SAP HANA Platform SPS 08
Document Version: 1.0 - 2014-05-28

SAP HANA Technical Operations Manual

Table of Contents

1 About Technical Operations

This Technical Operations Manual provides an end-to-end picture of the administration tools available with SAP HANA and the key tasks that a system administrator needs to perform. Links to the relevant administration documentation of each of the components included in the SAP HANA solution are provided for details and step procedures.

> ⚠ Caution
>
> This guide does not replace the daily operations handbook that we recommend customers to create for their specific production operations.

1.1 What Is SAP HANA?

SAP HANA is a modern, in-memory database and platform that is deployable on-premise or in the cloud.

The **SAP HANA platform** is a flexible data source agnostic in-memory data platform that allows customers to analyze large volumes of data in real-time. It is also a development platform, providing an infrastructure and tools for building high-performance applications based on SAP HANA Extended Application Services (SAP HANA XS). It is the foundation of various SAP HANA editions, like the SAP HANA Platform Edition, providing core database technology, and the SAP HANA Enterprise Edition, bundling additional components for data provisioning. The SAP HANA Platform Edition integrates a number of SAP components, including the SAP HANA database, SAP HANA studio, and SAP HANA clients.

1.2 Software Download

In the SAP Software Download Center, you have access to the installation media and components for SAP HANA.

You can find the installation media and components for SAP HANA in the following locations:

- Installation media for an SAP HANA SPS:
  ▶ *SAP Software Download Center* ❯ *Installations and Upgrades* ❯ *A-Z Index* ❯ *Installations and Upgrades - H* ❯ *SAP In-Memory (SAP HANA)* ❯ *HANA Platform Edition* ❯ *SAP HANA PLATFORM EDITION* ❯ *SAP HANA PLATFORM EDIT. 1.0* ❯ *Installation* ◗
- Support Packages and Patches for SAP HANA:
  ▶ *SAP Software Download Center* ❯ *Support Packages and Patches* ❯ *A-Z Index* ❯ *Support Packages and Patches - H* ❯ *SAP HANA PLATFORM EDITION* ❯ *SAP HANA PLATFORM EDIT. 1.0* ◗

The responsibility for acquiring and installing SAP HANA depends on the chosen deployment model:

- If a customer chooses the **SAP HANA tailored data center integration**, the components of SAP HANA have to be installed on validated hardware by a certified administrator.

- If a customer chooses an **SAP HANA appliance**, the components of SAP HANA can only be installed by certified hardware partners on validated hardware running a specific operating system. Any other system or content developed with systems of this type is not supported by SAP. For more information, see the information page of the product version. Support Package Stacks (SPS) can be downloaded and applied to appliances in accordance with agreements with the respective hardware partner.

Related Information

SAP Software Download Center

1.2.1 Download Patches and Updates

> i Note
>
> For more information on applying patches and updates see the section on Change Management.

Related Information

Change Management [page 11]
 Change management describes the permitted changes allowed to the standard configuration of an SAP HANA appliance.

2 Operating SAP HANA

You can operate SAP HANA on-premise using the appliance delivery model or using the tailored data center integration model.

Up to SPS 06 SAP HANA followed a strict appliance delivery model. With SPS 07 you can decide to implement SAP HANA using the appliance delivery model, meaning preconfigured software and hardware bundled by an SAP hardware partner, or you can opt for the SAP HANA tailored data center integration approach, which allows you more flexibility when integrating your SAP HANA system with your existing storage solution.

With the appliance model SAP co-ordinates all support requests for all components of the system including hardware with the responsible partners.

Tailored data center integration gives you the flexibility to install SAP HANA yourself on the same validated hardware as used for appliances but you are responsible for all aspects of the system, including managing support with all the involved partners.

Table 1: SAP HANA Appliance compared to SAP HANA Tailored Datacenter Integration

SAP HANA appliance	SAP HANA tailored data center integration
Fast implementation	More flexibility
Support fully provided by SAP	Customer aligns with the hardware partner on individual support model
Solution validation done by SAP and partners	Installation is done by the customer
Preconfigured hardware setup	Reuse existing infrastructure and thereby reduce costs
Preinstalled software	SAP HANA installation certification required for customer to perform SAP HANA software installation

Related Information

SAP HANA Appliance [page 7]
    SAP HANA comes as an appliance combining software components from SAP optimized on proven hardware provided by SAP's hardware partners.

SAP HANA Tailored Data Center Integration [page 14]
    SAP HANA tailored data center integration offers you more openness and freedom of choice to configure the layer for SAP HANA depending on your existing data center layout.

2.1 SAP HANA Appliance

SAP HANA comes as an appliance combining software components from SAP optimized on proven hardware provided by SAP's hardware partners.

As an appliance, SAP HANA combines software components from SAP optimized on proven hardware provided by SAP's hardware partners.

This approach offers you well-defined hardware designed for the performance needs of an in-memory solution out of the box. The appliance delivery is the first choice if you are looking for a preconfigured hardware set-up and a preinstalled software package for a fast implementation done by your chosen hardware partner and fully supported by both, the partner and SAP.

> **i Note**
>
> With SPS 07 you can decide to implement SAP HANA using the appliance delivery model, meaning preconfigured software and hardware bundled by an SAP hardware partner, or you can opt for the SAP HANA tailored data center integration approach, which allows you more flexibility when integrating your SAP HANA system with your existing storage solution.

2.1.1 Roles and Responsibilities

Using the appliance delivery model it is necessary to define which tasks, concepts, and procedures are necessary for the ongoing operation of the SAP HANA appliance software and who is responsible in each case.

- *Tasks and Responsibilities*
  The following table gives an overview of tasks and who is responsible for the task. It includes information on:

  - What SAP is responsible for and what SAP supports
  - What the SAP hardware partners do
  - What the customer is asked to do.

Phase	Task	Responsible
Provisioning/Setup		
	Installation of Hardware	Hardware Partner
	Installation of Operating System	Hardware Partner
	Installation of SAP HANA Platform	Hardware Partner
	Adding additional SAP HANA database instances (MCOS)	Customer
	Data Source Connectivity	Customer
	Installation of SMD Agent	Customer
Maintenance		
	Patching of Firmware	Customer (*)

Phase	Task	Responsible
	Patching of Operating System	Customer (*)
	Patching of SAP HANA platform components	Customer
	Patching of peripheral components	Customer
	Upgrade of Operating System	Customer (*)
	Update of SAP HANA platform components	Customer
Operations		
	General administration of SAP HANA database	Customer
	Backup and Recovery	Customer
	SAP HANA System Monitoring	Customer
	SAP HANA Database Monitoring	Customer
	Installation of third party software components	Customer
Support		
	Issue Handling Process	SAP, Customer (*), (**)

(*) The customer is generally responsible for maintenance of the SAP HANA system. If the customer has a special support agreement with the hardware partner, maintenance may be the responsibility of the hardware partner.

(**) SAP is the main point of contact and distributes all issues within the support organization by default, as is the case for other SAP applications. If the customer has defined special support agreements with the hardware partner, the customer should contact the hardware partner directly in the case of obvious hardware or operating system issues. If no agreements have been made neither SAP nor the hardware partner are responsible for the installation, maintenance and possible adjustment of external software installed on the SAP HANA system.

- *Schedule* - When and how often should each task be performed? What tasks are to be performed during ongoing operation, such as regular checks and periodic tasks? What tasks must be performed on demand and in special situations?

- *Procedures* - How can each task be performed? Several tools and transactions can be used for administering and monitoring SAP HANA.

A SAP HANA appliance can only be obtained from one of the following certified SAP hardware partners:

- Hewlett Packard (*http://www.hp.com/go/sap/hana*)

- IBM (*http://www.ibm.com/solutions/sap/hana*)

- Fujitsu Computers (*http://ts.fujitsu.com/hana*)

- CISCO systems (*http://www.cisco.com*)

- DELL (*http://www.dell.com*)

- Hitachi (*http://www.hds.com/solutions/applications/sap-application/sap-hana.html*)

- NEC (*http://www.nec.com*)

- Lenovo (*http://www.lenovo.com*)

- Huawei (*http://enterprise.huawei.com/en/*)

- VCE (*http://www.vce.com/products/specialized/sap-hana*)

For more information on supported software and hardware platforms refer to the SAP Product Availability Matrix (search for "HANA"), the *SAP HANA Master Guide* on the SAP Help Portal.

Related Information

http://help.sap.com/hana

2.1.2 Service and Support

If errors occur in any software component of SAP HANA, SAP is the main point of contact. SAP distributes all issues within the support organization by default, as is the case for other SAP applications.

If the customer has defined special support agreements with the hardware partner (such as a special Service Level Agreement (ESA)), the customer should contact the hardware partner directly in the case of obvious hardware or operating system issues.

To investigate SAP HANA-related problems, SAP support requires a support connection to all servers in the SAP HANA landscape. For more information about establishing a service connection for SAP HANA, see SAP Service Marketplace. See *SAP Note 1635304*.

> **i** Note
>
> You need a valid SAP Service Marketplace user to access SMP and related SAP Notes.

Related Information

SAP Note 1635304

2.1.3 Installation of the SAP HANA Appliance Software

The installation of the SAP HANA appliance software is performed by the SAP hardware partner.

The SAP hardware partners deliver the SAP HANA software together with corresponding, validated hardware as a prepackaged solution with the most recent support package stack (at the time of shipment of the SAP HANA appliance) to the customer site. The SAP hardware partner may add specific best-practices and SAP HANA software configuration. The installation is finalized by an onsite setup of the SAP HANA components.

This includes deploying the SAP HANA system in the customer data center, connectivity to the network, SAP system rename, and SAP Solution Manager connectivity, if applicable.

Once the network connection to the SAP HANA system has been established successfully and the SAP HANA system with the most recent support packages is reachable under its desired host name, IP address and SAP system ID from within the customer network, the establishment of data replication/connectivity to data source systems (including the deployment of additional replication components) and the installation of potential BI clients (such as SAP BusinessObjects Explorer on top of BOE) is again the responsibility of the customer or the consulting organization carrying out the SAP HANA project onsite for the customer.

> **i Note**
>
> With the appliance model the installation of the SAP HANA server components, in particular the installation of the SAP HANA database shall only be performed by certified hardware partners, on validated hardware, running a specific operating system, as published in the *SAP HANA Product Availability Matrix* on SAP Service Marketplace. Any other system or content developed with such systems is not supported by SAP for use in production. SAP HANA studio and clients need to be installed outside of the SAP HANA appliance and can therefore be downloaded and installed by SAP customers directly.
>
> Further note, there is a product "SAP HANA (platform|enterprise edition)", where the "SAP HANA database" is one component of the product. The product is also comprised of other components. In consequence, do not to use the individual installation tools of the various components to setup a SAP HANA system. Otherwise Lifecycle Management processes like the Software Update Manager (SUM) for SAP HANA, a tool for updating and patching of the SAP HANA product might fail if the product was not installed completely as a whole.

Installation of additional software

To ensure the smooth operation of an SAP HANA system and to fulfill and be compliant with IT policies, you need to use external software in the SAP HANA appliance. External software is software that was not delivered by SAP or by your SAP HANA appliance hardware partner.

SAP permits the installation and operation of external software that is required to fulfill the IT compliance and IT policy that are determined in operation, provided the prerequisites in the following SAP Notes are met.

SAP Note Number	Title
1730929	Using external tools in an SAP HANA appliance
1730932	Using backup tools with Backint
1730930	Using antivirus software in an SAP HANA appliance

Each of these SAP Notes refers to additional documentation about software and software versions that, in the experience of SAP Support, have caused problems in the customer environment or in the laboratories of SAP or the hardware suppliers and are therefore not recommended for use in the HANA appliance.

Neither SAP nor the hardware supplier of your SAP HANA appliance are responsible for the installation, maintenance and possible adjustment of the external software if no agreements have been made for this purpose.

Related Information

http://service.sap.com/pam

2.1.4 Change Management

Change management describes the permitted changes allowed to the standard configuration of an SAP HANA appliance.

The following sections provide information on permitted configuration changes to the SAP HANA software, operating system, and so on. You will also find information about the changes that are not supported by SAP or SAP hardware partners.

SAP allows changes in configuration parameters in the SAP HANA database, when this is recommended in SAP documentation, SAP Notes or by an SAP employee (for example, Consulting, Development, or Support).

To guarantee optimal performance and the highest stability, SAP HANA hardware and technology partners can deliver systems with settings that differ from the SAP standard system.

Customer-specific changes to operating system parameters and ports are permitted in agreement with the relevant SAP HANA hardware or technology partners.

SAP Note 1731000 describes configuration changes that resulted in problems in customer environments and in SAP or hardware partner lab environments.

To guarantee supportability from SAP, all changes that differ from the SAP standard system should be documented. This facilitates later root cause analysis if errors occur.

Related Information

SAP Note 1731000

2.1.4.1 Updating and Patching the Operating System

The customer is generally responsible for implementing operating system patches. If the customer has a special agreement with the hardware partner, support for operating system patching may be the responsibility of the hardware partner.

SAP supports additional changes to the configuration of the operating system, provided these changes do not contradict the currently applicable SAP HANA documentation, or their incompatibility with SAP HANA is already known. These incompatible configuration changes, which are known in customer environments and/or are known to have caused problems to the laboratories of SAP or to SAP HANA hardware or to technology partners, are described in *SAP Note 1731000*.

OS security patches may be installed immediately after they are available; however, the original packages of the distributors must be used, that a customer is entitled to within the framework of a valid support contract with the distributor or an authorized OEM.

Customer-specific changes to operating system parameters and ports are permitted in agreement with the relevant SAP HANA hardware or technology partners. To guarantee supportability on the SAP side, all changes that deviate from the SAP standard system should be documented. This facilitates later root cause analysis if errors occur.

Any updates related to kernel or runtime libraries (glibc) need to be validated and approved by SAP beforehand. Support Package Stacks (SPSs) can be downloaded and applied to the SAP HANA system only according to agreements with the respective hardware partner.

On rare occasions, SAP HANA might require a certain operating system patch. In this case, you must ensure that the configuration settings of the operating system persist. Do not change configuration settings when you patch the operating system unless explicitly stated in the corresponding SAP release note. SAP will state any dependencies in the relevant SAP Note published when a revision is released which requires such modifications.

> **i Note**
>
> In some cases, customers may have outsourced the administration of their landscapes. In such a case, the outsourced service provider should only change any settings after consultation with the hardware support provider. The customer is responsible for any complete upgrade of the operating system and the use of tools for distributing operating system patches in a distributed system landscape.

SUSE offers enhanced support for SUSE Linux in collaboration with SAP Linux Lab ('SUSE Priority Support for SAP'). This support offer facilitates communication and ensures high quality.

For more information, see *SAP Note 1056161 - SUSE Priority Support for SAP applications* and *SAP Note 171356 – SAP software on Linux: Essential information*. (SAP Service Marketplace user required).

Related Information

SAP Note 1056161

SAP Note 171356

SAP Note 1731000

2.1.4.2   Updating and Patching SAP HANA Appliance Software

The SAP HANA appliance is shipped with the latest available support package stack and the customer is usually responsible for implementing any subsequent updates and patches.

The SAP hardware partners ships the SAP HANA appliance pre-configured with the most recent support package stack that is available. The customer is responsible for the subsequent implementation of SAP HANA patches, revisions, or support packages and support package stacks.

In general the following points apply:

- Any system connected to the SAP HANA system, as a source or as a client, may require a minimum SAP HANA support package stack to be implemented. Unless explicitly stated, we recommend applying the latest SAP HANA support package stack after applying an update to the connected system.
- The SAP HANA support packages are cumulative. You can update directly from any older SAP HANA support package to any newer SAP HANA support package.
- New SAP HANA support packages mainly cover software improvements. SAP recommends that you download and install these new SAP HANA support packages if SAP support has to resolve issues with SAP HANA versions previously installed in your landscape. However, customers do not necessarily need to update their SAP HANA landscape with the newest SAP HANA support packages if their SAP HANA landscape works correctly.
- New SAP HANA support packages and patches will be produced and shipped at SAP's sole discretion. There is no periodic cycle for releasing and patching the SAP HANA software. The new SAP HANA support package stacks will be communicated to customers and hardware partners by means of SAP Notes.

An automated update of your SAP HANA system is performed using SAP HANA lifecycle manager. You can access the tool in SAP HANA studio. An additional command line interface is also be available.

The following scenarios are supported:

- SAP HANA standalone maintenance (online): In the SAP HANA studio, start the lifecycle management perspective to set up a connection to the SAP Service Marketplace, automatically download the requested SAP HANA SP-Stack to your SAP HANA system and apply the update.
- SAP HANA standalone maintenance (offline): Download the requested SP-Stack in advance from SAP Service Marketplace and make it available to your SAP HANA system. In the SAP HANA studio, enter the location of the downloaded SP-Stack and proceed directly with implementing the update.
- SAP Solution Manager based maintenance: Use Maintenance Optimizer in your local SAP Solution Manager to calculate and download SP-Stack in advance and make it available to your SAP HANA system. In the SAP HANA studio, enter the location of the downloaded SPS and proceed with applying the update.
- For more information, see the *SAP HANA Master Update Guide*. Additionally see the section Backing Up and Recovering the SAP HANA Database in the *SAP HANA Administration Guide*.

> **i Note**
>
> It is strongly recommended after the initial setup and before every system update, that you perform a full data and file system backup (including a configuration backup).

For the most current information on SAP HANA support packages and patches, see *SAP Note 1514967* – SAP HANA 1.0: Central Note (SAP Service Marketplace user required). Up to now there are no known dependencies between SAP HANA support package stacks and the SAP HANA hardware and installed firmware. The process for implementing SAP HANA support packages is described on *SAP Service Marketplace* (SAP Service Marketplace user required).

> **i Note**
>
> Special support agreements may exist with the hardware partner which include services such as patching.

Related Information

SAP HANA Master Update Guide
SAP HANA Administration Guide
SAP Note 1514967
SAP Service Marketplace

2.1.4.3 Updating and Patching File System Components

The customer is responsible for implementing patches for the shared file system.

If the customer has a special agreement with the hardware partner, the hardware partner is responsible for providing support for file system patching.

2.1.4.4 Updating and Patching Storage Components

The customer is responsible for implementing patches for the storage components.

If the customer has a special agreement with the hardware partner, the hardware partner is responsible for providing support for the storage components.

2.1.4.5 Updating and Patching Firmware of Underlying Hardware Components

The customer is responsible for implementing patches for the firmware of underlying hardware components.

If the customer has a special agreement with the hardware partner, the hardware partner is responsible for providing support for the firmware of underlying hardware components.

2.2 SAP HANA Tailored Data Center Integration

SAP HANA tailored data center integration offers you more openness and freedom of choice to configure the layer for SAP HANA depending on your existing data center layout.

In addition to SAP HANA as standardized and highly optimized appliance, SAP now offers the opportunity to run the SAP HANA server with your preferred storage solution. A SAP HANA server has the same bill of material as the certified SAP HANA appliance but without storage. This option enables a reduction in hardware and operational costs at through the reuse of existing hardware components and operational processes.

Roles and Responsibilities

With the appliance model SAP distributes all support requests regarding any component of SAP HANA to the correct part of the support organization. With tailored data center integration the customer is responsible for defining support agreements with the various partners and organizing all aspects of support.

Service and Support

Customers should work with their hardware partners to ensure hardware support requirements are being fulfilled.

A supportability tool called the SAP HANA HW Configuration Check Tool is provided by SAP, which allows you to check if the hardware is optimally configured to meet the requirements of SAP HANA. For more information see the *SAP HANA Administration Guide*.

Installation

A number of requirements have to be fulfilled before proceeding with installation of SAP HANA tailored data center integration.

- Only servers listed in the Product Availability Matrix for SAP HANA are supported. However, certain configuration adjustments are permitted for tailored data center integration. Note the following information:

  o No local disks and no flash cards are required.
  o Additional Fibre Channel adapters for a boot from a storage-area network are allowed.

- All storage devices must have successfully passed the hardware certification for SAP HANA.
- The exam "SAP Certified Technology Specialist (Edition 2013) – SAP HANA Installation" (E_HANAINS131) needs to be needs to be successfully passed for a person to perform SAP HANA software installations.

Change Management

- Updating and patching the Operating System
  With tailored data center integration the customer is responsible for updating and patching the Operating System.
- Updating and patching the SAP HANA Software
  With tailored data center integration the customer is responsible for installing, updating, and patching the SAP HANA software.

Related Information

Product Availability Matrix for SAP HANA

Partner Information Center

SAP Training and Certification Shop
SAP HANA Administration Guide

2.3 SAP HANA in Virtualized Environments

Further information about running SAP HANA in a virtualized environment can be found at following sources.

The main source of current information is *SAP Note 1788665 - SAP HANA running on VMware vSphere Environments*. See the Related Links section for more sources.

Related Information

SAP Note 1788665 - SAP HANA running on VMware vSphere Environments

SAP on VMware

SAP HANA Guidelines for being virtualized with VMware vSphere

2.4 Monitoring the SAP HANA Database

There are a number of tools available to you for monitoring the SAP HANA database, including the SAP HANA studio.

The SAP HANA database can be monitored with the SAP HANA studio. It enables you to check if your SAP HANA system runs as expected and provides possibilities to analyze problems. The SAP HANA studio also houses the SAP HANA modeler and the user management for the SAP HANA database.

SAP Solution Manager can be used for basic administration and monitoring of SAP HANA systems within existing SAP landscapes. It is used by SAP support for early problem analysis and as backbone for CTS+ transport integration.

SAP HANA also integrates with SAP Landscape and Virtualization Manager (LVM) for basic operation of larger/more complex SAP landscapes (including starting and stopping and management of dependencies).

For more information, see the section Monitoring SAP HANA Systems in the *SAP HANA Administration Guide*.

Related Information

SAP HANA Administration Guide

2.4.1 Integrating SAP HANA into a Central Operations Concept

You can integrate SAP HANA into an overall operations concept supported through SAP Solution Manager.

SAP HANA and SAP HANA-based applications are supported by SAP Solution Manager. If you already use SAP Solution Manager, the effort for integrating SAP HANA into your existing operations concept is relatively low. You can use its functions for SAP HANA related activities.

> ℹ **Note**
>
> SAP recommends to update your SAP Solution Manager to release 7.1, SP05 or higher to optimize the support for SAP HANA.

Related Information

Using SAP Solution Manager for SAP HANA
SAP Solution Manager documentation
SAP Solution Manager homepage on SMP
Knowledge transfer by Online Knowledge Products
SAP Solution Manager SCN community

2.4.1.1 Connecting SAP Solution Manager to SAP HANA

You can configure a connection to SAP HANA in SAP Solution Manager.

If you want to use capabilities of SAP Solution Manager, you have to make sure that the two systems know each other. Prerequisite for this is the registration of the SAP HANA system in the System Landscape Directory. From there, SAP Solution Manager gets the information that the SAP HANA system exists. The communication between the systems is based on a central agent infrastructure. The pre-configured agents are delivered by SAP and deployed on the SAP HANA appliance by the hardware partner.

The configuration of the connection itself is done as part of the basic configuration of SAP Solution Manager. In the guided procedure for Managed Systems Configuration you just need to set up the correct connection, assign the right agents, enter some parameters, create required users, and do a few more configurations. After this, you can start the collaboration of SAP HANA and SAP Solution Manager.

Some of the processes in SAP Solution Manager require additional configuration to specify how they should handle the SAP HANA database. For example, you have to specify in system monitoring which metrics you want to control. Or you have to define your transport landscape (development system -> quality assurance system -> productive system) for change control.

Related Information

SAP Note 1747682: Managed System Setup for HANA

2.4.1.2   Central Monitoring and Administration

The monitoring and alerting infrastructure of SAP Solution Manager is based on a central agent infrastructure.

SAP delivers pre-configured agents for SAP HANA. If these agents have been deployed on SAP HANA and have been connected to SAP Solution Manager, SAP Solution Manager receives all alerts of the SAP HANA studio. These alerts will be displayed in the unified alert inbox.

SAP Solution Manager also provides an automated regular status check of your SAP solutions: SAP EarlyWatch Alert (EWA). This tool monitors the essential administrative areas of SAP components and keeps you up to date on their performance and stability. If you have any ABAP-based system besides or on SAP HANA, you can include the SAP HANA database information into the EWA report of this system: availability of services; size and growth; performance; important alerts; and correct setting of configuration parameters.

Besides basic administration tasks, SAP provides specific task lists for SAP HANA, SAP NetWeaver Business Warehouse (BW) on SAP HANA, and SAP LT (Landscape Transformation) Replication Server. These lists describe additional activities for administering these components and provide best practices for individual tasks.

2.4.1.3   Analyzing the Root Cause of Problems

You can use SAP Solution Manager to analyze the root cause of problems in your system landscape.

Often, SAP HANA systems are integrated with business applications that either call application accelerators in the SAP HANA database or are deployed on the database. If a problem occurs in this setup, you have to take a systematic approach to identify the precise cause of the fault. This standardized process must first broadly identify the components concerned and then analyze them more closely to arrive at the cause of the fault by process of elimination (top-down approach).

End-to-end root cause analysis in SAP Solution Manager provides your support organization with tools and methods that reliably identify the affected component while minimizing the time needed to solve the problem. Besides your business applications, it supports also your SAP HANA database. This is the foundation for resolving problems in a holistic approach. The DBA Cockpit is used in SAP Solution Manager to get a detailed insight into the status of the database. Basically, this is about the same data that you can see in the SAP HANA studio for your in-memory database. But the DBA Cockpit supports other databases as well. 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.

Trace analysis records performance-related and functional faults in a specific user activity from the user's browser to the data stored in memory or on a storage sub-system. The measurement is triggered in the user interface and automatically activates recording of the relevant traces on every component processed by the user query.

After the root cause of a problem has been located, you can use detailed analysis to find a resolution. Workload analysis comprises server-related workload statistics for the connected systems. Exception analysis allows you to centrally analyze all exceptions from connected systems, such as serious error messages in logs or dumps. From here, you can start component-specific tools. And change analysis creates transparency for all changes (technical configuration, code, content) that have been made in the connected system landscape. This is particularly useful in the event of faults that occur once changes have been made to the production landscape.

2.4.1.4 Controlling Change

In integrated system landscapes SAP Solution Manager allows you to manage all changes centrally.

It is important to manage all changes in system landscapes using a central system. This is the only way to synchronously execute changes that affect more than one production system. For example, if you use SAP HANA application-specific accelerators (for example, CO-PA) in combination with your SAP ERP back-end, you might need to synchronize changes of the data models on SAP HANA with changes for accessing the changed models from the transactions on SAP ERP. SAP Solution Manager provides central transport functions for the entire system landscape as well as additional support for quality management.

The process for transporting changes from a SAP HANA development system to a quality assurance and further on to the productive system starts in SAP HANA Studio. There you retrieve content from the SAP HANA source system and export it to enhanced Change and Transport System (CTS+) in SAP Solution Manager. The exported delivery units are attached automatically to a transport request of CTS+. After you have released this transport request, CTS+ triggers the automatic activation of the change in the SAP HANA repository on the target system (quality assurance or production).

Besides supporting the transport and activation of new or changed data models, views, stored procedures, or analytic privileges, SAP Solution Manager offers additional possibilities to improve the quality of changes. You can use change analysis and reporting for getting information about the current status and history of changes. Configuration validation checks whether the systems of your system landscape (for example, development, quality assurance, and production) are configured consistently and according to the requirements. Quality Gate Management provides an additional quality inspection for projects and ensures changes are transported correctly and synchronously to the production systems. And Change Request Management within SAP Solution Manager controls the entire change execution with detailed process and workflow control. Out of these options, you can select those for controlling changes on SAP HANA that fit best to the quality strategy of your organization.

Related Information

CTS+ How To Guide on SCN

2.5   Administering the SAP HANA Database

You administer the SAP HANA database using the Administration Console perspective of the SAP HANA studio.

The Administration Console perspective of the SAP HANA studio allows technical users to manage the SAP HANA database as well as to create and manage user authorizations.

For more details, see the *SAP HANA Administration Guide*.

For further reference information, see *SAP HANA SQL and System Views Reference*.

Related Information

SAP HANA Administration Guide
SAP HANA SQL and System Views Reference

2.5.1   Initial Administrative Tasks

Once you have set up your new system there are initial administrative tasks you carry out to safeguard your SAP HANA implementation.

It is strongly recommended after the initial setup and after the initial load, that you perform a full data and filesystem backup (including a configuration backup). For more information, see the section Backing Up and Recovering the SAP HANA Database in the SAP HANA Administration Guide.

Related Information

SAP HANA Administration Guide

2.5.2   Regular Administrative Tasks

You can regularly check on the SAP HANA system status using the SAP HANA studio.

- Regularly check the system status on the *Overview* tab page of the administration editor in SAP HANA studio. This displays the most important system information:
  - Overall system state
  - General system information (software versions and so on).
  - The Alerts section shows the latest warnings and messages generated by the statistics server, which is a monitoring tool for the database. It collects statistical and performance information using SQL statements.

- The bar views provide an overview of important system resources: the amount of available memory, CPUs and storage space is displayed as well as the used amount of these resources. In a distributed landscape, the amount of available resources is aggregated over all servers. Additionally, the resource information of the server with the highest resource consumption is displayed.
- The Services tab page (under the Landscape tab) of the administration screen in SAP HANA studio, lets you check that all services belonging to your system are running: preprocessor, name server and index server for each host and one statistics server. A statistics server exists on each host of a distributed landscape; however, the service actually runs on the master host only. Running services are indicated by a green icon. Information about resource usage and possible bottlenecks is also displayed.
- Regularly perform data backups including configuration backups. There are no general guidelines for backup frequency (this depends on the usage scenario). For more information, see the section Backing Up and Recovering the SAP HANA Database in the SAP HANA Administration Guide.

Related Information

SAP HANA Administration Guide

2.5.3   On Demand Administrative Tasks

There are a number of administrative tasks that can be run on demand such as diagnostic utilities and backups.

- In case of problems with the SAP HANA database, you can check log and trace files for errors. These log files are available in the SAP HANA studio on the *Diagnosis Files* tab page of the administration screen.
- You can turn on and configure several traces from the *Trace Configuration* tab. For more details see Configuring Traces in the SAP HANA Administration Guide.
- In certain situations, you have to restart the system (for example, after a power failure). This can be done using the *SAP HANA Systems* view of the SAP HANA studio. After confirming the restart, you need to enter the operating system user and password before the system is actually restarted. For more details see the SAP HANA Administration Guide.
- Avoiding LOG FULL (file system full) situations
  When the log is backed up, the backed up log segments remain on disk until they have been released automatically after a savepoint. After the log has released, the oldest unused log segment can be 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 has completed. Automatic log backup prevents log full situations from arising. For more information, see the section Backing Up and Recovering the SAP HANA Database in the SAP HANA Administration Guide.
- Avoid log backup area becoming full.
  Regularly archive old log backups to a different location (using operating system commands).
- Monitor disk space that is used for diagnosis files (these files record errors and changes in system behavior) and can be viewed on the *Diagnosis Files* tab in SAP HANA studio. We recommend using SAP HANA studio to delete diagnosis files that are no longer needed.

> ⚠ **Caution**
>
> Do not delete log segments on the operating system level, as the log area will become unusable and the database may stop working.

- Before the SAP HANA database or SAP HANA software is upgraded to a new version, perform a data backup including a configuration backup. This allows for the recovery of the system in case the software upgrade fails.
- After the initial load has finished successfully, a data backup is strongly recommended.
- For more information, see the section Backing Up and Recovering the SAP HANA database in the *SAP HANA Administration Guide*.

> ℹ **Note**
>
> For replication scenarios during the initial data load - from ERP into the SAP HANA database – log writing should not be switched off.

> ℹ **Note**
>
> There is no system table or log file that records the information that log writing has been switched off. For this reason, SAP does not recommend switching off the log during initial load. It is not possible to check when log writing has been switched OFF/ON on customer systems.

Related Information

SAP HANA Administration Guide

2.6 Updating and Configuring SAP HANA

For information about updating SAP HANA see the *SAP HANA Server Installation and Update Guide*.

For information about configuring SAP HANA see the *SAP HANA Administration Guide*.

Related Information

SAP HANA Administration Guide
SAP HANA Server Installation and Update Guide

2.7 Backup and Recovery

Although SAP HANA is an in-memory database it still uses persistent disk-based storage for backup and recovery.

The backup and recovery functionality in SAP HANA is quite straightforward. The payload of all data volumes of all services (name server, index server, statistics server) that belong to a database is written to service-specific backups. Backups can be written to the file system or using a third-party backup tool. An administrator can back up the database with only a few clicks and the system does the rest automatically: all the relevant service-specific data is backed up, regardless of whether your system has a single node or multiple nodes.

In-Memory Data and Persistent Storage

SAP HANA holds the bulk of its data in memory for maximum performance, but still uses persistent storage to provide a fallback in case of failure. During normal database operation, data is automatically saved from memory to disk at regular intervals called savepoints. Additionally, all data changes are captured in the redo log. With each committed database transaction, the redo log is saved from memory to disk and also log backups for SAP HANA are available.

If a failure occurs, for example a power outage, the database can be restarted in the same way as any disk-based database, and it is returned to its last consistent state by replaying the redo log since the last savepoint.

Backups

Data backups contain the current payload of the data volumes. All services that persist data on disk have a data volume and are included in the data backup. SAP HANA automatically handles the synchronization of the data backup of all services across all hosts in distributed (scale-out) systems. No user interaction is required for synchronization. You can either start data backups manually (for example, using SAP HANA studio or SQL commands), or schedule them (for example, using DBA Cockpit).

Log backups contain the redo log information. They are carried out automatically whenever a log segment is full or the timeout for log backups has elapsed.

Information on all backups is recorded in the backup catalog, which itself is also backed up.

SAP HANA provides native backup and recovery functions as well as an integrating with third party tools:

- Backups to the file system, For example, an NFS share
- Backups to a third party backup server using the "Backint for SAP HANA" API (certification for third party tool required)
- Storage snapshots using a third party storage tool

Recovery

A recovery may be necessary in one of the following situations:

- Data area is unusable
- Log area is unusable
- Logical error.

You can recover SAP HANA to:

- A point in time in the past (log replay required), or
- A specific data backup or storage snapshot (no log replay).

Depending on the recovery target, SAP HANA uses data backups or storage snapshots in combination with log backups (and if available log entries from the log volumes) for the recovery.

> **i Note**
>
> For the recovery SAP HANA needs to be shut down. During recovery, SAP HANA is therefore not accessible for end-users or applications.

For more information, see the section Backing Up and Recovering the SAP HANA database in the *SAP HANA Administration Guide*.

Related Information

SAP HANA Administration Guide

2.8 Security

How you implement SAP HANA determines what you need to consider from a security perspective.

SAP HANA can be implemented in the following main ways:

- As a data mart for reporting and analytics
- As the primary database in a traditional 3-tier architecture, for example, in SAP NetWeaver Business Warehouse (SAP NetWeaver BW) or SAP Business Suite installations
- As a platform for providing database and application services to native SAP HANA-based applications

For more information about the security-relevant information that applies to SAP HANA in your scenario, see the *SAP HANA Security Guide*.

Related Information

SAP HANA Security Guide

2.8.1 Security Features and Functions

SAP HANA provides security-related features that enable you to implement different security policies and meet compliance requirements.

Depending on the scenario in which SAP HANA is used, only some of these features might actually be needed, others might be provided in other architecture layers.

SAP HANA supports standard interfaces to enable integration with customer security network and data center infrastructures.
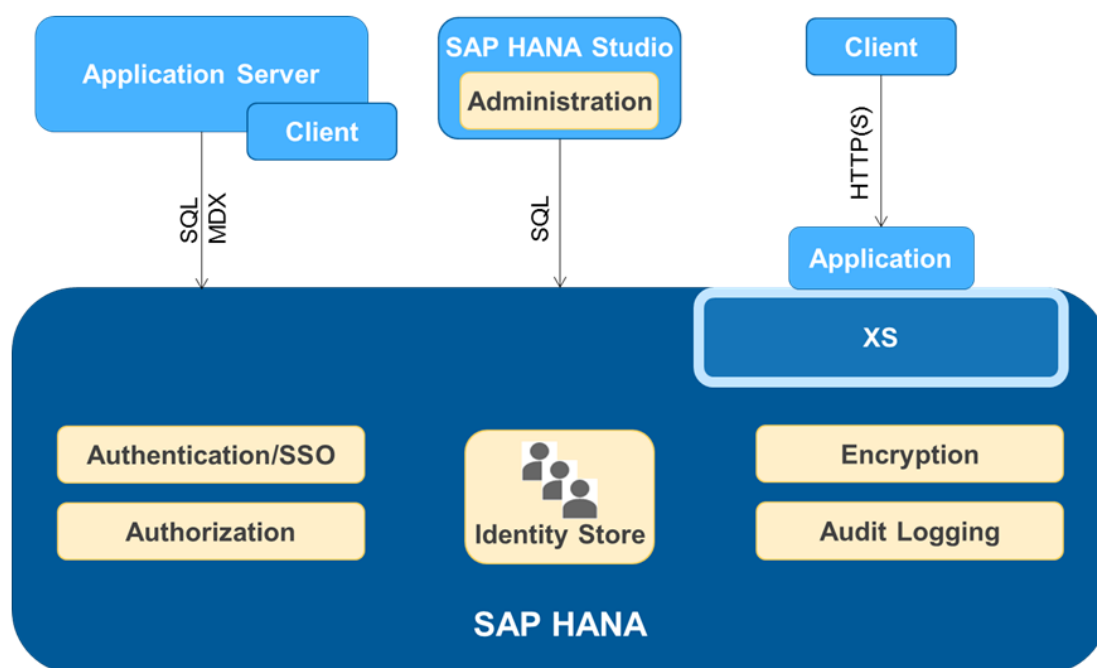
Authentication and Single Sign-On (SSO)

The identity of every database user accessing the database is verified through a process called authentication. The SAP HANA database supports several authentication mechanisms:

- JDBC/ODBC access: User name and password (including password policy), Kerberos, Security Assertion Markup Language (SAML bearer token), and SAP logon/assertion tickets
- HTTP access (SAP HANA XS): User name and password, SAP logon/assertion tickets, SAML, X.509 certificates, and Kerberos with Simple and Protected GSSAPI Negotiation Mechanism (SPNEGO)

SAP HANA never stores passwords in plain text. To enforce password quality, a password policy can be defined for user name/password access.

User and Role Management

Every user who wants to work directly with the SAP HANA database must have a database user with the necessary privileges.

Depending on the scenario, the user accessing the SAP HANA database may either be a technical user or an individual end user. Independent of the SAP HANA scenario, there will always be direct access to SAP HANA by database administrators. For user and role management, you can use the SAP HANA studio, SQL commands or user provisioning tools such as SAP Identity Management and SAP Access Control.

Authorization

All access to data and execution of actions in the database requires authorization. The user must have both the privilege(s) to perform the operation and the 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. A role is a set of privileges. Roles are the standard mechanism of granting privileges as they allow you to implement both fine-grained and coarse-grained reusable hierarchies of user access that can be modeled on business roles.

There are two types of roles available in SAP HANA: catalog and repository roles.

* Catalog roles follow traditional SQL concepts and are created directly in the production system. They are well suited for scenarios where the role creation and assignment process is carried out using a technical account, for example using a user provisioning tool.
* Repository roles are created in the development system and then transported to the production system, where they are assigned to end users by user/role administrators. Repository roles allow for a better separation of the role creation process from the role assignment process.

The following privilege types exist:

* System privileges: Authorize execution of administrative actions for the entire SAP HANA database
* SQL privileges: Authorize access to data and operations on database objects
* Analytic privileges: Authorize read access to analytic views at runtime and provide row-level access control based on the dimensions of the relevant view
* Application privileges: Authorize access to SAP HANA XS application functions, apply on top of the other privileges
* Package privileges: Authorize access of developers to the repository

Secure Communication

Network traffic can be encrypted using the Secure Sockets Layer (SSL) protocol. Communication between the SAP HANA database and clients can be secured, as well as communication between hosts in distributed SAP HANA systems and between sites in system replication (high availability) scenarios.

Encryption

For backup encryption, you have the following options:

1. Use a certified third party backup tool that supports encryption. For a list of certified backup tools, see

   http://www.sap.com/partners/directories/SearchSolution.epx (SAP-defined integration scenarios = HANA-BRINT 1.1 – HANA Backint Interface)
2. If you have encryption enabled for the SAP HANA data volumes and use storage snapshots for backup, the storage snapshots will contain these encrypted data volumes.

Auditing

The auditing feature of the SAP HANA database allows you to track actions performed in the database: who did what (or tried to do what), and when. SAP HANA provides audit actions for critical security events and for access to sensitive data. Both successful and unsuccessful events can be logged. The audit trail can be written to the logging system of the Linux operating system (syslog) or an internal database table.

Auditing can be configured in the SAP HANA studio or using SQL statements. Auditing can be enabled and disabled for the entire database only. User-defined audit policies specify which actions are logged.

Security Administration

Many security-related features are integrated into the SAP HANA studio, for example, user and role management, definition of analytic privileges, creation of audit policies, and configuration of password policy. Most administration tasks can also be carried out using SQL commands.

Application Services Security

The security concept of SAP HANA XS is fully integrated with the SAP HANA security model.

- User and role management: SAP HANA database users and privileges/roles are used.
- Authorization: User needs access to SAP HANA database objects. To provide even finer grained access control in SAP HANA XS, application privileges can also be defined.
- Authentication and single sign-on: User name and password, SAP logon/assertion tickets, SAML, X.509 certificates, and Kerberos with SPNEGO

- Communication encryption: Secure Sockets Layer (SSL)
- Data encryption: Inherits SAP HANA data volume encryption
- Audit logging: The SAP HANA audit logging infrastructure is used.

There are additional functions available to secure Web applications developed on SAP HANA XS, including protection against XSRF, SQL injection, and XSS.

2.8.2   Data Center Integration

SAP HANA supports standard interfaces to enable integration with customer security network and data center infrastructures.

There is detailed information available for secure set-up of the SAP HANA appliance and network integration.
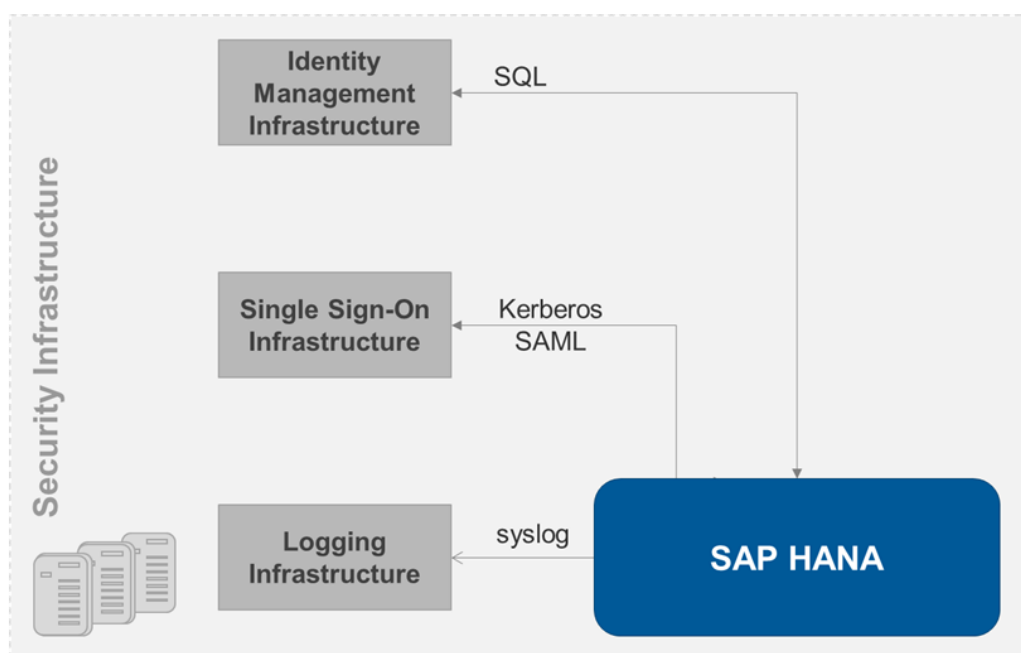
Operating System

SAP HANA is based on SUSE Linux Enterprise 11 for SAP and includes security pre-configurations (for example, minimal network services). The SAP HANA database runs with a restricted operating system user. Operating system access using SSH, for example, is not required for the day-to-day operations of SAP HANA.

Security Patches

SAP HANA security patches are published as part of the general SAP security patch strategy (SAP Security Notes). SAP HANA security patches are delivered as SAP HANA revisions and can be applied via SUM for SAP HANA. Operating system security patches are provided and published by SUSE and can be applied by customers as soon as they become available.

Network Integration

Network communication (purpose, ports) used by SAP HANA are documented in the *SAP HANA Security Guide* and in the *SAP HANA Master Guide*. This guide includes recommendations for the use of firewalls and how to separate internal and external communication. SSL is supported for network communication channels.

Related Information

SAP HANA Security Guide
SAP HANA Master Guide

2.8.3   Compliance

SAP HANA takes an end-to-end approach to enable application compliance.

This approach is based on the built-in database security features and the secure pre-configuration of the whole software stack (for example, operating system). This is extended by integration with existing security infrastructures via standard/documented interfaces and the option to use third-party tools that are required for data center operations. This is complemented by an end-to-end documentation covering the whole software life cycle, including an exhaustive security guide and recommendations for a secure setup.

Deletion

As part of the technical infrastructure for applications, SAP HANA provides delete functions for records in all types of tables. These functions can be leveraged by applications to implement their specific deletion requirements.

2.9 High Availability for SAP HANA

SAP HANA is fully designed for high availability. It supports recovery measures ranging from faults and software errors, to disasters that decommission an entire data center. High availability is the name given to a set of techniques, engineering practices and design principles that support the goal of business continuity.

High availability is achieved by eliminating single points of failure (fault tolerance), and providing the ability to rapidly resume operations after a system outage with minimal business loss (fault resilience). Fault recovery is the process of recovering and resuming operations after an outage due to a fault. Disaster recovery is the process of recovering operations after an outage due to a prolonged data center or site failure. Preparing for disasters may require backing up data across longer distances, and may thus be more complex.

The key to achieving high availability is redundancy, including hardware redundancy, network redundancy and data center redundancy. SAP HANA provides several levels of defense against failure-related outages:

1. Hardware Redundancy – SAP HANA appliance vendors offer multiple layers of redundant hardware, software and network components, such as redundant power supplies and fans, enterprise grade error-correcting memories, fully redundant network switches and routers, and uninterrupted power supply (UPS). Disk storage systems use batteries to guarantee writing even in the presence of power failure, and use striping and mirroring to provide redundancy for automatic recovery from disk failures. Generally speaking, all these redundancy solutions are transparent to SAP HANA's operation, but they form part of the defense against system outage due to single component failures.

2. Software – SAP HANA is based on SUSE Linux Enterprise 11 for SAP and includes security pre-configurations (for example, minimal network services). Additionally, the SAP HANA system software also includes a watchdog function, which automatically restarts configured services (index server, name server, and so on), in case of detected stoppage (killed or crashed).

3. Persistence – SAP HANA persists transaction logs, savepoints and snapshots to support system restart and recovery from host failures, with minimal delay and without loss of data.

4. Standby and Failover – Separate, dedicated standby hosts are used for failover, in case of failure of the primary, active hosts. This improves the availability by significantly reducing the recovery time from an outage.

SAP HANA High Availability Support

As an in-memory database, SAP HANA is not only concerned with maintaining the reliability of its data in the event of failures, but also with resuming operations with most of that data loaded back in memory as quickly as possible.

SAP HANA supports the following recovery measures from failures:

* Disaster recovery support:
    * Backups: Periodic saving of database copies in safe place.
    * Storage replication: Continuous replication (mirroring) between primary storage and backup storage over a network (may be synchronous).
    * System replication: Continuous update of secondary systems by primary system, including in-memory table loading.
* Fault recovery support:

- Service auto-restart: Automatic restart of stopped services on host (watchdog).
- Host auto-failover: Automatic failover from crashed host to standby host in the same system.
- System replication: Continuous update of secondary systems by primary system, including in-memory table loading.

System replication is flexible enough that it can also be used for both fault and disaster recovery to achieve high availability. The data pre-load option can be used for fault recovery to enable a quicker takeover than with Host Auto-Failover. You can build a solution with single node systems and do not need a scale out system and the additional storage and associated costs.

Using Secondary Servers for Non-Productive systems

With SAP HANA system replication you can use the servers on the secondary system for non-productive SAP HANA systems under the following conditions:

- Table pre-load is turned off in the secondary system.
- The secondary system uses its own disk infrastructure. In the case of single node systems this means, the local disk infrastructure needs to be doubled.
- The non-productive systems are stopped with the takeover to the productive secondary.

Related Information

SAP HANA Administration Guide

3   Operating SAP HANA Data Provisioning Technologies

You operate the extended SAP HANA system landscape by operating the SAP HANA components as well as the SAP HANA database.

The main components for data replication include:

- Trigger-based replication (SAP Landscape Transformation (LT) Replication Server)
- SAP HANA Direct Extractor Connection (DXC)
- ETL-based replication (SAP DATA Services)
- Log-based replication (SAP Replication Server)
- SAP Event Stream Processor.

Links to the relevant administration documentation of each of the components are included in the following sections.

3.1   Monitoring the SAP HANA System Landscape

You monitor the extended SAP HANA system landscape by monitoring these SAP HANA components:

- Monitoring components for trigger-based replication: SAP Landscape Transformation (LT) Replication Server
- Monitoring components for SAP HANA Direct Extractor Connection (DXC)
- Monitoring components for ETL-based replication: SAP Data Services

3.1.1   Monitoring Components for Trigger-Based Data Replication (SAP LT Replication Server)

For information about monitoring the SAP LT (Landscape Transformation) Replication Server, see the *Application Operations Guide – Trigger-Based Data Replication Using SAP LT Replication Server for SAP HANA*.

Related Information

Application Operations Guide – Trigger-Based Data Replication Using SAP LT Replication Server for SAP HANA

3.1.2 Monitoring components for SAP HANA Direct Extractor Connection (DXC)

For more information about monitoring the SAP HANA Direct Extractor Connection, see *SAP HANA Direct Extractor Connection Implementation Guide*.

Related Information

SAP HANA Direct Extractor Connection Implementation Guide

3.1.3 Monitoring Components for ETL-Based Replication (SAP Data Services)

For information about monitoring SAP Data Services, see:

- *SAP Data Services Management Console Guide*
  Information about how to use SAP Data Services Administrator and SAP Data Services Metadata Reports.
- *SAP Data Services Administrator's Guide*
  Information about administrative tasks such as monitoring, life cycle management, security, and so on.
- *SAP Data Services Performance Optimization Guide*
  Information about how to improve the performance of SAP Data Services.

3.2 SAP HANA Administration Overview

You administer the SAP HANA system landscape by administering the SAP HANA components:

- Administering components for Trigger-Based Replication: SAP LT (Landscape Transformation) Replication Server
- Administering components for ETL-Based Replication: SAP Data Services
- Administering components for SAP HANA Direct Extractor Connection (DXC)

3.2.1 Administrating Components for Trigger-Based Replication (SAP LT Replication Server)

For information about administering the SAP LT (Landscape Transformation) Replication Server, see the *Application Operations Guide – Trigger-Based Data Replication Using SAP LT Replication Server for SAP HANA*

Related Information

Application Operations Guide – Trigger-Based Data Replication Using SAP LT Replication Server for SAP HANA

3.2.2 Administrating Components for SAP HANA Direct Extractor Connection (DXC)

For information about administering the SAP HANA Direct Extractor Connection (DXC), see, *SAP HANA Direct Extractor Connection Implementation Guide*

3.2.3 Administrating Components for ETL-Based Replication (SAP Data Services)

For information about administrating SAP Data Services components, see the following guides:

* *Management Console Guide SAP Data Services XI 4.0 (14.0.0)*
* *SAP Data Services Administration Guide*
* *Information platform services Administration Guide*
* *SAP Data Services Upgrade Guide*

3.3 Backup and Recovery of SAP HANA Replication Technologies

You backup and recover the SAP HANA system landscape by extending the backup and recovery strategy to include the SAP HANA components as well as the SAP HANA database:

* Backup and Recovery for the Trigger-Based Replication components: SAP LT (Landscape Transformation) Replicator
* Backup and Recovery for SAP HANA Direct Extractor Connection (DXC)
* Backup and Recovery for ETL-Based Replication components: Data Services

3.3.1 Backup and Recovery for the Trigger-Based Replication components (SAP LT Replication Server)

For information about backup and recovery for the SAP LT (Landscape Transformation) Replication Server, observe the following:

- The SAP LT (Landscape Transformation) Replication Server creates a schema in the SAP HANA database called SYS_REPL where the replication control tables are stored (RS_* tables, to store replication status for example).
- These tables are part of the normal database backup/recovery.
- Recovery scenarios:
  - If the database is recovered to the latest consistent state, no further action is necessary, the replication automatically resumes (as the jobs are not running on the SAP HANA hardware).
  - If the database is recovered to an earlier point in time (recovery of data backup with clear log option), initial load and replication have to be started again.

3.3.2 Backup and Recovery for the SAP HANA Direct Extractor Connection (DXC)

The architecture of DXC involves components implemented in the SAP Business Suite system and also in SAP HANA. Thus, a significant amount of configuration data, metadata, and the state of the mechanisms for delta handling (change data capture) are stored in the SAP Business Suite system. Of course, your SAP Business Suite system must be backed up regularly.

- A schema is created for each SAP Business Suite system connected to SAP HANA via DXC, and for each DataSource a corresponding In-Memory DataStore Object is created in SAP HANA.
- These In-Memory DataStore Objects are be part of any normal SAP HANA database recovery.
- Recovery scenarios:
  - In any recovery scenario, data consistency must be maintained between the SAP Business Suite system and SAP HANA.
  - If the SAP Business Suite system is recovered to a point-in-time, in order to ensure that data in the In-Memory DataStore Objects in SAP HANA are consistent with the SAP Business Suite system, the data in the In-Memory DataStore Objects in SAP HANA should be deleted and a delta initialization should be performed for all DataSources in use in SAP HANA from that SAP Business Suite system.
  - Similarly, if the SAP HANA database is recovered to a point-in-time, in order to ensure that data in the In-Memory DataStore Objects in SAP HANA are consistent with the SAP Business Suite system, the data in the In-Memory DataStore Objects in SAP HANA should be deleted and a delta initialization should be performed for all DataSources in use in SAP HANA from that SAP Business Suite system.

> i Note
>
> Delta initialization typically involves loading historical data, and can be significantly time-consuming. This should be factored into business continuity planning

Related Information

SAP HANA Direct Extractor Connection Implementation Guide
Information on Delta Initialization

3.3.3 Backup and Recovery for the ETL-Based Replication components (SAP Data Services)

For information about backup and recovery for SAP Data Services, observe the following:

- SAP Data Services stores its own metadata in a relational database (currently not the SAP HANA database). This metadata is the definition of the source and target connections, the mappings and transformations, the workflows with load order and so on. Backup this database regularly.
- If the SAP HANA database needs to be recovered, replication can resume without any further user interaction. Data Services provides batch loading for SAP HANA database, so no further user interaction would be needed. The next time a batch load is started, it will be loaded into the database again.

3.4 Security for the SAP HANA System Landscape

You ensure security for your SAP HANA system landscape by ensuring security of the SAP HANA components:

For more information, see:

- *SAP HANA Security Guide*
- *SAP HANA Security Guide - Trigger-Based Replication (SLT)*
- SAP HANA Security Guides - ETL-Based Replication (Data Services)
    - *SAP Data Services Administrator's Guide*
      -> Security
    - *SAP Data Services Administrator's Guide*
      -> User Rights and management

4　SAP liveCache on SAP HANA

You can run SAP liveCache on SAP HANA, lowering total cost of ownership and taking advantage of SAP HANA features.

SAP liveCache is an in-memory object store technology that is used to speed up material planning scenarios in SAP Supply Chain Management (SCM). It is now offered as an add-on for SAP HANA.

There are a number of advantages to running SAP liveCache as part of SAP HANA:

- Only one database needs to be administered, which reduces setup and maintenance
- Backup and recovery is simplified:
  - There is no longer a need for separate backups
  - Inconsistencies that could arise from separate backups can be avoided.
- SAP liveCache can take advantage of existing SAP HANA capabilities like performance and high availability.

Supported Scenarios

The following scenarios are supported:

- Single host scenario: The SAP HANA database runs on a single host including SAP liveCache
- Scale out scenario: SAP liveCache runs on a dedicated, separate slave host in a distributed SAP HANA system.
  In a scale-out system, a standby host must be equipped so that it can take over the operation of any of the failed hosts in the failover group. In the case of a failed host, which was a dedicated liveCache host, with a reduced set of memory or CPU features, it may not be reliably used as a standby host for any hosts in the failover group.

> **i Note**
>
> There is no standalone SAP liveCache based on SAP HANA technology. This means that you cannot replace your standalone SAP liveCache based on MaxDB technology with a standalone SAP liveCache based on SAP HANA technology.

Monitoring SAP liveCache in SAP HANA Studio

You can use both the DBACockpit for SAP HANA and the SAP HANA studio to monitor SAP liveCache. The following system views allow you to monitor run time information in the SAP HANA studio:

- M_LIVECACHE_OMS_VERSIONS
- M_LIVECACHE_PROCEDURE_STATISTICS
- M_LIVECACHE_CONTAINER_STATISTICS
- M_LIVECACHE_SCHEMA_STATISTICS
- M_LIVECACHE_LOCKS

System Copy

For detailed information on the procedure for copying an SAP Supply Chain Management system with SAP liveCache on SAP HANA see the Administrator's Guide for SAP SCM 7.0, Version for SAP HANA.

Related Information

SAP HANA SQL and System Views Reference
SAP HANA Administration Guide
DBACockpit for SAP HANA

Administrator's Guide for SAP SCM 7.0, Version for SAP HANA

5   Important Disclaimers on Legal Aspects

This document is for informational purposes only. Its content is subject to change without notice, and SAP does not warrant that it is error-free. SAP MAKES NO WARRANTIES, EXPRESS OR IMPLIED, OR OF MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE.

Coding Samples

Any software coding and/or code lines / strings ("Code") included in this documentation are only examples and are not intended to be used in a productive system environment. The Code is only intended to better explain and visualize the syntax and phrasing rules of certain coding. SAP does not warrant the correctness and completeness of the Code given herein, and SAP shall not be liable for errors or damages caused by the usage of the Code, unless damages were caused by SAP intentionally or by SAP's gross negligence.

Accessibility

The information contained in the SAP documentation represents SAP's current view of accessibility criteria as of the date of publication; it is in no way intended to be a binding guideline on how to ensure accessibility of software products. SAP specifically disclaims any liability with respect to this document and no contractual obligations or commitments are formed either directly or indirectly by this document.

Gender-Neutral Language

As far as possible, SAP documentation is gender neutral. Depending on the context, the reader is addressed directly with "you", or a gender-neutral noun (such as "sales person" or "working days") is used. If when referring to members of both sexes, however, the third-person singular cannot be avoided or a gender-neutral noun does not exist, SAP reserves the right to use the masculine form of the noun and pronoun. This is to ensure that the documentation remains comprehensible.

Internet Hyperlinks

The SAP documentation may contain hyperlinks to the Internet. These hyperlinks are intended to serve as a hint about where to find related information. SAP does not warrant the availability and correctness of this related information or the ability of this information to serve a particular purpose. SAP shall not be liable for any damages caused by the use of related information unless damages have been caused by SAP's gross negligence or willful misconduct. Regarding link classification, see: *http://help.sap.com/disclaimer*

www.sap.com/contactsap

SAP