount of time, both with a large or a smaller database.

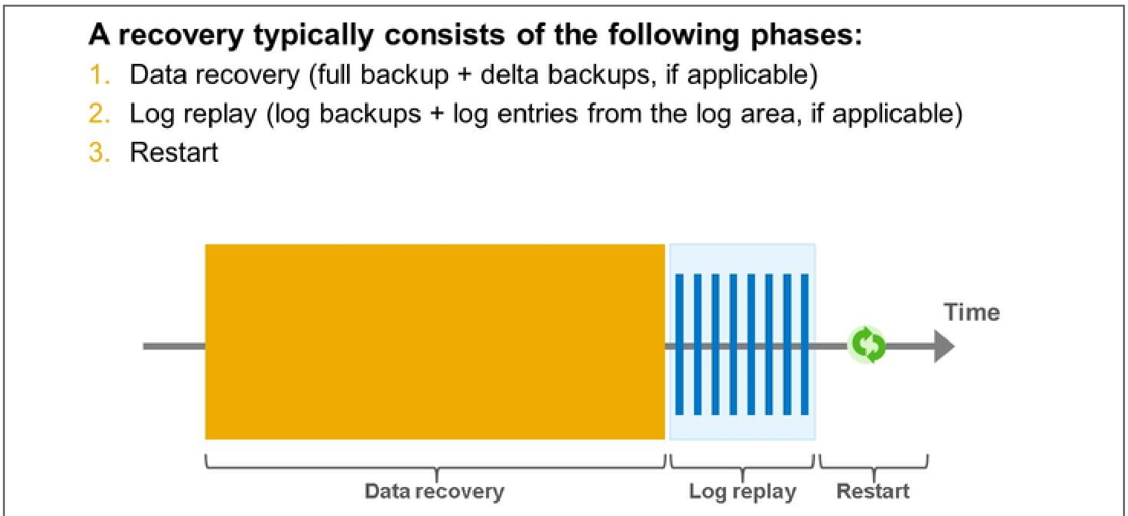

Figure 340: Recovery Phases

Resume Recovery After Error

Data recovery takes up most of the time of a recovery. If a recovery fails during delta data backup recovery or during log replay, SAP HANA can resume the recovery, thus shortening the outage significantly.

A typical example that can cause a failure during log replay is a temporary outage of the backup network.

SAP HANA can resume database recovery if outages occur during delta data backup recovery, or after successfully replayed logs, thus further reducing potential outages. A data recovery can be resumed from the last successfully recovered full backup, or from the successfully recovered delta backups that rely on additional fallback points set between full backups and delta backups during recovery. Failures during log replay can benefit from preserved changes at frequent savepoint intervals, so that only the most recent set of logs after the savepoint must be replayed.

Note:
A recovery from only a full data backup cannot be resumed. If a recovery from only a full data backup is interrupted, repeat the recovery from the beginning.

During a recovery, SAP HANA automatically defines fallback points, which mark the point after which you can resume a recovery.

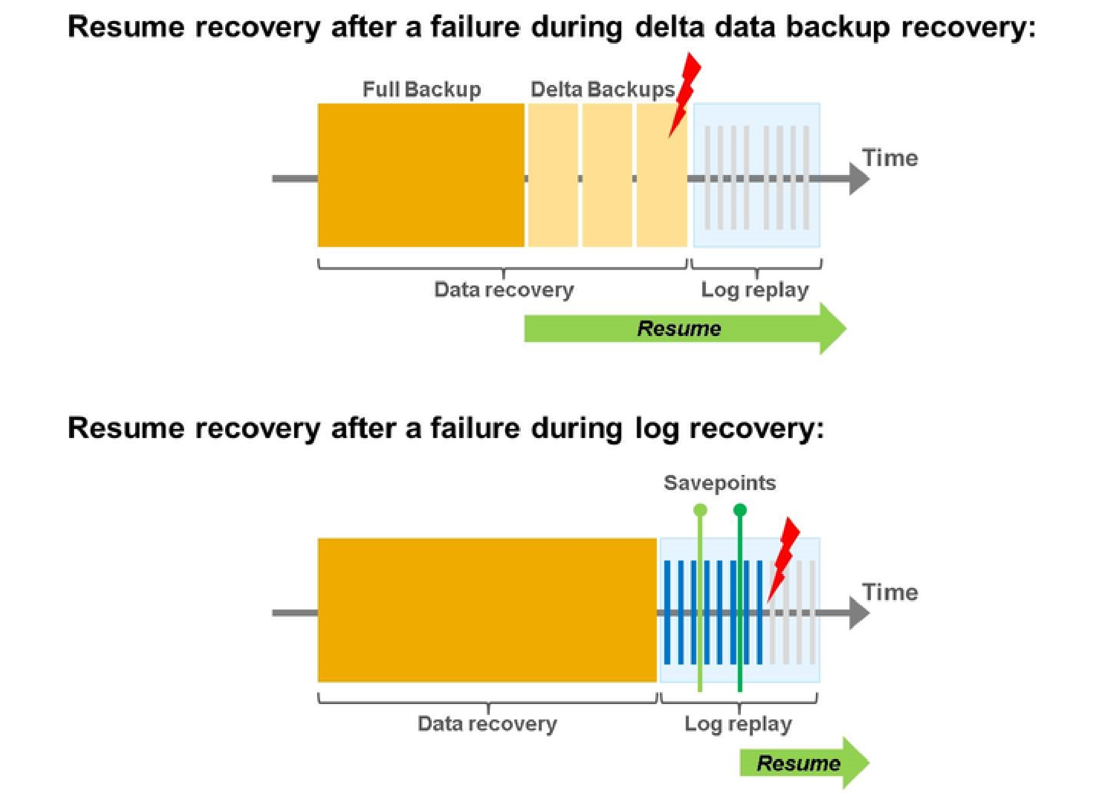

Figure 341: Resume Recovery After Error

If you resume database recovery after a failure during delta data backup recovery, the following occurs:

- Accelerate subsequent recovery using successful recovered full data backups, differential and incremental data backups, and log backups.
- Successfully recovered data backups are kept implicitly as an internal database snapshot (fallback point) in the data area.
- If a recovery failure occurs during delta data backup recovery, the full backup recovery is preserved. Only the delta data backups and the log backups are reprocessed.

If you resume database recovery after a failure during log recovery, the following occurs:

- Data changes caused by replaying log backups during a database recovery are periodically stored in the log area with fallback points.
- The frequency of fallback points can be set using the log_recovery_resume_point_interval configuration parameter (default: 1,800 sec; max: 18,000 sec, 0 = disabled).
- Once the interval expires, the next fallback point log entry is replayed.
- If a recovery failure occurs after a fallback point, only the log backups after most recent fallback point must be reprocessed.

After a recovery has been successfully completed, the fallback points are invalidated. It is then no longer possible to perform a new recovery based on those fallback points.

> Note:
> The fallback points are recorded in the `backup.log` file. The fallback points indicate whether you can resume a recovery.

Perform Recovery After Error

A recovery that can be resumed is also recorded in the backup catalog. In SAP HANA cockpit and SAP HANA studio, a partially completed recovery that can be resumed is given the backup prefix RESUME. To resume the recovery from that backup, proceed as follows:

1. Start the recovery from SAP HANA cockpit.
2. In the recovery dialog, select the backup with the prefix RESUME.
3. Follow the steps described on-screen to complete the recovery.

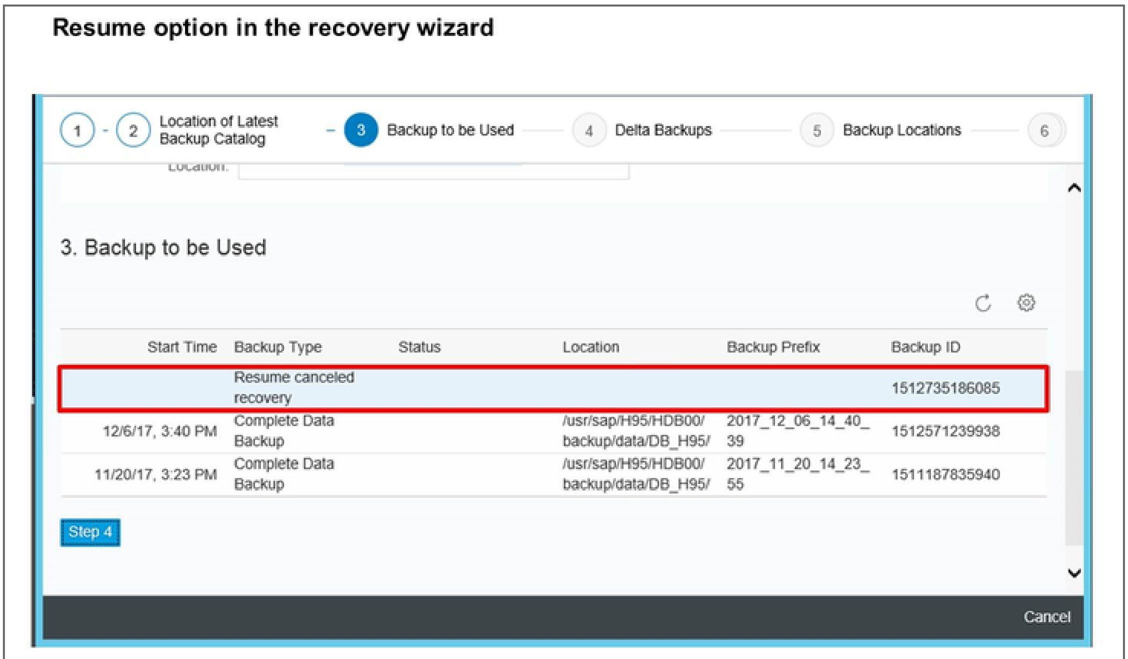

Figure 342: Perform Recovery After Error

LESSON SUMMARY

You should now be able to:

- Perform database recovery

Unit 10
Lesson 6

Explaining Backup and Recovery Using Data Snapshots

LESSON OVERVIEW

This lesson gives a short overview on performing backup and recovery using a storage snapshot.

Business Example

You have to perform backups for the SAP HANA database. Therefore, you need to know how storage snapshots are integrated in the backup concept.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explain backup and recovery using data snapshots

Data Snapshots

Data snapshots offer an additional option to safeguard the SAP HANA data area and to recover an SAP HANA database.

A data snapshot captures the data persisted in the data area at a particular point in time. A data snapshot includes all the data that is required to recover SAP HANA to a consistent state.

Data Snapshots Benefits

Data snapshots have the following benefits:

- They can be created with minimal impact on the system.

  This is because data snapshots are created in the storage system and do not consume database resources.

- Recovery from a data snapshot is faster than a recovery from a data backup.

  The data snapshot only needs to be available in the data area of the storage system. For a recovery based on a data snapshot, you can also use delta backups and log backups in the same way as with a recovery based on a data backup.

If you are planning a backup strategy that makes use of data snapshots, you should be aware of several important points.

- Currently, a data snapshot of an SAP HANA database with more than one tenant is not supported. You can only create and recover a data snapshot of an SAP HANA single-tenant system. To back up SAP HANA systems with more than one tenant database, use data backups.

- Data snapshots can only be created through the system database. It is not possible to create data snapshots for the tenant database separately.
- To create a data snapshot, you need to use native SQL. Recovery from a data snapshot is supported by SAP HANA cockpit and SAP HANA studio.

To create a data snapshot, first create an internal database snapshot. This internal database snapshot provides a view of the database at the point in time that it was started.

The internal database snapshot ensures the consistent state of the data snapshot. This is particularly important if multiple storage volume groups are involved.

> Note:
> As with the data backup types supported by SAP HANA (File or BACKINT), a data snapshot is created while the SAP HANA database is running.

Whereas a data backup is written to a separate storage location, a data snapshot must be manually stored in a location that is physically separate from the SAP HANA data area.

> Note:
> The internal database snapshot reflects a consistent state. When a data snapshot is created, the integrity of the data is not checked.
>
> With data backups, the integrity of the data to be backed up is checked automatically while the backups are created.

An SAP HANA database can be recovered in a single procedure, either using a data snapshot, or using a data snapshot in combination with log backups. You can replay log backups after the database has been recovered with a data snapshot.

Creation of Data Snapshots

You can only create a data snapshot for an SAP HANA single-tenant system. A data snapshot for a database with more than one tenant is not supported. To back up SAP HANA databases with more than one tenant, use data backups.

The SAP HANA database is online, and all the configured services are running.

There is a weak link between SAP HANA and the storage tool. Data snapshots are recorded in the SAP HANA backup catalog.

Creating a Data Snapshot

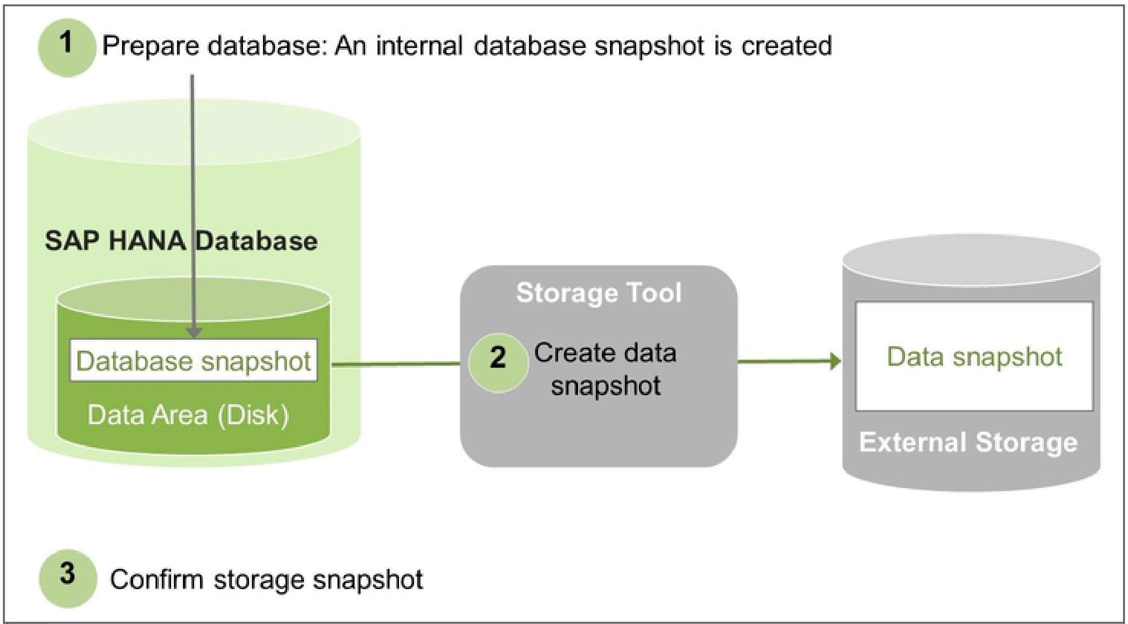

Figure 343: Creating a Data Snapshot

To create a snapshot, proceed as follows:

1. Prepare the database for the data snapshot.

   This creates an internal database snapshot.

2. Create the data snapshot of the SAP HANA data area with the tool provided by your storage vendor.

3. Confirm the data snapshot as successful. An entry with the external backup ID is written to the backup catalog. If the data snapshot cannot be created, or if confirmation fails, choose Abandon .

After you have confirmed or abandoned a data snapshot, it is recorded in the backup catalog as either successful or unsuccessful.

The SAP HANA database automatically deletes the internal snapshot from SAP HANA data area after it has been either confirmed or abandoned.

Creating a Data Snapshot Using SQL Command

Alternatively, you can use the SQL commands to create a data snapshot and to confirm the successful data snapshot. You can enter the external snapshot ID using the following commands:

```
BACKUP DATA CREATE SNAPSHOT COMMENT snapshot_test'
```

```
BACKUP DATA CLOSE SNAPSHOT BACKUP_ID 3456789 SUCCESSFUL
'storage_id_12345'
```

For information on creating data snapshots using SQL commands, see the SAP HANA Administration Guide.

Prepared data snapshots only exist until the data snapshot is executed using the storage tool. When a data snapshot is prepared but not confirmed for a longer period of time, an alert occurs. For more information, see SAP Note 1991615 .

Procedure to Create a Data Snapshot

1. Create a new internal database snapshot. Use the following SQL statement:

   ```
   BACKUP DATA FOR FULL SYSTEM CREATE SNAPSHOT [COMMENT <STRING>]
   ```

   Optionally, add a comment. This comment helps you to identify the data snapshot in the backup catalog.

2. Find out the backup ID of the internal database snapshot in the state PREPARED.

   ```
   SELECT * FROM M_BACKUP_CATALOG WHERE ENTRY_TYPE_NAME = 'data
   snapshot'
   ```

3. In the storage system, make all the files and directories from the data area available together in a separate storage location. To create the data snapshot, you can use the tool provided by your storage vendor.

4. Confirm or abandon the data snapshot.

   Confirm:
   ```
   BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID
   <BACKUP_ID> SUCCESSFUL <STRING>
   ```

   Abandon:
   ```
   BACKUP DATA FOR FULL SYSTEM CLOSE SNAPSHOT BACKUP_ID
   <BACKUP_ID> UNSUCCESSFUL <STRING>
   ```

> Note:
> It is strongly recommended to confirm or abandon a data snapshot as soon as possible after it has been created.
>
> While the data snapshot is being prepared or created, the snapshot-relevant data is frozen. While the snapshot-relevant data remains frozen, changes can still be made in the database. Such changes will not cause the frozen snapshot-relevant data to be changed. Instead, the changes are written to positions in the data area that are separate from the data snapshot. Changes are also written to the log.

Recovery with a Data Snapshot

If you are using a data snapshot for your recovery, first transfer it back to the data area of the SAP HANA database using the storage tool.

To use a snapshot for recovery, proceed as follows:

1. Using the storage tool, transfer the data snapshot to the data area of the SAP HANA database.

2. Using SAP HANA cockpit or SAP HANA studio, recover the database with the data snapshot. This is available in the recovery wizard.

   Note that all recovery options are available, including point-in-time recovery using log backups from the log area.

> Note:
> You can also call up the recovery wizard before transferring the data snapshot to the data area of SAP HANA. In this case, the recovery wizard shows all of the data snapshots recorded in the SAP HANA backup catalog. You can decide which one to transfer to the data area of SAP HANA.

After the recovery, SAP HANA automatically deletes the internal data snapshot from the data area, which was contained in the transferred data snapshot.

The system does not currently support recovery of delta backups using a data snapshot.

LESSON SUMMARY

You should now be able to:

- Explain backup and recovery using data snapshots

Unit 10
Lesson 7

Explaining Database Copy

LESSON OVERVIEW

This lesson describes how you can clone the database.

Business Example

To set up a three-system landscape, you have to clone your SAP HANA database.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explain the scenarios for a database copy

Copying of a Database Using Backup and Recovery

You can create a homogeneous copy of an SAP HANA database by using backups of that database to recover to the same system or a different system. A homogeneous database copy is a quick way to set up a cloned database, for example, for training, testing, or development.

You can use backup and recovery to copy a system database or a tenant database within the same system or to a different system.

The following combinations of source database and target database can be used to create a database copy:

Source Database	Target Database
System database	The system database of the same system The system database of a different system
Tenant database	A different tenant database in the same system A tenant database in a different system
Single-container system	Tenant database Note: An SAP HANA backup created with SAP HANA 1.0 SPS10 (single-container system) or newer can be used to recover a tenant database.

Figure 344: Combinations to Create a Database Copy

Prerequisites for a Copying a Database Using Backup and Recovery

Before you can create a copy of an SAP HANA database, some important preparations are needed.

Prerequisites for a Database Copy

- The version of the SAP HANA target database is the same or higher than the SAP HANA source database.
- You can copy a database to machines from different vendors and with different hardware configurations, provided that both the source and target machines are compliant with the SAP HANA appliance specifications.

  Special requirements may apply to ensure the compatibility of SAP HANA backups with IBM Power Systems.
- To copy a complete SAP HANA system, the system database needs to be recovered first, and then all the tenant databases are recovered individually.
- For the system database, you must have the logon credentials of the operating system user (<sid>adm).

  For a tenant database, the system database user must have the authorization DATABASE ADMIN.
- If you expect a different set of volumes to be recovered, before you start the recovery for a database copy, you should remove existing data and log volumes.

  After a recovery to create a database copy, the system may include different volumes, or volumes may be assigned to different hosts.

  Existing volumes that are not used for the new system will not be overwritten or removed. Any additional disk space is not released. This may lead to unexpected disk full situations.
- A valid license key is available for the target database.
- The target database has sufficient disk space and memory.
- For a database copy using SAP HANA cockpit, the target database must be at least SAP HANA 2.0 SPS 01.

Copy a Database Using a Database Backup

You can copy an SAP HANA database using file-based backups or backups created using third-party tools.

> Note:
> You can copy an SAP HANA database using file-based backups or backups created using third-party tools.
>
> The backup catalog, the data backups, and the log backups must be from either only a third-party backup tool or only the file system.
>
> (For a standard database recovery, it is possible to use a combination of backups from a third-party tool and the file system, provided that the backups originate from the same SAP HANA database.)

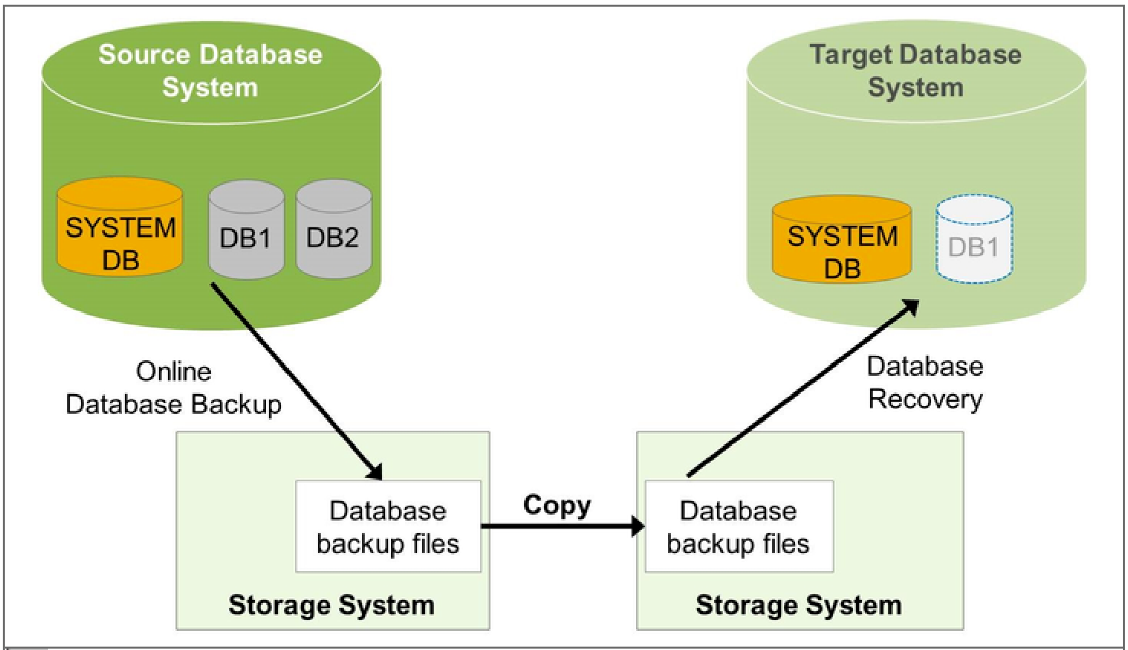

Figure 345: Copy a Tenant Database using a Database Backup

Using SAP HANA cockpit, you can create a copy of an SAP HANA database by using backups of that database to recover to the same system or a different system.

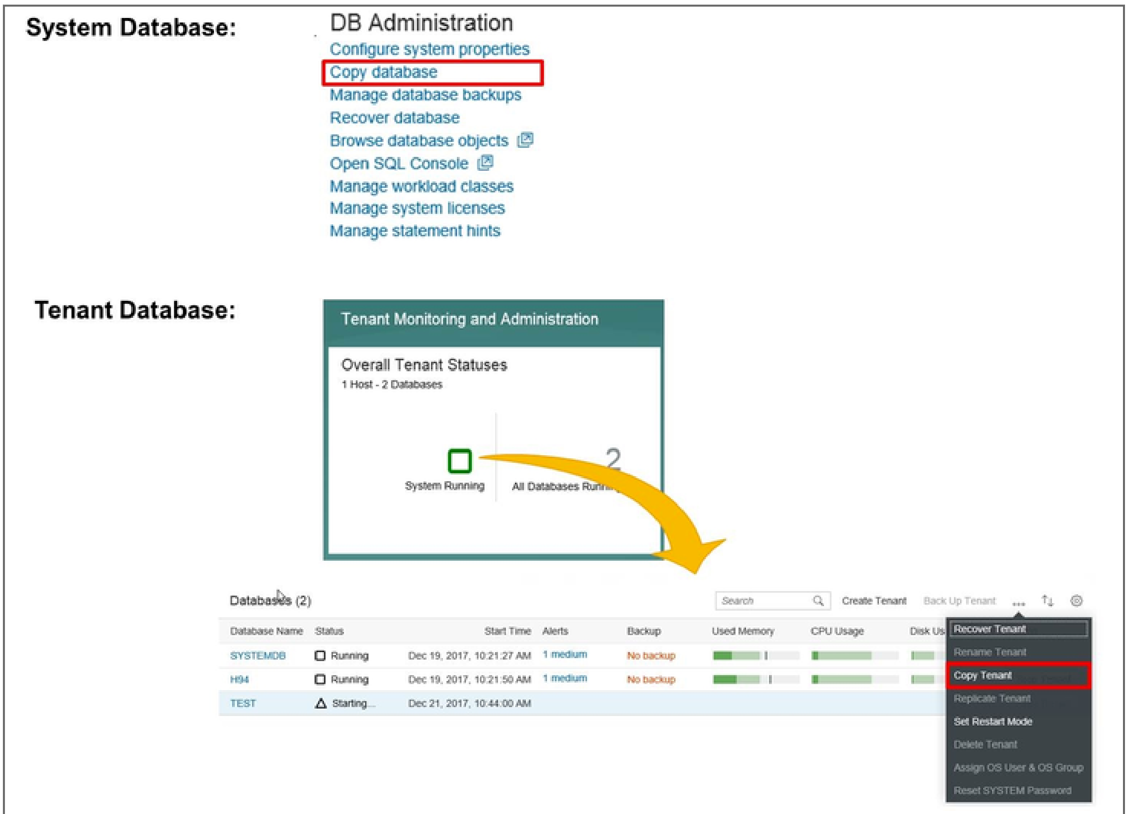

Figure 346: Copy Database in SAP HANA Cockpit

If you are copying a database using a full data backup only, you can either select the data backup from the backup catalog, or specify its location without using a backup catalog.

A copy to a point in time is not possible if the full data backup is not recorded in the backup catalog.

Copy a Database Using a Data Snapshot

> Caution:
> Using a data snapshot, it is only possible to copy an SAP HANA system with a single tenant database.
>
> If you attempt to use a data snapshot to copy an SAP HANA system with more than one tenant database, this may make the data area for all the tenant databases unusable.

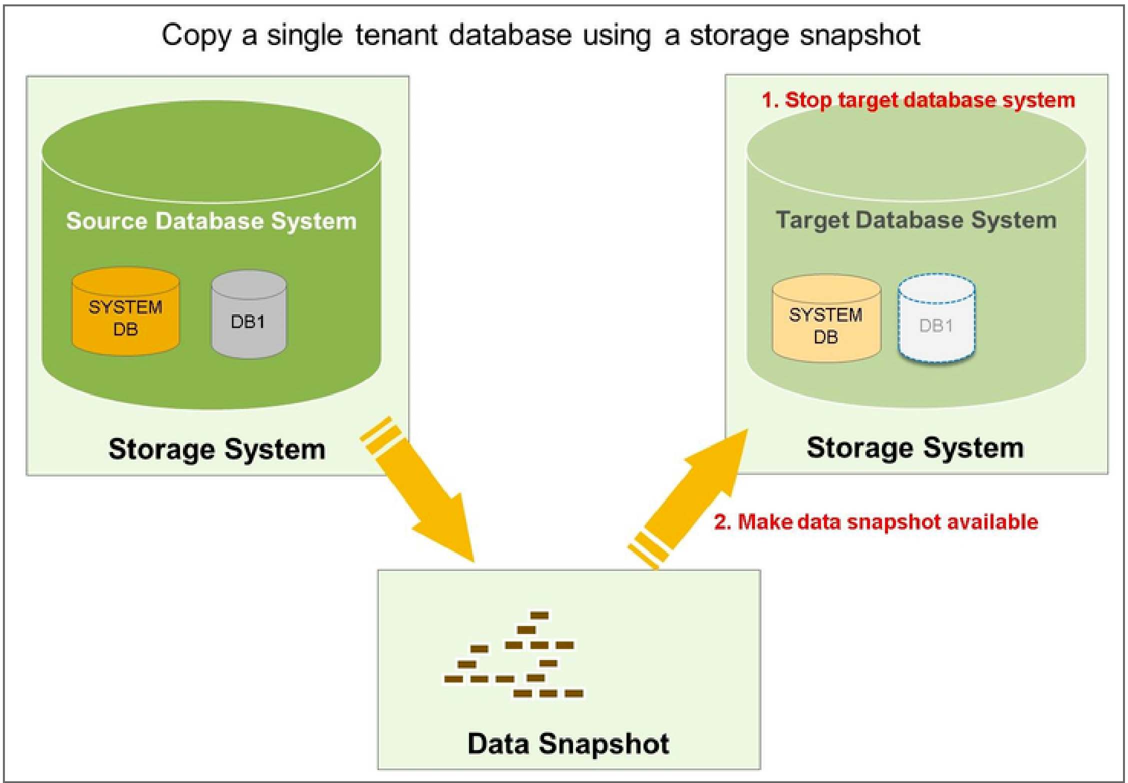

Figure 347: Database Cloning Using a Storage Snapshot

For a database copy using data snapshots, the number of hosts and the number and type of services assigned to each host must be the same for the source database and the target database, and the mountpoint IDs must be identical.

Database Copy from System with m Nodes to a System with n Nodes

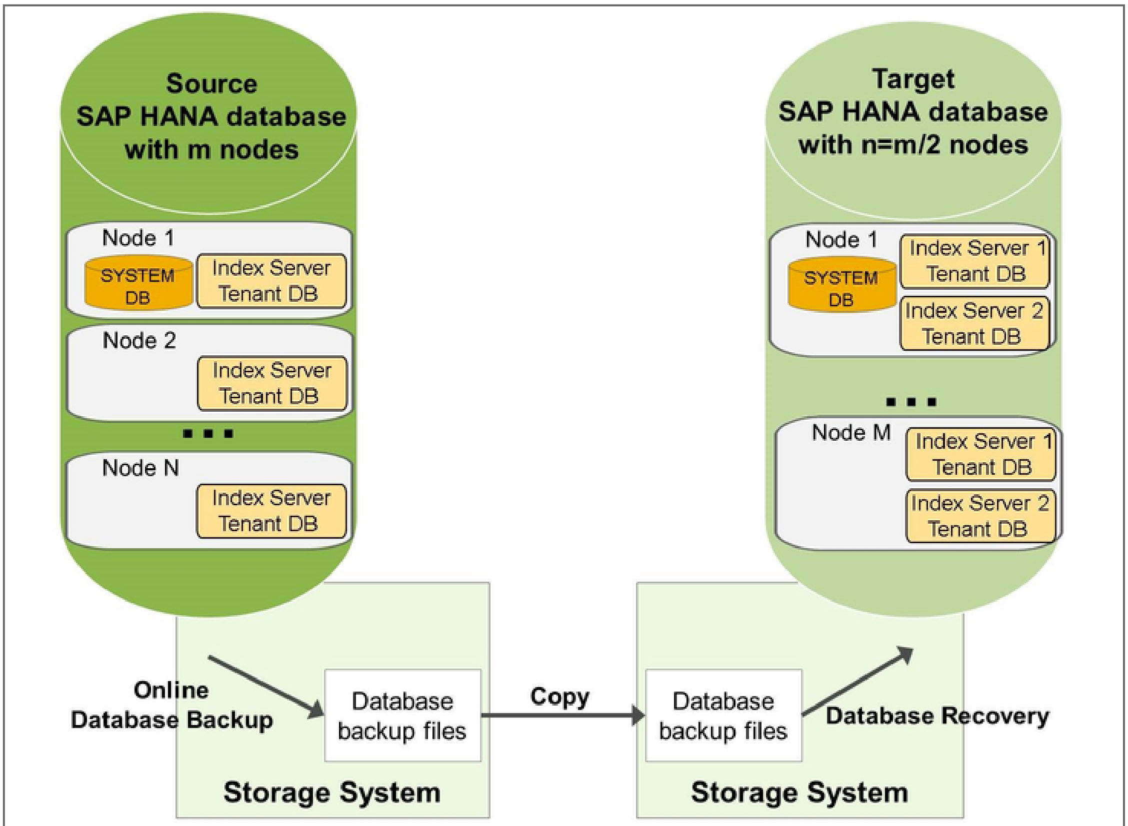

Figure 348: Database Copy from System with m Nodes to a System with n Nodes

You can copy a scale-out SAP HANA database with m nodes to an SAP HANA database with n nodes (m > n). This is useful for when you want to use a copy of your production system for tests on a smaller QA system, for example.

To perform a database copy from a system with m nodes to a system with n nodes, proceed as follows:

1. Create a data backup of the source database.
2. In the target database, configure (m-n) additional index servers to match the source system configuration (.ini file parameter).

   You can choose how you want to distribute these index servers across the available nodes.
3. Recover the data backup of the source database into the target database.

   Note that, before the recovery is executed on the target system, SAP HANA checks if it has been configured correctly.

Copying or Cloning of an SAP HANA System

You can use the SAP HANA database lifecycle manager (HDBLCM) to make a copy or a clone of an SAP HANA system by copying the file system containing the SAP HANA database installation from an old storage solution to a new storage solution, and registering the copied SAP HANA system on new hosts.

Before cloning the SAP HANA system, you must create a physical copy of the SAP HANA system (storage snapshot, file systems copy). The source system must be offline or a database snapshot must have been taken on the source system before the physical copy of the SAP HANA system is created.

> Note:
> An SAP HANA system can only be cloned or copied to a target system that runs on the same hardware platform as the source system.

Cloning an SAP HANA system produces a new SAP HANA system, identical to the existing one. Copying an SAP HANA system produces a new SAP HANA system with the same landscape as the existing one, but slightly different parameter settings. If the interactive parameter defaults are accepted during host registration, the system is effectively cloned. If the new system parameters are set to different values, the new system is similar, but not identical to the source system.

You could, for example, copy an existing production system, and accept all parameter defaults during host registration except **`system_usage`**, which would be specified as "test". This configuration would allow you to have an almost identical copy of the existing system for test or quality assurance purposes.

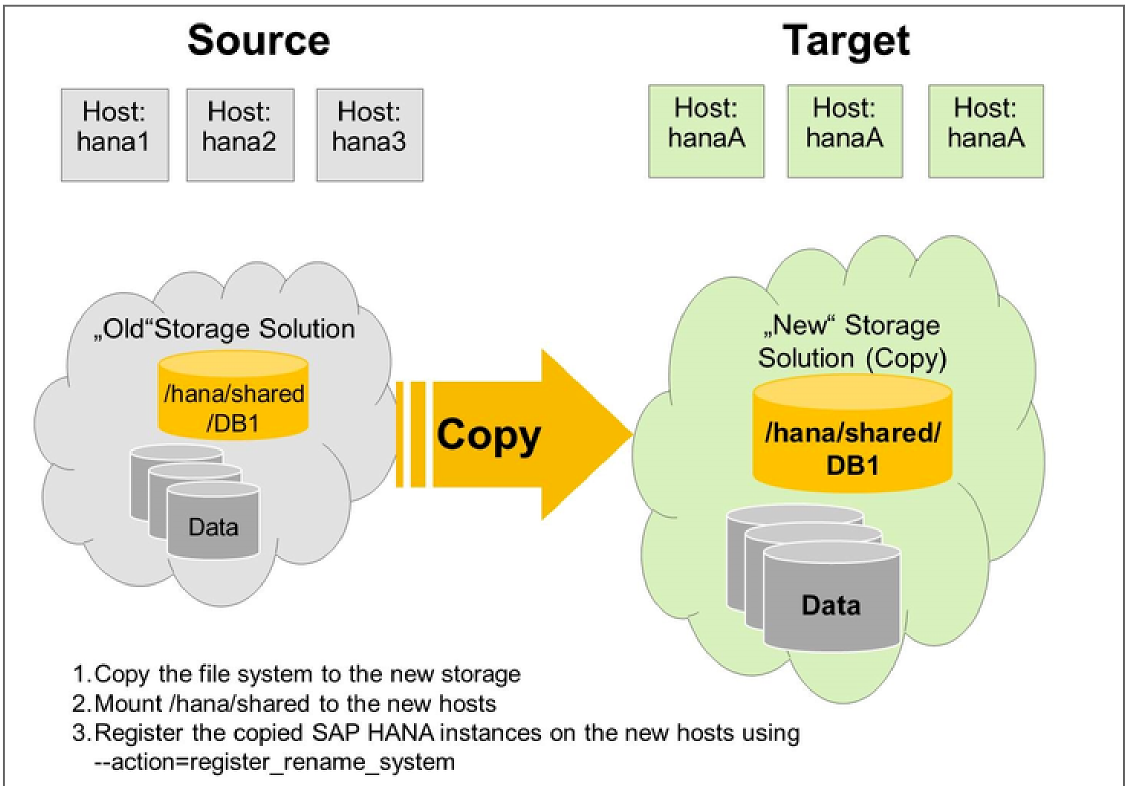

Figure 349: SAP HANA System Clone

Copying and Movement of Tenant Databases Between Systems

Copying and moving tenant databases between systems allows you to manage your system landscape efficiently and to respond flexibly to changing resource requirements. Therefore, SAP HANA uses system replication mechanisms. These mechanisms allow you to copy and move tenant databases securely and conveniently with near-zero downtime.

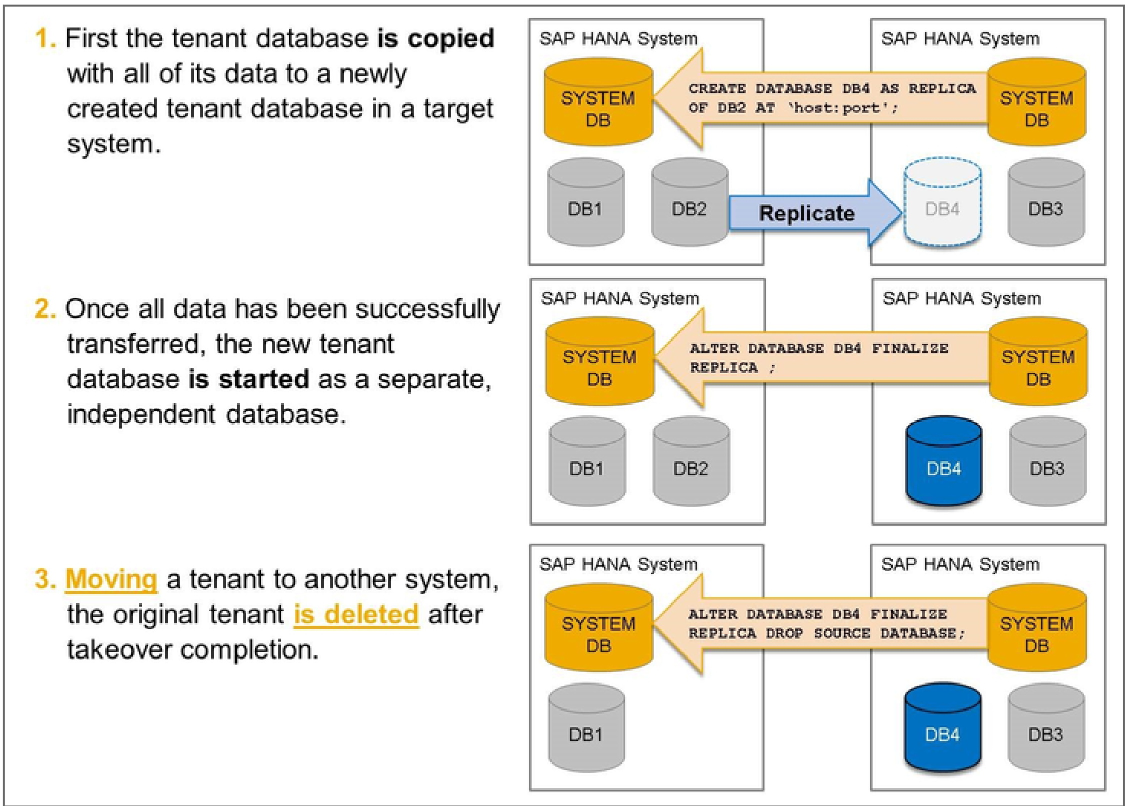

Figure 350: Copying and Moving Tenant Databases

The only difference between copying and moving a tenant database therefore is what happens to the original tenant database after all data has been transferred to the new tenant database in the target system.

In both cases, the new tenant database starts running as a fully separate, independent database.

Several tenant databases can be copied or moved to a system at the same time. It is also possible to copy or move a tenant database to a system with a different isolation level than the source system.

Use Cases for Copying and Movement of Tenant Databases

Copying and moving a tenant database from one system to another in this way has several applications.

Use cases

- Load balancing between systems

  For example, a tenant database is running a more demanding workload than anticipated, so you move it to a system running on a host with more CPU resources.

- Management of deployment environment

  For example, you want to copy a tenant database running in your test system to the live production system.

- Tenant-database-specific upgrades

  For example, you want to upgrade a single tenant database but not the entire system, so you move the tenant database to a system already running the higher version.

- Template databases

  For example, you create a tenant database with a default configuration that you want to reuse as the basis for new tenant databases in other systems. You can simply copy the tenant database as a template to other systems.

LESSON SUMMARY
You should now be able to:

- Explain the scenarios for a database copy

Unit 10

Learning Assessment

1. Which of the following tasks requires the BACKUP ADMIN authorization?

   Choose the correct answers.

   - [ ] **A** Backing up the database without a user interface.
   - [ ] **B** Recovering the database without a user interface.
   - [ ] **C** Physically deleting data and log backups and obsolete versions of the backup catalog from the backup location.
   - [ ] **D** Administration tasks executed on a tenant database through the system database.

2. Which of the following elements are valid rules for the location of the backup files?

   Choose the correct answers.

   - [ ] **A** The Volume ID is automatically added to a specified file name.
   - [ ] **B** The backup files should never be in the same directory and on the same file system as the data.
   - [ ] **C** The backup ID is unique for each persistence only.
   - [ ] **D** The backup folder will be created automatically when the backup process is started.

3. Your SAP HANA database is running in log mode “normal” and automatic log backups “enabled”. With this configuration, in which situations are the log segments backed up automatically?

   Choose the correct answers.

   - [ ] **A** The log segment is full.
   - [ ] **B** The database is started.
   - [ ] **C** The log segment is closed after exceeding the configured time threshold.
   - [ ] **D** The database is stopped.

4. To diagnose backup errors you can use the backup.log and backint.log files.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

5. Which of the following recovery types is not supported for delta backups?

   Choose the correct answer.

   - [ ] **A** Recover the database to its most recent state.
   - [ ] **B** Recover the database to a specific point in time.
   - [ ] **C** Recover the database to a specific data backup.

6. The creation of storage snapshots has a big impact on the system, and the recovery from storage snapshots takes longer.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

7. Which of the following elements are required for implementing a database copy?

   Choose the correct answers.

   - [ ] **A** The target database must have a the same or higher version than the source database.
   - [ ] **B** You don’t have to be concerned about disk space and memory of the target database.
   - [ ] **C** Source and target databases can have different number and types of services.
   - [ ] **D** A data backup (file-based backups or backups using a third-party backup tool) or a storage snapshot of the source system is available.

Unit 10

Learning Assessment - Answers

1. Which of the following tasks requires the BACKUP ADMIN authorization?

   Choose the correct answers.

   [X] **A** Backing up the database without a user interface.

   [ ] **B** Recovering the database without a user interface.

   [X] **C** Physically deleting data and log backups and obsolete versions of the backup catalog from the backup location.

   [ ] **D** Administration tasks executed on a tenant database through the system database.

   Correct! The required authorizations for backing up the database without a user interface are BACKUP ADMIN or BACKUP OPERATOR (recommended for batch users only). The required authorization for physically deleting data and log backups and obsolete versions of the backup catalog from the backup location is BACKUP ADMIN. Recovering the database without a user interface is supported for an SAP HANA single container system or the system database in an SAP HANA multitenant data-base container. The recovery is executed as the operating system user (<sid>adm). You therefore require the logon credentials of this user. The required authorization for administration tasks executed on a tenant database through the system database is DATABASE ADMIN. Read more on this in the lesson Explaining Backup and Recovery (Unit 10, Lesson 1) of the course HA200_14.

2. Which of the following elements are valid rules for the location of the backup files?

   Choose the correct answers.

   [X] **A** The Volume ID is automatically added to a specified file name.

   [X] **B** The backup files should never be in the same directory and on the same file system as the data.

   [X] **C** The backup ID is unique for each persistence only.

   [ ] **D** The backup folder will be created automatically when the backup process is started.

   Correct! The system adds a unique <suffix> to each backup file name that indicates the volume ID and the partition ID. For improved data safety, specify a path to an external backup location. The backup location should never be on the same file system as the data or log areas. The backup folder must already exist before the backup process is started. Read more on this in the lesson Performing Data Area Backup (Unit 10, Lesson 2) of the course HA200_14.

3. Your SAP HANA database is running in log mode “normal” and automatic log backups “enabled”. With this configuration, in which situations are the log segments backed up automatically?

   Choose the correct answers.

   - [X] **A** The log segment is full.
   - [X] **B** The database is started.
   - [X] **C** The log segment is closed after exceeding the configured time threshold.
   - [ ] **D** The database is stopped.

   Correct! After a log segment has been backed up, SAP HANA can overwrite the space in the log area that the log segment occupied with new log entries. If the log area becomes full and no more log segments can be created in the file system, the database freezes. No more log entries can be written until a log backup has been completed and the log segments are no longer needed to restart the database. The log_backup_timeout_s parameter forces log backups at a fixed time interval, specified in seconds. Log backups triggered by log_backup_timeout_s are performed in addition to the log backups that are performed when a log segment becomes full. There is no backup process when the database is stopped. Read more on this in the lesson Configuring a Log Area Backup (Unit 10, Lesson 3) of the course HA200_14.

4. To diagnose backup errors you can use the backup.log and backint.log files.

   Determine whether this statement is true or false.

   - [X] True
   - [ ] False

   Correct! The backup.log records information about data backups, log backups, the progress of a backup, and the backup catalog. It also records information about recovery operations. The backint.log file contains information about the activities of the BACKINT agent. The BACKINT agent is part of a third-party backup tool. The information recorded can be used to diagnose errors. Read more on this in the lesson Describing Additional Backup Topics (Unit 10, Lesson 4) of the course HA200_14.

5. Which of the following recovery types is not supported for delta backups?

   Choose the correct answer.

   - [ ] **A** Recover the database to its most recent state.
   - [ ] **B** Recover the database to a specific point in time.
   - [x] **C** Recover the database to a specific data backup.

   Correct! The recover the database to a specific data backup option uses the specified data backup and is not supported for delta backups. Log entries are not replayed, neither from the log backups nor from the log area. All log entries that still exist in the log area are deleted. The recover the database to its most recent state option recovers the database to as close as possible to the current time. For this it uses the most recent data backup, all log backups made since the most recent data backup and the log area. The recover the database to a specific point in time option uses the last data backup available before the specified point in time, the log backups made since the data backup to be used and the log area. Read more on this in the lesson Performing Database Recovery (Unit 10, Lesson 5) of the course HA200_14.

6. The creation of storage snapshots has a big impact on the system, and the recovery from storage snapshots takes longer.

   Determine whether this statement is true or false.

   - [ ] True
   - [x] False

   Correct! They can be created with minimal impact on the system. This is because storage snapshots are created in the storage system and do not consume database resources. Recovery from a storage snapshot is faster than a recovery from a data backup. The storage snapshot only needs to be available in the data area of the storage system. For a recovery based on a storage snapshot, you can also use delta backups and log backups in the same way as with a recovery based on a data backup. Read more on this in the lesson Explaining Backup and Recovery Using Data Snapshots (Unit 10, Lesson 6) of the course HA200_14.

7. Which of the following elements are required for implementing a database copy?

   Choose the correct answers.

   - [x] **A** The target database must have a the same or higher version than the source database.
   - [ ] **B** You don’t have to be concerned about disk space and memory of the target database.
   - [ ] **C** Source and target databases can have different number and types of services.
   - [x] **D** A data backup (file-based backups or backups using a third-party backup tool) or a storage snapshot of the source system is available.

   Correct! The target database must have a the same or higher version than the source database, as backward compatibility is not possible. A data backup (file-based backups or backups using a third-party backup tool) or a storage snapshot of the source system is available to bring the target database to the same state as the source database or to recover it to a desired point in time. The target database should have sufficient disk space and memory. If the target system has fewer resources, for example, less CPU and RAM, you cannot expect the performance to be the same as in the source system. For a database copy using storage snapshots, the number of hosts and the number and type of services assigned to each host must be the same for the source database and the target database. The mountpoint IDs must also be identical. Read more on this in the lesson Explaining Database Copy (Unit 10, Lesson 7) of the course HA200_14.

UNIT 11 Security

UNIT OBJECTIVES

- Describe the SAP HANA security functions
- Explain encryption
- Describe auditing

Unit 11
Lesson 1

Describing Security Functions

LESSON OVERVIEW

This lesson gives an overview of the security functions in SAP HANA.

Business Example

Depending on the implementation scenario, the SAP HANA database facilitates the integration of different security functions. Therefore, you need an overview of the supported security functions.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Describe the SAP HANA security functions

Security Perspective in Different Implementation Scenarios

The way in which you implement SAP HANA determines your security requirements.

Implementation Scenarios

The following main SAP HANA scenarios are possible:

- SAP HANA as a data mart for reporting and analytics
- SAP HANA in a classic three-tier architecture as the primary database. For example, in SAP BW/4HANA or SAP S/4HANA installations.
- SAP HANA as a platform for providing database and application services to native SAP HANA based applications

For more information about security for SAP HANA in the different scenarios, see the SAP HANA Security Guide.

Data Mart Scenario

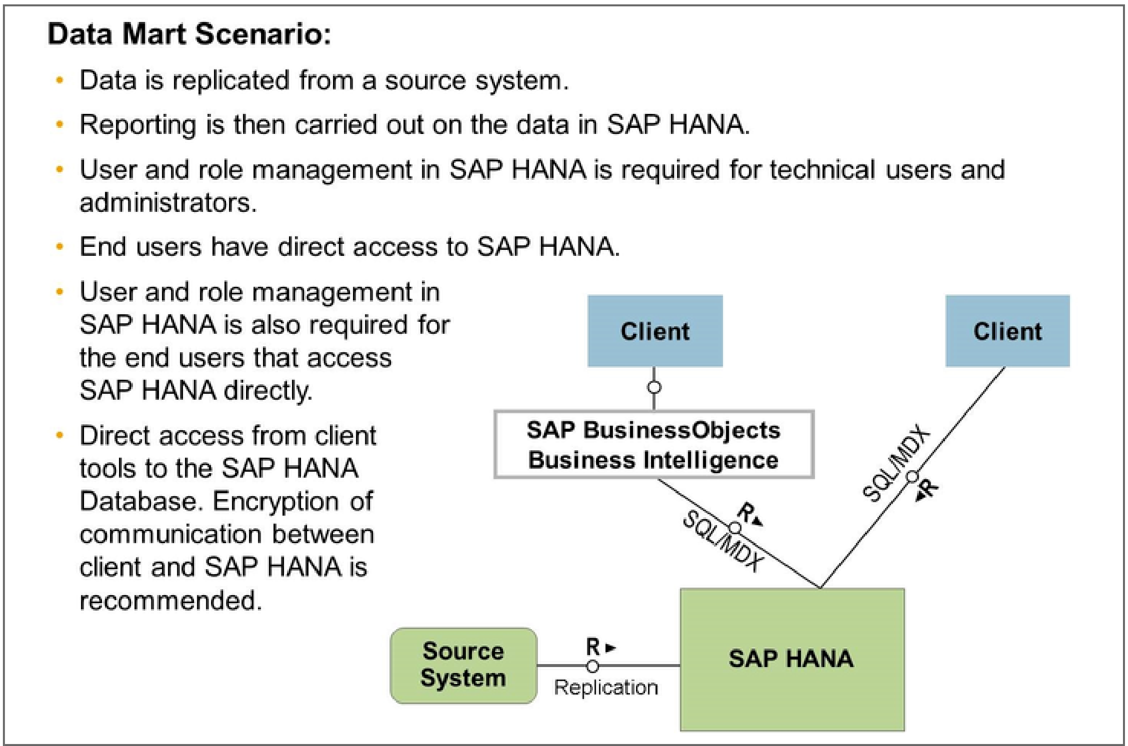

Figure 351: Data Mart Scenario

In a data mart scenario, data is replicated from a source system, such as SAP Business Suite, into the SAP HANA database. Reporting is then carried out on the data in SAP HANA (for example, using read-only views, dashboards, and so on). The following architectures can be used in this scenario:

- The implemented architecture determines how security issues are handled in SAP HANA.
- Some end users usually have direct access to SAP HANA. Therefore, user and role management in SAP HANA is required for these end users, as well as for technical users and administrators.
- SAP HANA security features are required for other security aspects, such as end user authorization for views.

SAP HANA in a Classic Three-Tier Architecture

SAP HANA in a Classic Three-Tier Architecture:

- Security-related features are located an enforced primarily in the application server layer.
- The database is used as a data store only.
- End users do not have direct access to database.
- The same security model for user access applies as with other databases.
- Security in the database layer is mainly focused on securing administrative access to the database.
- User and role management is only required for administrators.
- SAP HANA Database is not accessible from the client network directly.

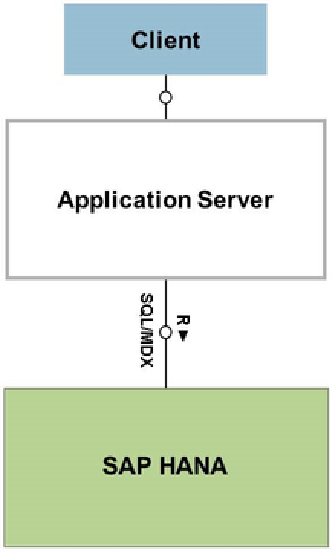

Figure 352: SAP HANA in a Classic Three-Tier Architecture

You can use SAP HANA as a relational database in a classic three-tier architecture (client, application server, and database). Security-related features are located and are enforced in the application server layer. These security features include authentication, authorization, encryption, and auditing. The database is used as a data store only.

The classic three–tier architecture has the following features:

- The same security model for user access applies as for other databases.
- End users do not have direct access to either the database itself or the database server on which it is running.
- Security in the database layer is mainly focused on securing administrative access to the database.
- Specific SAP HANA security features are needed to control access of administrators to the database.

SAP HANA as a Platform

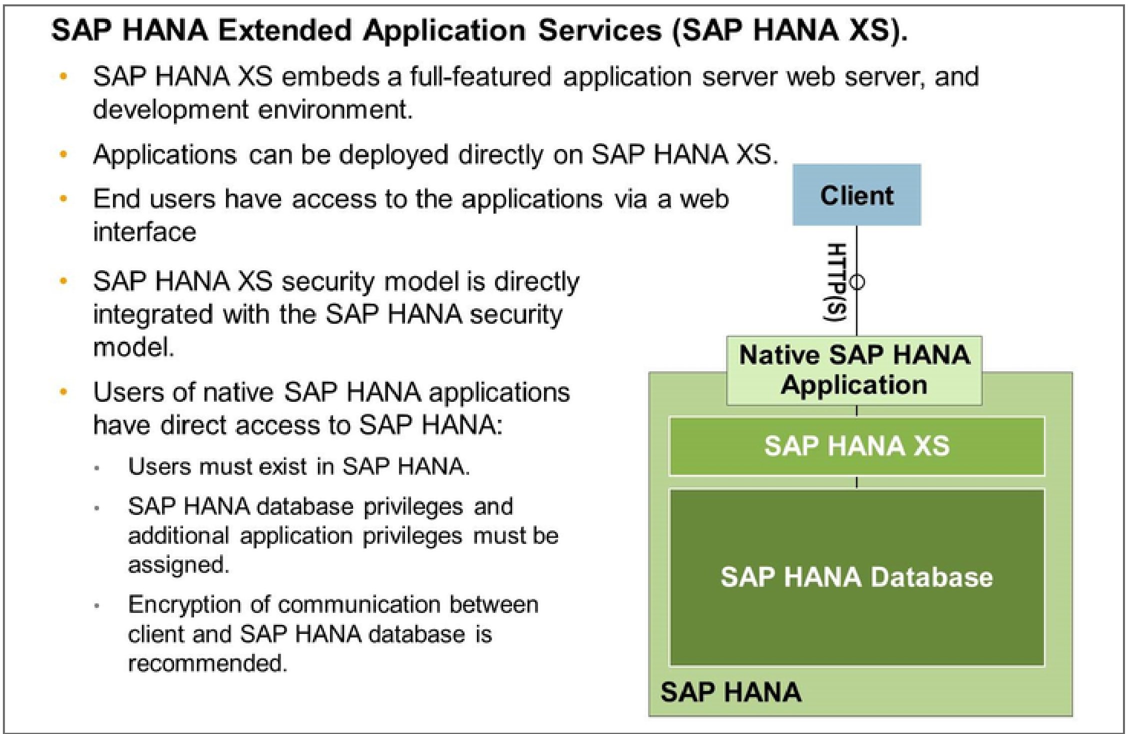

Figure 353: SAP HANA as a Platform

SAP HANA includes SAP HANA extended application services. SAP HANA extended application services embeds a full-featured application server, Web server, and development environment within SAP HANA. Applications can be deployed directly on SAP HANA extended application services. It exposes these applications to end users through a web interface.

It has the following requirements:

- The security model for SAP HANA extended application services is directly integrated with the SAP HANA security model.
- Users of native SAP HANA applications have direct access to SAP HANA:
  - Users must exist in SAP HANA.
  - SAP HANA database privileges and additional application privileges must be assigned.

Security Functions in SAP HANA

The security functions in SAP HANA include the following features:

- SAP HANA provides security features that enable you to implement different security policies and meet compliance requirements.
- Depending on the implementation scenario in which SAP HANA is used, only some of these features might be needed; others might be provided in other architecture layers.
- SAP HANA supports standard interfaces so that the customer security network and data center infrastructures can be integrated.

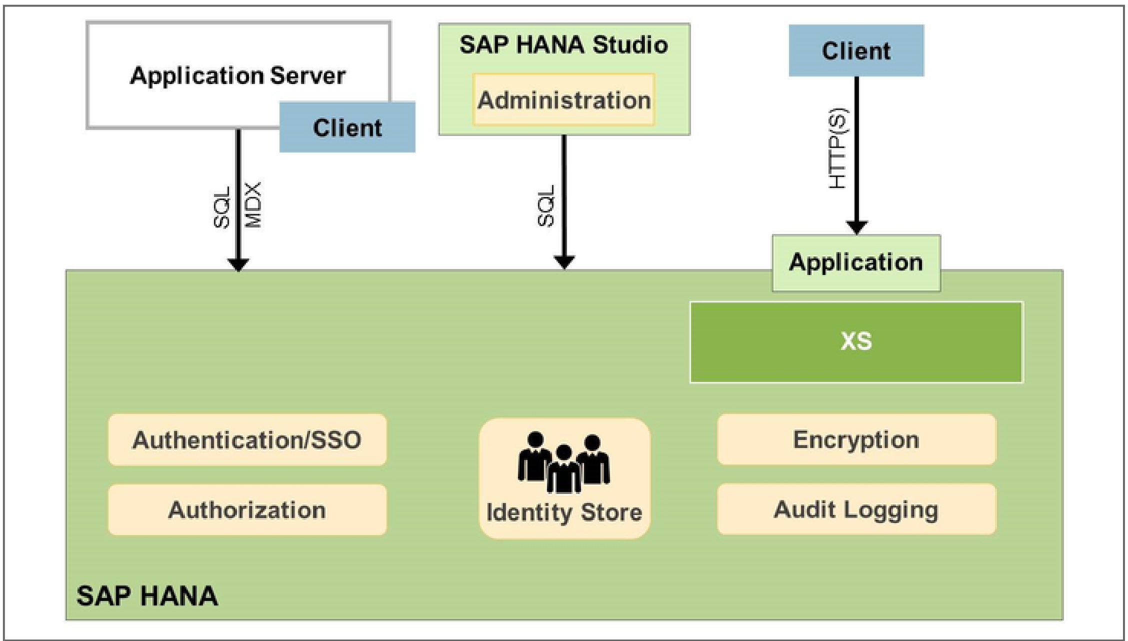

Figure 354: Overview: Security Functions

Data Center Integration

The security infrastructure of SAP HANA facilitates integration in the data center.

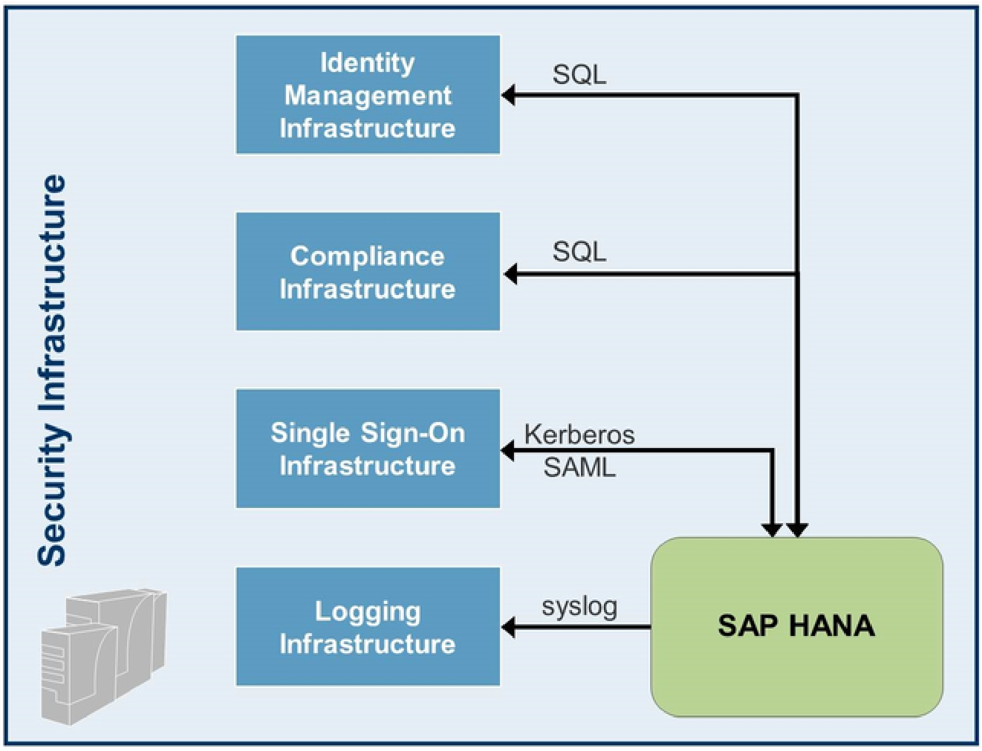

Figure 355: Security Infrastructure

The security infrastructure of SAP HANA includes the following features:

- User and role provisioning solutions

  - Out-of-the-box connector for SAP Identity Management
  - SQL interface for integration with other identity management solutions
- Compliance infrastructure
  - Out-of-the-box connector for SAP Access Control 10.1
- Standard-based single sign-on infrastructures
  - For example, Microsoft Active Directory
- Existing logging infrastructures
  - Database audit trail

Security Administration and Monitoring

Many security-related features are integrated into SAP HANA cockpit.

The security dashboard in SAP HANA Cockpit provides an overview of the important security KPIs during operation of your system.

SAP HANA Cockpit can be used for the following security administration, configuration, and monitoring tasks:

- User and role management
- Management of privileges
- Management of audit policies
- Configuration of password policy
- Certificate management

> Note:
> Most administration tasks can also be carried out using SQL commands.

SAP HANA Security Dashboard:

- View **Data Encryption** status for:
  - Data and Log Volume Encryption
  - Backup Encryption
- View **Auditing** status and the number of enabled/disabled audit policies
- View **Authentication** information
  - Configure the password policy
- View **Security Related** Information
  - Manage certificated
  - View network security information

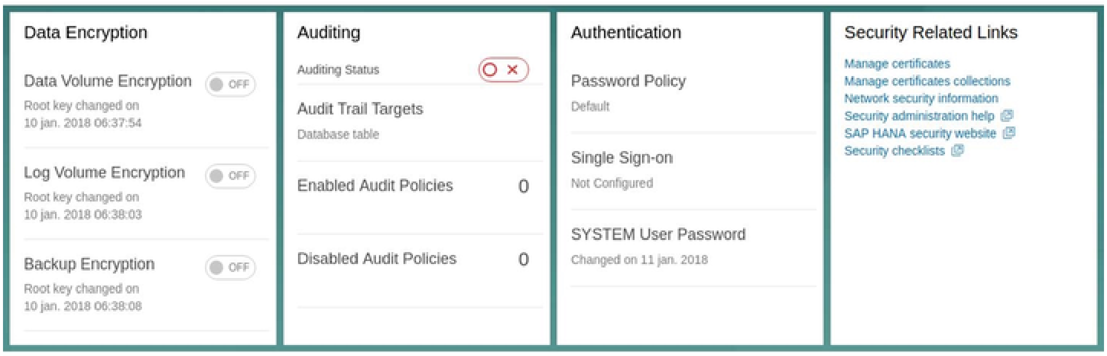

Figure 356: Security Dashboard

Related Information

For more information, see the following:

- SAP HANA Security Guide at http://help.sap.com/hana/SAP_HANA_Security_Guide_en.pdf.
- SAP HANA Master Guide at http://help.sap.com/hana/SAP_HANA_Master_Guide_en.pdf.
- SAP HANA Technical Operations Manual at http://help.sap.com/hana/SAP_HANA_Technical_Operations_Manual_en.pdf.
- SAP HANA Installation Guide at http://help.sap.com/hana/SAP_HANA_Installation_Guide_en.pdf.
- SAP HANA Administration Guide at http://help.sap.com/hana/SAP_HANA_Administration_Guide_en.pdf.
- SAP HANA Update and Configuration Guide at http://help.sap.com/hana/SAP_HANA_Update_and_Configuration_Guide_en.pdf.
- SAP Identity Management at http://help.sap.com/nwidm.

LESSON SUMMARY

You should now be able to:

- Describe the SAP HANA security functions

Unit 11
Lesson 2

Explaining Encryption

LESSON OVERVIEW

In this lesson, you will learn about SSL connection encryption and data volume encryption.

Business Example

In order to protect against security breaches or outside attacks, companies prefer to protect the data using encryption. The encryption can be done on data transfers and also on the data stored on a system.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explain encryption

Encryption Overview

- Secure Communication — Encryption of data communication in the network
  - Network traffic can be encrypted using Transport Layer Security (TLS), both between the SAP HANA database and clients, as well as between hosts in a distributed SAP HANA system.
- Encryption of the data persistence layer
  - The SAP HANA database can encrypt data at rest.
  - Encryption works at the page level and uses the AES256 encryption algorithm.

A cryptographic service provider on the server offers the following functions:

- The configuration of secure communication using Transport Layer Security
- The encryption of the persistence layer

Cryptograhic Libraries

SAP HANA supports the following cryptographic libraries:

- CommonCryptoLib (default)

  CommonCryptoLib (libsapcrypto.so) is installed by default as part of the SAP HANA server installation at `$DIR_EXECUTABLE`.
- OpenSSL

  The OpenSSL library is installed by default as part of the operating system installation.

SAP CommonCryptoLib is the successor of SAPCRYPTOLIB and it is the default cryptographic library for SAP HANA. CommonCryptoLib is installed as part of the SAP HANA server installation at the default location for library lookup: `/usr/sap/<SID>/SYS/exe/hdb/libsapcrypto.so`.

> Note:
> The OpenSSL library is also installed as part of the operating system installation. In most cases, you can use OpenSSL instead of CommonCryptoLib. However, there are some features in SAP HANA that are only supported by CommonCryptoLib; future features might also only be supported by CommonCryptoLib. For more information, see SAP Note 2093286 .

Secure Communication

The network communication channels used by SAP HANA can be categorized as follows:

- Channels used for database clients connecting to SAP HANA
- Channels used for internal database communication

Use encrypted communication channels where possible.

Types of Network Communication Channels

To support the different SAP HANA scenarios and setups, SAP HANA provides the following types of network communication channels:

- Channels used for external access to SAP HANA functionality by end-user clients, administration clients, application servers, and for data provisioning through SQL or HTTP
- Channels used for SAP HANA internal communication within the database, between hosts in multiple-host systems, and between systems in system-replication scenarios

Network Integration

The connections between SAP HANA and external components and applications come under the following categories:

External Communication Channels

- Connections used for administrative purposes
- Connections used for data provisioning
- Connections from database clients that access the SQL/MDXinterface of the SAP HANA database
- Connections from HTTP(S) clients
- Outgoing connections

Connections Between SAP HANA and External Components

The figure, Network Integration, shows an example of what these connections look like. Network connections are depicted by dotted arrows. The start of the arrow indicates which component is the initiator, while the end point of the arrow indicates the listener. The blue dotted arrows indicate administrative access to and from SAP HANA through the SAP HANA cockpit. Port numbers have a yellow background. The xx in the port numbers represents the number of your SAP HANA instance.

Although the figure only shows a single-host installation of SAP HANA, the connections in the figure apply equally to a distributed scenario.

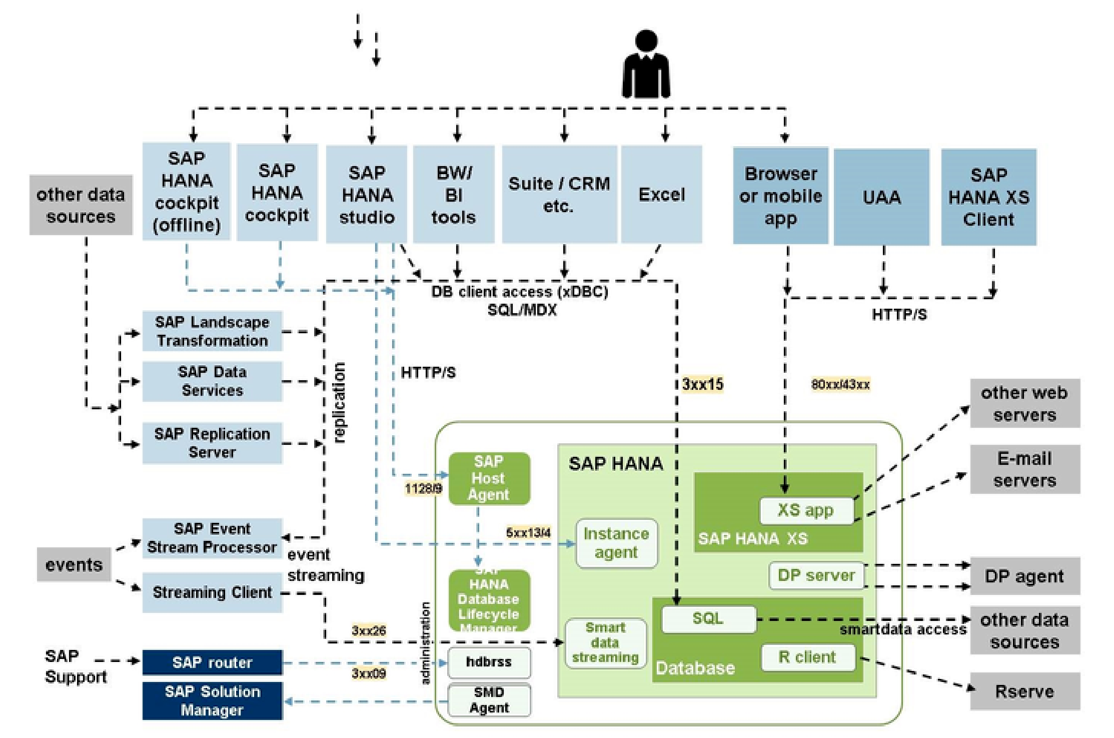

Figure 357: Connections Between SAP HANA and External Components

Furthermore, the different components of SAP HANA, as well as the hosts in a distributed scenario, also communicate with each other through the internal SAP HANA connections. These connections are also used in system replication scenarios for communication between a primary site and a secondary site. This ensures high availability in the event of a data center failure.

The following internal communication channels can be secured:

Internal Communication Channels

- Internal database communication
- Internal communication between hosts in a distributed (multiple-host) SAP HANA system
- Internal communication between systems at the different sites in a system replication (high availability) scenario
- Internal communication between the SAP HANA database and server components, such as extended storage (SAP HANA dynamic tiering).

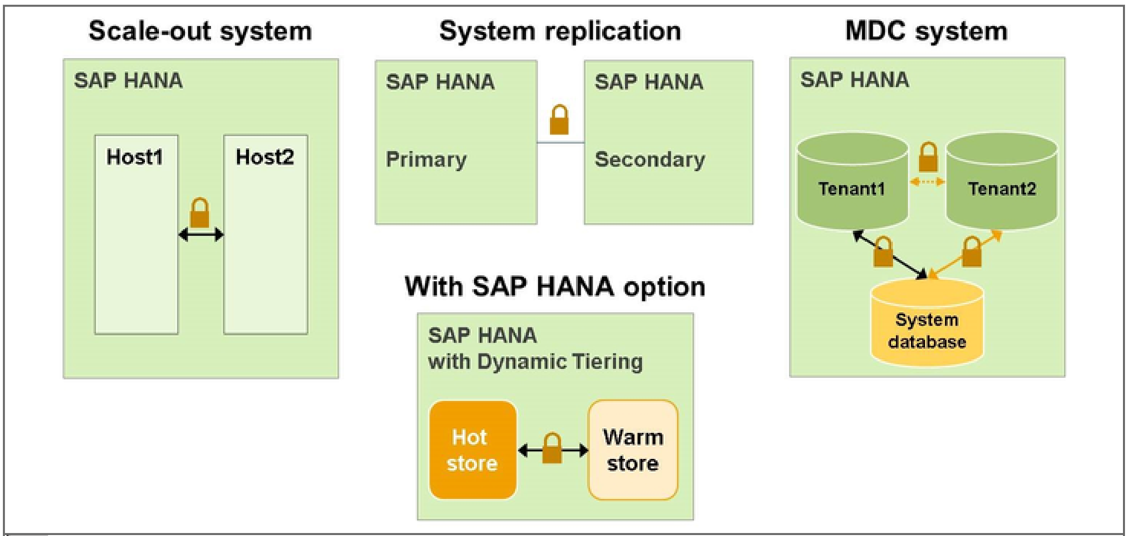

Figure 358: Internal Communication Channels

Communications Using Transport Layer Security Protocol

SAP HANA supports encrypted communication for network communication channels. Use encrypted channels wherever network attacks such as eavesdropping are not protected by other network security measures, for example, access from end-user networks. Alternatively, you can use virtual private network (VPN) tunnels to transfer encrypted information.

The network communication can be secured with the Transport Layer Security (TLS) protocol, as follows:

- Communication between the SAP HANA database and clients that access the SQL interface of the database
- Internal network communication between the individual components of an SAP HANA system on a single host and also between multiple hosts if the system is distributed
- For Client Application Access, the SAP Web Dispatcher can be configured to use HTTPS (TLS) for incoming requests from UI front ends and applications, for example, SAP HANA applications. The requests are then forwarded to SAP HANA.
- Communication between the SAP HANA Lifecycle Management tools and SAP HANA Cockpit, SAP Service Marketplace, and SAP Host Agent
- Communication between SAP HANA Cockpit and sapstartsrv
- Communication between SAP HANA Cockpit and SAP Host Agent

Separate certificate collections are supported for internal communication and external communication.

A certificate collection is also referred to as a personal security environment or PSE. It is a secure location where the public information (public-key certificates) and private information (private keys) of the SAP HANA server are stored. A certificate collection can also contain the public information (public-key certificates) of trusted communication partners or root certificates from trusted Certification Authorities.

By default, certificate collections for client-server communication over JDBC/ODBC are stored within the database. However, to maintain compatibility with previous releases, certificate collections (PSEs) can also be stored in the file system. You can create the certificate collections in the database directly.

Use of Communication in Keys and Certificates

The keys and certificates in the certificate collection for internal communication are used for the following communications:

- Communication between database services
- Communication between hosts in a multiple-host system
- Communication between hosts and sites in a system replication scenario

Certificates for external communication (for example, JDBC client access, HTTP access) are typically signed by an externally available Certification Authority (CA). This is because the CA certificates need to be integrated in the relevant clients.

The Transport Layer Secure (TLS)/Secure Sockets Layer (SSL) protocol secures communication between the SAP HANA database and the clients that access the SQL interface of the database. To use this function, configure the TLS/SSL on both the server and the client.

Secure Communication Between SAP HANA and JDBC/ODBC Clients

Server certificate validation is provided by enabling TLS/SSL for client-server communication. The server identifies itself to the client when the connection is established. This reduces the risk of man-in-the-middle attacks and prevents fake servers gaining information from clients.

If the identity of the client connecting to SAP HANA should be validated, you can also enable client certificate validation.

> Note:
> You can force all clients who communicate with the SAP HANA database through the SQL interface to use a secured connection. To do this, set the `sslEnforce` parameter in the `communication` section of the `global.ini` configuration file to **`true`**. The database then refuses SQL connection attempts that don't use SSL.

Secure Communication Between SAP HANA and JDBC/ODBC Clients

- TLS/SSL Configuration on the SAP HANA Server
- Server-Side TLS/SSL Configuration Properties for External Communication (JDBC/ODBC)
- TLS/SSL Configuration on the Client
- Client-Side TLS/SSL Configuration Properties (ODBC)
- Client-Side TLS/SSL Configuration Properties (JDBC)
- Configure SSL for SAP HANA Cockpit Connections

When you configure TLS/SSL on the SAP HANA server, the following general prerequisites apply:

- The SAP Cryptographic Library CommonCryptoLib is available on the server.

- The SAP HANA server possesses a public and private key pair, and a public-key certificate.

Connection Encryption Configuration: Server-Side Configuration

- TLS is installed (for example, Open SSL)
- For connecting with Studio using TLS, you need a certificate on the server side also
- Certificate hierarchy and certificate for the SAP HANA database server are available

  The corresponding root certificate has been deployed on all clients that are to use a TLS-encrypted connection to the server.
- Configuration can be customized for SQLDBC/JDBC-based clients using parameters in the indexserver.ini file → section Communication

Connection Encryption Configuration: TLS Configuration for SAP HANA Studio

- Prerequisite: the SAP HANA database is online, but has not yet been added to SAP HANA studio

The procedure for TLS configuration for the SAP HANA Studio is as follows:

1. In the SAP HANA studio, choose Add System... in the navigator tree.
2. Enter your user credentials and choose Connect using TLS .
3. If you want to validate the certificate, select the corresponding checkbox.
4. If you want to check the host name in the certificate, select the corresponding checkbox.
5. All connections from the SAP HANA studio to the database are now encrypted.

Note:
The procedure for configuring TLS is described in detail in the SAP HANA Security Guide.

Certificate Management

SAP HANA uses X.509 client certificates to secure internal and external communication channels, as well as for several user authentication mechanisms. Certificates can be stored and managed in files in the file system and, in some cases, directly in the SAP HANA database.

Table 14: Certificates Required for Securing Client-Server Communication using TLS

Server private key	• Used by the server to encrypt the connection • Stored securely using the internal data encryption service of the SAP HANA database
Server certificate	• Used by the client to authenticate the server

Root certificates from trusted CAs or certificates from trusted communication partners	• Optional • Used by the client for trust validation (certificate chains)

Table 15: Certificates Required for User Authentication

User certificates	• Used by the server to authenticate the connecting user • SAML assertions, SAP logon or assertion tickets, X.509 certificates (HTTP access over XS only)

All certificate-based user authentication mechanisms in SAP HANA rely on X.509 client certificates for authentication and verification of digital signatures. Secure communication between SAP HANA and clients that access the SQL interface of the database also rely on these certificates. To improve management, you can store these certificates and configure their use directly in the SAP HANA database.

Certificate Store

In systems that support multitenant database containers, in-database certificates are also used to secure communication when copying or moving a tenant database between two systems.

Although we recommend using in-database storage where possible, you can store and manage certificates in trust and key stores located in the file system. These are personal security environments or PSEs.

However, not all certificates can be stored in the database. In particular, the certificates required to secure internal communication channels with the system public key infrastructure (system PKI) and HTTP client access using SAP Web Dispatcher cannot be stored there. These certificates are contained in PSE files in the file system. Do not delete these files from the file system.

The following figure shows how in-database certificates stored in certificate collections can be used and how certificates stored in PSEs on the file system are used.

In-database certificates and certificate collections can be fully managed certificates in the SAP HANA cockpit.

Certificates can be stored in the SAP HANA database or in the file system

Certificates can be stored for…	…in the database	…in the file system
TLS (client-server communication over JDBC/ODBC)	YES	YES
TLS (client-server communication over HTTP)	NO	YES
TLS (internal communication)	NO	YES
TLS (tenant database replication)	YES	NO
TLS (LDAP server communication)	YES	NO
Authentication (SAML, SAP Logon and Assertion Tickets, X.509)	YES	YES

Recommendation: Store certificates in the database where possible.

Figure 359: Certificates Store

Certificate Management using SAP HANA Cockpit

The management of certificates in the SAP HANA database follows a typical workflow. User authorization allows for a full separation of duties. The full workflow is supported by the SAP HANA cockpit.

In-Database Certificate Management Workflow

To manage in-database certificates, proceed as follows:

1. Import the public-key certificates of trusted communication partners, and the root certificates of trusted Certification Authorities, into the certificate store.
2. Complete the following steps:
    a. Create the required certificate collections.
    b. Add trusted certificates from the certificate store to certificate collections.
    c. Add the SAP HANA server certificates to those collections that will be used for server authentication (for example, secure client-server communication over JDBC/ODBC).
3. Set the purpose of individual collections. The privilege that is required depends on the purpose that is set.

Certificate Management with SAP HANA Cockpit

Certificates in the database can be managed using SAP HANA Cockpit or SQL. There are certificates that are managed in the file system, for example, TLS/SSL for HTTP, and TLS/SSL for internal communication (automatic setup via SystemPKI). These cannot be managed with the SAP HANA Cockpit functionality. Simplified configuration for these scenarios is achieved by other means (SystemPKI).

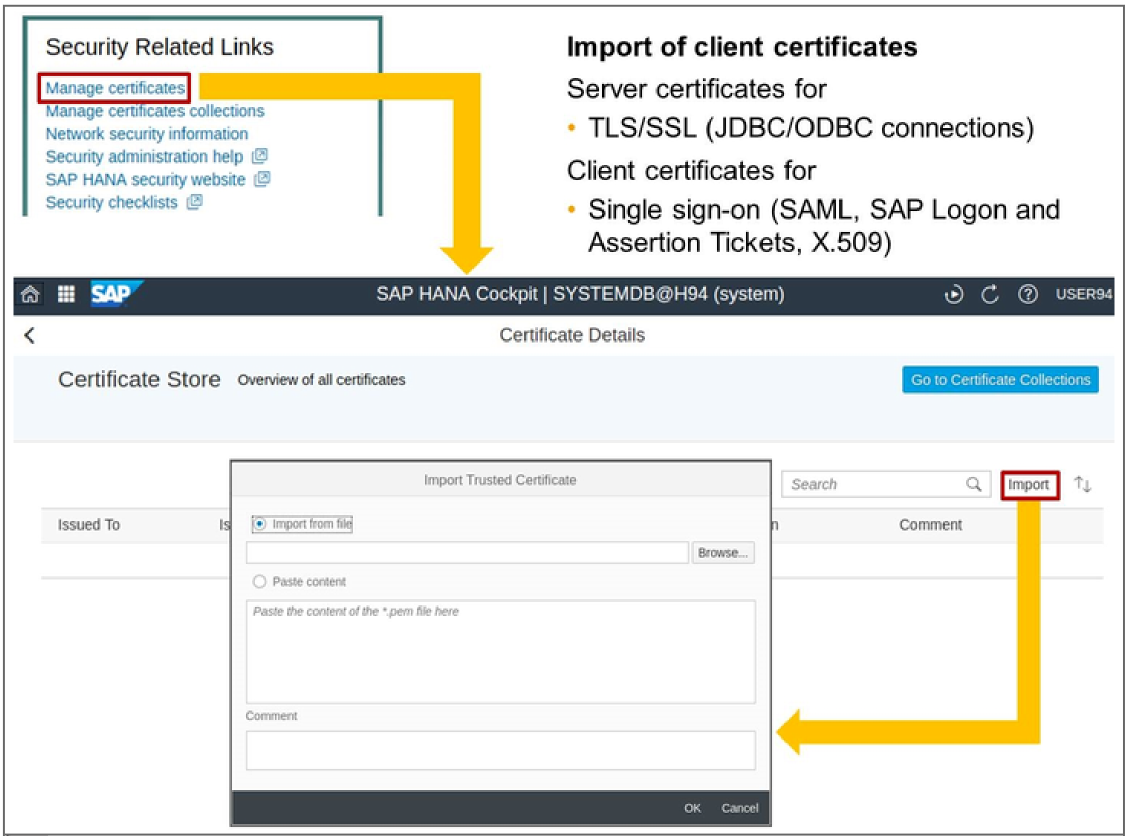

Figure 360: Certificate Management with SAP HANA Cockpit

> Note:
> For information on migrating file-based certificates into the database, see SAP Note 2175664 .

Automatic Generation of PKI/Certificates for Internal Communication Channels

You can secure the following internal communication channels:

- Between databases in a multiple-container system

  For an MDC system, only encryption is available; tenant authentication is unavailable.

- Between hosts in a scale-out system

  Also between processes in a single-host system.

- Between SAP HANA systems in system replication scenarios

  Metadata + data channel.

- Between the SAP HANA database and additional server components

  For example, an extended storage server (SAP HANA dynamic tiering option) and smart data streaming server (SAP HANA Smart Data Streaming option).

A public-key infrastructure (PKI system) for securing internal communication channels using TLS is set up automatically during installation. No user interaction is required for the setup.

> Note:
> The PKI system cannot be used to secure the communication between the SAP HANA server and clients.

PKI System Keys and Certificate Features

The keys and certificates used by the PKI system include the following features:

- Each component (host, database, additional server, and so on) receives a public/private key pair and a public-key certificate for mutual authentication.
- The certificates are signed by a dedicated trusted certificate authority (CA), which is unique for each SAP HANA system.
- The certificates are renewed automatically.
- CommonCryptoLib is used as the cryptographic library.

Depending on the communication channel, you might need to enable TLS explicitly.

> Note:
> For information on the migration from manual system SSL configuration to the PKI System, see SAP Note [2175672](2175672) .

Monitoring of Network Security Settings

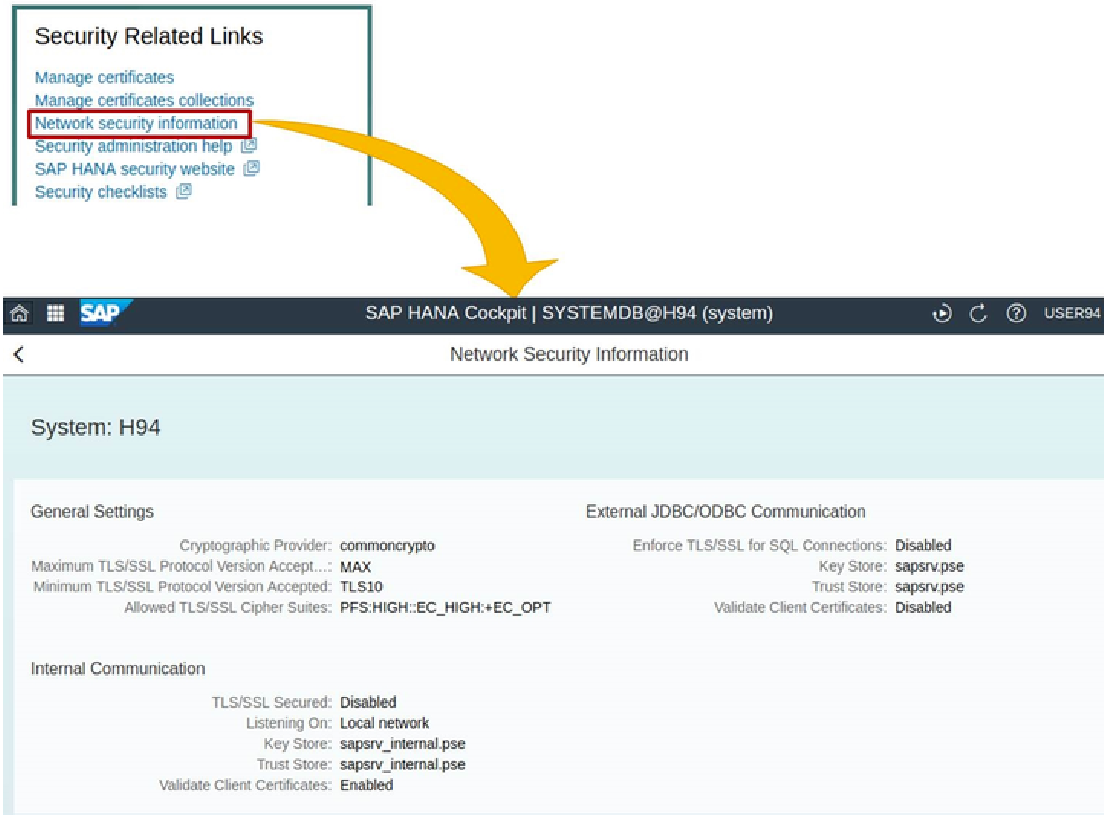

Figure 361: Monitor Network Security Settings

Network Security Settings

The Network Security Information screen shows the following information:

- Internal and external network security configuration
- Certificate and private key stores

Data Volume and Redo Log Encryption

To protect data saved to disk from unauthorized access at the operating system level, the SAP HANA database supports data encryption in the persistence layer. Data volume encryption protects the data area on-disk, while redo log encryption protects the log area on-disk. SAP HANA data and log backups can be encrypted too.

The SAP HANA database holds most of its data in-memory for maximum performance. However, it still uses persistent disk storage as a fallback in case of failure. During normal operation, data is automatically saved from memory to disk at regular savepoints. Additionally, all data changes are captured in redo log entries. A redo log entry is written to disk with each committed database transaction. After a power failure, SAP HANA can be restarted like any disk-based database. It returns to its last consistent state by replaying the redo log entries since the last savepoint.

Data Rest Encryption

Data volume, redo logs and backups (data and log) can be encrypted as follows:

- Data Volume Encryption

  If data volumes are encrypted, all pages that reside in the data area on-disk are encrypted using the AES-256-CBC algorithm. Pages are transparently decrypted as part of the load process into memory. When pages reside in-memory, they are therefore not encrypted and there is no performance overhead for in-memory page accesses. When changes to data are persisted to disk, the relevant pages are automatically encrypted as part of the write operation.

  Pages are encrypted and decrypted using 256-bit page encryption keys. Page keys are valid for a certain range of savepoints and can be changed by executing SQL statements. After data volume encryption is enabled, an initial page key is automatically generated. Page keys cannot be read in plain text, but are encrypted themselves with a dedicated data volume encryption root key. This key is generated randomly during installation.

- Redo Log Encryption

  If redo logs are encrypted, log entries are encrypted using the AES-256-CBC algorithm before they are written to disk. Log entries are encrypted and decrytped using a 256-bit long root key, which is generated randomly during installation.

- Backup Encryption

  Switch On Backup Encryption and all subsequent data backups, delta backups, and log backups will be encrypted. Note that data snapshots are not encrypted unless data volume encryption is enabled.

During start-up, administrator interaction is not required. The data volume encryption, redo log and backup encryption root keys are stored using the secure storage in the file system (SSFS) functionality, and are automatically retrieved from there.

Data volume encryption

- Page level encryption of data volumes on disk.

Redo log encryption

- Log entries are encrypted before they are written to log volumes on disk.

Data and Log backup encryption

- Full data backups, delta data backups, and log backups.

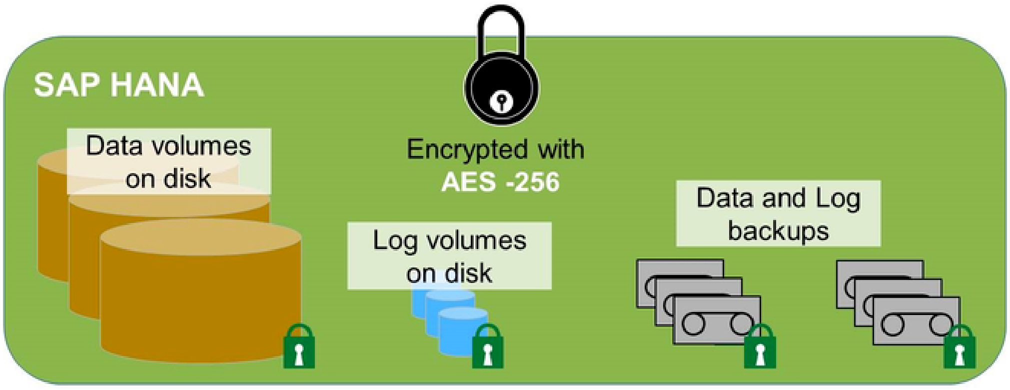

Figure 362: Data Volume and Redo Log Encryption

SAP HANA allows you to encrypt data volumes and redo logs independently of each other. However, if you require full protection in the persistence layer, you should enable both.

Enabling encryption does not increase data size.

> Note:
> SAP HANA uses the SSFS instance to protect the encryption root keys from unauthorized access. These root keys protect all encryption keys used in the SAP HANA system, and are encrypted using the SSFS master key.

Unencrypted Data

The persistence encryption feature does not encrypt the database traces.

For security reasons, do not run the system with extended tracing for more than short-term analysis. This is because tracing might expose security-relevant data that is encrypted in the persistence layer, but not in the trace. Therefore, do not keep such trace files on disk beyond the respective analysis task.

Configuration and Monitoring of Data Volume Encryption

Data at rest encryption on-disk can be configured using SAP HANA cockpit, or SQL commands.

After activating encryption, new data that is saved to disk is encrypted, starting with the next savepoint, log write or backup. Because of the shadow memory nature of SAP HANA persistence, outdated versions of pages can remain unencrypted on-disk. All future redo log entries persisted to log volumes are encrypted.

Caution:
For complete protection, enable data volume encryption after you reinstall the system. All your data is encrypted only after you have completed this process completed. This also ensures that a new root encryption key is generated.

Enable Data Encryption using SQL:

- Data Volume encryption: ALTER SYSTEM PERSISTENCE ENCRYPTION ON
- Redo Log encryption: ALTER SYSTEM LOG ENCRYPTION ON
- Data and Log backup encryption: ALTER SYSTEM PERSISTENCE BACKUP ON

Disable Data Encryption using SQL:

- Data Volume encryption: ALTER SYSTEM PERSISTENCE ENCRYPTION OFF
- Redo Log encryption: ALTER SYSTEM LOG ENCRYPTION OFF
- Data and Log backup encryption: ALTER SYSTEM BACKUP ENCRYPTION OFF

Figure 363: Enable Data Volume Encryption

You can change the page encryption key for the data volume encryption with SAP HANA cockpit or SQL commands.

After you change the page encryption key, you can choose if you also want to re-encrypt the existing encrypted data with the new key. This happens in the background.

You can monitor the encryption progress in SAP HANA cockpit or with SQL commands.

Note:
Although you can encrypt data volumes and redo logs independently of each other with SAP HANA, if you require full protection in the persistence layer, enable both.

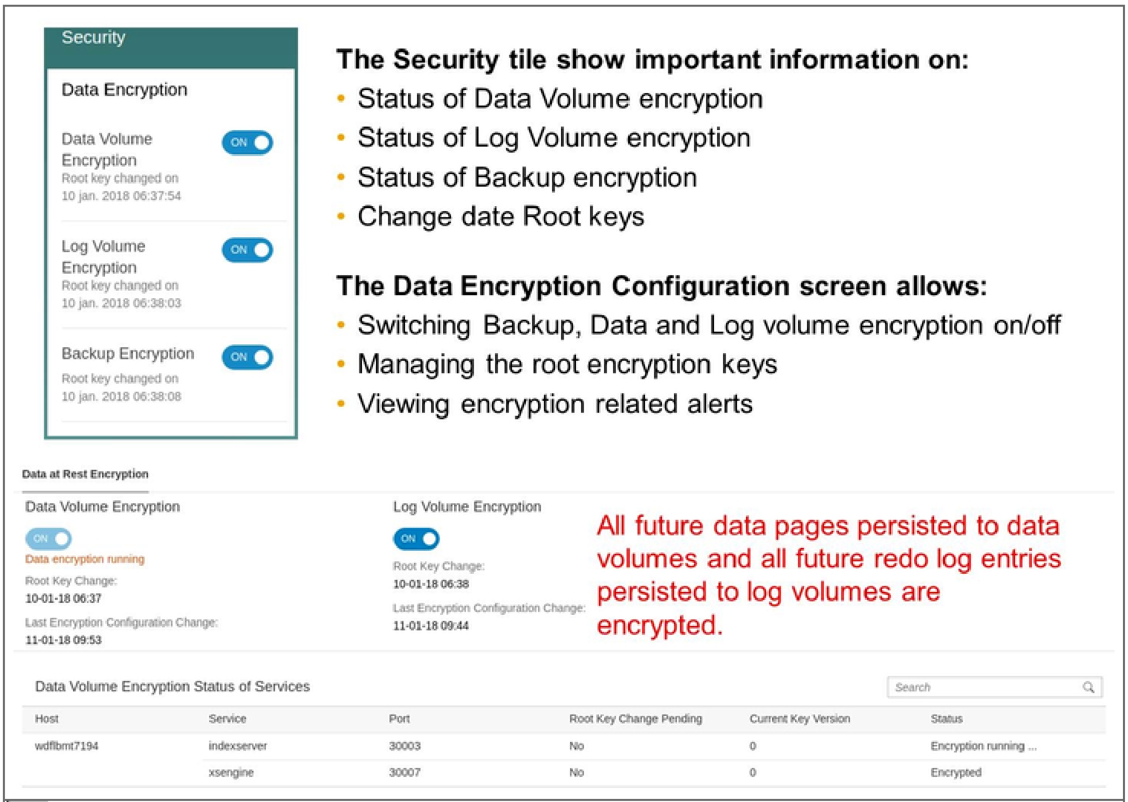

Figure 364: Monitor Data Volume Encryption

Note:
You can view the status of data volume encryption and redo log encryption in the system view M_ENCRYPTION_OVERVIEW.

SSFS Encryption Keys

SAP HANA includes encryption services for encrypting data at rest. It also has an internal encryption service for applications with data encryption requirements. SAP HANA uses the secure store in the file system (SFFS) functionality to protect all encryption root keys.

Secure Stores in the File System (SSFS)

SAP HANA uses two secure stores in the file system: the instance SSFS, and the system PKI SSFS. The instance SSFS protects the root keys used for all data-at-rest encryption services and the internal application encryption service. The system PKI SSFS protects system-internal root certificates that are required for secure internal communication.

SAP HANA uses the instance SSFS to protect the following encryption root keys:

- The root key used for data volume encryption
- The root key used for redo log encryption
- The root key used for the internal application encryption service of the database
- The password of the root key backup
- Encryption configuration information

These root keys protect all encryption keys and data used in the SAP HANA database from unauthorized access.

Note:
The application encryption root key is used by the secure internal credential store. This is needed in some scenarios, such as smart data access, to store additional user credentials securely (for example, for access to remote systems).

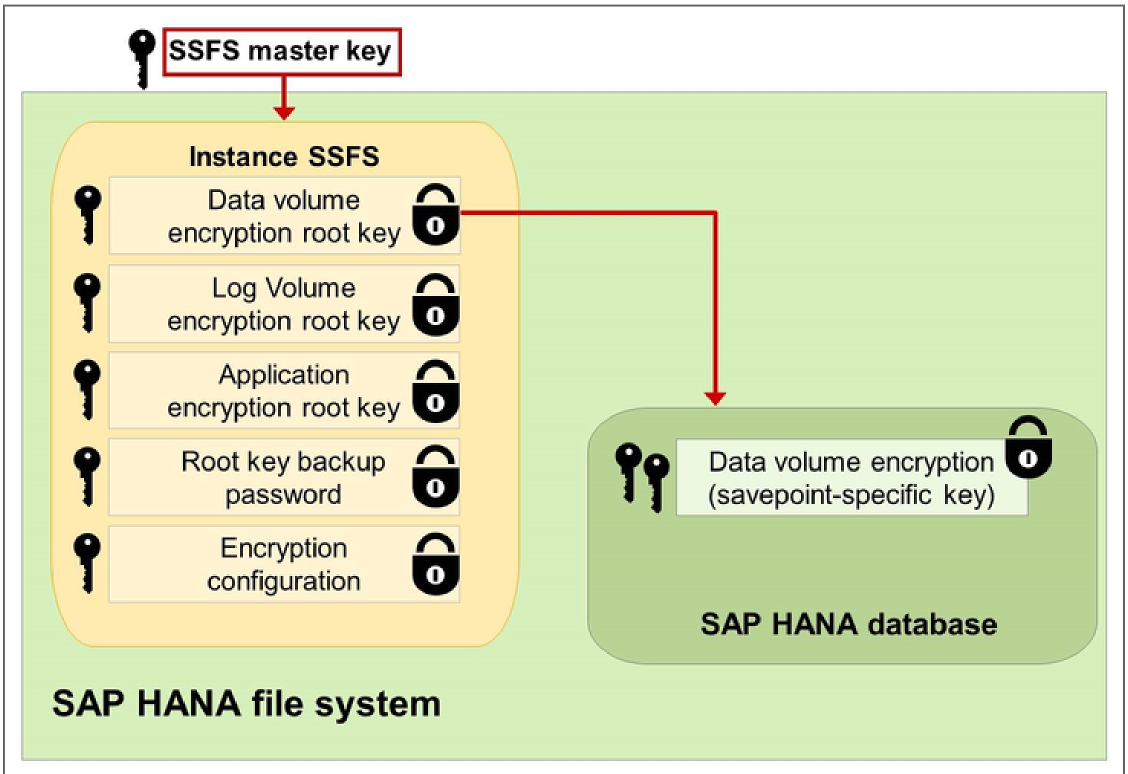

Figure 365: Encryption Keys Used in SAP HANA

To prevent data encrypted in the SAP HANA database from becoming inaccessible, the content of the instance SSFS and key information in the database must remain consistent. If this is not case, for example if the instance SSFS becomes corrupted, the database issues an alert (check 57). Contact SAP Support to resolve these issues.

The page encryption keys used for data volume encryption are encrypted themselves by the data volume encryption root key. The root key is generated randomly during installation. The page keys are created when data volume encryption is enabled.

This secure store, which is used by SAP HANA to store internal root keys, is protected by the SSFS master key. To support automatic unattended start-up of the SAP HANA system, the key store and the SSFS master key are stored on the file system. They are protected by operating system permissions, which require operating system access with the <sid>adm operating system user.

Encryption Key Management

SAP HANA generates unique keys during installation. However, if you received SAP HANA pre-installed from a hardware vendor, you can change them to ensure that they are not known outside of your organization. Perform this immediately after hand-over from your hardware partner.

The following master and root keys can be changed:

- Instance SSFS master key
- System PKI SSFS master key
- Data volume encryption root key
- Redo log encryption root key
- The password of the root key backup
- Application encryption service root key

You can change all master and root keys by re-installing your system. You can also change keys manually and individually.

The following figure shows the recommended process for configuring encryption in your SAP HANA system for the first time.

Overview: Managing Data at Rest Encryption

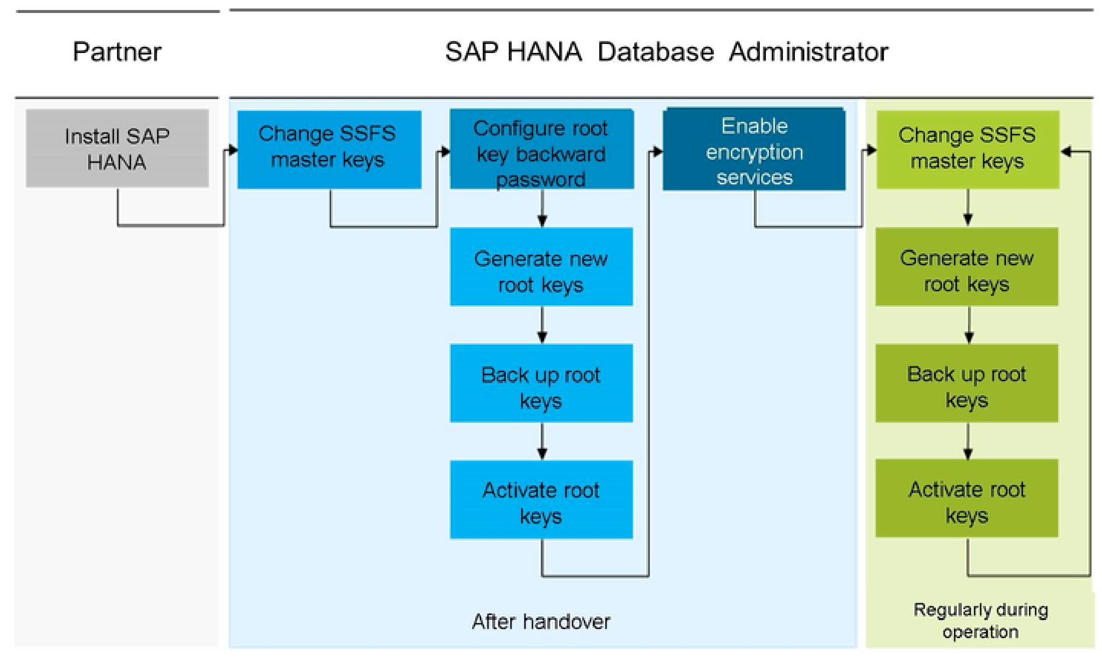

Figure 366: Overview: Managing Data at Rest Encryption

To manage data encryption, proceed as follows:

1. Change the master keys of the instance SSFS and the system PKI SSFS.

   Unique master keys are generated during installation or update. You can also change the master keys any time.

   An administrator can change the SSFS master key using the command line tool rsecssfx and the credentials of the operating system user <sid> adm. Therefore, the SAP HANA system has to be stopped. For special scenarios like snapshot-based backup and restore, or system replication, see SAP Note 2194396 .

2. Configure the password for the root key backup.

   This password is required to back up the root keys and to restore the backed-up root keys during data recovery.

3. Change the encryption root keys for all encryption services, including the following:

   - Data volume encryption
   - Redo log encryption
   - Internal application encryption

4. Enable the following required services:

   - Data volume encryption
   - Redo log encryption

5. Periodically change the SSFS master keys and encryption root keys according to your security policy.

Change of Secure Storage in the File System (SSFS) MasterKeys

Changing the Secure Storage in the File System (SSFS) Master Keys

The initial default master keys that protect the two secure stores in the file system (SSFS) used by SAP HANA are changed during installation or upgrade. If you received your system pre-installed from a hardware or a hosting partner, change them immediately after hand-over to ensure that they are not known outside of your organization.

Prerequisites

The following are prerequisites for changing the SSFS Master Keys:

- You have the credentials of the operating system user <sid>adm that was created when the system was installed.
- You have the system privilege INIFILE ADMIN.

SAP HANA uses SSFS to protect the root encryption keys. These root keys protect all encryption keys used in the SAP HANA database from unauthorized access. The root encryption keys are as follows:

- The root key used for the internal data encryption service of the database
- The root key used for data volume encryption

In a system that supports multitenant database containers, the system database and all tenant databases have their own root encryption keys for both the data encryption service and data volume encryption.

You can change the SSFS master keys using the command line tool `RSECSSFX`, which is installed with SAP HANA. It is available at `/usr/sap/<SID>/HDB<instance>/exe`.

Before changing the SSFS master keys, note the following:

- Changing SSFS master keys requires system downtime.
- In a distributed SAP HANA system, every host must be able to access the file location of the instance SSFS master key.

- In a system that supports multitenant database containers, you only need to change the SSFS master keys once for whole instance and not per tenant database.

Change the SSFS Master Keys

To change the SSFA master keys, proceed as follows:

1. Log on to the SAP HANA system host as the operating system user <sid>adm.
2. Shut the system down using the sapcontrol program:

   `/usr/sap/hostctrl/exe/sapcontrol -nr <instance_no> -function Stop`

3. Change the master key of the instance SSFS as follows:

   a. Re-encrypt the instance SSFS with a new key with the command:

```
RSEC_SSFS_DATAPATH=/usr/sap/<SID>/SYS/global/hdb/security/ssfs
RSEC_SSFS_KEYPATH=<path to key file> rsecssfx changekey $(rsecssfx
generatekey -getPlainValueToConsole)
```

   b. Configure the specified key file location in the global.ini configuration file at `/usr/sap/<SID>/SYS/global/hdb/custom/config/global.ini`.

      If the file does not exist, create it. Add the following lines:

      `[cryptography] ssfs_key_file_path = < path to key file>`

4. Re-encrypt the system PKI SSFS with a new key using the following command:

```
RSEC_SSFS_DATAPATH=/usr/sap/<SID>/SYS/global/security/rsecssfs/data
RSEC_SSFS_KEYPATH=/usr/sap/<SID>/SYS/global/security/rsecssfs/key
rsecssfx changekey $(rsecssfx generatekey -getPlainValueToConsole
```

5. Restart the system using the following command:

   `/usr/sap/hostctrl/exe/sapcontrol -nr <instance_no> -function Start`

SSFS and System Replication

In a system-replication setup, configure the location of the instance SSFS master key file on the secondary system or systems. The file itself is copied automatically.

For file system-based copies of SAP HANA database installations, save and restore the instance SSFS master key file manually; otherwise data loss can occur. In regular backup and recovery scenarios (including snapshots), you do not have to take any actions for the master key. This is because only the content of the SSFS, not the master key, is contained in the backup.

> Note:
> It is not necessary to save the system PKI SSFS key file. The system generates a new system PKI SSFS automatically, if required.

LESSON SUMMARY

You should now be able to:

- Explain encryption

Unit 11
Lesson 3

Describing Auditing

LESSON OVERVIEW

This lesson covers the audit logging infrastructure.

Business Example

Many regulatory requirements require audit logging.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Describe auditing

Auditing Overview

Auditing shows you who did what in the SAP HANA database (or tried to do what) and when. This allows you, for example, to log and monitor read access to sensitive data.

Audit Logging: Introduction

- Audit logging records critical system events
  - User management: for example, user changes, role granting
  - System access and configuration: for example, failed logons, parameter changes
  - Data access: for example. read and write access to tables and views, execution of procedures
  - "Log all": firefighter logging, for example, for support cases
- Audit policies
  - Include events to be recorded
  - If audit logging is enabled, some critical events are always logged
- Audit trails
  - Audit entries are created in one or more audit trails when an audit policy is triggered

The auditing feature of the SAP HANA database allows you to monitor and record selected actions that are performed in your database. To use this feature, first activate it for the database. You can then create and activate the required audit policies.

An audit policy defines the actions that are to be audited. It also outlines the conditions under which the action must occur for it to be relevant for auditing. When an action occurs, the policy is triggered and an audit event is written to the audit trail. Audit policies are specific to each database.

If the audit trail target is a database table, the audit table might increase continuously. To prevent this, delete the audit entries that were created up until a certain time and date.

You can use the SAP HANA cockpit or SQL to enable auditing, configure audit trail targets, and create audit policies.

Audit logging is not enabled by default.

Events that Can be Audited

An action corresponds to the execution of an action in the database by SQL statement. For example, to track user provisioning in your system, create an audit policy that audits the execution of the SQL statements `CREATE USER` and `DROP USER`.

Although most actions correspond to the execution of a single SQL statement, some actions cover the execution of multiple SQL statements. For example, the action GRANT ANY audits the granting of multiple entities for the SQL statements `GRANT PRIVILEGE`, `GRANT ROLE`, `GRANT STRUCTURED PRIVILEGE`, and `GRANT APPLICATION PRIVILEGE`.

Actions That Can Be Audited

- The following actions are typically audited:
  - Changes to user authorization
  - Creation and deletion of database objects
  - Authentication of users
  - Changes to system configuration
  - Access to or changing of system data
- **Both successful and unsuccessful actions can be recorded**

You can audit the following actions:

- Changes to user authorization
  - Create or drop user, and create or drop role
  - Grant or revoke role
  - Grant or revoke SQL privilege, system privilege, and analytical privilege
  - Create or drop analytical privilege
  - Create or drop and alter structured privilege
- Authentication of users

  Connection attempts of users to the database
- Changes to system configuration
  - Changes to system configuration, for example, ini file
  - Install license key
  - Set system license or unset system license for all
  - Changes to the data volume encryption

When you enable audit logging for configuration changes, the previous values of parameters are written to the audit trail.

- Access to or changing of sensitive data

  You can specify the following database objects to be audited:

  - Tables
  - Views
  - Procedures
  - Schema

  Both write and read access to data can be recorded as follows:

  - SELECT
  - INSERT
  - UPDATE
  - DELETE
  - EXECUTE

> Note:
> Currently, only the statements that were executed are logged.

> Hint:
> You can only audit actions that occur inside the database engine. If the database engine is not online when an action occurs, it cannot be detected and, therefore, cannot be audited. These actions include, for example, an upgrade of an SAP HANA database instance, or direct changes to system configuration files using operating system commands.

Mandatory Audit Actions

If auditing is active, certain actions are always audited and are therefore not available for inclusion in user-defined audit policies. In the audit trail, these actions are labeled with the internal audit policy Mandatory Audit Policy.

Mandatory audit actions include the following:

- Creation, modification, or deletion of audit policies
- Deletion of audit entries from the audit trail

  This only applies if audit entries are written to column store database tables.
- Changes to auditing configuration:

  - Enabling or disabling auditing
  - Changing the audit trail target

- Changing the location of the audit trail target if it is a CSV text file

Unauditable Events

Only actions that take place inside the database engine can be audited. If the database engine is not online when an action occurs, it cannot be detected and therefore cannot be audited.

This is important to consider in the following cases:

- Upgrade of an SAP HANA database instance

  Upgrade is triggered when the instance is offline. When it becomes available online again, you cannot determine which user triggered the upgrade and when.

- Direct changes to system configuration files using operating system commands

  Only changes that are made using SQL are visible to the database engine. You can change configuration files when the system is offline.

Audit Policies

An audit policy defines the actions to be audited, as well as the conditions under which the action must be performed to be relevant for auditing. When an action occurs, the policy is triggered and an audit event is written to the audit trail. Audit policies are database-specific.

An audit policy can specify any number of actions to be audited, but not all actions can be combined together in the same policy.

In addition to the actions to be audited, an audit policy specifies parameters that further narrow the number of events actually audited.

Audit Policies

The audit policies are as follows:

- Specifies any number of actions to be audited.
- Specifies parameters that further narrow the number of events actually audited.
  - Audited action status (successful or unsuccessful)
  - Target objects (schemata, tables, ...)
  - Audited users
  - Audit level

When the audit policy is triggered, an audit entry of the corresponding level is written to the audit trail. This allows tools that check audited actions to find the most important information, for example.

Audit Logging: Infrastructure

Audit logging infrastructure:

- Auditing can be configured in SAP HANA cockpit, or using SQL statements
- Audit trails are written to the Linux syslog or an internal database table
- Auditing can be enabled / disabled per tenant database

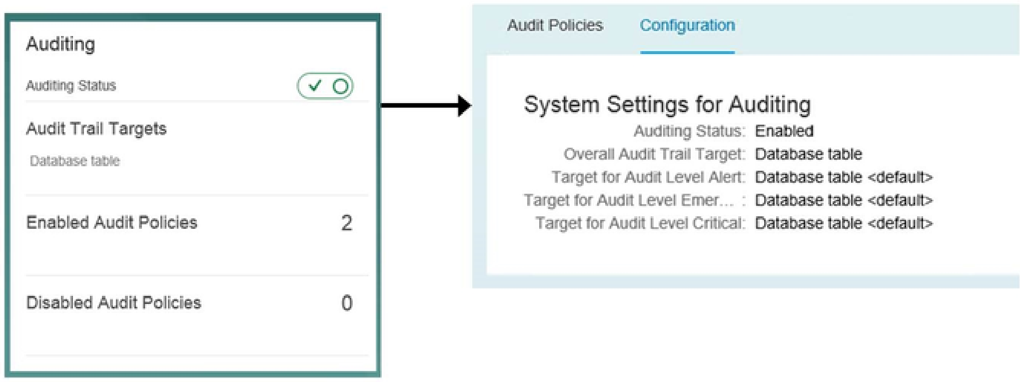

- Audit policies define which actions in the system are logged:
  - They can be explicitly enabled or disabled
  - Audited users can be specified
  - They are stored in the database catalog

Figure 367: Audit Logging: Infrastructure

When an action in the policy occurs under the conditions defined in the policy, an audit policy is triggered. When this occurs, an audit entry is created in one or more audit trails.

The following audit trail targets are supported for production systems:

- Linux syslog

  The logging system of the Linux operating system (syslog) is a secure storage location for the audit trail because not even the database administrator can access or change it. The syslog has numerous storage possibilities, including storing it on other systems. In addition, the syslog is the default log daemon in UNIX systems. The syslog therefore provides a high degree of flexibility and security, as well as integration into a larger system landscape. For more information about how to configure syslog, see the documentation of your operating system.

- Database table

  To query and analyze auditing information quickly, you can use an SAP HANA database table as the target for the audit trail. It also provides a secure and tamper-proof storage location.

  The internal column store table is in the _SYS_AUDIT schema of the SAP HANA database.

  You can only access audit entries through the public system view AUDIT_LOG. Only users with system privilege AUDIT ADMIN or AUDIT OPERATOR can perform SELECT operations on this view.

  To avoid the audit table growing too large, you can delete old audit entries.

Note:
For test purposes in non-production systems, you can also use a CSV text file as the audit trail. A separate CSV file is created for every service that executes SQL.

Hint:
You can configure multiple audit trail targets for different audit levels and per audit policy, as follows:

- **System-wide default**

  If no other trail target has been configured per audit level, audit entries are written to the audit trail targets configured for the system.

- **Audit level (optional)**

  Audit entries from audit policies with the audit level EMERGENCY, CRITICAL, or ALERT are written to the specified audit trail targets. If no audit trail target is configured, entries are written to the audit trail target configured for the system.

- **Audit policy (optional)**

  Audit entries from a particular policy are written to the specified audit trail target or targets. If no audit trail target is configured for an audit policy, entries are written to the audit trail target for the audit level if configured, or the audit trail target configured for the system. You can configure several audit trail targets for each individual policy.

Activation of Audit Policies

Auditing is implemented through the creation and activation of audit polices. An audit policy defines the actions to be audited, as well as the conditions under which the action must be performed to be relevant for auditing. For example, actions in a particular policy are audited only when they are performed by a particular user on a particular object. When an action occurs, the audit policy is triggered and an audit event is written to the audit trail.

Enable Auditing in SAP HANA Cockpit

The figure, Configuring Audit Logging, outlines how to configure and switch on audit logging for the SystemDB using SAP HANA cockpit.

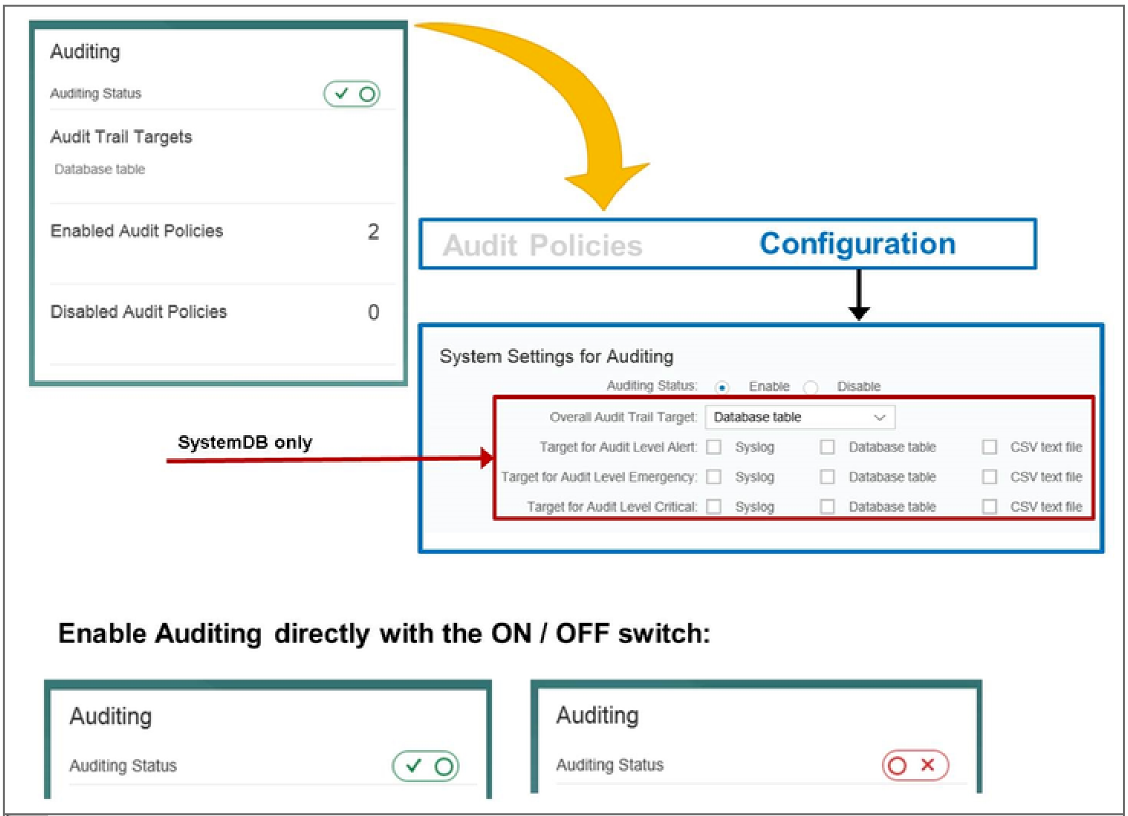

Figure 368: Configuring Audit Logging

> Note:
> Auditing can be activated or disabled per tenant in a multitenant SAP HANA database. In the SystemDB, you can also configure the audit trail targets. In a tenant database this cannot be changed.

Managing Audit Policies in SAP HANA Cockpit

You can manage audit policies as follows:

- Define the audit policy, and specify the following settings:
  - If only successful actions, only unsuccessful actions, or both, should be recorded
  - For the SystemDB, the audit trail target
  - Objects to which actions apply, for example, a schema or a table
  - Users to be audited or excluded from auditing
- Activate an audit policy, set the Status to Enable and save
- Switch auditing on or off

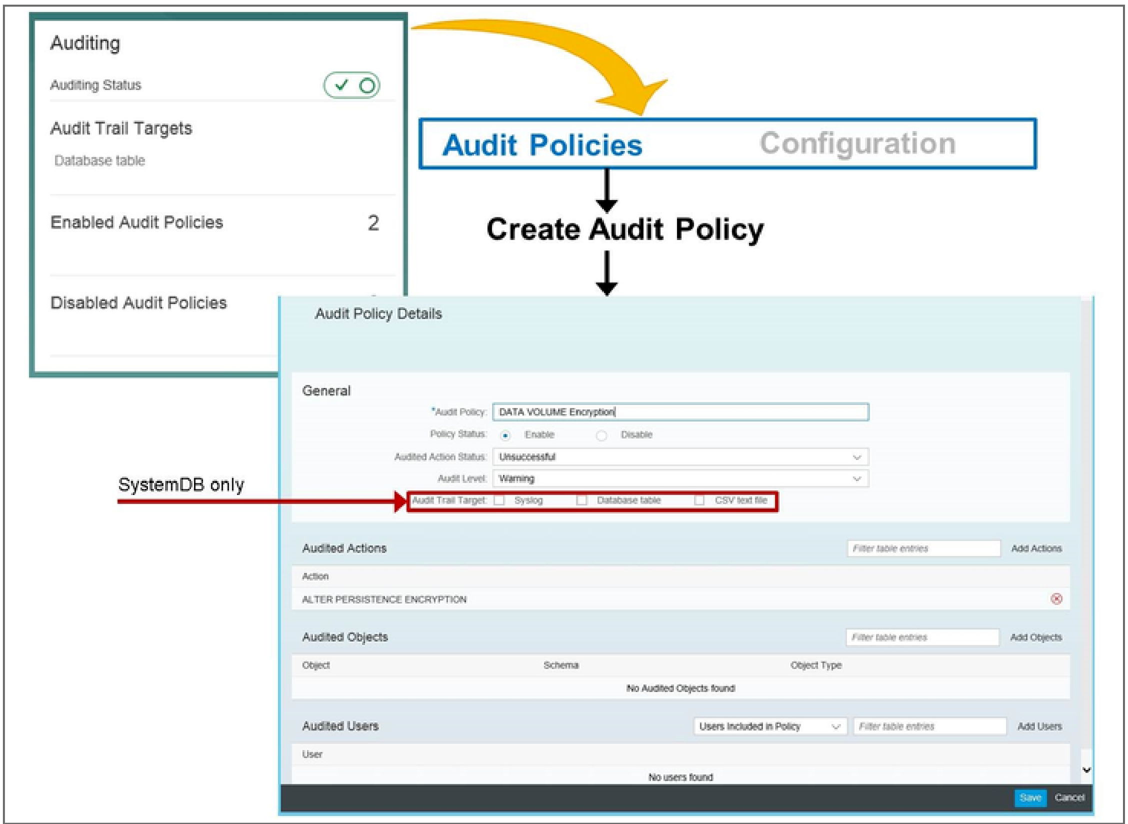

Figure 369: Create Audit Policies in SAP HANA Cockpit

You can specify any number of actions to audit in an audit policy. Not all actions can be combined together in the same policy; therefore, compatible audit actions are grouped together. When you select an action, any actions that are incompatible with the selected action are unavailable for selection.

If you want to select two incompatible audit actions, you need to create two separate audit policies.

In addition to the actions to be audited, an audit policy also specifies additional parameters that further narrow the number of events actually audited, as follows:

- Audited action status
  - On successful execution
  - On unsuccessful execution
  - On both successful and unsuccessful execution
- Target object or objects
  - Tables
  - Views
  - Procedures
- Audited user or users

  Individual users can be included or excluded from an audit policy.
- Audit level

- EMERGENCY
- ALERT
- CRITICAL
- WARNING
- INFO

When an audit policy is triggered (that is, when an action in the policy occurs under the conditions defined in the policy) an audit entry is created in the audit trail.

Firefighter logging logs all actions performed by a specific user. This covers all actions that can be audited individually, as well as actions that cannot otherwise be audited. Such a policy is useful if you want to audit the actions of a particularly privileged user.

> Note:
> Some actions cannot be audited using database auditing, even with a policy that includes all actions, in particular, system restart and system recovery.

> Caution:
> Firefighter logging can generate many audit entries, so only enable it if required.

Monitor Audit Logging Status and Check Policies

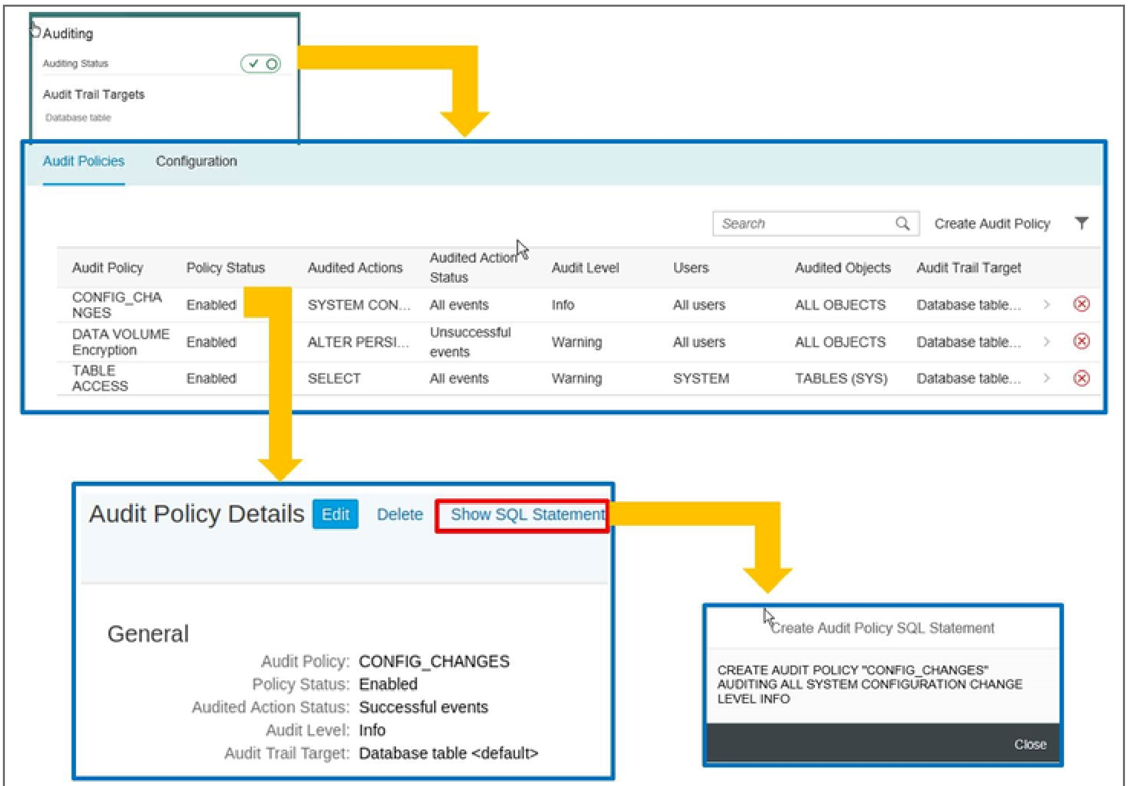

Figure 370: Monitor Audit Logging Status and Check Policies

The Auditing tile in SAP HANA Cockpit allows you to view the audit logging status, and check which audit policies are active.

> Note:
> Using SAP HANA cockpit, the SQL statement that corresponds to an existing audit policy can be shown in the audit policy details.

Viewing the Audit Trail

You can query and analyze auditing information quickly using an SAP HANA database table as audit trail target. It provides a secure and tamper-proof storage location. You can only access audit entries through the public system view AUDIT_LOG. This view is read-only; only a user with system privilege AUDIT OPERATOR can delete old entries from the underlying internal table.

Viewing the audit trail from the database table:

- In the SQL Console of the SAP HANA Cockpit open the view *PUBLIC.AUDIT_LOG*
- Alternatively, display the system view using SQL command:
  `SELECT * FROM "PUBLIC"."AUDIT_LOG"`

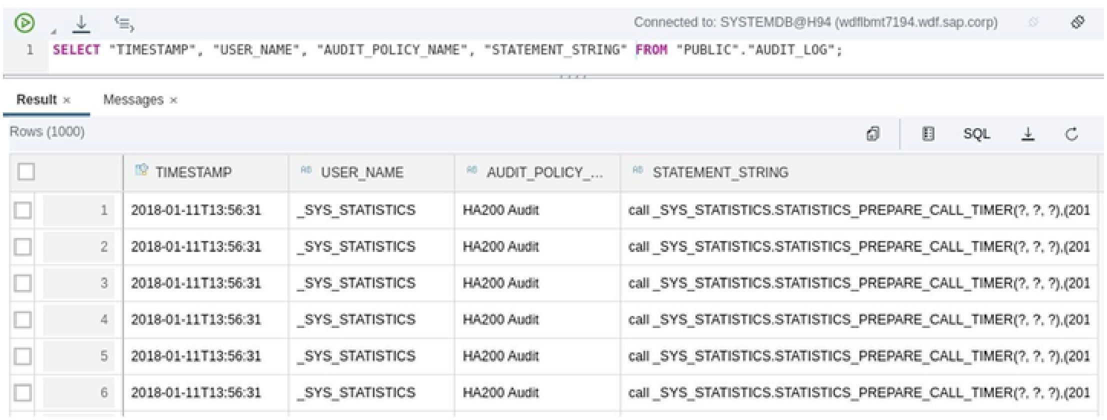

Figure 371: Viewing the Audit Trail

Monitor the Size of the Audit Trail Table

To avoid the audit table growing indefinitely, you can delete old audit entries by truncating the table. The system monitors the size of the table in relation to the overall memory allocation limit of the system. It issues an alert when it reaches defined values, which are 5%, 7%, 9%, and 11% of the allocation limit, by default. Configure this behavior can be configured with check 64: “Total memory usage of table-based audit log” . Only users with the system privilege AUDIT OPERATOR can truncate the audit table.

> Note:
> This alert only applies if you select a database table as an audit trail target (not for syslog).

You can use the SAP HANA cockpit to delete audit entries created up until a certain time from the audit table.

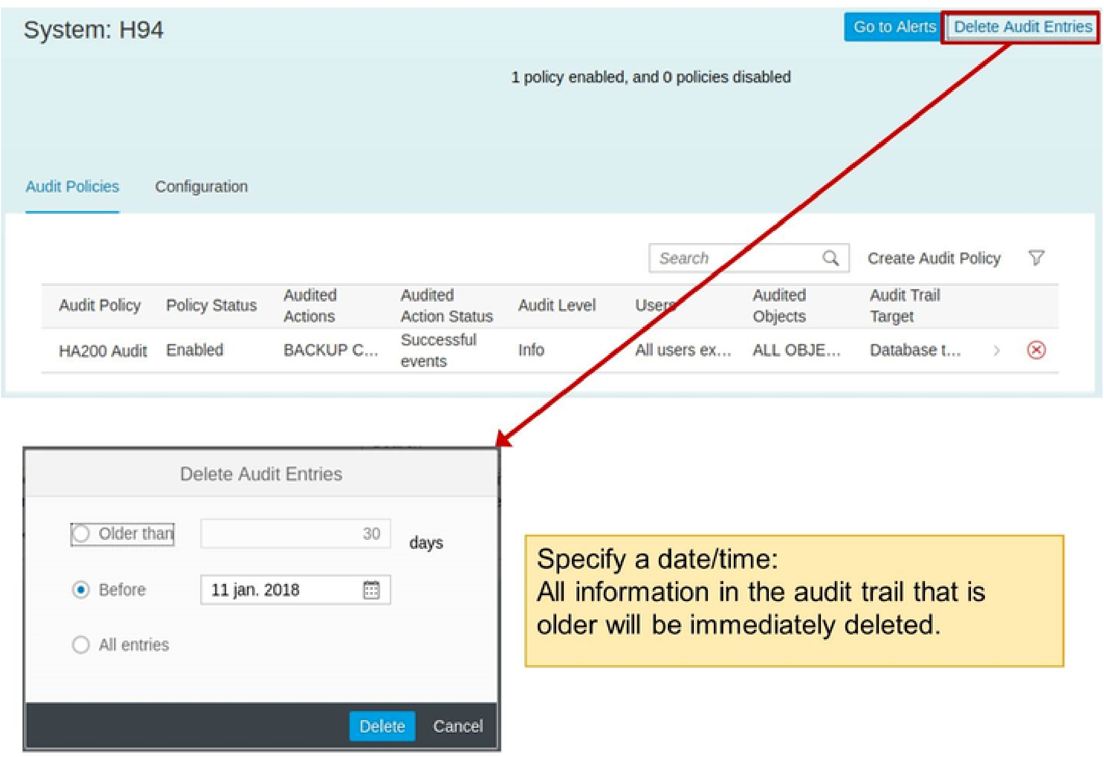

Figure 372: System Settings for Auditing

Caution:
If the table has grown so large that there is not enough memory to delete old entries, use the following SQL command to empty the table completely: `ALTER SYSTEM CLEAR AUDIT LOG ALL`.

Example for Setting Up an Audit Policy with SQL

The figure, Audit Policy Example, shows an example for setting up an audit policy using an SQL statement. It also shows what the audit logging output (audit trail written via Linux syslog) looks like.

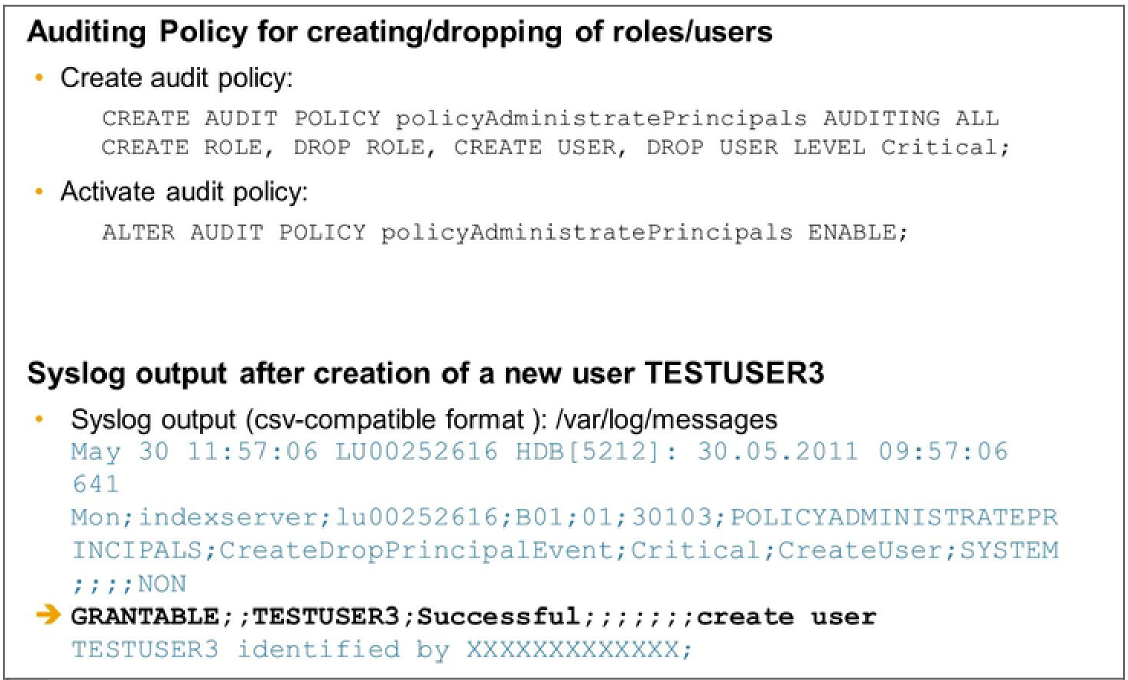

Figure 373: Audit Policy Example

Note:
To create and activate the audit policy, you need root-authorization.

Column header names are not written to the audit trail, so you need to add them manually.

An audit entry appears as follows:

```
<Event Timestamp>;<Service Name>;<Hostname>;<SID>;<Instance
Number>;<Port Number>;<Client IP Address>;<Client Name>;<Client
Process ID>;<Client Port Number>;<Audit Level>;<Audit Action>;<Active
User>;<Target Schema>;<Target Object>;<Privilege
Name>;<Grantable>;<Role Name>;<Target Principal>;<Action
Status>;<Component>;<Section>;<Parameter>;<Old Value>;<New
Value>;<Comment>;<Executed Statement>;<Session Id>;
```

For more information, see the SAP HANA Security Guide at http://help.sap.com/hana.

LESSON SUMMARY
You should now be able to:

- Describe auditing

Learning Assessment

1. In which of the following SAP HANA scenarios is it recommended to encrypt communication between the client software and SAP HANA?

   Choose the correct answers.

   - [ ] **A** SAP HANA as a data mart for reporting and analytics.
   - [ ] **B** SAP HANA as a classic three-tier architecture as the primary database (for example, SAP NetWeaver Business Warehouse, SAP Business Suite Installation).
   - [ ] **C** SAP HANA as a platform for providing database and application services to native SAP HANA-based applications.

2. You cannot force all clients who communicate with the SAP HANA database through the SQL interface to use a secured connection.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

3. What are the prerequisites for configuring audit logging?

   Choose the correct answers.

   - [ ] **A** AUDIT ADMIN system privilege
   - [ ] **B** INIFILE ADMIN system privilege
   - [ ] **C** AUDIT OPERATOR system privilege
   - [ ] **D** INIFILE OPERATOR system privilege

4. Which of the following actions can be audited by SAP HANA?

   Choose the correct answers.

   - [ ] **A** Changes to system configuration.
   - [ ] **B** Upgrade of an SAP HANA database instance.
   - [ ] **C** Create or drop user, and create or drop role.

Unit 11

Learning Assessment - Answers

1. In which of the following SAP HANA scenarios is it recommended to encrypt communication between the client software and SAP HANA?

   Choose the correct answers.

   [X] **A** SAP HANA as a data mart for reporting and analytics.

   [ ] **B** SAP HANA as a classic three-tier architecture as the primary database (for example, SAP NetWeaver Business Warehouse, SAP Business Suite Installation).

   [X] **C** SAP HANA as a platform for providing database and application services to native SAP HANA-based applications.

   Correct! In a data mart scenario, data is replicated from a source system, such as SAP Business Suite, into the SAP HANA database. Reporting is then carried out on the data in SAP HANA. Some end users usually have direct access to SAP HANA. Therefore, user and role management in SAP HANA is required for these end users, as well as for technical users and administrators. Encryption of communication between client and SAP HANA is recommended. SAP HANA extended application services embeds a full-featured application server, Web server, and development environment within SAP HANA. Users of native SAP HANA applications have direct access to SAP HANA. Users must exist in SAP HANA. SAP HANA database privileges and additional application privileges must be assigned. If SAP HANA runs as a relational database in a classic three-tier architecture (client, application server, and database), security-related features are located and are enforced in the application server layer. These security features include authentication, authorization, encryption, and auditing. The database is used as a data store only. End users do not have direct access to either the database itself or the database server on which it is running. Security in the database layer is mainly focused on securing administrative access to the database. Read more on this in the lesson Describing Security Functions (Unit 11, Lesson 1) of the course HA200_14.

2. You cannot force all clients who communicate with the SAP HANA database through the SQL interface to use a secured connection.

   Determine whether this statement is true or false.

   [ ] True

   [X] False

   Correct! To do this, set the “sslEnforce” parameter in the “communication” section of the “global.ini” configuration file to true. The database then refuses SQL connection attempts that don't use SSL. Read more on this in the lesson Explaining Encryption (Unit 11, Lesson 2) of the course HA200_14.

3. What are the prerequisites for configuring audit logging?

   Choose the correct answers.

   [X] **A** AUDIT ADMIN system privilege

   [X] **B** INIFILE ADMIN system privilege

   [ ] **C** AUDIT OPERATOR system privilege

   [ ] **D** INIFILE OPERATOR system privilege

   Correct! AUDIT ADMIN system privilege and INIFILE ADMIN system privilege are the prerequisites for configuring audit logging. Read more on this in the lesson Describing Auditing (Unit 11, Lesson 3) of the course HA200_14.

4. Which of the following actions can be audited by SAP HANA?

   Choose the correct answers.

   [X] **A** Changes to system configuration.

   [ ] **B** Upgrade of an SAP HANA database instance.

   [X] **C** Create or drop user, and create or drop role.

   Correct! Since SPS08, if you enable audit logging for configuration changes, the previous values of parameters are written to the audit trail. An action corresponds to the execution of an action in the database by SQL statement. For example, to track user provisioning in your system, create an audit policy that audits the execution of the SQL statements CREATE USER and DROP USER. You can only audit actions that occur inside the database engine. If the database engine is not online when an action occurs, it cannot be detected and, therefore, cannot be audited. These actions include, for example, an upgrade of an SAP HANA database instance, or direct changes to system configuration files using operating system commands. Read more on this in the lesson Describing Auditing (Unit 11, Lesson 3) of the course HA200_14.

UNIT 12 Maintaining Users and Authorization

UNIT OBJECTIVES

- Understand User Management
- Understand Authentication and Authorization
- Explain the different types of privileges
- Understand roles
- Perform Administrative Tasks
- Explain Information Sources for Administrators

Unit 12 Lesson 1

SAP HANA Authentication and Authorization

LESSON OVERVIEW

In this lesson, you will learn about authentication models and authorizations.

Business Example

The SAP HANA database facilitates the integration of different authentication methods. To integrate the SAP HANA database in your environment, you need an overview of the supported authentication methods.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Understand User Management
- Understand Authentication and Authorization

User Management in Multitenant Database Container Systems

In a system with multitenant database containers, each tenant database has its own database administrator and end users.

The system database and all tenant databases each have their own SYSTEM user. The SYSTEM user of the system database has additional privileges for managing tenant databases, for example, creating and dropping databases, changing configuration (*.ini) files of databases, and performing database-specific data backups.

In a multiple-container system, system privileges granted to users in a particular multitenant database container authorize operations in that database only. The only exception is the system privilege DATABASE ADMIN. This system privilege can only be granted to users of the system database. It authorizes the execution of operations on individual tenant databases. For example, a user with DATABASE ADMIN can create and drop tenant databases, change the database-specific properties in configuration (*.ini) files, and perform database-specific backups.

Overview: User Administration

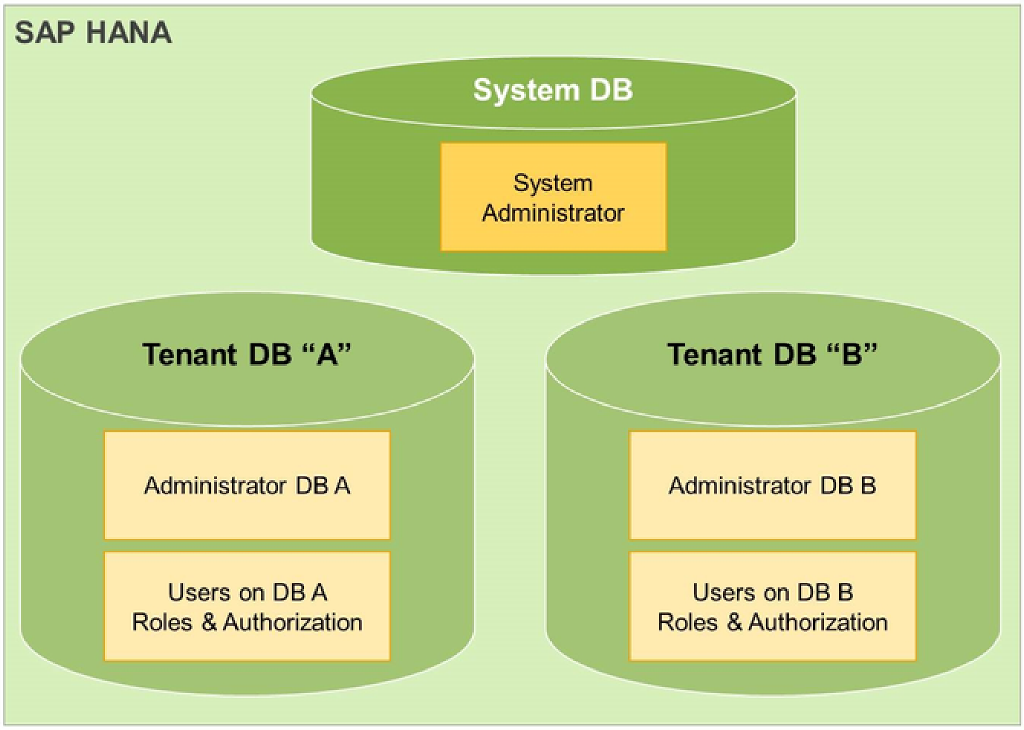

Figure 374: Overview: User Administration

In a multiple-container system, privileges granted to users in a particular database authorize access to and modification of database objects in that database only. That is, unless cross-database access has been enabled for the user. This is made possible through the association of the requesting user with a remote identity on the remote database.

User Management and Security in SAP HANA: Overview

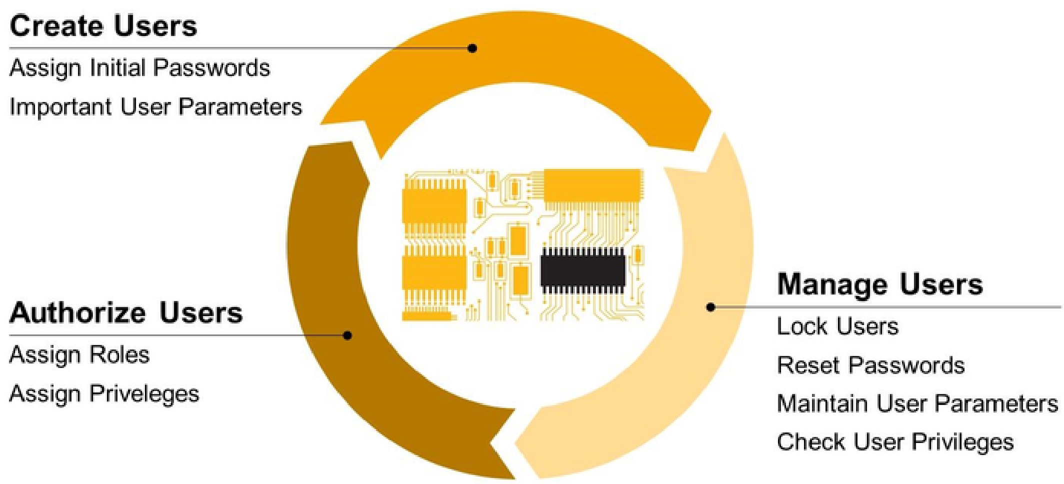

Figure 375: User Management and Security in SAP HANA

A security concept is required in SAP HANA for the following reasons:

- To restrict database administration to skilled (and empowered) people
- To restrict access to SAP ERP tables
- To restrict editing of SAP HANA data models to owners of the model

Security is important in SAP HANA as user administration plays a significant role, as follows:

- Several front end tools offer direct access into SAP HANA.
- Access to objects and to data model content must be controlled within SAP HANA.
- Information Consumers need named users in SAP HANA.

An exception to the security concept is when Information Consumers do not require user management. This occurs in the following situations:

- Access to data does not need to be controlled.
- All data access occurs through the SAP Business Intelligence (BI) semantic layer, and security is implemented in SAP BusinessObjects Enterprise.

Relationships Between Entities

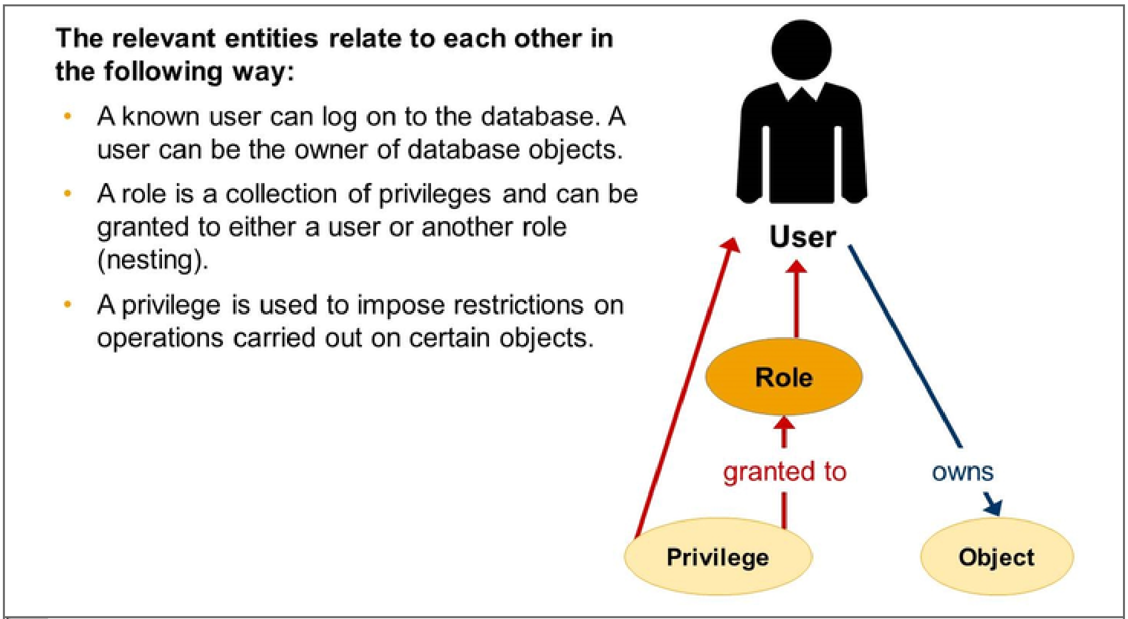

Figure 376: Relationships Between Entities

Privileges can be assigned to users directly or indirectly using roles. Privileges are required to model access control. Roles can be used to structure the access control scheme and to model reusable business roles.

You can manage authorization for users using roles. You can nest roles so that role hierarchies can be implemented. This increases their flexibility, allowing for both precise and broadscale authorization management for individual users.

All the privileges granted directly or indirectly to a user are combined. This means that whenever a user tries to access an object, the system performs an authorization check using the user, the user's roles, and directly allocated privileges.

You cannot explicitly deny privileges. This means that the system does not need to check all the user's roles. Once all the requested privileges are located, the system aborts the check and grants access.

Several predefined roles exist in the database. Some of them are templates that need to be customized; others can be used as they are.

Authorization

All access to data and the execution of actions in the database requires authorization. Every user who wants to work directly with the SAP HANA database must have a database user with the necessary authorizations.

Authorization Framework

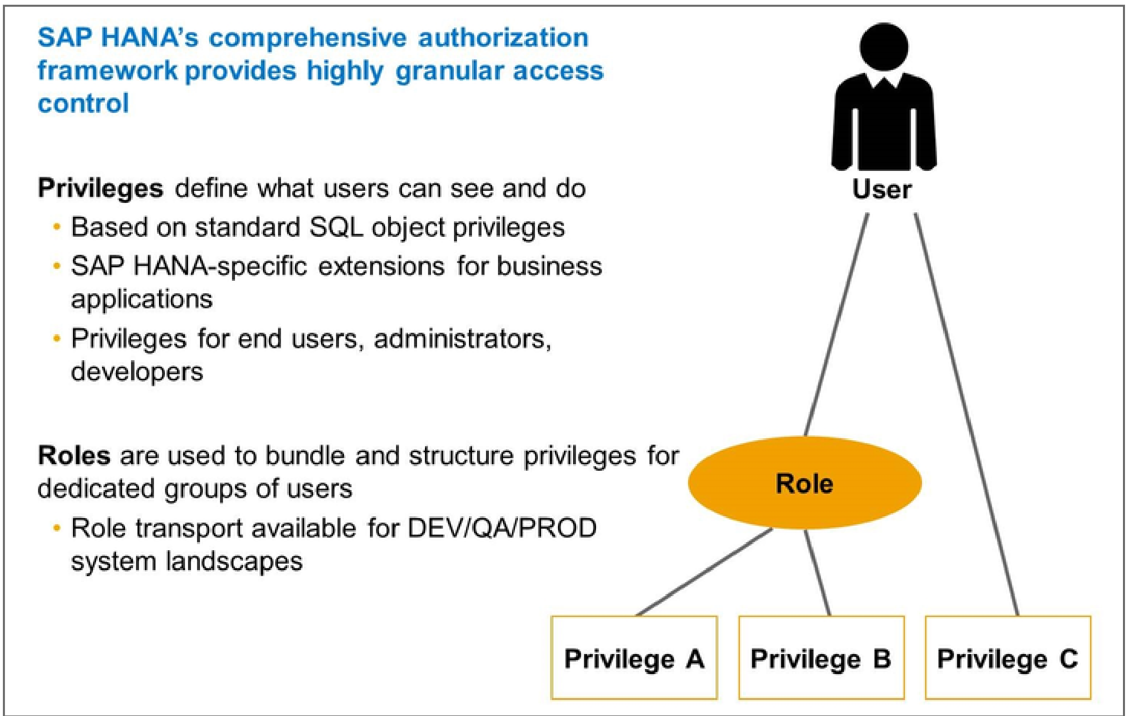

Figure 377: Authorization Framework

After successful logon, the system verifies the user's authorization to perform the requested operations on the requested objects. This is determined by the privileges that the user has been granted. The user must have both the privilege to perform the operation and the privilege to access the object (for example, a table) to which the operation applies.

Privileges can be granted to database users either directly, or indirectly through roles. A role is a set of privileges. Roles are the standard mechanism of granting privileges because they allow you to implement both fine-grained and coarse-grained reusable authorization concepts that can be modeled on business roles. Several standard roles are also delivered with the SAP HANA database (for example, MODELING, MONITORING). You can use these as templates for creating your own roles.

Authorizations Assigned by Privileges and Roles

Privilege(s)	➔ A user in SAP HANA must have the following privilege(s) assigned: • Privilege(s) to perform the operation • Privilege(s) to access the object (for example, a table) to which the operation applies ➔ Privileges can be granted to database users either directly, or indirectly through roles. Note that it is recommended to use roles. ➔ There is a privilege concept for both design time (developer use case) and runtime. (See next slide for details).
Role(s)	➔ A role is a set of privileges. ➔ Roles are the standard mechanism of granting privileges as they allow you to implement fine-grained user access that can be modelled on business roles. ➔ Several standard roles are also delivered with the SAP HANA database (for example, MODELING, MONITORING). These can be used as templates for creating your own roles. ➔ It is also possible to create roles as design-time objects in the repository of the SAP HANA database. Design-time roles can be transported from the development or QA system to the production system, where they are activated to be available in runtime.

Figure 378: Authorizations Assigned by Privileges and Roles

The figure, Authorizations Assigned by Privileges and Roles, shows the privileges and roles of a user in SAP HANA.

User Provisioning

Roles allow grouping privileges

- Create roles for specific tasks, for example:
  - Create data models (on a given subset of the data)
  - Activate data models
  - Consume models
- All types of privileges can be granted to a role
  - Individual privileges
  - Roles (→ create a hierarchy of roles)

Roles / privileges can be assigned to users

- User / Role management are closely related
- Reflected in almost identical editor

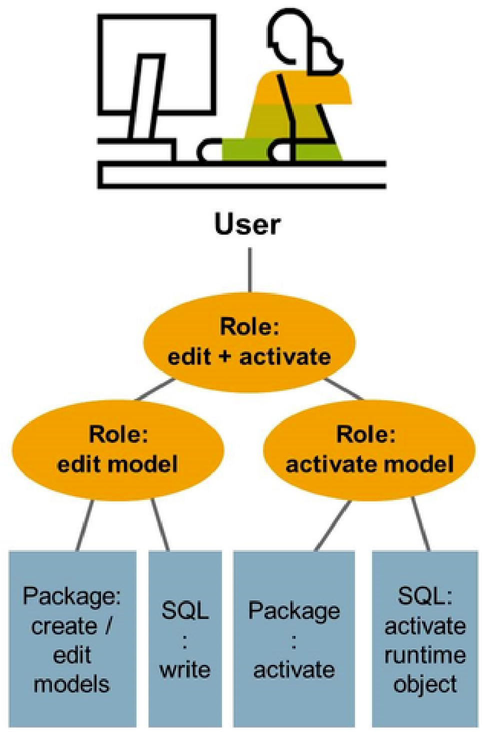

Figure 379: User Provisioning

When accessing the SAP HANA database using a client interface (such as Open Database Connectivity [ODBC], Java Database Connectivity [JDBC], and Multidimensional Expressions [MDX]), any access to data must be backed by corresponding privileges. Different schemes are implemented. On a higher level, this concept provides authorization for the data contained in the database when it is accessed using client interfaces. In the SAP HANA database system, the regular SQL authorization concept is implemented.

For each SQL statement type (for example, SELECT, UPDATE, and CALL), the executing user needs to have a corresponding privilege. Additionally, objects in the database (such as tables, views, or stored procedures) have an owner who can access the objects and grant privileges for them.

No user can access this particular object, other than the owner of an object and users that the owner has provided with a privilege. This authorization operates at the object level, where the smallest entities that can be privileged are, for example, a table or a view.

In addition, analytic privileges provide row-level authorization for certain database objects, such as analytic views.

Managing Users and Roles

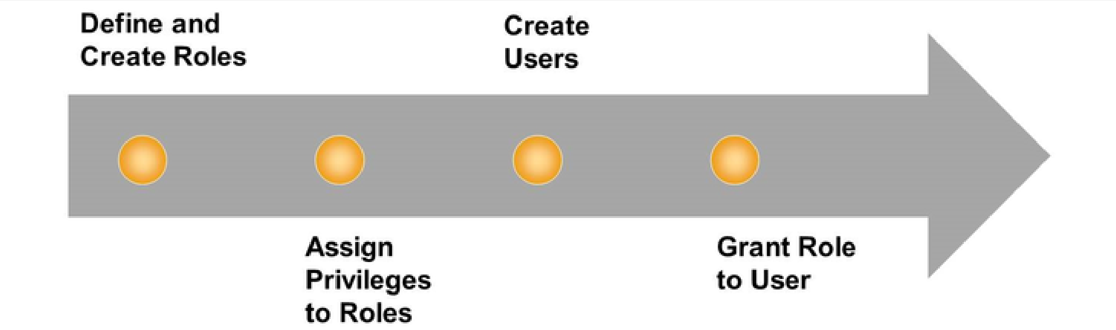

Figure 380: Managing Users and Roles

The process flow for user management is as follows:

1. Define and create privileges.
2. Define and create roles.

   Use the SAP HANA studio or run the following SQL statement: `CREATE ROLE <Role_Name>`
3. Assign privileges to roles.
4. Create users.

   Choose the following authentication methods:

   - Define the initial password
   - Define the external User ID (for example, Kerberos to set up SSO)
   - Other user settings

   Define default client is used as an implicit filter value when reading from SAP HANA data models.
5. Grant roles to users.

Use the SAP HANA studio or run the following SQL statement: `GRANT <Role_Name> TO <user>`

To revoke roles, use the following SQL statement: `REVOKE <Role_Name> FROM <user>`

User Management

User administrators create and configure users in SAP HANA. User administrators have the system privilege USER ADMIN which authorizes them to manage all users.

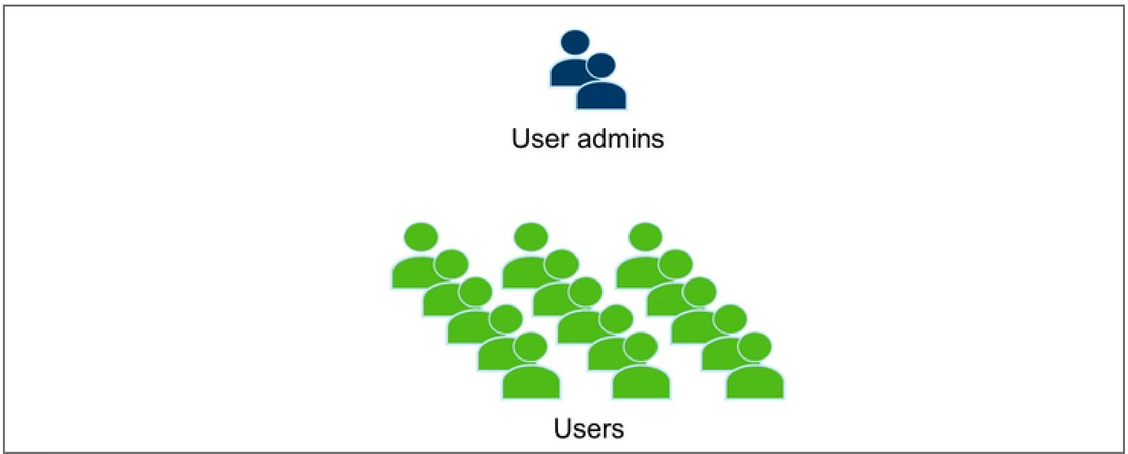

Figure 381: User Management

User Groups

User groups allow you to manage related users together. Dedicated group administrators can be assigned to manage individual user groups exclusively and independently of each other.

In the example scenario shown in the figure below, the user groups were set up for exclusive administration. Only the group admins for the restricted user group Sales can create or delete users in this user group, and manage security properties of the Sales users. The Sales group admins can only manage users of the Sales group, but not users from other user groups like Research or Training.

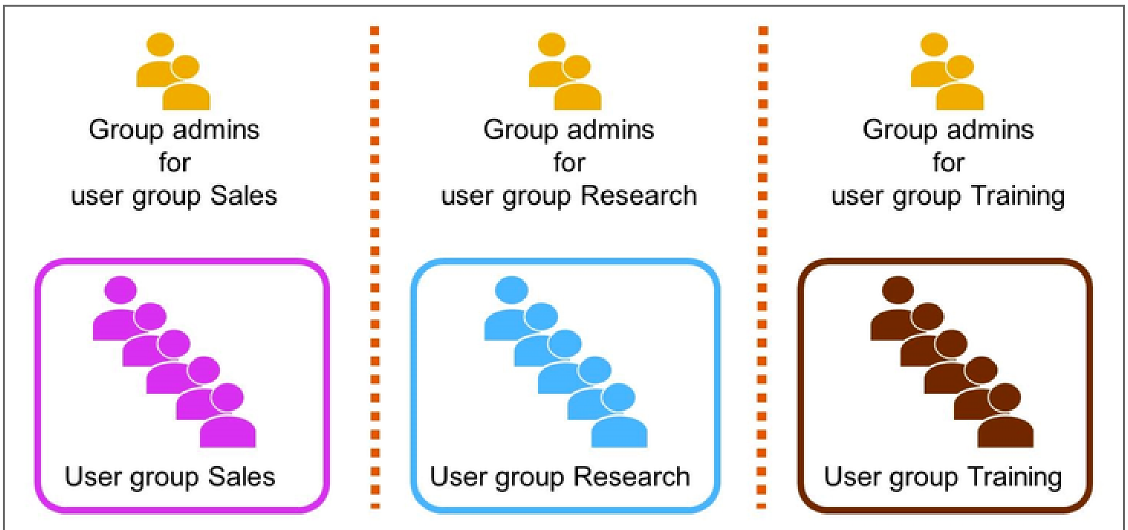

Figure 382: User Groups

The privilege USERGROUP OPERATOR authorizes a user to change the settings for a user group, and to add and remove users to/from a user group. Users with the USERGROUP

OPERATOR privilege can also create and drop users, but only within the user group they have the USERGROUP OPERATOR privilege on (`CREATE USER <user_name> SET USERGROUP <usergroup_name>`).

A user can have the USERGROUP OPERATOR privilege on more than one user group, and a user group can have more than one user with the USERGROUP OPERATOR privilege on it.

Creating user groups

1. A global user administrator (system privilege USER ADMIN) creates a user group.

   Optionally the user administrator can set up the user group for “exclusive administration”.

2. A role developer includes the new object privilege USERGROUP OPERATOR on the user group in a role.

3. A role administrator assigns the role to one or more users, who will be the group operators of the new user group.

Creating new users

When a new user is created, optionally a user group can be specified. A user can belong to only one user group. But users do not need to belong to any user groups.

User Types

You can specify different security policies for different types of users. In the SAP HANA database, you can differentiate between the following user types:

- Database users that correspond to real people

  The database administrator creates a database user for every person who works in the SAP HANA database. Database users that correspond to real people are dropped when the person leaves the organization. This means that the database objects that they own are also automatically dropped, and privileges that they granted are automatically revoked.

- Technical database users

  Technical database users do not correspond to real people. They are therefore not dropped if a person leaves the organization. Therefore, you can use them for administrative tasks such as creating objects and granting privileges for a particular application. Some technical users are available as standard, for example, the users SYS and _SYS_REPO. Other technical database users are application-specific. For example, an application server can log on to the SAP HANA database using a dedicated technical database user.

> Note:
> All user names can now contain unicode characters.

Database users that correspond to real people

- User for every person who works in the SAP HANA database

Technical database users

- Do not correspond to real people
- Used for administrative tasks
- Application-specific technical users

Figure 383: User Types

Technically, these user types are the same. The only difference between them is conceptual.

Database users that correspond to real people can be grouped according to different tasks.

> Hint:
> For a complete overview of the technical users of the SAP HANA database, see the SAP HANA Security Guide.

User Tasks

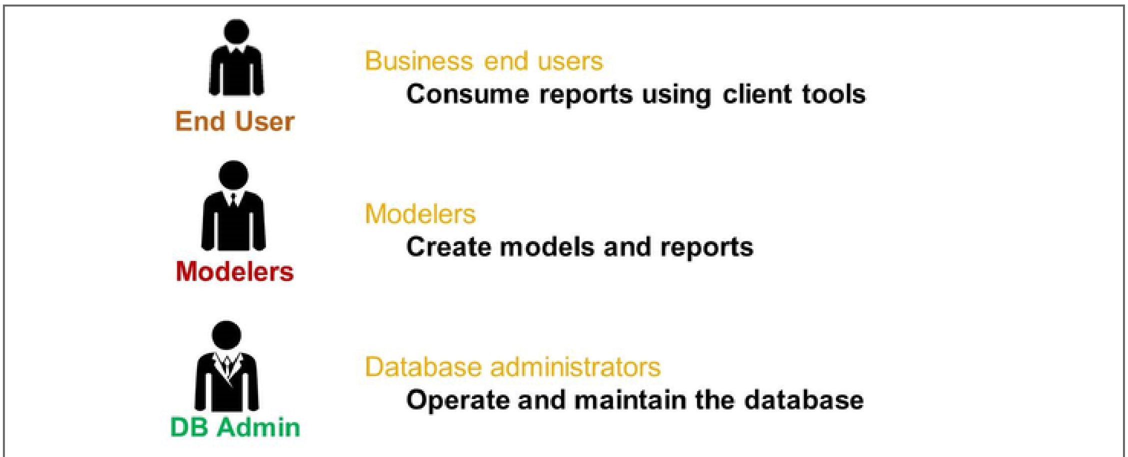

Figure 384: User Tasks

The SAP HANA database user and role concept allows for a fine granularity of access control based on the users' tasks. Some examples are as follows:

- Business end users reading reports using client tools, for example, Microsoft Excel
- Modelers creating models and reports using the SAP HANA studio
- Database administrators operating and maintaining the database, and users using the SAP HANA studio

Standard Users Installation, Upgrade, and Operation

Database User	**SYSTEM**	**Overall System Administrator**
Operating System User	**<sid>adm**	**SAP System User**
	ROOT	**User Installation and Upgrade**

Figure 385: Standard Users Installation, Upgrade, and Operation

To install, upgrade, and operate the SAP HANA database, the following standard users are required:

- Database Users

  When you install the SAP HANA database, a database user, called SYSTEM, is created by default. The database user SYSTEM has irrevocable system privileges, such as the ability to create other database users, access system tables, and so on.

> Note:
> For security reasons, do not use user SYSTEM for day-to-day activities. Use SYSTEM to create administration users with the minimum privilege required for their duties set. Then use those users for day-to-day administrative activities.

  Several internal database users are also created, such as SYS and _SYS_STATISTICS. These users cannot log on to the SAP HANA database.

- Operating System User

  The installation process also creates an external operating system user <sid>adm (for example, sp1adm or xyzadm). This operating system user, referred to as the operating system administrator , simply exists to provide an operating system context. From the operating system perspective, the operating system administrator is the user that owns all SAP HANA files and all the related operating system processes. Within the SAP HANA studio, the operating system administrators credentials are required, for example, to start or stop database processes or to execute a recovery. The operating system administrator is not an SAP HANA database user.

For installation and upgrade, the ROOT user is used. Do not use the ROOT user for day-to-day activities.

Cross-Tenant Database Access

There are situations where queries should run across tenant databases. In multiple-container systems, read-only queries across database containers are supported but are not enabled by default.

Read-only queries between multitenant database containers are possible through the association of the requesting user with a remote identity on the remote database or databases. Cross-database queries (federation) are supported in SQL engine and Calculation engine.

Cross Database Queries Between Multitenant Database Containers

Every tenant database in a multiple-container system is self-contained with its own isolated set of database users and isolated database catalog. However, to support cross-application reporting, cross-database SELECT queries are possible. This means that database objects such as tables and views can be local to one database, but be read by users from other databases in the same system.

A user in one database can run a query that references objects in another database, if the user is associated with a sufficiently privileged user in the remote database. This associated user is called a remote identity. This is the user who executes the query (or part of the query) in the remote database, and therefore the user whose authorization is checked.

Cross-database access is not enabled by default and must be configured before such user mappings can be set up.

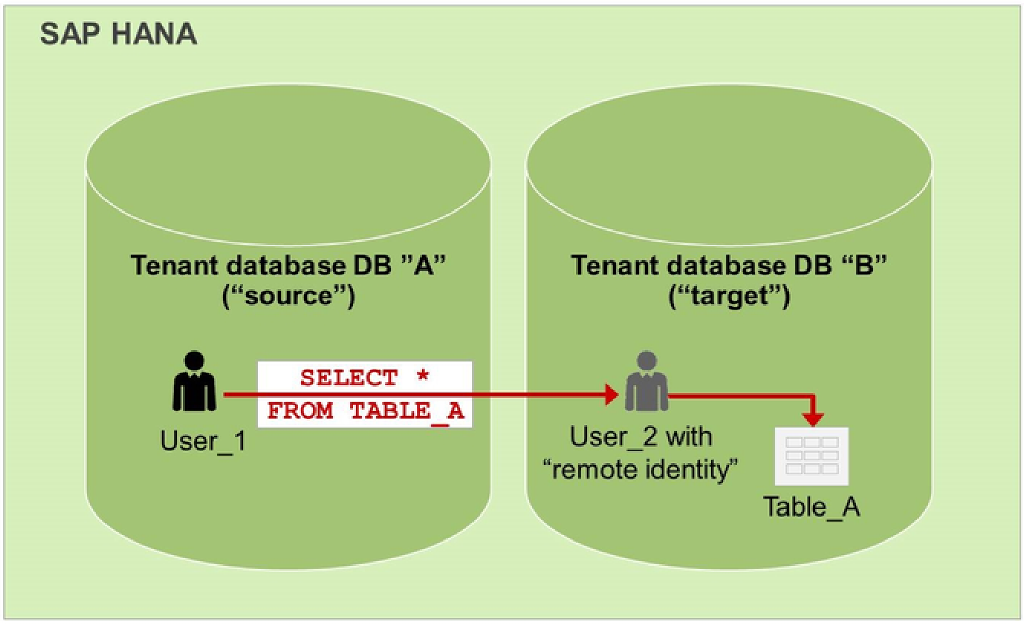

Figure 386: Cross Database Queries Between Multitenant Database Containers

By default, cross database access between tenants is inactive. To run queries spanning multiple tenant databases, the global cross database access switch has to be turned on. A whitelist of databases that are allowed to communicate with each other also has to be set up.

Activation of Cross-Tenant Database Access

To access a cross-tenant database, the following must be activated:

- Turn on cross-tenant database communication (run this from SYSTEM database only)

- Whitelisting a cross-tenant database communication channel (from SYSTEM database only)
- Add a remote identity to the requesting user on the remote database

> Hint:
> Communication channels are unidirectional by default. They can be made bidirectional by explicitly defining the configuration in reverse.

If enabled, a user from one tenant database can execute queries in another tenant database if this user is mapped to a user with "remote identity" there, as follows:

- A user in the target database can only be associated with one user in the source database.
- The association is unidirectional.
- Only the SELECT privileges of the user in the target database are considered during a cross-database query, all other privileges of the remote user are ignored.

User Administration Tools

User management is configured using the SAP HANA cockpit or the SAP HANA studio.

There is no replication of existing authorizations from the source system.

SAP HANA cockpit and SAP HANA studio

- Creating users
- Deleting, deactivating, and reactivating users
- Creating roles
- Assigning roles and privileges to users

User & Role Management
Manage users
Assign role to users
Assign privileges to users
Manage roles

Command line interface (hdbsql or other SQL Tool)

- Performing all administration tasks using SQL commands
- For example, run the following statement:

```
CREATE USER <user_name> PASSWORD
<password>
CREATE ROLE <role_name>;
```

SAP NetWeaver Identity Management

- Creating and deleting user accounts
- Assigning roles
- Setting passwords of users

Figure 387: User Administration Tools

All the user management functions can also be executed from the command line using SQL requests. This is useful when using scripts for automated processing.

SAP Identity Management provides additional support for user provisioning in the SAP HANA database.

The SAP Identity Management 7.2 SP 3 contains a connector to the SAP HANA database (IDM connector). With the SAP Identity Management, you can perform the following actions in the SAP HANA database:

- Create and delete user accounts
- Assign roles
- Set passwords for users

For more information about the SAP Identity Management and the IDM connector, see the SAP Community Network at http://www.sdn.sap.com, under SAP NetWeaver Releases.

Create User

You can create a standard database user for every person who works directly with the SAP HANA database. When you create a user, you also configure how the user is authenticated. You can do this on the User page of the SAP HANA cockpit.

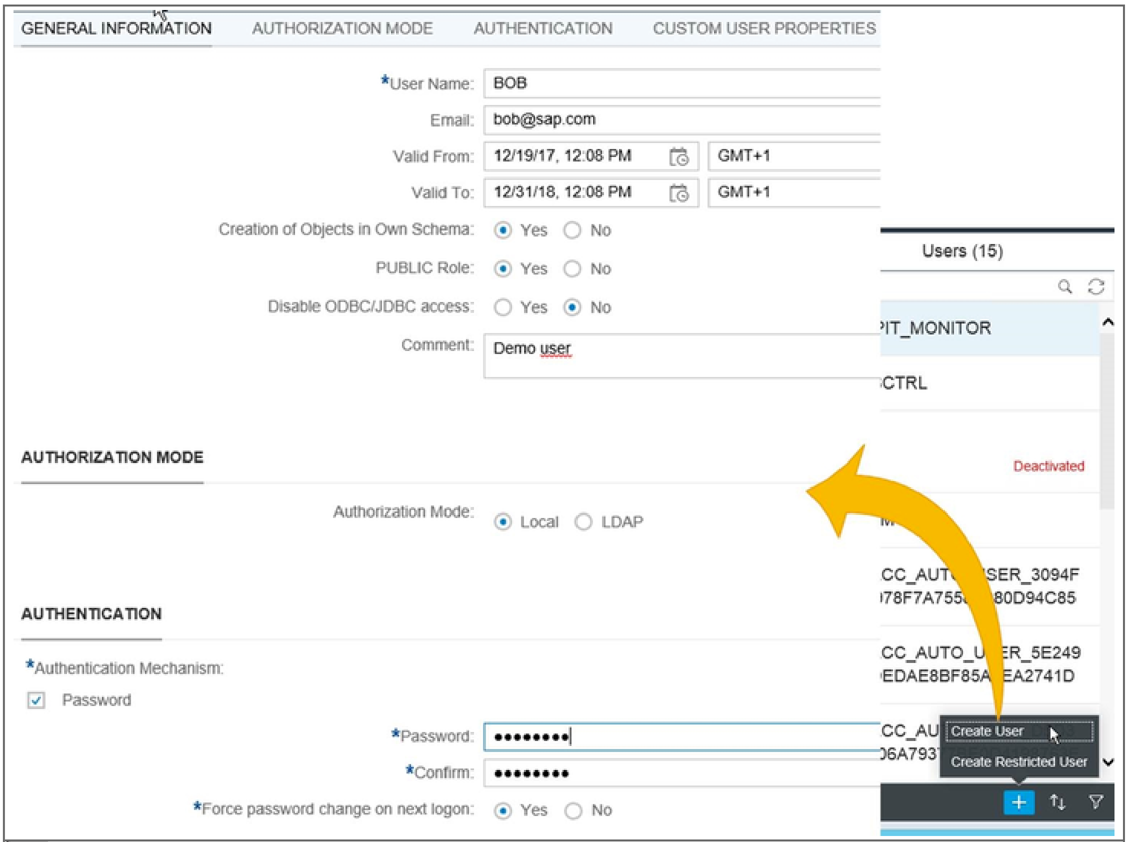

Figure 388: Create Users in SAP HANA Cockpit

When you create a user, specify the following properties:

- General Information
  - User Name

    Enter a unique user name.
  - Optional: Email address

Enter the user's email address.

- Optional: Validity Period, including the appropriate time zone

  Enter the validity period of the user. If the user account is not currently within its validity period, the user is inactive and cannot log on. If no validity period is configured, the user is indefinitely valid.

- Optional: Creation of Objects in Own Schema.

  Prevent the user from being able to create objects in his own database schema.

- Optional: PUBLIC Role.

  Prevent the user from being granted the standard PUBLIC role.

  The PUBLIC role contains the privileges for filtered read-only access to the system views. To see data in a particular view, the user also needs the SELECT privilege on the view.

- Optional: Disable ODBC or JDBC Access.

  This indicates whether or not the user can connect to the database via ODBC or JDBC. By default, ODBC and JDBC access is enabled for standard users and is disabled for restricted users. This means that restricted users can only connect via HTTP or HTTPS.

- Optional: Set the authorization mode to LDAP if the user's authorization is based on LDAP group membership.

- Specify how the user can be authenticated.

Authentication Methods

Authentication

- Password

  This indicates whether or not authentication for the user name password is enabled.

- Force password change on next logon

  This indicates if the user must change a password set by a user administrator when they log on for the first time.

- Kerberos
- SAP Logon Ticket, SAP Assertion Ticket
- SAML
- X509

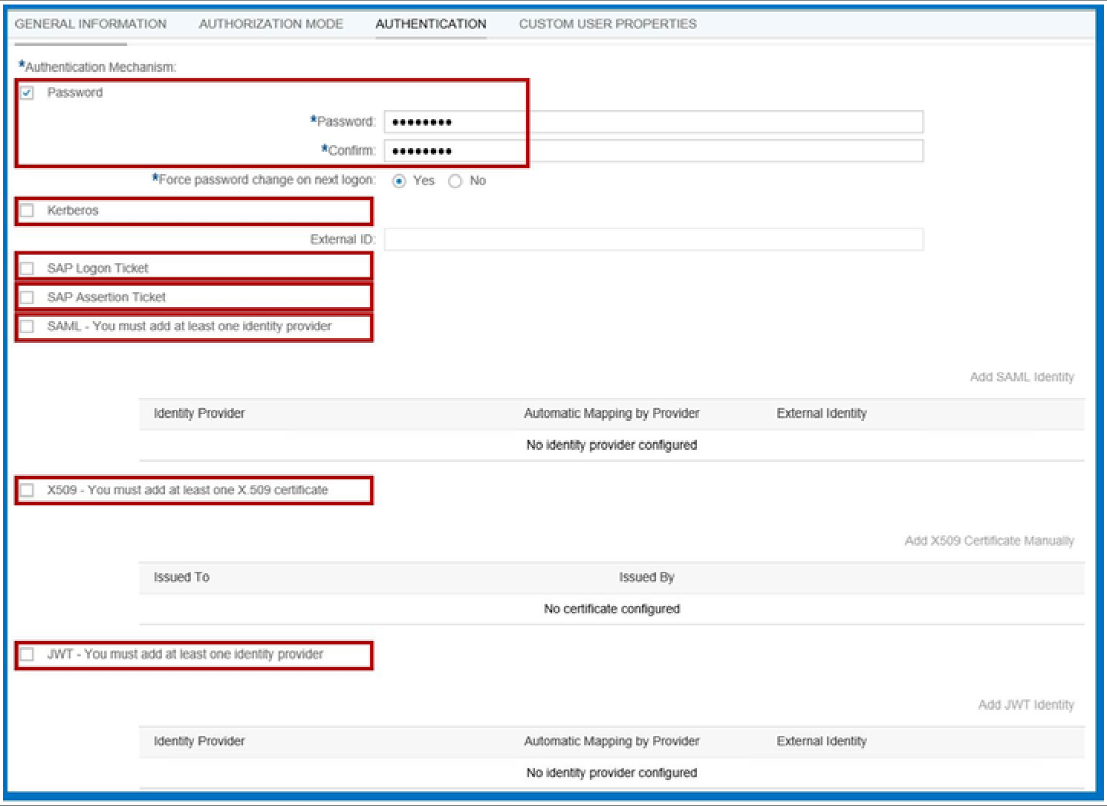

Figure 389: Authentication Methods

Additional User Parameter

You can configure additional user properties for client applications. The following Custom User properties are available by default:

- LOCAL

  This is the user's locale.

- PRIORITY

  This is the priority with which the thread scheduler handles statements executed by the user. The priority can be in the range 0-9, with 9 representing the highest priority. 5 is the default priority.

- STATEMENT MEMORY LIMIT

  This is the maximum memory (in GB) that can be used by a statement executed by the user.

- STATEMENT THREAD LIMIT

  This is the maximum number of threads that can be used by a statement executed by the user.

- TIME ZONE

  This is the user's timezone. The standard database formats for locale and timezone are supported.

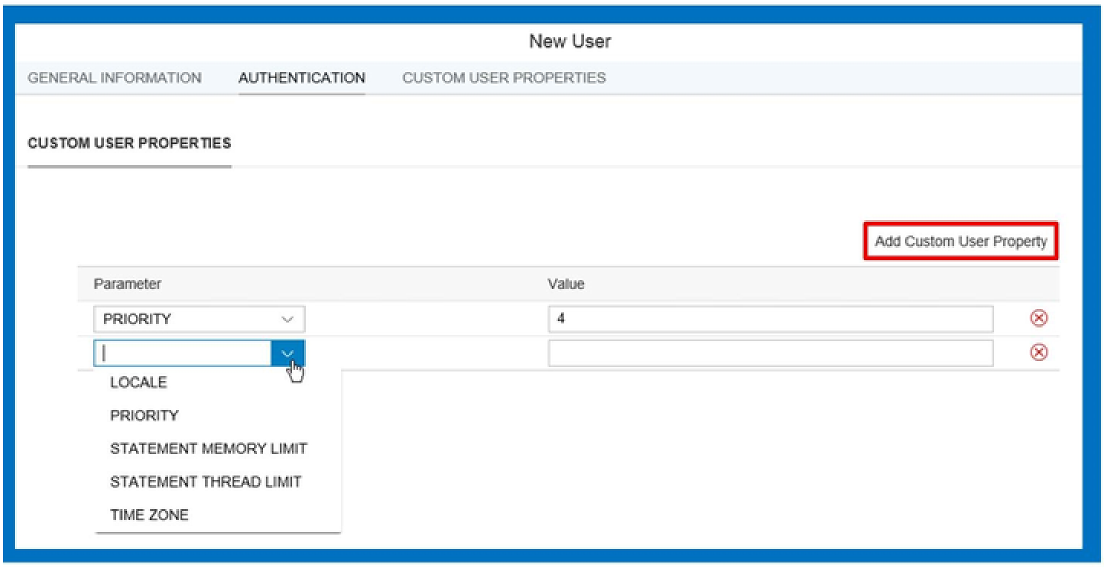

Figure 390: Additional User Parameters

Copy a User

If you are implementing user authorization through roles created in the SAP HANA repository, you can create a new user by copying an existing user. The repository roles granted to the existing user are automatically granted to the new user. SQL roles and individual privileges are not granted. You can copy a role in this way using the SAP HANA studio.

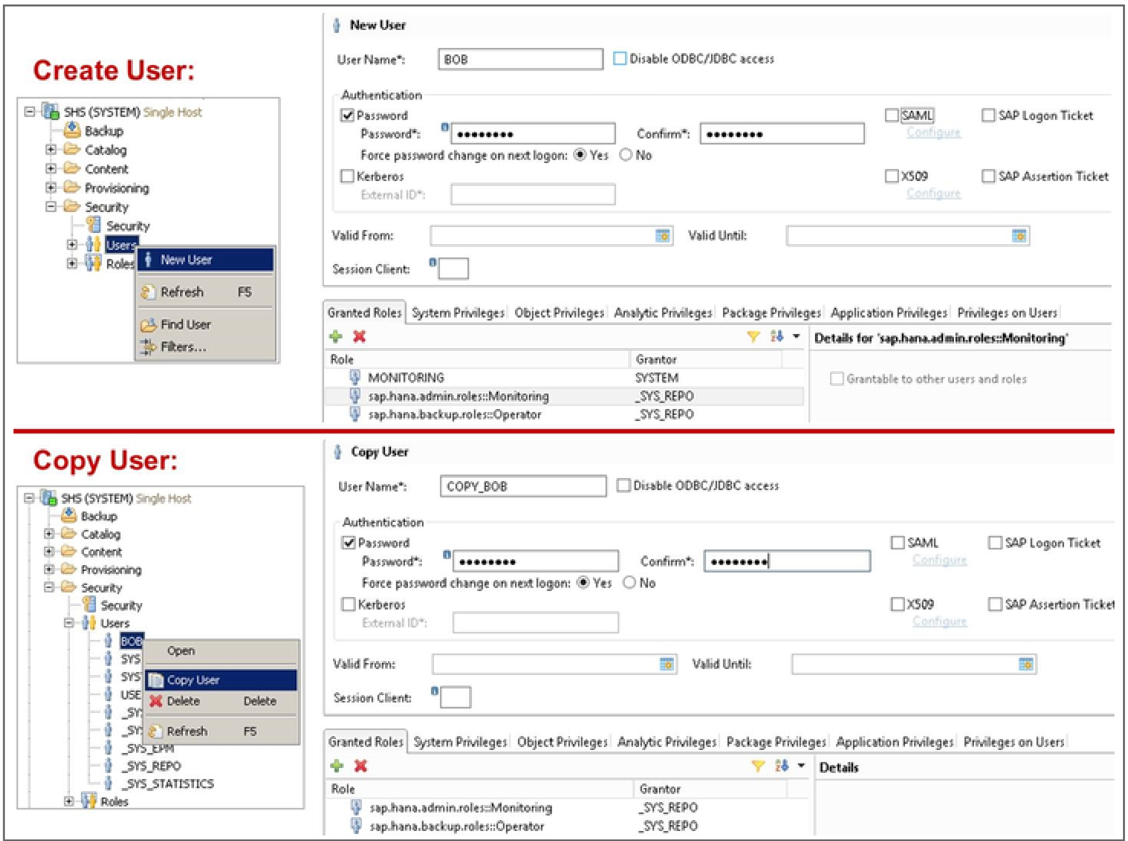

Figure 391: Ways To Define and Copy a User with SAP HANA Studio

Note:
Only roles created in the SAP HANA repository are granted. SQL roles, including the standard roles delivered with the SAP HANA database (MONITORING, MODELING, and so on), and individual privileges are not granted.

User Access Channels

Users can connect to SAP HANA through JDBC or ODBC, or through HTTP/S.

The protocol for database client access (SQLDBC (ODBC/JDBC)) is used in the following scenarios:

- Application servers that use SAP HANA as a database
- End-user clients that access the SAP HANA database directly
- SAP HANA cockpit and SAP HANA studio

HTTP/S client access is used for a connection between a Web browser or a mobile device to applications based on SAP HANA extended application services.

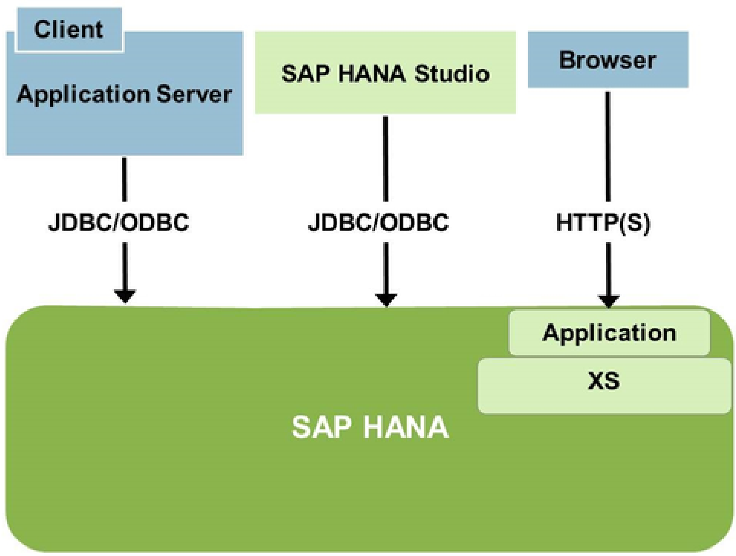

Figure 392: User Access Channels

User Types in SAP HANA

Database users are created as either standard users or as restricted users.

Standard users can create objects in their own schema and read data in system views. The PUBLIC role grants read access to system views, and this is granted to every standard user.

Restricted users initially have no privileges. Restricted users are for users who access SAP HANA through client applications and who do not have full SQL access through an SQL console. If the privileges that they need to use the application are included within an application-specific role, then you only have to grant the user this role. In this way, you can ensure that users only have those privileges that are essential to their work.

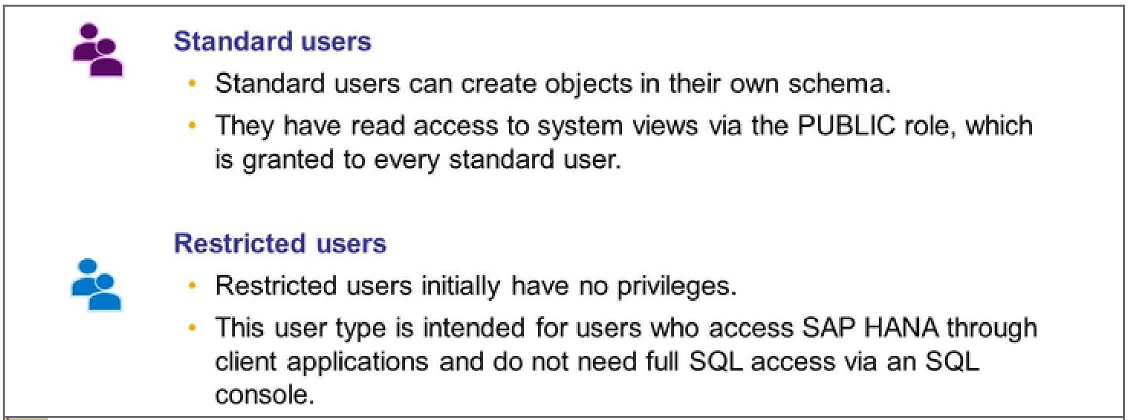

Figure 393: User Types in SAP HANA

Limitations of Restricted Users

Compared to standard database users, restricted users are initially limited in the following ways:

- They cannot create objects in the database because they are not authorized to create objects in their own database schema.
- They cannot view any data in the database because they are not granted (and cannot be granted) the standard PUBLIC role.
- They are only able to connect to the database using HTTP/HTTPS.

To enable a restricted user to use an application, grant the required application-specific roles. Initially restricted users can only connect to the database using HTTP or HTTPS.

> Note:
> To allow restricted users to connect through ODBC or JDBC, enable ODBC or JDBC access explicitly.

Authentication Methods

The identity of every database user accessing the database is verified through a process called authentication.

Authentication Methods

The SAP HANA database provides the following options for authentication:

- Direct logon to the SAP HANA database with user name and password
- Authentication using Kerberos (third-party authentication provider)
- Authentication using Security Assertion Markup Language (SAML) bearer token
- Authentication using SAP Logon Ticket and SAP Assertion Ticket
- Authentication using X.509 certificates
- JSON Web Tokens (JWT)

Overview of Authentication Methods for SQLDBC and HTTP Access

Authentication method	SQLDBC access (JDBC/ODBC)	HTTP access (SAP HANA XS)
Username/password	yes	yes
Kerberos	yes	yes (SPNEGO)
SAML (version 2)	yes	yes
SAP logon ticket	yes	yes
SAP assertion ticket	yes	yes
X.509	-	yes
JWT	yes	SAP HANA XSA

Figure 394: Overview of Authentication Methods for SQLDBC and HTTP Access

When using direct logon to the SAP HANA database with user name and password, the SAP HANA database authenticates the user.

By default, all authentication mechanisms are enabled. However, you can disable those that are not used in your environment. You do this by configuring the parameter `[authentication] authentication_methods` in the `global.ini` configuration file. The value of this parameter specifies all enabled methods as a comma-separated list.

The default value is `password,kerberos,spnego,saml,saplogon,x509xs,sessioncookie`.

> Note:
> Some administrative operations, such as database recovery, also require the credentials of the SAP operating system user (<sapsid>adm).

Kerberos

A user connecting to SAP HANA through Kerberos must have an SAP HANA database user. This database user is mapped to the external identity in a Key Distribution Center (KDC), such as Microsoft Active Directory.

For integration into Kerberos-based SSO scenarios, SAP HANA supports Kerberos version 5 based on Active Directory (Microsoft Windows Server) or Kerberos authentication servers. For HTTP access using SAP HANA extended application services, classic model, Kerberos authentication is enabled with Simple and Protected GSSAPI Negotiation Mechanism (SPNEGO).

Kerberos is a network authentication protocol that provides authentication for client-server applications across an insecure network connection using secret-key cryptography.

ODBC and JDBC database clients support the Kerberos protocol, for example, the SAP HANA studio. You can also implement access from front end applications (for example, SAP BusinessObjects XI applications) using Kerberos delegation. Support for constrained delegation and protocol transition is limited to scenarios in which the middle-tier application connects to SAP HANA as the database layer via JDBC.

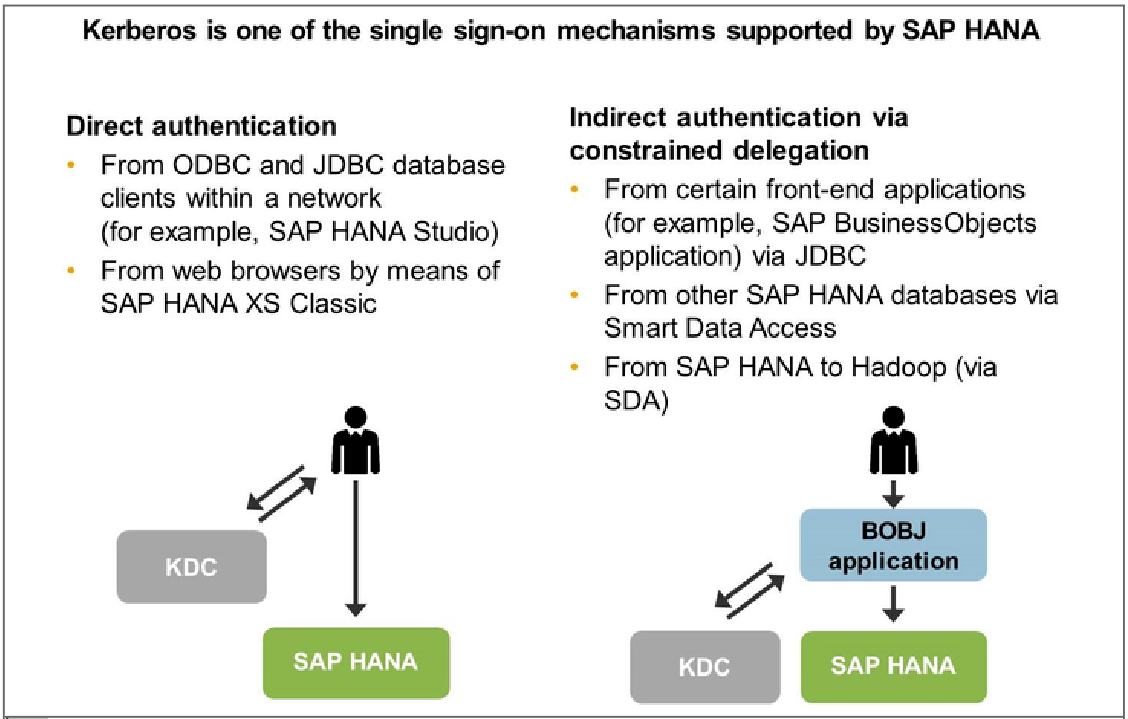

Figure 395: Authentication via Kerberos

Note:
For more information, see SAP Note 1837331 : How-To: HANA DB SSO Kerberos/ Active Directory.

Security Assertion Markup Language (SAML)

The SAP HANA database supports the login of users to the SAP HANA database using the Security Assertion Markup Language (SAML).

You can select SAML as a user authentication method when creating users in the SAP HANA studio.

SAML is a standard protocol for authentication

- Requires a trusted identity provider (SAP NetWeaver SSO) that can issue SAML assertions for clients (for example, browser)

Use Cases

- Single sign-on in middleware/application server scenarios
  - Whenever the application server needs to connect to the SAP HANA database on behalf of a user, it requests a SAML assertion from the client
  - The SAML assertion is issued by the identity provider after the client was successfully authenticated there, and is then sent to the application server. From there it is forwarded to the SAP HANA database

Restrictions

- The SAP HANA database can only act as a SAML service provider
- Assertions can be used for authentication only (no support for further properties)

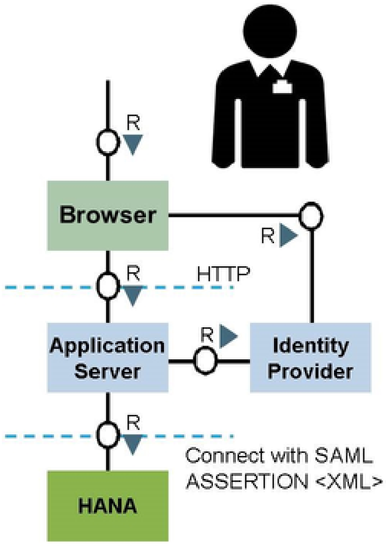

Figure 396: Additional Authentication Method, SAML

Features of SAML

The features of SAML are as follows:

- SAML, Security Assertion Markup Language, is the XML-based standard for communicating identity information between organizations. The primary function of SAML is to provide Internet Single Sign-On (SSO) for organizations. Organizations use SAML to securely connect Internet applications that exist both inside and outside the organization's firewall.
- SAML is a standard protocol for authentication. Internet SSO is a secure connection that communicates identity and trust from one organization to another. For users, Internet SSO eliminates additional logins to external resources. For system administrators, it improves security and reduces costs.

  SAML requires a trusted third-party (identity provider) that can issue SAML assertions for clients (for example, a browser).
- SSO in middleware or application server scenarios are as follows:
  - Whenever the application server needs to connect to the SAP HANA database on behalf of a user, it requests an SAML assertion from the client.
  - The SAML assertion is issued by the identity provider after the client is authenticated successfully, and is then sent to the SAP HANA database.
- SAML restrictions are as follows:
  - The SAP HANA database can only act as an SAML service provider.
  - Assertions can be used for authentication only; there is no support for further properties.

You cannot use SAML for authorization.

The configuration page for SAML identity providers is located in the Security editor in SAP HANA studio.

> Note:
> Previously, this configuration page was available in the system properties.

Configure SAML in SAP HANA

You can configure SAML providers for ODBC or JDBC-based SAML authentication using the SAP HANA studio or SQL statements. However, always use the SAP HANA extended application services administration tool to configure SAML providers that will be used for HTTP access via the classic SAP HANA extended application services server.

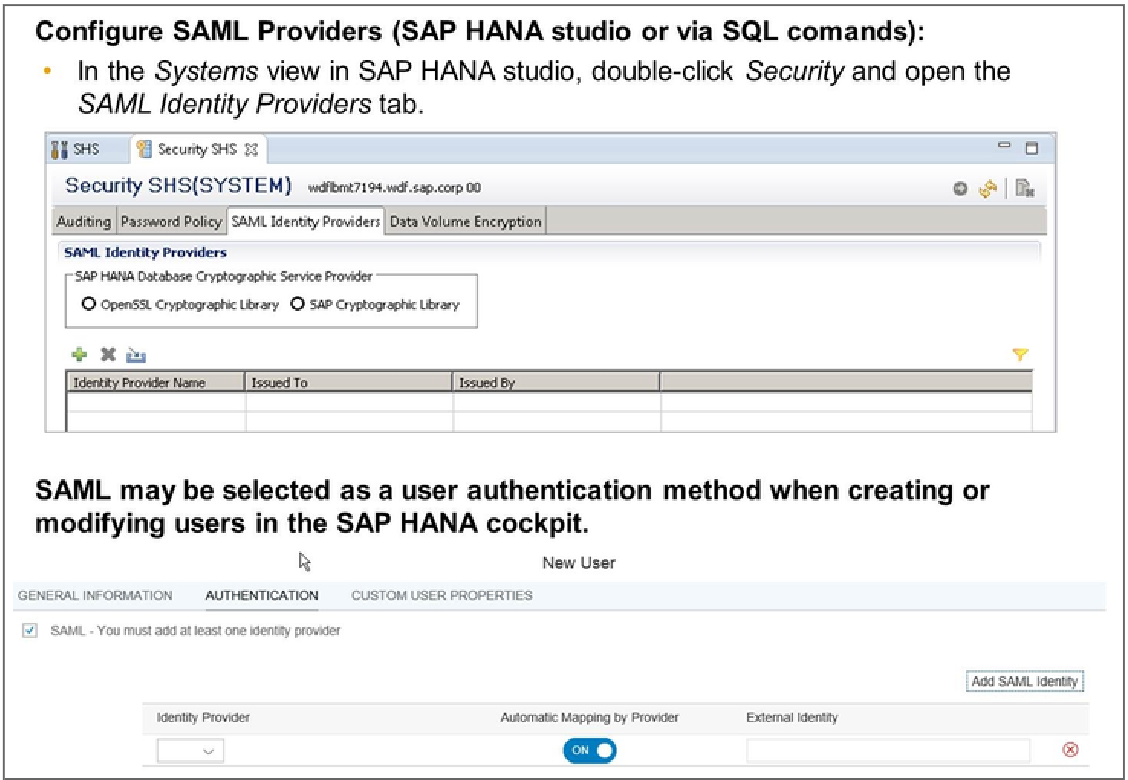

Figure 397: Configure SAML in SAP HANA

The main purpose of SAML for SAP HANA is to support scenarios where clients are not directly connected to the SAP HANA database, but to a middle-tier application server (SAP HANA extended application services engine, for example).

This middle-tier application server runs an HTTP server. Whenever the application server needs to connect to the database on behalf of the user, it requests an SAML assertion from the client.

The assertion is issued by an identity provider after the client is authenticated successfully. The assertion is then forwarded to the SAP HANA database, which grants access based on the previously established trust to the identity provider.

SAML Configuration in Administration of SAP HANA Extended Application Services

SAP HANA extended application services includes a Web-based administration tool. This tool enables you to configure several security-related aspects of applications for SAP HANA extended application services, including authentication (for example, enforced authentication mechanism, trust store configuration and management, and SAML configuration).

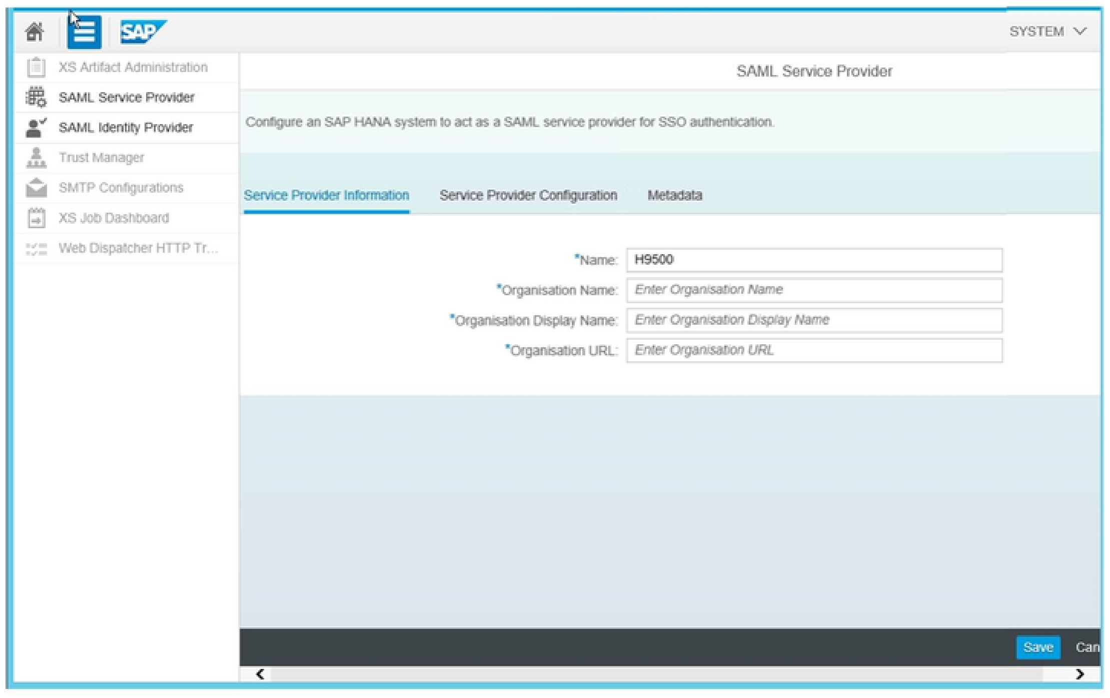

Figure 398: SAML Configuration in Administration of SAP HANA Extended Application Services

SAP Logon Ticket and SAP Assertion Ticket

You can authenticate users in SAP HANA with logon or assertion tickets. These tickets are issued to them when they log on to an SAP system configured to create tickets (for example, the SAP Web Application Server or SAP Enterprise Portal).

If you want to integrate an SAP HANA system into a landscape that uses SAP logon or assertion tickets for user authentication, configure SAP HANA to accept logon or assertion tickets.

SAP HANA validates incoming logon or assertion tickets against certificates signed by a trusted Certification Authority (CA) stored in a dedicated trust store. This trust store must contain all root certificates used to validate logon and assertion tickets.

The user named in an incoming SAP logon ticket must exist as a database user. You must also configure the database user for authentication using logon or assertion tickets. You can configure the database users in the user editor of the SAP HANA studio.

For more information about using logon tickets, see the SAP NetWeaver Library on the SAP Help Portal.

JSON Web Tokens

JSON Web Tokens (JWT) can be used for user authentication in single sign-on environments. The identity of users can be authenticated by JWT tokens issued by a trusted identity provider.

A JWT can be used to authenticate users accessing SAP HANA directly from ODBC/JDBC database clients or indirectly through SAP HANA extended application services, advanced model (SAP HANA XS, advanced).

> Note:
> A user who connects to the database using an external authentication provider must also have a database user known to the database. The external identity is mapped to the identity of an internal database user.

Lightweight Directory Access Protocol (LDAP) Group Authorization

The Lightweight Directory Access Protocol (LDAP) is an application protocol for accessing directory services. If you use an LDAP-compliant identity management server to manage users and their access to resources, you can use LDAP group membership to authorize SAP HANA users.

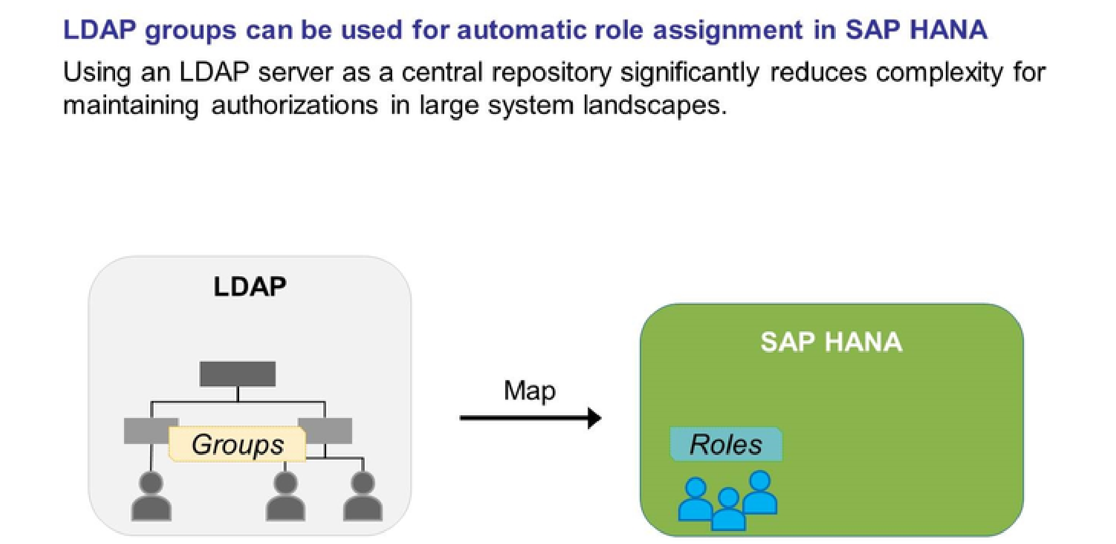

Figure 399: LDAP Group Authorization

Overview

You can map LDAP groups to SAP HANA roles. This means that if SAP HANA users are configured for LDAP group authorization, SAP HANA can determine which roles to assign them based on their membership in one or more LDAP groups. The privileges defined in the SAP HANA role determines users' access to requested resources.

How Does LDAP Group Authorization Work?

The procedure for LDAP group authorization is as follows:

1. The LDAP-enabled user logs on.
2. SAP HANA queries the LDAP directory for group memberships.
   - The logon to SAP HANA succeeds if the user’s LDAP groups map to at least one SAP HANA role.
   - The logon to SAP HANA fails if the user is not a member of any LDAP groups, or if the groups are not mapped to any SAP HANA roles.
3. SAP HANA grants the user roles according to the defined mapping.

LDAP group memberships are cached (the default is 4 hours). However, you can configure the caching duration, for example, force the LDAP group membership to be re-evaluated upon each user logon.

LDAP Group Authorization: Configuration

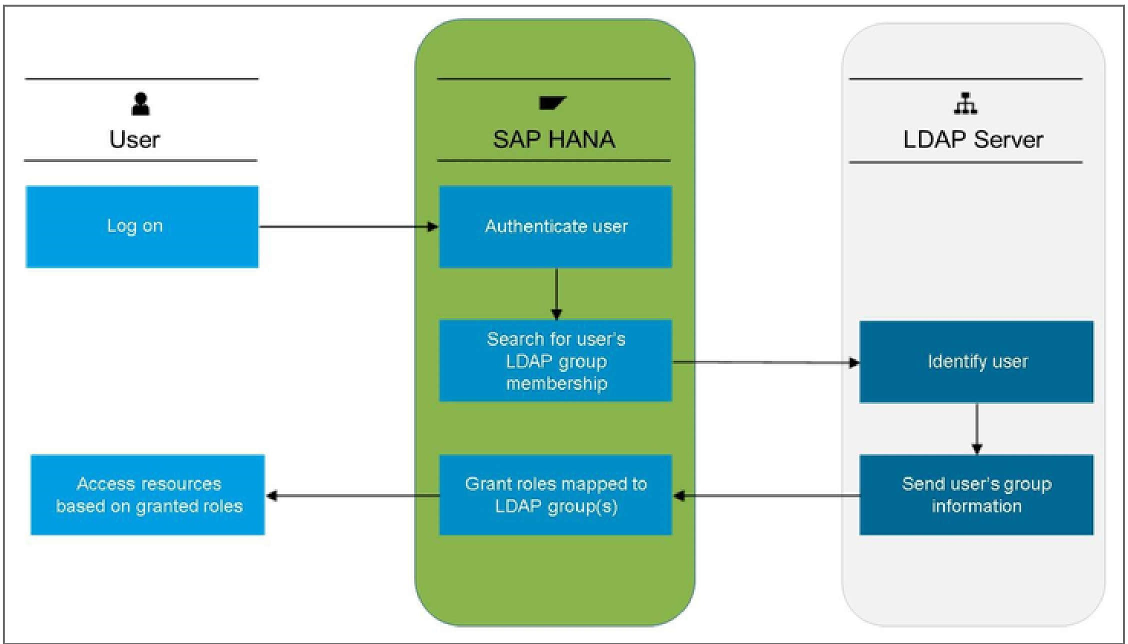

Figure 400: Authorization of User Configured for LDAP Group Authorization

Note:
You must enable LDAP group authorization explicitly for users.

LDAP Group Authorization: Configuration

The procedure for the configuration of LDAP group authorization is as follows:

1. Map LDAP groups to SAP HANA roles.
2. Configure the connection to the LDAP server.
3. Configure authorization mode LDAP for SAP HANA users.

Note:
A role that has an LDAP group mapping can also be granted to users and other roles as usual. If the role is deleted, it is also revoked from users as usual. Mappings of LDAP groups to this role are also deleted.

> Note:
> TLS or SSL-secured communication between SAP HANA and an LDAP server uses OpenSSL on the SAP HANA server side. The OpenSSL library is installed by default as part of the operating system installation.

LESSON SUMMARY

You should now be able to:

- Understand User Management
- Understand Authentication and Authorization

Unit 12
Lesson 2

Types of Privileges

LESSON OVERVIEW

In this lesson, you will learn about authorization, object privileges, SYSTEM privileges, package privileges, analytic privileges, and application privileges.

Business Example

The authorization concept is based on different types of privileges. To grant the users the right privileges, a sound understanding of the different types of privileges is necessary.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explain the different types of privileges

Privileges

When a user accesses the SAP HANA database using a client interface (for example, ODBC, JDBC, or HTTP), their ability to perform database operations on database objects is determined by the privileges that they have been granted.

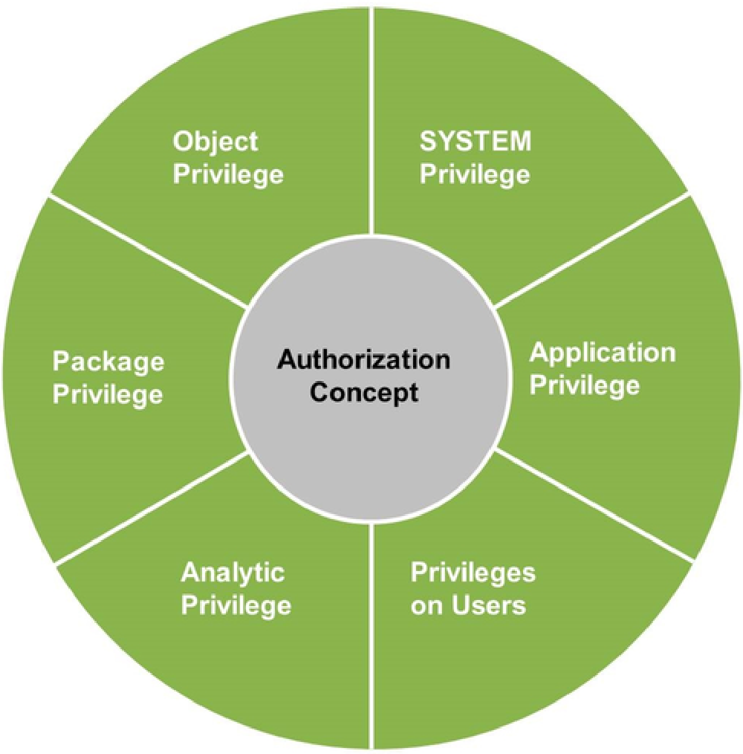

Figure 401: Authorization Management

Several privilege types are used in SAP HANA (system, object, analytic, package, and application), as follows:

System Privileges

System privileges control general system activities. They are mainly used for administrative purposes, such as creating schemas, creating and changing users and roles, performing data backups, managing licenses, and so on.

System privileges are also used to authorize basic repository operations.

Object Privileges

Object privileges allow access to and modification of database objects, such as tables and views. Depending on the object type, different actions can be authorized (for example, SELECT, CREATE ANY, ALTER, DROP, and so on).

Analytic Privileges

Analytic privileges allow read access to data in SAP HANA information models (that is, analytic views, attribute views, and calculation views) depending on certain values or combinations of values. Analytic privileges are evaluated during query processing.

Package Privileges

Package privileges authorize actions on individual packages in the classic SAP HANA repository. The privileges can grant you read access to the repository, or give you full access (read/write) to the classic SAP HANA repository.

Packages contain design time versions of various objects, such as analytic views, attribute views, calculation views, and analytic privileges.

Application Privileges

Developers of applications for the classic model of SAP HANA extended application services can create application privileges to authorize user and client access to their application. These apply in addition to other privileges, for example, object privileges on tables.

You can grant application privileges directly to users or roles in runtime in the SAP HANA studio. However, make sure that you grant application privileges to roles created in the repository in design time.

Privileges on Users

In the SAP HANA studio, an additional privilege type can be granted. Privileges on users are SQL privileges that users can grant on their user.

SQL Privilege

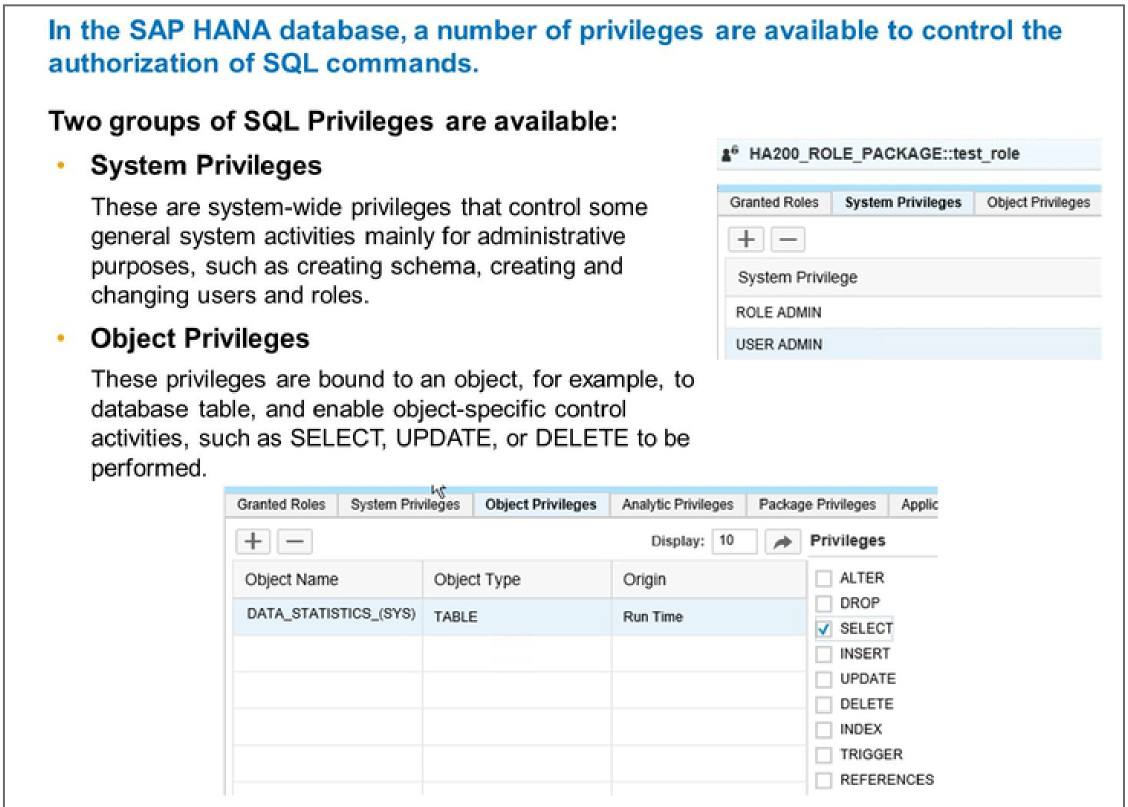

Figure 402: SQL Privilege

Object Privileges

Object privileges are SQL privileges that allow access to and modification of database objects.

For each SQL statement type (for example, SELECT, UPDATE, or CALL), a corresponding object privilege exists. If a user wants to execute a particular statement on a simple database object (for example, a table), they must have the corresponding object privilege for either the actual object itself, or the schema in which the object is located. This is because the schema is an object type that contains other objects. A user who has object privileges for a schema automatically has the same privileges for all objects currently in the schema and for any objects created there in the future.

You can grant object privileges for database catalog objects such as tables, views, and procedures. You can also grant object privileges for noncatalog objects such as development objects in the repository of the SAP HANA database.

Object Privilege Activities

Object privilege activities also include the following:

- CREATE ANY

  This privilege allows the creation of all kinds of objects, in particular, tables, views, sequences, synonyms, SQL script functions, or database procedures in a schema. You can only grant this privilege on a schema.

- ALL PRIVILEGES

This is a collection of all Data Definition Language (DDL) and Data Manipulation Language (DML) privileges that the grantor currently has and is allowed to grant, and that can be granted on this particular object. This collection is evaluated dynamically for the given grantor and object. ALL PRIVILEGES is not applicable to a schema, but only to a table, view, or table type.

- DROP and ALTER

  These are DDL privileges and they authorize the DROP and ALTER SQL commands. While the DROP privilege is valid for various objects, the ALTER privilege is not valid for sequences and synonyms. This is because their definitions cannot be changed after creation.

- SELECT, INSERT, UPDATE, and DELETE

  These are DML privileges and they authorize respective SQL commands. SELECT is valid for all kinds of objects, except for functions and procedures. However, INSERT, UPDATE, and DELETE are only valid for schemas, tables, table types, and table views.

- INDEX

  This special DDL privilege authorizes the creation, alteration, or revocation of indexes for an object using the CREATE INDEX, ALTER INDEX, and DROP INDEX commands. This privilege can only be applied to a schema, table, and table type.

- EXECUTE

  This special DML privilege authorizes the execution of an SQL script function or a database procedure using the CALLS or CALL command, respectively.

> Note:
> For more information about the object privileges available in SAP HANA and for which objects they are relevant, see Object Privileges (Reference) in the SAP HANA Security Guide.

System Privilege

System privileges control general system activities.

System privileges are mainly used to authorize users to perform administrative actions, including the following:

- Creating and deleting schemas
- Managing users and roles
- Performing data backups
- Monitoring and tracing
- Managing licenses

System privileges are also used to authorize basic repository operations. This includes the following examples:

- Importing and exporting content
- Maintaining Delivery Units (DUs)

Types of System Privilege

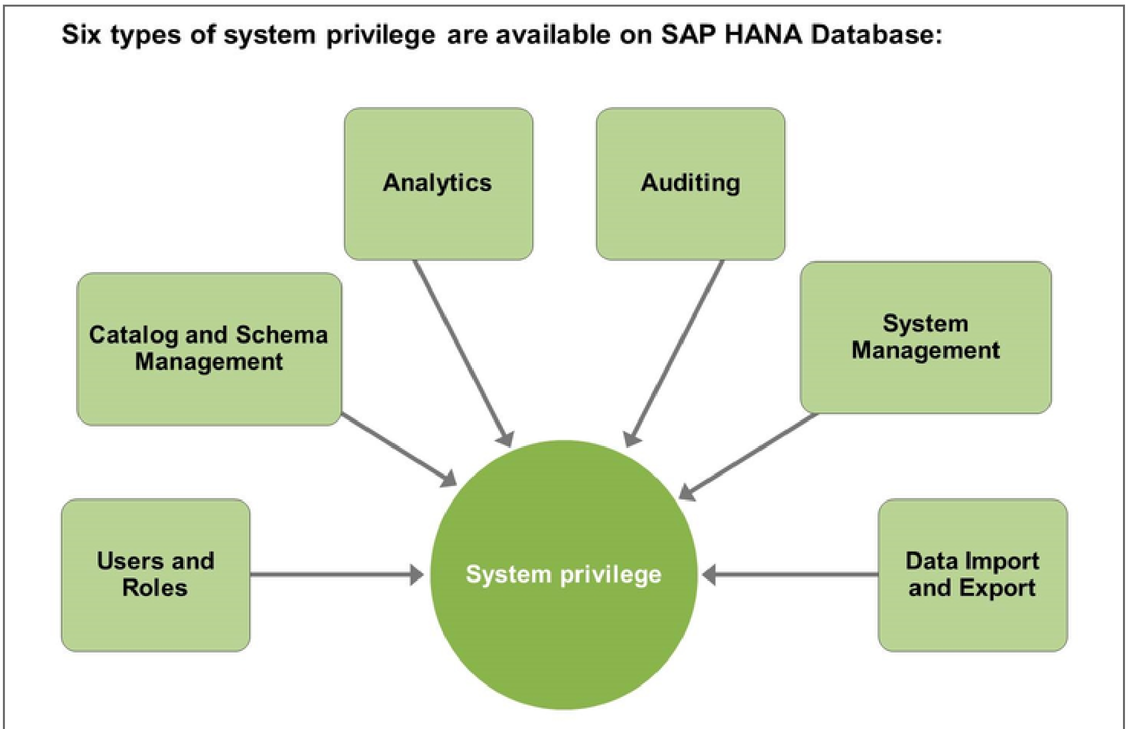

Figure 403: System Privilege

As shown in the figure, System Privilege, the following system privileges are available on SAP HANA Database:

- Users and Roles, which include the following:
  - USER ADMIN

    This privilege authorizes the creation and changing of users with the CREATE USER, ALTER USER, and DROP USER SQL commands.
  - ROLE ADMIN

    This privilege authorizes the creation and deletion of roles with the CREATE ROLE and DROP ROLE SQL commands. It also authorizes the granting and cancellation of roles with the GRANT and REVOKE SQL commands.
- Catalog and Schema Management, which include the following:
  - CREATE SCHEMA

    This privilege authorizes the creation of database schemas with the CREATE SCHEMA SQL command.
  - DATA ADMIN

    This privilege authorizes all users to have unfiltered read-only access to the full content of all system and monitoring views. It also authorizes users to execute all DDL – and only DDL – commands in the SAP HANA database. The content of those views is filtered based on the privileges of the user.

  - CATALOG READ

    This privilege authorizes all users to have unfiltered read-only access to the full content of all system and monitoring views. The content of those views is filtered based on the privileges of the accessing user.

- System Management

  These privileges authorize the various system activities that can be performed using the ALTER SYSTEM SQL commands. Because of the high level of impact on the system, these privileges are not designed for a normal database user. Proceed with caution when granting these privileges (for example, only grant them to a support user or role.)

- Data Import and Export

  The following system privileges are available for the authorization of the data import and export in the database:

  - IMPORT

    This privilege authorizes the import activity in the database using the IMPORT or LOAD TABLE SQL commands. Note that, besides this privilege, the user needs the INSERT privilege on the target tables to be imported.

  - EXPORT

    This privilege authorizes the export activity in the database via the EXPORT or LOAD TABLE SQL commands. Note that, besides this privilege, the user needs the SELECT privilege on the source tables to be exported.

> Note:
> For more information about the individual system privileges, see System Privileges (Reference) in the SAP HANA Security Guide.

Package Privilege

Package privileges authorize actions on individual packages in the classic SAP HANA repository.

> Note:
> With SAP HANA extended application services, advanced model, source code, and web content are not versioned and stored in the SAP HANA database. Therefore, package privileges are not used in this context.

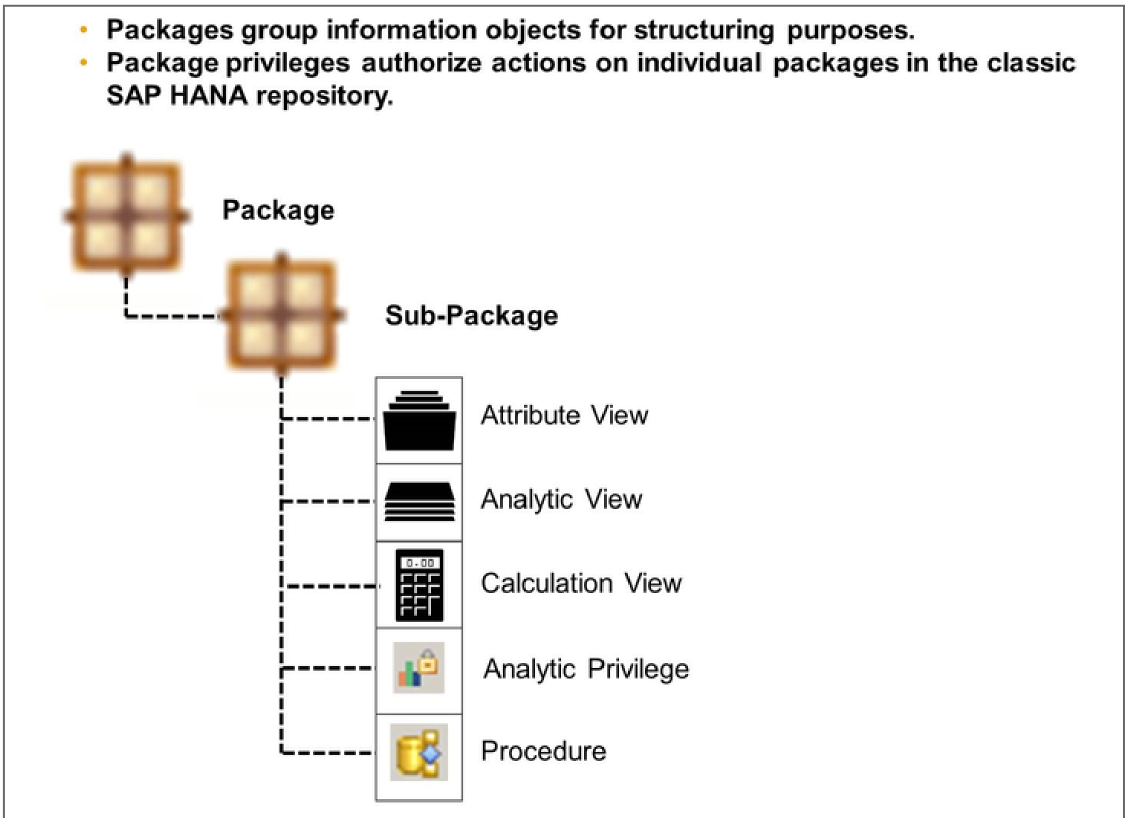

Figure 404: Package Privilege: Concept

The SAP HANA database repository is structured hierarchically with packages assigned to other packages as subpackages.

Packages contain design-time versions of various objects, such as Analytic, Attribute, and Calculation Views, as well as Analytic Privileges, and functions. To work with packages, the respective package privileges must be granted.

Privileges granted on a repository package are assigned implicitly to the design-time objects in the package, as well as to all subpackages. Users are only allowed to maintain objects in a repository package if they have the necessary privileges for the package in which they want to perform an operation, for example to read or write to an object in that package. To perform operations in all packages, a user must have privileges on the root package REPO_PACKAGE_ROOT.

For a requested operation in a specific package, the authorization check is repeated on the parent package and recursively up the package hierarchy to the root level of the repository. If the user does not have the necessary privileges for any of the packages in the hierarchy chain, the authorization check fails, and the user is not permitted to perform the requested operation.

In the context of repository package authorizations, there is a distinction between native packages and imported packages.

If you grant privileges to a user for a package, the user is also authorized automatically for all corresponding subpackages.

Analytic Privilege

Analytic privileges grant different users access to different portions of data in the same view based on their business role. Within the definition of an analytic privilege, the conditions that control which data users see is either contained in an XML document or is defined using SQL.

Standard object privileges (SELECT, ALTER, DROP, and so on) implement coarse-grained authorization at object level only. Users either have access to an object, such as a table, view or procedure, or they don't. While this is often sufficient, there are cases when access to data in an object depends on certain values or combinations of values. Analytic privileges in the SAP HANA database provide fine-grained control at row level of the type of data that individual users can see within the same view.

Analytic Privilege View

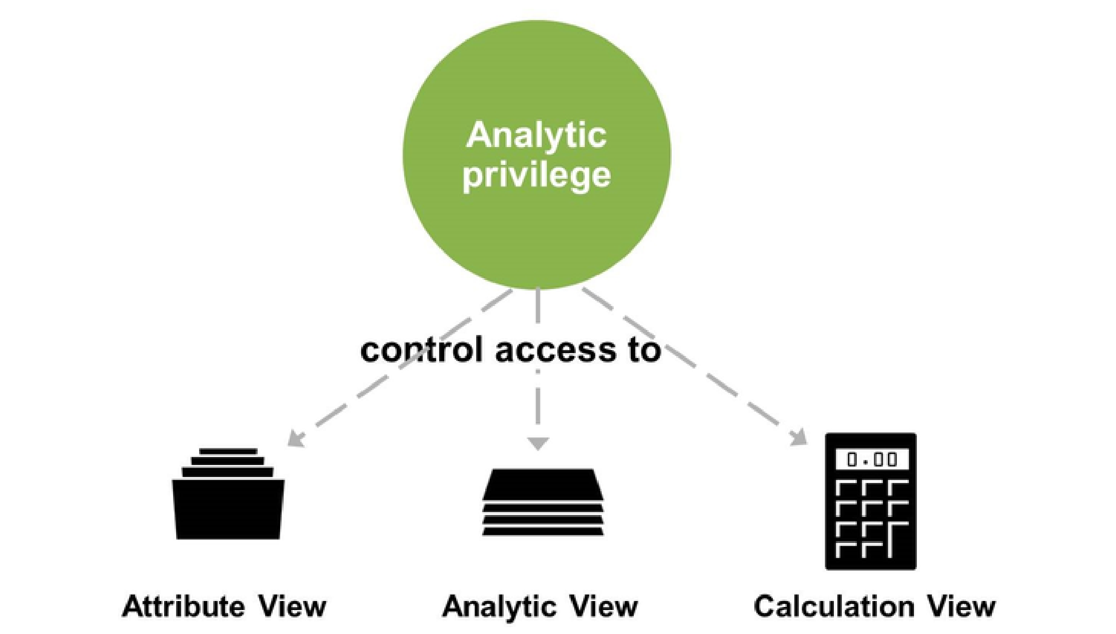

Figure 405: Analytic Privileges

All Attribute Views, Analytic Views, and Calculation Views, which have been designed in the modeler and have been activated from the modeler of the SAP HANA studio, are supported automatically by the Analytic Privilege mechanism.

If you are already familiar with the authorization model of SAP Business Warehouse (SAP BW), you will see that there are many similarities between the two models.

The overall idea behind analytic privileges is the reuse of Analytic Views by different users. However, the different users may not be allowed to see the same data. For example, different regional sales managers, who are only allowed to see sales data for their regions, can reuse the same Analytic View. They obtain the analytic privilege to see only data for their region, and their queries on the same view return the corresponding data. This is a major difference compared with the SAP BW model. While the concept is very similar, if you execute a query that returns values that you are not authorized to see, SAP BW forwards an error message. With the SAP HANA database, the query is executed and, corresponding to your authorization, only values that you are entitled to see are returned.

Manage Analytic Privileges

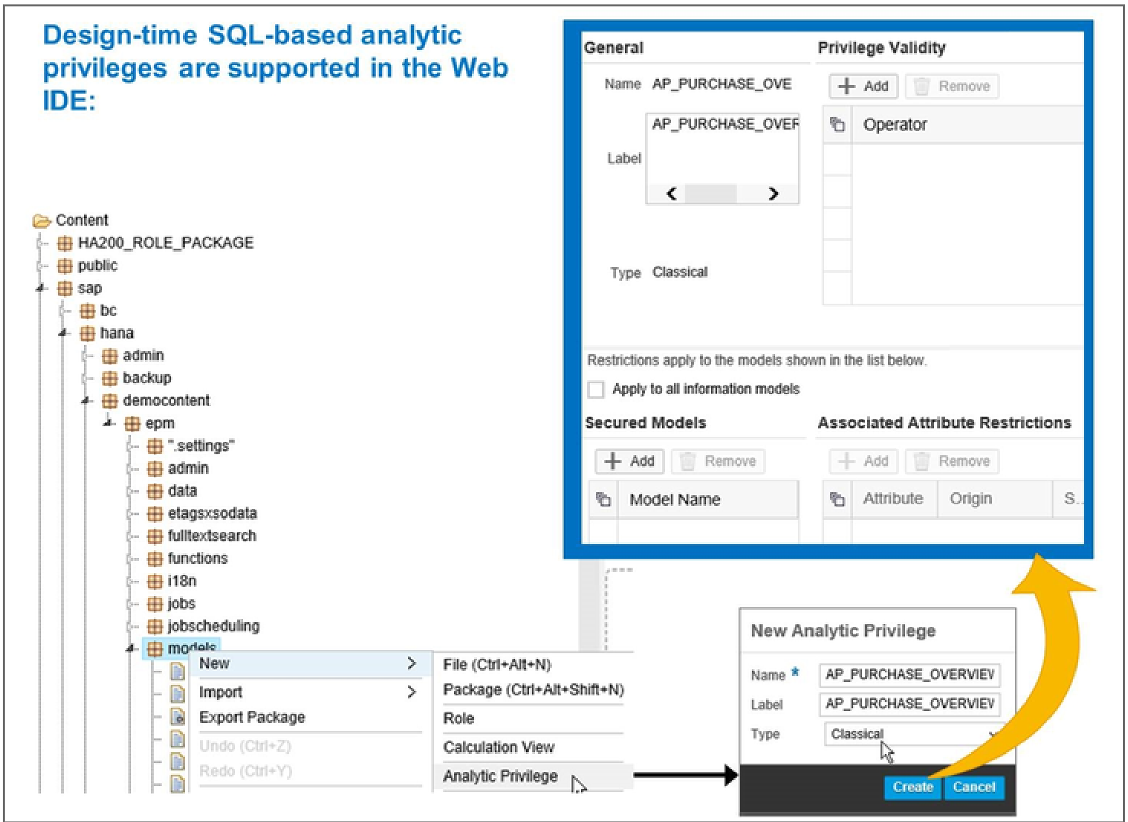

Figure 406: Manage Analytic Privileges

The user has access to an individual, independent view (Attribute, Analytic, or Calculation View), if the following prerequisites are met:

- The user was granted the SELECT privilege on the view or the containing schema.
- The user was granted an analytic privilege that is applicable to the view. An analytic privilege is applicable to a view if it contains the view in the Cube restriction and it contains at least one filter on one attribute of this view.

No SELECT privilege on the underlying base tables or views of this view is required.

You must implement row-level security with analytic privileges.

Restrict access to a given data container to the following selected attribute values:

- Field from Attribute View
- Field from Attribute View used in Analytic View
- Private dimension of Analytic View
- Attribute field in Calculation View
- Combinations of the previous values
- Single value, range, IN-list

Application Privilege

Developers of applications for SAP HANA extended application services can create application privileges to authorize user and client access to their application.

These privileges authorize user and client access to the application. For example, they authorize users and clients to start the application or to perform administrative actions in the application.

You can grant and revoke application privileges through the GRANT_APPLICATION_PRIVILEGE procedures and through the REVOKE_APPLICATION_PRIVILEGE procedure in the _SYS_REPO schema.

You can grant application privileges to users or roles in runtime in the SAP HANA studio. However, we recommend that you grant application privileges to roles created in the repository.

You can grant and revoke application privileges in the SAP HANA studio.

Privileges on Users

Privileges on users are SQL privileges that users can grant to other users. ATTACH DEBUGGER is the only privilege that can be granted on a user.

For example, User A can grant User B the privilege ATTACH DEBUGGER. This allows User B to debug SQLScript code in the session of User A. User A is the only user who can grant this privilege.

> Note:
> User B also needs the object privilege DEBUG on the relevant SQLScript procedure.
>
> You cannot grant the ATTACH DEBUGGER privilege on behalf of other users.

LESSON SUMMARY

You should now be able to:

- Explain the different types of privileges

Unit 12
Lesson 3

SAP HANA Roles

LESSON OVERVIEW

This module covers the following topics:

- Predelivered role
- Template role
- Support role

Business Example

For special tasks, standard roles are delivered. You need to know in which cases you can use these roles.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Understand roles

Overview of Roles

A role is a collection of privileges that can be granted to either a user or to another role in runtime.

A role typically contains the privileges required for a particular function or task, such as the following:

- Business end users reading reports using client tools
- Modelers creating models and reports in the modeler of the SAP HANA studio
- Database administrators operating and maintaining the database and the users

You can grant privileges directly to users of the SAP HANA database. However, roles are the standard mechanism of granting privileges. This is because they allow you to implement complex, reusable authorization concepts that can be modeled on business roles. Several standard roles are delivered with the SAP HANA database (for example, MODELING, MONITORING). You can use these as templates for creating your own roles.

A role can contain any number of the following privileges:

- System privileges for administrative tasks (for example, AUDIT ADMIN, BACKUP ADMIN, CATALOG READ)
- Object privileges on database objects (for example, SELECT, INSERT, UPDATE)
- Analytic privileges on SAP HANA information models
- Package privileges on repository packages (for example, REPO.READ, REPO.EDIT_NATIVE_OBJECTS, REPO.ACTIVATE_NATIVE_OBJECTS)

- Application privileges for enabling access to applications for SAP HANA extended application services

A role can also contain other roles.

Types of Roles

Roles in the SAP HANA database can exist as runtime objects only (catalog roles), or as design-time objects (repository roles) that become catalog objects on deployment (database artifact with the file suffix .hdbrole).

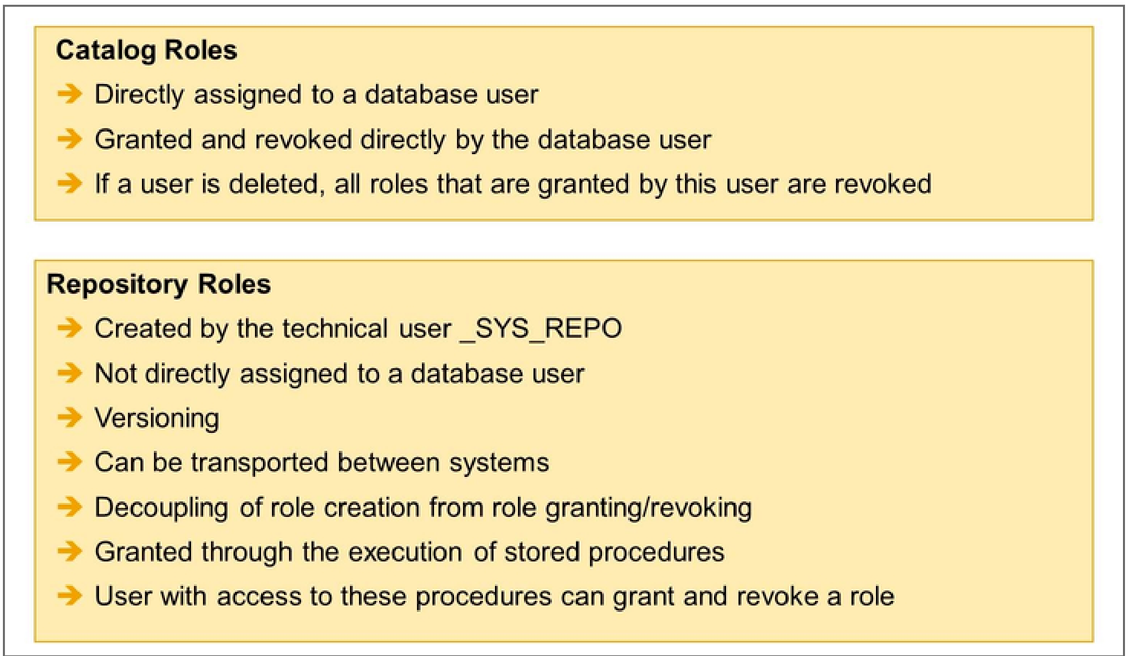

Figure 407: TypesOfRoles_Image.ppt

Model roles as design-time objects for the following reasons:

- First, unlike roles created in runtime, roles created as design-time objects can be transported between systems.

  This is important for application development because it means that developers can model roles as part of their application's security concept and then ship these roles or role templates with the application.

  Being able to transport roles is also advantageous for modelers implementing complex access control on analytic content. They can model roles in a test system and then transport them into a productive system. This avoids unnecessary duplication of effort.

- Second, catalog roles are created as design-time objects are not directly associated with a database user.

  Roles are created by the technical user _SYS_REPO and are granted through the execution of stored procedures. Any user with access to these procedures can grant and revoke a role. Roles created in runtime are granted directly by the database user and can only be revoked by the same user. Additionally, if the database user is deleted, all roles that they granted are revoked. Because database users correspond to real people, this can impact the implementation of your authorization concept, for example, if an employee leaves the organization or is on vacation.

Note:

- **Catalog roles** are created in runtime and have typical SQL object behavior (for example, ownership). They are mainly used in conjunction with high-level user/identity management tools where roles are managed by a technical user in SAP HANA.
- **Repository roles** are recommended to use for other scenarios where roles are managed by users directly in SAP HANA.

Hint:
Make sure that the design-time version of a role in the repository and its activated runtime version always contain the same privileges. In particular, do not grant additional privileges to the activated runtime version of a role created in the repository. Although there is no method for preventing a user from doing this, the next time the role is activated in the repository, any changes made to the role in runtime are reverted. Therefore, do not change the activated runtime version of a role in runtime.

You can create database roles in the built-in repository of the SAP HANA database using either the SAP Web IDE or the SAP HANA studio.

Catalog Roles

You can create catalog roles in the SAP HANA system. A role administrator creates the role in the runtime of the SAP HANA system. The database user grants catalog roles directly, and they can only be revoked by the same user.

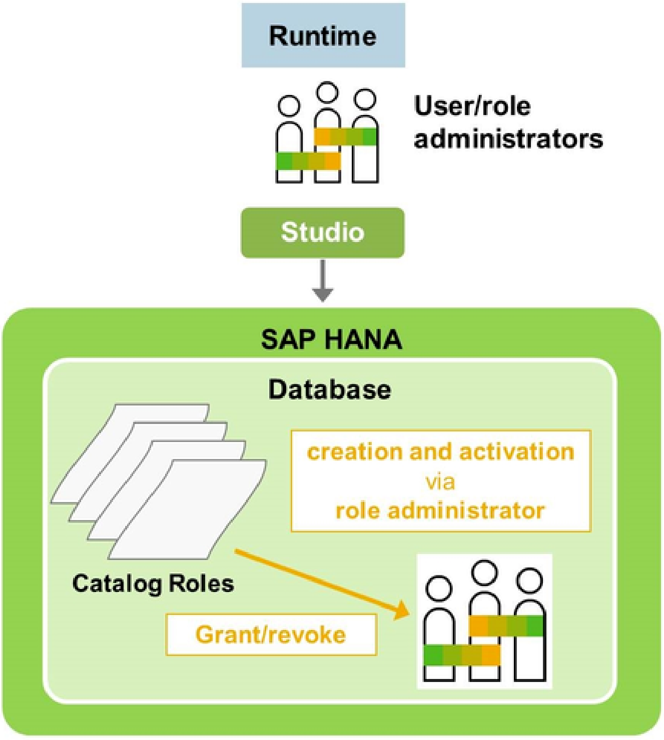

Figure 408: RuntimeRoles_Image.ppt

Create Catalog Roles

To create a catalog role, open the catalog role editor in SAP HANA cockpit. Then select the roles and privileges that you want to include, and save the role.

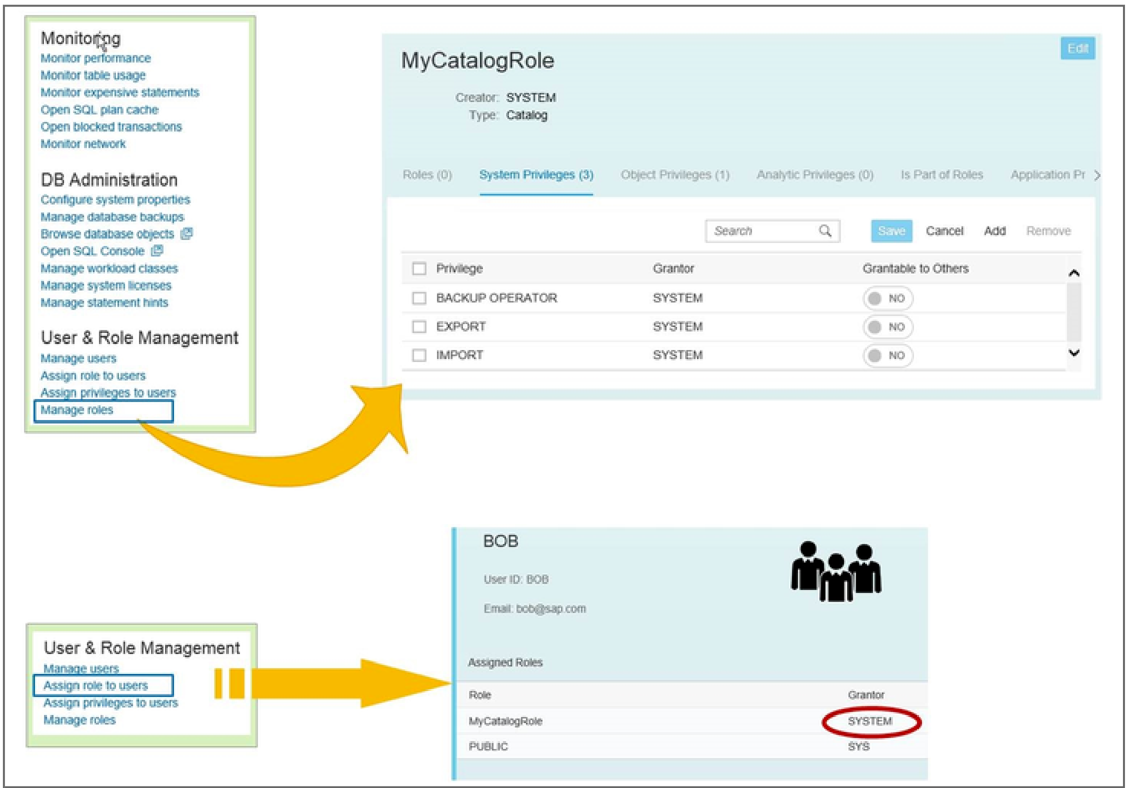

Figure 409: Create Catalog Roles

Repository Roles

You can create repository roles in the development system. A developer or role designer creates the role in the repository of the development system and tests it. Therefore, the following prerequisites have to be fulfilled:

- Authorization assigned: sap.hana.xs.ide.roles::EditorDeveloper role
- A shared project must exist with a suitable package for storing roles
- Package privileges on the required packages

The role is transported to the production system, for example, using HALM or CTS+.

In the production system, a user administrator grants the role to end users.

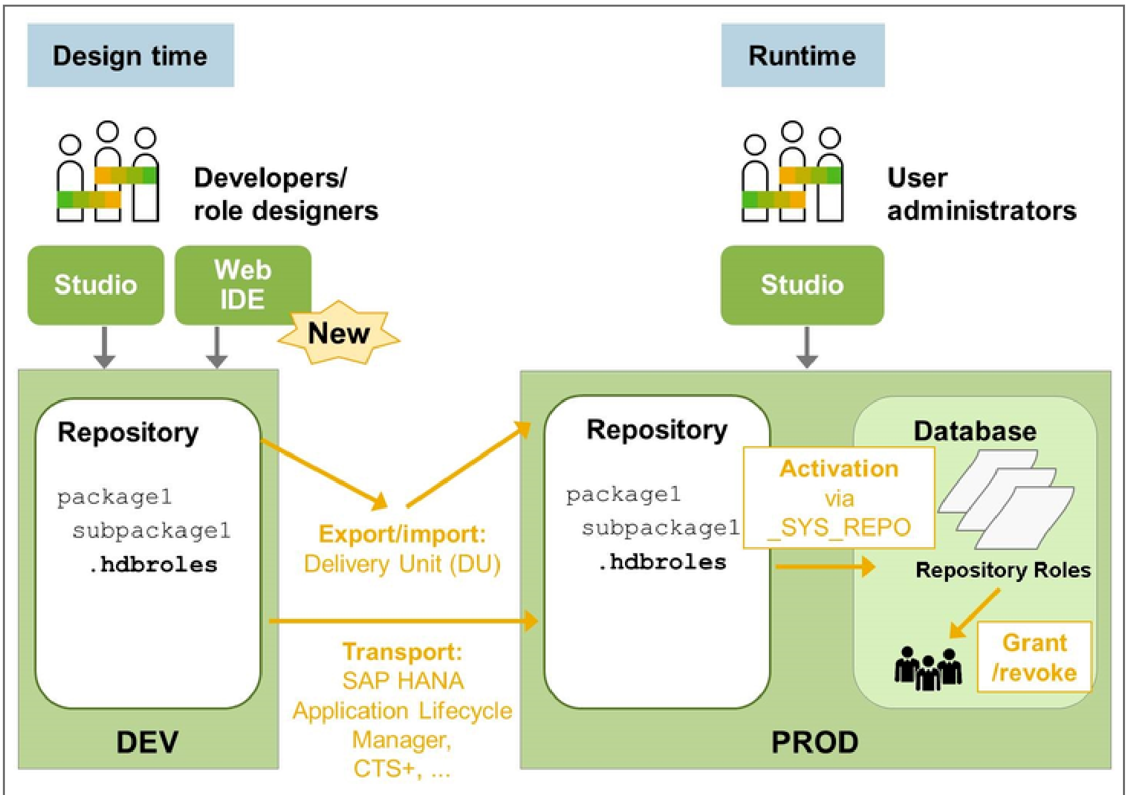

Figure 410: Repository Roles

Create Repository Roles

To create a repository role, open the Editor of the SAP Web IDE in your Web browser by entering the following URL:

https://<database_server>:43<instance_no>/sap/hana/xs/ide/editor

Create the new role in the Content tree by right clicking the folder where you want to create the role. Then select the roles and privileges that you want to include, and save the role.

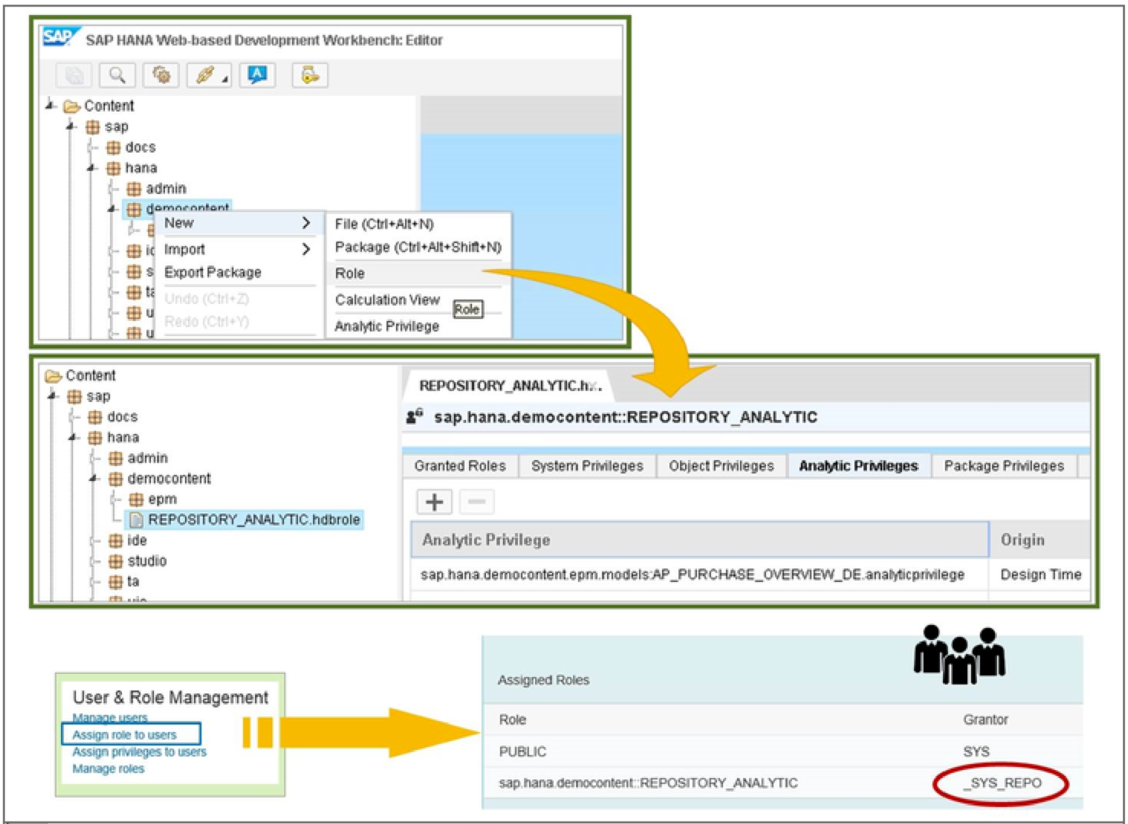

Figure 411: Create Repository Roles

Note:
The role is saved and activated in one step. If you only want to save the role, choose Settings and select Enable inactive save . An additional icon displays in the toolbar Save without Activating .

Predelivered Roles

Several roles are delivered with the SAP HANA database. You can use these as templates for creating your own roles.

Predelivered Catalog Roles

Several catalog roles are available by default in the SAP HANA database.

Several predefined catalog roles are delivered with the SAP HANA database. Do not use these roles directly, but instead use them as templates for creating your own roles.

Note:
These roles do not exist in the SAP HANA repository.

PUBLIC role

PUBLIC

- Privileges for a user to work with the database at all
- Contains privileges for filtered read-only access to the system views
- Only objects for which the users have access rights are visible
- Is implictly granted to every user, except restricted users
- **This role schould (must) be used unchanged**

Figure 412: Predelivered Catalog Role for Standard SQL Users

PUBLIC

This role contains privileges for filtered read-only access to the system views. Only objects for which the users have access rights are visible. By default, this role is granted to every user, except restricted users.

SAP HANA Template Roles

SAP HANA comes with several template roles:

MONITORING

- Role with full read-only access to all metadata, monitoring, and statistics

MODELING

- A very richly privileged role that enables
  - Creation and activation of Information Models
  - Creation and activation of Analytic Privileges

CONTENT_ADMIN

- Contains the same privileges as the MODELING role
- The only role in the system with vital privileges, for example:
  - SQL Privileges on Schema _SYS_BIC
  - SQL Privileges on Schema _SYS_BI
- Allows to grant these privileges to other users

Figure 413: Predelivered Catalog Roles

Hint:
Regard these roles as templates. Do not grant these roles, build your own roles instead, as follows.

MONITORING

This contains privileges for full read-only access to all metadata, the current system status in system and monitoring views, and the data of the statistics server.

MODELING

This role contains all privileges required for using the information modeler in the SAP HANA studio.

It contains the database authorization for a modeler to create various views and analytic privileges.

It allows access to all data in activated views without any filter (_SYS_BI_CP_ALL Analytic Privilege). However, this is restricted by missing SQL privileges on those activated objects.

Use this predefined role as a template.

CONTENT_ADMIN

This role contains all the privileges required for using the information modeler in the SAP HANA studio, as well as the additional authorization to grant these privileges to other users. It also contains system privileges for working with imported objects in the SAP HANA repository. Use this role as a template for creating roles for content administrators.

> Caution:
> These predefined roles are privileged and should not be granted to users, particularly in production systems. Only use these roles as a template.

> Note:
> For more information about predefined database roles in SAP HANA, see Predefined Database Roles in the SAP HANA Security Guide.

Support Role

This role contains system privileges (for example, CATALOG READ) and object privileges (for example, SELECT on SYS schema) that allow access to certain low-level internal system views needed by SAP HANA development support in support situations. All access is read only. This role does not allow access to any customer data.

The definition of the low-level internal system views to which this role allows access is not part of the stable end-user interface. It might change from revision to revision. To avoid administrators and end users accidentally accessing these internal system views in applications or scripts, this role has several usage restrictions. Only grant this role to SAP HANA development support users for their support activities.

SAP HANA comes with a preinstalled role for support cases:

SAP_INTERNAL_HANA_SUPPORT

Contains privileges for full read-only access to

- All metadata
- The current system status
- Data of the statistics server
- Contains the privileges for accessing low-level internal system views (this information is otherwise only available to the SYSTEM user)

For security reasons, the following restrictions apply:

- It cannot be granted to user SYSTEM
- It can be granted to a configurable number of users (parameter: internal_support_user_limit)
- It cannot be granted to another role
- No role can be granted to it
- Only system privileges can be granted to this role

Figure 414: Support Role

Roles Restrictions

This role contains privileges for read-only access to all metadata, the current system status, and the data of the statistics server. Additionally, it contains the privileges for accessing low-level internal system views. Without the SAP_INTERNAL_HANA_SUPPORT role, this information can be selected only by the SYSTEM user.

To avoid accidental use of this role in day-to-day activities, the following restrictions apply to the SAP_INTERNAL_HANA_SUPPORT role:

- It cannot be granted to the SYSTEM user.
- It can only be granted to a limited number of users at the same time.

  The maximum number of users to which the role can be granted can be configured with the parameter **`internal_support_user_limit`** in the authorization section of the **`indexserver.ini`** configuration file. The default value is 1.
- It cannot be granted to another role.
- It cannot be granted another role.
- It cannot be granted further object privileges.
- It can be granted only further system privileges.
- With every upgrade of the SAP HANA database, it is reset to its default privileges.

> Hint:
> An alert notifies administrators when a user is granted the SAP_HANA_INTERNAL_SUPPORT role (see SAP Note 1991615).

Predefined Repository Roles

SAP HANA is delivered with SAP HANA content, a set of preinstalled software components implemented as SAP HANA Web applications, libraries, and configuration data. The privileges required to use a software component delivered as SAP HANA content are contained within repository roles delivered with the component itself.

For more information about the repository roles delivered with SAP HANA content, see Components Delivered as SAP HANA Content in the SAP HANA Security Guide..

> Note:
> No user has any predefined repository roles initially, except the user _SYS_REPO (as the owner of all repository content).

LESSON SUMMARY

You should now be able to:

- Understand roles

Unit 12
Lesson 4

Administrative Tasks

LESSON OVERVIEW

This lesson explains how to create and copy users, deactivate and reactivate users, manage connection attempts, set initial passwords, and manage the password policy.

Business Example

The user administrations includes tasks to deactivate and reactivate users and to manage the password policy. To enhance the logon security of your SAP HANA database, you can configure password rules by using specific parameters.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Perform Administrative Tasks

User Management

Users can be automatically deactivated for security reasons, for example, if they violate password policy rules. However, as a user administrator, you might need to deactivate a user, for example, if an employee temporarily leaves the company or if a security violation is detected.

After you deactivate the user account, the user cannot log on to the SAP HANA database until the administrator resets the user’s password or activates the user again.

As a user administrator, you might need to reactivate a user, for example, after you explicitly deactivated the user or when the user makes too many invalid log-on attempts.

You can deactivate or activate a user as follows:

- On the SAP HANA cockpit User page
- With the SAP HANA studio User editor
- With SQL command

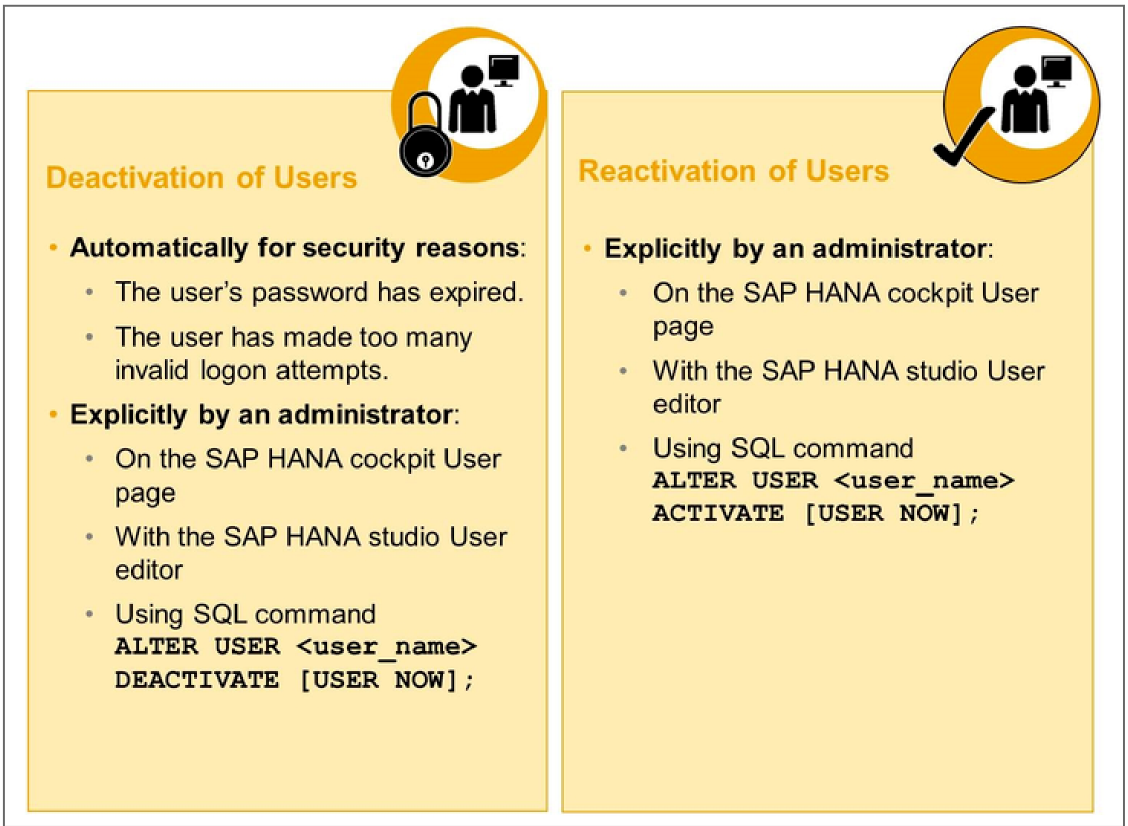

Figure 415: Deactivation and Reactivation of Users

Password Policy Management

If the user's password has expired, the user must change the password to a new value.

Passwords for the user name password authentication of database users are subject to certain rules. These are defined in the password policy and the password blacklist. You can change the default password policy according to your organization's security requirements. You cannot deactivate the password policy.

Password Policy

The password policy has the following conditions:

- Password quality (length, complexity)
- Blacklist of forbidden terms

  The password quality is defined by several parameters, which are described in detail in the SAP HANA Security Guide.

  The password blacklist is a list of words that are not allowed as passwords or parts of passwords. SAP HANA checks this blacklist whenever a password is created or changed. You can specify if the words in the blacklist are case-sensitive and if the check applies to whole words or parts of words.

  The password blacklist in SAP HANA is implemented with the table _SYS_PASSWORD_BLACKLIST (_SYS_SECURITY). This table is empty when you create a new instance. The _SYS_SECURITY schema and the _SYS_PASSWORD_BLACKLIST table are owned by the SYSTEM user.

During the initial system setup, the SYSTEM user grants change privileges for this table to a dedicated administrator user.

> Caution:
> For security reasons, and to prevent users from viewing sensitive information such as a password, manage the privilege to select carefully.

Password Policy Parameters

The password policy is defined by parameters in the password policy section of the `indexserver.ini` system properties file.

> Note:
> The password policy parameters for the system database of a multiple-container system are maintained in the `namesever.ini file`, not the `indexserver.ini` file.

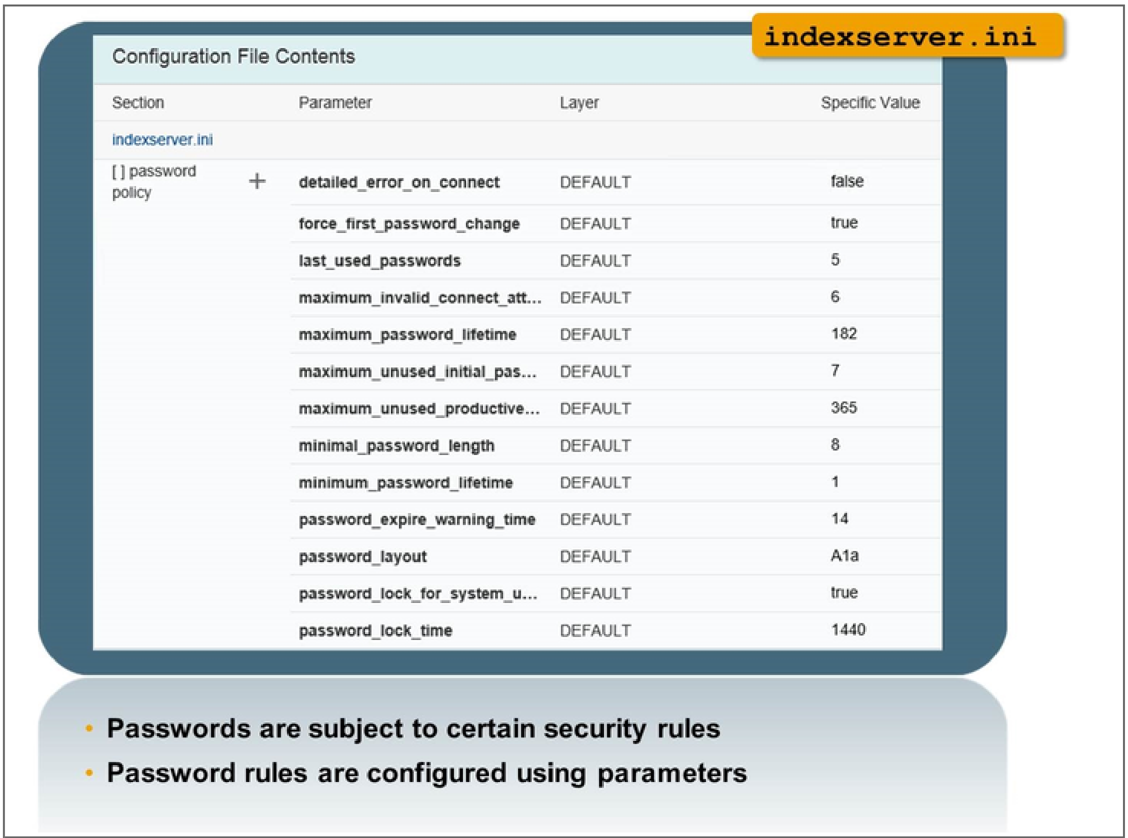

Section		Parameter	Layer	Specific Value
indexserver.ini				
[ ] password policy	+	detailed_error_on_connect	DEFAULT	false
		force_first_password_change	DEFAULT	true
		last_used_passwords	DEFAULT	5
		maximum_invalid_connect_att...	DEFAULT	6
		maximum_password_lifetime	DEFAULT	182
		maximum_unused_initial_pas...	DEFAULT	7
		maximum_unused_productive...	DEFAULT	365
		minimal_password_length	DEFAULT	8
		minimum_password_lifetime	DEFAULT	1
		password_expire_warning_time	DEFAULT	14
		password_layout	DEFAULT	A1a
		password_lock_for_system_u...	DEFAULT	true
		password_lock_time	DEFAULT	1440

Figure 416: Password Policy Parameters

Note:
The actual parameters are contained in the password policy section of the `indexserver.ini` system properties file. You can configure the password policy using the Password Policy and Blacklist page in the SAP HANA cockpit and the SAP HANA studio Security editor. You can also do this by editing the `indexserver.ini` directly.

If a parameter is set to a value outside the value range, either the minimum value or the maximum value of the value range, whichever is appropriate, is used instead.

Configure Password Policy

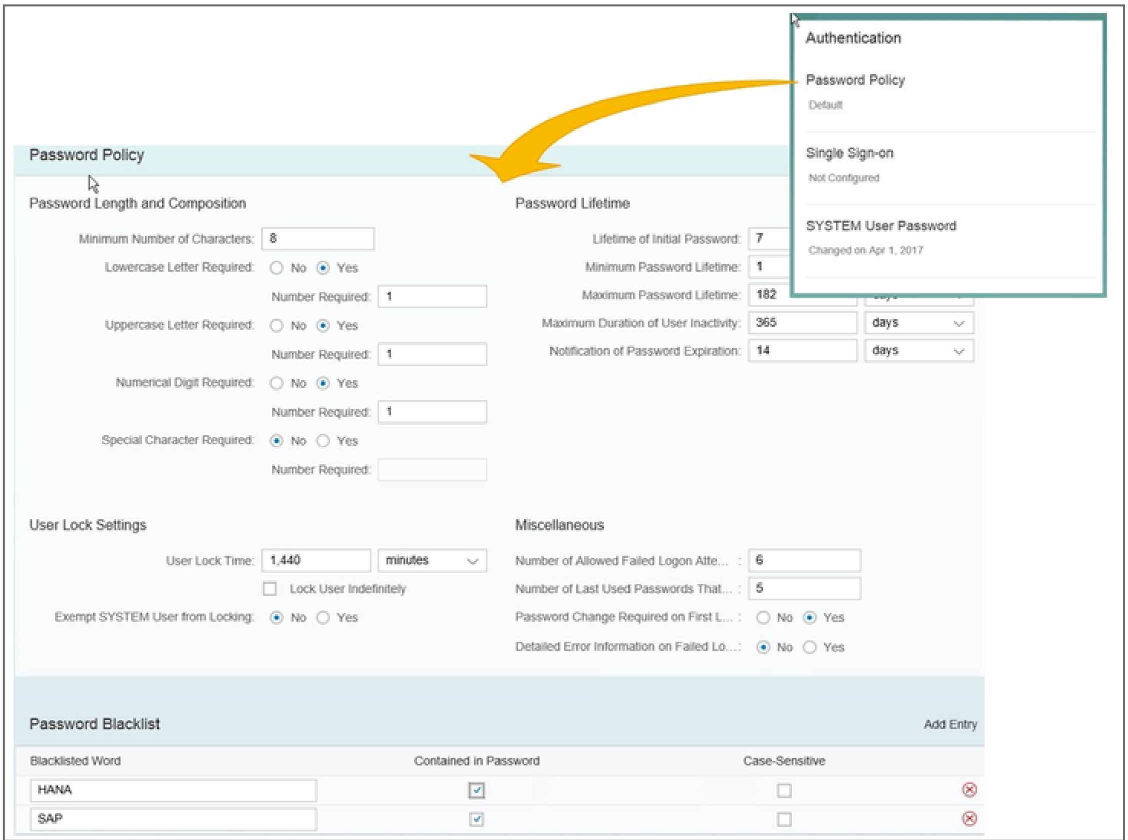

Figure 417: Configure Password Policy

Note:
The User Lock Settings (parameter `password_lock_time`) define the duration for which a user is locked after the maximum number of failed logon attempts. If you select the Lock indefinitely checkbox, the user is locked indefinitely. This corresponds to a parameter value of -1. The value 0 unlocks the user immediately.

The prerequisites for changing parameters are as follows:

- System privilege INIFILE ADMIN

- (For blacklist) INSERT and DELETE privileges for either the _SYS_PASSWORD_BLACKLIST table or the _SYS_SECURITY schema

To view the contents of the INI file, use the M_INIFILE_CONTENTS view.

The password policy parameters can be found in the M_PASSWORD_POLICY view. For more information about the parameter values, see the password policy parameters in the SAP HANA Security Guide.

Management of Connection Attempts

This section outlines a subset of password policy parameters.

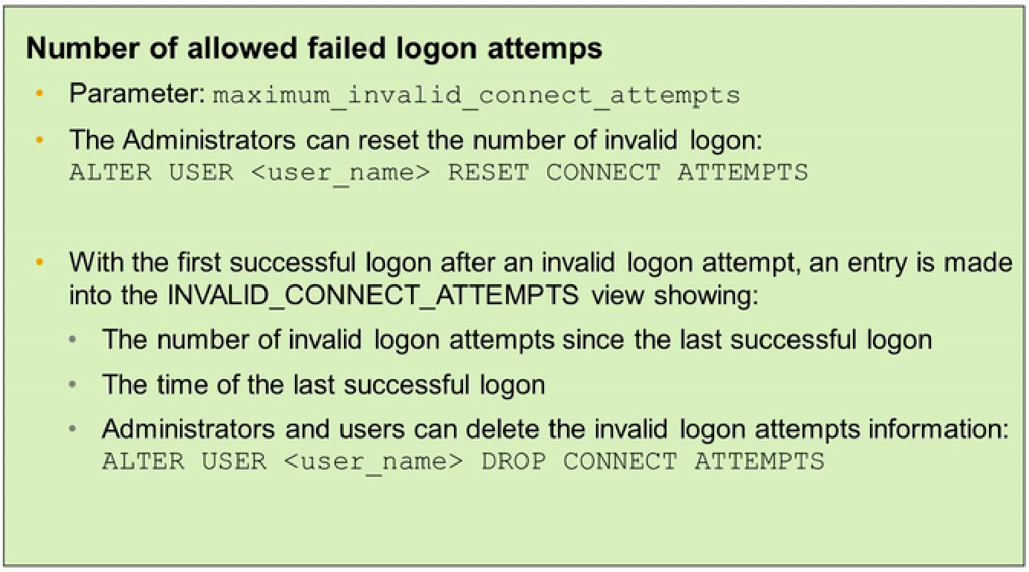

Figure 418: Manage Connection Attempts

Managing User Passwords

The `force_first_password_change` parameter defines whether users have to change their initial passwords immediately the first time they log on.

If this parameter is set to true, users can still log on with the initial password. But every action they try to perform returns an error message, which states that they must change their password.

If this parameter is set to false, users are not forced to change their initial password immediately the first time they log on. However, if a user does not change the password before the number of days specified in the parameter `maximum_unused_initial_password_lifetime`, then the password still expires and must be reset by a user administrator.

A user administrator (that is, a user with the system privilege USER ADMIN) can force a user to change his or her password at any time with the following SQL statement: `ALTER USER <user_name> FORCE PASSWORD CHANGE`.

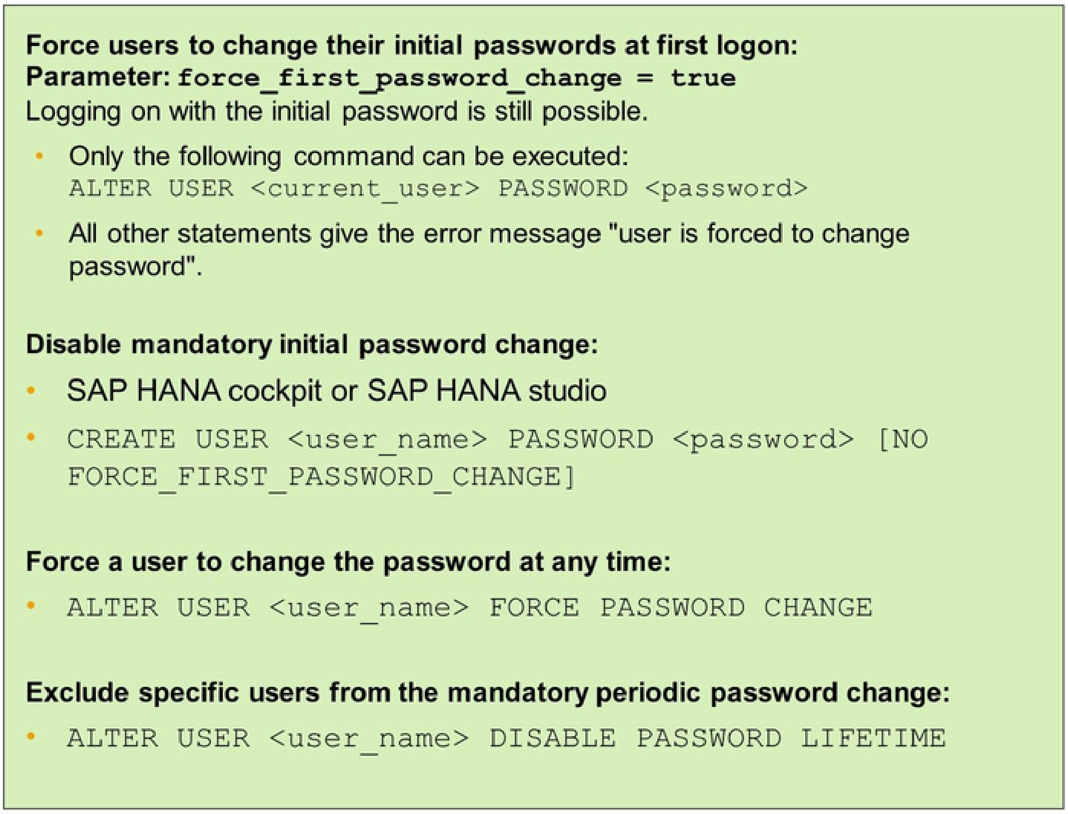

Figure 419: Manage User Password

SQL Statement for Password Policy

A user administrator can override this password policy setting for individual users in SAP HANA cockpit or SAP HANA studio or with the following SQL statement:

- CREATE USER <user_name> PASSWORD <password> [NO FORCE_FIRST_PASSWORD_CHANGE]
- Alter USER <user_name> PASSWORD <password> [NO FORCE_FIRST_PASSWORD_CHANGE]

This option is useful in the following situations:

- New technical user account with a generated password, which is not meant to be entered interactively.
- New users generated and managed by an identity management system.

Exclude Specific Users from the Mandatory Periodic Password Change

For connectivity purposes, you might create a technical user. This technical user should never change the password. You can re-enable the mandatory periodic password change with the following SQL commands:

- ALTER USER <user_name> DISABLE PASSWORD LIFETIME
- ALTER USER <user_name> ENABLE PASSWORD LIFETIME

Management of SYSTEM User

The SYSTEM database user is the initial user that is created during the creation of the SAP HANA database.

SYSTEM is the database superuser. It has irrevocable system privileges, such as the ability to create other database users, access system tables, and so on. In addition, to ensure that the administration tool SAP HANA cockpit can be used immediately after database creation, SYSTEM is automatically granted several roles the first time the cockpit is opened with this user. Note, however, that SYSTEM does not automatically have access to objects created in the SAP HANA repository.

In a system with multitenant database containers, the SYSTEM user of the system database has additional privileges for managing tenant databases, for example, creating and dropping databases, changing configuration (*.ini) files of databases, and performing database-specific data backups.

Do not use SYSTEM for day-to-day activities in production systems. Instead, use it to create database users with the minimum privilege required for their duties set (for example, user administration, system administration). Then, deactivate SYSTEM.

If the SYSTEM user's password is lost, you can reset it using the operating system user (`<sid>adm` user).

Managing SYSTEM User

You can manage the SYSTEM user as follows:

- Deactivate the SYSTEM User

  `ALTER USER SYSTEM DEACTIVATE USER NOW`

- Exempt SYSTEM User from locking

  Parameter: `password_lock_for_system_user`

- Reset the SYSTEM user's password

  - Reset the SYSTEM User Password of the System Database

    Reset it using the `<sid>adm` user (see SAP HANA System Administration Guide)

  - Reset the SYSTEM User Password of a Tenant Database

    The system administrator can reset it from the system database. Stop the tenant database and execute the following command: `ALTER DATABASE <database_name> SYSTEM USER PASSWORD <new_password>`

The parameter `password_lock_for_system_user` indicates whether or not the user SYSTEM is locked for the specified lock time (`password_lock_time`) after the maximum number of failed logon attempts (`maximum_invalid_connect_attempts`).

Deactivation of the SYSTEM User

SYSTEM is the database superuser. It has irrevocable system privileges, such as the ability to create other database users, access system tables, and so on. Do not use SYSTEM for day-to-day activities in production systems. Instead, use it to create database users with the minimum privilege required for their duties set (for example, user administration, system administration). Then, deactivate SYSTEM.

Execute the following statement, for example, in the SQL console of the SAP HANA studio:

```
ALTER USER SYSTEM DEACTIVATE USER NOW
```

The SYSTEM user is deactivated and can no longer connect to the SAP HANA database.

You can verify that this is the case in the user's system view. For user SYSTEM, check the values in the columns USER_DEACTIVATED, DEACTIVATION_TIME, and LAST_SUCCESSFUL_CONNECT.

> Note:
> You can still use the SYSTEM user as an emergency user, even if it has been deactivated. Any user with the system privilege USER ADMIN can reactivate SYSTEM by using the statement `ALTER USER SYSTEM ACTIVATE USER NOW`. To ensure that an administrator does not do this secretly, create an audit policy to monitor ALTER USER statements.

Reset of the SYSTEM User's Password

If the SYSTEM user's password is lost, you can reset it using the index server in emergency mode and the credentials of the operating system user. The complete procedure is described in detail in the SAP HANA System Administration Guide.

After performing this procedure, the password for the SYSTEM user is reset. Because you are logged on as the SYSTEM user in this console, you do not have to change this new password the next time you log on with this user, regardless of your password policy configuration.

The system database and all tenant databases each have their own SYSTEM user. The system administrator can reset the password of any SYSTEM user if it has been irretrievably lost.

Reset the SYSTEM User Password of the System Database

If the password of the SYSTEM user of the system database is lost, you can reset it as the operating system administrator by starting the name server in emergency mode.

To reset the SYSTEM user password of the system database the credentials of the operating system user <sid>adm are needed. Log on to the server on which the name server of the system database is running and execute the commands described in the SAP HANA Administration Guide.

The password of the SYSTEM user of the system database is reset. You have to change the new password the next time you log on with this user.

If you previously deactivated the SYSTEM user, it is now also reactivated. This means you will need to deactivate it again.

Reset the SYSTEM User Password of a Tenant Database

If the password of the SYSTEM user in a tenant database is lost, you as the system administrator can reset it from the system database.

As the system administrator of the system database you can stop the tenant database and execute the command: `ALTER DATABASE <database_name> SYSTEM USER PASSWORD <new_password>`. The password for the SYSTEM user is reset and the tenant database is started.

The password of the SYSTEM user of the tenant database is reset. You have to change the password the next time you log on with this user, this time in line with the password policy of the tenant database.

If the SYSTEM user was previously deactivated, locked, or expired, it is now activated again. We recommend that you deactivate it.

If auditing is enabled, the password change is automatically logged in both the system and tenant database audit trails.

> Note:
> If you can log on as SYSTEM or another user with the system privilege USER ADMIN, do not use the procedure described here to change the password of the SYSTEM user. Instead, log on to the tenant database directly and either execute the command `ALTER USER SYSTEM PASSWORD <new_password>` or change the password using SAP HANA cockpit.

LESSON SUMMARY
You should now be able to:

- Perform Administrative Tasks

Unit 12
Lesson 5

Information Sources for Administrators

LESSON OVERVIEW

This lesson covers the tables and views that support user management.

Business Example

After creating users and grant authorizations to them, you need to evaluate which privileges and roles are granted to the users.

LESSON OBJECTIVES

After completing this lesson, you will be able to:

- Explain Information Sources for Administrators

Information Sources for Administrators

You can query several system views to get detailed information about exactly which privileges and roles users have and how they come to have them. This can help you to understand why a user is authorized to perform particular actions, access particular data, or not.

You must have the system privilege CATALOG READ to query the following views.

System Tables and Monitoring Views

SQL | Result

```
select * from effective_roles where user_name='student50'
```

	USER_NAME	GRANTEE	GRANTEE_TYPE	GRANTOR	ROLE_NAME	IS_GRANTABLE
1	student50	STUDENT50	USER	SYS	PUBLIC	FALSE
2	student50	STUDENT50	USER	SYSTEM	TRAINING_USER_ROLE	FALSE
3	student50	TRAININ...	ROLE	SYSTEM	CONTENT_ADMIN	FALSE
4	student50	TRAININ...	ROLE	SYSTEM	INFORMATION_CON...	FALSE
5	student50	TRAININ...	ROLE	SYSTEM	MONITORING	FALSE
6	student50	TRAININ...	ROLE	SYSTEM	MODELING	FALSE
7	student50	TRAININ...	ROLE	SYSTEM	USER_ADMIN_ROLE	FALSE
8	student50	TRAININ...	ROLE	SYSTEM	SELECT_SYS_BI	TRUE

Column name	Description
USER_NAME	Name of the user for whom the effective roles are shown
GRANTEE	Shows to which user or role the role was granted
GRANTEE_TYPE	Shows whether the role was granted directly ('USER') or is inherited from another role ('ROLE')
GRANTOR	User that granted the role
ROLE_NAME	Role that was granted
IS_GRANTABLE	Shows whether the role was granted ‘WITH ADMIN OPTION‘: ‘TRUE or FALSE'

Figure 420: System Tables and Monitoring Views

System tables and monitoring views query information about the system using SQL commands. The results appear as tables in SYS Schema.

The system view M_CONNECTIONS contains additional information about the authentication method: `SELECT USER_NAME, AUTHENTICATION_METHOD FROM M_CONNECTIONS`. By default, users can only query information about themselves.

Display Roles Granted to a User

The system view EFFECTIVE_ROLES displays the roles of the currently logged-on user. It shows both the roles that were granted directly to the user, and the roles that were inherited from other roles. This system view complements the system view EFFECTIVE_PRIVILEGES.

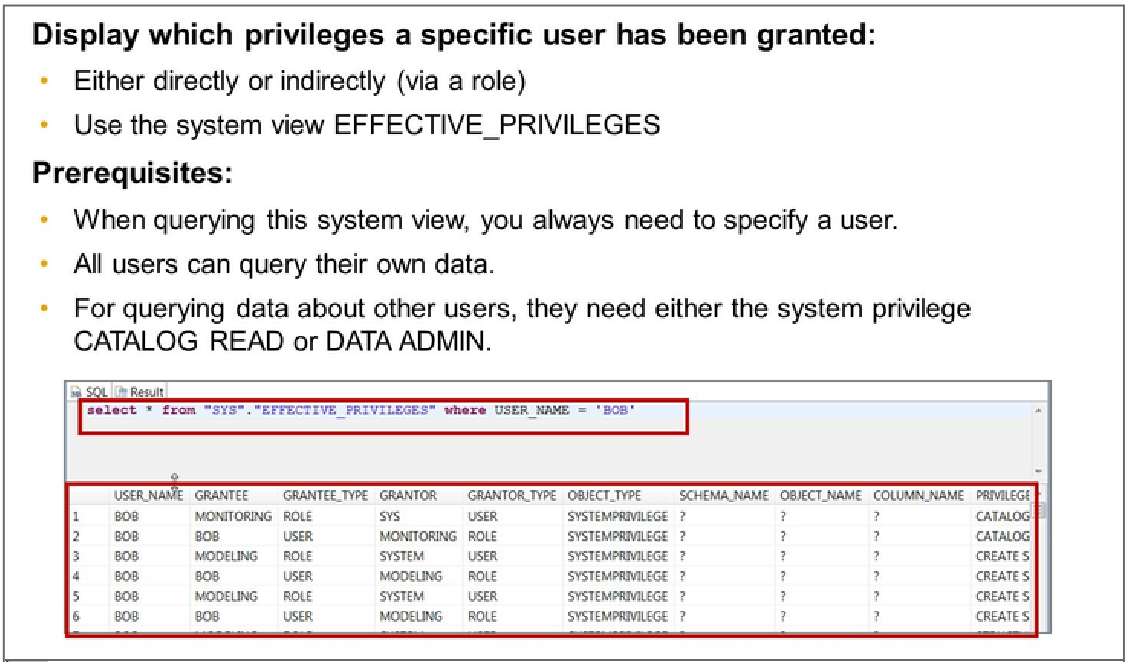

	USER_NAME	GRANTEE	GRANTEE_TYPE	GRANTOR	GRANTOR_TYPE	OBJECT_TYPE	SCHEMA_NAME	OBJECT_NAME	COLUMN_NAME	PRIVILEGE
1	BOB	MONITORING	ROLE	SYS	USER	SYSTEMPRIVILEGE	?	?	?	CATALOG
2	BOB	BOB	USER	MONITORING	ROLE	SYSTEMPRIVILEGE	?	?	?	CATALOG
3	BOB	MODELING	ROLE	SYSTEM	USER	SYSTEMPRIVILEGE	?	?	?	CREATE S
4	BOB	BOB	USER	MODELING	ROLE	SYSTEMPRIVILEGE	?	?	?	CREATE S
5	BOB	MODELING	ROLE	SYSTEM	USER	SYSTEMPRIVILEGE	?	?	?	CREATE S
6	BOB	BOB	USER	MODELING	ROLE	SYSTEMPRIVILEGE	?	?	?	CREATE S

Figure 421: Display Roles Granted to a User

Display Privileges Granted to a User

Because you can assign privileges directly or inherit them via roles, it can be difficult to see which privileges a user has been granted.

To provide better support, the view EFFECTIVE_PRIVILEGES was created.

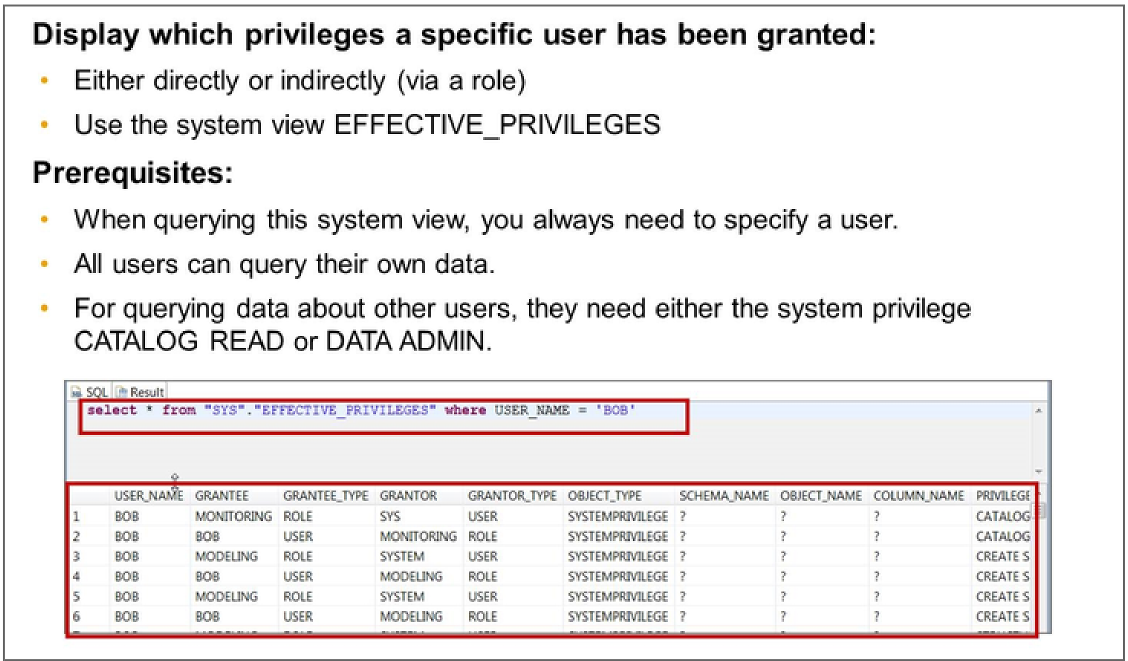

	USER_NAME	GRANTEE	GRANTEE_TYPE	GRANTOR	GRANTOR_TYPE	OBJECT_TYPE	SCHEMA_NAME	OBJECT_NAME	COLUMN_NAME	PRIVILEGE
1	BOB	MONITORING	ROLE	SYS	USER	SYSTEMPRIVILEGE	?	?	?	CATALOG
2	BOB	BOB	USER	MONITORING	ROLE	SYSTEMPRIVILEGE	?	?	?	CATALOG
3	BOB	MODELING	ROLE	SYSTEM	USER	SYSTEMPRIVILEGE	?	?	?	CREATE S
4	BOB	BOB	USER	MODELING	ROLE	SYSTEMPRIVILEGE	?	?	?	CREATE S
5	BOB	MODELING	ROLE	SYSTEM	USER	SYSTEMPRIVILEGE	?	?	?	CREATE S
6	BOB	BOB	USER	MODELING	ROLE	SYSTEMPRIVILEGE	?	?	?	CREATE S

Figure 422: Display Privileges Granted to a User

When selecting from EFFECTIVE_PRIVILEGES, you always need the condition USER_NAME = 'something' in the WHERE clause. Otherwise the query returns with an error.

Dependency Viewer

Use the authorization dependency viewer as a first step in troubleshooting authorization and invalid object errors for stored procedures and for calculation views with complex dependency structures.

The authorization dependency viewer helps you to identify invalid authorization dependencies in your object's structure. This is useful for objects with large and complex dependency structures. The authorization dependency viewer in the SAP HANA studio visualizes the object dependency structure of stored procedures and views, together with the SQL authorization status of the object owner along the dependency paths.

Use the authorization dependency viewer as a first step in troubleshooting the following authorization errors for column views and procedures:

- NOT AUTHORIZED (258)
- INVALIDATED VIEW (391)
- INVALIDATED PROCEDURE (430)

Authorization or invalid object errors occur if the object owner does not have all the required privileges for all underlying objects on which the object depends (for example, tables, views, and procedures). The object owner must have both the appropriate SQL object privilege (for example, EXECUTE, SELECT) and the authorization to grant the object privilege to others (that is, WITH GRANT OPTION is set).

The authorization dependency viewer helps you to identify invalid authorization dependencies in the object structure. This is useful for objects with large and complex dependency structures.

Hint:
Use the authorization dependency viewer only with procedures that have the DEFINER security mode. Procedures with the INVOKER security mode are not validated correctly.

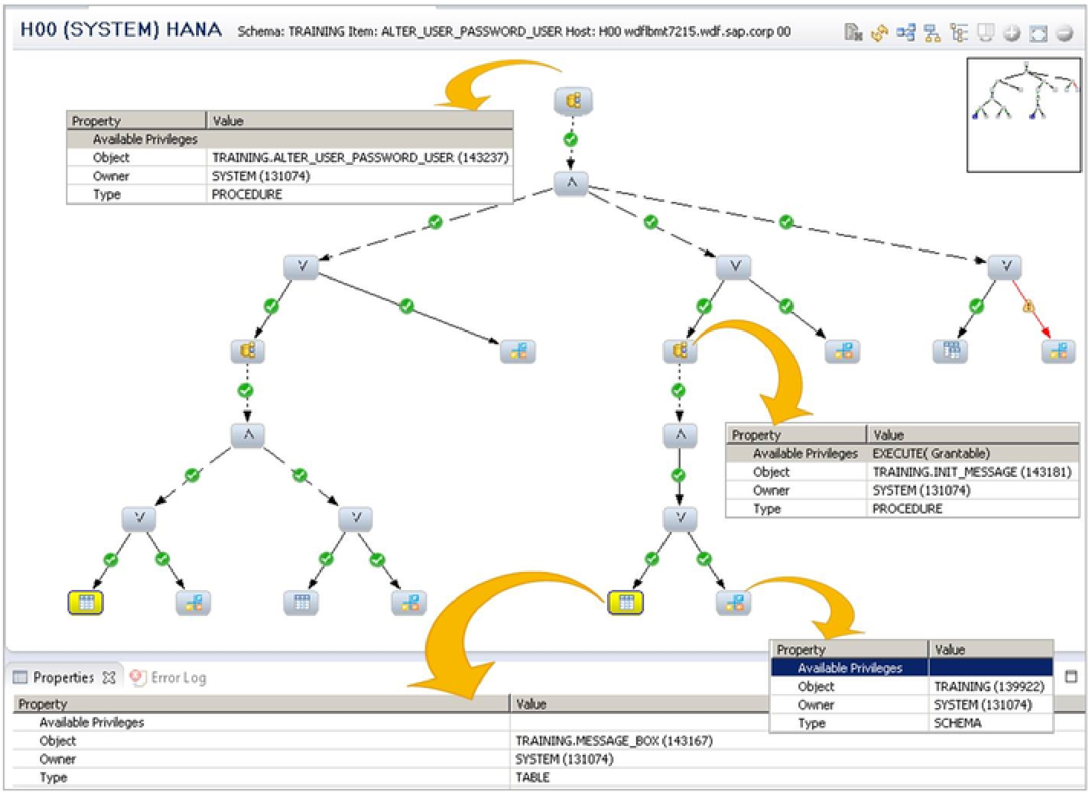

Figure 423: Dependency Viewer

LESSON SUMMARY

You should now be able to:

- Explain Information Sources for Administrators

Unit 12

Learning Assessment

1. Which of the following tools can be used for User Administration?

   Choose the correct answers.

   - [ ] **A** SAP HANA Cockpit
   - [ ] **B** hdbsql
   - [ ] **C** Resident hdblcm
   - [ ] **D** SAP NetWeaver Identity Management

2. Which of the following privileges allow read access to data in SAP HANA information models depending on certain values or combination of values?

   Choose the correct answer.

   - [ ] **A** System Privileges
   - [ ] **B** Object Privileges
   - [ ] **C** Analytic Privileges
   - [ ] **D** Package Privileges

3. Which of the following statements is true for the SAP HANA preinstalled role (SAP_INTERNAL_HANA_SUPPORT) for support cases?

   Choose the correct answer.

   - [ ] **A** It can be granted to the SYSTEM user.
   - [ ] **B** It can be granted to another role.
   - [ ] **C** It can only be granted further system privileges.
   - [ ] **D** It can only be granted further object privileges.

4. The password policy parameters for the system database of a multiple-container system are maintained in the “indexserver.ini” file.

   Determine whether this statement is true or false.

   - [ ] True
   - [ ] False

5. What are the prerequisites for querying the system view EFFECTIVE_PRIVILEGES to get data about other users and roles granted privileges?

   Choose the correct answers.

   - [ ] **A** CATALOG READ
   - [ ] **B** DATA ADMIN
   - [ ] **C** MONITOR ADMIN
   - [ ] **D** USER ADMIN

Unit 12

Learning Assessment - Answers

1. Which of the following tools can be used for User Administration?

   Choose the correct answers.

   [X] **A** SAP HANA Cockpit

   [X] **B** hdbsql

   [ ] **C** Resident hdblcm

   [X] **D** SAP NetWeaver Identity Management

   Correct! User management is configured using the SAP HANA Cockpit or the SAP HANA studio. You can create, delete, deactivate or reactivate users. Some other aspects are creating and assigning roles and privileges to users, etc. All the user management functions can also be executed from the SQL command line tool. SAP Identity Management provides additional support for user provisioning in the SAP HANA database. SAP Identity Management 7.2 SP 3 contains a connector to the SAP HANA database (IDM connector). With SAP Identity Management, you can create, delete user accounts, assign roles and set passwords for users in the SAP HANA database. The Resident HDBLCM Tool is not a user administration tool. You perform many complex database tasks, like renaming an SAP HANA system, converting an SAP HANA database to a multitenant database, add or remove Host Roles, add additional hosts to the SAP HANA database, configure the Interservice communication in a scale-out environment, configure the connection to the System Landscape Directory (SLD), convert an SAP HANA database system to a Multitenant Database Containers (MDC) system, uninstall the complete SAP HANA database or individual components, install or update additional SAP HANA components. remove an SAP HANA compute node from the SAP HANA storage to scale-up the compute node or rename the System ID or Instance Number of an SAP HANA system. Read more on this in the lesson SAP HANA Authentication and Authorization (Unit 12, Lesson 1) of the course HA200_14.

2. Which of the following privileges allow read access to data in SAP HANA information models depending on certain values or combination of values?

   Choose the correct answer.

   - [ ] **A** System Privileges
   - [ ] **B** Object Privileges
   - [x] **C** Analytic Privileges
   - [ ] **D** Package Privileges

   Correct! Analytic privileges allow read access to data in SAP HANA information models (that is, analytic views, attribute views, and calculation views) depending on certain values or combinations of values. Analytic privileges are evaluated during query processing. System privileges control general system activities. They are mainly used for administrative purposes, such as creating schemas, creating and changing users and roles, performing data backups, managing licenses, and so on. System privileges are also used to authorize basic repository operations. Object privileges allow access to and modification of database objects, such as tables and views. Depending on the object type, different actions can be authorized (for example, SELECT, CREATE ANY, ALTER, DROP, and so on). Package privileges allow access to and the ability to work in packages in the classic repository of the SAP HANA database. Packages contain design time versions of various objects, such as analytic views, attribute views, calculation views, and analytic privileges. Read more on this in the lesson Types of Feedback (Unit 12, Lesson 2) of the course HA200_14.

3. Which of the following statements is true for the SAP HANA preinstalled role (SAP_INTERNAL_HANA_SUPPORT) for support cases?

   Choose the correct answer.

   - [ ] **A** It can be granted to the SYSTEM user.
   - [ ] **B** It can be granted to another role.
   - [x] **C** It can only be granted further system privileges.
   - [ ] **D** It can only be granted further object privileges.

   Correct! For security reasons only system privileges can be granted to this role. To avoid accidental use of this role in day-to-day activities, this restrictions applies to the SAP_INTERNAL_HANA_SUPPORT role. Therefore it cannot be granted to the SYSTEM user or to another role and it cannot be granted further object privileges. Read more on this in the lesson SAP HANA Roles (Unit 12, Lesson 3) of the course HA200_14.

4. The password policy parameters for the system database of a multiple-container system are maintained in the “indexserver.ini” file.

   Determine whether this statement is true or false.

   [ ] True

   [X] False

   Correct! The password policy parameters for the system database of a multiple-container system are maintained in the "namesever.ini" file, not the "indexserver.ini" file. The actual parameters are contained in the password policy section of the indexserver.ini system properties file. You can configure the password policy using the Password Policy and Blacklist page in the SAP HANA cockpit and the Security editor of the SAP HANA studio. You can also do this by editing the indexserver.ini directly. Read more on this in the lesson Administrative Tasks (Unit 12, Lesson 4) of the course HA200_14.

5. What are the prerequisites for querying the system view EFFECTIVE_PRIVILEGES to get data about other users and roles granted privileges?

   Choose the correct answers.

   [X] **A** CATALOG READ

   [X] **B** DATA ADMIN

   [ ] **C** MONITOR ADMIN

   [ ] **D** USER ADMIN

   Correct! CATALOG READ authorizes unfiltered read-only access to all system views. Normally, the content of these views is filtered based on the privileges of the accessing user. DATA ADMIN authorizes the reading of all data in the system views. This privilege also enables the execution of any Data Definition Language (DDL) commands in the SAP HANA database. Users with this privilege cannot select or change data stored in tables for which they do not have access privileges. However, they can drop tables or modify table definitions. MONITOR ADMIN authorizes the execution of the ALTER SYSTEM commands for EVENTs. USER ADMIN authorizes the creation and modification of users using the CREATE USER, ALTER USER, and DROP USER commands. Read more on this in the lesson Information Sources for Administrators (Unit 12, Lesson 5) of the course HA200_14.

