PUBLIC
SAP HANA Platform 2.0 SPS 08
Document Version: 1.0 – 2025-04-22

SAP HANA Server Installation and Update Guide

Content

1 SAP HANA Server Installation and Update Guide

This guide describes how to install and update an SAP HANA system with the SAP HANA lifecycle management (LCM) tools.

2 SAP HANA Installation and Update Overview

This SAP HANA Installation and Update Guide describes how to install or update an SAP HANA system and its components with the SAP HANA database lifecycle manager (HDBLCM).

The SAP HANA database lifecycle manager is used to install either individual or multiple SAP HANA components in combination with the server. Compared to previous SAP HANA installers, the SAP HANA database lifecycle manager provides an efficient and optimized installation path, allowing you to customize your SAP HANA installation by selecting which components should be installed and updated. It is also now possible to perform installation in interactive graphical or interactive command-line interfaces, as well as configure the SAP HANA installation to be automated, using command line, the configuration file, and batch mode.

It is possible to also perform post-installation configurations tasks, such as renaming the system, adding or removing hosts, and reconfiguring the system, using the SAP HANA database lifecycle manager. For more information, see the *SAP HANA Administration Guide*.

The SAP HANA database lifecycle manager offers three user interfaces: a graphical user interface, a command-line interface and a Web user interface. The Web user interface can be accessed in a standalone Web browser or in the Platform Lifecycle Management view in the SAP HANA studio. Therefore, the SAP HANA database lifecycle manager is capable of performing all actions offered in the SAP HANA platform lifecycle management portfolio and is the only recommended tool for these actions.

Before starting the installation of SAP HANA, make sure that you have reviewed the *SAP HANA Master Guide*.

Related Information

SAP HANA Administration Guide
SAP HANA Master Guide

2.1 SAP HANA Platform Software Components

SAP HANA platform is composed of several components.

For information about the capabilities available for your license and installation scenario, refer to the Feature Scope Description for SAP HANA.

The **SAP HANA platform edition** is the technical foundation of the SAP HANA platform and various SAP HANA editions. The *SAP HANA platform edition* comprises among others:

- SAP HANA Database
- SAP HANA Client

- SAP HANA Studio
- SAP HANA XS advanced runtime
- SAP HANA XS Engine
- SAP HANA Advanced Data Processing
- SAP HANA Spatial
- SAP HANA Graph
- SAP HANA Local Secure Store

> ⓘ Note
>
> For information about the feature capability of the SAP HANA server on Intel-based hardware platforms or on IBM Power servers, see *SAP HANA Hardware and Software Requirements* in the *SAP HANA Master Guide*.

Related Information

SAP HANA Hardware and Software Requirements [page 15]

2.2 Software Download

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

In the SAP Software Download Center, you find media required to install a new SAP HANA system or to upgrade an existing one. Please note that all SAP HANA media on SAP Software Download Center are self-contained full installation media. This applies to the media available in the section **Installations & Upgrades** as well as to the media available in section **Support Packages & Patches**. The section **Installations & Upgrades** only contains media for the very first revision of a Support Package Stack (SPS). The section **Support Packages & Patches** only contains the latest revision of an SPS.

We strongly recommend using the most recent revision of an SPS to avoid running into issues, which are already known and fixed. Therefore, by default, download media for all components of SAP HANA for an installation or upgrade from the section **Support Packages & Patches**.

> ⚠ Caution
>
> It is essential to keep a copy of all installation media which may be required, for example, for disaster recovery purposes. Installation media, which has been superseded is routinely removed from the Software Center and there is no guarantee that all database revisions will be permanently available.

Installation Media and Components for SAP HANA

Installation Media for an SAP HANA SPS

1. Open the *SAP Software Downloads*.
2. Go to *INSTALLATIONS & UPGRADES*, if not already chosen.
3. Open *> By Alphabetical Index (A-Z)*.
4. Choose *H*.
5. Choose *SAP HANA PLATFORM EDITION*.
6. Go to *DOWNLOADS*, if not already opened.
7. Choose *SAP HANA PLATFORM EDITION 2.0*.
8. Open *DOWNLOADS*, if not already opened.
9. Choose *INSTALLATION*.
10. Download the items you need.

> ⓘ Note
>
> The items you have downloaded must be available on the host where the SAP HANA system will be installed or already is installed.

Support Packages and Patches for SAP HANA

1. Open the *SAP Software Downloads*.
2. Choose *SUPPORT PACKAGES & PATCHES*, if not already chosen.
3. Open *> By Alphabetical Index (A-Z)*.
4. Choose *H*.
5. Choose *SAP HANA PLATFORM EDITION*.
6. Choose *DOWNLOADS*, if not already chosen.
7. Choose *SAP HANA PLATFORM EDITION 2.0*.
8. Choose *DOWNLOADS*, if not already chosen.
9. Open the required component and download the items you need.

> ⓘ Note
>
> The items you have downloaded must be available on the host where the SAP HANA system will be installed or is already installed.

Responsibilities

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

- If a customer chooses the **SAP HANA Tailored Datacenter Integration**, then the components of SAP HANA must be installed on validated hardware by a certified administrator or official SAP HANA hardware partner.
- If a customer chooses an **SAP HANA appliance**, then 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 for the product version. Support Package Stacks (SPS) can be downloaded and applied to appliances in accordance with agreements with the respective hardware partner.

> **Note**
>
> The SAP HANA Client is also available separately under the terms of the SAP Developer License Agreement, from the SAP Development Tools public download site. The Developer Agreement licenses the software "as is", without warranty or any other obligations on the part of SAP. Other than what is made available on the SAP Community Website (SCN) by SAP at its sole discretion and by SCN members, SAP does not offer support for software that is the subject of the Developer Agreement.

Related Information

SAP Software Downloads
SAP Developer License Agreement
SAP Development Tools

2.3 Software Authenticity Verification

Before the installation and update of SAP HANA software components, the authenticity and integrity of the software should be verified.

Installing an SAP HANA System

An SAP HANA system can be installed using the SAP HANA database lifecycle manager (HDBLCM). Since the installation software is downloaded from outside your network, it cannot be trusted. Therefore, you should first make sure that the components are authentic, before starting the SAP HANA database lifecycle manager (HDBLCM).

To verify the authenticity of a SAR archive, use the following command:

```
/usr/sap/hostctrl/exe/SAPCAR -dVf <archive name> /usr/sap/hostctrl/exe/
libsapcrypto.so
```

To verify the signature the additional components, run `hdblcm` with the parameter `verify_signature`. For more information, see *SAP Note 2577617*.

Updating an SAP HANA System

An SAP HANA system can be updated using the SAP HANA resident HDBLCM.

The SAP HANA resident HDBLCM is part of the system and is therefore considered secure. It validates the authenticity of the database installation kit and delegates the execution to the SAP HANA database lifecycle manager (HDBLCM) that is located inside the database installation kit. The SAP HANA database lifecycle manager (HDBLCM) validates the authenticity of the selected additional components.

When updating SAP HANA using the SAP HANA resident HDBLCM, the software is verified for authenticity by default. In case you need to perform an update without verifying the authenticity, run `hdblcm` with the parameter `verify_signature=off`.

Installing or Updating SAP HANA Components

SAP HANA components can be installed and updated using the SAP HANA resident HDBLCM or the SAP HANA database lifecycle manager (HDBLCM) shipped with the SAP HANA database installation kit. The latter is not recommended as the database installation kit and its version must match the version of the installed SAP HANA system.

When using the SAP HANA resident HDBLCM to install or update components, the authenticity verification is only enabled by default if the SAP HANA database was installed or updated with an authentic signature. The signature file is located under `<inst_path>/<SID>/hdblcm/SIGNATURE.SMF`). If you are not sure whether the SAP HANA system was installed with a valid signature, you can run the SAP HANA resident HDBLCM with the parameter `verify_signature`.

Related Information

verify_signature [page 287]
SAP Note 2577617
SAP Note 2178665

3 Concepts and Requirements for an SAP HANA System

Before installing or updating an SAP HANA system, it is important to understand the basic system concepts and SAP HANA database lifecycle manager (HDBLCM) features in order to optimize the installation or update process and avoid unnecessary reconfiguration.

An SAP HANA system is made up of the SAP HANA server and its components. The system can be installed on one or multiple system hosts, which are configured to operate as worker or standby hosts. As of SAP HANA 2.0 Support Package Stack (SPS) 01, the SAP HANA system can only be installed or configured to be a tenant database system. A tenant database system contains one system database and can contain multiple tenant databases. A single-container system will be automatically converted to a tenant database system during the upgrade to SPS 01.

The SAP HANA database lifecycle manager offers three user interfaces: graphical, command-line, and Web. System installation can be performed using the graphical user or command-line interface. System update, or component installation or update can be performed using any of the three user interfaces.

The SAP HANA database lifecycle manager can be run interactively, requiring step-by-step input, or it can be run in batch mode, requiring no subsequent input. Defining installation and update parameters can be entered interactively, in a configuration file, or in combination with the call to the program on the command line.

> ⓘ Note
>
> The components of SAP HANA can only be installed by certified hardware partners, or any person holding the required certification, on validated hardware running an approved operating system.

Related Information

Recent changes in the SAP HANA Technology certification program 2016
Updating a Single-Container System [page 158]

3.1 SAP HANA Hardware and Software Requirements

A number of hardware and software requirements apply to SAP HANA.

> ⓘ Note
>
> You can find a complete list of all SAP HANA components and the corresponding SAP HANA hardware and software requirements in the Product Availability Matrix (PAM), in the SAP HANA Hardware Directory, and in the SAP Community Network.

Software Requirements

> **Note**
>
> Only software installed by certified hardware partners, or any person holding certification, is recommended for use on the SAP HANA system. Do not install any other software on the SAP HANA system. The components of SAP HANA can only be installed by certified hardware partners, or any person holding certification. Furthermore, it must be installed on validated hardware running an approved operating system.
>
> For more information, see *C_DBADM - SAP Certified Associate - Database Administrator - SAP HANA* and *The Future Has Been Written – the Rise of the Database Administrator* in the *Related Information* section.

End of Maintenance

The software component version SAPUI5 TOOLS IDE PLUGIN 3.00 is out of maintenance. SAPUI5 TOOLS IDE PLUGIN 3.00 is assigned to the activated instance HANA Studio SAPUI5 Plug-in of SAP HANA PLATFORM EDITION 2.0 SPS05.

For alternatives for SAPUI5 TOOLS IDE PLUGIN 3.00, see SAPUI5 Tools for Eclipse – Now is the Time to Look for Alternatives.

Supported Hardware Platforms

SAP HANA is available for:

- Intel-based hardware platforms
- IBM Power Systems

> **Note**
>
> You can perform a system copy of an SAP system with SAP HANA database as the source database and also SAP HANA database as the target database. This is relevant if you want to change the hardware platform on the SAP HANA system.
>
> For more information, see the *SAP NetWeaver Documentation* on **System Copy** for your SAP NetWeaver release.

The following features are **supported on Intel-based hardware platforms only**:

- SAP HANA Accelerator for SAP ASE

For detailed information about the supported hardware, see *Related Information* in *On-Premise* in the *SAP HANA Master Guide*.

Supported Operating Systems for SAP HANA

For information about supported operating systems for SAP HANA, see *SAP Note 2235581 - SAP HANA: Supported Operating Systems*.

Hardware Requirements

The supported hardware for SAP HANA depends on the deployment method (appliance or TDI). For more information, see the *Related Information* in this section and in *On-Premise* in the *SAP HANA Master Guide*.

Network Time Protocol (NTP)

We strongly recommend setting up an NTP server for the SAP HANA system landscape.

> **Note**
>
> If an NTP sever is not available, this means, for example, that trace files from distributed hosts cannot be displayed in the correct chronological order.

Hardware Requirements for SAP HANA Network Connection

For information about hardware requirements for SAP HANA network connections, see *SAP HANA Network Requirements*.

For installations on IBM Power Servers, Ethernet virtualization using dual VIOS is normally deployed. Natively attached Ethernet cards can also be used however.

Related Information

SUSE Linux Enterprise Server (SLES)

SAP Note 1944799

SAP Note 2205917

SAP Note 1984787

Red Hat Enterprise Linux (RHEL)

SAP Note 2009879

SAP Note 2292690

Supported Hardware Platforms

SAP HANA Hardware - Certified Appliances

On-Premise Deployment

SAP Note 2493172

SAP Note 2055470

SAP Note 2218464

SAP Note 2188482

General Links

SAP HANA Tailored Data Center Integration (TDI) Overview

C_DBADM - SAP Certified Associate - Database Administrator - SAP HANA

The Future Has Been Written – the Rise of the Database Administrator

SAP Note 52505

SAP Note 2235581

Product Availability Matrix

SAP HANA Network Requirements

SAP Note 1900823

SAP Note 2618154

3.2 Recommended File System Layout

The SAP HANA database lifecycle manager (HDBLCM) requires certain file systems in order to successfully install an SAP HANA system.

The file systems must be created and mounted manually before installation. The SAP HANA database lifecycle manager (HDBLCM) will create all required subdirectories during installation. It is important that the file systems listed below are available and have the recommended disk space before starting the SAP HANA database lifecycle manager.

Default File Systems

File System	Default Path	Recommendations
Root	/	

File System	Default Path	Recommendations
Installation path (`sapmnt`)	`/hana/shared`	The mount directory is used for shared files between all hosts in an SAP HANA system. This directory needs to be accessible to each of the servers in the SAP HANA cluster. Subdirectories: • `/hana/shared/<SID>` - Contains executable programs (`exe`), globally shared data (`global`), instance profiles (`profile`), and SAP HANA configuration files. • `/hana/shared/<SID>/hdbclient` - The SAP HANA client installation path. The following paths apply when the SAP HANA studio is installed: • `/hana/shared/<SID>/hdbstudio_update` - The studio repository installation path. The studio repository is used to update the local SAP HANA studio installation. • `/hana/shared/<SID>/hdbstudio` - The SAP HANA studio installation path. The following paths apply when the SAP HANA Local Secure Store is installed: • `/hana/shared/<SID>/lss` - Link to the SAP HANA Local Secure Store installation path.
System instance	`/usr/sap`	This is the path to the local SAP system instance directories. Subdirectories: • `/usr/sap/hostctrl` - The SAP host agent directory. • `/usr/sap/<SID>` - Contains symbolic links to system executables (`SYS`) and the home directory (`home`) of the `<sid>adm` user. **Note** `/usr/sap/<SID>` must not be a mount point. This directory is exclusively managed by the SAP HANA database lifecycle manager (HDBLCM). Files must not be saved or mounted inside this directory as they may be deleted.
Data volume	`/hana/data`	The default path to the data directory is `/hana/data/<SID>`.
Log volume	`/hana/log`	The default path to the log directory is `/hana/log/<SID>`.

File System	Default Path	Recommendations
LSS installation path (`lss_inst_path`)	`/lss/shared`	The mount directory is used for SAP HANA Local Secure Store files shared by all hosts in an SAP HANA system. This directory needs to be accessible to each of the servers in the SAP HANA cluster. It is only applicable for SAP HANA systems with SAP HANA Local Secure Store. For more information, see *Local Secure Store (LSS)* and *lss_inst_path*. Subdirectories: • `/lss/shared/<SID>` - Contains SAP HANA Local Secure Store executable programs and data.

It is strongly recommended to use the SAP HANA file system layout shown in the figure below:

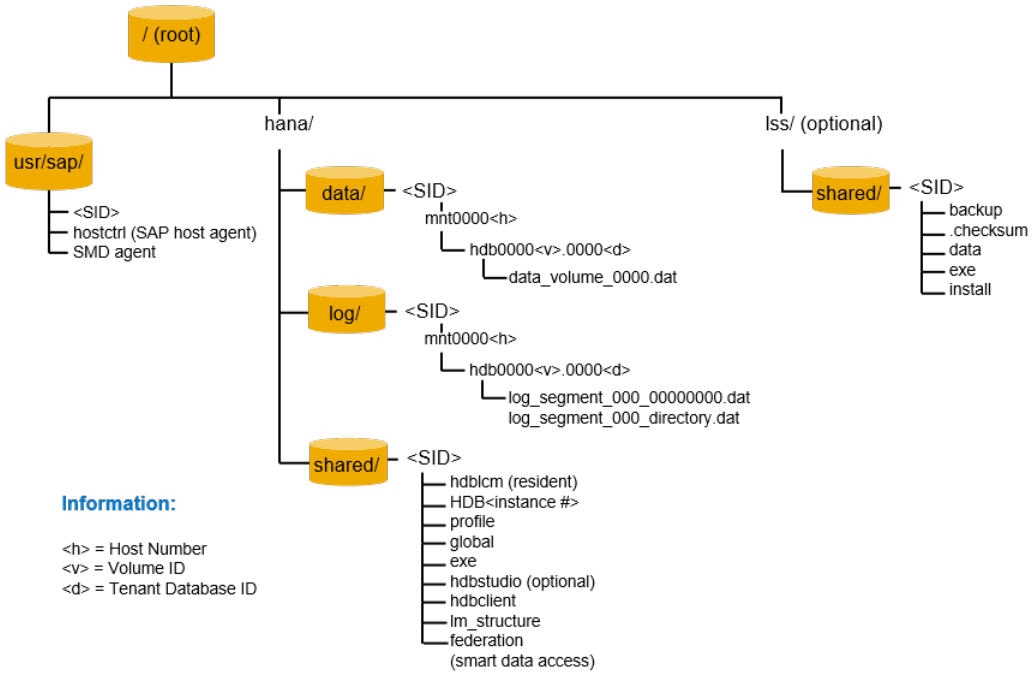

Note

SAP does support running multiple SAP HANA systems (SIDs) on a single production SAP HANA hardware installation. This is restricted to single host / scale-up scenarios only. Please keep in mind that multi-SID requires significant attention to various detailed tasks related to system administration and performance management.

SAP does support running multiple SAP HANA systems (SIDs) on a single non-production (DEV, QA, test, production fail-over, etc) single node or scale-out SAP HANA environment. For more information, see *SAP Note 1681092*.

The installation path, data path, and log path must not point to the same directory.

A SAP HANA system cannot be installed under `/usr/sap/<SID>`. The directory `/usr/sap` must not be shared across other hosts.

`/usr/sap/<SID>` is exclusively managed by the SAP HANA database lifecycle manager (HDBLCM). Files must not be saved or mounted inside this directory as they may be deleted.

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 must become exclusively available for the production system and no longer shared. For more details, refer to the high availability information in the *SAP HANA Administration Guide*.

Directory names inside the data path and log path will not show the database ID for single-container systems that were converted to tenant database systems.

`/hana/shared/<SID>` must be mounted with the option `suid`.

The file systems `/hana/data/<SID>` and `/hana/log/<SID>` may use shared file systems like NFS, or block storage using the SAP HANA storage connector API with non-shared file systems. For more details, see Related Information.

The installation path (`/hana/shared`) is visible on all hosts. By default, the installation path is also used for backup. However, backup directories should be manually configured, and must belong to a shared file system. For more details, refer to the backup information in the *SAP HANA Administration Guide*.

Related Information

Multiple-Host System Concepts [page 85]
SAP Note 405827
SAP Note 1681092
db_isolation [page 250]
Database Isolation [page 37]
Local Secure Store (LSS) [page 28]
lss_inst_path [page 264]

3.3 SAP HANA System Concepts

An SAP HANA system is composed of three main components: the host, the system, and the instance.

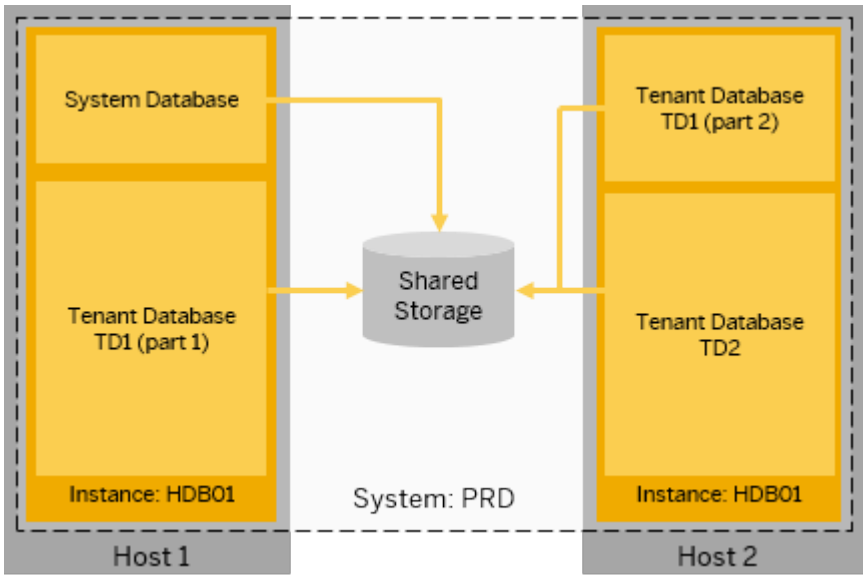

Host

A host is the operating environment in which the SAP HANA database runs. The host provides all the resources and services (CPU, memory, network, and operating system) that the SAP HANA database requires. The storage for an installation does not have to be on the host. For multiple-host systems, a shared storage or a storage that is accessible on-demand from all hosts is required.

For more information about the restrictions that apply to host names in SAP systems, see SAP Note 611361 in Related Information.

System

A system is one or more instances with the same instance number. If a system has more than one instance, they must be dispersed over several hosts as a multiple-host system. Every system has a unique SAP system ID (SID).

Instance (HDB)

An SAP HANA instance (HDB) is the smallest operational unit on a host. It is the set of SAP HANA system components that are installed on one host. A single-host system contains one instance on the one host. A multiple-host system contains several instances distributed across the multiple hosts (one per host). Every system has an instance number, which is a two-digit identifier. Each instance in a multiple-host system must have the same instance number.

System Database

The system database stores overall system landscape information, including knowledge of 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.

Tenant Databases

SAP HANA supports multiple isolated databases in a single SAP HANA system. These are referred to as tenant databases.

Related Information

SAP Note 611361

3.4 SAP HANA System Types

An SAP HANA system can be configured as a single-host or multiple-host system using the SAP HANA database lifecycle manager.

The SAP HANA system type definitions are as follows:

- Single-host system - One SAP HANA instance on one host.
- Multiple-host (distributed) system - Multiple SAP HANA instances distributed over multiple hosts, with one instance per host.

A single-host system is the simplest system installation type. It is possible to run an SAP HANA system entirely on one host and then scale the system up as needed. The SAP HANA database lifecycle manager can be used to install an SAP HANA single-host system in one of the program interfaces, and with a combination of parameter specification methods.

The following graphic shows the file system for a single-host installation:

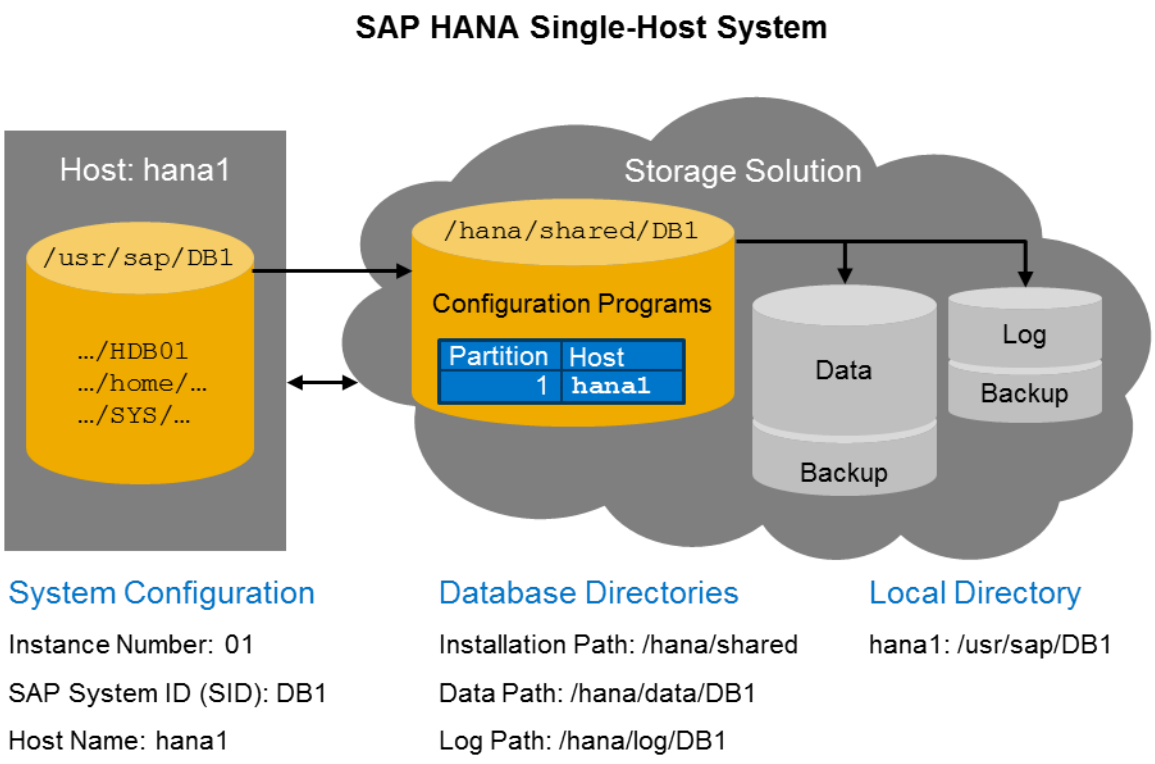

A multiple-host system is a system with more than one host, which can be configured as active worker hosts or idle standby hosts. The SAP HANA database lifecycle manager can be used to install an SAP HANA multiple-host system in one of the program interfaces, and with a combination of parameter specification methods. To add hosts to an existing system, use the SAP HANA resident HDBLCM. For more information about installing a multiple-host system, see the Related Information.

The following graphic shows the file system for a multiple-host system installed on a shared file system with three hosts:

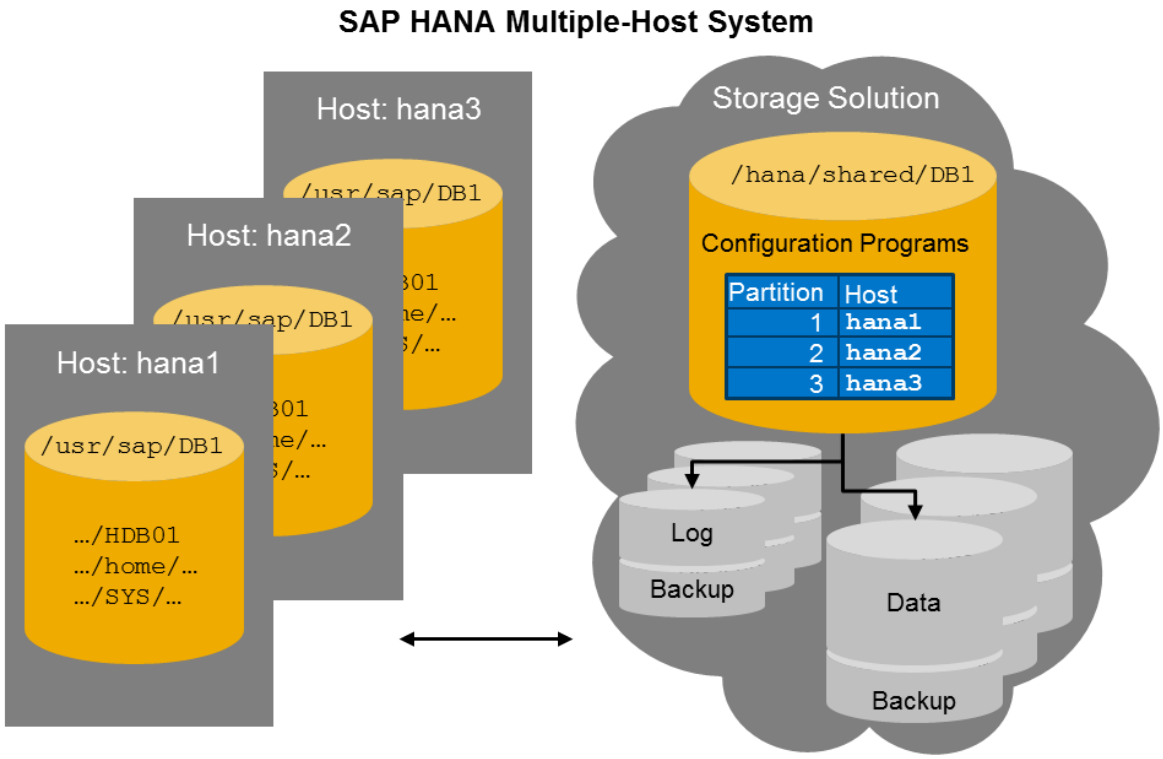

The server software is based on a flexible architecture that enables a distributed installation. This means that load can be balanced between different hosts. The server software has to be installed in a shared file system. This file system has to be mounted by all hosts that are part of the system.

Related Information

Installing a Multiple-Host System [page 84]

3.5 SAP HANA Deployment Types

SAP HANA can be deployed in a number of configurations that are approved in varying degrees for production environments (or not approved for production at all).

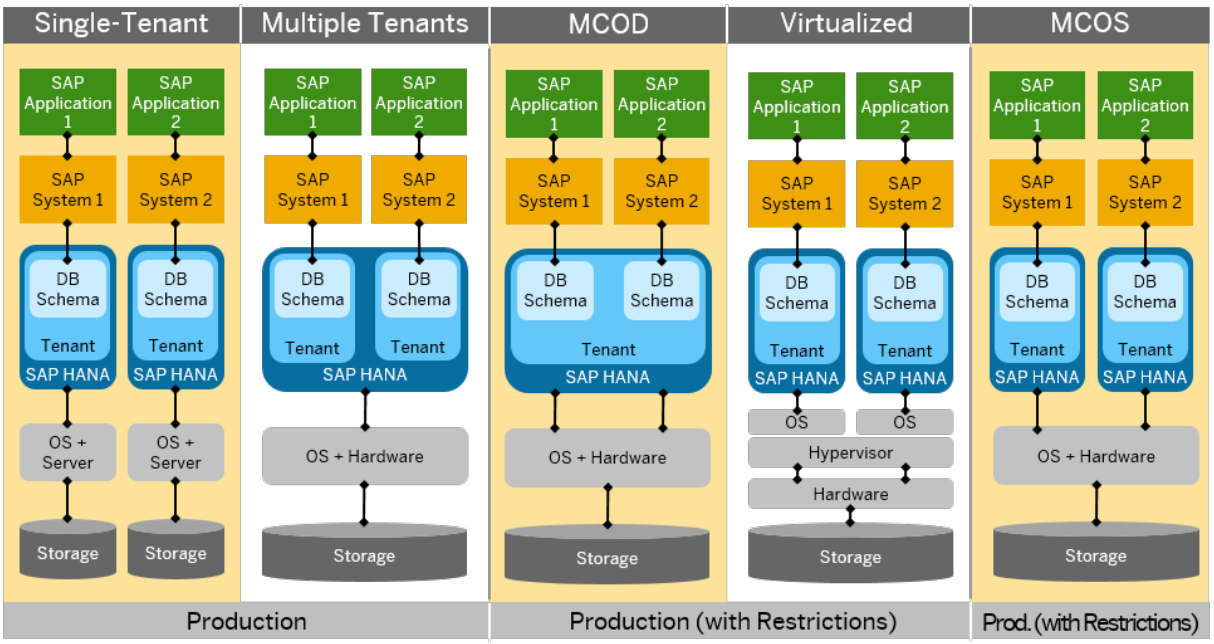

The server installation documentation is mainly written for SAP HANA deployment types which are completely approved for production environments, that is to say, an SAP HANA system running on dedicated hardware, or tenant database systems.

Tenant Databases

The tenant database deployment type makes it possible to run several SAP HANA instances on the same hardware in a production environment. It provides an alternative to a virtualized deployment, which is only production approved in some scenarios, and the MCOS (Multiple Components One System) deployment, which is approved for production environments on a single production SAP HANA hardware installation. This is restricted to single host / scale-up scenarios only. Please keep in mind that multi-SID requires significant attention to various detailed tasks related to system administration and performance management. For more information, see *SAP Note 1681092*.

An SAP HANA system is installed with support for tenant databases by default using the SAP HANA database lifecycle manager (HDBLCM). You can specify a database isolation type during installation setting the optional parameter `db_isolation`. High isolation can be configured to protect against unauthorized access at the operating-system level by separating system administrator users and providing authenticated communication within databases. By default, the isolation type is set to low isolation. For more information, see *Database Isolation* in the *SAP HANA Administration Guide* or the *SAP HANA Security Guide* and *db_isolation* in Related Information.

Multiple Components on One Database (MCOD)

MCOD deployments are characterised by multiple applications on one SAP HANA system. SAP supports deploying and running multiple applications on a single SAP HANA production database **only** for packaged applications and scenarios listed on the "White List" included in SAP Note 1661202. If a particular packaged

application or scenario is not on the "White List", then it is not supported to run together on the same SAP HANA database with any other packaged application or scenario. For more information, see SAP Note 1661202 in Related Information.

Virtualized

SAP HANA systems can be run on virtual machines with restrictions to the hypervisor (including logical partitions). For more information about running SAP HANA virtualized, see SAP Note 1788665, 2230704 and 2024433 in Related Information.

Multiple Components on One System (MCOS)

MCOS deployments are characterized by multiple SAP HANA systems on one host. This configuration is approved for production environments as of SAP HANA Support Package Stack (SPS) 09. This is restricted to single host / scale-up scenarios only. Please keep in mind that multi-SID requires significant attention to various detailed tasks related to system administration and performance management. For more information about running SAP HANA virtualized, see SAP Note 1681092 in Related Information.

It is approved for production environments for SAP HANA systems to share hardware between the SAP HANA server and SAP HANA options. As of SPS 10, you have the option to install SAP HANA systems with multiple host roles - including database server roles and SAP HANA option host roles - on one host, or give an existing SAP HANA host additional roles during system update. For more information about configuring additional host roles, see *add_local_roles* and *add_roles* in Related Information.

> ⚠ Caution
>
> For information about the capabilities available for your license and installation scenario, refer to the Feature Scope Description for SAP HANA.

Related Information

Converting an SAP HANA System to Support Tenant Databases
Database Isolation [page 37]
SAP Note 1661202
SAP Note 1681092
SAP Note 1788665
SAP Note 2230704
SAP Note 2024433
add_local_roles [page 235]
add_roles [page 236]
db_isolation [page 250]

3.6 SAP HANA and Virtualization

SAP HANA is supported on bare-metal and virtualized platforms.

> **Note**
>
> Some of the virtualization platforms that are available for SAP HANA are only supported in non-production environments. For more information on supported versions, see *SAP Note 1788665 - SAP HANA Support for virtualized / partitioned (multi-tenant) environments*.

VMware vSphere

SAP HANA is supported on VMware vSphere for scale-up, multi-VM and scale-out production environments, including SAP HANA Tailored Datacenter Integration deployments. For more information, see *Best Practices and Recommendations for Scale-up Deployments of SAP HANA on VMware vSphere* and *Best Practices and Recommendations for Scale-Out Deployments of SAP HANA on VMware vSphere*.

Hitachi LPAR

Hitachi Unified Compute Platform for the SAP HANA Platform with logical partitioning (LPAR) in a scale-up or multi-VM configuration is a pre-configured virtual appliance ready to plug into a network to provide real-time access to operational data for use in analytic models. SAP HANA on Hitachi solutions with logical partitioning are based on a number of bare metal appliance configurations with modifications that are required for running logical partitions (LPARs) in dedicated mode. For more information, see *Hitachi Unified Compute Platform for the SAP HANA Platform with Logical Partitioning in a Scale-Up Configuration Using Hitachi Compute Blade 2500 and Hitachi Virtual Storage Platform*.

Huawei FusionSphere

SAP HANA is supported for production and non-production use in a virtualized environment using Huawei FusionSphere. For more information, see *Best Practices for Deploying SAP HANA on Huawei FusionSphere Virtualization Platform* and *Guide for Deploying SAP HANA on Huawei FusionSphere Virtualization*.

Red Hat Enterprise Virtualization (RHV)

For information on supported configurations of SAP HANA in a virtualized environment using Red Hat Enterprise Virtualization (RHV), see *Best Practices / Deployment Guide SAP HANA on Red Hat Virtualization* and *SAP Note 2599726 - SAP HANA on Red Hat Virtualization*.

SUSE Linux Enterprise Hypervisor (KVM)

For information on supported configurations of SAP HANA in a virtualized environment using SUSE Linux Enterprise Hypervisor (KVM), see *Best Practices / Deployment Guide SAP HANA on SUSE Virtualization* and *SAP Note 1788665 - SAP HANA Support for virtualized and partitioned (multi-tenant) environments*.

IBM PowerVM

SAP HANA can be deployed on IBM PowerVM for IBM Power Systems. PowerVM is a combination of hardware, PowerVM Hypervisor, and software, which includes other virtualization features, such as the Virtual I/O Server. For more information, see *IBM PowerVM Best Practices*.

Related Information

SAP Note 1788665
Best Practices and Recommendations for Scale-up Deployments of SAP HANA on VMware vSphere
Best Practices and Recommendations for Scale-Out Deployments of SAP HANA on VMware vSphere
Hitachi Unified Compute Platform for the SAP HANA Platform with Logical Partitioning in a Scale-Up Configuration Using Hitachi Compute Blade 2500 and Hitachi Virtual Storage Platform
Best Practices for Deploying SAP HANA on Huawei FusionSphere Virtualization Platform
Guide for Deploying SAP HANA on Huawei FusionSphere Virtualization
Best Practices / Deployment Guide SAP HANA on Red Hat Virtualization
SAP Note 2599726
Best Practices / Deployment Guide SAP HANA on SUSE Virtualization
SAP Note 3430656
IBM PowerVM Best Practices
SAP HANA with Virtualization

3.7 Local Secure Store (LSS)

The local secure store (LSS) is a separate, lightweight utility for storing and securely managing encryption keys, encryption root keys, and other similarly sensitive data, such as security-relevant configuration settings. It helps protect sensitive server-side data from illegitimate or fraudulent usage.

> **Note**
>
> SAP HANA Local Secure Store will be installed and activated by default during an SAP HANA system installation and update. To opt out of Local Secure Store activation during SAP HANA system installation or update, run `hdblcm` with the parameter `secure_store=ssfs`. For more information on how to activate and use LSS, see *Activate the Local Secure Store (LSS)*.

To opt out of Local Secure Store installation during SAP HANA system installation or update, do not select it as a component.

As of SAP HANA 2.0 Support Package Stack (SPS) 05, the LSS is approved for use in production environments.

⚠ Restriction

The LSS is not supported by SAP HANA Dynamic Tiering. For SAP HANA systems configured with dynamic tiering (with LSS installed and activated), set the *Secure Store* type for them to SSFS (secure store in the file system).

For details, see Local Secure Store (LSS) under *Security Administration* > *Data Encryption* in the *SAP HANA Dynamic Tiering: Administration Guide 2.0 SP04*.

If LSS is installed and activated with an encryption_algorithm other than AES-256-CBC, migration back to SSFS is not possible.

For more information, see the section on the local secure store in the *SAP HANA Security Guide*.

Related Information

Local Secure Store (LSS)
SAP HANA Security Guide
Managing SAP HANA System Components [page 180]
Activate the Local Secure Store (LSS) [page 197]
SAP Note 2935272

4 Overview of SAP HANA Tenant Databases

SAP HANA supports multiple isolated databases in a single SAP HANA system. These are referred to as tenant databases.

An SAP HANA system is capable of containing more than one tenant database.

A system always has exactly one system database, used for central system administration, and any number of tenant databases (including zero). An SAP HANA system 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 cockpit, connect to specific databases.

All the databases 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:

- Set of database users
- Database catalog
- Repository
- Persistence
- Backups
- Traces and logs

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.

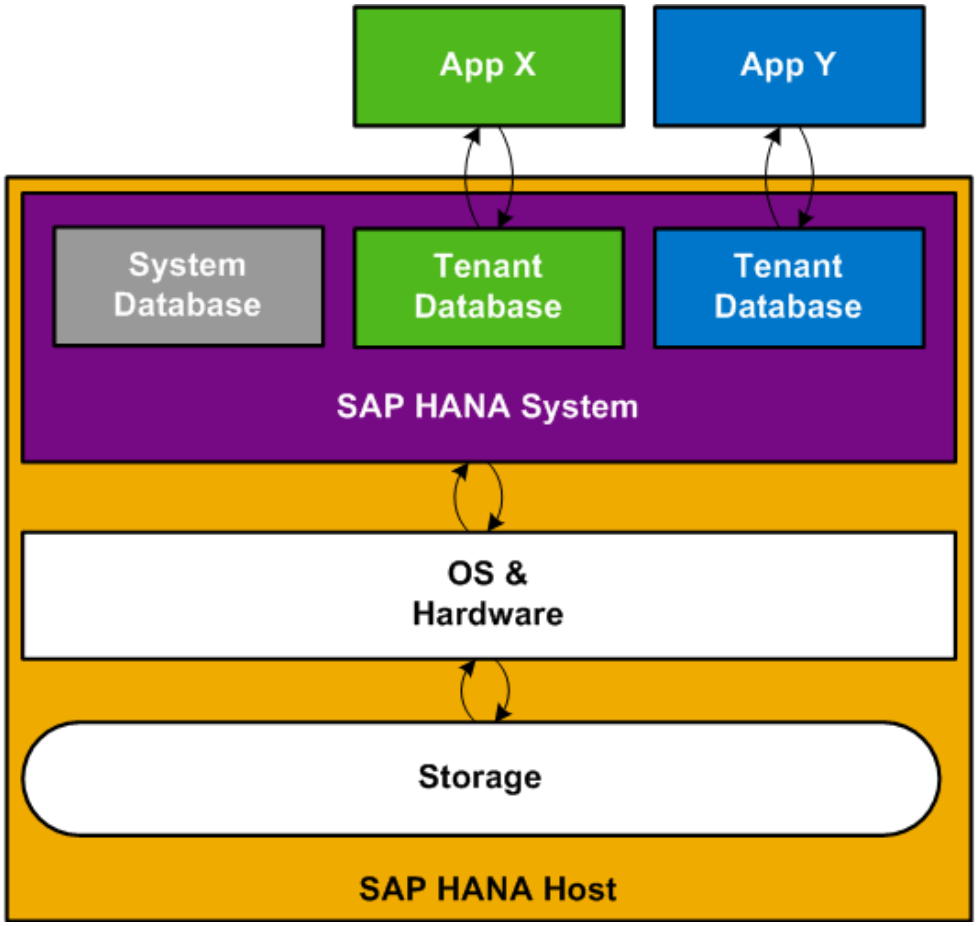

On-premise Deployment Using SAP HANA Tenant Databases

4.1 Server Architecture of Tenant Databases

An SAP HANA database consists of multiple servers, for example, name server, index server, preprocessor server, and so on. The databases in an SAP HANA system run different combinations of these servers. The most important server is the index server. It contains the actual data stores and the engines for processing the data and runs in every tenant database.

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 index server functionality for the system database. The name server does not own information about the location of tables and table partitions in tenant databases. Database-related information is stored in the relevant tenant database catalog.

Tenant databases require only an 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.

> ⓘ Note
>
> For a full list and description of all SAP HANA servers, see *Server Components of the SAP HANA Database*.

The following figure shows a sample system with three databases (system database and three tenant databases) on a single host.

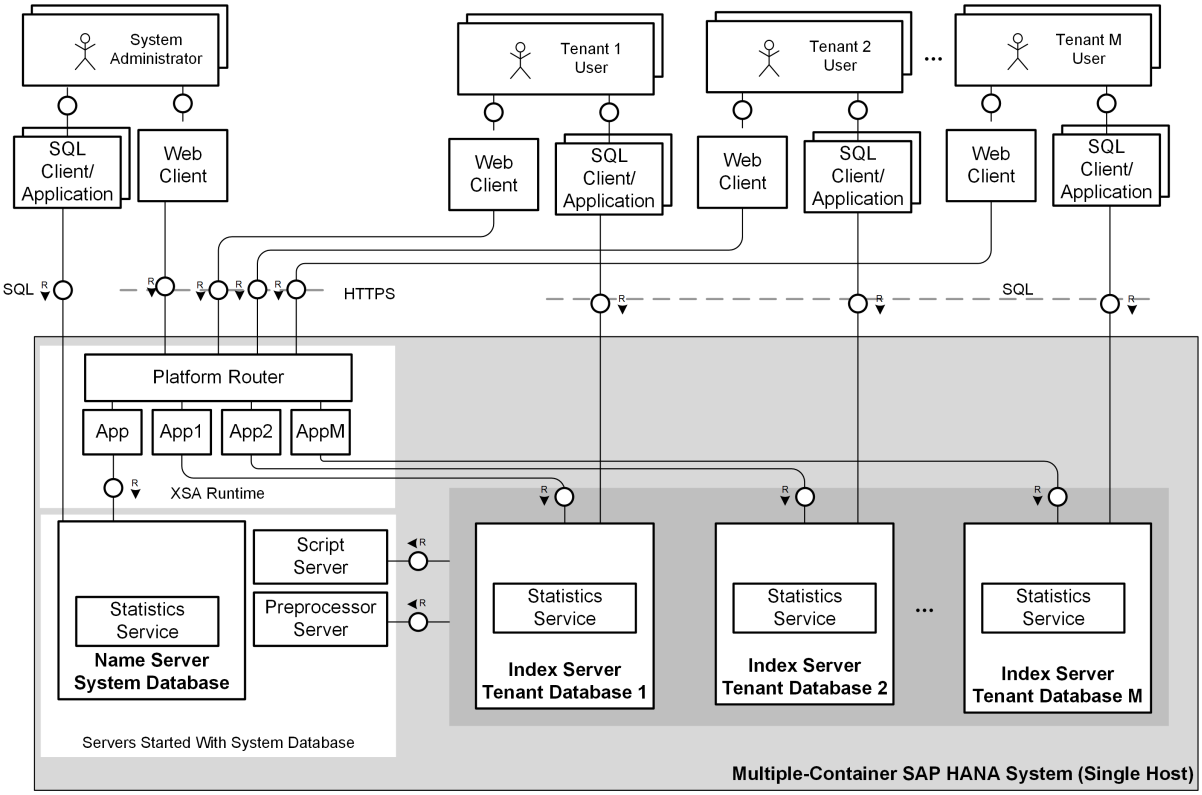

Single-Host SAP HANA System with Tenant Databases

> **Note**
>
> If the SAP HANA XS classic server is available, it runs embedded in the (master) index server of the tenant database by default, although it can be added as a separate service if necessary. The **SAP Web Dispatcher**, which runs as a separate database service on the host of the system database, is used to route incoming HTTP requests from clients to the correct XS classic server based on virtual host names. This is part of network configuration. In addition to the system-internal Web Dispatcher, you can implement an external Web Dispatcher for load distribution. See the section on using the SAP Web Dispatcher for load balancing with tenant databases.

Related Information

Server Components of the SAP HANA Database
Connections from Database Clients and Web Clients to SAP HANA
Port Assignment in Tenant Databases
Scale-Out Architecture of Tenant Databases [page 33]
Using SAP Web Dispatcher for Load Balancing with Tenant Databases

4.2 Scale-Out Architecture of Tenant Databases

Tenant databases can be distributed across several hosts in a multiple-host system.

To ensure system availability, an instance of the system database runs on all hosts (worker and standby) in a single master and multiple workers configuration. 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. Like in a single-host system, the master candidate for a failing host is determined. On that host the system database is restarted, if necessary. Up to three hosts can be configured to act as the master host of a system. These three hosts can be set up in the clients with the database name to be reconnected to a tenant database even in the case of a host auto-failover of the master host with the system database.

The following figure shows a tenant database system with three tenant databases distributed across three hosts. Tenant database DB1 has only one index server on host 1, while DB2 and DB3 are distributed across several hosts. Tenant database DB2, for example, is divided into three database shards, each of them with its own index server on a different host. In this context, a database shard is the union of all tables, partitions and replicas of one database that reside on one index server. Tenant database DB3 consists of two shards, one on host 2 and one on host 3. System administrators can specify the host when they create the tenant database, or they can let SAP HANA chose an appropriate host based on load-balancing algorithms.

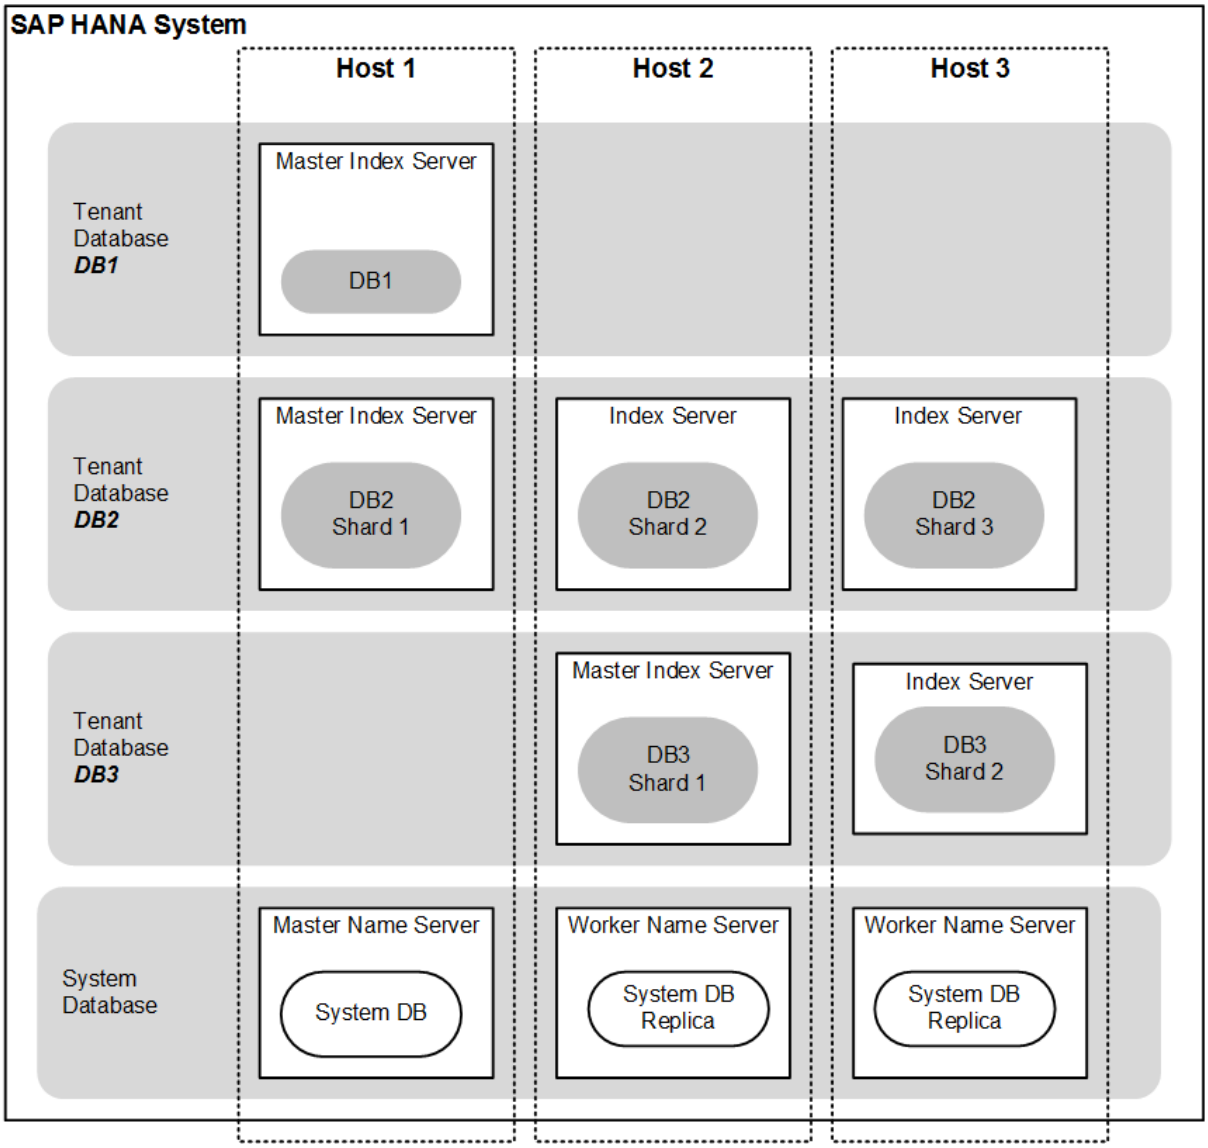

Multiple-Host System with Tenant Databases

Scale-Out Recommendations

When planning your SAP HANA deployment with tenant databases, various options exist with regard to scale-up versus scale-out.

In general, scaling up offers some performance advantages over scaling out, as memory access is local and minor overhead associated with inter-node network communication is avoided.

Note the following with regard to scale-out:

- It is possible to distribute tenant databases across several hosts in a scale-out system.
- The primary reason to distribute tenant databases generally is when their size is larger than the capacity of a single host. However, other reasons for distributing tenant database may exist, for example, a large SAP Business Warehouse (BW) system requires a scale-out configuration in accordance with its sizing rules.
- If tenant databases are distributed in a scale-out configuration due to sizing requirements, caution is advised when deploying additional tenant databases on the same host as a distributed tenant database

shard. The rationale is this: Workload in distributed scenarios can be somewhat volatile and less predictable. Therefore in many cases, it can be advantageous to dedicate maximum resources of the host to the distributed tenant database shard in order to maintain expected performance.

- In certain cases, more than one distributed tenant database shard may share the same host. In these cases, in order to dedicate maximum resources for a master node (for performance reasons), it is advisable to avoid deploying other tenant databases on the master node. For example, the following deployment should offer performance advantages:
  - Host 1: Master for tenant database 1
  - Host 2: Worker for tenant database 1 and worker for tenant database 2
  - Host 3: Master for tenant database 2
  - Host 4: Standby host for failover

Related Information

Scaling SAP HANA

4.3 The System Database

The system database is created during either installation or conversion from a single-container system to a tenant database 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 system has exactly one system database. It contains the data and users for system administration. System administration tools, such as the SAP HANA cockpit, can connect to this database. The system database stores overall system landscape information, including knowledge of 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 all of its databases (for example, system-level configuration settings), or can target specific tenant databases (for example, backup of a tenant database). For more information, see *Administration of Tenant Databases*.

Things to Remember About the System Database

- The system database does not have the same functionality as a tenant database.
- The system database is not a database with full SQL support.
- The system database cannot be distributed across multiple hosts, in other words, scale-out is not possible.
- If you need a full-featured SAP HANA database, you always have to create at least one tenant database.
- The system database does not support Application Function Libraries (AFL) and SAP liveCache applications.

- Cross-database access between the system database and a tenant database is not possible. The system database can show monitoring data from tenant databases (views in the schema SYS_DATABASES) but can never show actual content from tenant databases.
- The system database cannot be copied or moved to another host.
- SAP HANA options can only run in tenant databases.
- Tenant-specific configurations cannot be set in the system database. Only global settings are allowed.
- Features can only be restricted or disabled at high level for tenant databases.

Related Information

Administration of Tenant Databases [page 39]
Memory and CPU Usage for Tenant Databases
Cross-Database Authorization in Tenant Databases
Restricted Features in Tenant Databases

4.4 Cross-Database Access

Read-only queries between tenant databases in the same SAP HANA system are possible. This supports cross-application reporting. Cross-database access must be explicitly enabled.

Every tenant database is self-contained with its own isolated set of database users and isolated database catalog. However, to support for example 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.

The following object types on remote databases can be accessed using cross-database access:

- Schemas
- Rowstore and columnstore tables (not including virtual tables)
- SQL views (not including monitoring views)
- Graphical calculation views
  - If they only use supported object types as data sources
  - If they don't use procedure-based analytic privileges
- Synonyms

The following object types on the local tenant database can access database objects on the remote tenant database:

- SQL views
- Scripted and graphical calculation views
- Procedures
- Synonyms

The SAP HANA modeler supports modeling of graphical calculation views using tables and other graphical calculation views as data sources from different tenant databases. For more information, see *Modeling Graphical Calculation Views With Tenant Databases* in the *SAP HANA Modeling Guide (For SAP HANA Studio)*.

For more information about how to enable and configure cross-database access, see *Enable and Configure Cross-Database Access*.

Related Information

Enable and Configure Cross-Database Access
Cross-Database Authorization in Tenant Databases (SAP HANA Security Guide)
Troubleshooting Error Situations Related to Cross-Database Access
Workload Management and Cross-Database Queries
Modeling Graphical Calculation Views With Tenant Databases (SAP HANA Modeling Guide)
Import/Export Catalog Objects with Dependencies for Multi-TenantDB (SAP Community Blog)

4.5 Database Isolation

Every tenant database is self-contained and isolated in terms of users, database catalog, repository, logs, and so on. However, to protect against unauthorized access at the operating system (OS) level, it's possible to increase isolation further through OS user separation and authenticated communication within databases.

OS User Separation

By default, all database processes run under the default OS user `<sid>adm`. If it's important to mitigate against cross-database attacks through OS mechanisms, you can 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, instead of all database processes running under `<sid>adm`. Database-specific data on the file system is then protected using standard OS file and directory permissions.

> **Note**
>
> `<sid>adm` is the OS user for the system database.

Authenticated Communication

In addition, once high isolation has been configured, internal database communication is secured using the Transport Layer Security (TLS)/Secure Sockets Layer (SSL) protocol. Certificate-based authentication is used to ensure that only the processes belonging to the same database can communicate with each other. It's also possible to configure internal communication so that all data communication within databases is encrypted.

> **Note**
>
> If cross-database access is enabled, communication between configured tenant databases is allowed.

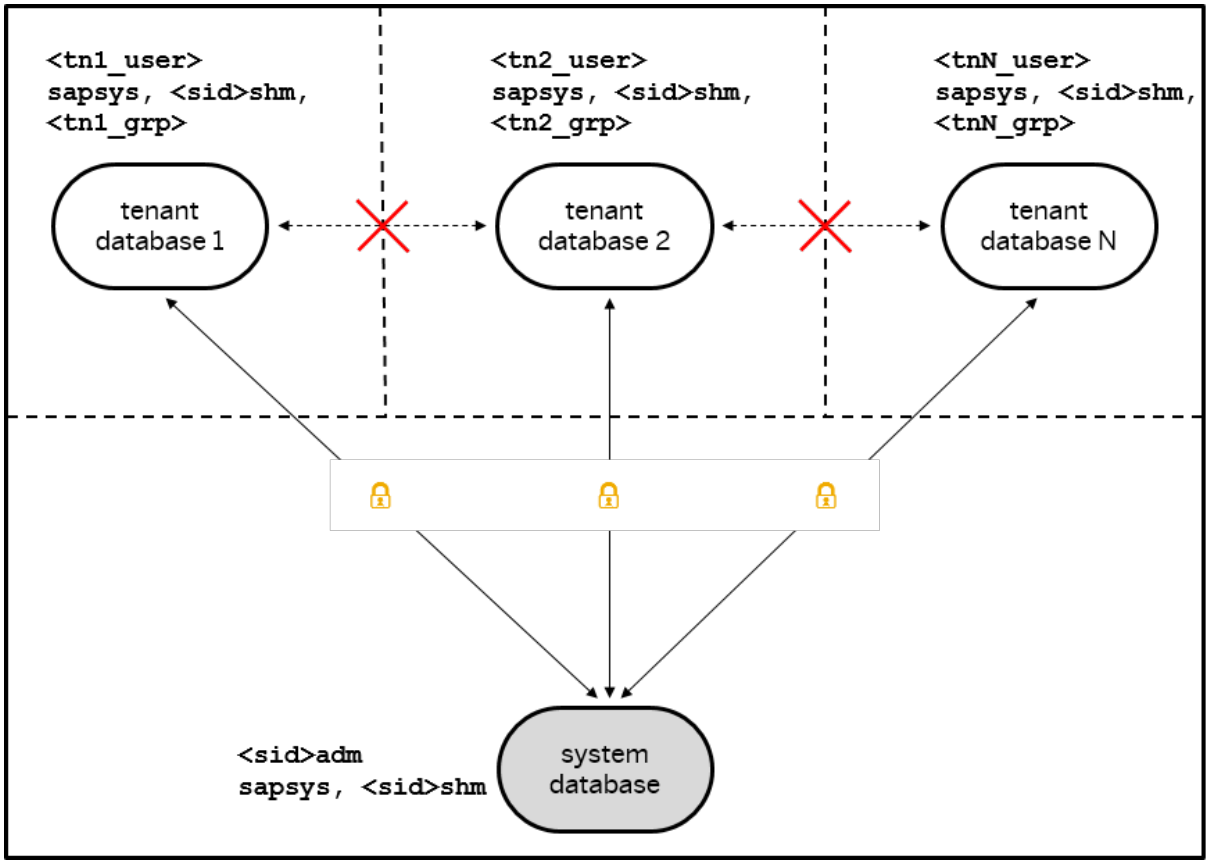

High Database Isolation

Configuration

You can specify the isolation level of the system during installation. The default isolation level is low. It's also possible to change the isolation level of an existing system (from low to high or from high to low) at any time. For more information, see *Increase the System Isolation Level* in the *SAP HANA Administration Guide*. Once high isolation has been configured, a dedicated OS user and group must exist for every tenant database. Otherwise, it's not possible to create or start a tenant database.

Internal database communication is secured with the same mechanism used for securing other internal SAP HANA communication channels. Once high isolation has been configured, authenticated communication within databases is enabled without any change required to the default TLS/SSL configuration for internal communication. However, encryption of data communication may need to be configured explicitly.

Related Information

File and Directory Permissions with High Isolation

Secure Internal Communication

Increase the System Isolation Level
SAP HANA Administration Guide

4.6 Administration of Tenant Databases

In SAP HANA systems there is a distinction between administration tasks performed at system level and those performed at database level.

System Versus Database Administration

Tenant database systems have two levels of administration.

Some administration tasks are performed in the system database and apply globally to the system and all its databases. They include for example:

- Starting and stopping the whole system
- Monitoring the system
- Configuring parameters in configuration (*ini) files at system level
- Setting up and configuring tenant databases, for example:
  - Creating and dropping tenant databases
  - Disabling features on tenant databases
  - Configuring system- and database-specific parameters in configuration (*ini) files
  - Scaling out tenant databases by adding services
- Backing up tenant databases
- Recovering tenant databases

Some administration tasks are performed in the tenant database and apply only to that database. They include for example:

- Monitoring the database
- Provisioning database users
- Creating and deleting schemas, tables, and indexes in the database
- Backing up the database
- Configuring database-specific parameters in configuration (*ini) files

Administration Tools

Several tools are available for the administration of SAP HANA. While all tools support database-level administration, system-level administration of tenant databases requires the SAP HANA cockpit (for example, monitoring availability of tenant databases, creating and deleting tenant databases).

For more information about the SAP HANA cockpit and other administration tools, see the section on administration tools in the *SAP HANA Administration Guide*.

Related Information

Tenant Databases
The System Database [page 35]
Creating and Configuring Tenant Databases
SAP HANA Administration Tools
Monitoring and Managing Tenant Databases

5 Using the SAP HANA Platform LCM Tools

The SAP HANA database lifecycle manager (HDBLCM) is used to perform SAP HANA platform lifecycle management (LCM) tasks, including installing, updating, and configuring an SAP HANA system. The SAP HANA database lifecycle manager is designed to accommodate hardware partners and administrators, and so it offers a variety of usage techniques.

The SAP HANA database lifecycle manager is used by means of program interface type, program interaction mode, and parameter entry mode. Before using the SAP HANA database lifecycle manager, you should choose which user interface you prefer to use and how you want to modify the platform LCM task to achieve your desired result. You modify the actions of the platform LCM tools 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 you need to use and the program interaction mode you choose.

Platform LCM Tools and Program Interaction Modes

	Interactive Mode	Advanced Interactive Mode	Batch Mode
Graphical User Interface	■	■	◇
Command-Line Interface	■	■	■
Web User Interface	■	◇	◇

The first choice to make is which SAP HANA database lifecycle manager (HDBLCM) interface type you prefer to use. The SAP HANA HDBLCM program can be run as a graphical user interface, a command-line interface, or as Web user interface in a Web browser (the Web user interface is not available for all platform LCM tasks).

Once you've chosen the graphical user, command-line, or Web user interface, you can decide if you prefer to interactively enter parameter values, or give all required parameters with the call to the platform LCM tool, and let it run unattended to completion. Interactive mode is available for all user interfaces, and is the default mode for program interaction. To use interactive mode, you simply call the SAP HANA HDBLCM user interface, and enter 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'd like to modify parameter default values which are not requested in interactive mode. Batch mode is an advanced platform LCM interaction method because all required parameters must be provided with the call to the LCM program on the command line. Batch mode is designed for large-scale platform LCM tasks, which would be time consuming to perform interactively.

Platform LCM parameters can be entered interactively (only available for interactive mode or advanced interactive mode), 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, you must enter parameter values either as call options to the SAP HANA database lifecycle manager or from a configuration file. The syntax for the parameters as call options can be found in the *Parameter Reference*. The configuration file is generated as a blank template, then edited, and called as a call option.

Related Information

Use Interactive Mode to Perform Platform LCM Tasks [page 50]
Use Advanced Interactive Mode to Perform Platform LCM Tasks [page 51]
Use Batch Mode to Perform Platform LCM Tasks [page 53]

5.1 Choosing the Correct SAP HANA HDBLCM for Your Task

It is important to distinguish between the version of the SAP HANA database lifecycle manager (HDBLCM) that is available on the installation medium and the version that is unpacked during installation, and subsequently used to perform administration and configuration tasks after the SAP HANA system has been installed.

The SAP HANA database lifecycle manager is available in two varieties - an installation medium version to perform installation and update, and a resident version for update and configuration that is unpacked on the SAP HANA host during installation or update. The SAP HANA resident HDBLCM has been designed to be version-compatible. That means, every time you install or update an SAP HANA system, you can be sure that any subsequent configuration tasks performed with the SAP HANA database lifecycle manager will work as expected because the installation or update tool and the configuration tool are of the same version and have been tested together. The SAP HANA resident HDBLCM is located at `<sapmnt>/<SID>/hdblcm`.

5.2 Performing LCM Tasks by Program Interface

SAP HANA platform lifecycle management tasks can be performed from a graphical, command-line and Web user interface.

Related Information

Use the Graphical User Interface to Perform Platform LCM Tasks [page 43]
Use the Command-Line Interface to Perform Platform LCM Tasks [page 44]
Using the Web User Interface [page 45]

5.2.1 Use the Graphical User Interface to Perform Platform LCM Tasks

SAP HANA platform lifecycle management tasks can be performed from a graphical interface.

Procedure

1. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
Installation Medium (IBM Power Systems)	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`
Installation Archive downloaded from SAP Support Portal (SAP Service Marketplace)	`cd SAP_HANA_DATABASE`
SAP HANA resident HDBLCM	`cd <sapmnt>/<SID>/hdblcm`

   In general, installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM. For more information about the two SAP HANA database lifecycle manager types, see Related Information.

2. Start the SAP HANA platform lifecycle management tool:

   ```
   ./hdblcmgui
   ```

3. Enter parameter values in the requested fields. In addition, you can specify parameter key-value pairs as call options or in the configuration file template.

   > ⓘ Note
   >
   > If parameter key-value pairs are specified as command-line options, they override the corresponding parameters in the configuration file. Parameters in the configuration file override default settings.
   >
   > **Order of parameter precedence:**
   >
   > Command Line > Configuration File > Default

   For more information about program interaction modes and parameter values entry methods, see Related Information.

Related Information

Choosing the Correct SAP HANA HDBLCM for Your Task [page 42]
Entering Platform LCM Parameters as Call Options from the Command Line [page 57]

5.2.2 Use the Command-Line Interface to Perform Platform LCM Tasks

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

Procedure

1. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
Installation Medium (IBM Power Systems)	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`
Installation Archive downloaded from SAP Support Portal (SAP Service Marketplace)	`cd SAP_HANA_DATABASE`
SAP HANA resident HDBLCM	`cd <sapmnt>/<SID>/hdblcm`

   In general, installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM. For more information about the two SAP HANA database lifecycle manager types, see Related Information.

2. Start the SAP HANA platform lifecycle management tool:

```
./hdblcm
```

3. Enter parameter values in one of the following ways.
   - **Interactive parameter entry** - If you call the SAP HANA platform LCM tool only, the program runs in interactive mode. Parameter default values are suggested in brackets, and can be accepted with *Enter*. Otherwise, enter a non-default value, then select *Enter*.
   - **Command-line parameter entry as call options** - If you enter parameter key-value pairs as call options with the call to the SAP HANA platform LCM tool, the program runs in advanced interactive mode and requests values for any parameter values which you didn't specify in the original input. If you entered the batch mode call option, the program runs to completion without any further requests, unless a mandatory parameter was left out of the original input, in which case, the program fails to perform the platform LCM task.

- **Configuration file parameter entry** - If you enter parameter key-value pairs in the configuration file template, and enter the configuration file path as a call option with the call to the SAP HANA platform LCM tool, the program runs in advanced interactive mode and requests values for any parameter values which you didn't specify in the original input. If you entered the batch mode call option, the program runs to completion without any further requests, unless a mandatory parameter was left out of the original input, in which case, the program fails to perform the platform LCM task.

> **Note**
>
> If parameter key-value pairs are specified as command-line options, they override the corresponding parameters in the configuration file. Parameters in the configuration file override default settings.
>
> **Order of parameter precedence:**
>
> Command Line > Configuration File > Default

For more information about program interaction modes and parameter values entry methods, see Related Information.

Related Information

Choosing the Correct SAP HANA HDBLCM for Your Task [page 42]
Performing LCM Tasks by Parameter Entry Method [page 54]
Performing LCM Tasks by Program Interaction Mode [page 50]
Entering Platform LCM Parameters as Call Options from the Command Line [page 57]

5.2.3 Using the Web User Interface

SAP HANA platform lifecycle management tasks can be performed using the SAP HANA database lifecycle manager (HDBLCM) Web user interface.

5.2.3.1 About the Web User Interface

The SAP HANA database lifecycle manager (HDBLCM) Web user interface is hosted by the SAP Host Agent, which is installed on the SAP HANA host. When installing or updating the SAP HANA system, as part of the SAP HANA resident HDBLCM configuration, the SAP HANA system deploys its artifacts on the SAP Host Agent, thus enabling the Web user interface.

All Web user interface actions are always performed in the context of an already installed and registered SAP HANA system. In order to access the SAP HANA database lifecycle manager Web user interface you need to log on as the system administrator user `<sid>`adm.

The communication between the Web browser and the SAP Host Agent is always done over HTTPS, which requires that the SAP Host Agent has a secure sockets layer (SSL) certificate (PSE) in its security directory. For more information about SSL certificate handling, see Related Information.

The backend is provided by the special executable `hdblcmweb`, which is started automatically by the SAP Host Agent as soon as an action is triggered from the Web user interface and terminates after the action completes.

> **Note**
>
> You should never start `hdblcmweb` manually. For security reasons, `hdblcmweb` is always started with system administrator user `<sid>`adm privileges. If you require logging with individual users (to ensure personalized logging), use the SAP HANA database lifecycle manager graphical user or command-line interface.

> **Note**
>
> Make sure that the system administrator user `<sid>`adm has permissions to read the paths, passed as parameters in the Web user interface (for example, the SAP HANA database installation kit or locations with SAP HANA components).

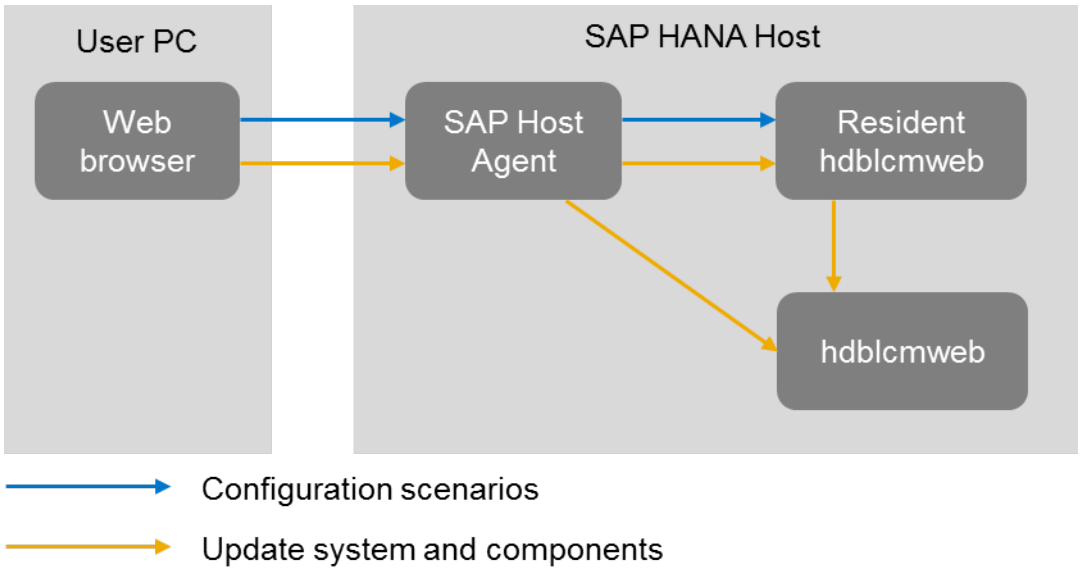

One platform LCM task, which is worth special attention is the update of the SAP HANA system and components. The SAP HANA system updates are always performed by the installation kit SAP HANA database lifecycle manager in the graphical user and command-line interfaces, (and not the SAP HANA resident HDBLCM). This is because the SAP HANA database lifecycle manager, in the graphical user and command-line interfaces, is not forward compatible. Meaning that only the new version of the tool knows how to update an older system.

On the other hand, all scenarios in the Web user interface are handled by the SAP HANA resident HDBLCM, which is part of the system. For this reason, as a first step before even starting the update process, you are required to enter a location of an SAP HANA database installation kit. After detecting the kit, the update Web user interface is loaded from the installation kit and the installation kit SAP HANA database lifecycle manager starts serving as backend until the update process finishes. It is as if you start the SAP HANA database lifecycle manager directly from the installation kit in graphical user or command-line interface.

Related Information

Secure Sockets Layer (SSL) Certificate Handling [page 60]

5.2.3.2 Use the Web User Interface to Perform Platform LCM Tasks

The SAP HANA database lifecycle manager (HDBLCM) can be accessed as a Web user interface in either a standalone browser or in the SAP HANA cockpit.

Prerequisites

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
  Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
  - Microsoft Windows
    - Internet Explorer - Version 9 or higher
      If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing *Tools* > *Compatibility View Settings*.
    - Microsoft Edge
    - Mozilla Firefox - Latest version and Extended Support Release
    - Google Chrome - Latest version
  - SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
  - Mac OS - Safari 5.1 or higher

  > **Note**
  >
  > For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.
- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

Context

The Web user interface supports only the following SAP HANA platform lifecycle management tasks:

- View system information
- Update system and components
- Install or update additional components
- Configure System Landscape Directory (SLD) registration
- Configure inter-service communication

When performing installation and update tasks, various parameters can be set in the *Advanced Parameters Configuration* dialog. To access the *Advanced Parameters Configuration* dialog, click on the gear icon in the footer bar of the SAP HANA HDBLCM Web user interface.

Procedure

Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: https://<hostname>:1129/lmsl/HDBLCM/<SID>/index.html **Note** The URL is case sensitive. Make sure you enter upper and lower case letters correctly.
SAP HANA cockpit	1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. `https://<host_FQDN>:<port>` **Note** FQDN = fully qualified domain name 2. Drill down on the name of the system from *My Resources* or from a group. 3. The links in *Platform Lifecycle Management* each launch additional functionality, giving you expanded capabilities for managing the resource.

Results

The SAP HANA database lifecycle manager is displayed as a Web user interface in either a standalone browser or in the SAP HANA cockpit.

Related Information

SAPUI5 Developer Guide
Add an SAP HANA System

5.2.3.3 Log Off From an SAP HANA System

In the SAP HANA database lifecycle manager (HDBLCM) Web user interface, you can log off from an SAP HANA system and close all connections to the system. To be able to connect to system again, you must log on.

Procedure

- To log off from a system click the *Log out* button.
  All open connections to the system are closed.

  > **Note**
  >
  > Currently, this feature is not available for browsers on mobile devices.

5.2.3.4 Troubleshooting the Web User Interface

If you have problems with the Web user interface, see SAP Note 2078425 for steps you can take to troubleshoot and resolve them.

> **Note**
>
> The Web browser used to render the platform lifecycle management Web user interface in the SAP HANA studio **cannot** be changed via *Windows* > *Preferences* > *General* > *Web Browser*.

Related Information

SAP Note 2078425

5.3 Performing LCM Tasks by Program Interaction Mode

SAP HANA platform lifecycle management tasks can be performed in interactive mode, advanced interactive mode and batch mode.

5.3.1 Use Interactive Mode to Perform Platform LCM Tasks

Interactive mode is a method for running SAP HANA platform lifecycle management (LCM) tools which starts the program and requires you to enter parameter values successively before the program is run. Interactive mode is the default mode for the SAP HANA platform LCM tools.

Context

In general, installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM. For more information about the different SAP HANA database lifecycle manager types, see Related Information.

The SAP HANA platform LCM tools offer a wide variety of parameters which can modify the platform LCM task you are performing. Some parameters can be modified in interactive mode when the graphical user, command-line, or Web user interface requests a value for a given parameter.

Procedure

1. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
Installation Medium (IBM Power Systems)	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`
SAP HANA resident HDBLCM	`cd <sapmnt>/<SID>/hdblcm`

   To access the SAP HANA database lifecycle manager Web user interface, see Related Information.

2. Start the SAP HANA platform lifecycle management tool:

Option	Description
Graphical Interface	`./hdblcmgui`
Command-line Interface	`./hdblcm`

To start the SAP HANA platform LCM tools in interactive mode, simply **do not** enter the parameter for batch mode (`--batch` or `-b`) as a call option. You can enter any other required parameters as call options or load a configuration file. The program runs in interactive mode and requests any missing parameters values, which must be verified or changed. You are provided with a summary of parameter values, which you can accept to run the program to completion, or reject to exit the program.

Related Information

Choosing the Correct SAP HANA HDBLCM for Your Task [page 42]
Use the Web User Interface to Perform Platform LCM Tasks [page 47]

5.3.2 Use Advanced Interactive Mode to Perform Platform LCM Tasks

Interactive mode is a method for running SAP HANA platform lifecycle management (LCM) tools which starts the program and requires you to enter parameter values successively before the program is run. If you would like to enter call options not available in interactive mode, or make use of the configuration file, you can use a combination of interactive mode and advanced parameter entry methods.

Context

In general, installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM. For more information about the different SAP HANA database lifecycle manager types, see Related Information.

The SAP HANA platform LCM tools offer a wide variety of parameters which can modify the platform LCM task you are performing. Some parameters can be modified in interactive mode when the graphical user, command-line, or Web user interface requests a value for a given parameter. However, some parameters are not available in interactive mode, and must be specified either as a call option with the call to the platform LCM tool, or from within a configuration file.

Procedure

1. Review which parameters are offered in interactive mode.

   If the parameter you want to configure is not available in interactive mode, you have two options. You can either enter the parameter key-value pair as a call option with the call to the platform LCM tool. Alternatively, you can generate a configuration file template, and edit the parameters value in the configuration file. Then call the configuration file as a call option with the call to the platform LCM tool.

   Using the configuration file for interactive mode is recommended if you plan to perform the exact same platform LCM task multiple times.

2. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
Installation Medium (IBM Power Systems)	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`
SAP HANA resident HDBLCM	`cd <sapmnt>/<SID>/hdblcm`

3. If you plan to use a configuration file, prepare it with the following steps:
   a. Generate the configuration file template using the SAP HANA platform lifecycle management tool:

      Run the SAP HANA platform LCM tool using the parameter `dump_configfile_template` as a call option. Specify an action and a file path for the template. A configuration file template and a password file template are created.

      ```
      ./hdblcm --action=<LCM action> --dump_configfile_template=<file path>
      ```

   b. Edit the configuration file parameters. Save the file.
   c. Edit the password file. Save the file.
4. Start the SAP HANA platform lifecycle management tool:

   Start the SAP HANA database lifecycle manager in either the graphical user interface or in the command-line interface, with a call option:

   ```
   ./hdblcmgui --<parameter key>=<parameter value>
   ```

   or

   ```
   ./hdblcm --<parameter key>=<parameter value>
   ```

   If you are using a configuration file, you must use the call option `--configfile=<file path>`.

Related Information

Choosing the Correct SAP HANA HDBLCM for Your Task [page 42]

5.3.3 Use Batch Mode to Perform Platform LCM Tasks

Batch mode is a method for running the SAP HANA database lifecycle manager which starts the program and runs it to completion without requiring you to interact with it any further. All required parameter values must be passed as call options or from a configuration file.

Prerequisites

- When using batch mode, passwords must either be defined in the configuration file, or passed to the installer using an XML password file and streamed in via standard input. In both cases, it is necessary to prepare the passwords. For more information, see *Specifying Passwords*.

Context

In general, installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM. For more information about the different SAP HANA database lifecycle manager types, see Related Information.

If you are new to performing the desired SAP HANA platform LCM task in batch mode, it is recommended to run some tests before using batch mode in a production environment.

Procedure

1. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
Installation Medium (IBM Power Systems)	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`
SAP HANA resident HDBLCM	`cd <sapmnt>/<SID>/hdblcm`

2. Start the SAP HANA platform lifecycle management tool:

```
./hdblcm --batch <additional parameters>
```

or

```
./hdblcm -b <additional parameters>
```

It is mandatory to provide an SAP HANA system ID (SID) and user passwords during installation. In batch mode, you are restricted to providing these parameter values as call options on the command line (for passwords, by means of an XML file) or in a configuration file. If you don't provide parameter values for the other required parameters, you implicitly accept the default values.

Example

The following example installs the SAP HANA server and client as a single-host system. The SAP system ID and instance number are also specified from the command line. The system passwords are read from a standard input stream by the installer. All other parameter defaults are automatically accepted and no other input is requested in order to complete the installation.

```
cat ~/hdb_passwords.xml | ./hdblcm --batch --action=install --
components=client,server --sid=DB1 --number=42  --read_password_from_stdin=xml
```

If a configuration file is used in combination with batch mode, an identical system can be installed with a simplified call from the command line. In the following example, passwords are defined in the configuration file, in addition to the action, components, SAP system ID, and instance number.

```
./hdblcm --batch --configfile=/var/tmp/H01_configfile
```

Related Information

Choosing the Correct SAP HANA HDBLCM for Your Task [page 42]
Specifying Passwords [page 77]
Use LCM Configuration Files to Enter Parameters [page 55]
Entering Platform LCM Parameters as Call Options from the Command Line [page 57]

5.4 Performing LCM Tasks by Parameter Entry Method

SAP HANA platform lifecycle management (LCM) parameter values can be entered in a variety of methods: interactively by iteratively providing values in either the graphical interface or command prompt, as command-line options with the call to the platform LCM tool, or in a configuration file.

SAP HANA platform lifecycle management parameter values allow you to customize your SAP HANA installation, update, or configuration. Parameter values can be entered by **one or more** of the following methods:

Interactively (Default)	Using either command line interface, the graphical interface or the Web user interface, most parameters are requested interactively. Default parameter values are proposed in brackets and can be changed or confirmed. Parameters that are not requested (or specified via another method) accept the default value.

Command Line Options	Parameters are given in their accepted syntax as a space delimited list after the program call (for example, `hdblcm` or `hdblcmgui`). The specified parameters replace the defaults. If any mandatory parameters are excluded, they are requested interactively (unless batch mode is specified). All parameters can be entered from the command line. For more details about the accepted parameter syntax, see the inline help output (`--help`) for the individual SAP HANA lifecycle management tool.
Configuration File	The configuration file is a plain text file, for which a template of parameter key-value pairs can be generated, edited, and saved to be called in combination with the program call. If any mandatory parameters are not specified, they are requested interactively (unless batch mode is used). All parameters can be entered in the configuration file. For more information about the configuration file, see Related Information.

> ### Note
>
> If parameters are specified in the command line, they override the corresponding parameters in the configuration file. Parameters in the configuration file override default settings.
>
> **Order of parameter precedence:**
>
> Command Line > Configuration File > Default

5.4.1 Entering Platform LCM Parameters Interactively

SAP HANA platform LCM interactive mode is default interaction mode for all platform LCM programs and interfaces.

You can run the graphical, command-line, or Web user interface in interactive mode by simply starting the program, and entering parameter values as they are requested by the program. In interactive mode, parameter default values are suggested in brackets and can be accepted with `Enter`.

Not all parameters are requested in interactive mode. If you would like to configure a parameter not offered in interactive mode, you must enter it as a call option with the call to the platform LCM program, or use corresponding configuration file for the platform LCM task.

5.4.2 Use LCM Configuration Files to Enter Parameters

By defining a prepared configuration file during installation, specified parameter values are used by the SAP HANA platform lifecycle management (LCM) tools to build a customized SAP HANA system.

Context

The configuration file is a plain text file of specified parameters, written in the same syntax as in the command line (except without the leading two dashes --). A configuration file template can be generated, edited, and saved to be called with the call to the SAP HANA database lifecycle manager (HDBLCM).

The configuration file template provides a brief, commented-out summary of each parameter. Each parameter is set to its default value.

Procedure

1. Change to the directory where the SAP HANA database lifecycle manager is located:

Option	Description
Installation Medium (Intel-Based Hardware Platforms)	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
Installation Medium (IBM Power Systems)	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`
Installation Archive downloaded from SAP Support Portal (SAP Service Marketplace)	`cd SAP_HANA_DATABASE`
SAP HANA resident HDBLCM	`cd <sapmnt>/<SID>/hdblcm`

   In general, installation and update is carried out from the installation medium. Configuration tasks are performed using the SAP HANA resident HDBLCM. For more information about the two SAP HANA database lifecycle manager types, see Related Information.

2. Generate the configuration file template using the SAP HANA platform lifecycle management tool:

   Run the SAP HANA platform LCM tool using the parameter `dump_configfile_template` as a call option. Specify an action and a file path for the template. A configuration file template and a password file template are created.

   ```
   ./hdblcm --action=<LCM action> --dump_configfile_template=<file path>
   ```

3. Edit the configuration file parameters. Save the file.

   It is recommended that at least the SAP system ID (`sid`) and the instance number (`number`) are uniquely defined. There are several required parameters, that are provided default values in case they are not customized. For more information, refer to the default values.

   Some file path parameters have automatic substitution values as part of the default file path, using the `sid` (SAP HANA system ID) and `sapmnt` (installation path) parameters, so that the substituted values create file paths that are unique and system-specific. For example, the default for the data file path is: `datapath=/hana/data/${sid}`, where `sid` is automatically replaced by the unique SAP HANA system ID.

4. Start the SAP HANA platform lifecycle management tool:

   Run the SAP HANA platform LCM tool using the parameter `configfile` as a call option. Specify the file path of the edited template.

   ```
   ./hdblcm --configfile=<file path>
   ```

You can specify the path to a directory in which custom configuration files are saved using the parameter `custom_cfg` as a call option.

Related Information

configfile [page 246]
custom_cfg [page 248]
Choosing the Correct SAP HANA HDBLCM for Your Task [page 42]

5.4.3 Entering Platform LCM Parameters as Call Options from the Command Line

Call options are available for every SAP HANA platform LCM program.

You can use call options for a number of reasons:

- The parameter is not available in interactive mode, but can be entered as a call option.
- You are using batch mode.
- You are using a configuration file, but would like to override a parameter in the configuration file with a new value.
- You are installing an SAP HANA multiple-host system from the command line.

A call option is entered with the following notation:

```
./<program call> --<parameter1 key>=<parameter1 value> --<parameter2
key>=<parameter2 value>
```

Call options start with a double dash (`--`) if they are written in long-form syntax. Some parameters also have short-form syntax, in which they are preceded with a single dash (`-`). For more information about call option syntax, see the *Parameter Reference* topics.

Related Information

Parameter Reference [page 233]

5.5 Executing Platform LCM Tasks

SAP HANA platform lifecycle management tasks can be performed on multiple-host systems centrally, by running the SAP HANA database lifecycle manager (HDBLCM) from any worker host and using remote execution to replicate the call on all remaining system hosts. Otherwise, the platform LCM tasks can be

executed first on a worker host, and then re-executed manually on each remaining host. This method is considered decentralized execution.

The following is an example of an SAP HANA system update performed centrally and decentrally.

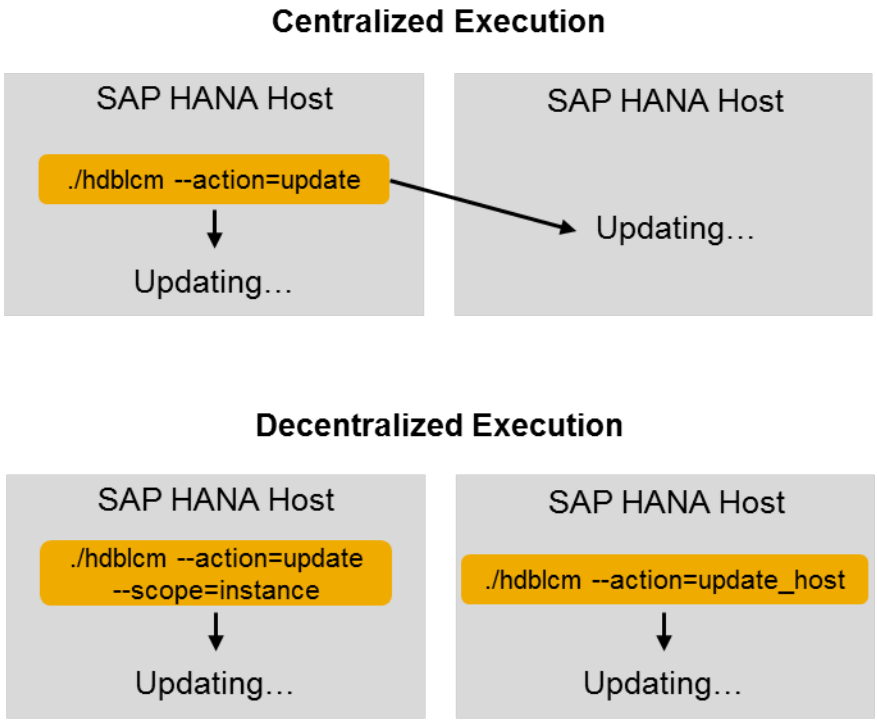

Related Information

Centralized Execution of Platform LCM Tasks [page 59]
Decentralized Execution of Platform LCM Tasks [page 63]

5.5.1 Centralized Execution of Platform LCM Tasks

SAP HANA platform lifecycle management (LCM) tasks can be performed centrally on multiple-host SAP HANA systems in a number of ways depending on the available certificate keys and the remote execution configuration.

5.5.1.1 Using Secure Shell (SSH) to Execute Platform LCM Tasks

An SAP HANA system must be installed with root user credentials. During installation a secure shell (SSH) key is configured so that future platform LCM tasks can be performed remotely on multiple-host SAP HANA systems without requiring the root user password.

By default, the SAP HANA database lifecycle manager (HDBLCM) uses SSH during SAP HANA system installation or update. In order to use SSH, the SFTP subsystem must be active. Install the SAP Host Agent on all system hosts to perform platform LCM tasks without root credentials. Once the SAP Host Agent is installed, it is used to perform any platform LCM tasks executed from the Web user interface or as the system administrator user `<sid>`adm.

> **Note**
>
> Platform LCM tasks cannot be executed remotely via SSH as the system administrator user `<sid>`adm.

The root user name must be the same for all hosts in a multiple-host system. The password of the root user must be identical on all hosts. If the passwords are not identical on all hosts or if you do not want to pass the passwords to the SAP HANA database lifecycle manager (HDBLCM), SSH keys can be distributed between the hosts in advance. Alternatively, platform LCM tasks can be executed without root credentials by using the SAP Host Agent.

Related Information

SAP Note 1944799
SAP Note 2009879
Using SAP Host Agent to Execute Platform LCM Tasks [page 60]

5.5.1.2 Using SAP Host Agent to Execute Platform LCM Tasks

Platform LCM tasks can be executed without root credentials by using the SAP Host Agent. The SAP Host Agent is installed and updated by default during SAP HANA system installation and update.

The SAP HANA database lifecycle manager (HDBLCM) relies on the SAP Host Agent for the following functionality to work:

- Execution as the system administrator user `<sid>`adm
- Connectivity to remote hosts via HTTPS (when no SSH or root user credentials are available)
- Execution from the SAP HANA database lifecycle manager Web user interface

> ⓘ Note
>
> The SAP HANA cockpit uses the SAP Host Agent to execute tasks as the system administrator user `<sid>`adm, for example, stopping and starting the system, or troubleshooting a system experiencing performance problems.

If execution on the remote hosts is done via SSH (default, `--remote_execution=ssh`), the SAP HANA database lifecycle manager is able to connect to a remote host via SSH and install and configure the SAP Host Agent. In contrast, the remote execution via SAP Host Agent (`--remote_execution=saphostagent`) requires that the SAP Host Agent is installed and configured on all involved hosts in advance, which includes:

- Install SAP Host Agent
- Configure a Secure Sockets Layer (SSL) certificate for the SAP Host Agent, so that the HTTPS port 1129 is accessible. If you don't want to configure HTTPS, it is also possible to use the call option `--use_http`. It tells the SAP HANA database lifecycle manager to communicate with the SAP Host Agent via HTTP. During the addition of new hosts to an SAP HANA system (also during the installation of a multiple-host system), the HTTPS of the SAP Host Agent is automatically configured by the SAP HANA database lifecycle manager.

> ⚠ Caution
>
> Use the call option `--use_http` with caution, because passwords are also transferred in plain text via HTTP.

Related Information

SAP Host Agent
Configuring SSL for SAP Host Agent on UNIX

5.5.1.2.1 Secure Sockets Layer (SSL) Certificate Handling

To enable secure communication with the SAP Host Agent over HTTPS, the SAP Host Agent needs a secure sockets layer (SSL) certificate in its security directory. This certificate is also used by the SAP HANA database

lifecycle manager (HDBLCM) Web-based user interface because the Web pages are served by the SAP Host Agent.

The SAP HANA database lifecycle manager handles certificate management during system installation, update, or rename, as well as during the addition of new hosts as follows:

- If there is no certificate in the SAP Host Agent security directory, the SAP HANA database lifecycle manager generates one. The SAP HANA host name is used as the default certificate owner. The certificate owner can be changed by using the call option `--certificates_hostmap=<fully_qualified_domain_name>`.
- If there is an existing certificate, the following applies:
  - If the certificate host name is not passed to the SAP HANA database lifecycle manager, or if the certificate host name is the same as the owner of the current certificate, the current certificate is preserved.
  - If the certificate host name is passed via the call option `--certificates_hostmap` and it differs from the owner of the current certificate, a new certificate is generated.
  - During update of an SAP HANA system, if the certificates on all hosts are in place, the call option `--certificates_hostmap` is ignored and the current certificates are preserved.

If you want to use your own SSL certificates, see the SAP Host Agent documentation in Related Information.

Related Information

certificates_hostmap [page 241]
SAP Note 1907566
SSL Configuration for the SAP Host Agent

5.5.1.2.2 Starting Platform LCM Tasks as the System Administrator User <sid>adm

When starting platform LCM tasks as the system administrator user `<sid>`adm, the SAP HANA database lifecycle manager (HDBLCM) requires the usage of SAP Host Agent for execution of remote and local operations.

The following tasks in the SAP HANA database lifecycle manager can be performed as the system administrator user `<sid>`adm:

- System update from the installation medium
- Installation or update of additional components from the SAP HANA resident HDBLCM
- Host addition and host removal
- System Landscape Directory (SLD) registration configuration
- Inter-service communication configuration

Make sure that SAP Host Agent is installed and configured (HTTPS-enabled) on all hosts of the SAP HANA system.

> **Note**
>
> Platform LCM tasks cannot be executed remotely via SSH as the system administrator user <sid>adm.

> **Note**
>
> Make sure that the system administrator user <sid>adm has permissions to read the paths passed as parameters (for example, the locations of the SAP HANA components).

5.5.1.2.3 Add Hosts Using SAP Host Agent

You can add hosts to an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program in combination with the SAP Host Agent in the command-line interface.

Prerequisites

- The SAP HANA system has been installed with its server software on a shared file system (export options `rw`, `no_root_squash`).
- The host which is to be added has access to the installation directories <sapmnt> and <sapmnt>/<SID>.
- The latest supported version of SAP Host Agent is installed on the host which is to be added. The SAP Host Agent will create the<sapsys> group, if it does not exist prior to installation. Make sure that the group ID of the <sapsys> group is the same on all hosts. For information about installing or updating the SAP Host Agent individually, see *Installing SAP Host Agent Manually* and *Upgrading SAP Host Agent Manually*.
- A Secure Sockets Layer (SSL) certificate is configured for the SAP Host Agent, so that the HTTPS port 1129 is accessible and the Personal Security Environment (PSE) for the server is prepared. For more information about SSL configuration for the SAP Host Agent, see *Configuring SSL for SAP Host Agent on UNIX*.
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running.
- You are logged on as root user or as the system administrator user <sid>adm.
- The difference between the system time set on the installation host and the additional host is not greater than 180 seconds.
- The operating system administrator (<sid>adm) user may exist on the additional host. Make sure that you have the password of the existing <sid>adm user, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) resident program will not modify the properties of any existing user or group.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

   ```
   cd <sapmnt>/<SID>/hdblcm
   ```

   By default, `<sapmnt>` is `/hana/shared`.
2. Start the SAP HANA database lifecycle manager interactively in the command line:

   ```
   ./hdblcm --remote_execution=saphostagent
   ```
3. Select the index for the *add_hosts* action.
4. Enter the names of the hosts to be added.
5. Enter the SAP Host Agent administrator (`sapadm`) password.
6. Define additional system properties.
7. Review the summary, and select *y* to finalize the configuration.

Results

You have added one or more new hosts to an SAP HANA system. The SAP HANA system you have configured is a multiple-host system.

The new hosts have been added to the SAP HANA landscape information. The new hosts have been added to the landscape information of the system database.

This configuration task can also be performed in batch mode and using a configuration file. For more information about the available configuration methods, see *Using the SAP HANA Platform LCM Tools*.

Related Information

remote_execution [page 271]
Host Addition Concepts
Using the SAP HANA Platform LCM Tools [page 41]
Using SAP Host Agent to Execute Platform LCM Tasks [page 60]
SAP Host Agent Installation
SAP Host Agent Upgrade
SSL Configuration for the SAP Host Agent

5.5.2 Decentralized Execution of Platform LCM Tasks

In some circumstances platform LCM actions must be executed on each individual host of the multiple-host system. This is referred to as **decentralized execution**.

Typically, SAP HANA platform lifecycle management actions, such as update, rename, and inter-service communication configuration, can be performed on a multiple-host system from one host. This is referred

to as **centralized execution** and requires SSH or root credentials. For more information, see Centralized Execution of Platform LCM Tasks in Related Information.

In some circumstances, a secure shell (SSH) key may not be installed or root credentials are not available. In this case, the platform LCM actions must be executed on each individual host of the multiple-host system, which is also known as **decentralized execution**. For more information about decentralized execution, see SAP Note 2048681 in Related Information.

Related Information

SAP Note 2048681
Executing Platform LCM Tasks [page 57]
Centralized Execution of Platform LCM Tasks [page 59]

5.6 Additional Information About Using the SAP HANA Platform LCM Tools

If you have already familiarized yourself with the way the SAP HANA database lifecycle manager (HDBLCM) works, you may be interested in additional information like log and trace files, Linux kernel parameter settings, or troubleshooting.

Related Information

Logging [page 64]
Linux Kernel Parameters [page 65]
General Troubleshooting for the SAP HANA Platform LCM Tools [page 67]

5.6.1 Logging

SAP HANA platform lifecycle management processes are logged by the system. The log files are stored in the following path:

`/var/tmp/hdb_<SID>_<action>_<time stamp>`

where `<action> :: = 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 display in the installation tool with the GUI

- `<hostname>_tracediff.tgz`: provides a delta analysis of the original trace files, makes a detailed analysis easier

You can also view diagnostic files in the SAP HANA database explorer using the administration function. For more information, see *View Diagnostic Files in the SAP HANA Database Explorer* in the *SAP HANA Administration Guide*.

Instant Logging

If an LCM action crashes or hangs before the execution is finished, even if no LCM action trace is enabled, HDBLCM writes a trace, which has the function of a preliminary (unformatted) log file. Upon program completion, this preliminary logfile is removed and replaced by the real, formatted log file.

The environment variable `HDB_INSTALLER_TRACE_FILE=<file>` enables the trace.

The environment variable `HDBLCM_LOGDIR_COPY=<target directory>` creates a copy of the log directory.

Log Collection

If you perform platfom LCM actions on multiple-host SAP HANA systems, all log files are collected to a local folder to make error analysis more convenient.

To collect log files for multiple-host SAP HANA systems, an HDBLCM action ID is passed to each sub-program (underlying LCM tool) working on a remote host. Each sub-program writes a copy of the log file in to the following directory: `<installation path>/<SID>/HDB<instance number>/<host name>/trace`

Related Information

View Diagnostic Files in the SAP HANA Database Explorer

5.6.2 Linux Kernel Parameters

The following table describes the parameters and limits that are set by the SAP HANA database lifecycle manager (HDBLCM) during the installation or update of an SAP HANA database. The actual values may differ, depending on your system configuration.

> **Note**
>
> The SAP HANA database lifecycle manager (HDBLCM) will only change existing values if they are lower than the default values.

The SAP Host Agent can automatically optimize the following Linux Kernel Parameters:

- `net.ipv4.ip_local_port_range`
- `net.ipv4.ip_local_reserved_ports`

To configure the SAP Host Agent, make sure that the `/etc/sysctl.conf` configuration does not contain any of these two parameters. Afterwards, configure the SAP Host Agent profile parameters as described in *SAP Note 401162*.

Parameter	Description	Value	Location
`nofile`	Open file descriptors per user	`1048576`	`/etc/security/limits.conf`
`fs.file-max`	Open file descriptors per host	`20000000`	`/etc/sysctl.conf`
`fs.aio-max-nr`	Maximum number of asynchronous I/O requests	`18446744073709551615 (= 2^64-1 = ULONG_MAX)`	`/etc/sysctl.conf`
`vm.memory_failure_early_kill`	Method for killing processes when an uncorrected memory error occurs	`1`	`/etc/sysctl.conf`
`kernel.shmmax`	Maximum shared memory segment size (the default minimum value is 1 GB)	`1073741824`	`/etc/sysctl.conf`
`kernel.shmmni`	Maximum number of shared memory segments	`32768`	`/etc/sysctl.conf`
`kernel.shmall`	System-wide limit of total shared memory, in 4k pages	• RAM >= 35.5 TB: `(shmmax * shmmni) / 65536` • RAM < 35.5 TB: `(0.9 * <RAM in bytes>) / 4096`	`/etc/sysctl.conf`
`net.ipv4.ip_local_port_range`	Lower limit of ephemeral port range **Note** No changes are applied if the SAP HANA database lifecycle manager (HDBLCM) detects or installs SAP Host Agent version 7.20.162 or greater. The SAP Host Agent takes care of adjusting this parameter and setting it manually is neither recommended nor required. For more information, see SAP Note 2382421	`40000`	`/etc/sysctl.conf`
`vm.max_map_count`	Maximum number of Virtual Memory Areas (VMAs) that a process can own	`2147483647`	`/etc/sysctl.conf`

Related Information

SAP Note 2382421

SAP Note 401162
SAP Note 941735

5.6.3 General Troubleshooting for the SAP HANA Platform LCM Tools

The SAP HANA database lifecycle manager (HDBLCM) is a wrapper tool that calls the underlying HDB tools to perform the platform LCM action. If something unexpected happens when using HDBLCM, and the LCM action cannot be completed, you can check the logs and separately run the affected underlying tools.

> ⚠ Caution
>
> We only recommend the following underlying tools to be used for troubleshooting purposes.

Program Name	Description	Location
hdbinst	Command-line tool for installing the software	Installation media
hdbsetup	Installation tool with a graphical interface for installing or updating the software	Installation media
hdbuninst	Command-line tool for uninstalling the software and removing a host	Installation media and `<installation path>/<SID>/global/hdb/install/bin`
hdbaddhost	Command-line tool for adding a host to a system	`<installation path>/<SID>/global/hdb/install/bin`
hdbupd	Command-line tool for updating the software	Installation media

hdbrename	Command-line tool for renaming a system	`<installation path>/<SID>/global/hdb/install/bin`
		and
		`/usr/sap/<SID>/SYS/global/hdb/install/bin`

Program Name	Description	Location

hdbreg	Command-line tool for registering an SAP HANA system	`<installation path>/<SID>/global/hdb/install/bin`
		and
		`/usr/sap/<SID>/SYS/global/hdb/install/bin`

hdbremovehost	Command-line tool for removing a host	`<installation path>/<SID>/global/hdb/install/bin`
		and
		`/usr/sap/<SID>/SYS/global/hdb/install/bin`

hdbmodify	This command line tool removes and adds remote hosts.
	Furthermore, the listen interface can be changed ('local', 'global', 'internal').
		`<installation path>/<SID>/global/hdb/install/bin`
		and
		`/usr/sap/<SID>/SYS/global/hdb/install/bin`

hdbupdrep	Command-line tool for upgrading a repository by loading delivery units into the database	`<installation path>/<SID>/global/hdb/install/bin`
		and
		`/usr/sap/<SID>/SYS/global/hdb/install/bin`

5.6.4 Managing SAP HANA System Components

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

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

- SAP HANA mandatory components

    - SAP HANA server
    - SAP HANA client
    - Local secure store (LSS)
- SAP HANA additional components
    - SAP HANA studio
    - Application Function Libraries (AFL and the product-specific AFLs IBP, RTL, TRP, VCH, XRP)
    - SAP liveCache applications (SAP LCA or LCAPPS-Plugin)
    - SAP HANA smart data access (SDA)

> **Note**
>
> To install or uninstall the Solution Manager Diagnostics Agent, use Software Provisioning Manager (SWPM). For more information about the setting up the Solution Manager Diagnostics Agent using SWPM, see SAP Note 1858920 in Related Information.

> **Note**
>
> SAP LT replication configuration is a part of SL Toolset 1.0. For more information about configuring SAP LT replication, see SAP Note 1891393 in Related Information.

- SAP HANA options
    - SAP HANA dynamic tiering
    - SAP HANA streaming analytics
    - SAP HANA accelerator for SAP ASE

For more information about installing, updating, and uninstalling the SAP HANA mandatory components and SAP HANA additional components, see the *SAP HANA Server Installation and Update Guide*. For more information about installing, updating, and uninstalling the SAP HANA options, see SAP HANA option documentation in Related Information.

> **Caution**
>
> Be aware that you need additional licenses for SAP HANA options. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities* in Related Information.

Related Information

SAP Note 1858920
SAP Note 1891393
Important Disclaimer for Features in SAP HANA

5.6.5 Check the Installation Using the Command-Line Interface

You can check the installation of an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) resident program in the command-line interface for troubleshooting.

Prerequisites

- You are logged in as root user.
- Any user has read and execute permissions for the directory that contains the installation medium.
- Depending on the storage solution, set the export options `rw`,`no_root_squash` for the installation directory.
- The operating system administrator (`<sid>adm`) user and other operating system users may exist prior to installation. Make sure that you have the passwords of the existing users, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) will not modify the properties of any existing user or group. The `<sid>adm` user passwords must be identical on all hosts.
- The SAP HANA system has been installed with its server software on a shared file system (export options `rw`, `no_root_squash`).
- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

   ```
   cd <sapmnt>/<SID>/hdblcm
   ```

   By default, `<sapmnt>` is `/hana/shared`.
2. Start the SAP HANA database lifecycle manager interactively in the command line:

   ```
   ./hdblcm --action=check_installation
   ```
3. Enter the required credentials.
4. Review the summary, and select *y* to finalize the configuration.

Results

The check tool outputs basic information about the configuration of the file system, system settings, permission settings, and network configuration. The checks are based on the property file stored in the following path:

```
<sapmnt>/<SID>/global/hdb/install/support/hdbcheck.xml
```

Use the generated log files as a reference in the case of troubleshooting. The log file is stored in the following path:

```
/var/tmp/hdb_<SID>_hdblcm_check_installation_<time stamp>/hdblcm.log
```

Related Information

check_installation [page 243]

6 Installing an SAP HANA System

The SAP HANA database lifecycle manager (HDBLCM) is used to install an SAP HANA system, including server, client, studio, and additional components, in a graphical user interface or the command-line interface. The SAP HANA system can be uninstalled or configured using the resident version of the SAP HANA database lifecycle manager (HDBLCM).

Related Information

Getting Started with SAP HANA System Installation [page 72]
Installing a Single-Host System [page 79]
Installing a Multiple-Host System [page 84]
Installing an SAP HANA System Including the XS Advanced Runtime [page 97]

6.1 Getting Started with SAP HANA System Installation

The SAP HANA database lifecycle manager (HDBLCM) is used to install an SAP HANA system. Installation parameters are a fundamental aspect of the SAP HANA database lifecycle manager (HDBLCM), and provide the opportunity to optimize and customize the system during installation.

Installation parameters are used for all methods of parameter specification - interactively, as command line options, or with the configuration file. The following information is relevant for either installation mode (interactive mode or batch mode).

6.1.1 Mandatory Installation Values

The only mandatory parameter without a default value is the SAP system ID (SID, `sid`), which must be specified. However, it is also recommended to define an instance number (`number`), since these two parameters are meaningful to the identity of the system.

If the installation is run in batch mode from the installation medium, the minimum required parameters are the SID and the passwords (specified in XML syntax and streamed in, or specified in the configuration file). In the case that only the SID and passwords are entered as parameters, the other parameters automatically take their default value.

In batch mode, if one of the mandatory parameters, without a default, is not specified, the installation fails with an error. If an installation is not run in batch mode, but in interactive mode (default) instead, the missing mandatory parameters are requested in the console.

When building a multiple-host system, the `action` and `addhosts` parameters are mandatory command-line specifications for the SAP HANA database lifecycle manager.

When installing SAP HANA Local Secure Store, which is the default behavior during SAP HANA system installations & updates, it is required to define an LSS installation path (`lss_inst_path`), as well as `<sid>crypt` user password (`lss_user_password`) and Local Secure Store backup password (`lss_backup_password`).

Related Information

lss_inst_path [page 264]

6.1.2 Changeable Default Values for Installation

The SAP HANA database lifecycle manager (HDBLCM) uses the following default values during installation unless you change them.

Some default values are based on the predefined values on the current host. In a multiple-host system, it is recommended to manually check the mandatory values on each hosts **before** installation.

Changeable Parameter Defaults

Parameter	System Default Value	Interactive Mode Availability
`autoadd_xs_roles`	`1` (on)	■
`autostart`	`0` (off)	■
`certificates_hostmap`	`<current host>`	■
`client_path`	`<sapmnt>/<SID>/hdbclient`	◇
`components`	`client,server,studio,lss` (dependent on the installer finding installation sources for the components)	■
`configure_python`	`python3`	◇
`copy_repository`	`/hana/shared/<SID>/hdbstudio_update`	◇
`create_initial_tenant`	`1` (on)	◇
`datapath`	`/hana/data/<SID>`	■

Parameter	System Default Value	Interactive Mode Availability
db_isolation	low	◇
groupid	79	◇
home	/usr/sap/<SID>/home	■
hostname	<current host>	■
import_xs_content	1 (on)	◇
install_hostagent	y (on)	◇
install_ssh_key	y (on)	◇
logpath	/hana/log/<SID>	■
lss_inst_path	/lss/shared	■
lss_userid	<SID>crypt	■
lss_user_home	/usr/sap/<SID>/lss/home	■
lss_user_shell	/bin/sh	■
max_mem	0	■
number	<next successive un-used instance number on the host>	■
org_manager_user	XSA_ADMIN	■
org_name	orgname	■
prod_space_name	PROD	■
remote_execution	ssh	◇
repository	y (on)	◇
restrict_max_mem	(off)	■
root_user	root	■
sapmnt	/hana/shared	■
secure_store	localsecurestore	◇

Parameter	System Default Value	Interactive Mode Availability
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>	■
workergroup	default	■
xs_components	<xsac_monitoring,xsac_services,xsac_ui5_fesv3,xsac_portal_serv, xsac_alm_pi_ui,xsac_xsa_cockpit>	■
xs_components_nostart	none	◇
xs_customer_space_isolation	1 (on)	■
xs_routing_mode	<ports>	■
xs_sap_space_isolation	1 (on)	◇

> **Note**
>
> To substitute parameters in configuration files and batch mode, they must be written in the form `${<parameter>}`. Substitution also occurs in interactive mode in order to create a suggested path. The advantage of substitution is that the SAP system ID (`SID`) and the installation path (`sapmnt`, which is `/hana/shared`, by default) only need to be specified once, and are then substituted in to the other parameter values. This ensures that the system has unique file system paths if multiple systems are installed on the same host. However, if it is preferred to deviate from the default paths, it is necessary to pay attention to the settings, especially in the configuration file, and when installing in batch mode.

Related Information

Parameter Reference [page 233]

6.1.3 Users Created During Installation

The following users are automatically created during the installation: <sid>adm, sapadm, SYSTEM, and <sid>crypt.

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 identifier (UID) of the <sid>adm user is defined during the system installation. • The password of the <sid>adm user is set during installation with the password parameter. • If you do not want the operating system user <sid>adm and its primary group to be created automatically, you can create it before installation. This might be the case if you use central user management such as Lightweight Directory Access Protocol (LDAP) or Network Information System (NIS). The SAP HANA database lifecycle manager (HDBLCM) will not modify the properties of any existing user or group. The following requirements apply: • The name of the user must follow the schema <sid>adm. All letters must be lowercase. • The user should have a UID greater than 999. • The primary group of the user must be sapsys. The default GID of the sapsys group is 79. • The UID of this operating system user and GID of its primary group must be unique and identical on each host of a multiple-host system.
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. • If you do not want the user sapadm and its primary group to be created automatically, you can create it before installation. The following requirements apply: • The primary group of the user must be sapsys. The default GID of the sapsys group is 79. • The GID of the primary group of the sapadm user must be unique and identical on each host of a multiple-host system.

User	Description
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. • Two SYSTEM users area created: one for the system database and one for the tenant database. • The password of the SYSTEM user is set during installation with the `system_user_password` parameter.
`<sid>crypt`	The trusted local secure store (LSS) user. • The user `<sid>crypt` owns the storage of the encryption keys and other similarly sensitive data. • The user `<sid>crypt` is the only trusted user of the local secure store. Only processes called by a trusted user are accepted by the LSS right away.

Related Information

Predefined Database Users
Operating System User <sid>adm
Predefined XS Advanced Users

6.1.4 Specifying Passwords

Passwords are a mandatory parameter for installing an SAP HANA system. There are three methods for configuring passwords.

Interactive Mode

Interactive installation is available for the SAP HANA database lifecycle manager in the graphical user interface, the command-line interface, and the Web user interface. Passwords are entered manually one-by-one as they are requested by the installer. This method is preferred for quick, individual system installations.

Command Line

In the command line, passwords can be specified by passing them to the standard input stream of the SAP HANA database lifecycle manager. The SAP HANA database lifecycle manager is called using the `read_password_from_stdin` parameter in the command line with batch mode. For security reasons, you may want to consider providing encrypted passwords along with a tool that decrypts the passwords and

passes them to the standard input stream. Parameters specified in the command-line override parameters specified in the configuration file. Since this method is the most powerful and flexible method, it is often the preferred method for installing multiple SAP HANA systems at one time.

Configuration File

It is possible to specify passwords in the configuration file. A configuration file template is created with all the parameters set to their default values. The configuration file is edited to the preferred parameter values, then it is saved, and the values are read by the installer during installation. This method is preferred for a one-step installation that can be re-created several times. If passwords are specified in the configuration file, its permission settings should limit access to the root user, for security reasons.

Example

The following is an example of the configuration file, with configured password parameters:

```
configfile1.cfg
# Root User Password
root_password=Root1234
...
# SAP Host Agent (sapadm) Password
sapadm_password=Agent1234
...
# System Administrator Password
password=Adm1234
....
# Database User (SYSTEM) Password
system_user_password=Sys1234
```

Now, the configuration file (stored in the root user's home directory) is called from the command line using the `configfile` parameter:

```
./hdblcm --sid=DB1 --configfile=~/configfile1.cfg
```

Related Information

configfile [page 246]
read_password_from_stdin [page 270]
sid [page 275]
Tutorial: Installing a Single-Host System with Passwords Read from XML Standard Input Stream [page 224]

6.2 Installing a Single-Host System

The SAP HANA database lifecycle manager can be used to install an SAP HANA single-host system in one of the program interfaces, and with a combination of parameter specification methods.

A single-host system is the simplest system installation type. It is possible to run an SAP HANA system entirely on one host and then scale the system up as needed.

The following graphic shows the file system for a single-host installation:

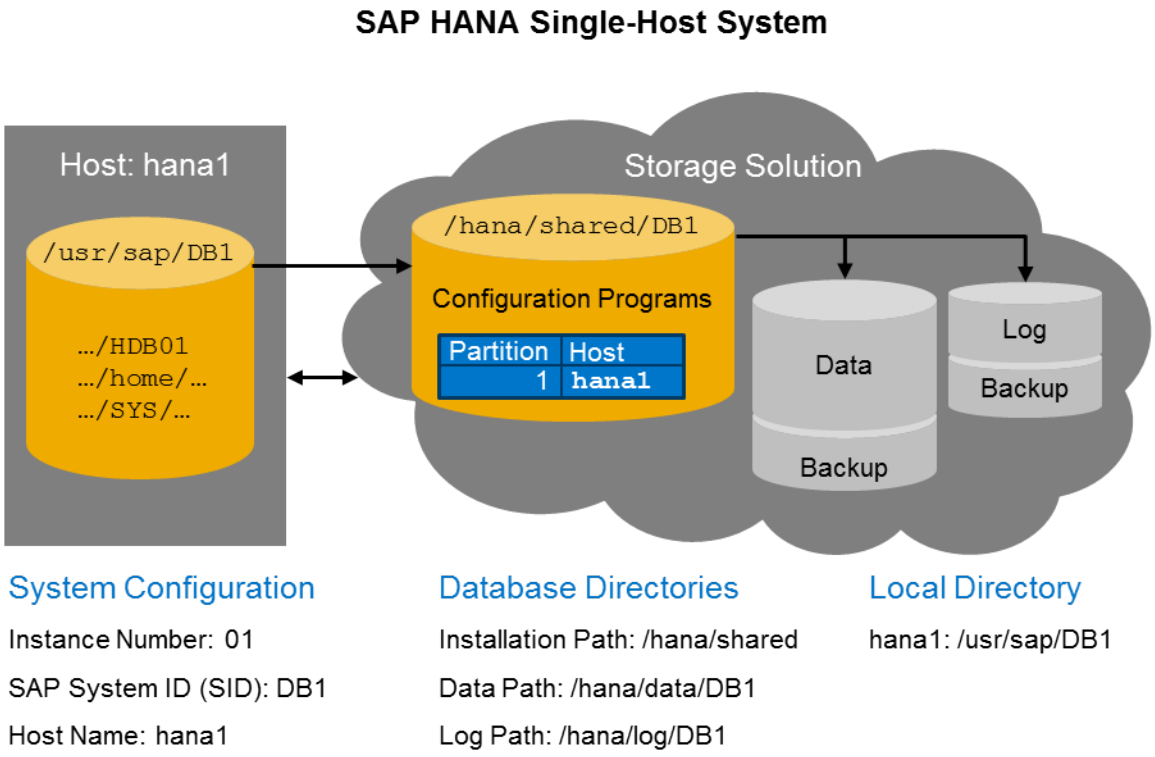

6.2.1 Install a Single-Host SAP HANA System Using the Graphical User Interface

A single-host SAP HANA system can be installed using the SAP HANA database lifecycle manager (HDBLCM) graphical user interface.

Prerequisites

- You are logged in as root user.
- Any user has read and execute permissions for the directory that contains the installation medium.

- Depending on the storage solution, set the export options `rw,no_root_squash` for the installation directory.
- The operating system administrator (`<sid>adm`) user and other operating system users may exist prior to installation. Make sure that you have the passwords of the existing users, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) will not modify the properties of any existing user or group. The `<sid>adm` user passwords must be identical on all hosts.
- The root user is able to execute graphical applications.

Context

The following procedure describes the installation of an SAP HANA system in interactive mode by entering parameters interactively. This procedure may also be performed in advanced interactive mode, with parameters entered as call options or from a configuration file. For more information about interaction modes and parameter entry methods, see *Using the SAP HANA Platform LCM Tools* in Related Information.

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. You can add additional tenant databases later using the SAP HANA cockpit. SAP HANA options hosts or host roles can be automatically provisioned to a system that is installed with a single tenant. If the SAP HANA system contains multiple tenant databases, the SAP HANA options host or host role must be manually provisioned to the tenant.

> **Note**
>
> Not all parameters are requested interactively. Some parameters have default values that do not require confirmation in interactive mode. Those parameters must be specified as call options or from a configuration file. For more information about changeable default values, see Related Information.

Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
IBM Power Systems	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`

> **Note**
>
> If you downloaded the components to a different directory, change to the directory where you unpacked the archive.

2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

   ```
   ./hdblcmgui
   ```

   The SAP HANA database lifecycle manager graphical user interface appears.

   > **Note**
   >
   > To activate the local secure (LSS) store during installation, run `hdblcmgui` with the parameter `secure_store=localsecurestore`.

3. Select a detected software component or add a software component location by selecting *Add Component Location*. Then select *Next*.
4. Select *Install New System*, then select *Next*.
5. Select the components you would like to install, then select *Next*.
6. Select *Single-Host System* as the *System Type*, then select *Next*.
7. Specify the SAP HANA system properties.

   For a list of all system properties, see *System Properties* in Related Information.
8. After specifying all system properties, review the summary, and select *Install*.

Results

A single-host SAP HANA system is installed. A log file is available.

After installing the SAP HANA system, you may want to perform configuration tasks. For more information, see *Managing the SAP HANA System After Installation* or the platform lifecycle management section of the *SAP HANA Administration Guide*.

Related Information

System Properties [page 142]
Using the SAP HANA Platform LCM Tools [page 41]
Changeable Default Values for Installation [page 73]
Managing the SAP HANA System After Installation [page 209]
SAP HANA Platform Lifecycle Management
SAP Note 2243156
Activate the Local Secure Store (LSS) [page 197]

6.2.2 Install a Single-Host SAP HANA System Using the Command-Line Interface

A single-host SAP HANA system can be installed using the SAP HANA database lifecycle manager (HDBLCM) command-line interface.

Prerequisites

- You are logged in as root user.
- Any user has read and execute permissions for the directory that contains the installation medium.
- Depending on the storage solution, set the export options `rw,no_root_squash` for the installation directory.
- The operating system administrator (`<sid>adm`) user and other operating system users may exist prior to installation. Make sure that you have the passwords of the existing users, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) will not modify the properties of any existing user or group. The `<sid>adm` user passwords must be identical on all hosts.

Context

The following procedure describes the installation of an SAP HANA system in interactive mode and entering parameters interactively. This procedure may also be performed in advanced interactive mode or batch mode, with parameters entered as call options or from a configuration file. For more information about interaction modes and parameter entry methods, see *Using the SAP HANA Platform LCM Tools* in Related Information.

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. You can add additional tenant databases later using the SAP HANA cockpit. SAP HANA options hosts or host roles can be automatically provisioned to a system that is installed with a single tenant. If the SAP HANA system contains multiple tenant databases, the SAP HANA options host or host role must be manually provisioned to the tenant.

> **Note**
>
> Not all parameters are requested interactively. Some parameters have default values that do not require confirmation in interactive mode. Those parameters must be specified as call options or from a configuration file. For more information about changeable default values, see Related Information.

Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
IBM Power Systems	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`

> **Note**
>
> If you downloaded the components to a different directory, change to the directory where you unpacked the archive.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblcm
```

> **Note**
>
> To activate the local secure (LSS) store during installation, run `hdblcm` with the parameter `secure_store=localsecurestore`.

3. Select the index for *Install New System*, then select `Enter`.
4. Select the components you would like to install as a comma-separated list, then select `Enter`.
5. Specify the SAP HANA system properties.

   For a list of all system properties, see *System Properties* in Related Information.
6. After specifying all system properties, review the summary, and select *y*.

Results

A single-host SAP HANA system is installed. A log file is available.

After installing the SAP HANA system, you may want to perform configuration tasks. For more information, see *Managing the SAP HANA System After Installation* or the platform lifecycle management section of the *SAP HANA Administration Guide*.

Related Information

System Properties [page 142]
Using the SAP HANA Platform LCM Tools [page 41]
Changeable Default Values for Installation [page 73]
Managing the SAP HANA System After Installation [page 209]
SAP HANA Platform Lifecycle Management
SAP Note 2243156

Activate the Local Secure Store (LSS) [page 197]

6.3 Installing a Multiple-Host System

The SAP HANA database lifecycle manager can be used to install an SAP HANA multiple-host system in one of the program interfaces, and with a combination of parameter specification methods.

A multiple-host system is a system with more than one host, which can be configured as active worker hosts or idle standby hosts. The server software is based on a flexible architecture that enables a distributed installation. This means that load can be balanced between different hosts. The server software has to be installed in a shared file system. This file system has to be mounted by all hosts that are part of the system.

The following graphic shows the file system for a multiple-host installation using a shared file system:

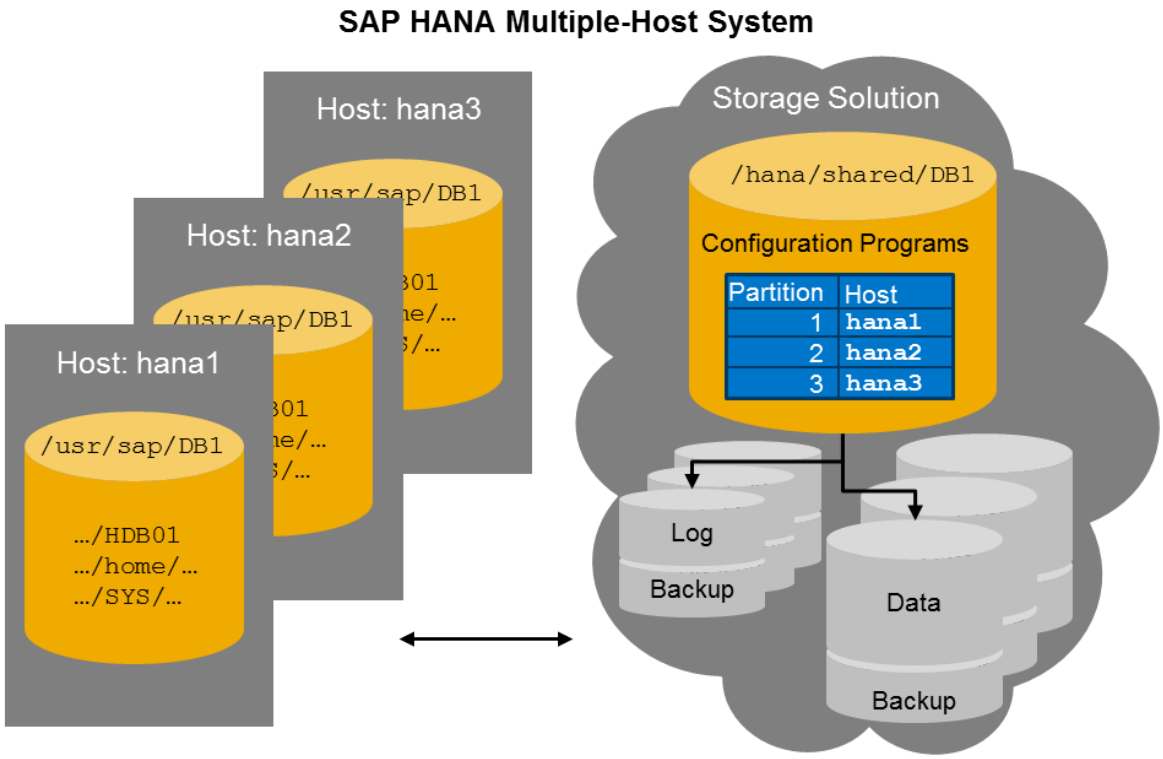

System Configuration	Database Directories	Local Directories
Instance Number: 01	Installation Path: /hana/shared	hana1: /usr/sap/DB1
SAP System ID (SID): DB1	Data Path: /hana/data/DB1	hana2: /usr/sap/DB1
Host Names: hana1, hana2, hana3	Log Path: /hana/log/DB1	hana3: /usr/sap/DB1

To create a multiple-host system after installing a single-host system, hosts must be added to the SAP HANA system. To add hosts to an existing system, use the SAP HANA resident HDBLCM. For more information about host addition, see Related Information or the *SAP HANA Administration Guide*.

The following information only covers building a multiple-host system **during** installation.

Related Information

Adding Hosts to an SAP HANA System

6.3.1 Multiple-Host System Concepts

It is important to review multiple-host system concepts like host grouping and storage options before installing a multiple-host system.

Host Types

When configuring a multiple-host system, the additional hosts must be defined as **worker** hosts or **standby** hosts (worker is default). Worker machines process data; standby machines do not handle any processing and instead wait to take over processes in the case of worker machine failure.

Auto-Failover for High Availability

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. Host auto-failover is a local fault recovery solution that can be used as a supplemental or alternative measure to system replication. One (or more) standby hosts are added to a SAP HANA system, and configured to work in standby mode.

Before installing a multiple-host system, it is important to consider whether high availability is necessary and how hosts should be grouped to ensure preferred host auto-failover. For host auto-failover to be successful, if the active (worker) host fails, the standby host takes over its role by starting its database instance using the persisted data and log files of the failed host. The name server of one of the SAP HANA instances acts as the cluster manager that pings all hosts regularly. If a failing host is detected, the cluster manager ensures that the standby host takes over the role and the failing host is no longer allowed write access to the files (called fencing) so that they do not become corrupted. The crash of a single service does not trigger failover since services are normally restarted by `hdbdaemon`. For more information, see *Setting Up Host Auto-Failover* in the *SAP HANA Administration Guide*.

Host Grouping

Host grouping does not affect the load distribution among worker hosts - the load is distributed among all workers in an SAP HANA system. If there are multiple standby hosts in a system, host grouping should be considered, because host grouping decides the allocation of standby resources if a worker machine fails. If no host group is specified, all hosts belong to one host group called "default". The more standby hosts in one host group, the more failover security.

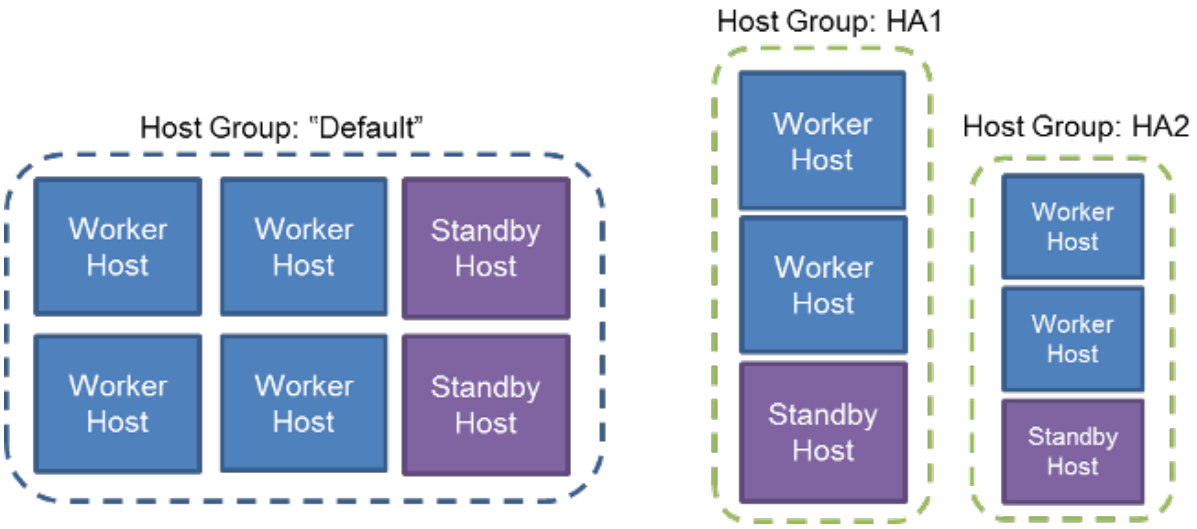

If the standby hosts are each in a different host group, the standby host in the same group as the failing worker host is preferred. Only if no standby host is available in the same host group, the system will try to fail over to a standby host, which is part of another host group. The advantage of this configuration is that in an SAP HANA system with mixed machine resources, similar sized machines can be grouped together. If a small worker host fails, and a small standby in the same group takes over, the processes are moved to a machine with similar resources, which allows processing to continue as usual with optimal resource allocation.

Worker Host Grouping for Warm Data

If you use SAP Business Warehouse to apply a temperature-based data strategy you can significantly optimize the usage of memory and hardware resources by reserving one node of the scaled-out HANA landscape exclusively for warm data. Due to information lifecycle management, multi-temperature strategies are often applied, whereby data is classified by access frequency as either hot, warm, or cold. Depending on this classification and data usage, this data is stored in different memory areas.

A multi-temperature memory strategy may be required for different reasons, for example:

- Storage of historical data
- Clickstream logs for multiple years of Web data and detailed machine logs
- Guidelines for saving company data, such as the need to save all data for at least seven years due to legal reasons.

The standard SAP HANA sizing guidelines allow for a data footprint of 50% of the available RAM. This ensures that all data can be kept in RAM at all times and there is sufficient space for intermediate result sets. These sizing guidelines can be significantly relaxed on the extension group, since 'warm' data is accessed:

- less frequently
- with reduced performance SLAs
- with less CPU-intensive processes
- only partially at the same time.

To implement a multi-temperature memory strategy, you can assign hosts to worker groups. Hot and warm data is then distributed across hosts. To increase performance and memory usage, a slave node is assigned

to a separate *Extension Node*. Unlike the standard nodes (master and slave), the extension node is intended exclusively for data that is not accessed as frequently (warm) as other data (hot).

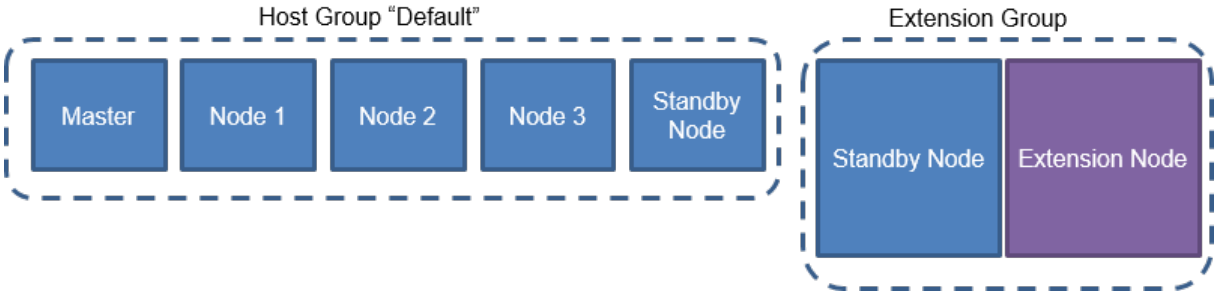

For more information, see *Data Temperature: Extension Node for Business Warehouse* in the *SAP HANA Administration Guide* and SAP Note 2453736.

Storage and File System Options

In single-host SAP HANA systems, it is possible to use local file systems residing on direct-attached internal or external storage devices, such as SCSI hard drives, SSDs, SAN storage, or NAS. However, in order to build a multiple-host system with failover capabilities this is not sufficient. Either the chosen file system type or the SAN Infrastructure along with a SAP HANA functionality capable of disc fencing must ensure the following:

- The standby host has file access to data and log volumes of the failed host.
- The failed worker host no longer has access to write to files - called fencing.

There are two fundamentally different storage configurations which meet the two conditions above: **shared storage devices** or **separate storage devices with failover reassignment**. Do not confuse "shared storage" with the installation directory `/hana/shared` that must be shared across all hosts.

Shared File Systems

A shared storage subsystem, which is accessed using file systems such as NFS or IBM's GPFS, makes it easy to ensure that the standby host has access to all active host files in the system. In a shared storage solution, the externally attached storage subsystem devices are capable of providing dynamic mount points for hosts. Since 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 which is specific for the file system used to access that storage subsystem. An NFSv3 storage solution must be used in combination with the storage connector supplied by the hardware partner. NFSv4 and GPFS storage solutions can optionally be used with a storage connector.

A shared storage system could be configured as in the diagram below, however mounts may differ among hardware partners and their configurations. For more information, see the *SAP HANA Storage Whitepaper* available in SAP Note 1900823 in Related Information.

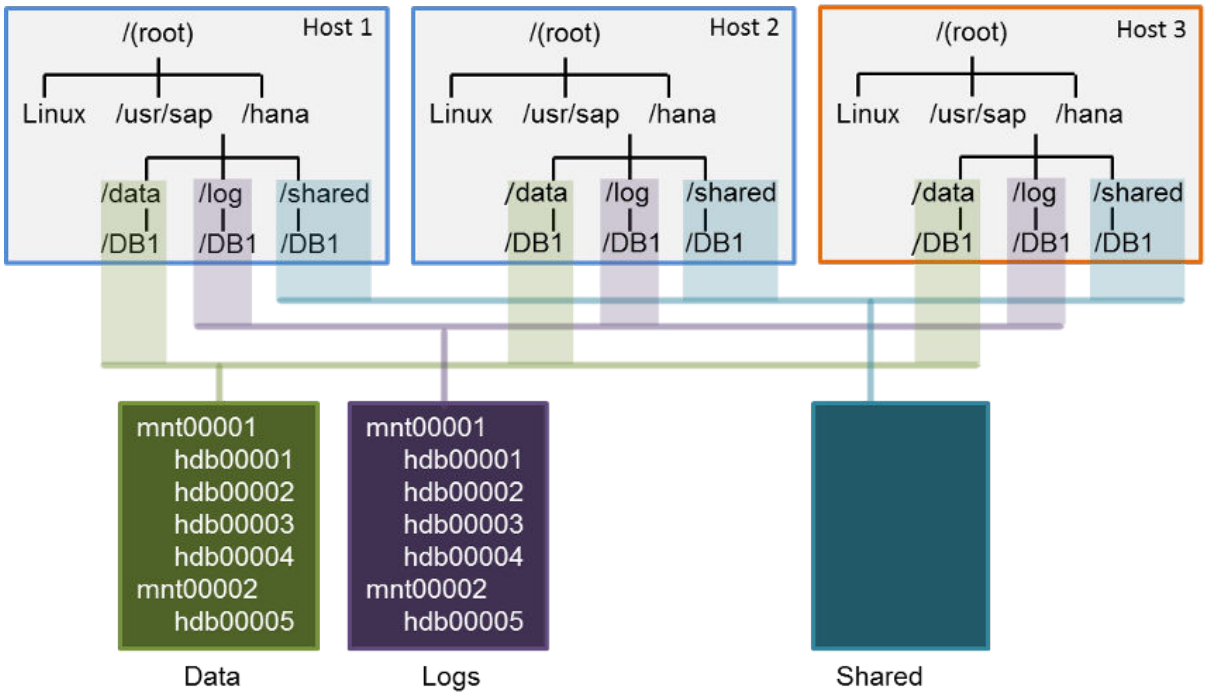

Non-shared Storage

It is also possible to assign every SAP HANA host a separate storage, which has nothing mounted except the shared area. A SAN storage must be used in combination with the SAP Fiber Channel Storage Connector, which SAP HANA offers storage technology vendors. During failover, SAP HANA uses the storage connector API to tell the storage device driver to remount the required data and logs volumes to the standby host and fence off the same volumes from the failed host.

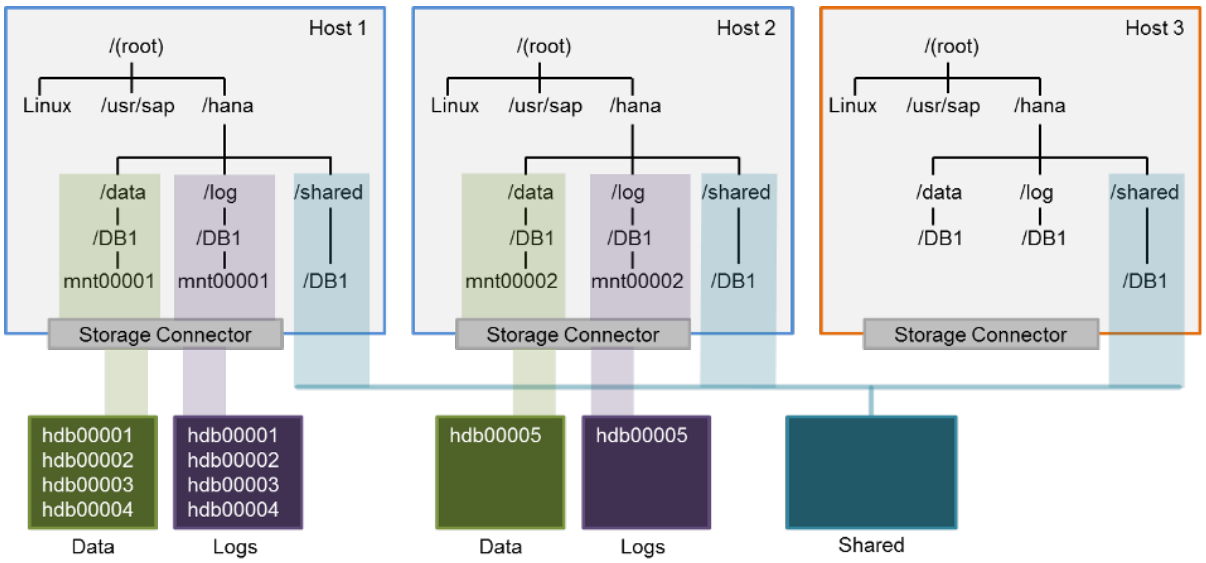

In a non-shared environment, separate storage is used in combination with the storage connector API. For more information about the storage connector API, see the *SAP Fiber Channel Storage Connector Admin Guide* available in SAP Note 1900823 in Related Information.

Related Information

Recommended File System Layout [page 18]
SAP Note 405827
Setting Up Host Auto-Failover
SAP Note 1900823
Extension Node
SAP Note 2453736
More Details – HANA Extension Nodes for BW-on-HANA
workergroup [page 288]
ROUTE_TO Hint

6.3.2 Install a Multiple-Host SAP HANA System Using the Graphical User Interface

A multiple-host SAP HANA system can be installed using the SAP HANA database lifecycle manager (HDBLCM) graphical user interface.

Prerequisites

- You are logged in as root user.
- Any user has read and execute permissions for the directory that contains the installation medium.
- Depending on the storage solution, set the export options `rw,no_root_squash` for the installation directory.
- The operating system administrator (`<sid>adm`) user and other operating system users may exist prior to installation. Make sure that you have the passwords of the existing users, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) will not modify the properties of any existing user or group. The `<sid>adm` user passwords must be identical on all hosts.
- The root user is able to execute graphical applications.
- The SAP HANA system must be installed with its server software on a shared file system:
  - Create an installation directory, e.g. `/hana/shared/`.
  - Depending on the storage solution, set the export options `rw,no_root_squash` for the installation directory.
  - Mount the installation directory on all hosts.

- To install the Local Secure Store (LSS), the installation path must be created manually by the root user before the installation and mounted on all hosts, e.g. `/lss/shared/`.

  > ⓘ Note
  >
  > The local secure store (LSS) must be installed and activated by default during an SAP HANA system installation or update.

- Depending on the desired storage solution, **shared storage devices** or **separate storage devices with failover reassignment**, different configurations apply:
  - In a shared file system, the data files and log files are configured so that they are present and mounted on all hosts, including the primary host.
  - In a system that uses separate storage devices, each host only has access to its own data files and log files.

  The suggested locations for the file systems are as follows:
  - `/hana/data/<SID>`
  - `/hana/log/<SID>`
- (Optional) Additional storage is configured.
- Root user name must be the same for all hosts in a multiple-host system. The password of the root user must be identical on all hosts.
- If the root user name is **not** root, it must be specified as a parameter during installation using the parameter `root_user`.

Context

The following procedure describes the installation of an SAP HANA system in interactive mode by entering parameters interactively. This procedure may also be performed in advanced interactive mode, with parameters entered as call options or from a configuration file. For more information about interaction modes and parameter entry methods, see *Using the SAP HANA Platform LCM Tools* in Related Information.

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. You can add additional tenant databases later using the SAP HANA cockpit. SAP HANA options hosts or host roles can be automatically provisioned to a system that is installed with a single tenant. If the SAP HANA system contains multiple tenant databases, the SAP HANA options host or host role must be manually provisioned to the tenant.

> ⓘ Note
>
> Not all parameters are requested interactively. Some parameters have default values that do not require confirmation in interactive mode. Those parameters must be specified as call options or from a configuration file. For more information about changeable default values, see Related Information.

Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
IBM Power Systems	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`

> ### Note
>
> If you downloaded the components to a different directory, change to the directory where you unpacked the archive.

2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

```
./hdblcmgui
```

The SAP HANA database lifecycle manager graphical user interface appears.

> ### Note
>
> To activate the local secure (LSS) store during installation, run `hdblcmgui` with the parameter `secure_store=localsecurestore`.

3. Select a detected software component or add a software component location by selecting *Add Component Location*. Then select *Next*.
4. Select *Install New System*, then select *Next*.
5. Select the components you would like to install, then select *Next*.
6. Select *Multiple-Host System* as the *System Type*, and select *Add Host* to specify host parameters for the additional hosts.

Field Name	Description
Installation Path	Specifies the path to the SAP mount directory, which can be used as a shared directory between multiple hosts.
Non-standard Shared File System	Specifies a non-standard shared file system, which can be accessed by all hosts during installation. This parameter is typically used when the SID is included in the mountpoint.
Host Name	Specifies the host name of the machine.

Field Name	Description
Role	Specifies the purpose of the SAP HANA host. Although multiple host roles may be assigned, check the corresponding documentation for the SAP HANA option for what configurations are supported in production environments. • Database Worker (worker) - A worker host (default) is used for database processing. • Database Standby (standby) - A standby host is idle and available for failover in a high-availability environment. • Database Elastic Compute Server (compute) - Database elastic compute server • Dynamic Tiering Worker (extended_storage_worker) - Worker host for SAP HANA dynamic tiering • Dynamic Tiering Standby (extended_storage_standby) - Standby host for SAP HANA dynamic tiering • Accelerator for SAP ASE Worker (ets_worker) - Worker host for SAP HANA accelerator for SAP ASE • Accelerator for SAP ASE Standby (ets_standby) - Standby host for SAP HANA accelerator for SAP ASE • Streaming Analytics (streaming) - Host for SAP HANA streaming analytics • XS advanced runtime worker (xs_worker) - Host for SAP HANA XS advanced runtime • XS advanced runtime standby (xs_standby) - Standby host for SAP HANA XS advanced runtime
High-Availability Group	Specifies the host group ID for failover scenarios. If undefined, the host group is named "default".
Worker Group	Specifies the worker group of the host. If undefined, the worker group is named "default". If you are using extension node for Business Warehouse, you must name the worker group "worker_dt".
Storage Partition	Specifies the storage partition number, which is a logical role number assigned to non-shared storage devices in a storage connector API. Standby hosts do not have a storage partition.

7. Specify the SAP HANA system properties.

   For a list of all system properties, see *System Properties* in Related Information.
8. After specifying all system properties, review the summary, and select *Install*.

Results

A multiple-host SAP HANA system is installed. A log file is stored in the following path:

```
/var/tmp/hdb_<SID>_<action>_<time stamp>
```

After installing the SAP HANA system, you may want to perform configuration tasks. For more information, see *Managing the SAP HANA System After Installation* or the platform lifecycle management section of the *SAP HANA Administration Guide*.

Related Information

System Properties [page 142]
Using the SAP HANA Platform LCM Tools [page 41]
Changeable Default Values for Installation [page 73]
Managing the SAP HANA System After Installation [page 209]
SAP HANA Platform Lifecycle Management
SAP Note 2243156
Local Secure Store (LSS) [page 28]
lss_inst_path [page 264]
Activate the Local Secure Store (LSS) [page 197]

6.3.3 Install a Multiple-Host SAP HANA System Using the Command-Line Interface

A multiple-host SAP HANA system can be installed using the SAP HANA database lifecycle manager (HDBLCM) command-line interface.

Prerequisites

- You are logged in as root user.
- Any user has read and execute permissions for the directory that contains the installation medium.
- Depending on the storage solution, set the export options `rw,no_root_squash` for the installation directory.
- The operating system administrator (`<sid>adm`) user and other operating system users may exist prior to installation. Make sure that you have the passwords of the existing users, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) will not modify the properties of any existing user or group. The `<sid>adm` user passwords must be identical on all hosts.
- The SAP HANA system must be installed with its server software on a shared file system:
  - Create an installation directory, e.g. `/hana/shared/`.
  - Depending on the storage solution, set the export options `rw,no_root_squash` for the installation directory.
  - Mount the installation directory on all hosts.
  - To install the Local Secure Store (LSS), the installation path must be created manually by the root user before the installation and mounted on all hosts, e.g. `/lss/shared/`.

    > **Note**
    >
    > The local secure store (LSS) must be installed and activated by default during an SAP HANA system installation or update.

- Depending on the desired storage solution, **shared storage devices** or **separate storage devices with failover reassignment**, different configurations apply:
  - In a shared file system, the data files and log files are configured so that they are present and mounted on all hosts, including the primary host.
  - In a system that uses separate storage devices, each host only has access to its own data files and log files.

  The suggested locations for the file systems are as follows:
  - `/hana/data/<SID>`
  - `/hana/log/<SID>`
- (Optional) Additional storage is configured.
- Root user name must be the same for all hosts in a multiple-host system. The password of the root user must be identical on all hosts.
- If the root user name is **not** root, it must be specified as a parameter during installation using the parameter `root_user`.

Context

The following procedure describes the installation of an SAP HANA system in interactive mode and entering parameters interactively. This procedure may also be performed in advanced interactive mode or batch mode, with parameters entered as call options or from a configuration file. For more information about interaction modes and parameter entry methods, see *Using the SAP HANA Platform LCM Tools* in Related Information.

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. You can add additional tenant databases later using the SAP HANA cockpit. SAP HANA options hosts or host roles can be automatically provisioned to a system that is installed with a single tenant. If the SAP HANA system contains multiple tenant databases, the SAP HANA options host or host role must be manually provisioned to the tenant.

> ⓘ Note
>
> Not all parameters are requested interactively. Some parameters have default values that do not require confirmation in interactive mode. Those parameters must be specified as call options or from a configuration file. For more information about changeable default values, see Related Information.

Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`

Option	Description
IBM Power Systems	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`

> **Note**
>
> If you downloaded the components to a different directory, change to the directory where you unpacked the archive.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblcm
```

> **Note**
>
> To activate the local secure (LSS) store during installation, run `hdblcm` with the parameter `secure_store=localsecurestore`.

> **Note**
>
> A *Non-standard Shared File System* which can be accessed by all hosts during installation can be set during installationwith the `checkmnt` parameter. This parameter is typically used when the SID is included in the mountpoint.

3. Select the index for *Install New System*, then select `Enter`.
4. Select the components you would like to install as a comma-separated list, then select `Enter`.
5. Specify the installation path, and the local host name:

SAP HANA System Properties

Field Name	Description
Installation Path	Specifies the path to the SAP mount directory, which can be used as a shared directory between multiple hosts.
Local Host Name	Specifies the host name of the machine.

6. Select *y* to the question *Do you want to add additional hosts to the system?*, and enter the following details for the additional host:

Field Name	Description
Host Name	Specifies the host name of the machine.

Field Name	Description
Role	Specifies the purpose of the SAP HANA host. Although multiple host roles may be assigned, check the corresponding documentation for the SAP HANA option for what configurations are supported in production environments. • Database Worker (worker) - A worker host (default) is used for database processing. • Database Standby (standby) - A standby host is idle and available for failover in a high-availability environment. • Database Elastic Compute Server (compute) - Database elastic compute server • Dynamic Tiering Worker (extended_storage_worker) - Worker host for SAP HANA dynamic tiering • Dynamic Tiering Standby (extended_storage_standby) - Standby host for SAP HANA dynamic tiering • Accelerator for SAP ASE Worker (ets_worker) - Worker host for SAP HANA accelerator for SAP ASE • Accelerator for SAP ASE Standby (ets_standby) - Standby host for SAP HANA accelerator for SAP ASE • Streaming Analytics (streaming) - Host for SAP HANA streaming analytics • XS advanced runtime worker (xs_worker) - Host for SAP HANA XS advanced runtime • XS advanced runtime standby (xs_standby) - Standby host for SAP HANA XS advanced runtime
Host Failover Group	Specifies the host group ID for failover scenarios. If undefined, the host group is named "default".
Worker Group	Specifies the worker group of the host. If undefined, the worker group is named "default". If you are using extension node for Business Warehouse, you must name the worker group "worker_dt".
Storage Partition	Specifies the storage partition number, which is a logical role number assigned to non-shared storage devices in a storage connector API. Standby hosts do not have a storage partition.

7. Specify the SAP HANA system properties.

   For a list of all system properties, see *System Properties* in Related Information.
8. After specifying all system properties, review the summary, and select *y*.

Results

A multiple-host SAP HANA system is installed. A log file is stored in the following path:

```
/var/tmp/hdb_<SID>_<action>_<time stamp>
```

After installing the SAP HANA system, you may want to perform configuration tasks. For more information, see *Managing the SAP HANA System After Installation* or the platform lifecycle management section of the *SAP HANA Administration Guide*.

Related Information

System Properties [page 142]
Using the SAP HANA Platform LCM Tools [page 41]
Changeable Default Values for Installation [page 73]
checkmnt [page 243]
Managing the SAP HANA System After Installation [page 209]
SAP HANA Platform Lifecycle Management
SAP Note 2243156
Local Secure Store (LSS) [page 28]
lss_inst_path [page 264]
Activate the Local Secure Store (LSS) [page 197]

6.4 Installing an SAP HANA System Including the XS Advanced Runtime

The SAP HANA database lifecycle manager can be used to install the XS Advanced Runtime.

> **Note**
>
> From SPS 11, SAP HANA includes an additional run-time environment for application development: SAP HANA extended application services (XS), advanced model. SAP HANA XS 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 (XS), classic model. SAP recommends that customers and partners who want to develop new applications use SAP HANA XS advanced model. If you want to migrate existing XS classic applications to run in the new XS advanced run-time environment, SAP recommends that you first check the features available with the installed version of XS advanced; if the XS advanced features match the requirements of the XS classic application you want to migrate, then you can start the migration process.

SAP HANA extended application services, advanced model provides a comprehensive platform for the development and execution of native data-intensive applications. It requires the installation of the XS advanced runtime.

In support of this data-integrated application paradigm, SAP HANA Extended Application Services provide a comprehensive set of embedded services that provide end-to-end support for Web-based applications. This includes a lightweight web server, configurable OData support, JavaScript execution and, of course, full access to SQL and SQLScript.

These SAP HANA Extended Application Services are provided by the SAP HANA XS server, which provides lightweight application services that are fully integrated into SAP HANA. It allows clients to access the SAP HANA system via HTTP. Controller applications can run completely natively on SAP HANA, without the need for an additional external application server.

The application services can be used to expose the database data model, with its tables, views and database procedures, to clients. This can be done in a declarative way using OData services or by writing native

application-specific code that runs in the SAP HANA context . Also, you can use SAP HANA XS to build dynamic HTML5 UI applications.

For more information about SAP HANA XS advanced, see the *SAP HANA Developer Guide (For SAP HANA XS Advanced Model)*.

Downloading XS Advanced from the SAP Support Portal

SAP HANA Extended Application Services, advanced model, is available not only on the SAP HANA media but also as a separate component on the SAP Support Portal. Users with the required S-User ID can download the latest version of XS advanced component in the package `SAP EXTENDED APP SERVICES 1` from the following location:

SAP Support Portal > *Software Downloads [Downloads]* > *SUPPORT PACKAGES & PATCHES* > *By Alphabetical Index (A-Z)* > *H* > *SAP HANA PLATFORM EDITION*:

- *SAP HANA PLATFORM EDITION 2.0* > *SAP EXTENDED APP SERVICES 1*

> → Tip
>
> SAP HANA Extended Application Services, advanced model, is backwards compatible; you can provide access to new features by installing the latest version of the XS advanced component even on older versions of SAP HANA. To download the package `SAP EXTENDED APP SERVICES 1`, see *SAP Software Download Center* in *Related Information* below.

Installing XS Advanced from the XS Advanced Installation Media

Although XS Advanced is included on the SAP HANA revision release media for installation at the same time as the SAP HANA database, you can also use the XS Advanced installation media (also known as the *XS Advanced Collection*) to install or update SAP HANA XS Advanced and any additional components at any time independently from the SAP HANA database.

For more information about where to find and download the XS Advanced Installation Media, see SAP Note 2711421 or *Related Information below*.

Related Information

SAP Software Download Center (Logon required)
SAP Note 2711421
SAP Note 2347931
SAP HANA Developer Guide for XS Advanced Model (SAP Web IDE)

6.4.1 System Concepts for XS Advanced Runtime Installations

It is important to review single-host and multiple-host system concepts before installing a SAP HANA system that includes the XS advanced runtime.

Host Types

When configuring a multiple-host system with XS advanced runtime, the additional hosts must be defined as **xs_worker** hosts or **xs_standby** hosts. Worker machines run XS applications; standby machines do not handle any processing and instead just wait to take over processes in the case of worker machine failure.

Automatic Role Assignment

XS advanced runtime host roles can be assigned automatically during installation. Assign worker and standby roles to all hosts and choose *Automatically assign XS Advanced Runtime roles to the hosts with database roles* when prompted. The installer will then assign the role xs_worker to every worker host and xs_standby to every standby host. To create a multiple-host system with dedicated xs_worker and xs_standby hosts, assign host roles to each host individually during installation. Do not choose the option to assign XS Advanced host roles automatically.

Single-Host Setup

A single-host system is the simplest system installation type. It is possible to run an SAP HANA system entirely on one host and then scale the system up as needed. The host must have the database worker and xs_worker host roles assigned.

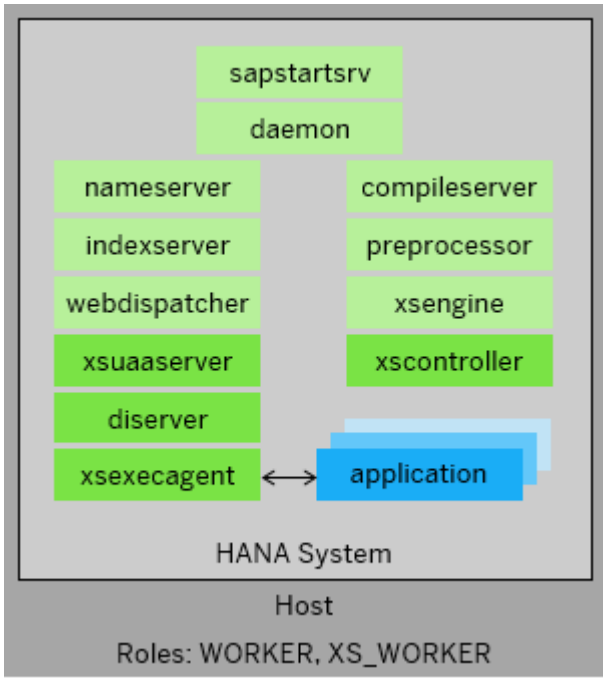

Multi-Host Setup

A multiple-host system is a system with more than one host, which can be configured as active worker hosts or idle standby hosts. The server software is based on a flexible architecture that enables a distributed installation. This means that load can be balanced between different hosts. The server software has to be installed in a shared file system. This file system has to be mounted by all hosts that are part of the system.

In a basic multi-host system all worker hosts also act as XS worker hosts.

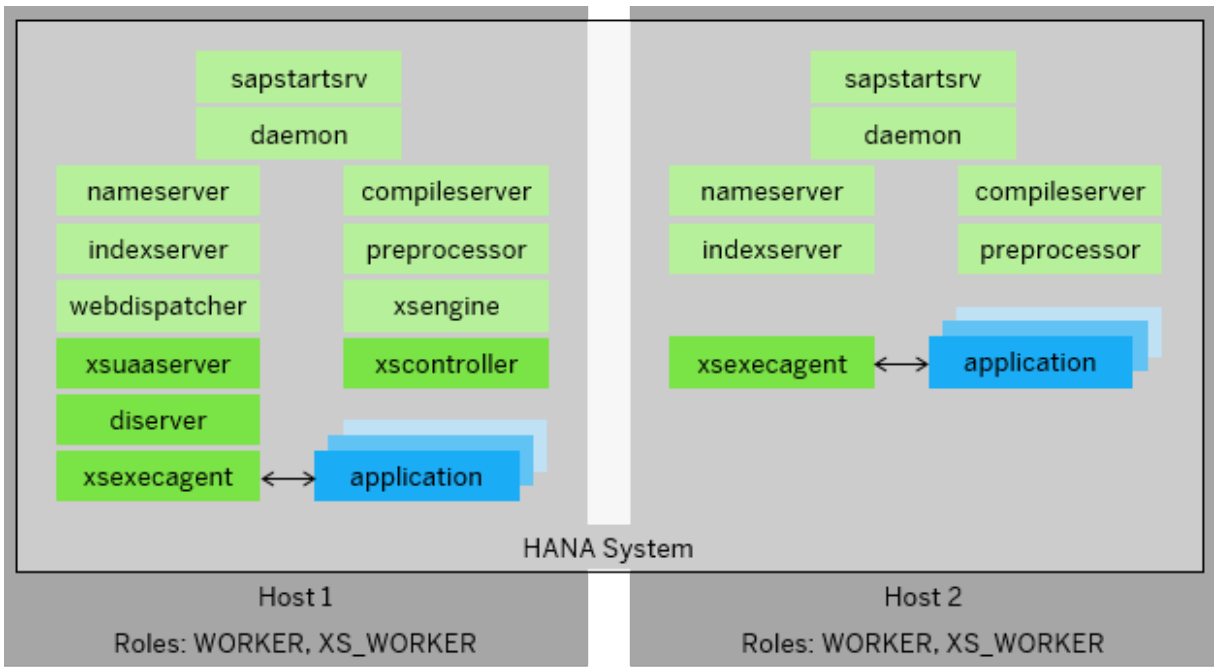

Multi-Host System with Standby Host

A dedicated standby host can be added to the basic multi-host setup. It will act as a failover host in the case of worker machine failure for both worker and XS worker hosts.

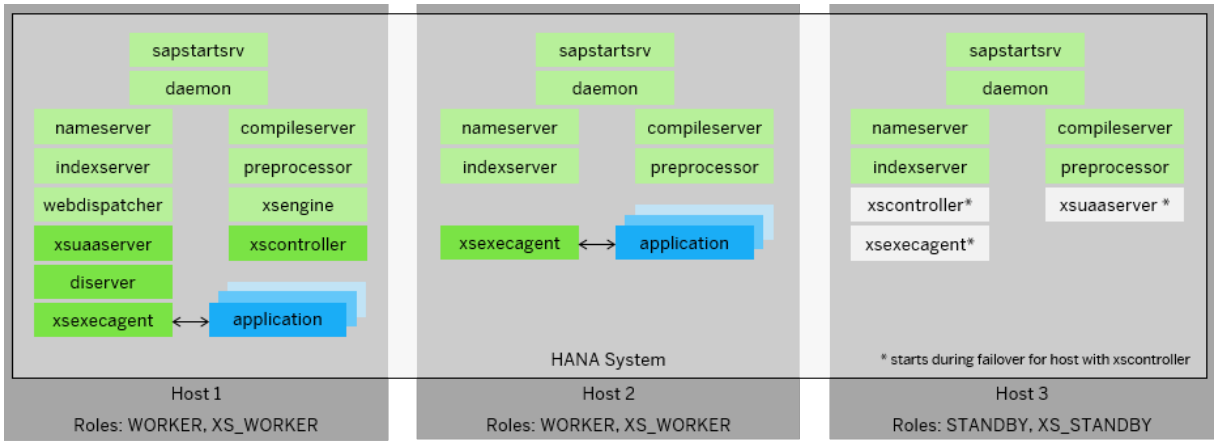

Dedicated XS advanced runtime Host

SAP HANA extended application services, advanced model provides a comprehensive platform for the development and execution of native data-intensive applications. Worker and XS worker host roles can therefore be assigned to different hosts. To create a multiple-host system with a dedicated XS worker host, hosts must be assigned manually during installation.

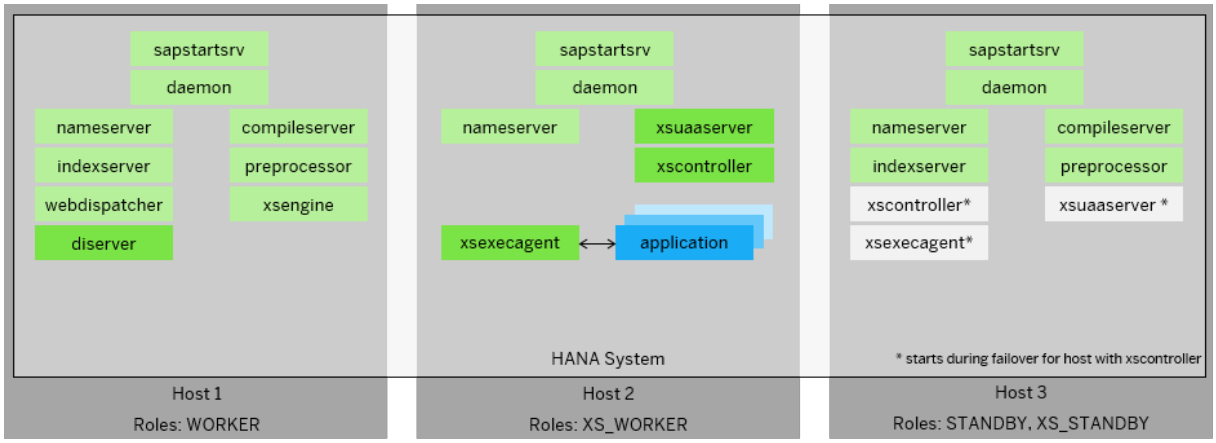

Tenant Database Installation

Starting from SAP HANA SPS03 (Rev. 34) and XS advanced model SPS03 (Rev. 1), you can choose the location of the XS advanced platform data during installation, for example, in the (default) System database or in a specific tenant database.

Installing XS Advanced in a tenant database allows you to keep all XS Advanced related data in a single database. Compared to installing XS Advanced in the system database, this setup helps to overcome restrictions of backup and recovery and when moving XS Advanced related data to other SAP HANA systems. Please note, that when the tenant database containing the XS Advanced platform data is deleted, XS Advanced services will not continue to work and must be stopped. For more information, see *Installing XS Advanced in a Tenant Database*.

Default Domain Configuration

Since the XS advanced default domain is the domain by which the XS advanced platform can be reached, the XS advanced default domain must be mapped within the Domain Name System (DNS). For example, it must point to the SAP HANA host which is running the '`xscontroller`' and '`xsuaaserver`' services.

You can configure the default domain by using the property `default_domain` in the `xscontroller.ini` file, as illustrated in the following example:

Sample Code

```
default_domain=example.org
```

Routing Mode Configuration

During the installation of XS advanced, you have to select one of the following routing modes:

- Port-based routing mode
  The port-name routing mode is **not** recommended for productive use.
- Hostname-based routing mode
  Application URLs in host-name routing mode are more user-friendly since they contain the name of the application.

Note

It is **not** recommended or supported to change the routing mode configuration after installation.

The routing mode you choose has an influence on the URLs by which applications are exposed. For more details about which end-points are opened for each routing mode, see *Public Endpoints* in *Related Information* below.

Port-based Routing Mode

In port-based routing mode, different applications use different application ports while sharing the same domain name. In this routing mode, the URL of an XS advanced application takes the following form, where `<application port>` is a port number out of the router port range:

Sample Code

```
<protocol>://<default domain>:<application port>
```

For example, assuming the applications **application1** and **application2** are deployed, XS advanced exposes these applications using the following URLs in port-based routing mode:

- **application1**
  https://example.org:51000
- **application2**
  https://example.org:51001

The XS advanced installation automatically sets the property `routing_mode` in the file `xscontroller.ini`. If this property does not exist or is set to ports, then you have configured port-based routing mode. In `xscontroller.ini`, the router port range is determined by the properties `router_portrange_begin` and `router_portrange_end` , and the default port range is 51000-51500. For more information about configuration parameters for XS advanced, see *XS Advanced System Configuration Parameters* in *Related Information* below.

The advantage of the port-based approach to routing is that it allows you to use a single, non-wildcard DNS entry for all application and platform endpoints. However, it is important to note the following disadvantages:

- Browsers share cookies among all endpoints on a particular domain.
  This means that one application might "see" cookies set by another application. For this reason, from a security point of view, port-based routing mode is not recommended for productive use!
- Changing the router port range only affects newly created routes.
  Existing routes with assigned ports are not automatically re-assigned in this case.
- Setting up XS advanced behind a reverse proxy is more complex in this mode, as traffic on several ports needs to be forwarded by the reverse proxy.

> ### Note
>
> For more information about setting up XS advanced behind a reverse proxy, see *Related Information* below.

Hostname-based Routing Mode

In hostname-based routing mode, different applications use different sub-domains (hostnames) while sharing a single port. In this routing mode, the URL of an XS advanced application takes the following form:

> ### Sample Code
>
> ```
> <protocol>://<hostname>.<default domain>:<router port>
> ```

For example, assuming the applications **application1** and **application2** are deployed, XS advanced exposes these applications using the following URLs in hostname-based routing mode:

- **application1**
  https://application1.example.org:443
- **application2**
  https://application2.example.org:443

The XS advanced installation automatically sets the property `routing_mode` in the `xscontroller.ini` file. If this property is set to `hostnames`, then you have configured hostname-based routing mode. In the `xscontroller.ini`file, the shared `<router port>` is determined by the property `router_port`. The default router port is `3<instance nr>33`, as described in *XS Advanced System Configuration Parameters* in *Related Information* below.

> **Note**
>
> When setting a router port lower than 1024 in Unix, make sure you remember to configure the `icmbnd` binary for the SAP HANA XS advanced Web Dispatcher correctly in the directory `<SAP HANA Installation>/xs/router/webdispatcher` on your SAP HANA host, as described in *Binding Ports < 1024* in *Related Information* below.

The following advantages apply to using hostname-based routing:

- Using hostname routing mode is recommended for productive use.
  In this mode, browsers do not share cookies amongst different applications.
- In this mode, application URLs are user friendly; they contain the name of the application.
- It is possible to change the router port after installing XS advanced:
  Changing the value of the setting `router_port` in `xscontroller.ini` will also change existing application URLs when XS advanced is restarted.

> **Note**
>
> To allow the XS advanced Platform Router to dispatch requests to these URLs to the target application, a wildcard DNS entry is required for the XS advanced default domain.

In the following example, the default domain "`example.org`" is mapped to the IP address 10.10.10.10, which would be the SAP HANA host running the '`xscontroller`' and '`xsuaaserver`' services.

> **Sample Code**

```
example.org IN A 10.10.10.10
*.example.org IN CNAME example.org
```

> **Restriction**
>
> With hostname-based routing, the default domain must be fully-qualified (it must contain at least one separator dot).

Implications for Backup and Restore

When making the decision about the routing mode to use, bear in mind the following implications for your backup and restore strategy:

- Restoring a backup from a system with port-based routing mode into a system with hostname-based routing mode (or vice versa) will take over the routing mode of the source system.
- Restoring a backup from a system with port-based routing mode will take over the application ports from the source system, independently of the router port range in the target system. However, application routes adapt to the default domain setting in target system.
- Restoring a backup from a system with hostname-based routing mode will take over application host names from the source system. However, application routes adapt to the router port and default domain settings in the target system.

File-System Setup

By default, XS advanced uses `/hana/shared/xs` as the root directory for static files such as binaries but also for all the files that it creates dynamically. The following directories are very likely to experience very high file-system write-throughput:

- `app_working`
  Contains the file-system sandboxes for application instances
- `fss`
  Contains the files of file-system service instances

```
/hana/shared/xs
   |-- app_working
   |-- controller_data/fss
   |-- ...
```

It is possible that the `/hana/shared` directory does not provide optimal performance for XS advanced, for example when `/hana/shared` points to a network share (for example, NFS). This can have a huge impact on the startup and staging performance of XS advanced applications.

Increasing the Performance of Application Startup and Staging

To increase application start performance, you can configure a different location for the `app_working` directory, for example, by setting the property `basepath_xsa_appworkspace` in the `global.ini` file, as shown in the following example:

```
basepath_xsa_appworkspace = /local_ssd/app_working
```

This location you specify should point to a local device with high throughput capability, for example, a local SSD, as illustrated in the following example:

```
/local_ssd/
   |-- app_working
/hana/shared/xs
   |-- controller_data/fss
   |-- ...
```

Bear in mind the following points:

- It is not necessary to use a device that is shared between the hosts in a multi-host setup.
- If the directory does not exist on a host, XS advanced will try to create it during startup. Make sure, that the `<sidadm>` user has the permissions required to create a directory in the specified location.
- Make sure that the directory has executable right for others on the complete path. This is necessary because applications running as restricted operating system users will access the configured location.
- You can run the command `XSA diagnose` to check if your current configuration provides enough performance.
- See also: "*How to configure the application working directory to improve application start performance?*" in SAP Note 2596466 (*FAQ: SAP HANA XS Advanced*).

Reconfiguring the Location of the File-System Service

It is also possible to re-configure the location of the file-system service to a custom network share, as illustrated in the following example:

```
/local_ssd/
```

```
    |-- app_working
/custom_nfs/
    |-- fss
/hana/shared/xs
    |-- ...
```

Bear in mind the following points:

- The new location needs to be shared between all hosts in a multi-host setup.
- For more information, see *Configuring the XS Advanced File-System Service* in *Related Information* below.

Related Information

Installing a Single-Host System [page 79]
Installing a Multiple-Host System [page 84]
Setting Up Host Auto-Failover
Install an SAP HANA System Including XS Advanced Runtime Using the Graphical Interface [page 116]
Install an SAP HANA System Including XS Advanced Runtime Using the Command-Line Interface [page 121]
SAP Note 2245631
Maintaining the SAP HANA XS Advanced Model Run Time
Installing XS Advanced in a Tenant Database [page 108]
Public Endpoints (SAP HANA Security Guide)
XS Advanced System Configuration Parameters (SAP HANA Administration Guide for SAP HANA Platform)
Setting Up the XS Advanced Runtime Behind a Reverse Proxy [page 127]
Binding Ports < 1024 on UNIX (SAP Web Dispatcher Guide)
Configuring the XS Advanced File-System Service (SAP HANA Administration Guide for SAP HANA Platform)

6.4.2 XS Advanced Database Setup Options

Configure the location of XS-advanced-related data.

Starting from SAP HANA 2.0 SPS 03 (Rev. 34) and the latest XS advanced revision, it is possible to choose the location of the XS advanced platform data during installation. The location of the XS advanced platform data has implications on the following:

- The location of SAP HANA users (either platform or XS advanced application users)
- Backup and recovery of the XS advanced platform

> → Tip
>
> For more information about where to find and download the latest revision of the XS advanced run-time platform, see *Installing an SAP HANA System Including the XS Advanced Run Time* in *Related Information* below.

When you are installing and setting up the XS advanced-model platform in an SAP HANA MDC system, the setup options you choose have implications on the location on the data stored and used by the XS advanced platform. The following table lists the type of data involved:

XS Advanced Data Types

XS Advanced Data Type	Description
Users	SAP HANA users who access the XS advanced platform, for example, with the following tools: • The `xs` command-line-interface (CLI) client • The *XS Advanced Cockpit* • XS advanced business-application users
Platform data	Data uploaded to the XS advanced platform, for example: applications, build packs, application run time environments, application droplets Metadata, for example, describing the current state of applications
System application data	Data stored by **System** applications, for example, the audit-log service, the deploy service, and the product installer
Custom application data	Data stored by **custom** applications deployed to the XS advanced platform

The information provided in this section explains how to install XS advanced model into either the system database or a tenant database; discusses the various setup options, and describes the implications of the configuration choices available.

Related Information

Installing XS Advanced in the System Database [page 113]
Installing XS Advanced in a Tenant Database [page 108]
SAP Note 2542036
SAP Software Download Center (Logon required)
Installing an SAP HANA System Including the XS Advanced Runtime [page 97]

6.4.2.1 Installing XS Advanced in a Tenant Database

Install XS Advanced in a tenant database.

As of SAP HANA SPS05, XS advanced is installed in the default tenant database by default. During first-time installation of SAP HANA and XS advanced, the SAP HANA database lifecycle manager (HDBLCM), displays the following system prompt in the command output:

Output Code

`hdblcm` installation prompt for XS advanced data location

```
Install XS Advanced in the default tenant database? (y/n) [y]: y
```

If you are installing XS advanced in an existing SAP HANA system that already contains tenant databases, you can specify the name of an existing tenant database as the target location to install the XS advanced persistence data, as illustrated in the following example command-line prompt:

Output Code

Selecting the Target Tenant Database for XS Advanced Data

```
Enter System Database User Name [SYSTEM]:
Enter System Database User (SYSTEM) Password:
  Index | XS Advanced Tenant Database
  ----------------------------------
  1     | SYSTEMDB
  2     | MY_TENANT1
  3     | MY_TENANT2
  [...]
Select XS Advanced Tenant Database / Enter Index [1]: 2
```

Note

This setup is also automatically used when migrating an SAP HANA system from Single Database Container (SDC) configuration to a Multiple Database Container (MDC) configuration, where XS advanced is already installed.

To enable you to back up and restore XS-advanced-related content independently, it is recommended to install XS advanced in a different tenant database than other SAP applications (for example. S/4 HANA). as illustrated in the following diagram:

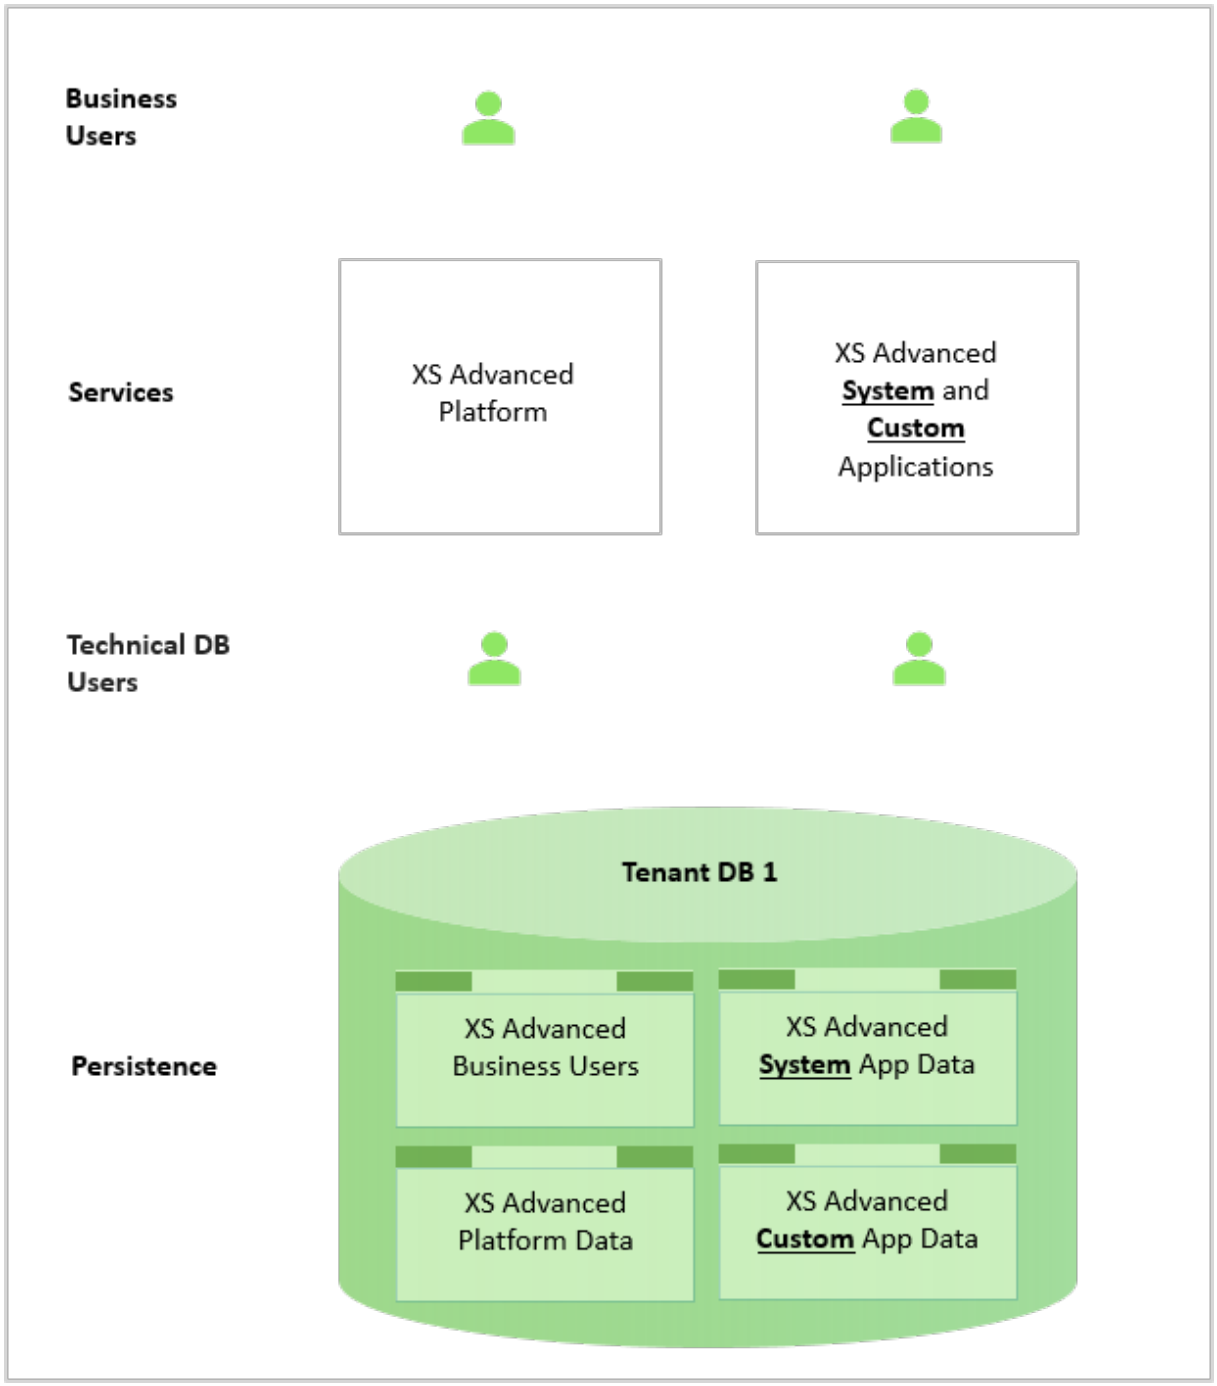

XS Advanced Installed in a Single Tenant Database

After installation, you can register additional tenant databases for XS advanced, as illustrated in the following diagram:

> **Note**
>
> This setup enables you to separate data from different XS advanced applications in different tenant databases. For more information, see *Maintaining Tenant Databases in XS Advanced* in *Related Information* below.

XS Advanced Installed in a Tenant Database with Additional Tenants

Data Location

Installing XS advanced in a tenant database allows you to keep all XS advanced related data in a single database. Compared to installing XS advanced in the system database, this tenant-database setup helps you to overcome the restrictions mentioned in the System-database setup concerning backup and recovery and moving XS-advanced-related data to another SAP HANA system.

> **Note**
>
> If the tenant database containing the XS advanced platform data is deleted, XS advanced services will not continue to work. For this reason, the XS advanced services must be stopped.

Implications for Backup and Recovery

Installing XS advanced into a tenant database helps you to overcome restrictions that exist for backup and recovery compared to installing XS advanced in the System database. Where XS advanced is installed in a tenant database, note the following important points:

- Keeping all XS-advanced-related data in a single tenant database ensures the consistency of the data contained in the backup. For this reason, neither XS advanced services nor XS advanced applications need to be stopped during the backup operation, whereas these services and applications would need to be stopped if XS advanced is installed in the System database, to ensure the consistency of the backup.
- In this tenant-database setup, the restrictions regarding the recovering of the system database described in the system-database scenario do not apply.
- Make sure you back up and restore together at the same time all tenant databases that registered for XS advanced.

> ### Note
>
> If you configure custom XS advanced applications to use additional tenant databases, the cross-database consistency of the backup cannot be ensured if XS advanced continues to run during the backup operation. For this reason, it is recommended to stop XS advanced services during the backup operation.

For more information about backup and recovery tasks, see *Backup and Recovery in XS Advanced* in *Related Information* below.

Implications for Moving Tenant Databases

Installing XS advanced in a tenant database enables the copying or moving of XS-advanced-related data to other SAP HANA systems. You can choose to move all tenant databases or only those tenant databases that are used by XS advanced; tenant databases in the same system that are not registered for use by XS advanced do not have to be moved. Bear in mind, however, that if you move the tenant databases used by XS advanced, the same version of XS advanced must be installed on both the source and target SAP HANA system. In addition, **all** tenant databases registered for XS advanced (that is; containing XS-advanced-related data) must be moved together to the target system.

> ### Note
>
> For more information about how to find out which tenant databases contain XS advanced related data and how to select the XS advanced platform data in the target system, see Displaying the XS Advanced Database Setup [page 112] and Selecting the XS Advanced Database After Backup and Recovery [page 112] below.

Displaying the XS Advanced Database Setup

To find out which tenant databases are relevant to backup XS advanced, open a command shell, log on to the SAP HANA system as `<sid>adm`, and run the following command:

```
XSA list-tenants
```

The `list-tenants` command displays a list of all tenant databases that contain XS advanced related data; these databases must be included in an XS advanced backup. The `list-tenants` command also displays information indicating if the tenant database contains XS advanced **platform** data, as illustrated in the following excerpt of the command output:

{·} Output Code

`XSA list-tenants` Command Output

```
...
  DB name: MYTENANT
  [...]
  XS advanced platform persistence: YES
  [...]
```

→ Tip

The tenant database containing the XS advanced platform data is the database you need to select after restoring XS advanced data to an SAP HANA system.

Selecting the XS Advanced Database After Backup and Recovery

After restoring XS-advanced-related data to an SAP HANA system, for example, during a tenant-database move or after a backup-and-recovery operation, you have to inform the system-wide XS advanced services where to find the XS advanced platform data that was recovered. You can use the command `XSA select-xsa-runtime-db` with the -n (database name) option after logging on the SAP HANA system as `<sid>adm` in a command shell, as shown in the following example:

```
XSA select-xsa-runtime-db -n <tenant_database_name>
```

→ Tip

Use the name of the tenant database displayed in the output of the `XSA list-tenants` command.

When using the command `XSA select-xsa-runtime-db`, bear in mind the following important points:

- All XS advanced services are restarted
- After restarting, the XS advanced services use the XS advanced platform data that was recovered
- XS advanced platform data residing in a different database than the one selected in this step is ignored but not deleted.

Related Information

XS Advanced Database Setup Options [page 106]
Installing XS Advanced in the System Database [page 113]
Maintaining Tenant Databases in XS Advanced
SAP Note 2542036
Backup and Recovery in XS Advanced (SAP HANA Admin Guide)

6.4.2.2 Installing XS Advanced in the System Database

Install XS advanced in the system database.

The XS advanced installation provides the option to keep XS advanced platform data in the system database, while custom applications are configured to use a tenant database as storage, as illustrated by the following diagram:

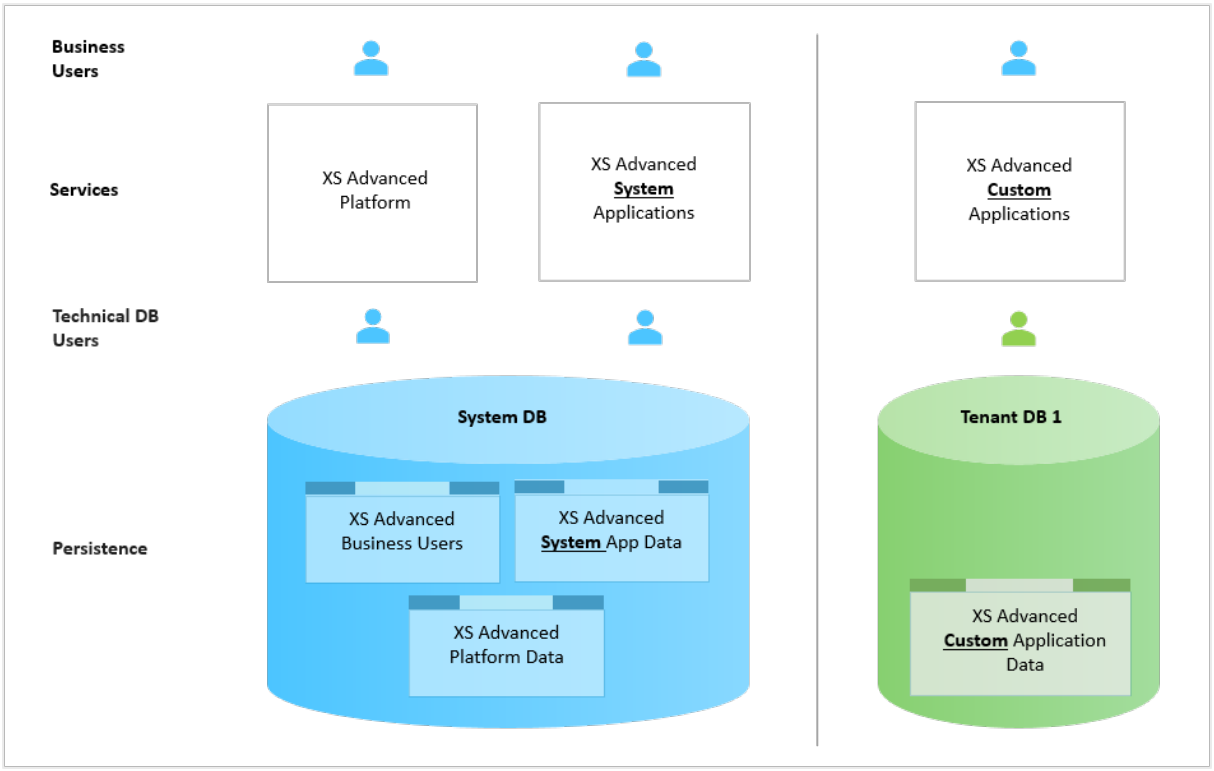

XS Advanced Data in the System Database (Default)

After installation, XS advanced allows you to register additional tenant databases in the SAP HANA system, as illustrated in the following diagram, which enables you to ensure that data belonging to different XS advanced applications can be stored in separate tenant databases:

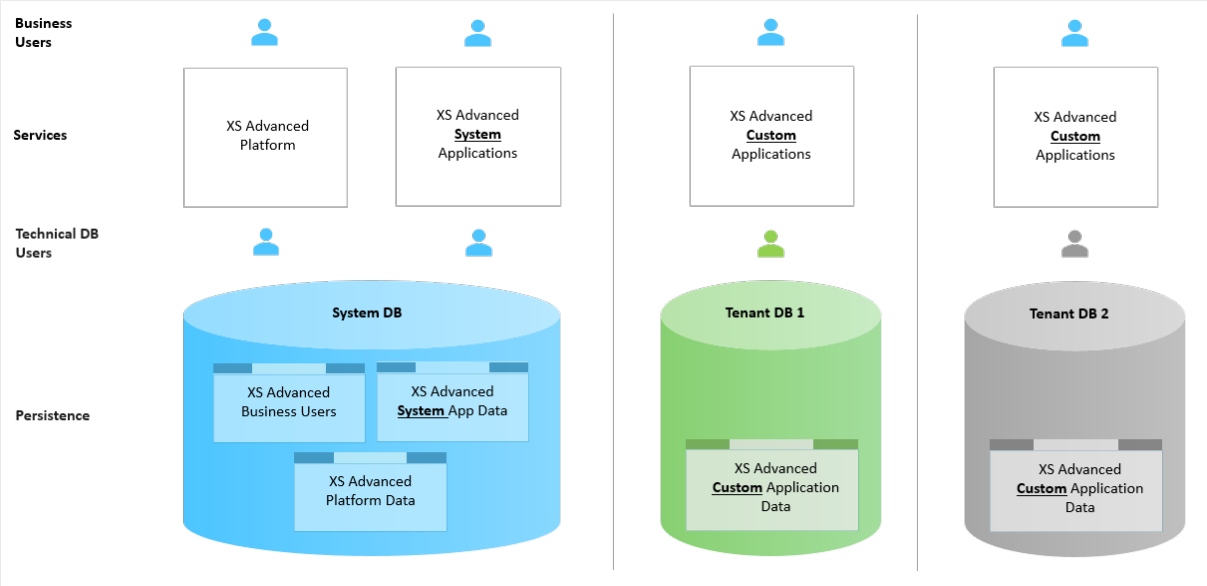

XS Advanced Applications using Additional Tenant Databases

> ### Note
>
> For more information about tenant databases, see *Maintaining Tenant Databases in XS Advanced* in *Related Information* below.

Data Location

In this setup, XS advanced platform data is kept in the System database. System applications store their data in the system database, too. Custom applications store their data in the default tenant database. This way, XS advanced platform data and custom application data are strictly separated. As a consequence, even when all tenant databases are deleted, XS advanced platform services are still functional.

> ### Caution
>
> If a tenant database containing XS advanced custom application data is deleted, the data of affected applications is deleted, too, and the corresponding applications need to be re-deployed.

Location of Platform and Business Users

In this (default) setup, only SAP HANA users created in the system database can be used to access the XS advanced platform or XS advanced applications, even if an XS advanced application stores its data in a tenant database.

> **Note**
>
> To access the actual data stored in the respective database, the XS advanced platform and XS advanced applications use technical SAP HANA users that are created in the corresponding tenant database.

Implications for Backup and Recovery

In order to save all XS advanced related data, the system database needs to be backed up together with all other tenant databases containing XS advanced data, as described in *Backup & Recovery* in *Related Information* below.

Please note the following:

- Complete system backup
  In order to backup and recover XS advanced related data, the system database must be backed up as well in this case. Therefore, a complete system backup must be created.
- System recovery
  When recovering XS advanced related data, the complete system backup including the system database must be restored; it is not recommended to restore only particular individual tenant databases. The reason for this recommendation is that recovering the system database might have implications on the state of the entire SAP HANA system, which also affects tenant databases not related to XS advanced. For example, the system database contains information about SAP HANA topology, etc. Restoring the system database re-establishes the state of the SAP HANA system as it existed at the time of the backup.
- Data Consistency
  In this default setup, XS advanced related data is distributed across several databases, and since no cross-database snapshot is supported at the moment, each database must be backed up separately. To ensure the consistency of the data contained in the backup across databases, it is recommend to stop the XS advanced services during the backup operation.

Implications for Tenant Database Move

It is essential to understand that it is not possible to restore a backup of the System database into a tenant database. Moving XS-advanced-related data to a different SAP HANA system is only possible by creating a complete system copy of the source SAP HANA system.

Related Information

Maintaining Tenant Databases in XS Advanced
Installing XS Advanced in a Tenant Database [page 108]
XS Advanced Database Setup Options [page 106]
Backup and Recovery in XS Advanced (SAP HANA Admin Guide)

6.4.3 Install an SAP HANA System Including XS Advanced Runtime Using the Graphical Interface

An SAP HANA system with XS advanced runtime can be installed in a single-host or multi-host environment using the SAP HANA database lifecycle manager (HDBLCM).

Prerequisites

- The most recent version of the SAP HANA and SAP HANA XS advanced runtime installation packages are downloaded, and all packages are of an equivalent support or revision level.
- You are logged in as root user.
- Any user has read and execute permissions for the directory that contains the installation medium.
- Depending on the storage solution, set the export options `rw,no_root_squash` for the installation directory.
- The operating system administrator (`<sid>adm`) user and other operating system users may exist prior to installation. Make sure that you have the passwords of the existing users, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) will not modify the properties of any existing user or group. The `<sid>adm` user passwords must be identical on all hosts.
- The root user is able to execute graphical applications.
- The SAP HANA system must be installed with its server software on a shared file system:
  - Create an installation directory, e.g. `/hana/shared/`.
  - Depending on the storage solution, set the export options `rw,no_root_squash` for the installation directory.
  - Mount the installation directory on all hosts.
  - To install the Local Secure Store (LSS), the installation path must be created manually by the root user before the installation and mounted on all hosts, e.g. `/lss/shared/`.

    > **Note**
    >
    > The local secure store (LSS) must be installed and activated by default during an SAP HANA system installation or update.
- Depending on the desired storage solution, **shared storage devices** or **separate storage devices with failover reassignment**, different configurations apply:
  - In a shared file system, the data files and log files are configured so that they are present and mounted on all hosts, including the primary host.
  - In a system that uses separate storage devices, each host only has access to its own data files and log files.

  The suggested locations for the file systems are as follows:
  - `/hana/data/<SID>`
  - `/hana/log/<SID>`
- (Optional) Additional storage is configured.
- Root user name must be the same for all hosts in a multiple-host system. The password of the root user must be identical on all hosts.

- If the root user name is **not** root, it must be specified as a parameter during installation using the parameter `root_user`.

Context

The following procedure describes the installation of an SAP HANA system in interactive mode by entering parameters interactively. This procedure may also be performed in advanced interactive mode, with parameters entered as call options or from a configuration file. For more information about interaction modes and parameter entry methods, see *Using the SAP HANA Platform LCM Tools* in Related Information.

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. You can add additional tenant databases later using the SAP HANA cockpit. SAP HANA options hosts or host roles can be automatically provisioned to a system that is installed with a single tenant. If the SAP HANA system contains multiple tenant databases, the SAP HANA options host or host role must be manually provisioned to the tenant.

If you want to add the XS advanced runtime to an existing SAP HANA system, see *Installing or Updating SAP HANA Components* in Related Information.

> **Note**
>
> Not all parameters are requested interactively. Some parameters have default values that do not require confirmation in interactive mode. Those parameters must be specified as call options or from a configuration file. For more information about changeable default values, see Related Information.

Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
IBM Power Systems	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`

2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

   ```
   ./hdblcmgui
   ```

   The SAP HANA database lifecycle manager graphical user interface appears.

   > **Note**
   >
   > To activate the local secure (LSS) store during installation, run `hdblcmgui` with the parameter `secure_store=localsecurestore`.

3. On the Select Software Component Locations page, if the SAP HANA XS advanced runtime component appears on the list, click *Next*; otherwise:
   a. Click *Add Component Location*.
   b. Type the path to the missing installation package and click *OK*.
4. Select *Install New System*, then select *Next*.
5. Select SAP HANA XS advanced runtime, then select *Next*.
6. Select *Single-Host System* or *Multiple-Host System* as the *System Type*, and select *Add Host* to specify host parameters for each additional host.

Field Name	Description
Installation Path	Specifies the path to the SAP mount directory, which can be used as a shared directory between multiple hosts.
Non-standard Shared File System	Specifies a non-standard shared file system, which can be accessed by all hosts during installation. This parameter is typically used when the SID is included in the mountpoint.
Host Name	Specifies the host name of the machine. **Note** For XS advanced runtime installations, the fully-qualified host name must be specified.
Role	Specifies the purpose of the SAP HANA host. Although multiple host roles may be assigned, check the corresponding documentation for the SAP HANA option for what configurations are supported in production environments. • Database Worker (worker) - A worker host (default) is used for database processing. • Database Standby (standby) - A standby host is idle and available for failover in a high-availability environment. • Database Elastic Compute Server (compute) - Database elastic compute server • Dynamic Tiering Worker (extended_storage_worker) - Worker host for SAP HANA dynamic tiering • Dynamic Tiering Standby (extended_storage_standby) - Standby host for SAP HANA dynamic tiering • Accelerator for SAP ASE Worker (ets_worker) - Worker host for SAP HANA accelerator for SAP ASE • Accelerator for SAP ASE Standby (ets_standby) - Standby host for SAP HANA accelerator for SAP ASE • Streaming Analytics (streaming) - Host for SAP HANA streaming analytics • XS advanced runtime worker (xs_worker) - Host for SAP HANA XS advanced runtime • XS advanced runtime standby (xs_standby) - Standby host for SAP HANA XS advanced runtime
High-Availability Group	Specifies the host group ID for failover scenarios. If undefined, the host group is named "default".

Field Name	Description
Worker Group	Specifies the worker group of the host. If undefined, the worker group is named "default". If you are using extension node for Business Warehouse, you must name the worker group "worker_dt".
Storage Partition	Specifies the storage partition number, which is a logical role number assigned to non-shared storage devices in a storage connector API. Standby hosts do not have a storage partition.

XS advanced runtime host roles can be assigned automatically during installation. Assign worker and standby roles to all hosts and choose *Automatically assign XS Advanced Runtime roles to the hosts with database roles* when prompted. The installer will then assign the role xs_worker to every worker host and xs_standby to every standby host. To create a multiple-host system with dedicated xs_worker and xs_standby hosts, assign host roles to each host individually during installation. Do not choose the option to assign XS Advanced host roles automatically.

7. Specify the SAP HANA system properties.

   For a list of all system properties, see *System Properties* in Related Information.
8. Select the XS Advanced components you would like to install, then select `Next`.

Component	Description
GUI for HALM for XSA	The XS Advanced Application Lifecycle Management graphical user interface allows you to install, update, and uninstall SAP HANA products and software components in the SAP HANA XS Advanced model.
SAP File Processing	SAP File Processing provides structured information from unstructured files. The rich set of HTTP APIs enables application programmers to integrate File Processing features in client applications.
SAP Enterprise Architecture Designer for SAP HANA	SAP Enterprise Architecture Designer for SAP HANA lets you capture, analyze, and present your organization's landscapes, strategies, requirements, processes, data, and other artifacts in a shared environment. Using industry-standard notations and techniques, organizations can leverage rich metadata and use models and diagrams to drive understanding and promote shared outcomes in creating innovative systems, information sets, and processes to support goals and capabilities.
SAP HANA Runtime Tools	SAP HANA provides a selection of tools to help in the various phases of the design-time development and run-time administration of Multi-Target Applications (MTA) on XS advanced.
SAPUI5 Frontend Server	The SAPUI5 Frontend Server provides access to the SAPUI5 Demo Kit as part of the XS Advanced environment. The Demo Kit also provides you with technical documentation and samples.
SAP Web IDE Web Client	SAP Web IDE for SAP HANA is a comprehensive browser-based IDE for the development of complex applications comprised of Web-based or mobile UIs, business logic, and extensive SAP HANA data models.
XS Advanced Portal Services	With XS Advanced Portal Services you can develop and run portal services for custom apps running on XS Advanced.
XS Monitoring	SAP HANA XS Advanced Model includes a Web-based tool that enables you to maintain important parts of the application-development environment, for example, security and authentication methods.

Component	Description
XS Services	In SAP HANA XS Advanced, application developers can make use of a catalog of services managed by a service broker, for example, for job schedules or user accounts and OAuth clients.
XSAC XMLA Interface For HANA	XMLA enables you to send MDX queries via http against SAP HANA calculation views following the XML/A-protocol. With XMLA, you write multi-dimensional-expressions (MDX) queries wrapped in an XMLA document. An XML for Analysis (XMLA) application running in SAP HANA application services (SAP HANA XS) is used to provide the consumption model for client applications exchanging MDX queries (wrapped in XMLA documents) with the SAP HANA database.
XS Advanced Cockpit	The XS Advanced Cockpit provides the central web user interface for XS advanced runtime. As an application developer or an administrator, cockpit enables you to manage your applications, services, and resources. Cockpit is an auto-content app delivered with XS advanced runtime. The Cockpit provides an overview of the Organizations and Spaces the logged in user has access to. Within the Spaces, the application developer can manage his/her applications and related resources like service instances. It also provides access to the Service Marketplace and the details of the Services enabled for the corresponding Organization.

9. After specifying all system properties, review the summary, and select *Install*.

Results

A SAP HANA system with XS advanced runtime is installed. A log file is available.

Next Steps

After installing the SAP HANA system, you may want to perform configuration tasks. For more information, see *Managing the SAP HANA System After Installation* or the platform lifecycle management section of the *SAP HANA Administration Guide*.

For certain versions of SAP Host Agent it is necessary to configure the TCP/IP port ranges (50000-50999, 51000-51500) manually after installation. This is explained in detail in SAP Note 401162.

In the XSA Public area the XSA administrator is responsible for deploying the domain-specific certificates. These can be either self-signed or issued by the global certificate authority (CA). The certificates can be deployed in the xs client using the `set-certificate` command. This is explained in detail in SAP Note 2243019. However, by default, the system generates self-signed certificates that the administrator can manually and securely distribute among the clients. For more information, see*Network and Communication Security with SAP HANA XS Advanced* in the *SAP HANA Security Guide*.

> **Note**
>
> In production XSA installations, replace the self-signed certificate with one trusted in your organization.

Related Information

System Properties [page 142]
Using the SAP HANA Platform LCM Tools [page 41]
Changeable Default Values for Installation [page 73]
Managing the SAP HANA System After Installation [page 209]
Maintaining the SAP HANA XS Advanced Model Run Time
SAP Note 2243156
Network and Communication Security with SAP HANA XS Advanced
Installing or Updating SAP HANA Components [page 181]

6.4.4 Install an SAP HANA System Including XS Advanced Runtime Using the Command-Line Interface

An SAP HANA system with XS advanced runtime can be installed in a single-host or multi-host environment using the SAP HANA database lifecycle manager (HDBLCM).

Prerequisites

- The most recent version of the SAP HANA and SAP HANA XS advanced runtime installation packages are downloaded, and all packages are of an equivalent support or revision level.
- You are logged in as root user.
- Any user has read and execute permissions for the directory that contains the installation medium.
- Depending on the storage solution, set the export options `rw,no_root_squash` for the installation directory.
- The operating system administrator (`<sid>adm`) user and other operating system users may exist prior to installation. Make sure that you have the passwords of the existing users, and that the user attributes and group assignments are correct. The SAP HANA database lifecycle manager (HDBLCM) will not modify the properties of any existing user or group. The `<sid>adm` user passwords must be identical on all hosts.
- The SAP HANA system must be installed with its server software on a shared file system:
  - Create an installation directory, e.g. `/hana/shared/`.
  - Depending on the storage solution, set the export options `rw,no_root_squash` for the installation directory.
  - Mount the installation directory on all hosts.
  - To install the Local Secure Store (LSS), the installation path must be created manually by the root user before the installation and mounted on all hosts, e.g. `/lss/shared/`.

    > **Note**
    >
    > The local secure store (LSS) must be installed and activated by default during an SAP HANA system installation or update.

- Depending on the desired storage solution, **shared storage devices** or **separate storage devices with failover reassignment**, different configurations apply:
  - In a shared file system, the data files and log files are configured so that they are present and mounted on all hosts, including the primary host.
  - In a system that uses separate storage devices, each host only has access to its own data files and log files.

  The suggested locations for the file systems are as follows:
  - `/hana/data/<SID>`
  - `/hana/log/<SID>`
- (Optional) Additional storage is configured.
- Root user name must be the same for all hosts in a multiple-host system. The password of the root user must be identical on all hosts.
- If the root user name is **not** root, it must be specified as a parameter during installation using the parameter `root_user`.

Context

The following procedure describes the installation of an SAP HANA system in interactive mode and entering parameters interactively. This procedure may also be performed in advanced interactive mode or batch mode, with parameters entered as call options or from a configuration file. For more information about interaction modes and parameter entry methods, see *Using the SAP HANA Platform LCM Tools* in Related Information.

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. You can add additional tenant databases later using the SAP HANA cockpit. SAP HANA options hosts or host roles can be automatically provisioned to a system that is installed with a single tenant. If the SAP HANA system contains multiple tenant databases, the SAP HANA options host or host role must be manually provisioned to the tenant.

If you want to add the XS advanced runtime to an existing SAP HANA system, see *Installing or Updating SAP HANA Components* in Related Information.

> ⓘ Note
>
> Not all parameters are requested interactively. Some parameters have default values that do not require confirmation in interactive mode. Those parameters must be specified as call options or from a configuration file. For more information about changeable default values, see Related Information.

Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
IBM Power Systems	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblcm
```

3. Select the index for *Install New System*, then select `Enter`.
4. Select *server*, *xs* and any other components you would like to install as a comma-separated list, then select `Enter`.
5. Specify the installation path, and the local host name:

SAP HANA System Properties

Field Name	Description
Installation Path	Specifies the path to the SAP mount directory, which can be used as a shared directory between multiple hosts.
Local Host Name	Specifies the host name of the machine.

6. If you want to install a single-host system, select *n* to the question *Do you want to add hosts to the system?*. If you want to install a multi-host system, select *y* and enter the following details for each additional host:

Field Name	Description
Host Name	Specifies the host name of the machine. **Note** For XS advanced runtime installations, the fully-qualified host name must be specified.

Field Name	Description
Role	Specifies the purpose of the SAP HANA host. Although multiple host roles may be assigned, check the corresponding documentation for the SAP HANA option for what configurations are supported in production environments. • Database Worker (worker) - A worker host (default) is used for database processing. • Database Standby (standby) - A standby host is idle and available for failover in a high-availability environment. • Database Elastic Compute Server (compute) - Database elastic compute server • Dynamic Tiering Worker (extended_storage_worker) - Worker host for SAP HANA dynamic tiering • Dynamic Tiering Standby (extended_storage_standby) - Standby host for SAP HANA dynamic tiering • Accelerator for SAP ASE Worker (ets_worker) - Worker host for SAP HANA accelerator for SAP ASE • Accelerator for SAP ASE Standby (ets_standby) - Standby host for SAP HANA accelerator for SAP ASE • Streaming Analytics (streaming) - Host for SAP HANA streaming analytics • XS advanced runtime worker (xs_worker) - Host for SAP HANA XS advanced runtime • XS advanced runtime standby (xs_standby) - Standby host for SAP HANA XS advanced runtime
Host Failover Group	Specifies the host group ID for failover scenarios. If undefined, the host group is named "default".
Worker Group	Specifies the worker group of the host. If undefined, the worker group is named "default". If you are using extension node for Business Warehouse, you must name the worker group "worker_dt".
Storage Partition	Specifies the storage partition number, which is a logical role number assigned to non-shared storage devices in a storage connector API. Standby hosts do not have a storage partition.

XS advanced runtime host roles can be assigned automatically during installation. Assign worker and standby roles to all hosts and choose *Automatically assign XS Advanced Runtime roles to the hosts with database roles* when prompted. The installer will then assign the role xs_worker to every worker host and xs_standby to every standby host. To create a multiple-host system with dedicated xs_worker and xs_standby hosts, assign host roles to each host individually during installation. Do not choose the option to assign XS Advanced host roles automatically.

7. Specify the SAP HANA system properties.

   For a list of all system properties, including those used by XS advanced, see *System Properties* in Related Information.

8. Select the XS Advanced components you would like to install as a comma-separated list, then press `Enter`.

Component	Description
xsac_alm_pi_ui	The XS Advanced Application Lifecycle Management graphical user interface allows you to install, update, and uninstall SAP HANA products and software components in the SAP HANA XS Advanced model.
xsac_file_proc	SAP File Processing provides structured information from unstructured files. The rich set of HTTP APIs enables application programmers to integrate File Processing features in client applications.
xsac_hana_ea_d	SAP Enterprise Architecture Designer for SAP HANA lets you capture, analyze, and present your organization's landscapes, strategies, requirements, processes, data, and other artifacts in a shared environment. Using industry-standard notations and techniques, organizations can leverage rich metadata and use models and diagrams to drive understanding and promote shared outcomes in creating innovative systems, information sets, and processes to support goals and capabilities.
xsac_hrtt	SAP HANA provides a selection of tools to help in the various phases of the design-time development and run-time administration of Multi-Target Applications (MTA) on XS advanced.
xsac_mess_srv	In SAP HANA XS, advanced model, the messaging service is an application service that enables consumer applications to receive messages sent by producer applications depending on the topics or queues that consumer applications are subscribed to. The messaging service represents a message broker for remote communication between systems using the Apache ActiveMQ library with JMS (Java Message Service) specification. It uses ActiveMQ as message-oriented middleware, which enables the service to provide scalability, high availability, reliability, and security for enterprise messaging.
xsac_monitoring	SAP HANA XS Advanced Model includes a Web-based tool that enables you to maintain important parts of the application-development environment, for example, security and authentication methods.
xsac_portal_serv	With XS Advanced Portal Services you can develop and run portal services for custom apps running on XS Advanced.
xsac_sap_web_ide	SAP Web IDE for SAP HANA is a comprehensive browser-based IDE for the development of complex applications comprised of Web-based or mobile UIs, business logic, and extensive SAP HANA data models.
xsac_services	In SAP HANA XS Advanced, application developers can make use of a catalog of services managed by a service broker, for example, for job schedules or user accounts and OAuth clients.
xsac_ui5_fesv4	The SAPUI5 Frontend Server provides access to the SAPUI5 Demo Kit as part of the XS Advanced environment. The Demo Kit also provides you with technical documentation and samples.
xsac_ui5_sb	The UI5 service provides resources to SAP UI5 applications which are needed to run their graphical user interfaces. In a typical XS advanced installation, one service is installed at any point in time and the service corresponds to an UI5 release. UI5 applications no longer need to refer to a particular UI5 service that they depend on; they can refer to the UI5 Service Broker, which serves the bootstrap URL of the service they require.

Component	Description
xsac_xmla_int	XMLA enables you to send MDX queries via http against SAP HANA calculation views following the XML/A-protocol. With XMLA, you write multi-dimensional-expressions (MDX) queries wrapped in an XMLA document. An XML for Analysis (XMLA) application running in SAP HANA application services (SAP HANA XS) is used to provide the consumption model for client applications exchanging MDX queries (wrapped in XMLA documents) with the SAP HANA database.
xsac_xsa_cockpit	The XS Advanced Cockpit provides the central web user interface for XS advanced runtime. As an application developer or an administrator, cockpit enables you to manage your applications, services, and resources. Cockpit is an auto-content app delivered with XS advanced runtime. The Cockpit provides an overview of the Organizations and Spaces the logged in user has access to. Within the Spaces, the application developer can manage his/her applications and related resources like service instances. It also provides access to the Service Marketplace and the details of the Services enabled for the corresponding Organization.

9. After specifying all system properties, review the summary, and select *y*.

Results

A SAP HANA system with XS advanced runtime is installed. A log file is available.

Next Steps

After installing the SAP HANA system, you may want to perform configuration tasks. For more information, see *Managing the SAP HANA System After Installation* or the platform lifecycle management section of the *SAP HANA Administration Guide*.

For certain versions of SAP Host Agent it is necessary to configure the TCP/IP port ranges (50000-50999, 51000-51500) manually after installation. This is explained in detail in SAP Note 401162.

In the XSA Public area the XSA administrator is responsible for deploying the domain-specific certificates. These can be either self-signed or issued by the global certificate authority (CA). The certificates can be deployed in the xs client using the `set-certificate` command. This is explained in detail in SAP Note 2243019. However, by default, the system generates self-signed certificates that the administrator can manually and securely distribute among the clients. For more information, see*Network and Communication Security with SAP HANA XS Advanced* in the *SAP HANA Security Guide*.

> **Note**
>
> In production XSA installations, replace the self-signed certificate with one trusted in your organization.

Related Information

System Properties [page 142]
Using the SAP HANA Platform LCM Tools [page 41]
Changeable Default Values for Installation [page 73]
Managing the SAP HANA System After Installation [page 209]
Maintaining the SAP HANA XS Advanced Model Run Time
SAP Note 2243156
Network and Communication Security with SAP HANA XS Advanced
Installing or Updating SAP HANA Components [page 181]

6.4.5 Setting Up the XS Advanced Runtime Behind a Reverse Proxy

The XS advanced runtime can work behind a reverse proxy, for example, for load balancing.

You can set up the XS advanced runtime behind a reverse proxy in such a way that XS advanced applications and system services can be reached both externally and internally between applications. Before you do so, it is important that you understand how routing works in an XS advanced landscape.

Default Domain and Routing Modes

XS advanced applications are exposed to the end user by URLs that are based on the XS advanced default domain and routing mode. This is also true for system components like the XS controller and UAA (User Account and Authentication). You can specify both the default domain and the routing mode during installation.

> → Tip
>
> To see the URL of each individual application use the `xs apps` command. For more information about the XS command-line interface, see the *SAP HANA Developer Guide*.

> ⓘ Note
>
> It's possible to create further domains for XS advanced applications in addition to the default domain. The information provided in this documentation applies equally to such additionally created domains.

Port-based Routing

If you select port-based routing mode during installation, applications are exposed using the XS advanced default domain. Different applications are distinguished by different TCP ports.

> ✤ Example
>
> If your XS advanced default domain is `hana.example.com`, your application URLs would have the format: `hana.example.com:50500`, `hana.example.com:50501`, and so on.

Hostname-based Routing

If you select hostname-based routing mode during installation, applications are exposed using a sub-domain of the XS advanced default domain (for example, the XS advanced default domain prefixed with a hostname unique for each application). In this case, the TCP port is the same for each application.

> **Example**
>
> If your XS advanced default domain is `hana.example.com`, your app URLs would have the format: `myapp1.hana.example.com`, `myapp2.hana.example.com`, and so on.

Changing the Default Domain

After installation, you can change the default domain with the property `default_domain` in the `xscontroller.ini` configuration file. After changing this property, you must restart the XS advanced runtime, for example, using the `XSA restart` command.

> **Note**
>
> Changing the default domain requires some downtime of your applications because all applications are re-staged during XS Controller startup. Changing the routing mode after installation is not supported.

Default Domain and External Domain

From XS Advanced 1.3.0, it is recommended to set up a reverse proxy with an external domain that is different from the XS Advanced default domain. With this setup, you can ensure that the system is still functional internally, even if the reverse proxy is not available.

As a second option, you can use a reverse proxy without an external domain. With this setup, the system is only functional when the reverse proxy is up and running.

Related Information

Using A Reverse Proxy with an External Domain [page 129]
Using A Reverse Proxy without an External Domain [page 132]
xs_routing_mode [page 292]
xs_domain_name [page 291]
XS CLI: Application Management
The XSA Command Reference
SAP Note 2245631

6.4.5.1 Using A Reverse Proxy with an External Domain

Set up XS advanced with a reverse proxy and external domain (recommended).

As of XS Advanced 1.3.0, it is possible (and recommended) to use an external domain, for example to route requests via a reverse proxy or to setup a failover router for SAP HANA system replication. For more information about creating an external domain, see *Maintaining Domains And Certificates* in *Related Information* below.

> ⚠ Restriction
>
> Using a reverse proxy with an external domain with the SAP Web IDE for SAP HANA is only supported as of XSA version 1.4.0 and SAP Web IDE for SAP HANA version 4.8.2. If you have deployed a previous version of the SAP Web IDE for SAP HANA, see the instructions in *Setup a Reverse Proxy Without an External Domain* in *Related Information* below, instead.

Install XS Advanced

In this setup, XS Advanced is installed without a reverse proxy by configuring the `xs_worker` SAP HANA host name as the default domain, for example, `internal.com`, as illustrated in the following example.

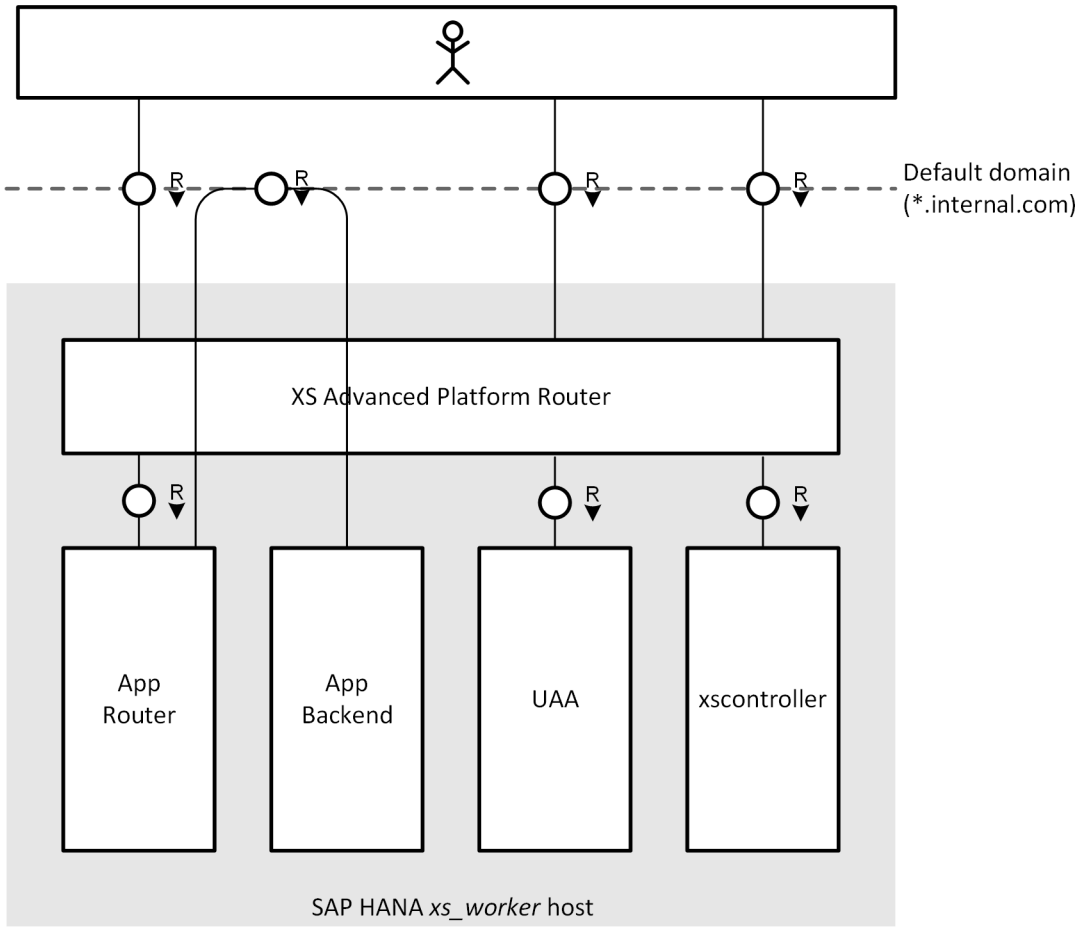

Setting up the Reverse Proxy

After installing XS Advanced, you can create an external domain to route incoming requests via the reverse proxy, as illustrated in the following example:

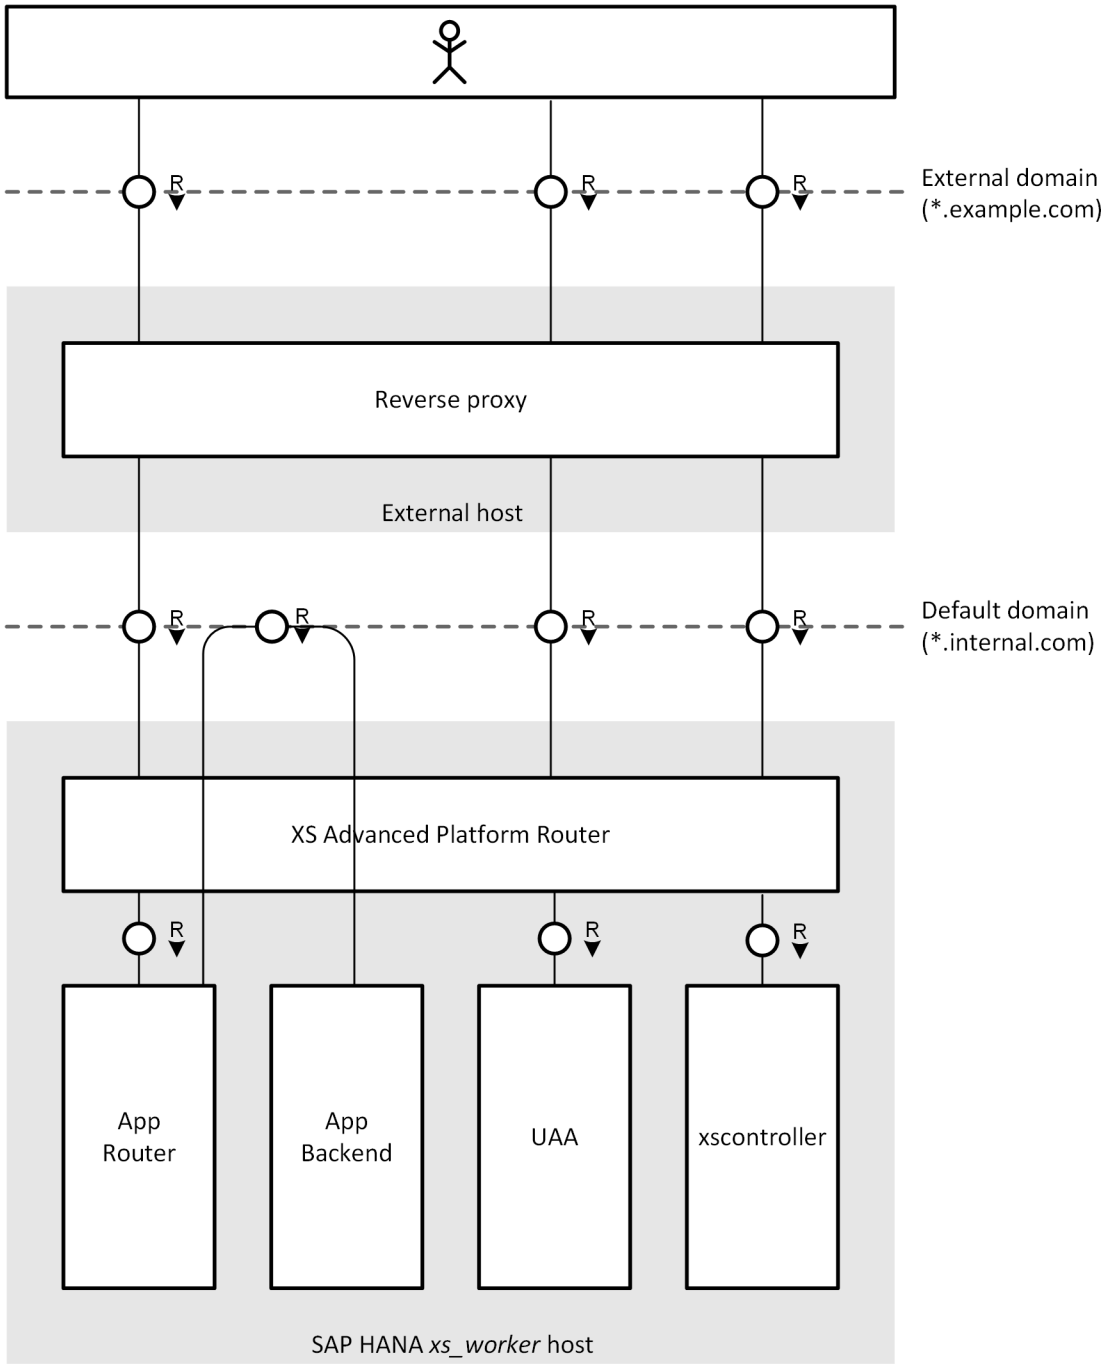

Setup Overview and Process

To set up an external domain for use with XS advanced, you need to perform the following high-level steps:

1. Configure an external domain that points to the reverse proxy.
   As user `XSA_ADMIN`, open a command console and run the `xs create-domain` command with the `--external` option to configure an external domain that points to the reverse proxy, for example, `example.com`:

   ```
   xs create-domain example.com --external
   ```

2. Configure certificate trust.
   - If the reverse proxy does not terminate SSL or the external domain is a virtual hostname.

     > → Tip
     >
     > You can skip this step unless you want to use a custom server certificate for the external domain. In this case, set a certificate for the external domain as described in *Maintaining Domains And Certificates* in *Related Information* below.

   - If the reverse proxy terminates SSL, configure the reverse proxy to trust the XS Advanced default domain root certificate. You can extract this certificate from the following file-system location on the SAP HANA system:

     ```
     $XSPATH/controller_data/controller/ssl-pub/router/default.root.crt.pem
     ```

3. Restart XS Advanced
   Run the following command as the `<sid>adm` user:

   ```
   XSA restart
   ```

4. Configure the reverse proxy to route requests on the external domain to the XS Advanced default domain. In the above example, the router forwards requests on `*.example.com` to `internal.com`.
   The configuration required depends on the routing mode configured, for example, **port-based** or **hostname-based** routing, as follows:
   - For port-based routing
     Forward all traffic on the ports `30<instance nr>30`, `30<instance nr>32` and all ports within the router port range (by default ports 51000-51500) on the external domain (for example `example.com:[30030,30032,51000-51500]`) to the same ports on the internal domain (for example `internal:[30030,30032,51000-51500]`)
   - For hostname-based routing
     Forward all traffic on the configured external port (by default `30<instance nr>33`) on all sub-domains (for example, `*.example.com:30033`) to the router port on the internal domain (for example `internal:30033`).

     > ⓘ Note
     >
     > Make sure that the reverse proxy does not modify the HTTP "host" header in this case.

Example: Using the SAP Web Dispatcher as a Reverse Proxy

It is possible to configure an instance of the SAP Web Dispatcher to work as a reverse proxy. For example, the following configuration to the SAP Web Dispatcher profile will ensure that all requests on port 30033 are forwarded from the reverse proxy to XS advanced (running on the host xsa-internal:30033 with hostname-based routing) without terminating SSL:

Sample Code

SAP Web Dispatcher Configuration

```
...
icm/server_port_0=PROT=TCP, PORT=30033, TIMEOUT=60, PROCTIMEOUT=600
wdisp/system_0=NAME=XSA, SID=XSA, EXTSRV=tcp://internal.com:30033,
SRCTCPSRV=*:30033
...
```

Related Information

Setting Up the XS Advanced Runtime Behind a Reverse Proxy [page 127]
Using A Reverse Proxy without an External Domain [page 132]
Maintaining Domains And Certificates (SAP HANA Administration Guide for SAP HANA Platform

6.4.5.2 Using A Reverse Proxy without an External Domain

Configure XS advanced to use a reverse proxy without an external domain.

The information in this topic explains how the request flows work in a basic (default) XS advanced landscape with no reverse proxy configured as well as in a landscape that does make use of a reverse-proxy setup.

Basic XS Advanced Landscape (no Reverse Proxy)

In a basic landscape with no reverse proxy in place, requests from the end user's browser are sent directly to the Platform Router, or the XS advanced Web Dispatcher, which dispatches the requests to the appropriate application.

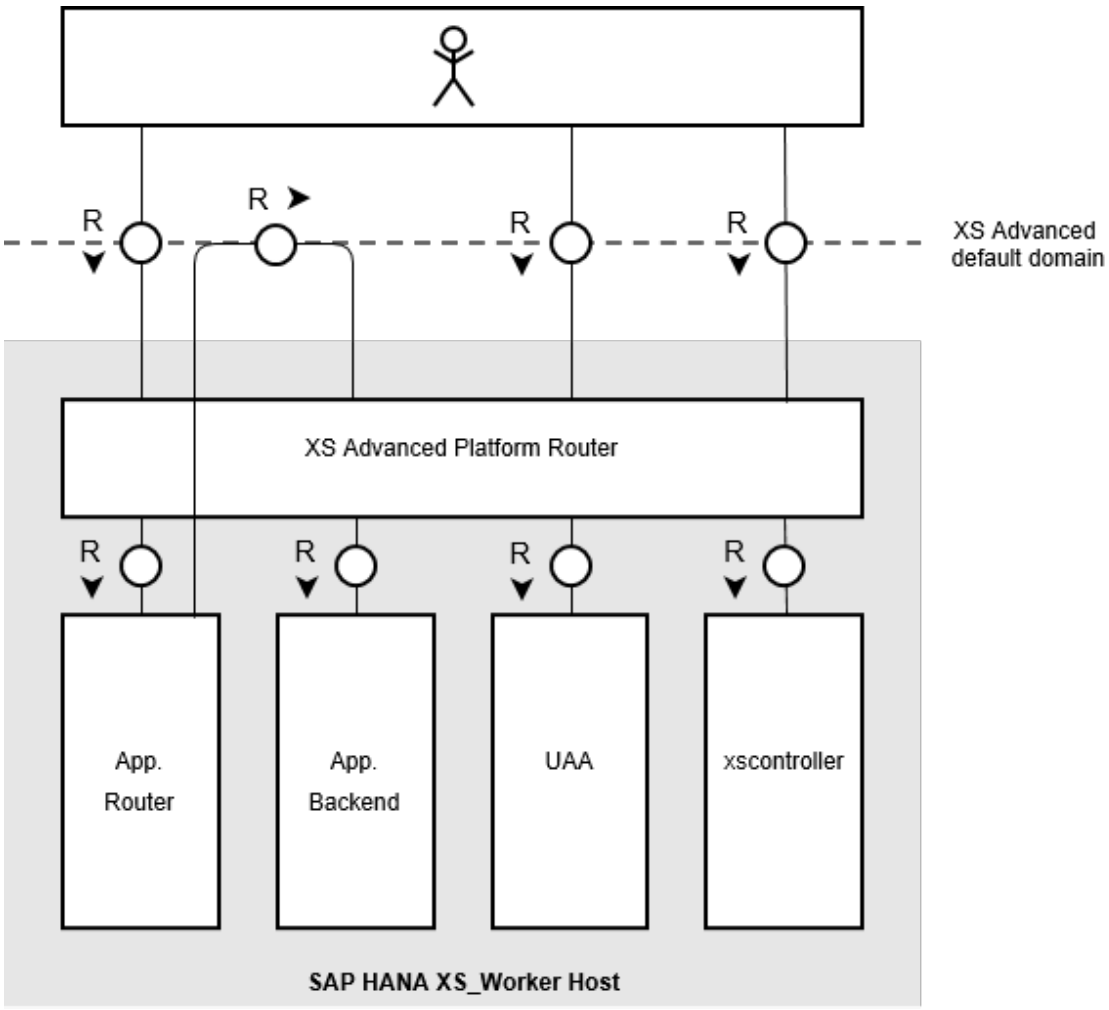

Basic Request Flow

This means that the XS advanced default domain must be mapped in your DNS to the SAP HANA host on which the XS advanced Web Dispatcher is running. This is the same host that runs as the `xscontroller` service.

> ⓘ Note
>
> If you are using hostname-based routing, the XS advanced default domain and all its subdomains must be mapped by means of a wildcard DNS entry.

XS Advanced Landscape with Reverse Proxy

If a reverse proxy (for example, a load balancer) is in place between the end user's browser and the XS advanced Web Dispatcher, the reverse proxy is responsible for routing all requests on all domains registered with XS advanced to the SAP HANA host on which the XS advanced Web Dispatcher is running. Internal requests from an application to its back end are also dispatched via the reverse proxy. However, it is possible to configure internal requests to be routed directly to the XS advanced Web Dispatcher. For more information, see *Routing Internal Requests to Backend Applications*.

> **Note**
>
> In a high-availability setup (SAP HANA System Replication or Host Auto-Failover) or in a multi-host system where several hosts have been assigned the role of `xs_worker`, the reverse proxy must forward requests to the currently active host or system. For more information, see the section on high availability in the *SAP HANA Administration Guide*.

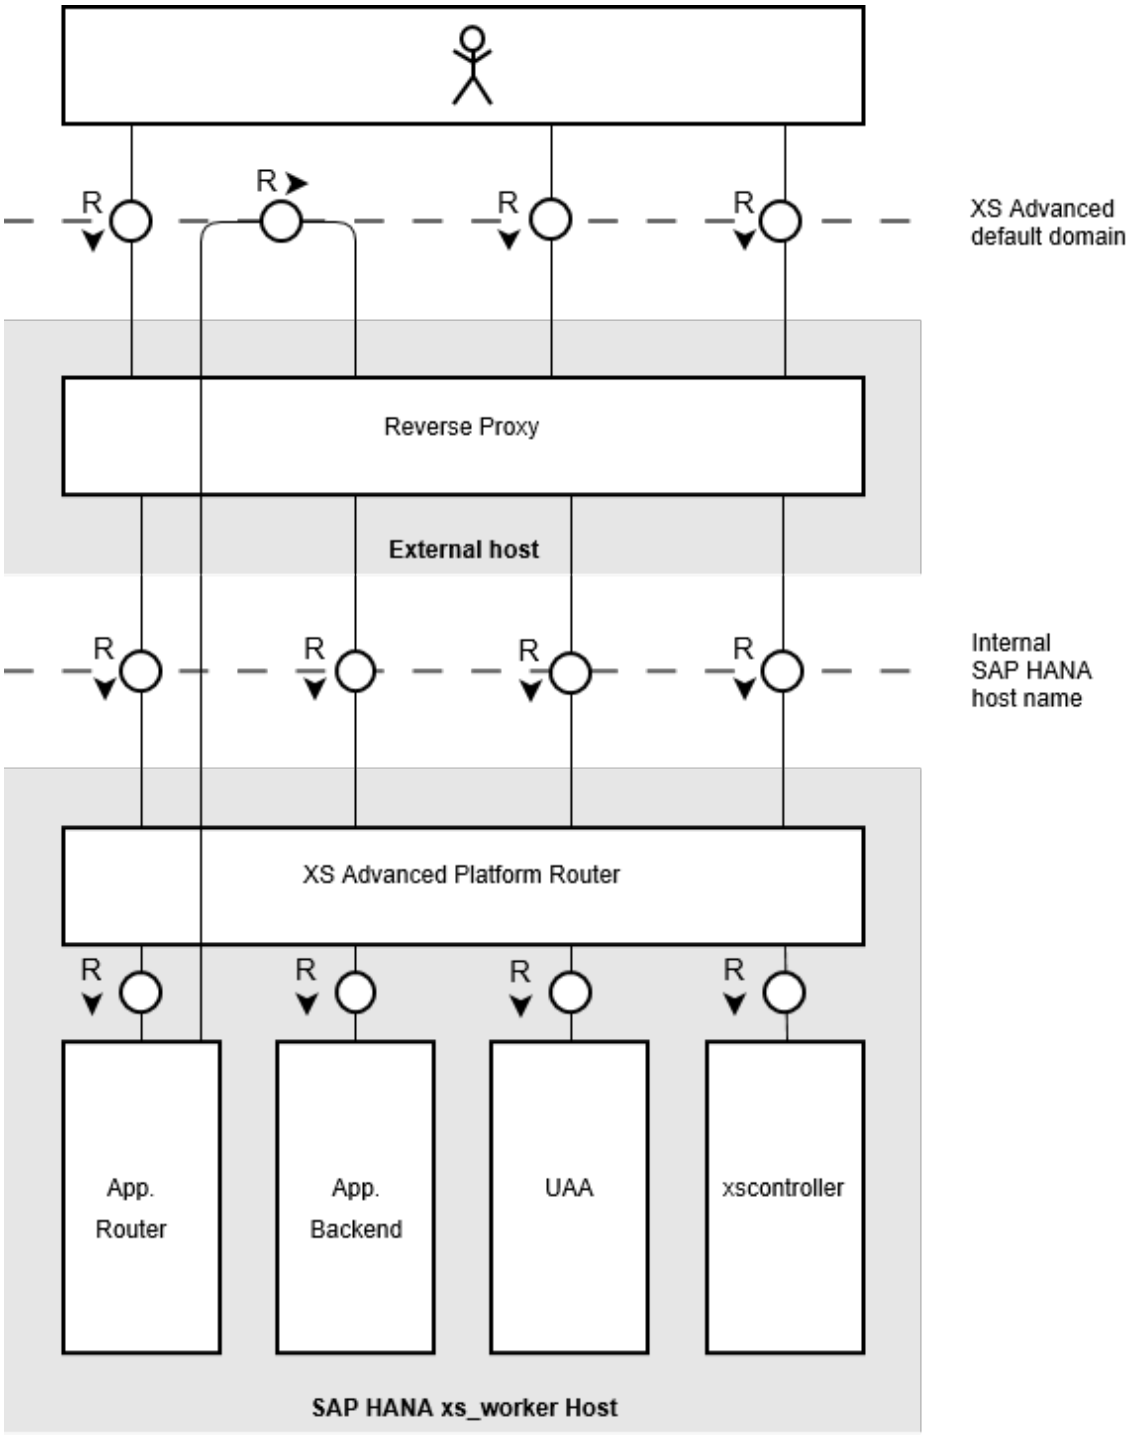

Request Flow with Reverse Proxy

As a result, the DNS entry for the XS advanced default domain must point to the host running the reverse proxy.

If you are using port-based routing, the reverse proxy forwards requests on each individual application port running on the default domain to the XS advanced Web Dispatcher. If you are using hostname-based routing, the reverse proxy forwards requests on the XS advanced default domain and each subdomain to the XS advanced Web Dispatcher.

Related Information

Using A Reverse Proxy with an External Domain [page 129]
Routing Internal Requests to Backend Applications [page 140]
Host Auto-Failover Setup with XS Advanced Run Time
SAP HANA System Replication Setup for XS Advanced Runtime

6.4.5.2.1 Installing XS advanced behind a Reverse Proxy

Configure XS advanced to run behind a reverse-proxy setup.

To operate XS advanced behind a reverse proxy, it is recommended to install XS advanced without a reverse proxy, first, and then **move** XS advanced behind the reverse proxy, as described in *Moving XS Advanced behind a Reverse Proxy* in *Related Information* below. However, it is also possible to install XS advanced behind a reverse proxy directly.

> ⓘ Note
>
> During installation, requests are made to the XS advanced default domain. For this reason, when the installation starts, the reverse proxy must already be configured to forward requests on that domain to the SAP HANA host that is assigned the `xs_worker` role.

Using a Reverse Proxy with SSL Termination

If you terminate SSL at the reverse proxy, certificate trust must already be set up between the reverse proxy and the XS advanced Web Dispatcher at this point in time. To establish this, the following solutions exist:

- Specify the Appropriate Certificates During XS Advanced Installation
- Temporarily Map the Default Domain to the XS advanced Web Dispatcher

Specify the Appropriate Certificates During XS Advanced Installation

This method requires that you obtain a certificate for the XS advanced default domain which is signed by a certificate authority (CA) that the reverse proxy trusts before you run the installation. Then you need to specify

the paths to the certificate files as non-interactive parameters with `HDBLCM`. The following list describes the relevant parameters:

- `xs_cert_pem`
  The path to the default domain certificate file (in PEM format) that the reverse proxy trusts
- `xs_cert_key`
  The path to the default domain certificate key file (in PKCS8 format) that the reverse proxy trusts
- `xs_trust_pem`
  The path to an additional trust certificate (to trust the reverse proxy)

Temporarily Map the Default Domain to the XS advanced Web Dispatcher

As an alternative to specifying certificates, you can re-map the default domain to the SAP HANA host that runs the XS advanced Web Dispatcher, for example, either by modifying the corresponding DNS entry, described here, or by using the methods described in *Routing Internal Requests to Back-end Applications* in *Related Information* below.

To re-map the default domain by modifying the corresponding DNS entry, proceed as follows:

1. Before installation, map the XS advanced default domain (and its sub-domains if you are using hostname-based routing) to the SAP HANA host running the XS advanced Web Dispatcher.
2. Run the XS advanced installation.
3. Disable XS advanced.
4. Establish the trust relationship between the reverse proxy and the XS advanced Web Dispatcher as described in the section *Configure Certificate Trust* of *Moving XS Advanced behind a Reverse Proxy* in *Related Information* below.

   > **Note**
   >
   > Make sure, that the certificate you install at XS advanced contains both the internal and external domains as the `Common Name` (CN) and `Subject Alternative Name` (SAN), respectively.

5. Enable XS advanced.
6. Map the DNS entry for the default domain back to the reverse proxy.

If you decide to map the default domain using the `/etc/hosts` file, it is sufficient to map the following domains temporarily during installation. After successful installation and certificate setup, you can remove the temporary mapping:

- `<xsa-default-domain>`
- auditlog-server.`<xsa-default-domain>`
- auditlog-broker.`<xsa-default-domain>`
- deploy-service.`<xsa-default-domain>`
- `<org>-<space>`-product-installer.`<xsa-default-domain>`
- subdomaintest.`<xsa-default-domain>`
- uaa-server.`<xsa-default-domain>`
- api.`<xsa-default-domain>`

Related Information

Move XS advanced Behind a Reverse Proxy [page 137]
Routing Internal Requests to Backend Applications [page 140]

6.4.5.2.2 Move XS advanced Behind a Reverse Proxy

Change the XS advanced configuration to enable it to run behind a reverse-proxy setup.

Prerequisites

- You have already installed XS advanced without a reverse proxy and you have chosen an internal default domain during installation (for example `xsa-internal`)
- It is recommended to use `hostname` routing mode, which is described in *Routing Mode Configuration* in *Related Information* below. Hostname routing mode simplifies the reverse proxy configuration because only the traffic on a single port needs to be forwarded.
- You need to have DNS entries for two different domains for this setup:
  - The internal domain (for example, `*.xsa-internal`):
    The domain of the SAP HANA host with `xs_worker` role.
  - The external domain (for example, `*.example.com`):
    The domain that is exposed to the public, i.e. the domain of the host running the reverse proxy
- When you use `hostname` routing mode, both domains must be so-called "wildcard domains".

Context

Although it is also possible to install XS advanced behind a reverse proxy directly, it is recommended first to install XS advanced without a reverse proxy, and then to move XS advanced behind the reverse proxy, as described here:

Procedure

1. Stop XS advanced.

   To connect the reverse proxy, all XS advanced services must be stopped. In the command shell, run the following command as `sidadm` user:

   ```
   XSA disable
   ```

2. Configure the reverse proxy to forward requests on the external domain to the internal domain.

The configuration required depends on the routing mode configured, for example, **port-based** or **hostname-based** routing, as follows:

- For port-based routing
  Forward all traffic on the ports `30<instance nr>30`, `30<instance nr>32` and all ports within the router port range (by default ports 51000-51500) on the external domain (for example `example.com:[30030,30032,51000-51500]`) to the same ports on the internal domain (for example `xsa-internal:[30030,30032,51000-51500]`)
- For hostname-based routing
  Forward all traffic on the router port (by default `30<instance nr>33`) on all sub-domains (for example, `*.example.com:30033`) to the same port on the internal domain (for example `xsa-internal:30033`).

  > **Note**
  >
  > Make sure that the reverse proxy does not modify the HTTP "host" header in this case.

3. Configure certificate trust.

   How you configure certificate trust with the reverse proxy depends on whether the reverse proxy terminates SSL or not, as follows:

   - Reverse Proxy **without** SSL termination:
     If the reverse proxy does **not** terminate SSL, you can skip this step.
   - Reverse Proxy **with** SSL termination:
     If the reverse proxy **does** terminate SSL, you must set up certificate trust between the reverse proxy and XS advanced, as follows:
     1. Create a signing request at a certificate authority (CA) for a certificate that includes the external domain (for example, `*.example.com`)
     2. Install the signed certificate for external domain (for example, `*.example.com`) at the reverse proxy
     3. Create a signing request at a CA for a certificate that includes the internal domain (for example, `*.xsa-internal`)
     4. Depending on the format of the certificate (`pem` or `pse` format), install the signed certificate for the internal domain as default certificate at XS advanced by running the following command in a command shell:

        ```
        XSA set-certificate --key <key> --cert <cert> --no-restart
        ```

        or

        ```
        XSA set-certificate --pse <pse file> --no-restart
        ```

     5. If the certificate for the external domain (for example, `*.example.com`) was **not** signed by a well-known Certificate Authority (CA), then you need to configure XS advanced to trust the CA's root certificate by running the following command in a command shell as the `sidadm` user:

        ```
        XSA trust-certificate --alias default --cert <CERTIFICATE_FILE>
        ```

     6. If the certificate for internal domain (for example, `*.xsa-internal`) was **not** signed by a well-known CA, configure the reverse proxy to trust the CA's root certificate.

4. Configure the default domain.

In the file `xscontroller.ini` configure the property `default_domain` to point to the external domain (for example, `default_domain=example.com`)

5. Start XS advanced.

   To restart all XS advanced services, run the following command in a command shell as the `sidadm` user:

   ```
   XSA enable
   ```

Example: Using the SAP Web Dispatcher as a Reverse Proxy

It is possible to configure an instance of the SAP Web Dispatcher to work as a reverse proxy. For example, the following configuration to the SAP Web Dispatcher profile will ensure that all requests on port `30033` are forwarded from the reverse proxy to XS advanced (running on the host `xsa-internal:30033` with hostname-based routing) without terminating SSL:

Sample Code

SAP Web Dispatcher Configuration

```
...
icm/server_port_0=PROT=TCP, PORT=30033, TIMEOUT=60, PROCTIMEOUT=600
wdisp/system_0=NAME=XSA, SID=XSA, EXTSRV=tcp://xsa-internal:30033,
SRCTCPSRV=*:30033
...
```

For more information about the SAP Web Dispatcher, see *Related Information* below.

Related Information

Installing XS advanced behind a Reverse Proxy [page 135]
Routing Internal Requests to Backend Applications [page 140]
Routing Mode Configuration (System Concepts for XS Advanced Runtime Installations)
SAP Note 2245631
SAP Web Dispatcher

6.4.5.2.3 Routing Internal Requests to Backend Applications

Route internal requests on the XS advanced default domain directly to the XS advanced Web Dispatcher to optimize request performance and enable app-to-app calls even when the XS advanced default domain cannot be reached from behind the reverse proxy.

To route internal requests on the XS advanced default domain directly to the XS advanced Web Dispatcher, you must map the XS advanced default domain to the SAP HANA host on which the XS advanced Web Dispatcher is running on every host with an XS advanced host role (`xs_*` , for example `xs_worker`).

> ### Note
>
> Routing internal requests as described here also allows you to avoid trust-related issues between the reverse proxy and the XS advanced Web Dispatcher during installation. For more information, see the section on establishing the trust relationship.

> ### Note
>
> The certificate you install on the XS advanced Web Dispatcher (see section on establishing the trust relationship) must have both the SAP HANA host running the XS advanced Web Dispatcher and the default domain in its “subjectAltName” section.

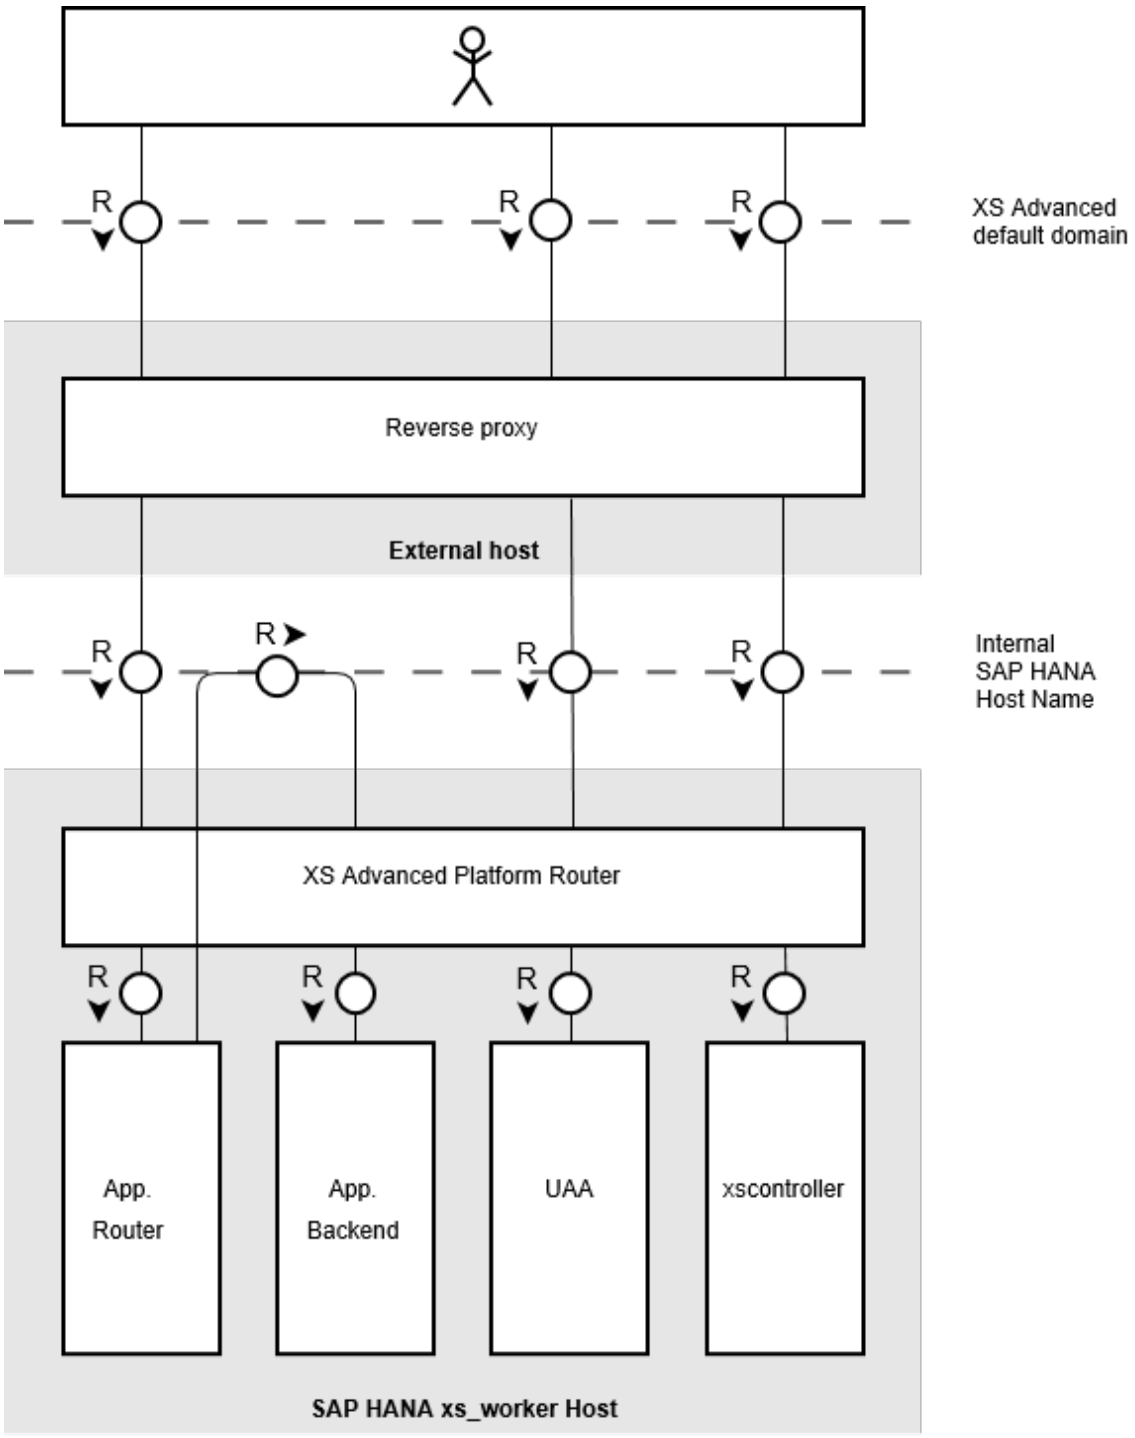

Request Flow with Direct Routing of Internal Requests

You can route internal requests directly to the XS advanced Web Dispatcher using a local DNS proxy or altering the `/etc/hosts` file.

Using a Local DNS Proxy

The most convenient way to route internal requests directly to the XS advanced Web Dispatcher is to set up a local DNS proxy server on each SAP HANA host with an `xs_*` role. Configure the DNS proxy on each host so that the following domains map to the host running the XS advanced Web Dispatcher:

- `<xsa-default-domain>`
- *.`<xsa-default-domain>`

Altering /etc/hosts

Another way to route requests directly to backend applications is by altering the `/etc/hosts` file on all SAP HANA hosts with an xs_* role. The mapping depends on the routing mode:

- Port-based routing mode: You just must map the XS advanced default domain.
- Hostname-based routing mode: You must map the XS advanced default domain and the sub-domains used by XS advanced applications and system services individually. This is because `/etc/hosts` mapping does not support wildcard mappings.

6.5 System Properties

SAP HANA system properties and XS advanced runtime properties.

System Properties

SAP HANA System Properties

Property	Description
Local Host Name	Specifies the host name of the machine.
SAP HANA System ID	Specifies a system ID. The SAP system ID (SID) is the identifier for the SAP HANA system.
Instance Number	Specifies the instance number of the SAP HANA system.
Local Host Worker Group	Specifies the worker group of the host. If undefined, the worker group is named "default". If you are using extension node for Business Warehouse, you must name the worker group "worker_dt".
System Usage	Specifies the usage type of the system to be installed. This setting is stored in the `global.ini` file, and can be used to identify the intended usage of the system.

Property	Description
Restrict maximum memory allocation?	Specifies whether maximum memory allocation is restricted for a new system. Specifies the maximum memory allocation for a new system in MB.
Restart system after machine reboot?	Restarts system after machine reboot.
Location of Data Volumes	Specifies the path to the data directory of the SAP HANA system.
Location of Log Volumes	Specifies the path to the log directory of the SAP HANA system.
Location of Persistent Memory Volume	Specifies the location of persistent memory volumes. This setting is stored in the `global.ini` file.
Edit Certificate Host...	Specifies the hostname used for generation of self-signed SSL certificates for the SAP Host Agent.
System Administrator User ID	Specifies the user ID of the system administrator. This parameter is relevant only if the operating system administrator (`<sid>adm`) does not exist prior to installation.
ID of User Group (sapsys)	Specifies the SAP system (sapsys) group ID. This parameter is relevant only if a sapsys group does not already exist on the host. If a sapsys group already exists, passing the `groupid` parameter does not alter the existing group.
System Administrator Login Shell	Specifies a system administrator login shell. This parameter is relevant only if the operating system administrator (`<sid>adm`) does not exist prior to installation.
System Administrator Home Directory	Specifies the home directory of the system administrator. This parameter is relevant only if the operating system administrator (`<sid>adm`) does not exist prior to installation.

XS Advanced Runtime Properties

The following table lists the options that you have to specify during the installation of XS advanced and provides a brief description of each option. You can also find the default value and the option to use to set a value manually, for example, using the `hdblcm` command:

XS Advanced Runtime Properties

Property	Description	Default	hdblcm Non-interactive Option
Automatically assign XS Advanced Runtime roles to the hosts with database roles	Assigns XS_WORKER and XS_STANDBY host roles. The host role XS_WORKER will be assigned to all worker hosts, the host role XS_STANDBY will be assigned to all standby hosts. To create a multiple-host system with dedicated xs_worker and xs_standby hosts, assign host roles to each host individually during installation. Do not choose the option to assign XS Advanced host roles automatically. **Note** The first XS_WORKER host added runs the `xscontroller` and `xsuaaserver` services. For more information, see *Multi-host Setup with XS Advanced* in *Related Information* below.	Yes	`--autoadd_xs_roles[=off]`
XS Advanced Tenant Database	The tenant database in which the XS advanced runtime is installed. **Note** For more information, see *XS Advanced Database Setup Options* in *Related Information* below.	Default tenant database	`--xs_runtime_db_tenant=<tenant name or SYSTEMDB>`
XS Advanced App Working Path	Specifies the XS advanced runtime app working path for runtime data of application instances. For best performance, specify a local directory, which is then created on all XS_WORKER hosts.	`/hana/shared/<SID>/app_working`	`--xs_app_working_path=<path>`

Property	Description	Default	hdblcm Non-interactive Option
Organization Name For Space "SAP"	Sets the initial name of the customer organization, which can then be changed if required. Organizations enable developers to collaborate by sharing resources, services, and applications. Access to the shared resources, services, and applications is controlled by the organization manager. **Note** For more information about maintaining organizations and spaces in XS Advanced, see *Related Information* below.	org	`--org_name=<org_name>`
XS Advanced Admin User	Creates an SAP HANA XS advanced runtime admin user. An admin user can add and manage users, view users, edit organization roles, view the organization quota, and perform other administrative tasks.	`XSA_ADMIN`	`--org_manager_user=<username>--org_manager_password=<password>`
Customer Space Name	Sets the initial name of the customer space for the SAP HANA XS advanced runtime. The space name can be changed later, if required. In an organization, spaces enable users to access shared resources that can be used to develop, deploy, and maintain applications.	PROD	`--prod_space_name=<prod_space_name>`
Routing Mode	Specifies the routing mode to be used for XS advanced runtime installations. **Note** The routing mode cannot be changed after installation. For more information on routing configurations, see *Routing Mode* in *System Concepts for XS Advanced Runtime Installations* in *Related Information* below.	hostname-based routing	`--xs_routing_mode=[hostnames|ports]`

Property	Description	Default	hdblcm Non-interactive Option
XS Advanced Domain Name	Specifies the domain name of an xs_worker host. The domain name has to resolve to the SAP HANA host which is running the **xscontroller** and **xsuaaserver** service. The default domain can also resolve to a host with a reverse proxy forwarding requests to the SAP HANA host that is running **xscontroller** and **xsuaaserver** services. **Note** For more information, see *Default Domain Configuration in XS Advanced* in *Related Information* below.	No default name	`--xs_domain_name=<Domain Name>`
Run Applications in Customer Space with Separate OS User	Run applications in customer space with a separate OS user. For more information, see SAP Note 2243156 and *Spaces and OS Users* in *Related Information* below. **Note** "If you specify “No”, the processes of applications provided by SAP are not separated from processes of applications provided by you.	Yes	`--xs_sap_space_isolation[=off]`
XS Advanced SAP Space OS User ID	OS user ID used for running XS advanced runtime applications in SAP space	1002	`--xs_sap_space_user_id=<uid>`
XS Advanced Customer Space OS User ID	OS user ID used for running XS Advanced applications in customer space.	1003	`--xs_customer_space_user_id=<uid>`

Local Secure Store Properties

Local Secure Store Properties

Property	Description
Installation Path for Local Secure Store	Specifies the installation path for local secure store.
Local Secure Store User ID	Specifies the user ID for the local secure store.
Local Secure Store User Group ID	Specifies the group ID for the local secure store.
Local Secure Store User Home Directory	Specifies the user home directory for the local secure store.

Property	Description
Local Secure Store User Login Shell	Specifies the user login shell for the local secure store.

Related Information

Parameter Reference [page 233]
XS Advanced Database Setup Options [page 106]
Multi-Host Setup with XS Advanced (Administration Guide for SAP HANA Platform)
System Concepts for XS Advanced Runtime Installations [page 99]
Default Domain Configuration in XS Advanced
Maintaining Organizations and Spaces in XS Advanced (SAP HANA Admin Guide)
Spaces and Operating System (OS) Users (SAP HANA Security Guide)

7 Updating the SAP HANA System

SAP HANA system components like the SAP HANA client, SAP HANA studio, and additional system components like Application Function Libraries (AFL, EML and the product-specific AFLs IBP, RTL, TRP, VCH, XRP), SAP liveCache applications (SAP LCA or LCAPPS plug-in), SAP HANA EPM-MDS (EPMMDS plug-in), XS advanced runtime applications, SAP HANA smart data access (SDA), or SAP HANA XS Advanced Runtime can be updated using the SAP HANA database lifecycle manager (HDBLCM).

To update an SAP HANA system, you need to first download the individual components from Service Marketplace (SMP). This can be done manually, from the SAP HANA studio, or 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.

Before starting the update for the components in your SAP HANA landscape several prerequisites have to be fulfilled. For more information, see *Prerequisites* in the *SAP HANA Master Guide*.

If you need to upgrade the operating system first, disable the automatic start of the SAP HANA system. For more information, see *autostart* in Related Information and *Upgrade an Existing SAP HANA System* in the *SAP HANA Master Guide*.

We recommend to perform an SAP HANA system update from a local host, a host that is part of the SAP HANA system, which you are logged on to. Performing an update from a local host minimizes the risk of a failed update due to network glitches.

For security reasons, the SYSTEM user might not be available during a system update. It is, therefore, required to create a lesser-privileged database user for updating a system. For more information, see *Create a Lesser-Privileged Database User for Update* in Related Information.

> **Note**
>
> If you have SAP HANA options installed, review the section about tenant databases in the administration guide of the corresponding option for additional information before proceeding. Be aware that you need additional licenses for SAP HANA options and capabilities. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities*.

Related Information

Before Updating [page 149]
Updating [page 165]
Create a Lesser-Privileged Database User for Update [page 157]
autostart [page 239]
Updating an SAP HANA System Landscape
SAP Note 2115815
SAP Note 2422421

7.1 Before Updating

An SAP HANA system can be updated from a local host using the SAP HANA database lifecycle manager.

It is possible to check for available updates of the SAP HANA system and components using the SAP HANA database lifecycle manager Web user interface or from the SAP HANA studio. If updates are available, the software must be downloaded, and the archive prepared before the update software can be installed.

There are three methods for downloading an SAP HANA installation medium:

- Download from SAP Support Portal
  - Download the installation medium - This method is only applicable when updating to a new SPS
  - Download individual components
- Use the *Check for Updates* functionality in the SAP HANA studio
  See Related Information, for more information about using the SAP HANA studio to check for available software component updates and download them from SAP Support Portal.
- Use the Maintenance Planner in the SAP Solution Manager
  If you have SAP Solution Manager, you can use it to update your SAP HANA system like other SAP systems in your landscape. In this case, you have to register your SAP HANA system using the System Landscape Directory (SLD). Once this configuration is performed, the SAP HANA database server regularly updates the SLD with data about your SAP HANA system. Afterward, SAP Solution Manager can access this data to calculate updates for this system. You can then use the Maintenance Planner in SAP Solution Manager to generate an SPS and download the necessary archives to a specified location. For more information about updating with the Maintenance Planner, see Related Information. You can then provide the SPS location to the SAP HANA database lifecycle manager during system update.

> ⚠ Caution
>
> The installation medium cannot be patched manually with newer versions of SAP HANA components and used for installation or update. The result is an error-prone installation.

> ⓘ Note
>
> Before updating your SAP HANA system to SAP HANA 2.0, ensure that all applications that use this system are released for SAP HANA 2.0.

Related Information

Configure the Connection to SAP Support Portal [page 150]
Download Components from SAP Support Portal Using the Web User Interface [page 150]
Prepare the Software Archive for the Update
Prepare SAP HANA Packages for Installation and Update [page 153]
Prepare an Update for Flexible System Downtime [page 154]
Create a Lesser-Privileged Database User for Update [page 157]
Upload and Extract SAP HANA Components Using the Web User Interface [page 192]
Maintenance Planner

7.1.1 Configure the Connection to SAP Support Portal

Before downloading software components using the SAP HANA studio, configure the connection to SAP Support Portal. As of SPS 12, configuring the connection is only required if you want to connect to and update an SAP HANA system version SPS 11 or lower.

Procedure

1. Start the SAP HANA studio.
2. Select *Window* > *Preferences* > *SAP HANA* > *Lifecycle Management* > *Service Marketplace*.
3. Select *Use SMP configuration*.

   The *Host Name* is set to *service.sap.com*. This is the URL of SMP, where the files that are used for the update are located. Enter your SMP S-User credentials in *User Name/Password*.

   > ⓘ Note
   >
   > You can obtain an S-User on SMP at `https://service.sap.com` under *Registration*.

4. Select *Apply*, then *OK* to complete the configuration.

   > ⓘ Note
   >
   > To reset your entries, choose *Restore Defaults*.

Next Steps

After configuring the connection from the SAP HANA studio to SAP Support Portal, download the software components from the SAP Support Portal. For more information about downloading software components using the SAP HANA studio, see Related Information.

7.1.2 Download Components from SAP Support Portal Using the Web User Interface

You can use your SAP HANA database lifecycle manager (HDBLCM) Web user interface to check for available software component updates and download them from SAP Support Portal.

Prerequisites

- The SAP HANA database server is up and running.

- The SAP HANA system has access to the Internet.

> **Note**
>
> Alternatively, you can download the components to a shared location to which the SAP HANA system has access or copy the downloaded components to the SAP HANA system manually.

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
  Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
  - Microsoft Windows
    - Internet Explorer - Version 9 or higher
      If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing *Tools* > *Compatibility View Settings*.
    - Microsoft Edge
    - Mozilla Firefox - Latest version and Extended Support Release
    - Google Chrome - Latest version
  - SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
  - Mac OS - Safari 5.1 or higher

> **Note**
>
> For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: https://`<hostname>`:1129/lmsl/HDBLCM/`<SID>`/index.html **Note** The URL is case sensitive. Make sure you enter upper and lower case letters correctly.

Option	Description
SAP HANA cockpit	1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. `https://<host_FQDN>:<port>` **Note** FQDN = fully qualified domain name 2. Drill down on the name of the system from *My Resources* or from a group. 3. The links in *Platform Lifecycle Management* each launch additional functionality, giving you expanded capabilities for managing the resource.

2. Select the *Download Components* tile.
3. Specify the download mode. Then select *Next*.

   You can choose between downloading the software archives to the SAP HANA host or via the Web browser to your local computer. If you choose to download the archives to your local computer, you will have to copy them to the host manually.
4. Specify the HTTPS proxy properties. Then select *Next*.
5. Provide the SAP Support Portal credentials, then select *Next*.
6. Select the components for download. Then select *Next*.

   Note that you can also select and download SAP HANA core components which are not currently, but can be installed on the system. To display all components available for download, select *Show Components (Only Updates)* and then select *All*.
7. Specify the download properties. Then select *Next*.
8. Select *Download* to download the components.
9. Once the downloads have finished, select *Close* to return to the main screen.

Next Steps

After downloading components from the SAP Support Portal using the SAP HANA database lifecycle manager (HDBLCM) Web user interface, prepare the software archive so that it is detected by the SAP HANA database lifecycle manager during update. For more information about preparing the software archive, see Related Information.

Related Information

Prepare the Software Archive for the Update

Upload and Extract SAP HANA Components Using the Web User Interface [page 192]

7.1.3 Prepare SAP HANA Packages for Installation and Update

In SAP HANA systems version SPS 10 or lower, the archive can be prepared manually for installation and update by making local package directories, that are recognized by the update tool.

Context

In order to ensure security during its operation, SAP HANA database lifecycle manager (hdblcm) needs to work with SAP HANA database component, officially signed by SAP. The signature of the component contents is contained in a file, called SIGNATURE.SMF, which is produced by SAPCAR during extraction of the archive.

If for some reason you cannot get this signature file or the SAP HANA database component, which you use had not been signed with the official SAP signature, you can still perform the installation or update. The database itself will not be affected by the missing signature. To do so, you need to use hdblcm or hdblcmgui as root user and pass the parameter --ignore=check_signature_file on the command line. If you do so, the following features of the database lifecycle management tools (hdblcm, hdbinst, hdbupd, etc.) cannot be used afterwards until the next database update with valid signature.

- Web user interface
- Execution as <sid>adm user
- Remote execution via SAP Host Agent

> **Note**
>
> If you extract more than one component SAR into a single directory, you need to move the `SIGNATURE.SMF` file to the subfolder (`SAP_HANA_DATABASE`, `SAP_HANA_CLIENT`, etc.), before extracting the next SAR in order to avoid overwriting the `SIGNATURE.SMF` file. For more information, see also SAP Note 2178665 in Related Information.

Procedure

1. Create a local directory for the downloaded packages.

   Example: `/hana/local/downloads`
2. Download SAP HANA packages from the SAP Support Portal to the local directory.

   The Debug Symbol Packages (*.TAR) that are also available for download are only needed for troubleshooting purposes and not required for a system update. For more information, see SAP Note 2145573 in Related Information.
3. Create a directory, in to which the package contents can be unpacked.

   Example: `/hana/local/downloads/install`
4. Unpack the SAP HANA database archive in to the local directory.

   Example:`/usr/sap/hostctrl/exe/SAPCAR -manifest SIGNATURE.SMF -xvf IMDB_SERVER<version number>.SAR`

5. Unpack the component archives to a local directory:

   Run the SAP HANA platform LCM tool using the parameter `extract_components` as a call option.

Next Steps

You can now update the SAP HANA system with the SAP HANA database lifecycle manager.

Related Information

SAP Note 2178665
SAP Note 2577617
SAP Note 2145573
extract_components [page 252]
Update an SAP HANA System Using the Graphical User Interface [page 166]
Update an SAP HANA System Using the Command-Line Interface [page 168]
Update an SAP HANA System Using the Web User Interface [page 171]

7.1.4 Prepare an Update for Flexible System Downtime

You can run an SAP HANA system update in two phases - an update preparation phase and a resume update phase. You can perform the prepare update phase and the update resume phase using either the SAP HANA database lifecycle manager graphical user interface, command-line interface or Web user interface.

Prerequisites

- You are updating to a new SPS from an installation medium or you have prepared for update, either in the SAP HANA studio or manually.
- With system replication active, you have updated the secondary system before the primary system. The version of the secondary system must be equal to or higher than the version of the primary system.
- You have performed a system backup. Also note that during the update there is a business downtime for your SAP HANA
- You know the `<sid>adm`, and database administrator passwords.
- You have applied a valid license key for the SAP HANA system.

Context

After downloading the SAP HANA software, and preparing the downloaded archives for update execution, you have the choice to update your SAP HANA system in one step, or to update it in a phased approach to minimize system downtime.

When you start the SAP HANA database lifecycle manager with the `prepare_update` flag set, the SAP HANA database lifecycle manager extracts the packages (like the SAP Host Agent, and delivery units) from the new source, but does not actually perform the update. During the preparation phase the system is not modified by the installer or restarted. The software switch occurs when the SAP HANA database lifecycle manager is run a second time, resuming the system update.

The phased update aims to:

- Lower the system downtime
- Reduce the chances of a failed system update due to preliminary steps like archive preparation or dependency conflicts

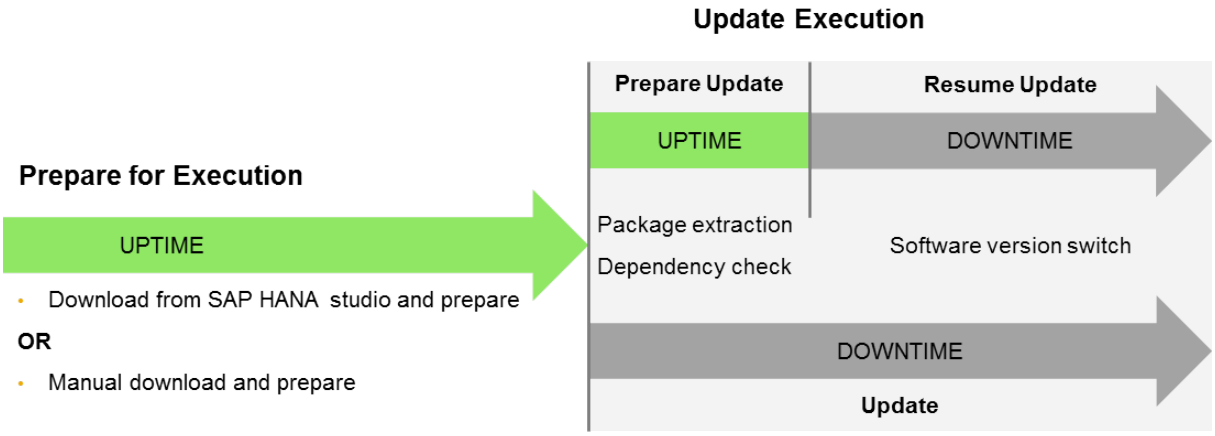

- Before Updating [page 149]
- Prepare an Update for Flexible System Downtime [page 154]
- Update an SAP HANA System Using the Graphical User Interface [page 166]
- Update an SAP HANA System Using the Graphical User Interface [page 166]

> **Note**
>
> It is also possible to use system replication to achieve near zero downtime upgrades. For more information, see *Use System Replication for Near Zero Downtime Upgrades* in the *SAP HANA Administration Guide*.

Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
IBM Power Systems	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`

> **Note**
>
> If you downloaded the components to a different directory, change to the directory where you unpacked the archive.

2. Perform the update preparation phase step with the SAP HANA database lifecycle manager using one of the following commands.

Option	Description
Graphical user interface	`./hdblcmgui --action=update --prepare_update`
Command-line interface	`./hdblcm --action=update --prepare_update`

Provide the required system update information and credentials. See Related Information for more details about SAP HANA system update.

Before triggering the update preparation phase, confirm that the following line is listed in the action summary under *Update Parameters*:

```
Stop update before software version switch, resumable: Yes
```

If you are using the SAP HANA HDBLCM Web user interface, open the *Advanced Parameters Configuration* dialog from the footer bar and select *Prepare Update Only* under *General Parameters*.

3. Resume the update.

During the planned maintenance window, you can resume the prepared update using any of the standard update procedures. For standard SAP HANA system update procedures, see Related Information.

Related Information

Before Updating [page 149]
Update an SAP HANA System Using the Graphical User Interface [page 166]
Update an SAP HANA System Using the Command-Line Interface [page 168]
Update an SAP HANA System Using the Web User Interface [page 171]
prepare_update [page 270]
SAP Note 2407186

7.1.5 Create a Lesser-Privileged Database User for Update

As the most powerful database user, SYSTEM is not intended for day-to-day activities in production environments. Create a lesser-privileged database user in the system database for updating an SAP HANA system.

Context

For security reasons, the SYSTEM user is not intended for day-to-day activities in production environments. It is, therefore, required to perform the SAP HANA system update as another user.

> **Note**
>
> A lesser-privileged database user cannot be used to update an XS advanced runtime installation. The SYSTEM user is required and needs to be temporarily reactivated for the duration of the update.

A system database user should be created with the following granted roles and object privileges:

- Granted Roles: CONTENT_ADMIN
- Object Privileges: _SYS_REPO (SELECT)

By calling one of the update LCM tools with the `system_user` option specified, the previously defined database user is used in place of SYSTEM to authenticate the configuration task.

Procedure

1. Add the SAP HANA system database in the SAP HANA cockpit.

   For more information, see *Working with Resources and Resource Groups* in the *SAP HANA Administration Guide*.
2. Create a role in the SAP HANA cockpit with the CONTENT_ADMIN role and the object privilege SELECT on the catalog object _SYS_REPO.
3. Create a user in the SAP HANA cockpit and grant the new user the role you created in step 2.

   For more information, see *Provisioning Users in SAP HANA Cockpit* in the *SAP HANA Administration Guide*.
4. Add the SAP HANA system database in the SAP HANA cockpit again as the new user.
5. Perform the SAP HANA system update as usual, specifying the new lesser-privileged user in place of the SYSTEM user when prompted.

Related Information

Registered Databases
Deactivate the SYSTEM User

system_user [page 282]

7.1.6 Updating a Single-Container System

During the update to SAP HANA 2.0 SPS 01 or greater, all single-container systems are converted to support tenant databases.

As of SAP HANA 2.0 SPS 01, the multi-container database mode is the only database mode. A single-container system will be automatically converted to a tenant database system during the update. The following sections describe the changes made during the update.

The database of a single-container system is converted into a system database and a tenant database. A new user (SYSTEM) is created in the system database (SYSTEMDB). During the update, a password has to be specified for this user. The database superuser (SYSTEM) of the single-container system becomes the SYSTEM user of the tenant database.

If you have additional questions about the automatic conversion of your single-container system, or wish to discuss migration support, please contact your SAP support team representative.

Updating a Single-Host, Single-Container system to Support Tenant Databases

In a single-host, single-container system the database is converted to a system database and a single tenant database.

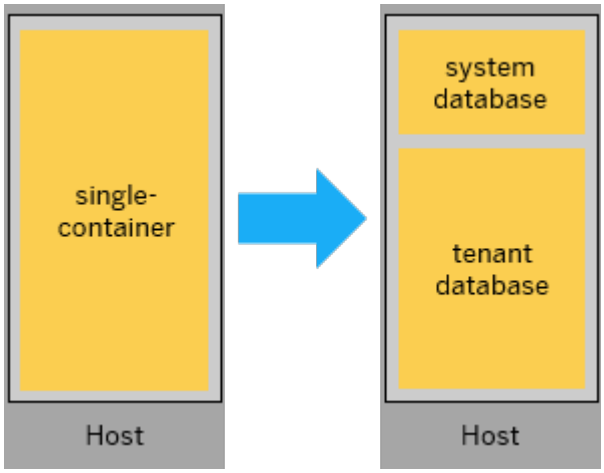

Updating a Multiple-Host, Single-Container system to Support Tenant Databases

In a multiple-host, single-container system the database is converted to a system database on the first master host and a single tenant database that is spread across the worker hosts.

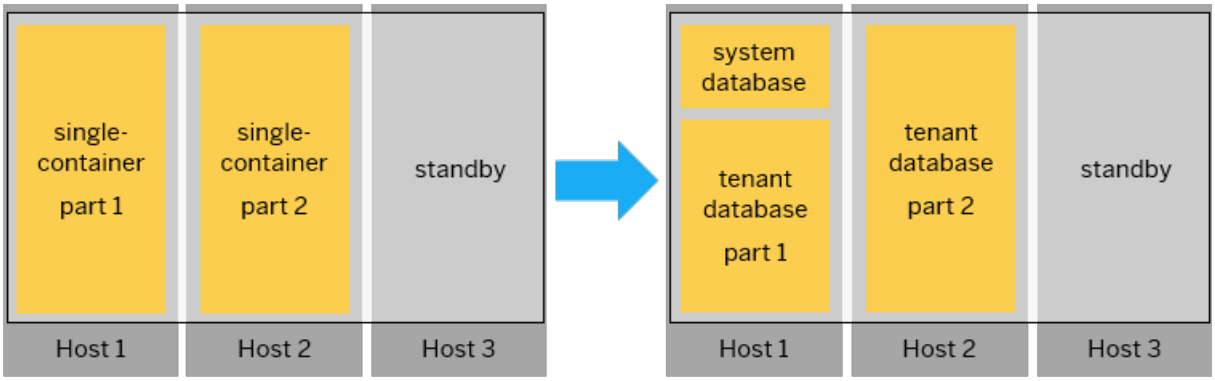

Related Information

SAP Note 2423367
Converting an SAP HANA System to Support Tenant Databases
Create a Tenant Database

7.1.6.1 Configuration

After update, you need to review and if necessary reconfigure certain settings.

Configuration Area	After Update
Database configuration	Parameters, which were changed are stored with the tenant database and become database-specific.
User administration	Users of the single-container database are now present in the tenant database. During the update, you are asked to provide a new password for the system user of the system database.
Network	Ports and URLs do not change after the update. The tenant database retains the port numbers of the original single-container system: 3<instance>03 (internal communication), 3<instance>15 (SQL), and 3<instance>08 (HTTP via SAP HANA classic server). The port number of the **system database** are fixed: 3<instance>01 (internal), 3<instance>13 (SQL), and 3<instance>14 (HTTP via XS classic server).
SAP HANA options	If you were running SAP HANA options on your single-container system, no configuration changes are required after the update. SAP HANA options hosts or host roles can be automatically provisioned to a system that is installed with a single tenant. If the SAP HANA system contains multiple tenant databases or if the single tenant was not running when you added the SAP HANA options host or host role, it must be manually provisioned to the tenant. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities*.

Configuration Area	After Update
XS advanced runtime	If you have XS advanced runtime installed, a separate **xsengine** process is created and the internal Web Dispatcher of the SAP HANA system routes by default to the single tenant.

Related Information

Important Disclaimer for Features in SAP HANA
SAP HANA Server Installation and Update Guide [page 9]
Port Assignment in Tenant Databases

7.1.6.2 Security

After update, you need to review and if necessary reconfigure certain security-related settings.

Configuration Area	After Update
User administration	You need to set up new users for administration (at least for recovery) in the system database.
Network security	The system database uses additional ports. These ports that might be firewalled for security reasons in the system. If this the case, you need to open these ports so that the system database can be accessed from the SAP HANA cockpit on other hosts. The system database is accessible via SQL port 3`<instance_no>`13.
TLS/SSL configuration for external communication	If TLS/SSL is enabled for both the system database and tenant database, the in-database certificate collection containing the certificates used for trust validation is available only in the tenant database. If you want to use the same certificates, you will need import them into the certificate store of the system database and add them to a certificate collection there. ⚠ Caution If TLS/SSL is being enforced for client connections (that is, parameter `[communication] sslEnforce` in *global.ini* set to `true`), it will not be possible to establish a connection to the system database. You have to set `sslEnforce` to `false` first. If the certificates used for trust validation are stored in a PSE in the file system, both the tenant database and the system database will have access, so no reconfiguration is required. However, you should validate that sharing the certificate stores for system database and tenant is actually intended.

Configuration Area	After Update
Database isolation	The default system isolation level is low. It is possible to change the isolation level (from low to high or from high to low) at any time after the update. Once high isolation has been configured, a dedicated OS user and group must exist for every tenant database. Otherwise, it's not possible to create or start a tenant database.
Auditing	Existing audit policies are available in the tenant database database only. You need to create new audit policies for administration tasks in the system database.

Related Information

User Administration Tools
TLS/SSL Configuration on the SAP HANA Server
Increase the System Isolation Level
Decrease the System Isolation Level

7.1.6.3 Backup and Recovery

After update, you need to review and if necessary reconfigure certain settings related to backup and recovery.

Configuration Area	After Update
System database	If you do not change your backup configuration, a backup of the tenant database is created by default. It can be restored into any existing tenant database. The system database has to be backed up separately.
Third-party tools	If your backup strategy is based on data snapshots taken by a third-party tool, we recommend that you get in touch with your snapshot tool vendor to check if the tool supports snapshots of SAP HANA tenant database systems. If you use a script that is based on the documented SQL statements, refer to the documentation to adapt it for use with tenant databases.
Data snapshots	As of SAP HANA 2.0 SPS 04, data snapshots are supported on single-tenant and multiple-tenant systems.
Retention policy	Without adjusting your backup retention policy, only the tenant database backups are maintained. As a consequence, the system database backup catalog grows unchecked, consumes main memory, and prolongs backup times. Furthermore, system database data and log backups will eat up your backup space.
Backup history	The migration from single-container mode to tenant databases does not break the backup history. Single database data and log backups can be used to recover a tenant database system.
Disaster recovery	In the event of a disaster, you first have to recover the system database in operational mode offline and then the tenant database using the system database connectivity. The system database requires to be in operational state online while recovering the tenant database.

Related Information

SAP HANA SQL Reference Guide for SAP HANA Platform
Recovering an SAP HANA Database

7.1.6.4 Landscape

After update, you need to review and if necessary reconfigure certain landscape-related settings.

Configuration Area	After Update
Host addition and removal	After the update, hosts can be added to or removed from a single-host or multiple-host SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM).
Scale-out	To scale out a tenant database or distribute it across multiple hosts, you can add further server components, for example, an additional index server or a separate XS server. You add a service to a tenant database using the ALTER DATABASE statement. The statement is executed on the system database for a specific tenant database.
Service addition and removal	Add or remove service statements are executed on the system database for a specific tenant database.

Related Information

Adding and Removing Hosts
Add Services in a Tenant Database

7.1.6.5 Perform an Offline Update

You can perform an offline update of an SAP HANA system replication landscape. A single-container system will be automatically converted to a tenant database system during the update. Converting an SAP HANA system to a tenant database system is permanent and cannot be reversed.

Prerequisites

- The statistics server is **not** running as a separate server process (`statisticsserver`), but instead as an embedded service in the master index server. If this is not the case, migrate the statistics server to the embedded statistics service as described in SAP Note 1917938.
- The SAP HANA system has been installed with its server software on a shared file system (export options `rw`, `no_root_squash`).

- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- You are logged on as the system administrator user `<sid>adm`.

Procedure

1. Stop all SAP HANA systems on all sites.
2. Update the primary system using the SAP HANA database lifecycle manager. The migration to a tenant database system is done automatically.
3. Wait until the update has finished and the system is active again.
4. Create a data backup of the system database.
5. Prepare the secondary system for authentication by copying the system PKI SSFS `.key` and the `.dat` file from the primary system to the secondary system. For more information, see SAP Note 2369981.

   The `.key` and `.dat` files can be found in the following location:

   ```
   /usr/sap/<SID>/SYS/global/security/rsecssfs/data/SSFS_<SID>.DAT
   /usr/sap/<SID>/SYS/global/security/rsecssfs/key/SSFS_<SID>.KEY
   ```
6. Repeat steps 2, 3, and 5 on all remaining secondary systems, following the replication chain.

Related Information

SAP Note 1917938
SAP Note 2369981
Configuring SAP HANA System Replication

7.1.6.6 Perform a Near-Zero Downtime Update

You can perform a near-zero downtime update of an SAP HANA system replication landscape. A single-container system will be automatically converted to a tenant database system during the update. Converting an SAP HANA system to a tenant database system is permanent and cannot be reversed.

Prerequisites

- The statistics server is **not** running as a separate server process (`statisticsserver`), but instead as an embedded service in the master index server. If this is not the case, migrate the statistics server to the embedded statistics service as described in SAP Note 1917938.
- The SAP HANA system has been installed with its server software on a shared file system (export options `rw`, `no_root_squash`).

- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- You are logged on as the system administrator user `<sid>adm`.
- Full data shipping has been successfully completed after setting up the system replication landscape. For more information, see *Initializing the Secondary* in the SAP HANA Administration Guide.

Procedure

1. Update the secondary system using the SAP HANA database lifecycle manager. The migration to a tenant database system is triggered automatically.
2. Wait until the update has finished and all systems are in sync again. The replication will be possible in this situation although the primary is still a single-container system.
3. Perform a takeover to the updated secondary system. Only now the migration to a tenant database system is finalized for the secondary.
4. Update the primary system using the SAP HANA database lifecycle manager. Set the `no_restart` option to prevent a restart of the primary system after the update. The migration to a tenant database system is done automatically.
5. Prepare the secondary system for authentication by copying the system PKI SSFS `.key` and the `.dat` file from the primary system to the secondary system. For more information, see SAP Note 2369981.

   The `.key` and `.dat` files can be found in the following location:

   ```
   /usr/sap/<SID>/SYS/global/security/rsecssfs/data/SSFS_<SID>.DAT
   /usr/sap/<SID>/SYS/global/security/rsecssfs/key/SSFS_<SID>.KEY
   ```

6. Register this former primary system as the new secondary to the new primary (former secondary) and start it. The conversion to a tenant database system is performed automatically.

Related Information

SAP Note 1917938
Configuring SAP HANA System Replication
Initializing the Secondary

7.2 Updating

The SAP HANA database lifecycle manager (HDBLCM) is the program used to update an SAP HANA system, including server, client, studio, and additional components, in a graphical user interface, the command-line interface, or the Web user interface.

Optimized Update

As of SPS 11, you can run an optimized update of an SAP HANA system to reduce the number of restarts and system downtime. You can perform the update using either the SAP HANA database lifecycle manager graphical user interface, command-line interface or Web user interface.

After downloading the SAP HANA software, and preparing the downloaded archives for update execution, you have the choice to update your SAP HANA system in standard mode, or to update it using optimized mode to minimize the number of restarts. The optimized update execution mode is enabled by default, if more than one component that supports the phased update process is selected to be installed or updated. Usually there is no need to specify this option explicitly.

The following components support the optimized update:

- SAP HANA server
- SAP HANA Accelerator for SAP ASE
- SAP HANA Dynamic Tiering
- SAP HANA Streaming Analytics

The optimized execution mode is enabled when you install or update more than one of these components. Components which do not support optimized update will be installed or updated in the most appropriate moment.

When you start the SAP HANA database lifecycle manager with `update_execution_mode` set to `optimized`, the SAP HANA database lifecycle manager:

1. Prepares the components for installation. In combination with the parameter `prepare_update` this step can be executed before the actual update to reduce the system downtime during the maintenance window and make sure that all checks pass.
2. Stops the system or individual instances.
3. Updates the software.
4. Starts the system or individual instances.
5. Resumes the update of the components.

The optimized update aims to:

- Reduce the number of system restarts
- Lower the system downtime
- Reduce the chances of a failed system update due to dependency conflicts or errors in preliminary steps, such as archive preparation
- Ensure consistency of component binaries during system start

Related Information

Before Updating [page 149]
Update an SAP HANA System Using the Graphical User Interface [page 166]
Update an SAP HANA System Using the Command-Line Interface [page 168]
Update an SAP HANA System Using the Web User Interface [page 171]
Changeable Default Values for Update [page 174]
Use SAP HANA System Replication for Near Zero Downtime Upgrades [page 177]
update_execution_mode [page 285]

7.2.1 Update an SAP HANA System Using the Graphical User Interface

All SAP HANA platform components can be updated using the SAP HANA database lifecycle manager (HDBLCM) graphical user interface.

Prerequisites

- You are updating to a new SPS from an installation medium or you have prepared for update, either in the SAP HANA studio or manually.
- With system replication active, you have updated the secondary system before the primary system. The version of the secondary system must be equal to or higher than the version of the primary system.
- You have performed a system backup. Also note that during the update there is a business downtime for your SAP HANA system.
- You know the `<sid>adm`, and database administrator passwords.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.
- You have applied a valid license key for the SAP HANA system.
- The `<sid>adm` user is able to execute graphical applications.
- You have uninstalled SAP HANA Remote Data Sync from your system. For more information, see SAP Note 2641466.

Context

The following procedure describes the update of an SAP HANA system in interactive mode and entering parameters interactively. This procedure may also be performed in advanced interactive mode, with parameters entered as call options or from a configuration file. For more information about interaction modes and parameter entry methods, see *Using the SAP HANA Platform LCM Tools* in Related Information.

> **Note**
>
> Not all parameters are requested interactively. Some parameters have default values, that do not require confirmation in interactive mode. Those parameters must be specified as call options or from a configuration file using the graphical user or command-line interface and advanced interactive mode. For more information about changeable default values, see Related Information.

Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
IBM Power Systems	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`

   > **Note**
   >
   > If you downloaded the components to a different directory, change to the directory where you unpacked the archive.

2. Run the SAP HANA database lifecycle manager:

   ```
   ./hdblcmgui
   ```

   The SAP HANA database lifecycle manager graphical user interface appears.
3. Select a detected software component or add a software component location by selecting *Add Component Location*. Then select *Next*.
4. Select *Update Existing System*, and choose the SID from the drop-down menu. Then select *Next*.
5. Select the components you would like to update, then select *Next*.
6. Specify the SAP HANA authorization information.

   When asked for database user, you have the opportunity to specify a lesser-privileged database user if you have previously created one. For more information about creating a database user for update, see Related Information.
7. Define additional properties, depending on which components are selected.
8. After specifying all system properties, review the summary, and select *Update*.

Next Steps

SAP HANA content that was installed as part of the SAP HANA database is updated automatically on tenant databases. Restarting the tenant databases is not necessary to trigger the deployment of the updated

content on the tenant databases. You can monitor the progress of DU deployment by executing the following statement:

```
SELECT * FROM "PUBLIC"."M_SERVICE_THREADS" WHERE THREAD_TYPE = 'ImportOrUpdate
Content';
```

Non-automated content and delivery units that were not installed together with the SAP HANA database and are not part of the SAP HANA database need to be updated manually on the database where the content was initially installed.

If your system is configured for system replication, you must perform the update on each host individually, starting with the secondary host. The version of the secondary system must be the same or higher as the one running on the primary system

Related Information

Using the SAP HANA Platform LCM Tools [page 41]
Changeable Default Values for Update [page 174]
Use Advanced Interactive Mode to Perform Platform LCM Tasks [page 51]
Create a Lesser-Privileged Database User for Update [page 157]
SAP HANA Content
SAP Note 2641466

7.2.2 Update an SAP HANA System Using the Command-Line Interface

All SAP HANA platform components can be updated using the SAP HANA database lifecycle manager (HDBLCM) command-line interface.

Prerequisites

- You are updating to a new SPS from an installation medium or you have prepared for update, either in the SAP HANA studio or manually.
- With system replication active, you have updated the secondary system before the primary system. The version of the secondary system must be equal to or higher than the version of the primary system.
- You have performed a system backup. Also note that during the update there is a business downtime for your SAP HANA system.
- You know the `<sid>adm`, and database administrator passwords.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.
- You have applied a valid license key for the SAP HANA system.

- You have uninstalled SAP HANA Remote Data Sync from your system. For more information, see SAP Note 2641466.

Context

The following procedure describes the update of an SAP HANA system in interactive mode and entering parameters interactively. This procedure may also be performed in advanced interactive mode or batch mode, with parameters entered as call options or from a configuration file. For more information about interaction modes and parameter entry methods, see *Using the SAP HANA Platform LCM Tools* in Related Information.

> **Note**
>
> Not all parameters are requested interactively. Some parameters have default values that do not require confirmation in interactive mode. Those parameters must be specified as call options or from a configuration file. For more information about changeable default values, see Related Information.

Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
IBM Power Systems	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`

   > **Note**
   >
   > If you downloaded the components to a different directory, change to the directory where you unpacked the archive.

2. Locate the SAP_HANA_DATABASE directory, and run the SAP HANA database lifecycle manager:

   ```
   ./hdblcm
   ```

   > **Note**
   >
   > If your SAP HANA system has XS advanced runtime installed, you can specify XS advanced runtime components to be updated using the `xs_components` parameter. MTA extension descriptors (`*.mtaext`) can be specified using the `xs_components_cfg` parameter.

3. Select the index for the system to be updated, then select `Enter`.
4. Select the components you would like to update as a comma-separated list, then select `Enter`.

5. Specify the SAP HANA authorization information.

   When asked for database user, you have the opportunity to specify a lesser-privileged database user if you have previously created one. For more information about creating a database user for update, see Related Information.
6. Define additional properties, depending on which components are selected.
7. After specifying all system properties, review the summary, and select *y*.

Next Steps

SAP HANA content that was installed as part of the SAP HANA database is updated automatically on tenant databases. Restarting the tenant databases is not necessary to trigger the deployment of the updated content on the tenant databases. You can monitor the progress of DU deployment by executing the following statement:

```
SELECT * FROM "PUBLIC"."M_SERVICE_THREADS" WHERE THREAD_TYPE = 'ImportOrUpdate
Content';
```

Non-automated content and delivery units that were not installed together with the SAP HANA database and are not part of the SAP HANA database need to be updated manually on the database where the content was initially installed.

If your system is configured for system replication, you must perform the update on each host individually, starting with the secondary host. The version of the secondary system must be the same or higher as the one running on the primary system

Related Information

Using the SAP HANA Platform LCM Tools [page 41]
Changeable Default Values for Update [page 174]
Create a Lesser-Privileged Database User for Update [page 157]
SAP HANA Content
SAP Note 2641466

7.2.3 Update an SAP HANA System Using the Web User Interface

All SAP HANA platform components can be updated using the SAP HANA database lifecycle manager (HDBLCM) Web user interface.

Prerequisites

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
  Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
  - Microsoft Windows
    - Internet Explorer - Version 9 or higher
      If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing *Tools* > *Compatibility View Settings*.
    - Microsoft Edge
    - Mozilla Firefox - Latest version and Extended Support Release
    - Google Chrome - Latest version
  - SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
  - Mac OS - Safari 5.1 or higher

  > **Note**
  >
  > For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.
- The installation medium must be owned by the root user and should not have write permissions for the group (except for when the group ID is 0) and others.
- You are updating to a new Support Package Stack (SPS) from an installation medium or you have prepared for update, either in the SAP HANA studio or manually. For more information, see *Prepare for Update Manually* in Related Information.
- With system replication active, you have updated the secondary system before the primary system. The version of the secondary system must be equal to or higher than the version of the primary system.
- You have performed a system backup. Also note that during the update there is a business downtime for your SAP HANA system.
- You have applied a valid license key for the SAP HANA system.

- You have uninstalled SAP HANA Remote Data Sync from your system. For more information, see SAP Note 2641466.

Context

SAP HANA content that was installed as part of the SAP HANA database is updated automatically on tenant databases. Restarting the tenant databases is not necessary to trigger the deployment of the updated content on the tenant databases. You can monitor the progress of DU deployment by executing the following statement:

```
SELECT * FROM "PUBLIC"."M_SERVICE_THREADS" WHERE THREAD_TYPE = 'ImportOrUpdate
Content';
```

Non-automated content and delivery units that were not installed together with the SAP HANA database and are not part of the SAP HANA database need to be updated manually on the database where the content was initially installed.

> **Note**
>
> Not all parameters are requested interactively. Some parameters have default values that do not require confirmation in interactive mode. Those parameters must be specified as call options or from a configuration file. For more information about changeable default values, see Related Information.

Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: https://<hostname>:1129/lmsl/HDBLCM/<SID>/index.html **Note** The URL is case sensitive. Make sure you enter upper and lower case letters correctly.
SAP HANA cockpit	1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. `https://<host_FQDN>:<port>` **Note** FQDN = fully qualified domain name 2. Drill down on the name of the system from *My Resources* or from a group. 3. The links in *Platform Lifecycle Management* each launch additional functionality, giving you expanded capabilities for managing the resource.

2. Select the *Update System and Components* tile.

3. Enter the file path of the installation medium in the location field:

Option	Description
Intel-Based Hardware Platforms	`<installation medium>/DATA_UNITS/HDB_SERVER_LINUX_X86_64`
IBM Power Systems	`<installation medium>/DATA_UNITS/HDB_SERVER_LINUX_PPC64`

   If you downloaded the components to a different directory, enter the file path to the directory where you unpacked the server archive.

4. Select *Proceed with Update*.

   The SAP HANA database lifecycle manager (HDBLCM) detects all available components for the given file path.

   If you would like to add more than one software location, select *Add Software Location*.
5. Select the components you would like to update, or install if they are not already available on your system.
6. Specify the SAP HANA authorization information.

   When asked for database user, you have the opportunity to specify a lesser-privileged database user if you have previously created one. For more information about creating a database user for update, see Related Information.
7. Define additional properties, depending on which components are selected.
8. After specifying all system properties, review the summary, and select *Update*.

Next Steps

SAP HANA content that was installed as part of the SAP HANA database is updated automatically on tenant databases. Restarting the tenant databases is not necessary to trigger the deployment of the updated content on the tenant databases. You can monitor the progress of DU deployment by executing the following statement:

```
SELECT * FROM "PUBLIC"."M_SERVICE_THREADS" WHERE THREAD_TYPE = 'ImportOrUpdate
Content';
```

Non-automated content and delivery units that were not installed together with the SAP HANA database and are not part of the SAP HANA database need to be updated manually on the database where the content was initially installed.

If your system is configured for system replication, you must perform the update on each host individually, starting with the secondary host. The version of the secondary system must be the same or higher as the one running on the primary system

Related Information

SAPUI5 Developer Guide
Add an SAP HANA System
Before Updating [page 149]
Prepare SAP HANA Packages for Installation and Update [page 153]
Changeable Default Values for Update [page 174]
Create a Lesser-Privileged Database User for Update [page 157]
SAP HANA Content
SAP Note 2641466

7.2.4 Changeable Default Values for Update

The SAP HANA database lifecycle manager (HDBLCM) uses the following default values during update unless you change them.

Some default values are based on the predefined values on the current host.

Changeable Parameter Defaults

Parameter	System Default Value	Interactive Mode Availability
`client_path`	`${sapmnt}/${SID}/hdbclient`	◇
`install_hostagent`	`y` (on)	◇
`install_ssh_key`	`y` (on)	◇
`remote_execution`	`ssh`	◇
`root_user`	`root`	■
`scope`	`system`	◇
`studio_path`	`${sapmnt}/${SID}/hdbstudio`	◇
`studio_repository`	`1` (on)	◇
`system_user`	`SYSTEM`	■

> **Note**
>
> The default paths written in the form `${<parameter>}` indicate that substitution of the parameter occurs in the configuration file and in batch mode. Substitution also occurs in interactive mode in order to create

a suggested path. The advantage of substitution is that the SAP system ID (`sid`) and the installation path (`sapmnt`, which is `/hana/shared`, by default) only need to be specified once, and are then substituted in to the other parameter values. This ensures that the system has unique file system paths if multiple systems are installed on the same host. However, if it is preferred to deviate from the default paths, it is necessary to pay attention to the settings, especially in the configuration file, and when installing in batch mode.

Related Information

client_path [page 244]
install_hostagent [page 259]
install_ssh_key [page 259]
remote_execution [page 271]
root_user [page 273]
scope [page 274]
studio_path [page 280]
studio_repository [page 281]
system_user [page 282]

7.2.5 Update SAP HANA Systems Running in a System Replication Setup

You can update your SAP HANA system with active system replication by updating the secondary and the primary system one after the other.

Prerequisites

System replication is configured and active between two SAP HANA systems.

Context

You must update your SAP HANA system running in a system replication setup by updating the secondary system first and then updating the primary system.

→ Remember

For system replication setups it is required that the secondary system has the same version as the primary system or a higher version. As such, the secondary system must always be updated before the primary system.

> **Note**
>
> Updating one system after the other results in some downtime. If you want to update your system with reduced downtime, see *Use SAP HANA System Replication for Near Zero Downtime Upgrades*.
>
> It is possible to reduce the time required to perform an update. For more information, see *Prepare an Update for Flexible System Downtime* .

Procedure

1. Upgrade the SAP HANA server software and all installed components on the secondary system.

   From your installation directory execute as root or as `<sid>adm`:

   ```
   ./hdblcm --action=update
   ```

2. With the secondary system online, use the SAP HANA lifecycle management tools to upgrade all the other components to the same revision as the server software.
3. Verify that system replication is active and that all services are in sync.

   You can check that the REPLICATION_STATUS column in M_SERVICE_REPLICATION has the value ACTIVE for all services.
4. Upgrade the SAP HANA server software and all installed components on the primary system.

   From your installation directory, execute as root or as `<sid>adm`:

   ```
   ./hdblcm --action=update
   ```

5. With the primary system online, use the SAP HANA lifecycle management tools to upgrade all other components to the same revision as the server software.
6. Verify that system replication is active and that all services are in sync.

Related Information

Use SAP HANA System Replication for Near Zero Downtime Upgrades [page 177]
Prepare an Update for Flexible System Downtime [page 154]
SAP Note 2407186
SAP Note 2599514

7.2.6 Use SAP HANA System Replication for Near Zero Downtime Upgrades

You can use SAP HANA system replication to upgrade your SAP HANA systems as the secondary system can run with a higher software version than the primary system.

Prerequisites

You configured a user in the local userstore under the SRTAKEOVER key. For more information, see *Configure a User Under the SRTAKEOVER Key*.

System replication is configured and active between two identical SAP HANA systems:

- The primary system is the production system.
- The secondary system will become the production system after the upgrade.

Context

With system replication active, you can first upgrade the secondary system to a new revision and have it take over in the role of primary system. The takeover is carried out in only a few minutes and committed transactions or data are not lost. You can then do an upgrade on the primary system, which is now in the role of secondary.

> **Note**
>
> It is possible to reduce the time required to perform an update. For more information, see *Prepare an Update for Flexible System Downtime* in the *SAP HANA Server Installation and Update Guide*.

The secondary system can be initially installed with the new software version or upgraded to the new software version when the replication has already been configured. After the secondary has been upgraded, all data has to be replicated to the secondary system (already having the new software version). When the secondary system is ACTIVE (all services have synced), a takeover has to be executed on the secondary system. This step makes the secondary system the production system running with the new software version.

If the installed system version on the primary is HANA 2.0 SPS 04 or greater then you are recommended to use the 'takeover with handshake' option to ensure a consistent handover. Using this option the primary continues to run but the writing of transactions on the primary system is suspended. The takeover is only executed when all redo log is available on the secondary system. See step 3 in the following procedure.

If you are upgrading from SAP HANA 1.0 to SAP HANA 2.0 note that system replication with SAP HANA 2.0 requires authentication for data and log shipping channels, this is done using the certificates in the system PKI SSFS store. You must therefore copy the system PKI SSFS key and the data file from the current primary system to the new to-be secondary system. Copy the files before registration when the secondary system is offline; the files can be found here:

```
/usr/sap/<SID>/SYS/global/security/rsecssfs/data/SSFS_<SID>.DAT
/usr/sap/<SID>/SYS/global/security/rsecssfs/key/SSFS_<SID>.KEY
```

For more information, see *SAP Note 2369981: Required configuration steps for authentication with HANA System Replication*.

In an Active/Active (read enabled) system replication setup, the version of the primary and the secondary systems must be identical. For the near zero downtime upgrade to work, the operation mode on the secondary system is automatically set to `logreplay`. Like this, the two systems can get back in sync before the takeover step. To establish again the Active/Active (read enabled) landscape at the end, the `logreplay_readaccess` operation mode must be explicitly specified during the former registration of the primary system as a new secondary system.

For more information about near zero downtime upgrades when using a multitarget system replication setup, see *Use Multitarget System Replication for Near Zero Downtime Upgrades*.

Procedure

1. Upgrade the secondary system's SAP HANA server software and all other components.

   From your installation directory execute as root:

   ```
   ./hdblcm --action=update
   ```

2. Verify that system replication is active and that all services are in sync.

   You can check that the column REPLICATION_STATUS in M_SERVICE_REPLICATION has the value ACTIVE for all services.

3. Depending on the version installed on the primary, perform a takeover by doing one of the following:

   - If the installed system version on the primary is HANA 2.0 SPS 04 or greater then you have the option to use the `--suspendPrimary` parameter for a 'takeover with handshake' which ensures that all redo logs are written to disk. In this case, execute the takeover as `<sid>adm` with the following command:

     ```
     hdbnsutil -sr_takeover --suspendPrimary
     ```

   - If the installed system version on the primary is less than HANA 2.0 SPS 04, then:
     - Stop the primary system.
     - Execute the takeover as `<sid>adm` with the following (default) command:

       ```
       hdbnsutil -sr_takeover
       ```

   You can then switch virtual IP addresses to the secondary system, and start using it productively.

4. If XS Advanced is being updated as well, update the XS Advanced applications.

   ```
   ./hdblcm --action=update
   ```

5. If the primary has not been stopped (takeover with handshake option), you can now stop the primary system.
6. Upgrade the original primary from the installation directory as root user using the 'nostart' option. This option is required because otherwise the primary has to be stopped again before it can be registered as the secondary:

   ```
   ./hdblcm --action=update --hdbupd_server_nostart
   ```

> **Note**
>
> For a fast synchronization of the sites – after registering again the original primary system – perform this failback within the time given by the `datashipping_snapshot_max_retention_time` parameter (default 300 minutes), otherwise, a full data shipping will be done. Furthermore, the optimized resync depends on the availability of the last snapshot.
>
> For more information about near zero downtime upgrades in multitier system replication, see *SAP Note 2386973*.

7. Register the original primary as secondary as `<sid>adm`.

```
hdbnsutil -sr_register --name=<secondary_alias>
 --remoteHost=<primary_host> --remoteInstance=<primary_systemnr>
 --replicationMode=[sync|syncmem|async] --operationMode=[delta_datashipping|
logreplay|logreplay_readaccess]
```

8. Start the original primary.

Related Information

Configure a User Under the SRTAKEOVER Key
Prepare an Update for Flexible System Downtime [page 154]
Updating the SAP HANA System [page 148]
Perform a Near-Zero Downtime Update [page 163]
Use Multitarget System Replication for Near Zero Downtime Upgrades
Deploy a Multi-Target Application with Zero-Downtime Maintenance
Takeover with Handshake
SAP Note 2369981
SAP Note 1984882
SAP Note 2386973
SAP Note 2494079
SAP Note 2407186
SAP Note 2300936

8 Managing SAP HANA System Components

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

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

- SAP HANA mandatory components
  - SAP HANA server
  - SAP HANA client
  - Local Secure Store (LSS)

  > **Note**
  >
  > As of SAP HANA 2.0 SPS 08, SAP HANA Local Secure Store will be installed and activated by default during an SAP HANA system installation and update. As a result, the LSS installation path (`lss_inst_path`), as well as the `<sid>crypt` user password (`lss_user_password`) and the LSS backup password (`lss_backup_password`) are required to be defined in this case.

  > **Restriction**
  >
  > The LSS is not supported by SAP HANA Dynamic Tiering. For SAP HANA systems configured with dynamic tiering (with LSS installed and activated), set the *Secure Store* type for them to SSFS (secure store in the file system).
  >
  > For details, see Local Secure Store (LSS) under *Security Administration* > *Data Encryption* in the *SAP HANA Dynamic Tiering: Administration Guide 2.0 SP04*.
  >
  > If LSS is installed and activated with an encryption_algorithm other than AES-256-CBC, migration back to SSFS is not possible.
- SAP HANA additional components
  - SAP HANA studio
  - Application Function Libraries (AFL, EML and the product-specific AFLs RME, RTL POS, SAL, SCA, SOP, TRP, VCH)
  - SAP liveCache applications (SAP LCA or LCAPPS-Plugin)
  - SAP HANA EPM-MDS (EPMMDS plug-in)
  - SAP HANA smart data access (SDA)
  - SAP HANA XS Advanced Runtime (For more information about installating XS advanced runtime, see Installing XS Advanced Runtime in Related Information.)
- SAP HANA options
  - SAP HANA dynamic tiering
  - SAP HANA streaming analytics
  - SAP HANA accelerator for SAP ASE

> **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*.

For more information about installing, updating, and uninstalling the SAP HANA mandatory components and SAP HANA additional components, see Related Information. For more information about installing, updating, and uninstalling the SAP HANA options, see SAP HANA option documentation in Related Information.

> **Caution**
>
> Be aware that you need additional licenses for SAP HANA options. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities* in Related Information.

Related Information

SAP Note 1858920
SAP Note 1891393
Installing or Updating SAP HANA Components [page 181]
SAP HANA Options in SAP Help Portal
Important Disclaimer for Features in SAP HANA [page 304]
Installing an SAP HANA System Including the XS Advanced Runtime [page 97]
Local Secure Store (LSS) [page 28]
lss_inst_path [page 264]

8.1 Installing or Updating SAP HANA Components

SAP HANA components can be installed or updated the following ways:

- From the installation medium, using
  - the graphical user interface,
  - the command-line interface.
- From the resident program, using
  - the graphical user interface,
  - the command-line interface.
- Using the Web user interface.

Related Information

Install or Update SAP HANA Components Using the Graphical User Interface [page 182]
Install or Update SAP HANA Components Using the Command-Line Interface [page 184]
Install or Update SAP HANA Components Using the Resident Program [page 186]
Install or Update SAP HANA Components Using the Web User Interface [page 189]
SAP Note 1858920
SAP Note 2293092
SAP Note 2014334

8.1.1 Install or Update SAP HANA Components Using the Graphical User Interface

You can install additional SAP HANA system components like the SAP HANA client, SAP HANA studio, and additional system components like Application Function Libraries (AFL, EML and the product-specific AFLs IBP, RTL, TRP, VCH, XRP), SAP liveCache applications (SAP LCA or LCAPPS plug-in), SAP HANA EPM-MDS (EPMMDS plug-in), XS advanced runtime applications, SAP HANA smart data access (SDA), or SAP HANA XS Advanced Runtime using the SAP HANA database lifecycle manager (HDBLCM) graphical user interface.

Prerequisites

The system component should have the same version as the SAP HANA database. Do one of the following:

- Patch the SAP HANA system component to a higher patch number within the same SP (revision).
- Update both the SAP HANA system component and the SAP HANA database to a higher SP (revision).

You cannot update the AFLs to a higher revision number unless you also update your SAP HANA database to the same revision number. The installation and update of XS advanced runtime components requires the installation/update of the XS advanced runtime.

- The `root` user is able to execute graphical applications.

Context

In order to install or update SAP HANA system components or additional components, you must start the SAP HANA database lifecycle manager (HDBLCM) and run an update. The update component menu offers the opportunity to update the components (if a more recent version is found) and to install additional components, which were not installed during the initial server installation. If you select the action *Install new system* instead of *Update existing system*, the SAP HANA database lifecycle manager (HDBLCM) also installs the SAP HANA server, by default. Therefore, if you would like to add only one additional system component to an existing system, it is necessary to select the action *Update existing system*, and select only that system component from the available component list.

Some AFL components can be updated without the need for a database restart and system downtime. A component supports this feature if its `manifest` file contains the entry

```
online-upgrade-plugin: 1
```

> **Note**
>
> Before updating such a component, verify the version information contained in the `required-components` entry inside the `manifest` file, for example:
>
> ```
> required-components: name="HDB"; vendor="sap.com";
> version="[1.00.110.00,1.00.110.1)"
> ```
>
> If the installed SAP HANA database matches the version requirements of the component, the component is updated and enabled without a database restart. In this example, any SAP HANA database including version `1.00.110.00` and excluding version `1.00.110.1` supports this feature. If the versions do not match, the component is updated, but the database must be updated to the required version and restarted before the component is enabled.

> **Note**
>
> The product-specific AFLs are released individually and are no longer released as part of SAP HANA AFL. Therefore, before updating AFL version SPS 07 to a current version, it is necessary to perform a migration. For more information, see SAP Note 2014334 in Related Information. You can update AFL version SPS 08 to a current version as described.

> **Note**
>
> Adding SAP liveCache applications (SAP LCA or LCAPPS-Plugin) is only supported for the integrated SAP liveCache, single-host scenario. The SAP LCA archive should be owned by the root user.

Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
IBM Power Systems	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`

> **Note**
>
> If you downloaded the components to a different directory, change to the directory where you unpacked the archive.

2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

   ```
   ./hdblcmgui
   ```

   The SAP HANA database lifecycle manager graphical user interface appears.

   > ⓘ Note
   >
   > To activate the local secure (LSS) store during installation, run `hdblcmgui` with the parameter `secure_store=localsecurestore`.

3. Select *Update existing system* from the activity options. Then select *Next*.
4. Select the components you would like to install or update as a comma-seperated list, then select *Next*.
5. Specify the SAP HANA system properties.

   You have the opportunity to specify a lesser-privileged database user if you have previously created one. For more information about creating a database user for update, see Related Information.
6. Review the summary, and select *Run* to finalize the configuration.

Related Information

SAP Note 2327295

8.1.2 Install or Update SAP HANA Components Using the Command-Line Interface

You can install additional SAP HANA system components like the SAP HANA client, SAP HANA studio, and additional system components like Application Function Libraries (AFL, EML and the product-specific AFLs IBP, RTL, TRP, VCH, XRP), SAP liveCache applications (SAP LCA or LCAPPS plug-in), SAP HANA EPM-MDS (EPMMDS plug-in), XS advanced runtime applications, SAP HANA smart data access (SDA), or SAP HANA XS Advanced Runtime using the SAP HANA database lifecycle manager (HDBLCM) command-line interface.

Prerequisites

The system component should have the same version as the SAP HANA database. Do one of the following:

- Patch the SAP HANA system component to a higher patch number within the same SP (revision).
- Update both the SAP HANA system component and the SAP HANA database to a higher SP (revision).

You cannot update the AFLs to a higher revision number unless you also update your SAP HANA database to the same revision number. The installation and update of XS advanced runtime components requires the installation/update of the XS advanced runtime.

Context

In order to install or update SAP HANA system components or additional components, you must start the SAP HANA database lifecycle manager (HDBLCM) and run an update. The update component menu offers the opportunity to update the components (if a more recent version is found) and to install additional components, which were not installed during the initial server installation. If you select the action *Install new system* instead of *Update existing system*, the SAP HANA database lifecycle manager (HDBLCM) also installs the SAP HANA server, by default. Therefore, if you would like to add only one additional system component to an existing system, it is necessary to select the action *Update existing system*, and select only that system component from the available component list.

Some AFL components can be updated without the need for a database restart and system downtime. A component supports this feature if its `manifest` file contains the entry

```
online-upgrade-plugin: 1
```

> **Note**
>
> Before updating such a component, verify the version information contained in the `required-components` entry inside the `manifest` file, for example:
>
> ```
> required-components: name="HDB"; vendor="sap.com";
> version="[1.00.110.00,1.00.110.1)"
> ```
>
> If the installed SAP HANA database matches the version requirements of the component, the component is updated and enabled without a database restart. In this example, any SAP HANA database including version `1.00.110.00` and excluding version `1.00.110.1` supports this feature. If the versions do not match, the component is updated, but the database must be updated to the required version and restarted before the component is enabled.

> **Note**
>
> The product-specific AFLs are released individually and are no longer released as part of SAP HANA AFL. Therefore, before updating AFL version SPS 07 to a current version, it is necessary to perform a migration. For more information, see SAP Note 2014334 in Related Information. You can update AFL version SPS 08 to a current version as described.

> **Note**
>
> Adding SAP liveCache applications (SAP LCA or LCAPPS-Plugin) is only supported for the integrated SAP liveCache, single-host scenario. The SAP LCA archive should be owned by the root user.

Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
IBM Power Systems	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`

> **Note**
>
> If you downloaded the components to a different directory, change to the directory where you unpacked the archive.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

   ```
   ./hdblcm
   ```

3. Select the index for the system to be updated, then select `Enter`.
4. Select the components you would like to install or update as a comma-seperated list, then select `Enter`.
5. Specify the SAP HANA system properties.

   You have the opportunity to specify a lesser-privileged database user if you have previously created one. For more information about creating a database user for update, see Related Information.

6. After specifying all system properties, review the summary, and select *y*.

Related Information

SAP Note 2327295

8.1.3 Install or Update SAP HANA Components Using the Resident Program

SAP HANA system components like the SAP HANA client, SAP HANA studio, HLM, Application Function Libraries (AFL), SAP liveCache applications (SAP LCA), XS advanced runtime, or SAP HANA smart data access (SDA) can be added to an SAP HANA system after installation from a local host using the SAP HANA lifecycle management tool hdblcm(gui).

Prerequisites

The system component should have the same version as the SAP HANA database. Do one of the following:

- Patch the SAP HANA system component to a higher patch number within the same SP (revision).

- Update both the SAP HANA system component and the SAP HANA database to a higher SP (revision).

You cannot update the AFLs to a higher revision number unless you also update your SAP HANA database to the same revision number. The installation and update of XS advanced runtime components requires the installation/update of the XS advanced runtime.

Context

In order to add SAP HANA system components, you must start the hdblcm(gui) installer and select update. The update component menu offers to update the server (if a more recent version is found) and also offers to install additional components, which were not installed during the initial server installation. If you select the action *Install new system* instead of *Update existing system*, hdblcm(gui) also installs the SAP HANA server, by default. Therefore, if you would like to add only one additional system component, it is necessary to select the action *Update existing system*, and select only that system component from the available component list.

Some AFL components can be updated without the need for a database restart and system downtime. A component supports this feature if its `manifest` file contains the entry

```
online-upgrade-plugin: 1
```

> **Note**
>
> Before updating such a component, verify the version information contained in the `required-components` entry inside the `manifest` file, for example:
>
> ```
> required-components: name="HDB"; vendor="sap.com";
> version="[1.00.110.00,1.00.110.1)"
> ```
>
> If the installed SAP HANA database matches the version requirements of the component, the component is updated and enabled without a database restart. In this example, any SAP HANA database including version `1.00.110.00` and excluding version `1.00.110.1` supports this feature. If the versions do not match, the component is updated, but the database must be updated to the required version and restarted before the component is enabled.

To install and update software components in SAP HANA XS Advanced, the `xs install` command is available in the XS Advanced command-line interface (CLI). For more information, see *Installing and Updating Software Components in SAP HANA XS Advanced Model* in the SAP HANA Administration Guide.

> **Note**
>
> The product-specific AFLs are released individually and are no longer released as part of SAP HANA AFL. Therefore, before updating AFL version SPS 07 to a current version, it is necessary to perform a migration. For more information, see SAP Note 2014334 in Related Information. You can update AFL version SPS 08 to a current version as described.

Procedure

1. Prepare the component location.

   The SAP HANA database lifecycle manager detects software components in nearby standard directories, but not in unique or distant directories. Therefore, you should make note of where the software components for installation or update are located, so that you can add the component location in the graphical user interface field, or specify the component location as a call option in the command-line interface using one of the following parameters:

   - `component_dirs`
   - `component_medium`
   - `component_root`

   For more information about these parameters, see Related Information.

2. Change to the SAP HANA resident HDBLCM directory:

   ```
   cd <sapmnt>/<SID>/hdblcm
   ```

   By default, `<sapmnt>` is `/hana/shared`.

3. Start the SAP HANA database lifecycle manager using the graphical user interface or the command-line interface.

Option	Description
Graphical User Interface	1. Start the SAP HANA HDBLCM graphical user interface: `./hdblcmgui` 2. Select the *Install or Update Additional Components* from the activity list. Then select *Next*. 3. Select a detected software component or add a software component location by selecting *Add Component Location...*. Then select *Next*. 4. Select the components you would like to install or update, then select *Next*. 5. Specify the SAP HANA system properties. You have the opportunity to specify a lesser-privileged database user if you have previously created one. For more information about creating a database user for update, see Related Information. 6. If you are installing or updating XS advanced runtime components, you can select which of the detected components are installed. 7. Review the summary and select *Upgrade*.
Command-Line Interface	1. Start the SAP HANA HDBLCM command-line interface: `./hdblcm` If the component location is not detected by the SAP HANA HDBLCM, rerun the program with one of the following call options specified `--component[_dirs|_medium|_root]=<component location>` 2. Select the index for the `update_components`, then select `Enter`. 3. Select the components you would like to install or update as a comma-seperated list of indexes, then select `Enter`. 4. Specify the SAP HANA system properties. You have the opportunity to specify a lesser-privileged database user if you have previously created one. For more information about creating a database user for update, see Related Information.

Option	Description
	5. If you are installing or updating XS advanced runtime components, you can select which of the detected components are installed. 6. Review the summary and select *y*.

Results

A system component has been added to the SAP HANA system. The component list has been updated. A log has been produced.

Related Information

Installing and Updating Products and Software Components in SAP HANA XS Advanced Model
component_dirs [page 244]
component_medium [page 245]
component_root [page 245]
Create a Lesser-Privileged Database User for Update [page 157]
SAP Note 2327295

8.1.4 Install or Update SAP HANA Components Using the Web User Interface

You can install additional SAP HANA system components like the SAP HANA client, SAP HANA studio, and additional system components like Application Function Libraries (AFL, EML and the product-specific AFLs IBP, RTL, TRP, VCH, XRP), SAP liveCache applications (SAP LCA or LCAPPS plug-in), SAP HANA EPM-MDS (EPMMDS plug-in), XS advanced runtime applications, SAP HANA smart data access (SDA), or SAP HANA XS Advanced Runtime using the SAP HANA database lifecycle manager (HDBLCM) Web user interface.

Prerequisites

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
  Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
  - Microsoft Windows

    - Internet Explorer - Version 9 or higher
      If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing *Tools* > *Compatibility View Settings*.
    - Microsoft Edge
    - Mozilla Firefox - Latest version and Extended Support Release
    - Google Chrome - Latest version
  - SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
  - Mac OS - Safari 5.1 or higher

> **Note**
>
> For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

The system component should have the same version as the SAP HANA database. Do one of the following:

- Patch the SAP HANA system component to a higher patch number within the same SP (revision).
- Update both the SAP HANA system component and the SAP HANA database to a higher SP (revision).

You cannot update the AFLs to a higher revision number unless you also update your SAP HANA database to the same revision number. The installation and update of XS advanced runtime components requires the installation/update of the XS advanced runtime.

Context

Some AFL components can be updated without the need for a database restart and system downtime. A component supports this feature if its `manifest` file contains the entry

```
online-upgrade-plugin: 1
```

> **Note**
>
> Before updating such a component, verify the version information contained in the `required-components` entry inside the `manifest` file, for example:
>
> ```
> required-components: name="HDB"; vendor="sap.com";
> version="[1.00.110.00,1.00.110.1)"
> ```
>
> If the installed SAP HANA database matches the version requirements of the component, the component is updated and enabled without a database restart. In this example, any SAP HANA database including version `1.00.110.00` and excluding version `1.00.110.1` supports this feature. If the versions do not match, the component is updated, but the database must be updated to the required version and restarted before the component is enabled.

> **Note**
>
> The product-specific AFLs are released individually and are no longer released as part of SAP HANA AFL. Therefore, before updating AFL version SPS 07 to a current version, it is necessary to perform a migration. For more information, see SAP Note 2014334 in Related Information. You can update AFL version SPS 08 to a current version as described.

Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: https://`<hostname>`:1129/lmsl/HDBLCM/`<SID>`/index.html **Note** The URL is case sensitive. Make sure you enter upper and lower case letters correctly.
SAP HANA cockpit	1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. `https://<host_FQDN>:<port>` **Note** FQDN = fully qualified domain name 2. Drill down on the name of the system from *My Resources* or from a group. 3. The links in *Platform Lifecycle Management* each launch additional functionality, giving you expanded capabilities for managing the resource.

2. Select the *Install or Update Additional Components* tile.
3. Select *Add Software Locations...* to add SAP HANA components for installation or update. Then select *Next*.
4. Select the components for installation or update. Then select *Next*.

   Various parameters can be set in the *Advanced Parameters Configuration* dialog. To access the *Advanced Parameters Configuration* dialog, click on the gear icon in the footer bar of the SAP HANA HDBLCM Web user interface.
5. Specify additional hosts or roles to be added. Then select *Next*.
6. Specify the SAP HANA authorization information.

   Depending on the component selection, if asked for database user, you have the opportunity to specify a lesser-privileged database user if you have previously created one.
7. After specifying all system properties, review the summary, and select *Update*.

Related Information

SAP Note 2327295

8.1.5 Upload and Extract SAP HANA Components Using the Web User Interface

You can upload and extract SAP HANA component archives that were downloaded from the SAP Support Portal for installation or update using the SAP HANA database lifecycle manager (HDBLCM) Web user interface.

Prerequisites

- The SAP HANA database server is up and running.

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
  Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
  - Microsoft Windows
    - Internet Explorer - Version 9 or higher
      If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing *Tools* > *Compatibility View Settings*.
    - Microsoft Edge
    - Mozilla Firefox - Latest version and Extended Support Release
    - Google Chrome - Latest version
  - SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
  - Mac OS - Safari 5.1 or higher

  > **Note**
  >
  > For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

The system component should have the same version as the SAP HANA database. Do one of the following:

- Patch the SAP HANA system component to a higher patch number within the same SP (revision).
- Update both the SAP HANA system component and the SAP HANA database to a higher SP (revision).

You cannot update the AFLs to a higher revision number unless you also update your SAP HANA database to the same revision number. The installation and update of XS advanced runtime components requires the installation/update of the XS advanced runtime.

Context

Some AFL components can be updated without the need for a database restart and system downtime. A component supports this feature if its `manifest` file contains the entry

```
online-upgrade-plugin: 1
```

> **Note**
>
> Before updating such a component, verify the version information contained in the `required-components` entry inside the `manifest` file, for example:
>
> ```
> required-components: name="HDB"; vendor="sap.com";
> version="[1.00.110.00,1.00.110.1)"
> ```
>
> If the installed SAP HANA database matches the version requirements of the component, the component is updated and enabled without a database restart. In this example, any SAP HANA database including version `1.00.110.00` and excluding version `1.00.110.1` supports this feature. If the versions do not match, the component is updated, but the database must be updated to the required version and restarted before the component is enabled.

> **Note**
>
> The product-specific AFLs are released individually and are no longer released as part of SAP HANA AFL. Therefore, before updating AFL version SPS 07 to a current version, it is necessary to perform a migration. For more information, see SAP Note 2014334 in Related Information. You can update AFL version SPS 08 to a current version as described.

Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: https://`<hostname>`:1129/lmsl/HDBLCM/`<SID>`/index.html **Note** The URL is case sensitive. Make sure you enter upper and lower case letters correctly.
SAP HANA cockpit	1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. `https://<host_FQDN>:<port>` **Note** FQDN = fully qualified domain name 2. Drill down on the name of the system from *My Resources* or from a group.

Option	Description
	3. The links in *Platform Lifecycle Management* each launch additional functionality, giving you expanded capabilities for managing the resource.

2. Select the *Upload/Extract Components* tile.
3. Select the upload method.

Option	Description
The Archives are Accessible from the SAP HANA Host	Use this option if the archives are located on a file system accessible from the SAP HANA host. Specify the directory which contains the component archives under *Location of SAP HANA Component Archives*. Then select *Next*.
Upload Archives to the SAP HANA Host	Use this option if the archives are accessible only from your local machine. Select one or more component archives that you want to upload to the SAP HANA host. Then select *Upload*.

4. Specify an empty target directory to extract the software component archives to under *Temporary Extract Directory*. Then select *Next*.
5. After specifying all system properties, review the summary, and select *Extract*.

8.1.6 Update XS Advanced Components

If an XS Advanced component needs to be updated (for example, due to a security vulnerability), the patch process depends on the category of the affected component.

The following list covers the update process for each category:

XS Advanced System Components

The XS Advanced core installation includes the following system components:

- XS Advanced Runtime
- UAA
- HANA Service Broker
- Instance Manager
- Audit-Log Service
- Product Installer
- Deploy Service
- File System Service

If one of these components needs to be updated, a new XS Advanced release has to be installed using the SAP HANA lifecycle tool `hdblcm`. Users with the required S-User ID can download the latest version of the XS

advanced component in the package `SAP EXTENDED APP SERVICES 1` from the following location in SAP Support Portal:

Software Downloads > *SUPPORT PACKAGES & PATCHES* > *By Alphabetical Index (A-Z)* > *H* > *SAP HANA PLATFORM EDITION* > *SAP HANA PLATFORM EDITION 2.0* > *SAP EXTENDED APP SERVICES 1*.

Download the latest installation archive and extract it to a target directory. Next, change to the directory `<sapmnt>/<SID>/hdblcm` (where `<sapmnt>` by default is `/hana/shared`) and execute the resident `hdblcm` command, as follows:

Output Code

```
./hdblcm --component_dirs <path to extracted installation archive>
```

For more information, see *Install or Update SAP HANA Components Using the Resident Program* and *SAP Note 2347931 - SAP HANA extended application services, advanced model components versioning* in *Related Information*.

XS Advanced Additional Components

To update XS Advanced System Components and any **additional** components, use the XS Advanced Installation Media. For more information about where to find and download the XS Advanced Installation Media, see SAP Note 2711421 (Installing SAP HANA Extended Application Services, advanced model using the XS Advanced installation media) and *Related Information* below.

Build Packs

Each new XS Advanced release includes the latest versions of the default build packs (for example, "Java Build Pack" and "Node.js Build Pack"). Updating XS Advanced will also update the default build packs.

If you want to update any **custom** build packs, use the `xs update-buildpack` command, as illustrated in the following example:

```
xs update-buildpack <BUILDPACK> [-p <PATH>] [-i <POSITION>] [--enable|--disable]
[--lock|--unlock]
```

→ Tip

`<PATH>` is the location of the updated package.

Updating a build pack does not imply an automatic change for any application droplets that have been built using a previous version of the updated build pack. To update the application droplet, you must trigger the restage and restart operations manually, for example using the `xs restage` and `xs restart` commands, as illustrated in the following example:

```
xs restage <APP>
xs restart <APP>
```

Application Run Times

XS Advanced comes with a set of default application run-time environments. Each new XS Advanced release comes with the latest versions of the default application run-time environments. Updating XS Advanced will also update the default application run-time environments.

To update a custom application run time environment, you must create and upload a new version of the application run time by using the `xs create-runtime` command, as illustrated in the following example:

```
xs create-runtime -p <PATH> [--inactive] [<DESCRIPTION>]
```

When you deploy new applications, the newly created run time environment is used. To enable any already staged applications to use the new run-time environment, you must restage the affected applications. You can then delete the old application run time using the `xs delete-runtime` command, as illustrated in the following example:

```
xs delete-runtime -i <ID> [-f]
```

The command `xs runtimes` displays a list of all available XS advanced run-times including information about the individual run-time ID.

> ⓘ Note
>
> -
> - Application run time environments might be pinned to one or more specific applications. The `xs delete-runtime` command removes all pinning settings for the deleted application run time. You can use the command `xs pinned-runtimes` to display a list of all pinned run time environments.

Service Brokers

Each new XS Advanced release includes the latest release of the system service brokers. Updating XS Advanced also updates the system service brokers.

To update custom service brokers manually, deploy the updated service broker to your landscape and use the command `xs update-service-broker` to register the changes for the updated service broker with the XS Controller, as illustrated in the following example:

```
xs update-service-broker <SERVICE_BROKER> <USERNAME> <PASSWORD> <URL>
```

> ⓘ Note
>
> Some service brokers are integrated in XS Advanced applications. If the service broker you update is integrated in an application, you must deploy a new version of the application.

XS Advanced Applications

If you need to update an XS Advanced application (for example, due to a security vulnerability in one of the application's components), deploy a new version of the application with the `xs install` command, as illustrated in the following example:

```
xs install <installation-archive>
```

For more information about how to download new versions of additional SAP XS Advanced applications, see SAP Note 2347931 (SAP HANA extended application services, advanced model components versioning) or *Related Information* below.

For more information about the individual `xs` commands, see the *SAP HANA Developer Guide for XS Advanced (XS CLI)* or the *SAP HANA Administration Guide*.

Related Information

SAP Note 2711421
SAP Note 2347931
SAP HANA Administration Guide
SAP HANA Developer Guide for XS Advanced Model (SAP Web IDE)
Install or Update SAP HANA Components Using the Resident Program [page 186]

8.1.7 Activate the Local Secure Store (LSS)

After the installation of Local Secure Store (LSS) has finished successfully and the SAP HANA secure store remains unchanged, LSS must be enabled from the command-line before it can be used.

Prerequisites

You are logged on with the required root user or system administrator user `<sid>adm` credentials.

> **⚠ Restriction**
>
> The local secure store (LSS) is not supported by SAP HANA Dynamic Tiering. For SAP HANA systems configured with dynamic tiering (with LSS installed and activated), set the *Secure Store* type for them to SSFS (secure store in the file system).
>
> For details, see Local Secure Store (LSS) under *Security Administration* > *Data Encryption* in the *SAP HANA Dynamic Tiering: Administration Guide 2.0 SP04*.

Context

The following procedure describes how to activate the LSS after installing or updating an SAP HANA system with the LSS component and keeping SSFS (Secure Store in the File System) as the SAP HANA secure store. It will also migrate SSFS to LSS.

> **Note**
>
> LSS is installed and activated by default during an SAP HANA system installation and update.

Procedure

1. Stop the SAP HANA system:

   ```
   /usr/sap/hostctrl/exe/sapcontrol -nr <instance number> -function StopSystem
   ```

2. Activate the LSS by executing the following command on the active master nameserver host (the host that is set as `active_master` in the system layer `nameserver.ini` configuration file):

   ```
   hdbnsutil -migrateSecureStore --target=LSS
   ```

3. Start the SAP HANA system:

   ```
   /usr/sap/hostctrl/exe/sapcontrol -nr <instance number> -function StartSystem
   ```

Related Information

Local Secure Store (LSS) [page 28]
Start and Stop the SAP HANA System [page 210]
secure_store [page 274]

8.2 Uninstalling SAP HANA Components

SAP HANA components - including system components and additional components - can be uninstalled the following ways:

- From the resident program, using
  - the graphical user interface,
  - the command-line interface.
- Using the Web user interface.

Related Information

Uninstall SAP HANA Components Using the Graphical User Interface [page 199]
Uninstall SAP HANA Components Using the Command-Line Interface [page 200]
Uninstall an SAP HANA Component on a System Missing the SAP HANA Resident Program [page 202]
Uninstall SAP HANA Components Using the Web User Interface [page 203]
SAP Note 1858920

8.2.1 Uninstall SAP HANA Components Using the Graphical User Interface

SAP HANA system components and additional system components can be removed from an SAP HANA system after installation using the SAP HANA database lifecycle manager (HDBLCM) graphical user interface.

Prerequisites

- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running. Otherwise, inconsistencies in the configuration occur.
- The user is able to execute graphical applications.

Context

> **Note**
>
> Using the SAP HANA database lifecycle manager, it is possible to remove the SAP HANA server, SAP HANA client, SAP HANA studio, HLM, Application Function Libraries (AFL, EML and the product-specific AFLs IBP, RTL, TRP, VCH, XRP), SAP liveCache applications (SAP LCA or LCAPPS plug-in), SAP HANA EPM-MDS (EPMMDS plug-in), Local Secure Store (LSS), XS advanced runtime applications, or SAP HANA smart data access (SDA), but it is not possible to remove the SAP host agent or the Solution Manager Diagnostics (SMD) agent. If you need to uninstall the SMD agent, see SAP Note 1858920 in Related Information.

> **Note**
>
> If you want to uninstall a component that requires specific host roles, you must first remove all related host roles or hosts which have these host roles assigned. For more information, see the removing hosts and removing host roles sections in the *SAP HANA Administration Guide*.

During the uninstallation of the Local Secure Store (LSS) the secure store is automatically migrated to SSFS. Uninstallation of Local Secure Store (LSS) is not possible if it has been installed with an encryption_algorithm other than AES-256-CBC.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

   ```
   cd <sapmnt>/<SID>/hdblcm
   ```

   By default, `<sapmnt>` is `/hana/shared`.
2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

   ```
   ./hdblcmgui
   ```

   The SAP HANA database lifecycle manager graphical user interface appears.
3. Select *Uninstall SAP HANA Components* from the activity options. Then select *Next*.
4. Select *Uninstall separate components*, and then choose the components to be uninstalled. Then select *Next*.
5. Review the summary, and select *Run* to finalize the configuration.

Results

The selected components are uninstalled. A log file is available.

Related Information

SAP Note 1858920
Removing Hosts from an SAP HANA System
Removing Host Roles

8.2.2 Uninstall SAP HANA Components Using the Command-Line Interface

SAP HANA system components and additional system components can be removed from an SAP HANA system after installation using the SAP HANA database lifecycle manager (HDBLCM) command-line interface.

Prerequisites

- The SAP HANA system has been installed with the SAP HANA database lifecycle manager (HDBLCM).
- The SAP HANA database server is up and running. Otherwise, inconsistencies in the configuration occur.

Context

> **Note**
>
> Using the SAP HANA database lifecycle manager, it is possible to remove the SAP HANA server, SAP HANA client, SAP HANA studio, HLM, Application Function Libraries (AFL, EML and the product-specific AFLs IBP, RTL, TRP, VCH, XRP), SAP liveCache applications (SAP LCA or LCAPPS plug-in), SAP HANA EPM-MDS (EPMMDS plug-in), Local Secure Store (LSS), XS advanced runtime applications, or SAP HANA smart data access (SDA), but it is not possible to remove the SAP host agent or the Solution Manager Diagnostics (SMD) agent. If you need to uninstall the SMD agent, see SAP Note 1858920 in Related Information.

> **Note**
>
> If you want to uninstall a component that requires specific host roles, you must first remove all related host roles or hosts which have these host roles assigned. For more information, see the removing hosts and removing host roles sections in the *SAP HANA Administration Guide*.

During the uninstallation of the Local Secure Store (LSS) the secure store is automatically migrated to SSFS. Uninstallation of Local Secure Store (LSS) is not possible if it has been installed with an encryption_algorithm other than AES-256-CBC.

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

   ```
   cd <sapmnt>/<SID>/hdblcm
   ```

   By default, `<sapmnt>` is `/hana/shared`.
2. Start the SAP HANA database lifecycle manager interactively in the command line:

   ```
   ./hdblcm
   ```
3. Select the index for *uninstall*, then select `Enter`.
4. Select the components to be uninstalled as a comma-separated list of indexes. Then select *Enter*.
5. Review the summary, and select *y* to finalize the configuration.

Results

The selected components are uninstalled. A log file is available.

Related Information

SAP Note 1858920

Removing Hosts from an SAP HANA System
Removing Host Roles

8.2.3 Uninstall an SAP HANA Component on a System Missing the SAP HANA Resident Program

If you would like to uninstall SAP HANA components from a system, which has been installed with the SAP HANA platform lifecycle management tool `hdbinst`, and has been updated with `hdbupd`, you cannot use the typical uninstallation procedures with the SAP HANA database lifecycle manager (HDBLCM). The reason for this is that the SAP HANA resident HDBLCM is missing from the system.

Prerequisites

- You are logged on to the host where the server software is installed.
- You are logged on as the root user.

Context

To uninstall an SAP HANA component on a system missing the SAP HANA resident HDBLCM:

Install the SAP HANA resident HDBLCM, then perform component uninstallation as usual using the newly available SAP HANA resident HDBLCM.

To install the SAP HANA resident HDBLCM, follow the procedure below.

Start the SAP HANA database lifecycle manager from an SAP HANA server installation kit, which has the same version as the installed SAP HANA database, with the following command:

```
./hdblcm --action=update --components=hdblcm
```

Then uninstall using one of the typical uninstallation procedures in *Uninstalling SAP HANA Components*.

During the uninstallation of the Local Secure Store (LSS) the secure store is automatically migrated to SSFS. Uninstallation of Local Secure Store (LSS) is not possible if it has been installed with an encryption_algorithm other than AES-256-CBC.

Procedure

1. Change to the following directory on the installation medium:

Option	Description
Intel-Based Hardware Platforms	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64`
IBM Power Systems	`cd <installation medium>/DATA_UNITS/HDB_LCM_LINUX_PPC64`

> **Note**
>
> If you downloaded the components to a different directory, change to the directory where you unpacked the archive.

2. Start the SAP HANA database lifecycle manager interactively in the command line:

```
./hdblcm --action=update --components=hdblcm
```

3. Optionally, specify additional components locations.
4. Select the index for the system you want to update, then select `Enter`.
5. Review the summary, and select *y* to finalize the configuration.

8.2.4 Uninstall SAP HANA Components Using the Web User Interface

SAP HANA system components and additional system components can be removed from an SAP HANA system after installation using the SAP HANA database lifecycle manager (HDBLCM) Web user interface.

Prerequisites

You should verify that the following prerequisites are fulfilled before trying to access the SAP HANA database lifecycle manager from a Web browser.

- The communication port 1129 is open.
  Port 1129 is required for the SSL communication with the SAP Host Agent in a standalone browser via HTTPS.
- The following Web browser requirements are fulfilled:
  - Microsoft Windows
    - Internet Explorer - Version 9 or higher
      If you are running Internet Explorer version 9, make sure that your browser is not running in compatibility mode with your SAP HANA host. You can check this in your browser by choosing *Tools* > *Compatibility View Settings*.
    - Microsoft Edge
    - Mozilla Firefox - Latest version and Extended Support Release

    - Google Chrome - Latest version
  - SUSE Linux - Mozilla Firefox with XULRunner 10.0.4 ESR
  - Mac OS - Safari 5.1 or higher

> **Note**
>
> For more information about supported Web browsers for the SAP HANA database lifecycle manager Web interface, see the browser support for `sap.m` library in the *SAPUI5 Developer Guide*.

- You are logged on as the system administrator user `<sid>adm`.
- The `<sid>adm` user has read and execute permissions for the directory that contains the installation medium.

The system component should have the same version as the SAP HANA database. Do one of the following:

- Patch the SAP HANA system component to a higher patch number within the same SP (revision).
- Update both the SAP HANA system component and the SAP HANA database to a higher SP (revision).

You cannot update the AFLs to a higher revision number unless you also update your SAP HANA database to the same revision number. The installation and update of XS advanced runtime components requires the installation/update of the XS advanced runtime.

> **Note**
>
> If you want to uninstall a component that requires specific host roles, you must first remove all related host roles or hosts which have these host roles assigned. For more information, see the removing hosts and removing host roles sections in the *SAP HANA Administration Guide*.

Context

> **Note**
>
> The product-specific AFLs are released individually and are no longer released as part of SAP HANA AFL. Therefore, before updating AFL version SPS 07 to a current version, it is necessary to perform a migration. For more information, see SAP Note 2014334 in Related Information. You can update AFL version SPS 08 to a current version as described.

During the uninstallation of the Local Secure Store (LSS) the secure store is automatically migrated to SSFS. Uninstallation of Local Secure Store (LSS) is not possible if it has been installed with an encryption_algorithm other than AES-256-CBC.

Procedure

1. Access the SAP HANA HDBLCM Web user interface.

Option	Description
Web browser	Enter the SAP HANA database lifecycle manager (HDBLCM) URL in an HTML5-enabled browser: https://<hostname>:1129/lmsl/HDBLCM/<SID>/index.html **Note** The URL is case sensitive. Make sure you enter upper and lower case letters correctly.
SAP HANA cockpit	1. Enter the URL of the SAP HANA cockpit administration and monitoring console in your browser. `https://<host_FQDN>:<port>` **Note** FQDN = fully qualified domain name 2. Drill down on the name of the system from *My Resources* or from a group. 3. The links in *Platform Lifecycle Management* each launch additional functionality, giving you expanded capabilities for managing the resource.

2. Select the *Uninstall Components* tile.
3. Select the components for uninstallation. Then select *Next*.
4. Specify the SAP HANA authorization information.
5. Select *Uninstall*.

9 Uninstalling the SAP HANA System

If required, you can uninstall the previously installed SAP HANA components by running the SAP HANA database lifecycle manager (HDBLCM) from the SAP HANA resident HDBLCM directory in the graphical user interface or command-line interface.

Related Information

Uninstall the SAP HANA System Using the Graphical User Interface [page 206]
Uninstall the SAP HANA System Using the Command-Line Interface [page 207]

9.1 Uninstall the SAP HANA System Using the Graphical User Interface

You can uninstall an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) graphical user interface.

Prerequisites

You are logged in as root user.

Context

> ⓘ Note
>
> Using the SAP HANA database lifecycle manager, it is possible to remove the SAP HANA server, SAP HANA client, SAP HANA studio, HLM, Application Function Libraries (AFL, EML and the product-specific AFLs IBP, RTL, TRP, VCH, XRP), SAP liveCache applications (SAP LCA or LCAPPS plug-in), SAP HANA EPM-MDS (EPMMDS plug-in), Local Secure Store (LSS), XS advanced runtime applications, or SAP HANA smart data access (SDA), but it is not possible to remove the SAP host agent or the Solution Manager Diagnostics (SMD) agent. If you need to uninstall the SMD agent, see SAP Note 1858920 in Related Information.

> ⚠ Caution
>
> Uninstalling the SAP HANA system removes all data volumes and log volumes. It is a permanent action that cannot be undone!

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

   ```
   cd <sapmnt>/<SID>/hdblcm
   ```

   By default, `<sapmnt>` is `/hana/shared`.
2. Start the SAP HANA database lifecycle manager interactively in the graphical user interface:

   ```
   ./hdblcmgui
   ```

   The SAP HANA database lifecycle manager graphical user interface appears.
3. Select *Uninstall SAP HANA Database Components* from the activity options. Then select *Next*.
4. Select *Uninstall SAP HANA Database version* `<version number>` *and all other components*. Then select *Next*.
5. Review the summary, and select *Uninstall* to finalize the configuration.

Results

The selected components are uninstalled. A log file is available.

Related Information

SAP Note 1858920

9.2 Uninstall the SAP HANA System Using the Command-Line Interface

You can uninstall an SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM) command-line interface.

Prerequisites

You are logged in as root user.

Context

> **Note**
>
> Using the SAP HANA database lifecycle manager, it is possible to remove the SAP HANA server, SAP HANA client, SAP HANA studio, HLM, Application Function Libraries (AFL, EML and the product-specific AFLs IBP, RTL, TRP, VCH, XRP), SAP liveCache applications (SAP LCA or LCAPPS plug-in), SAP HANA EPM-MDS (EPMMDS plug-in), Local Secure Store (LSS), XS advanced runtime applications, or SAP HANA smart data access (SDA), but it is not possible to remove the SAP host agent or the Solution Manager Diagnostics (SMD) agent. If you need to uninstall the SMD agent, see SAP Note 1858920 in Related Information.

> **Caution**
>
> Uninstalling the SAP HANA system removes all data volumes and log volumes. It is a permanent action that cannot be undone!

Procedure

1. Change to the SAP HANA resident HDBLCM directory:

   ```
   cd <sapmnt>/<SID>/hdblcm
   ```

   By default, `<sapmnt>` is `/hana/shared`.
2. Start the SAP HANA database lifecycle manager interactively in the command line:

   ```
   ./hdblcm
   ```
3. Select the index for *uninstall*, then select `Enter`.
4. Select the index for *all*. Then select `Enter`.
5. Review the summary, and select *y* to finalize the configuration.

Results

The selected components are uninstalled. A log file is available.

Related Information

SAP Note 1858920

10 Managing the SAP HANA System After Installation

After the installation has finished, it is recommended to perform the following tasks:

- Perform a system backup
  We strongly recommend that you perform an initial backup of your system once you have finished the installation. For more details, see the system backup information in the *SAP HANA Administration Guide*.
- Change the passwords
  If you are receiving an newly installed SAP HANA platform from a hardware provider, it is recommended to update the passwords so they comply with your security guidelines. For more information, see the *SAP HANA Security Guide*.
- Configure encryption on the SAP HANA server:
  - Change the SSFS master keys
    SAP HANA secures content in two secure stores in the file system (SSFS): the instance SSFS and the system PKI SSFS. Unique master keys are generated for the instance SSFS, if used, and the system PKI SSFS during installation or update. However, if you received your system pre-installed from a hardware or hosting partner, we recommend that you change them immediately after handover to ensure that they are not known outside of your organization. You can also change the master keys any time later.
  - Change and back up the encryption root keys used for SAP HANA data-at-rest encryption and application data encryption services.

    > ⚠ Caution
    >
    > Store the instance SSFS root key files in a safe location. If these files are lost, it may not be possible to recover the database. The default path of the key files is `/usr/sap/<sid>/SYS/global/hdb/security/ssfs`.

  - Enable data-at-rest encryption if required.

  For more information about sever-side data encryption services, as well as and how to configure and manage them, see the SAP HANA Security Guide and the SAP HANA Administration Guide.
- Finalize your customization
  Use the SAP HANA lifecycle management tools to adapt the existing configuration, if necessary. For more information, see the *SAP HANA Administration Guide*.

The following sections in this chapter are optional tasks that can be performed as part of installation management.

Related Information

Server-Side Data Encryption Services
Encryption Configuration
SAP Note 2600030

SAP Note 2400024

10.1 Start and Stop the SAP HANA System

After the installation has finished successfully, the SAP HANA system is up and running. So you do not need to start the SAP HANA system.

Context

However, if required, you can start and stop the SAP HANA system from the command line in one of the following ways:

Procedure

- By using the `sapcontrol` program:
  a. Log on to the SAP HANA system host as user `<sid>`adm or as a user with root permissions.
  b. Execute one of the following commands:
     - Start the SAP HANA system by entering the following command:

       ```
       /usr/sap/hostctrl/exe/sapcontrol -nr <instance number> -function
       StartSystem
       ```

     - Stop the SAP HANA system by entering the following command:

       ```
       /usr/sap/hostctrl/exe/sapcontrol -nr <instance number> -function
       StopSystem
       ```

- By using the `HDB` program:

  > ⚠ Caution
  >
  > The HDB program performs operations **only** on the local host. Its usage is discouraged for distributed SAP HANA systems. For more information, see Related Information.

  a. Log on to the SAP HANA system host as user `<sid>`adm.
  b. Execute one of the following commands:
     - Start the SAP HANA system by entering the following command:

       ```
       /usr/sap/<SID>/HDB<instance number>/HDB start
       ```

       Example:
       `/usr/sap/KB1/HDB26/HDB start`

- Stop the SAP HANA system by entering the following command:

  ```
  /usr/sap/<SID>/HDB<instance number>/HDB stop
  ```

  Example:
  `/usr/sap/KB1/HDB26/HDB stop`

  > **Note**
  >
  > The SAP HANA database does not start automatically by default when the SAP HANA system is started. But you can enable an automated start of this kind. For more information, see Related Information.

Related Information

Starting and Stopping SAP HANA Systems
autostart [page 239]

10.2 Display the Process List

It is possible to display the SAP HANA system processes from the command line.

Prerequisites

You are logged on with the required root user or system administrator user `<sid>adm` credentials.

Procedure

Display the SAP HANA system processes by running the following from the command line:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance number> -function GetProcessList
```

You can also display the SAP HANA system processes using the SAP Microsoft Management Console (SAP MMC) from a Microsoft Windows PC.

Results

Example: Displaying the Process List

```
/usr/sap/hostctrl/exe/sapcontrol -nr 00 -function GetProcessList
```

```
09.07.2015 14:09:20
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2015 07 06 13:38:00, 72:31:20, 1195
hdbnameserver, HDB Nameserver, GREEN, Running, 2015 07 06 13:38:03, 72:31:17,
1213
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2015 07 06 13:38:18,
72:31:02, 1279
hdbindexserver, HDB Indexserver, GREEN, Running, 2015 07 06 13:38:26, 72:30:54,
1317
hdbxsengine, HDB XSEngine, GREEN, Running, 2015 07 06 13:38:26, 72:30:54, 1320
hdbcompileserver, HDB Compileserver, GREEN, Running, 2015 07 06 13:38:18,
72:31:02, 1282
hdbwebdispatcher, HDB Web Dispatcher, GREEN, Running, 2015 07 06 13:39:10,
72:30:10, 1540
```

11 Tutorials

The following tutorials are a sampling of general use cases, which illustrate the functionality and versatility of the SAP HANA database lifecycle manager.

11.1 Tutorial: Automating Installation

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

Introduction to Installation Automation

With the SAP HANA database lifecycle manager (HDBLCM), it is now possible to automate installation by using a combination of the configuration file, command line options, and batch mode, so that an SAP HANA system installation can be executed once without any follow-up selection.

To illustrate the purpose of installation automation, let us assume there is a hardware partner who offers SAP HANA platforms (SAP HANA system installed on SAP verified hardware) in three sizes: Small, Large, and Extra Large.

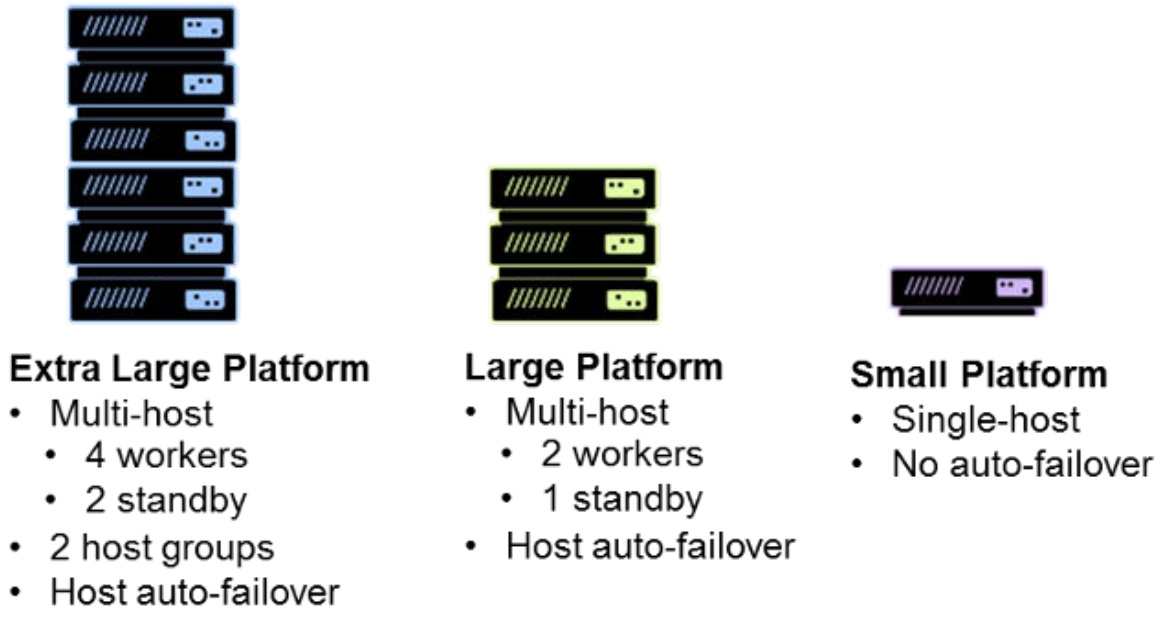

Since the hardware partner prepares the SAP HANA platform on-demand, he needs to be able to react quickly to orders as they are placed. During slow sales periods, manual installation is feasible, however during busy sales periods, manual installation could create too much overhead. In this case, installation automation would allow the hardware partner to start as many installation copies as required, without any further interaction with the installer. It would even be possible for the partner to start the installation near the end of the business day, leave the installation, and ship out the order the next morning.

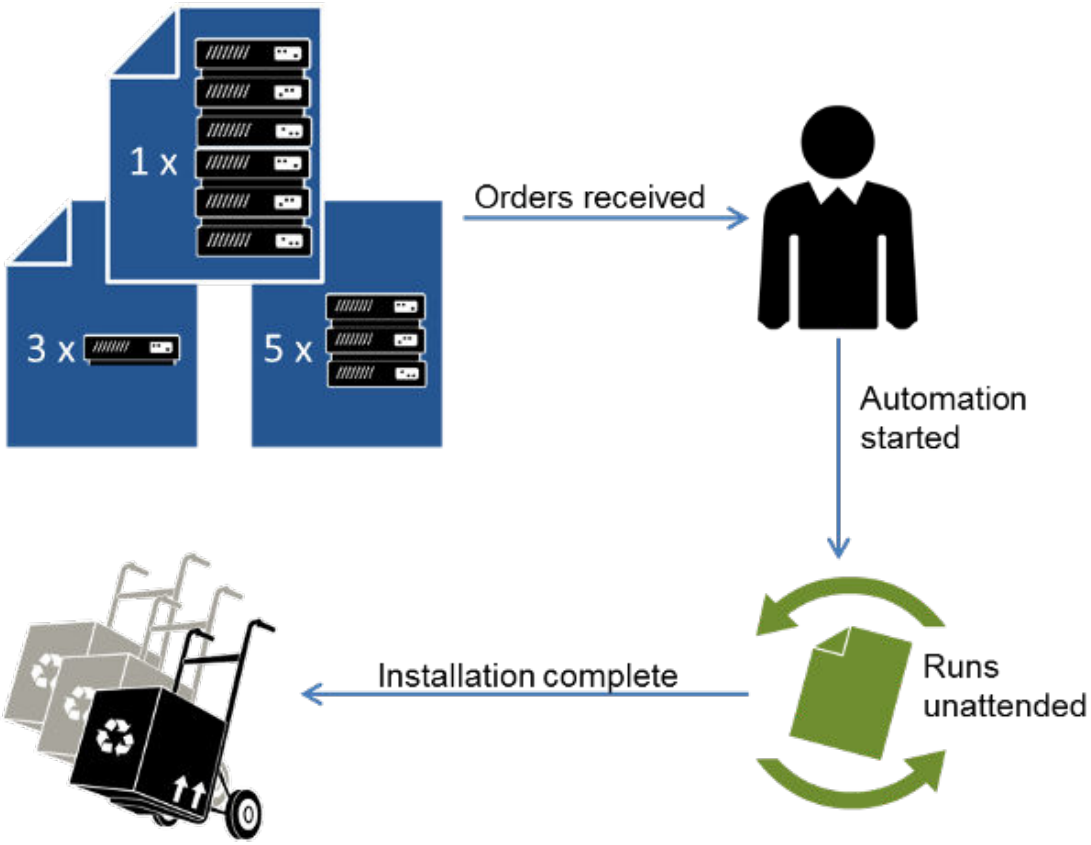

The hardware partner automates the installation of nine SAP HANA systems (1 Extra Large, 5 Large, and 3 Small). Previously, he has created configuration files for each of the three system types. So, when several orders come in at the same time, he fine tunes his installation script to include the number of systems and calls the SAP HANA database lifecycle manager using the command-line interface with the configuration file parameter in batch mode. When the installation script is run, SAP HANA is installed on both the single-host and multiple-host systems, without any additional input. By reusing the same configuration files, the installations are reliable, flexible, and efficient.

Step 1: Prepare for automation with the automation checklist.

Automation is ideal for installations that run unattended. Before you start an automated installation, it is recommended to consider the following.

For All System Types

Do the operating system and additional software components meet the requirements?	Refer to the *SAP HANA Hardware and Software Requirements*.
Do you know where the data volumes and the log volumes will be located?	For security reasons, the data device and log device should not be the same. Refer to the *SAP HANA System Types*.
Where will the required file systems be located and is the storage sufficient?	Refer to the *Recommended File System Layout*.
Have you performed a hardware check?	The installer performs a hardware check before installation, for automated installation it is recommended to perform a manual hardware check first. Refer to the *Hardware and Software Requirements*.
Where will you specify passwords?	SYSTEM, `<sid>`adm, and root passwords must be specified in the configuration file or in an XML file. Refer to the *Specifying Passwords* information.
For the Multiple-Host System Type Only	
How many worker hosts and standby hosts will there be? How will they be grouped?	System processes and data are distributed among worker hosts, including the original host. Standby hosts do not perform tasks. Refer to the *Multiple-Host System Concepts*.
How will storage devices be configured?	External storage can be configured so that hosts have shared or non-shared access. Refer to the *Multiple-Host System Concepts*.
The root user name must be the same for all hosts in a multiple-host system. Will the root user name for all hosts be "root"?	If not, the `root_user` parameter must be specified during installation. Refer to the *root_user* parameter information.

Step 2: Review the installation scenario.

In order to provide flexibility, it is possible to install the same SAP HANA system in several ways. The differences between installation methods are best depicted through a one-to-one comparison of the same system installed with each available method.

In the following example, there is a hardware partner who plans to install several SAP HANA single-host systems. His desired system has the following specifications:

- System name: P01
- Instance number: 01
- Installation path: `/hana/shared`
- Data path: `/hana/data/P01`
- Log path: `/hana/log/P01`
- User group ID: 110

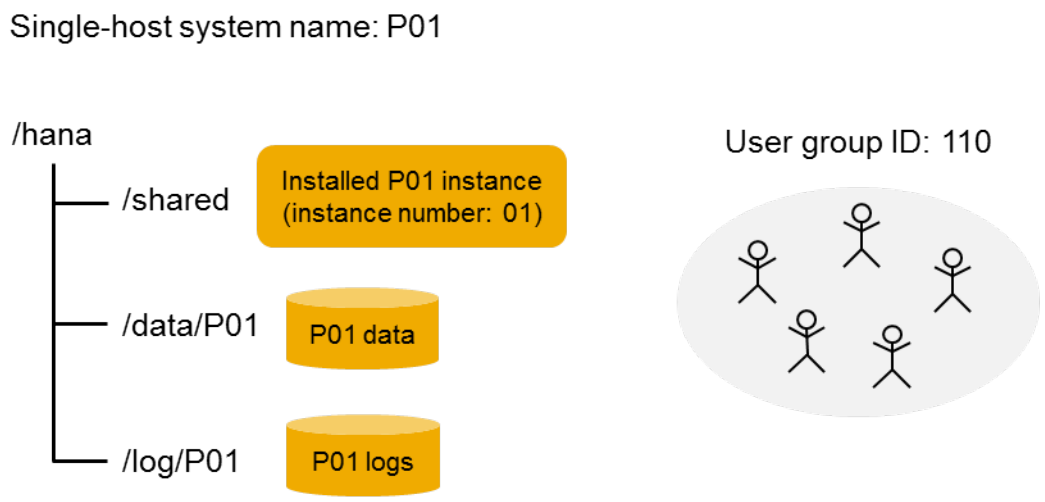

He has several customers who have pre-ordered the single-host SAP HANA P01 system and he is expecting more P01 orders. His goal is to iteratively improve his installation method to the point that he can automate his ideally configured system installation on several hosts at the same time, come back later, and the installed SAP HANA platforms will be finished and ready to ship. To reach his goal, he installs the same system (P01) three times, using:

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

Step 3: Create a test installation by installing the system using command line options.

Since the hardware partner is still new to installing SAP HANA, he tries out the installation first from the command line. He reviews the installation parameters and finds the corresponding parameter key-value pairs for his desired P01 single-host system.

System Detail	Specification	Command Line Parameter
System name	P01	`--sid=P01`
Instance number	01	`--number=01`
Installation path	/hana/shared	`--sapmnt=/hana/shared`
Data path	/hana/data/P01	`--datapath=/hana/data/P01`
Log path	/hana/log/P01	`--logpath=/hana/log/P01`
User group ID	110	`--groupid=110`

The hardware partner takes the parameter key-value pairs, and builds the command line input as follows:

```
./hdblcm --sid=P01 --number=01 --groupid=110 --sapmnt=/hana/shared --datapath=/
hana/data/P01 --logpath=/hana/log/P01
```

Upon review of the parameter syntax, he realizes he did not need to specify `sapmnt`, `datapath`, or `logpath` parameters because he chose the default values. The rest of the parameters could also have been given the short-form syntax. Therefore, the same P01 system installation could be simplified to the following syntax:

```
./hdblcm –s P01 –n 01 –G 110
```

System Detail	Specification	Simplified Command Line Parameter
System name	P01	`-s P01`
Instance number	01	`-n 01`
Installation path	/hana/shared	`<default>`
Data path	/hana/data/P01	`<default>`
Log path	/hana/log/P01	`<default>`
User group ID	110	`-G 110`

Step 4: Create a test installation by installing the system using the configuration file.

Now that the hardware partner successfully installed from the command line, he wants to make use of the configuration file, since his end goal is to automate the installation.

He generates a template of the configuration file:

```
./hdblcm --action=install --dump_configfile_template=/home/root/HANA_install.cfg
```

He opens the blank configuration file template and fills it in as follows:

```
HANA_install.cfg# SAP HANA System ID
sid=P01

# Instance Number (Default: 00)
number=01
…
# ID of User Group 'sapsys'
Groupid=110
```

If you want to run the installation with a configuration file and you do not want be asked interactively for specific parameters, you can either provide a fixed value in the configuration file or set the value to `USE_DEFAULT` to use the paramater's default value. Example `addhosts=USE_DEFAULT`.

You can specify a new default value for a parameter to be used in interactive mode using the following syntax:`<parameter>::DEFAULT=<new_default_value>`.

Since it was already discovered that only the SID, instance number, and group ID differ from the default values, only they are specified in the configuration file. The hardware partner also realizes that there is automatic substitution of the SID (`sid`) and installation path (`sapmnt`) throughout the configuration file according to default values (for example, `datapath` default: `/hana/data/${<sid>}`), so he does not need to append the SID (in this case, P01) to the paths.

Now, he can start the installer from the command line with the following simple command:

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

Step 5: Automate the system installation using the configuration file in batch mode

Now the hardware partner can take the last step of automation with the SAP HANA lifecycle management tool hdblcm and use batch mode. It is important to note, up this point the hardware partner has been interactively entering passwords and confirming other default parameters as part of interactive mode. Batch mode runs the installer without asking for any confirmation or parameter entry, thereby allowing installation to run to completion from one push of a button.

There are two choices for password entry in batch mode. Either the passwords can be stored in an XML file and passed to the installer as a stream by standard input, or they can be specified in the configuration file. Since the hardware partner is already making use of the configuration file, he decides to enter the passwords there. The only mandatory parameters in batch mode are the SID and passwords, so he checks the other defaults of the mandatory installation values before continuing the installation. In batch mode, the installer accepts default values for any unspecified parameters.

With the addition of passwords to the configuration file, it now looks like this:

```
HANA_install.cfg# Root User Password
root_password=Root1234
...
# SAP HANA System ID
sid=P01

# Instance Number (Default: 00)
number=01
...
# ID of User Group 'sapsys'
Groupid=110
...
# SAP Host Agent (sapadm) Password
sapadm_password=Agent1234
...
# System Administrator Password
password=Adm1234
....
# Database User (SYSTEM) Password
system_user_password=Sys1234
```

Now, the partner starts the installer, this time with the addition of the batch mode parameter, `-b` (or `--batch`):

```
./hdblcm --configfile=/home/root/HANA_install.cfg -b
```

Related Information

Getting Started with SAP HANA System Installation [page 72]
SAP HANA Hardware and Software Requirements [page 15]

SAP HANA System Types [page 23]
Recommended File System Layout [page 18]
Specifying Passwords [page 77]
root_user [page 273]
Multiple-Host System Concepts [page 85]

11.2 Tutorial: Using Custom Configuration Files

During system installation, customized SAP HANA configuration (*.ini) files can be placed inside a dedicated directory that is specified using the `custom_cfg` parameter.

Context

Custom configuration files can be placed inside a custom configuration folder. These files are used during installation to override default settings. This reduces the number of restarts during installation and facilitates configuration of the SAP HANA system.

Procedure

1. Create an empty directory that will contain your custom configuration files. The configuration files will be copied to the following directory before system start and will override the default settings: `<sapmnt>/<SID>/global/hdb/custom/config`
2. Create your own configuration files (*.ini) inside the custom configuration files directory. Alternatively, you can copy existing files from another system to re-use a tested configuration. For more information on configuration files, see *Configuring SAP HANA System Properties (INI Files)* in the *SAP HANA Administration Guide*.
3. Now the installer can be called from the command line.

```
./hdblcm --action=install --custom_cfg=<path to directory containing custom configuration files>
```

Related Information

custom_cfg [page 248]
Configuring SAP HANA System Properties (INI Files)

11.3 Tutorial: Installing a Multiple-Host System Using a Configuration File in Batch Mode

The following use case installs a multiple-host system, with two hosts (both worker). The installer reads the parameters from a configuration file. The installation is run in batch mode, so that once the installation is started, both host installations are configured without any further input required.

Procedure

1. The following are the relevant parameters for the SAP HANA server. They are specified in a combination of command line options and configuration file:

Parameter	Input	Location Specified

Passwords	`# Root User Password`
	`root_password=Root1234`
	`...`
	`# SAP Host Agent  (sapadm)`
	`Password sapadm_password=Agent1234`
	`...`
	`# System Administrator`
	`Password password=Adm1234`
	`....`
	`# Database User  (SYSTEM) Password`
	`system_user_password=Sys1234`
		Configuration file

`action`	`--action=install`	Call option
`sid`	`sid=ABC`	Configuration file
`number`	`number=01`	Configuration file
`root_user`	`root_user=sysroot`	Configuration file
`addhosts`	`addhosts=hananode1`	Configuration file
`configfile`	`--configfile=/home/sysroot/hdblcm.cfg`	Call option
`batch`	`-b`	Call option
`userid`	`userid=55`	Configuration file

Parameter	Input	Location Specified
groupid	groupid=110	Configuration file
storage_cfg	storage_cfg=/home/sysroot/storage	Configuration file

2. The reusable installation parameter values are saved in the following configuration file:

```
[Server]
# Additional Hosts
addhosts=hananode1
# Root User Name (Default: root)
root_user=sysroot
# Root User Password
root_password=Root1234
# SAP HANA System ID
sid=ABC
# Instance  Number (Default: <next available number>)
number=01
# SAP Host Agent (sapadm) Password
sapadm_password=Agent1234
#  System Administrator Password
password=Adm1234
#  System Administrator User ID (Default: <next available number>)
userid=55
# ID of User Group 'sapsys' (Default: 79)
groupid=110
# Directory containing a storage configuration
storage_cfg=/home/sysroot/storage
# Database User (SYSTEM) Password
system_user_password=Sys1234
```

3. Now the installer can be called from the command line, with the remaining parameters read from the configuration file. The installation is run in batch mode, so no follow-up confirmation is required.

```
./hdblcm --action=install --configfile=/home/sysroot/hdblcm.cfg -b
```

11.4 Tutorial: Overwriting Configuration File Parameters with Command Line Parameters

The following use case uses the same configuration file as above. However, this time, the desired system deviates slightly from the one defined the configuration file. By specifying parameters in the command line which are already specified in the configuration file, the configuration file settings are effectively overwritten. The command line parameters take precedence over the configuration file parameters.

Procedure

1. The following are the desired parameters for the SAP HANA server. They are specified in a combination of command-line and configuration file:

Parameter	Input	Location Specified	Used for Installation
Passwords	`# Root User Password root_password=Root1234 ... # SAP Host Agent (sapadm) Password sapadm_password=Agent1234 ... # System Administrator Password password=Adm1234 .... # Database User (SYSTEM) Password system_user_password=Sys1234`	Configuration file	yes
`action`	`--action=install`	CLI	yes
`sid`	`sid=ABC`	Configuration file	no
`sid`	`--sid=DB1`	CLI	yes
`number`	`number=01`	Configuration file	no
`number`	`-n 01`	CLI	yes
`root_user`	`root_user=sysroot`	Configuration file	yes
`addhosts`	`addhosts=hananode1`	Configuration file	no
`addhosts`	`--addhosts=hananode1,hananode2:role=standby`	CLI	yes
`configfile`	`--configfile=/home/sysroot/hdblcm_SPS7.cfg`	CLI	yes
`batch`	`-b`	CLI	yes
`userid`	`userid=55`	Configuration file	yes
`groupid`	`groupid=110`	Configuration file	yes
`storage_cfg`	`storage_cfg=/home/sysroot/storage`	Configuration file	no

Parameter	Input	Location Specified	Used for Installation
`storage_cfg`	`--storage_cfg=/home/sysroot/storage_new`	CLI	yes

2. The reusable installation parameter values are saved in the following configuration file:

```
[Server]
# Additional Hosts
addhosts=hananode1
# Root User Name (Default: root)
root_user=sysroot
# Root User Password
root_password=Root1234
# SAP HANA System ID
sid=ABC
# Instance  Number (Default: <next available number>)
number=01
# SAP Host Agent (sapadm) Password
sapadm_password=Agent1234
#  System Administrator Password
password=Adm1234
#  System Administrator User ID (Default: <next available number>)
userid=55
# ID of User Group 'sapsys' (Default: 79)
groupid=110
# Directory containing a storage configuration
storage_cfg=/home/sysroot/storage
# Database User (SYSTEM) Password
system_user_password=Sys1234
```

If you want to run the installation with a configuration file and you do not want be asked interactively for specific parameters, you can either provide a fixed value in the configuration file or set the value to `USE_DEFAULT` to use the paramater's default value. Example `addhosts=USE_DEFAULT`.

You can specify a new default value for a parameter to be used in interactive mode using the following syntax:`<parameter>::DEFAULT=<new_default_value>`.

> ### Example
>
> `sid::DEFAULT=ABC`

3. Now the installer can be called from the command line, with the remaining parameters read from the configuration file. The installation is run in batch mode, so no follow-up confirmation is required.

```
./hdblcm --action=install -n 01 --sid=DB1
--addhosts=hananode1,hananode2:role=standby --storage_cfg=/home/sysroot/
storage_new --configfile=/home/sysroot/hdblcm_SPS7.cfg -b
```

11.5 Tutorial: Installing a Single-Host System with Passwords Read from XML Standard Input Stream

The following use case installs a single-host system. The installer reads the parameters from the command line, and the passwords from a standard input stream. The installation is run in batch mode, so that once the installation is started, the host is configured without any further input required.

Context

This tutorial describes how passwords are supplied to the SAP HANA database lifecycle manager in an XML file in which passwords are stored as plain text. For security reasons, you may want to consider providing encrypted passwords along with a tool that decrypts the passwords and passes them to the standard input stream of the SAP HANA database lifecycle manager in XML format.

Procedure

1. The following are the desired parameters for the SAP HANA server, to be entered in command line in combination with the call to the installer.

Parameter	Input	Location Specified
Passwords	`--read_password_from_stdin=xml`	Call Option
`sid`	`--sid=P02`	Call Option
`number`	`--number=01`	Call Option
`root_user`	`--root_user=sysroot`	Call Option
`batch`	`--batch`	Call Option

2. The following password file is prepared with the accepted XML syntax:

```
hdb_passwords.xml
<?xml version="1.0" encoding="UTF-8"?>
<Passwords>
<password><![CDATA[Adm1234]]></password>
<sapadm_password><![CDATA[Agent1234]]></sapadm_password>
<system_user_password><![CDATA[Sys1234]]></system_user_password>
<root_password><![CDATA[Root1234]]></root_password>
</Passwords>
```

> **Note**
>
> Make sure to set the appropriate file permissions.

3. Now the installer can be called from the command line, with the passwords read from a standard input stream. The installation is run in batch mode, so no follow-up confirmation is required.

```
cat ~/hdb_passwords.xml | ./hdblcm --sid=P02 --number 01 --root_user=sysroot
--read_password_from_stdin=xml --batch
```

Related Information

read_password_from_stdin [page 270]

12 Troubleshooting

Troubleshooting should be referred to if the installation fails for an unknown reason, or for workarounds in special circumstances.

12.1 Accessing the Underlying Installer Components (pass_through_help)

Since `hdblcm` and `hdblcmgui` are wrapper tools, in some troubleshooting cases, it may be useful to 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. The `action` parameter and `--help` or `-h` must be specified in combination with `pass_through_help`.

Syntax

To view the help output for the installation or the update `pass_through_help` parameters, use the following syntax:

```
--action=[install|update] --pass_through_help --help
```

or

```
--action=[install|update] --pass_through_help -h
```

Available Parameters for pass_through_help

Parameter	Install	Update
`--hdbinst_client_ignore=<check1>[]...` Ignores failing prerequisite checks (`check_version`)	■	■
`--hdbinst_client_sapmnt=<installation_path>` Mount point for shared client installations [default: `--hdbinst_client_sapmnt=/hana/shared`]	■	■
`--hdbinst_plugin_ignore=<check1>[]...` Ignores failing prerequisite checks (`check_busy_files`, `check_version`)	■	■
`--hdbinst_plugin_nostart` Does not start the instance after installation	■	■

Parameter	Install	Update
`--hdbinst_plugin_system_user` Specifies the system user of the database [default: `--hdbinst_plugin_system_user=SYSTEM`]	■	■
`--hdbinst_server_prepare_update` Stops the update before software version switch	■	■
`--hdbinst_server_remote_execution` Specifies the connectivity method for multiple host operations [default: `ssh`]	■	■
`--hdbinst_server_ignore=<check1>[]...` Ignores failing prerequisite checks (`check_busy_files`, `check_diskspace`, `check_hosts`, `check_license`, `check_min_mem`, `check_pending_upgrade`, `check_plugin_dependencies`, `check_secondary_system`, `check_version`)	■	■
`--hdbinst_server_import_content[=off]` Imports delivery units [default: `--hdbinst_server_import_content`]	■	◇
`--hdbinst_server_xs_engine[=off]` Enables the XS engine [default: `--hdbinst_server_xs_engine`]	■	◇
`--hdbinst_server_xs_engine_http_port=<port>` Specifies the HTTP port of the XS engine	■	◇
`--hdbinst_server_xs_engine_https_port=<port>` Specifies the HTTPS port of the XS engine	■	◇
`--hdbinst_studio_features=all|<feat1>[,<feat2>]...` Specifies the features to be installed [default: `--hdbinst_studio_features=all`]	■	■
`--hdbinst_studio_path=<hdbinst_studio_path>` Installation path [default: `--hdbinst_studio_path=/usr/sap/hdbstudio`]	■	■
`--hdbupd_server_change_initial_ssfs_key[=off]` Changes the initial SSFS key [default: `off`]	◇	■

Parameter	Install	Update
`--hdbupd_server_ignore=<check1>[]...` Ignores failing prerequisite checks (`check_busy_files`, `check_diskspace`, `check_hosts`, `check_license`, `check_min_mem`, `check_pending_upgrade`, `check_plugin_dependencies`, `check_secondary_system`, `check_version`)	◇	■
`--hdbupd_server_import_content[=off]` Imports delivery units [default: `--hdbupd_server_import_content`]	◇	■
`--hdbupd_server_nostart` Does not start the instance after upgrade	◇	■
`--hdbupd_server_prepare_update` Stops the update before software version switch	◇	■
`--hdbupd_server_remote_execution` Specifies the connectivity method for multiple host operations [default: `ssh`]	◇	■
`--hdbupd_server_xs_engine[=off]` Enables the XS engine [default: `--hdbupd_server_xs_engine`]	◇	■
`--hdbupd_server_xs_engine_http_port=<port>` Specifies the HTTP port of the XS engine	◇	■
`--hdbupd_server_xs_engine_https_port=<port>` Specifies the HTTPS port of the XS engine	◇	■
XML password tag: `<hdbinst_plugin_password>` System administrator password	■	■

12.2 Locating all SAP HANA File System Components

In addition to the main components installed in the default file systems described in Recommended File System Layout [page 18], it may also be necessary to locate the temporary files from the SAP HANA system. They can be found in the following directories:

- SAP HANA DB files:
  - `/var/lib/hdb` - IPC data (volatile)
  - `/var/tmp` - Installer log files, `HDB_alive_*`

    - `/tmp` - `.hdb_*_lock` (volatile)
- `sapstartsrv` files:
    - `/etc/init.d/sapinit*` - Boot script
    - `/tmp/.sap*` - Sockets, lock files (volatile)

12.3 Enabling the Installer Trace

If the installer crashes or loops it may make sense to trace the installer until the problem occurs, open a CSS ticket, and attach the trace file for further analysis. You can switch on the installer trace by setting the environment variable `HDB_INSTALLER_TRACE_FILE` to `<tracefilename>`. The directory containing the trace file must already exist.

12.4 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>_<time_stamp>`. The time stamp depends on the time zone that was set as an environment variable for the user under which an action was executed.

Since the SAP HANA lifecycle management tools `hdblcm` and `hdblcmgui` are wrappers for underlying component installers, it is also possible to check the component logs. It is recommended to review and analyze the SAP HANA lifecycle management tools `hdblcm` and `hdblcmgui` logs first. Once the source of the problem is narrowed down to a specific component, then the component logs can be further analyzed.

The component log files are stored in the following path:

`/var/tmp/hdb_<SID>_hdblcm_<action>_<time_stamp>`

where `<action> :: = install | update | addhost | 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, makes a detailed analysis more easy

You can also view the last three log files in the SAP HANA studio using the administration function *Diagnosis Files*. For more information, see the *SAP HANA Administration Guide*.

12.5 Adding hdblcm to an Existing SAP HANA Installation

The SAP HANA lifecycle management tool `hdblcm` can be added to an existing installation that was performed using `hdbinst` by executing the following command:

```
./hdblcm --action=update --sid=<SID> --components=hdblcm
```

> **Note**
>
> The version of `hdblcm` must be identical with the version of the installed SAP HANA system.

12.6 Disabling the Optimized Update

As of SPS 11, an optimized update of an SAP HANA system is performed to reduce the number of restarts and system downtime. The optimized update execution mode is enabled by default, if more than one component that supports the phased update process is selected to be installed. For troubleshooting purposes, the update mode can be switched to standard. Perform a standard update with the SAP HANA database lifecycle manager using one of the following commands.

Graphical user interface	`./hdblcmgui --action=update --update_execution_mode=standard`
Command-line interface	`./hdblcm --action=update --update_execution_mode=standard`

If you are using the SAP HANA HDBLCM Web user interface, open the *Advanced Parameters Configuration* dialog from the footer bar and select the standard update *Update Execution Mode* under *General Parameters*.

12.7 Dealing with a Failed Update

If the update stops without fully installing, uninstall the SAP HANA system using the SAP HANA database lifecycle manager (HDBLCM). Then, reinstall the SAP HANA system and run the update. After successfully applying the update, recover the system from the last backup. Note that a recovery from backup is only possible on systems that are running a version that is equal to or greater than the version of the system on which the backup was created.

Related Information

Recovering an SAP HANA Database
Updating the SAP HANA System [page 148]

12.8 Importing Delivery Units Manually

If the import of delivery units fails during installation or update, you can import the delivery units manually. Use the command-line tool `hdbupdrep` to import the delivery units into the database. `hdbupdrep` is located in `/usr/sap/<SID>/SYS/global/hdb/install/bin`.

As of SAP HANA 2.0 Support Package Stack (SPS) 03, you can rerun the SAP HANA database lifecycle manager to resume the update.

Related Information

General Troubleshooting for the SAP HANA Platform LCM Tools [page 67]
SAP Note 1795885

12.9 Removing a Partially Installed System

If the installation stops without fully installing, and you would like to remove the components and start over, run the component uninstaller (`hdbuninst`) for each component individually. Then start the installation over again using `hdblcm`.

12.10 Adjusting the System Time

A host cannot be added using the SAP HANA lifecycle management tool `hdblcm` if the time difference between the system time set on the installation host and the additional host is greater than 180 seconds. For information about setting the system time, see the documentation of your Linux distribution.

12.11 Skipping the Import of Initial XS Advanced Runtime Content

The XS advanced runtime requires initial content to be imported during installation. However, this import can be skipped by executing the SAP HANA installation with the following parameter:

`--import_xs_content=off`

If the import of the initial content was skipped during installation, the content can be imported at a later time by executing the resident SAP HANA lifecycle management tool `hdblcm` with the following parameter:

`--load_initial_xs_content`

12.12 Proxy Server Settings for SAP HANA Installations with XS Advanced Runtime

An SAP HANA system with XS advanced runtime installed that is using a proxy server requires the values of the `http_proxy`, `https_proxy` and `no_proxy` environment variables to be set. To do so, execute the following commands:

```
export http_proxy=http://<proxy_host>:<proxy_port>
export https_proxy=https://<proxy_host>:<proxy_port>
export no_proxy="<full qualified name of the XS Advanced server>"
```

13 Parameter Reference

Reference information is provided for each installation parameter.

Each parameter section includes some or all of the following information, depending on the complexity of the parameter:

- **Description** - The function of the parameter.
- **Syntax** - The command line option format and configuration file format for the parameter. Interactive mode (for both the GUI and CLI) do not require the use of parameter syntax.
- **Options** - The sub-specifications for the highly configurable parameters.
- **Remarks** - Important information about the parameters, including default values and whether the parameter is offered in interactive mode. If it is not available in interactive mode, and the default value is not wanted, the parameter must be specified as a command line option or in the configuration file in combination with the call to the installation tool in interactive mode.
- **Examples** - The syntax in its common usage with acceptable parameter specifications.
- **Related Information** - Links to relevant sections in the current document or to other documents which contain more detailed information.

13.1 `action`

Specifies the action of hdblcm.

Syntax

In the command line, the following syntax is used:

```
--action=[extract_components|install|update|
print_detected_components]
```

Remarks

The default for this parameter is `--action=exit`.

This parameter is available in interactive mode.

13.2 `addhosts`

Specifies additional hosts for the SAP HANA system as a comma-separated list. Individual host options are specified by a colon-separated list. This parameter is used when configuring a multiple-host system during installation.

Requirements

If the root user has a user name other than "root", the `root_user` parameter must also be specified in combination with `addhosts`.

When used with command-line batch mode, the `action` parameter must be specified in combination with `addhosts`.

Syntax

In the command line, the following syntax is used:

```
--action=install --addhosts=<host>[,<host2>]
```

where the `<host>` syntax is as follows:

```
<host_name>[:role=worker|standby|extended_storage_worker|
extended_storage_standby| ets_worker|ets_standby|streaming|rdsync|
xs_worker|xs_standby][:group=<name>][:storage_partition=<number>]
```

Options

The following options can be used to configure the parameter:

- `role` - Specifies the purpose of the SAP HANA host. Although multiple host roles may be assigned, check the corresponding documentation for the SAP HANA option for what configurations are supported in production environments.
  - `worker` - A worker host (default) is used for database processing.
  - `standby` - A standby host is idle and available for failover in a high-availability environment.
  - `compute` - Database elastic compute server
  - `extended_storage_worker` - Worker host for SAP HANA dynamic tiering
  - `extended_storage_standby` - Standby host for SAP HANA dynamic tiering
  - `ets_worker` - Worker host for SAP HANA accelerator for SAP ASE
  - `ets_standby` - Standby host for SAP HANA accelerator for SAP ASE
  - `streaming` - Host for SAP HANA streaming analytics
  - `xs_worker` - Host for SAP HANA XS advanced runtime
  - `xs_standby` - Standby host for SAP HANA XS advanced runtime
- `workergroup` - Specifies the worker group of the host. If undefined, the worker group is named "default". If you are using extension node for Business Warehouse, you must name the worker group "worker_dt".
- `group` - Specifies the host group ID for failover scenarios. If undefined, the host group is named "default".
  (The host group ID is NOT the same as the sapsys group ID, which is specified by the parameter `groupid`).
- `storage_partition` - Specifies the storage partition number, which is a logical role number assigned to non-shared storage devices in a storage connector API. Standby hosts do not have a storage partition.

Remarks

This parameter is available in interactive mode.

The following SAP HANA options are supported on Intel-based hardware platforms only:

- SAP HANA Accelerator for SAP ASE

> ⚠ Caution
>
> Be aware that you need additional licenses for SAP HANA options. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities* in Related Information.

Related Information

Multiple-Host System Concepts [page 85]
Install a Multiple-Host SAP HANA System Using the Graphical User Interface [page 89]
root_user [page 273]
action [page 233]
Important Disclaimer for Features in SAP HANA [page 304]

13.3 `add_local_roles`

Specifies additional roles of the local host during SAP HANA system installation. Multiple host roles are **not** supported in production environments. However, if XS advanced runtime is installed, hosts can share multiple roles.

Syntax

In the command line, the following syntax is used:

```
--add_local_roles=<role1>[,<role2>]
```

where the following roles can be specified:

- `worker` - A worker host (default) is used for database processing.
- `standby` - A standby host is idle and available for failover in a high-availability environment.
- `compute` - Database elastic compute server
- `extended_storage_worker` - Worker host for SAP HANA dynamic tiering
- `extended_storage_standby` - Standby host for SAP HANA dynamic tiering
- `ets_worker` - Worker host for SAP HANA accelerator for SAP ASE
- `ets_standby` - Standby host for SAP HANA accelerator for SAP ASE
- `streaming` - Host for SAP HANA streaming analytics
- `xs_worker` - Host for SAP HANA XS advanced runtime
- `xs_standby` - Standby host for SAP HANA XS advanced runtime

Remarks

The following SAP HANA options are supported on Intel-based hardware platforms only:

- SAP HANA Accelerator for SAP ASE

This parameter is only available for the resident `HDBLCM` program.

> ⚠ Caution
>
> Be aware that you need additional licenses for SAP HANA options. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities* in Related Information.

Related Information

Important Disclaimer for Features in SAP HANA [page 304]

13.4 `add_roles`

Specifies additional roles for exisiting SAP HANA hosts during SAP HANA option installation. Multiple host roles are **not** supported in production environments. However, if XS advanced runtime is installed, hosts can share multiple roles.

Syntax

In the command line, the following syntax is used:

```
--add_roles=<host name>=<role> -R <host name>=<role>
```

where the following roles can be specified:

- `compute` - Database elastic compute server
- `extended_storage_worker` - Worker host for SAP HANA dynamic tiering
- `extended_storage_standby` - Standby host for SAP HANA dynamic tiering
- `ets_worker` - Worker host for SAP HANA accelerator for SAP ASE
- `ets_standby` - Standby host for SAP HANA accelerator for SAP ASE
- `streaming` - Host for SAP HANA streaming analytics
- `xs_worker` - Host for SAP HANA XS advanced runtime
- `xs_standby` - Standby host for SAP HANA XS advanced runtime

> ⚠ Caution
>
> Be aware that you need additional licenses for SAP HANA options. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities* in Related Information.

Related Information

Important Disclaimer for Features in SAP HANA [page 304]

13.5 `apply_system_size_dependent_parameters`

Applies a resource limit depending on the system size..

Syntax

In the command line, the following syntax is used:

```
--apply_system_size_dependent_parameters[=off]
```

Remarks

The default for this parameter is `--apply_system_size_dependent_parameters` (on).

The parameter is available during installation and update to a newer version. If you want to apply a resource limit without upgrading to a higher SAP HANA revision, specify `--ignore=check-version` when starting the update.

Related Information

SAP Note 3014176

13.6 `ase_datapath`

Specifies the path to the directory of the SAP HANA accelerator for SAP ASE data. Required for installation of SAP HANA accelerator for SAP ASE.

Syntax

In the command line, the following syntax is used:

```
--ase_datapath=<path to SAP HANA accelerator for SAP ASE data
directory>
```

Remarks

The default for this parameter is `--ase_datapath=/hana/data_ase/<SID>`.

This parameter is available in interactive mode.

This parameter is not supported by SAP HANA on IBM Power Systems.

> ⚠ Caution
>
> Be aware that you need additional licenses for SAP HANA options. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities* in Related Information.

Related Information

Important Disclaimer for Features in SAP HANA [page 304]

13.7 ase_logpath

Specifies the path to the directory of the SAP HANA accelerator for SAP ASE logs. Required for installation or update of SAP HANA accelerator for SAP ASE.

Syntax

In the command line, the following syntax is used:

```
--ase_logpath=<path to SAP HANA accelerator for SAP ASE log directory>
```

Remarks

The default for this parameter is `--ase_logpath=/hana/log_ase/<SID>`.

This parameter is available in interactive mode.

This parameter is not supported by SAP HANA on IBM Power Systems.

> ⚠ Caution
>
> Be aware that you need additional licenses for SAP HANA options. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities* in Related Information.

Related Information

Important Disclaimer for Features in SAP HANA [page 304]

13.8 ase_user

Specifies the administrator user of SAP HANA accelerator for SAP ASE.

Syntax

In the command line, the following syntax is used:

```
--ase_user=<administrator user name>
```

Remarks

The default for this parameter is `--ase_user=sa`.

This parameter is available in interactive mode.

This parameter is not supported by SAP HANA on IBM Power Systems.

> ⚠ Caution
>
> Be aware that you need additional licenses for SAP HANA options. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities* in Related Information.

Related Information

Important Disclaimer for Features in SAP HANA [page 304]

13.9 `autoadd_xs_roles`

Assigns XS_WORKER and XS_STANDBY host roles. The host role XS_WORKER will be assigned to all worker hosts, the host role XS_STANDBY will be assigned to all standby hosts. To create a multiple-host system with dedicated xs_worker and xs_standby hosts, assign host roles to each host individually during installation. Do not choose the option to assign XS Advanced host roles automatically.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --autoadd_xs_roles[=off]
```

Remarks

The default for this parameter is `--autoadd_xs_roles` (on).

This parameter is available in interactive mode.

13.10 `autostart`

Restarts system after machine reboot.

Syntax

In the command line, the following syntax is used:

```
--action=install --autostart=[0|1]
```

where 0 = off, and 1 = on

Remarks

The default for this parameter is `--autostart=0` (off).

> **Note**
>
> To avoid the start of the SAP HANA system during the operating system upgrade, set the `autostart` parameter in the profile to `0` (=off).
>
> The profile can be found here: `<sapmnt>/<SID>/profile/<SID>_HDB<instance_number>_<hostname>`

13.11 `backup_encryption`

Enables backup encryption.

Syntax

In the command line, the following syntax is used:

```
--backup_encryption[=off]
```

Remarks

The default for this parameter is `--backup_encryption` (on).

This parameter is available in interactive mode.

Related Information

Backup Encryption

13.12 `basepath_streaming`

Specifies the location of streaming logstores and runtime information.

Syntax

In the command line, the following syntax is used:

```
--basepath_streaming=<location of streaming logstores and
runtime information>
```

Remarks

The default for this parameter is `--basepath_streaming=/hana/data_streaming/<SID>`.

This parameter is available in interactive mode.

13.13 `batch`

Runs the SAP HANA lifecycle management tool from the command line in batch mode using default values for unspecified parameters. If mandatory values are omitted or if invalid values are specified, the program issues an error message.

Syntax

In the command line, the following syntax is used:

```
--action=[extract_components|install|update] --batch
```

or, in short form:

```
-b
```

Related Information

Specifying Passwords [page 77]

13.14 `certificates_hostmap`

Specifies the hostname used for generation of self-signed SSL certificates for the SAP Host Agent.

Requirements

The key-value pair `action=install` must be specified in combination with `certificates_hostmap`.

Syntax

In the command line, the following syntax is used:

```
--action=install --certificates_hostmap=<host name>=<certificate host name>
```

or

```
--action=install -C <host name>=<certificate host name>
```

Remarks

The default for this parameter is the current host.

This parameter is available in interactive mode.

Example

The following example generates certificates for two hosts in the long-form syntax:

```
./hdblcm --action=install --certificates_hostmap=hananode01=server1.company.com
--certificates_hostmap=hananode02=server2.company.com
```

The following example generates certificates for two hosts in the short-form syntax:

```
./hdblcm --action=install -C hananode01=server1.company.com -C
hananode02=server2.company.com
```

In this example, not all required installation parameters are specified in the command line. If this is the case, the remaining mandatory parameters are requested interactively.

Related Information

action [page 233]
components [page 245]

13.15 `change_lss_backup_password`

Changes the password of the Local Secure Store (LSS) backup.

Syntax

In the command line, the following syntax is used:

```
--change_lss_backup_password
```

Remarks

This parameter is available in interactive mode.

13.16 `change_system_user_password`

Changes the password of the database user (SYSTEM).

Syntax

In the command line, the following syntax is used:

```
--change_system_user_password
```

Remarks

This parameter is available in interactive mode.

13.17 `check_installation`

Checks the SAP HANA installation

Syntax

In the command line, the following syntax is used:

```
--action=check_installation
```

Remarks

This parameter is only available for the resident `HDBLCM` program.

13.18 `checkmnt`

Specifies a non-standard shared file system, which can be accessed by all hosts during installation. This parameter is typically used when the SID is included in the mountpoint.

Syntax

In the command line, the following syntax is used:

```
--action=install --checkmnt=<path>
```

13.19 `check_only`

Executes checks, but does not change the SAP HANA system.

Syntax

In the command line, the following syntax is used:

```
--action=update --check_only
```

13.20 client_path

Specifies the installation path for the client.

Syntax

In the command line, the following syntax is used:

```
--client_path=<path for client installation>
```

Remarks

The default for this parameter is `--client_path=<sapmnt>/<SID>/hdbclient`.

13.21 component_archives_dir

Specifies the directory of the SAP HANA database component archives.

Syntax

In the command line, the following syntax is used:

```
--action=extract_components --
component_archives_dir=<component_archives_path>
```

Remarks

This parameter is available in interactive mode.

13.22 component_dirs

Specifies the installer component directories as a comma-separated list.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --component_dirs=<component
directory>
```

Remarks

This parameter supports relative paths.

13.23 component_medium

Specifies the location of the SAP HANA installation medium.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --component_medium=<directory of the
installation medium>
```

13.24 component_root

Specifies the directory root to search for components.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --component_root=<component root
directory>
```

Remarks

This parameter supports relative paths.

13.25 components

Specifies the components to be installed in combination with the SAP HANA server: SAP HANA client, SAP HANA studio, and additional system components like Application Function Libraries (AFL, EML and the product-specific AFLs IBP, RTL, TRP, VCH, XRP), SAP liveCache applications (SAP LCA or LCAPPS plug-in), SAP HANA EPM-MDS (EPMMDS plug-in), or SAP HANA options. It is also possible to specify all components, or to specify a combination of components as a comma-separated list. The server is always installed, even if it is not explicitly specified.

Requirements

The parameter `action` must be specified in combination with `components`.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update|uninstall]
--components[=all|afl|client|es|ets|lcapps|pos|sal|sca|server|
smartda|sop|streaming|studio|udf|xs]
```

Remarks

The default for this parameter is `--components=client,server` and is dependent on the installer finding installation sources for the components.

> ⚠ Caution
>
> Be aware that you need additional licenses for SAP HANA options. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities* in Related Information.

This parameter is available in interactive mode.

Example

The following example installs the SAP HANA client, the SAP HANA studio, and the SAP HANA database (always installed, despite the specification):

```
./hdblcm --action=install --components=client,studio
```

In this example, not all required installation parameters are specified in the command line. If this is the case, the remaining mandatory parameters are requested interactively.

Related Information

action [page 233]
Important Disclaimer for Features in SAP HANA [page 304]

13.26 `configfile`

Loads a configuration file of call option key-value pairs to be passed to the SAP HANA lifecycle management program.

Syntax

In the command line, the following syntax is used:

```
--action=[extract_components|install|update] --
configfile=<file path>
```

Remarks

The configuration file makes installation and configuration tasks more efficient and reliable. For more information, see Related Information.

This complement to this call option is the call option `dump_configfile_template`.

Related Information

dump_configfile_template [page 250]

13.27 `configure_fast_restart_option`

Allows configuration of the Fast Restart Option.

Syntax

In the command line, the following syntax is used:

```
--configure_fast_restart_option
```

Remarks

This parameter is available in interactive mode.

13.28 `continue_update`

Continues the pending update with the persisted parameters. For more details about update planning and updating in a two-phase approach, see Related Information.

Syntax

In the command line, the following syntax is used:

```
--action=update --continue_update[=off]
```

Remarks

The default for this parameter is `--continue_update` (on).

This parameter is available in interactive mode.

Related Information

Prepare an Update for Flexible System Downtime [page 154]

13.29 copy_repository

Specifies the target path to which the SAP HANA studio repository should be copied.

Syntax

In the command line, the following syntax is used:

```
--copy_repository=<target path>
```

Remarks

The default for this parameter is `--copy_repository=/<sapmnt>/<SID>/hdbstudio_update`

13.30 create_initial_tenant

Creates the initial tenant database.

Syntax

In the command line, the following syntax is used:

```
--action=install --create_initial_tenant[=off]
```

Remarks

The default for this parameter is `--create_initial_tenant` (on).

13.31 custom_cfg

Specifies the path to the directory which contains custom configuration (*.ini) files.

Syntax

In the command line, the following syntax is used:

```
--action=install --custom_cfg=<path to directory containing
custom configuration files>
```

Remarks

Customized versions of SAP HANA configuration files for configuring the system as a whole and individual tenant databases, hosts, and services can be placed inside the directory. These configuration files will be copied to the following directory before system start and will override the default settings:

```
<sapmnt>/<SID>/global/hdb/custom/config
```

Parameter values passed by the SAP HANA database lifecycle manager (HDBLCM) take precedence over corresponding values that are specified in customized configuration files.

Related Information

Configuring SAP HANA System Properties (INI Files)
Tutorial: Using Custom Configuration Files [page 219]

13.32 `databackuppath`

Specifies the location of the data backups.

Syntax

In the command line, the following syntax is used:

```
--databackuppath=<path>
```

Remarks

This parameter is available in interactive mode.

The path must be specified if the new directory is located on a different physical storage. The contents of the directory must be moved manually to the new location.

13.33 `datapath`

Specifies the path to the data directory of the SAP HANA system.

Syntax

In the command line, the following syntax is used:

```
--action=install --datapath=<path to data directory>
```

Remarks

The default for this parameter is `--datapath=/hana/data/<SID>`.

This parameter is available in interactive mode.

The path must be specified if the new directory is located on a different physical storage. The contents of the directory must be moved manually to the new location.

13.34 db_isolation

Specifies the isolation of the tenant databases on operating system level for SAP HANA systems. By default, all database processes run under the default OS user `<sid>adm`. If it's important to mitigate against cross-database attacks through OS mechanisms, you can 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 subsequently protected using standard OS file and directory permissions.

Syntax

In the command line, the following syntax is used:

```
--action=install --db_isolation=high|low
```

Remarks

The default for this parameter is `--db_isolation=low`.

In high isolation mode, `/hana/shared/<SID>` must not be mounted with the option `nosuid`.

For more information about database isolation, see *Database Isolation* in the *SAP HANA Administration Guide* or the *SAP HANA Security Guide*.

Related Information

Database Isolation [page 37]
Recommended File System Layout [page 18]

13.35 dump_configfile_template

Specifies a file path to which a template configuration file is exported. The call options in the template configuration file are set to their default values, and can be edited.

Syntax

In the command line, the following syntax is used:

```
--action=[extract_components|install|update] --
dump_configfile_template=<file path>
```

Remarks

The configuration file makes installation and configuration tasks more efficient and reliable. For more information, see Related Information.

This complement to this call option is the call option `configfile`.

13.36 `encryption_algorithm`

Specifies the encryption algorithm used during system installation. It is applicable only if the secure store is Local Secure Store.

Syntax

In the command line, the following syntax is used:

```
--encryption_algorithm=[AES-256-CBC|AES-256-CTR|ARIA-256-CBC|
ARIA-256-CTR]
```

Related Information

Local Secure Store (LSS) [page 28]

13.37 `es_datapath`

Specifies the path to the directory of the SAP HANA dynamic tiering data. Required for installation of SAP HANA dynamic tiering.

Syntax

In the command line, the following syntax is used:

```
--es_datapath=<path to SAP HANA dynamic tiering data directory>
```

Remarks

The default for this parameter is `--es_datapath=/hana/data_es/<SID>`.

This parameter is available in interactive mode.

> ⚠ Caution
>
> Be aware that you need additional licenses for SAP HANA options. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities* in Related Information.

Related Information

Important Disclaimer for Features in SAP HANA [page 304]

13.38 es_logpath

Specifies the path to the directory of the SAP HANA data tiering logs. Required for installation or update of SAP HANA dynamic tiering.

Syntax

In the command line, the following syntax is used:

```
--es_logpath=<path to SAP HANA dynamic tiering log directory>
```

Remarks

The default for this parameter is `--es_logpath=/hana/log_es/<SID>`.

This parameter is available in interactive mode.

> ⚠ Caution
>
> Be aware that you need additional licenses for SAP HANA options. For more information, see *Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities* in Related Information.

Related Information

Important Disclaimer for Features in SAP HANA [page 304]

13.39 extract_components

Extracts content that was downloaded from the SAP Support Portal for installation or update. For more details about preparing software archives for update, see Related Information.

Syntax

In the command line, the following syntax is used:

```
--extract_components
```

Options

The following options are available:

`component_archives_dir` - Location of the SAP HANA component archives.
`extract_temp_dir` - The target directory to extract the software component archives to.
`sapcar_location` - Location to the SAPCAR executable.
`tar_executable_location` - Location of the tar executable.

Remarks

The default for this parameter is `--extract_components --component_archives_dir --extract_temp_dir=<component_archives_dir>/extracted --sapcar_location=<install path>/<SID>/global/hdb/saphostagent_setup/SAPCAR --tar_executable_location=</bin/tar>`.

This parameter is available in interactive mode.

Related Information

Prepare the Software Archive for the Update

13.40 `extract_temp_dir`

Specifies the temporary extraction directory.

Syntax

In the command line, the following syntax is used:

```
--action=extract_components --
extract_temp_dir=<extract_temp_dir>
```

Remarks

This parameter is available in interactive mode.

13.41 `groupid`

Specifies the SAP system (sapsys) group ID. This parameter is relevant only if a sapsys group does not already exist on the host. If a sapsys group already exists, passing the `groupid` parameter does not alter the existing group.

Syntax

In the command line, the following syntax is used:

```
--action=install --groupid=<sapsys group ID>
```

or, in short form:

```
-G <sapsys group ID>
```

Remarks

The default for this parameter is `--groupid=79`.

This parameter is available in interactive mode.

13.42 `help`

Displays the inline help information.

Syntax

In the command line, the following syntax is used:

```
--action=[extract_components|install|update] --help
```

or, in short form:

```
-h
```

Remarks

A general help output is available for all SAP HANA lifecycle management programs. Task-specific help output is available for some programs. Refer to the task documentation for more details.

13.43 `home`

Specifies the home directory of the system administrator. This parameter is relevant only if the operating system administrator (`<sid>adm`) does not exist prior to installation.

Syntax

In the command line, the following syntax is used:

```
--action=install --home=<home directory>
```

Remarks

The default for this parameter is `--home=/usr/sap/<SID>/home`.

This parameter is available in interactive mode.

The home directory must not be located inside a directory that is used by the SAP HANA installation. The home directory must be a local directory on each system host, i.e. not shared between the hosts.

13.44 `hostmap`

Specifies the host mapping to rename one host. The parameter has to be specified for each renamed host.

Syntax

In the command line, the following syntax is used:

```
--hostmap=<old host name>=<new host name>
```

or, in short form:

```
-H <old host name>=<new host name>
```

Remarks

This parameter is available in interactive mode.

Example

The following command renames two hosts, host1 and host2, to host3 and host4:

```
./hdblcm --action=rename_system -H host1=host3 -H host2=host4
```

13.45 `hostname`

Specifies the virtual host name of the system host.

Syntax

In the command line, the following syntax is used:

```
--hostname=<name of the host machine>
```

or, in short form:

```
-H <name of the host machine>
```

Remarks

Restrictions apply to host names in SAP HANA systems. Alphanumerical string of lowercase alpha characters [a-z] and digits [0-9] and the hyphen (or minus) character "-" are permitted. Although the newer RFCs permit hostnames beginning with digits we recommend hostnames to begin with an alpha character. The period character "." is only allowed to delimit components of domain names like (sapc11.sap.com). Host names with up to 64 characters are supported.

The default for this parameter is the host name of the current machine.

This parameter is available in interactive mode.

13.46 `https`

Specifies whether or not to use HTTPS.

Syntax

In the command line, the following syntax is used:

```
--https
```

Remarks

This parameter is available in interactive mode.

13.47 `ignore`

Specifies failing prerequisite checks that the SAP HANA platform lifecycle management tools should ignore.

Syntax

In the command line, the following syntax is used:

```
--ignore=<check1>[,<check2>]...
```

Remarks

Ignores the following prerequisite checks:

Check	Description
check_busy_files	Verify that the files to be modified are not locked by a process.
check_component_dependencies	Verify that all component dependencies are resolved.
check_diskspace	Verify that enough free disk space of the file system is available.
check_min_mem	Verify that enough physical memory on the host is available.
check_pending_upgrade	Verify if an upgrade is pending and can be resumed.
check_platform	Verify that the required gcc runtime and other libraries are present.
check_resume_hostname	Verify that the operation is resumed from the correct host.
check_signature_file	Verify that the signature file exists and has a correct format.

Check	Description
check_version	Verify that the version to be installed is newer than the current one.
verify_signature	Verify that the software packages are authentic.

13.48 `import_content`

Imports delivery units.

Syntax

In the command line, the following syntax is used:

```
--import_content[=off]
```

Remarks

The default for this parameter is `--import_content`.

Related Information

SAP HANA Content

13.49 `import_xs_content`

Imports SAP HANA XS advanced runtime content.

Syntax

In the command line, the following syntax is used:

```
--import_xs_content[=off]
```

Remarks

The default for this parameter is `--import_xs_content`.

13.50 `init_user`

Uses the properties of an existing system administrator (`<sid>adm`) if the user ID does not match the user ID that was used to set up the SAP HANA system.

Syntax

In the command line, the following syntax is used:

```
--init_user
```

13.51 `init_user_home_dir`

Creates a home directory for the system administrator on each host. This parameter is relevant only if the home directory of the operating system administrator (`<sid>adm`) does not exist.

Syntax

In the command line, the following syntax is used:

```
--init_user_home_dir
```

13.52 `install_execution_mode`

Specifies the install execution workflow. In optimized installation mode, the number of restarts is reduced during installation.

Syntax

In the command line, the following syntax is used:

`--install_execution_mode=[optimized|standard]`

Remarks

The default for this parameter is `--install_execution_mode=standard`.

The default value changes to `--install_execution_mode=optimized` if a server plug-in is being installed or if more than one component selected for installtion supports a phased installation.

The local secure store component (LSS) can only be installed in optimized installation mode.

13.53 `install_hostagent`

Enables the installation or update of the SAP host agent.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --install_hostagent[=off]
```

Remarks

The default for this parameter is `--install_hostagent` (on).

13.54 `install_ssh_key`

Installs SSH key to access remote hosts.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --install_ssh_key[=off]
```

Remarks

The default for this parameter is `--install_ssh_key` (on).

13.55 `internal_network`

Specifies the internal subset address in CIDR notation.

Syntax

In the command line, the following syntax is used:

```
--action=install --internal_network=<address>|none
```

Example

The following example specifies the internal network address in CIDR notation.

```
--internal_network=192.168.1.0/24
```

13.56 isc_mode

Specifies the inter-service communication mode.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --isc_mode=standard|ssl
```

Remarks

If SAP HANA XS Advanced Runtime is installed, SSL is used for inter-service communication by default. SSL can be disabled during the installation of the SAP HANA XS Advanced Runtime by passing the parameter `--isc_mode=standard` to the SAP HANA database lifecycle manager.

> **Note**
>
> Switching the inter-service communication mode of an SAP HANA system via the SAP HANA database lifecycle manager (HDBLCM) results in a system restart.

13.57 keep_lss_user

Prevents the Local Secure Store User (`<sid>crypt`) from being removed.

Syntax

In the command line, the following syntax is used:

```
--keep_lss_user[=on]
```

Remarks

The default for this parameter is `--keep_lss_user` (off).

This parameter is available in interactive mode.

13.58 keep_lss_user_group

Prevents the Local Secure Store User Group (`<sid>crypt`) from being removed.

Syntax

In the command line, the following syntax is used:

```
--keep_lss_user_group[=on]
```

Remarks

The default for this parameter is `--keep_lss_user_group` (off).

This parameter is available in interactive mode.

13.59 `keep_user`

Keeps the system administrator user (`<sid>adm`) from the source system to be used in the target system.

Syntax

In the command line, the following syntax is used:

```
--keep_user=[yes|no]
```

or, in short form:

```
-k=[yes|on]
```

Remarks

The default for this parameter is `--keep_user=no`.

This parameter is available in interactive mode.

13.60 `keep_user_home_dir`

Prevents the home directory of the source system administrator user (`<sid>adm`) from being removed.

Syntax

In the command line, the following syntax is used:

```
--keep_user_home_dir[=on]
```

Remarks

The default for this parameter is `--keep_user_home_dir` (off).

This parameter is available in interactive mode.

13.61 `listen_interface`

Specifies the listen interface for the internal network communication.

Syntax

In the command line, the following syntax is used:

```
--listen_interface=[local|global|internal]
```

Options

The following options are available:

`global` - Binds the processes to all interfaces. This option does not require an internal network address entry.

`internal` - Binds the processes to this address only and to all local host interfaces. This option requires an internal network address entry.

`local` - Opens the communication ports for internal usage on the local interfaces. This configuration is only an option for single installations as the server is not reachable from outside. This option does not require an internal network address entry.

Remarks

If you define a value other than `local`, the local interfaces will always be open.

To disable internal network binding when using the global option for `listen_interface`, set `internal_network` to `none`.

This parameter is available in interactive mode.

13.62 `list_systems`

Displays a list of the installed SAP HANA systems on the current host, including the SAP system ID (SID), instance number, version number, and hosts.

Syntax

In the command line, the following syntax is used:

```
--action=[extract_components|install] --list_systems
```

or, in short form:

```
-L
```

13.63 `load_initial_xs_content`

Imports SAP HANA XS advanced runtime content if --import_xs_content=off was selected during installation of the system.

Syntax

In the command line, the following syntax is used:

```
--load_initial_xs_content[=off]
```

Remarks

The default for this parameter is `--load_initial_xs_content`.

13.64 `logbackuppath`

Specifies the location of the log backups.

Syntax

In the command line, the following syntax is used:

```
--logbackuppath=<path>
```

Remarks

This parameter is available in interactive mode.

The path must be specified if the new directory is located on a different physical storage. The contents of the directory must be moved manually to the new location.

13.65 `logpath`

Specifies the path to the log directory of the SAP HANA system.

Syntax

In the command line, the following syntax is used:

```
--action=install --logpath=<path to log directory>
```

Remarks

The default for this parameter is `--logpath=/hana/log/<SID>`.

This parameter is available in interactive mode.

The path must be specified if the new directory is located on a different physical storage. The contents of the directory must be moved manually to the new location.

13.66 `lss_groupid`

Specifies the group ID for the local secure store.

Syntax

In the command line, the following syntax is used:

```
--action=install --lss_groupid=<lssgroupid>
```

Remarks

This parameter is available in interactive mode.

13.67 `lss_inst_path`

Specifies the installation path for local secure store.

Syntax

In the command line, the following syntax is used:

```
--action=install --lss_inst_path=<lssinst_path>
```

Remarks

The default for this parameter is `--lss_inst_path=/lss/shared`.

On a single-host system, the installation directory will be created automatically. On a multiple-host system, the installation path must be created manually before the installation of local secure store and mounted on all hosts.

This parameter is available in interactive mode.

13.68 `lss_trust_unsigned_components`

Instructs the Local Secure Store (LSS) to trust unsigned SAP HANA database, SAP HANA AFLs and SAP HANA server plugins.

Syntax

In the command line, the following syntax is used:

```
--lss_trust_unsigned_components
```

Remarks

When system usage is set to `production` or `test`, installation or update of the Local Secure Store (LSS) will fail if the SAP HANA server, SAP HANA AFLs or SAP HANA server plugins are installed or updated without a signature (`SIGNATURE.SMF`). Specify the `lss_trust_unsigned_components` parameter to trust the unsigned components' binaries.

> ⚠ Caution
>
> By specifying the `lss_trust_unsigned_components` parameter data and log volume encryption keys may be leaked if the server binaries have been tampered with.

Related Information

system_usage [page 281]

13.69 `lss_user_home`

Specifies the user home directory for the local secure store.

Syntax

In the command line, the following syntax is used:

```
--action=install --lss_user_home=<lssuser_home_dir>
```

Remarks

The default for this parameter is `--lss_user_home=/usr/sap/<SID>/lss/home`.

This parameter is available in interactive mode.

13.70 `lss_user_shell`

Specifies the user login shell for the local secure store.

Syntax

In the command line, the following syntax is used:

```
--action=install --lss_user_shell=<lssuser_shell>
```

Remarks

The default for this parameter is `--lss_user_shell=/bin/sh`.

This parameter is available in interactive mode.

13.71 `lss_userid`

Specifies the user ID for the local secure store.

Syntax

In the command line, the following syntax is used:

```
--action=install --lss_userid=<lssuserid>
```

Remarks

The default for this parameter is `--lss_userid=<SID>crypt`.

This parameter is available in interactive mode.

13.72 `max_mem`

Specifies the maximum memory allocation for a new system in MB.

Requirements

This parameter must be used in combination with `restrict_max_mem`. For more information, see Related Information.

Syntax

In the command line, the following syntax is used:

```
--action=install --max_mem=<MB>
```

Remarks

The default for this parameter is `--max_mem=0`.

This parameter is available in interactive mode.

Related Information

restrict_max_mem [page 272]

13.73 `nostart`

Prevents the SAP HANA system from being started.

Syntax

In the command line, the following syntax is used:

```
--nostart
```

13.74 nostart_tenant_db

Prevents the SAP HANA tenant databases from being started.

Syntax

In the command line, the following syntax is used:

```
--nostart_tenant_db
```

13.75 number

Specifies the instance number of the SAP HANA system.

Requirements

The instance number must be a two-digit number between 00 and 97.

Syntax

In the command line, the following syntax is used:

```
--action=install --number=<instance number>
```

or, in short form:

```
-n <instance number>
```

Remarks

The default value for this parameter is the next successive un-used instance number.

This parameter is available in interactive mode.

13.76 org_manager_user

Creates an SAP HANA XS advanced runtime admin user. An admin user can add and manage users, view users, edit organization roles, view the organization quota, and perform other administrative tasks.

Syntax

In the command line, the following syntax is used:

```
--org_manager_user=<admin user>
```

Remarks

The default for this parameter is `--org_manager_user=XSA_ADMIN`.

This parameter is available in interactive mode.

13.77 `org_name`

Sets the initial name of the customer organization, which can then be changed if required. Organizations enable developers to collaborate by sharing resources, services, and applications. Access to the shared resources, services, and applications is controlled by the organization manager.

Syntax

In the command line, the following syntax is used:

```
--org_name=<org_name>
```

Remarks

The default for this parameter is `--org_name=org`.

This parameter is available in interactive mode.

13.78 `overwrite_extract_dir`

Overwrite the existing files in the temporary extraction directory.

Syntax

In the command line, the following syntax is used:

```
--action=extract_components --overwrite_extract_dir
```

Remarks

This parameter is available in interactive mode.

13.79 `pmempath`

Specifies the path to persistent memory volumes.

Syntax

In the command line, the following syntax is used:

```
--action=install --pmempath=<path>
```

Remarks

The default for this parameter is `--pmempath=/hana/pmem/<SID>`.

The parameter cannot be used in combination with the parameter `pmem_mountpoints`.

This parameter is available in interactive mode.

Related Information

use_pmem [page 286]
pmem_mountpoints [page 269]
Persistent Memory
SAP Note 2618154

13.80 `pmem_mountpoints`

Specifies the mountpoints of the persistent memory volumes.

Syntax

In the command line, the following syntax is used:

```
--action=install --pmem_mountpoints=<mountpoint>
```

Remarks

The parameter cannot be used in combination with the parameter `pmempath`.

This parameter is available in interactive mode.

Related Information

use_pmem [page 286]
pmempath [page 268]
Persistent Memory
SAP Note 2618154

13.81 `prepare_update`

Stops the SAP HANA update before software version switch. The update is resumable. For more details about update planning and updating in a two-phase approach, see Related Information.

Syntax

In the command line, the following syntax is used:

```
--action=update --prepare_update
```

Related Information

Prepare an Update for Flexible System Downtime [page 154]

13.82 `prod_space_name`

Sets the initial name of the customer space for the SAP HANA XS advanced runtime. The space name can be changed later, if required. In an organization, spaces enable users to access shared resources that can be used to develop, deploy, and maintain applications.

Syntax

In the command line, the following syntax is used:

```
--prod_space_name=<prod_space_name>
```

Remarks

The default for this parameter is `--prod_space_name=PROD`.

This parameter is available in interactive mode.

13.83 `read_password_from_stdin`

Reads passwords in XML syntax via input stream in batch mode.

Requirements

When using `read_password_from_stdin` in batch mode, the program looks for the necessary passwords from the standard input stream. Therefore, the following two requirements apply:

- Passwords must be provided in XML syntax.

- Batch mode is specified from the command line.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --read_password_from_stdin=xml
```

Remarks

Passwords can be specified in several ways depending on your installation method of choice. For more details, refer to the information about passwords specification.

Example

The following example shows XML data containing the passwords:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- Replace the 3 asterisks with the password -->
<Passwords>
    <root_password><![CDATA[***]]></root_password>
    <sapadm_password><![CDATA[***]]></sapadm_password>
    <master_password><![CDATA[***]]></master_password>
    <password><![CDATA[***]]></password>
    <system_user_password><![CDATA[***]]></system_user_password>
    <lss_user_password><![CDATA[***]]></lss_user_password>
    <lss_backup_password><![CDATA[***]]></lss_backup_password>
    <streaming_cluster_manager_password><![CDATA[***]]></
streaming_cluster_manager_password>
    <ase_user_password><![CDATA[***]]></ase_user_password>
    <org_manager_password><![CDATA[***]]></org_manager_password>
</Passwords>
```

Related Information

Specifying Passwords [page 77]
Tutorial: Installing a Single-Host System with Passwords Read from XML Standard Input Stream [page 224]

13.84 `remote_execution`

Specifies the connectivity method for SAP HANA multiple-host system operations.

Syntax

In the command line, the following syntax is used:

```
--remote_execution=saphostagent|ssh
```

Remarks

The default for this parameter is `--remote_execution=ssh`.

13.85 `removehosts`

Specifies remote hosts to be removed from the SAP HANA system.

Syntax

In the command line, the following syntax is used:

```
removehosts=<hostname1>[,<hostname2>]...
```

Remarks

This parameter is available in interactive mode.

13.86 `repository`

Defines the source path from which the SAP HANA studio repository should be copied.

Syntax

In the command line, the following syntax is used:

```
--repository=<source path>
```

or, in short form:

```
-r <source path>
```

Remarks

If you do not specify this option, the repository contained in the installation kit is copied.

13.87 `restrict_max_mem`

Specifies whether maximum memory allocation is restricted for a new system.

Requirements

This parameter must be used in combination with `max_mem`. For more information, see Related Information.

Syntax

In the command line, the following syntax is used:

```
--restrict_max_mem
```

Remarks

This parameter is available in interactive mode.

Related Information

max_mem [page 266]

13.88 `restrict_tmpfs_global_allocation_limit`

Allows restriction on the overall TMPFS memory usage.

Syntax

In the command line, the following syntax is used:

```
--restrict_tmpfs_global_allocation_limit
```

Remarks

This parameter is available in interactive mode.

13.89 `root_user`

Specifies the root user name.

Syntax

In the command line, the following syntax is used:

```
--root_user=<user name>
```

Remarks

The default for this parameter is `--root_user=root`.

This parameter is available in interactive mode.

13.90 `sapcar_location`

Location of the SAPCAR executable.

Syntax

In the command line, the following syntax is used:

```
--action=extract_components --sapcar_location=<sapcar_location>
```

13.91 sapmnt

Specifies the path to the SAP mount directory, which can be used as a shared directory between multiple hosts.

Syntax

In the command line, the following syntax is used:

```
--sapmnt=<installation path>
```

Remarks

The default for this parameter is `--sapmnt=/hana/shared`.

This parameter is available in interactive mode.

This parameter can be specified once and automatically subsituted throughout all parameter defaults, which use `sapmnt` as part of their default value.

13.92 scope

Performs task on the SAP HANA System (all hosts) or only on the local instance.

Syntax

In the command line, the following syntax is used:

```
--action=update --scope=instance|system
```

Remarks

The default for this parameter is `--scope=system`.

13.93 secure_store

Specifies whether the SAP HANA database should use the standard SSFS secure store or Local Secure Store (LSS).

Syntax

In the command line, the following syntax is used:

```
--secure_store=localsecurestore|ssfs
```

Remarks

The default for this parameter is `--secure_store=localsecurestore`.

Related Information

Activate the Local Secure Store (LSS) [page 197]

13.94 `shell`

Specifies a system administrator login shell. This parameter is relevant only if the operating system administrator (`<sid>adm`) does not exist prior to installation.

Syntax

In the command line, the following syntax is used:

```
--action=install --shell=<admin login shell>
```

Remarks

The default for this parameter is `--shell=/bin/sh`.

This parameter is available in interactive mode.

13.95 `sid`

Specifies a system ID. The SAP system ID (SID) is the identifier for the SAP HANA system.

Requirements

- The ID must be unique throughout your organization and consistent throughout your SAP system installation landscape.
- If you want to install an additional application server instance, make sure that no gateway instance with the same SAP SID exists in your SAP system landscape.
- The ID must consist of exactly three alphanumeric characters. Only uppercase letters are allowed. The first character must be a letter (not a digit).
- The following IDs are reserved and cannot be used: ADD ALL AMD AND ANY ARE ASC AUX AVG BIT CDC COM CON DBA END EPS FOR GET GID IBM INT KEY LOG LPT MAP MAX MIN MON NIX NOT NUL OFF OLD OMS OUT PAD PRN RAW REF ROW SAP SET SGA SHG SID SQL SUM SYS TMP TOP UID USE USR VAR.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --sid=<SID>
```

or, in short form:

```
-s <SID>
```

Remarks

This parameter is available in interactive mode.

This parameter can be specified once and automatically substituted throughout all parameter defaults, which use `sid` as part of their default value.

13.96 `simplified_ui`

Enables simplified installation user interface.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --simplified_ui
```

13.97 `skip_hostagent_calls`

Skips all SAP Host Agent calls.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --skip_hostagent_calls
```

13.98 `skip_modify_sudoers`

Prevents the file `/etc/sudoers.d/sap_hdb_sudoers` from being modified.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --skip_modify_sudoers
```

or, in short form:

```
-M
```

13.99 sld_hostname

Specifies the name of the host where the SLD system is installed.

Syntax

In the command line, the following syntax is used:

```
--sld_hostname=<host name>
```

Remarks

This parameter is available in interactive mode.

13.100 sld_port

Specifies the standard HTTP access port of the SLD.

Syntax

In the command line, the following syntax is used:

```
--sld_port=<port number>
```

Remarks

This parameter is available in interactive mode.

13.101 sld_proxy

Specifies whether a client proxy for the HTTP connection to the SLD server is used.

Syntax

In the command line, the following syntax is used:

```
--action=configure_sld --sld_proxy
```

13.102 sld_proxy_host

Specifies the SLD proxy host.

Syntax

In the command line, the following syntax is used:

```
--sld_proxy_host=<sld proxy host>
```

Remarks

This parameter is available in interactive mode.

13.103 sld_proxy_port

Specifies the SLD proxy port number.

Syntax

In the command line, the following syntax is used:

```
--sld_proxy_port=<sld proxy port>
```

Remarks

This parameter is available in interactive mode.

13.104 sld_proxy_user

Specifies the proxy user.

Syntax

In the command line, the following syntax is used:

```
--sld_proxy_user=<sld proxy user>
```

Remarks

This parameter is available in interactive mode.

13.105 `sld_username`

Specifies the user of the SLD system. It must be a user that already exists on the host where the SLD system is running.

Syntax

In the command line, the following syntax is used:

```
--sld_username=<user name>
```

Remarks

This parameter is available in interactive mode.

13.106 `source_sid`

Specifies the source system ID if the `<sapmnt>/<SID>` is mounted under the new `<SID>`.

Syntax

In the command line, the following syntax is used:

```
--source_sid=<SID>
```

or, in short form:

```
-s <SID>
```

Remarks

This parameter is available in interactive mode.

13.107 `sso_cert`

Single sign-on (SSO) certificate file to authenticate `sapcontrol` and SAP Host Agent admin user.

Syntax

In the command line, the following syntax is used:

```
--action=update --sso_cert=<ssocertificate>
```

13.108 `storage_cfg`

Specifies a location where a `global.ini` is defined. It is possible to set up a storage connector, allowing SAP HANA to use hardware vendor-specific scripts for automated resource allocation and input/output fencing during failover.

Requirements

Resource allocation scripts are dependent on the hardware used. Therefore, only the hardware partners can provide correct scripts.

Syntax

In the command line, the following syntax is used:

```
--action=install --storage_cfg=<directory of the storage configuration>
```

Remarks

Parameter values passed by the SAP HANA database lifecycle manager (HDBLCM) take precedence over corresponding values that are specified in a storage configuration file.

Related Information

Multiple-Host System Concepts [page 85]

13.109 `studio_path`

Specifies the installation path for the SAP HANA studio.

Syntax

In the command line, the following syntax is used:

```
--studio_path=<installation path>
```

Remarks

The default for this parameter is `--studio_path=<sapmnt>/<SID>/hdbstudio`.

13.110 `studio_repository`

Enables the copying of the SAP HANA studio repository. When enabled (default), the SAP HANA studio repository is copied to the location defined by `copy_repository` from the location defined by `repository`.

Syntax

In the command line, the following syntax is used:

```
--studio_repository[=off]
```

Remarks

The default for this parameter is `--studio_repository` (on).

Related Information

copy_repository [page 248]
repository [page 272]

13.111 `system_usage`

Specifies the usage type of the system to be installed. This setting is stored in the `global.ini` file, and can be used to identify the intended usage of the system.

Syntax

In the command line, the following syntax is used:

```
--action=install --system_usage=[production|test|development|
custom]
```

Remarks

The default for this parameter is `--system_usage=custom`.

This parameter is available in interactive mode.

For information about implementing the next steps of the system usage type, see the *SAP HANA Administration Guide*.

Related Information

Configure System Usage Type

13.112 `system_user`

Specifies the system user of the database.

Syntax

In the command line, the following syntax is used:

```
--action=user --system_user=<name>
```

Remarks

The default for this parameter is `--system_user=SYSTEM`.

13.113 `systemdb_use_singledb_user_password`

Use the single-container database user password for the system database user that is created during the conversion to tenant databases.

Syntax

In the command line, the following syntax is used:

```
--action=update --systemdb_use_singledb_user_password
```

13.114 `tar_executable_location`

Location of the TAR executable.

Syntax

In the command line, the following syntax is used:

```
--action=extract_components --
tar_executable_location=<tar_executable_location>
```

13.115 `target_sid`

Specifies the SID for the target system.

Syntax

In the command line, the following syntax is used:

```
--target_sid=<new sid>
```

or, in short form:

```
-S <new sid>
```

Remarks

This parameter is available in interactive mode.

13.116 `tenantmap`

Specifies the tenant mapping to rename a tenant. The parameter has to be specified for each renamed tenant. It can only be used when changing the SID, instance number, or host names of a system.

Syntax

In the command line, the following syntax is used:

```
--tenantmap=<old_tenant_name>=<new_tenant_name> [--
tenantmap=<old_tenant_name2>=<new_tenant_name2>]
```

Example

The following command renames two tenants, tenant1 and tenant2, to tenant3 and tenant4:

```
./hdblcm --action=rename_system --target_sid=DEF --tenantmap
tenant1=tenant3 --tenantmap tenant2=tenant4
```

Related Information

Renaming a System

13.117 `timeouts`

Sets customized timeouts (`start_instance`, `start_service`, `stop_instance`, `stop_service`).

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --
timeouts=<name1>:<sec>[,<name2>:<sec>]...
```

13.118 `tmpfs_global_allocation_limit`

Sets the global allocation limit in MB.

Syntax

In the command line, the following syntax is used:

```
--tmpfs_global_allocation_limit=<tmpfs_global_allocation_limit>
```

Remarks

This parameter is available in interactive mode.

13.119 `tmpfs_mountpoints_path`

Specifies the location where the TMPFS mountpoints will be created. The value must be a valid absolute path.

Syntax

In the command line, the following syntax is used:

```
--tmpfs_mountpoints_path=<path>
```

Remarks

This parameter is available in interactive mode.

13.120 `unrar_executable_location`

Location of the UNRAR executable.

Syntax

In the command line, the following syntax is used:

```
--action=extract_components --
unrar_executable_location=<unrar_executable_location>
```

13.121 `update_component_list`

Updates the list of components detected on the system. Updating the component list may be necessary, if your technical system landscape is not reflected properly inside SAP Solution Manager landscape management database (LMDB).

Syntax

In the command line, the following syntax is used:

```
--action=update_component_list
```

13.122 `update_execution_mode`

Specifies the update mode of hdblcm to be either standard or optimized. If the optimized update mode is selected, the update process will be carried out in a phased approach to minimize system downtime. For more details about the update process, see Related Information.

Syntax

In the command line, the following syntax is used:

`--action=update --update_execution_mode=[standard|optimized]`

Remarks

The default for this parameter is `--update_execution_mode=standard`.

Related Information

Updating [page 165]

13.123 `use_master_password`

Controls the usage of a single master password for all users, created during installation.

Syntax

In the command line, the following syntax is used:

```
--action=install --use_master_password=[yes|no]
```

Remarks

The default for this parameter is `--use_master_password=no`.

13.124 use_pmem

Enables the use of persistent memory.

Syntax

In the command line, the following syntax is used:

```
--action=install --use_pmem
```

Remarks

If you enable this option, you can specify the mount point of your persistent memory volumes during installation.

Related Information

Persistent Memory
pmempath [page 268]
pmem_mountpoints [page 269]
SAP Note 2618154 - SAP HANA Persistent Memory (NVM) Release Information

13.125 userid

Specifies the user ID of the system administrator. This parameter is relevant only if the operating system administrator (<sid>adm) does not exist prior to installation.

Syntax

In the command line, the following syntax is used:

```
--action=install --userid=<ID number>
```

or, in short form:

```
-U <ID number>
```

Remarks

The default value for this parameter is the next successive un-used user ID number.

This parameter is available in interactive mode.

Related Information

Users Created During Installation [page 76]

13.126 `verify_signature`

Verifies the signature of SAP HANA components.

Syntax

In the command line, the following syntax is used:

```
--action=[install|update] --verify_signature
```

Related Information

SAP Note 2577617

13.127 `version`

Displays the version of the SAP HANA lifecycle management program.

Syntax

In the command line, the following syntax is used:

```
--version
```

or, in short form:

```
-v
```

13.128 `vm`

Specifies the path of the Java runtime file. This parameter is only relevant if the SAP HANA studio is selected for installation or update.

Syntax

In the command line, the following syntax is used:

```
--vm=<Java path>
```

Remarks

The default value for this parameter is the Java runtime that is found in the environment variable `PATH`, or the Java runtime specified with the environment variable `JAVA_HOME`.

13.129 `volume_encryption`

Enables data and log encryption.

Syntax

In the command line, the following syntax is used:

```
--volume_encryption[=off]
```

Remarks

The default for this parameter is `--volume_encryption` (on).

This parameter is available in interactive mode.

Related Information

Data and Log Volume Encryption

13.130 `workergroup`

Specifies the worker group of the host. If undefined, the worker group is named "default". If you are using extension node for Business Warehouse, you must name the worker group "worker_dt".

Syntax

In the command line, the following syntax is used:

```
--action=install --workergroup=<worker group>
```

Remarks

This parameter is available in interactive mode.

Related Information

Multiple-Host System Concepts [page 85]
SAP Note 2453736

13.131 xs_app_working_path

Specifies the XS advanced runtime app working path for runtime data of application instances. For best performance, specify a local directory, which is then created on all XS_WORKER hosts.

Syntax

In the command line, the following syntax is used:

```
--xs_app_working_path=<path>
```

Remarks

The default for this parameter is `--xs_app_working_path=<sapmnt>/<SID>/xs/app_working`

13.132 xs_cert_key

Specifies the path to the key of the certificate to be used by XS advanced (`*.key`)

Syntax

In the command line, the following syntax is used:

```
--xs_cert_key=<path>
```

13.133 xs_cert_pem

Specifies the path to the certificate to be used by XS advanced (`*.pem`)

Syntax

In the command line, the following syntax is used:

```
--xs_cert_pem=<path>
```

13.134 xs_components

Specifies the XS advanced runtime components to be installed or updated.

Syntax

In the command line, the following syntax is used:

```
--xs_components=<all>|<comp1>[,<comp2>]...
```

Remarks

The default for this parameter is `--xs_components=xsac_monitoring,xsac_services`.

This parameter is available in interactive mode.

13.135 xs_components_cfg

Specifies the path to the directory containing MTA extension descriptors (`*.mtaext`)

Syntax

In the command line, the following syntax is used:

```
--xs_components_cfg=<path>
```

13.136 xs_components_nostart

Do not start the selected xs advanced components after installation

Syntax

In the command line, the following syntax is used:

```
--xs_components_nostart=all|none|comp1[,comp2]...
```

Remarks

The default for this parameter is `--xs_components=none`.

13.137 xs_customer_space_isolation

Run applications in customer space with a separate OS user

Syntax

In the command line, the following syntax is used:

```
--xs_customer_space_isolation[=off]
```

Remarks

The default for this parameter is `--xs_customer_space_isolation` (on).

This parameter is available in interactive mode.

13.138 xs_customer_space_user_id

OS user ID used for running XS Advanced applications in customer space

Syntax

In the command line, the following syntax is used:

```
--xs_customer_space_user_id=uid
```

Remarks

This parameter is available in interactive mode.

13.139 xs_domain_name

Specifies the domain name of an xs_worker host. The domain name has to resolve to the SAP HANA host which is running the **xscontroller** and **xsuaaserver** service. The default domain can also resolve to a host with a reverse proxy forwarding requests to the SAP HANA host that is running **xscontroller** and **xsuaaserver** services.

Syntax

In the command line, the following syntax is used:

```
--xs_domain_name=domain_name
```

Remarks

This parameter is available in interactive mode.

Alternatively, the domain name can resolve to a host which is not part of the SAP HANA system. On this host a SAP Web Dispatcher must be installed and configured to act as a reverse proxy, forwarding the requests to the xs_worker hosts.

Related Information

SAP Note 2245631
xs_routing_mode [page 292]

13.140 `xs_routing_mode`

Specifies the routing mode to be used for XS advanced runtime installations.

Syntax

In the command line, the following syntax is used:

```
--xs_routing_mode=ports|hostnames
```

Remarks

The default for this parameter is `--xs_routing_mode=ports`.

This parameter is available in interactive mode.

Related Information

SAP Note 2245631
xs_domain_name [page 291]

13.141 `xs_runtime_db_tenant`

The tenant database in which the XS advanced runtime is installed.

Syntax

In the command line, the following syntax is used:

```
--xs_runtime_db_tenant=tenant database
```

Remarks

The default for this parameter is `--xs_runtime_db_tenant=SYSTEMDB`.

This parameter is available in interactive mode.

13.142 xs_sap_space_isolation

Run applications in SAP space with a separate OS user

Syntax

In the command line, the following syntax is used:

```
--xs_sap_space_isolation[=off]
```

Remarks

The default for this parameter is `--xs_sap_space_isolation` (on).

13.143 xs_sap_space_user_id

OS user ID used for running XS advanced runtime applications in SAP space

Syntax

In the command line, the following syntax is used:

```
--xs_sap_space_user_id=uid
```

Remarks

This parameter is available in interactive mode.

13.144 xs_trust_pem

Specifies the path to the trust certificate to be used by XS advanced (`*.pem`)

Syntax

In the command line, the following syntax is used:

```
--xs_trust_pem=<path>
```

13.145 `xs_use_default_tenant`

Install XS advanced in the default tenant database.

Syntax

In the command line, the following syntax is used:

```
--xs_use_default_tenant
```

Remarks

This parameter is available in interactive mode.

14 Important SAP Notes

SAP Notes contain important information that can help you to successfully install, update, administer, and work with an SAP HANA system.

SAP Note Number	Title
1514967	SAP HANA: Central Note
2380229	SAP HANA Platform 2.0 - Central Note
3411133	SAP HANA Platform 2.0 SPS 08 Release Note
3531605	SAP HANA 2 SPS08 Database Revision 080 00
2372809	Guideline for Upgrading a SAP HANA 1.0 to SAP HANA 2.0 System
1948334	SAP HANA Database Update Paths for Maintenance Revisions
2503043	Global temporary ROW table could not be dropped
2378962	SAP HANA 2.0 Revision and Maintenance Strategy
2272550	Remove Host from SAP HANA System
2380291	SAP HANA 2.0 Cockpit Central Release Note
2373065	SAP HANA Runtime Tools 2.0 Release Notes
3514077	SAP Web IDE for SAP HANA 2.0 SPS 08 - Central Release Note
3475293	SAP HANA Cockpit 2.0 SP 17
3385805	SAP HANA Runtime Tools 2.0 SP 16
2000003	FAQ: SAP HANA
2235581	SAP HANA: Supported Operating Systems
1944799	SAP HANA Guidelines for SLES Operating System
2009879	SAP HANA Guidelines for Red Hat Enterprise Linux (RHEL)
2055470	HANA on POWER Planning and Installation Specifics - Central Note

SAP Note Number	Title
2218464	Supported products when running SAP HANA on IBM Power Systems
52505	Support after end of mainstream/extended maintenance
3537501	Sunset of components XMLA for XSC/XSA
1681092	Support for multiple SAP HANA databases on a single SAP HANA appliance
1976729	Application Component Hierarchy for SAP HANA
1661202	Support for multiple applications on SAP HANA
1927949	Standard Behavior for SAP Logon Tickets
1906576	HANA client and server cross-version compatibility
1637145	SAP BW on HANA: Sizing SAP HANA Database
1793345	Sizing for Suite on HANA
2428875	Full-text index creation runs endlessly in Preprocessor service.
2435642	Deprecation of legacy text mining implementation.
2560918	Broken transport of changes in HALM after source system SID renaming
2078425	Troubleshooting note for SAP HANA platform lifecycle management tool hdblcm

Check the current SAP Notes for the various parts of SAP HANA by searching for any of the following application areas:

SAP HANA Native Applications

- *HAN-APP* - SAP HANA Native Applications
- *HAN-APP-DCI* - Please use HAN-APP-IOA
- *HAN-APP-DWS* - SAP HANA Data Warehouse Services
- *HAN-APP-DWS-DDO* - SAP HANA Data Distribution Optimizer
- *HAN-APP-DWS-DLM* - SAP HANA Data Lifecycle Manager
- *HAN-APP-DWS-DSO* - DataStore Object
- *HAN-APP-DWS-DWS* - Data Warehousing Scheduler

SAP HANA Application Services

- *HAN-AS* - SAP HANA Application Services
- *HAN-AS-INA* - SAP HANA InA Tools and Infrastructure

- *HAN-AS-INA-FL* - SAP HANA InA File Loader
- *HAN-AS-INA-UI* - SAP HANA InA Toolkit, Fiori Search UI
- *HAN-AS-RPO* - SAP HANA Repository
- *HAN-AS-RST* - SAP HANA Development Environment REST API
- *HAN-AS-RUL* - SAP HANA Rules Framework
- *HAN-AS-XSA* - SAP HANA XS Basis Applications
- *HAN-AS-XSA-LIB* - Please use HAN-AS-XS
- *HAN-AS-XSA-SHN* - SAP HANA Interactive Education (SHINE Model)
- *HAN-AS-XSA-WF* - HWF (Deprecated)
- *HAN-AS-XSA-TM* - SAP HANA Task management
- *HAN-AS-XS* - SAP HANA Extended Application Services
- *HAN-AS-XS-ADM* - SAP HANA XS Administration
- *HAN-AS-XS-JOB* - SAP HANA XS Scheduled Jobs

SAP HANA Tools - SAP Business Application Studio

- *HAN-BAS* - SAP HANA Tools - SAP Business Application Studio
- *HAN-BAS-DBX* - HANA Database Explorer in Business Application Studio
- *HAN-BAS-EDT* - HANA Editors in Business Application Studio
- *HAN-BAS-EDT-MOD* - Calculation View Editor in Business Application Studio
- *HAN-BAS-EIM* - Flowgrath, RepTasks and other SDA Tools
- *HAN-BAS-EXL* - HANA Explorer
- *HAN-BAS-SA* - SQL Analyzer
- *HAN-BAS-SPT* - HANA Supportability Tools
- *HAN-BAS-TPL* - HANA Templates in Business Application Studio

HANA Cloud Services

- *HAN-CLS* - HANA Cloud Services
- *HAN-CLS-CPT* - HANA Cloud Services Management Tool
- *HAN-CLS-DB* - HANA Database as a Service
- *HAN-CLS-DB-ALI* - HANA Service for Alicloud
- *HAN-CLS-DB-PSA* - PSA instances based on CF
- *HAN-CLS-HC* - HANA Cloud Services HANA Cloud
- *HAN-CLS-HC-ALI* - SAP HANA Cloud on AliCloud (operated by CDC)
- *HAN-CLS-HC-FILES* - SAP HANA Cloud, Data Lake Files service
- *HAN-CLS-HC-HDL* - HANA Cloud HANA Data Lake
- *HAN-CLS-HC-HSA* - HANA Cloud - HANA Streaming Analytics
- *HAN-CLS-HC-SPARK* - SAP HANA Cloud, Apache Spark service
- *HAN-CLS-MIG* - SAP HANA Cloud Migration
- *HAN-CLS-OLP* - The HANA Cloud Operator Launchpad
- *HAN-CLS-SRC* - Enterprise Search as a Service

SAP HANA Cockpit

- *HAN-CPT* - SAP HANA Cockpit

- *HAN-CPT-CPT1* - SAP HANA Cockpit version 1 - based on XSC
- *HAN-CPT-CPT1-ADM* - SAP HANA Cockpit 1 Administration Core
- *HAN-CPT-CPT1-ASE* - SAP HANA Cockpit 1 Accelerator for SAP ASE Administration
- *HAN-CPT-CPT1-BAC* - SAP HANA Cockpit 1 Backup and Recovery
- *HAN-CPT-CPT1-SDS* - SAP HANA Cockpit 1 Smart Data Streaming Administration
- *HAN-CPT-CPT1-SEC* - SAP HANA Cockpit 1 Security
- *HAN-CPT-CPT1-WA* - SAP HANA Cockpit 1 Workload Analyser
- *HAN-CPT-CPT2 SAP* - HANA Cockpit version 2 - based on XSA
- *HAN-CPT-CPT2-ADM* - SAP HANA Cockpit 2 (Administration Core)
- *HAN-CPT-CPT2-ASE* - SAP HANA Cockpit 2 (Accelerator for SAP ASE Administratio
- *HAN-CPT-CPT2-BAC* - SAP HANA Cockpit 2 Backup and Recovery
- *HAN-CPT-CPT2-CNR* - SAP HANA Cockpit 2 (Capture and Replay)
- *HAN-CPT-CPT2-DBX* - SAP HANA Cockpit 2 Database Explorer
- *HAN-CPT-CPT2-DYT* - SAP HANA Cockpit 2 (Dynamic Tiering Administration)
- *HAN-CPT-CPT2-EWA* - SAP HANA Cockpit Early Watch Alert Support
- *HAN-CPT-CPT2-LA* - SAP HANA Cockpit Landscape Administration
- *HAN-CPT-CPT2-MDC* - SAP HANA Cockpit 2 (MDC Tenant Handling)
- *HAN-CPT-CPT2-MGC* - SAP HANA Cockpit Manager
- *HAN-CPT-CPT2-PM* - SAP HANA Cockpit 2 (Performance Monitoring)
- *HAN-CPT-CPT2-REC* - SAP HANA Cockpit 2 (Recommendation Application)
- *HAN-CPT-CPT2-SA* - SAP HANA Cockpit 2 SQL Analyzer
- *HAN-CPT-CPT2-SDA* - SAP HANA Cockpit 2 (Smart Data Access)
- *HAN-CPT-CPT2-SDI* - SAP HANA Cockpit 2 (Smart Data Integration)
- *HAN-CPT-CPT2-SDS* - SAP HANA Cockpit 2 (Smart Data Streaming Administration)
- *HAN-CPT-CPT2-SEC* - SAP HANA Cockpit 2 (Security)
- *HAN-CPT-CPT2-SR* - SAP HANA Cockpit 2 (System Replication)
- *HAN-CPT-CPT2-TEL* - HANA Express Telemetry
- *HAN-CPT-CPT2-WA* - SAP HANA Cockpit 2 (Workload Analyzer)
- *HAN-CPT-DP* - Please use HAN-DP-SDI

HAN-DB - SAP HANA Database

- *HAN-DB* - SAP HANA Database
- *HAN-DB-AFL* - Appl. Function Library - SAP Note 2198403 for subcomponents
- *HAN-DB-AFL-APL* - Automated Predictive Library in HANA Cloud
- *HAN-DB-AFL-DQ* - SAP HANA Data Quality Library
- *HAN-DB-AFL-EML* - SAP HANA External Machine Learning Library
- *HAN-DB-AFL-GEN* - SAP HANA AFL Shipment and general AFL topics
- *HAN-DB-AFL-HIE* - SAP HANA AFL Hierarchies
- *HAN-DB-AFL-PAL* - SAP HANA Predictive Analysis Library
- *HAN-DB-AFL-POS* - SAP HANA On-Shelf Availability
- *HAN-DB-AFL-SAL* - SAP HANA Self Service Analytics Library
- *HAN-DB-AFL-SCA* - SAP HANA Supply Chain Algorithm Library

- *HAN-DB-AFL-SOP* - SAP HANA Sales and Operations Planning
- *HAN-DB-AFL-TEC* - SAP HANA AFL Technology and SDK
- *HAN-DB-AFL-UDF* - SAP HANA Unified Demand Forecast
- *HAN-DB-AFL-VCH* - Variant Configuration Library (VCH AFL)
- *HAN-DB-ANO* - SAP HANA Data Anonymization
- *HAN-DB-BAC* - SAP HANA Backup and Recovery
- *HAN-DB-CDS* - SAP HANA Activation of HDBDD-files (CDS Definitions)
- *HAN-DB-CLI* - SAP HANA Clients (JDBC, ODBC)
- *HAN-DB-CLI* -MLAPI HANA machine learning API
- *HAN-DB-DI* - SAP HANA DI (HDI): diserver, HDI plugins and client libs
- *HAN-DB-ENG* - SAP HANA DB Engines
- *HAN-DB-ENG-BW* - SAP HANA BW Engine
- *HAN-DB-ENG-ESH* - SAP HANA Enterprise Search Engine
- *HAN-DB-ENG-GPH* - SAP HANA Graph Engine
- *HAN-DB-ENG-IM* - Please use HAN-DB-SDQ
- *HAN-DB-ENG-JSON* - SAP HANA JSON Document Store
- *HAN-DB-ENG-MDS* - SAP HANA Multidimensional Services MDS / InA
- *HAN-DB-ENG-PLE* - SAP HANA Planning Engine
- *HAN-DB-ENG-SPA* - SAP HANA Spatial Engine
- *HAN-DB-ENG-SPA-E* - SAP HANA Spatial - Esri Geodatabase
- *HAN-DB-ENG-TRX* - TREX API for SAP HANA Database
- *HAN-DB-ENG-TXT* - SAP HANA Text Engine
- *HAN-DB-EPM* - SAP HANA Planning and Simulation Platform
- *HAN-DB-EPM-PLT* - SAP HANA EPM Platform
- *HAN-DB-EPM-XSL* - SAP HANA EPM XSJS library
- *HAN-DB-HA* - SAP HANA High Availability (System Replication, DR, etc.)
- *HAN-DB-MDX* - SAP HANA MDX Engine
- *HAN-DB-MON* - SAP HANA Monitoring
- *HAN-DB-NSE* - SAP HANA Native Storage Extension
- *HAN-DB-PER* - SAP HANA Database Persistence
- *HAN-DB-PERF* - SAP HANA Database Performance
- *HAN-DB-R* - SAP HANA Integration with R
- *HAN-DB-SCR* - SAP HANA SQL Script
- *HAN-DB-SDA* - SAP HANA Smart Data Access
- *HAN-DB-SDQ* - Information Mgmt Platform - smart data quality
- *HAN-DB-SEC* - SAP HANA Security and User Management

HAN-DP - Data Provisioning - please select one of the subcomponents

- *HAN-DP-BC* - SAP HANA Blockchain Adapter
- *HAN-DP-DXC* - SAP HANA Direct Extractor Connector
- *HAN-DP-ESS* - SAP HANA Enterprise Semantic Services (ESS)
- *HAN-DP-LTR* - Use CA-LT-SLT instead

- *HAN-DP-SDI* - SAP HANA smart data integration (SDI)
- *HAN-DP-SDI-DS* - SAP Datasphere specific SDI issues

HAN-DYT - SAP HANA Dynamic Tiering

- *HAN-DYT* - SAP HANA Dynamic Tiering

HAN-HCO - HANA Hardware and Cloud Optimization Tool

- *HAN-HCO* - HANA Hardware and Cloud Optimization Tool

HAN-LM - SAP HANA Lifecycle Management

- *HAN-LM* - SAP HANA Lifecycle Management
- *HAN-LM-APP* - SAP HANA Application Lifecycle Management
- *HAN-LM-INS* - SAP HANA Installation
- *HAN-LM-INS-DB* - Installation of HANA Database
- *HAN-LM-INS-SAP* - Installation of SAP Systems on HANA
- *HAN-LM-PLT* - SAP HANA Platform Lifecycle Management
- *HAN-LM-UPG* - SAP HANA Upgrade
- *HAN-LM-UPG-DB* - Upgrade of HANA Database
- *HAN-LM-UPG-SAP* - Upgrade of SAP Systems on HANA

HAN-SDS - SAP HANA Smart Data Streaming

- *HAN-SDS* - SAP HANA Smart Data Streaming

HAN-STD - SAP HANA Studio (Eclipse)

- *HAN-STD* - SAP HANA Studio (Eclipse)
- *HAN-STD-ADM* - SAP HANA Studio (Eclipse) Tooling
- *HAN-STD-ADM-BAC* - SAP HANA Studio Backup and Recovery
- *HAN-STD-ADM-DBA* - SAP HANA Studio (Eclipse) Admin Tooling
- *HAN-STD-ADM-PVZ* - SAP HANA Plan Visualizer
- *HAN-STD-ADM-SEC* - SAP HANA Security and User Management (Studio)
- *HAN-STD-DEV* - SAP HANA Development Tools
- *HAN-STD-DEV-DP* - SAP HANA Data Provisioning Modeler
- *HAN-STD-DEV-EPM* - SAP HANA EPM Modeler
- *HAN-STD-DEV-MOD* - SAP HANA Analytical Modeling
- *HAN-STD-DEV-MOD-CLT* - SAP HANA Analytical Modeling Client
- *HAN-STD-DEV-MOD-SRV* - SAP HANA Analytical Modeling - Server Component
- *HAN-STD-DEV-REF* - SAP HANA Tools for Where-used, Refactoring and Mass Copy
- *HAN-STD-DEV-RUL* - SAP HANA Rules Editor
- *HAN-STD-DEV-SCR* - SAP HANA SQL Script Editor/Debugger
- *HAN-STD-DEV-TP* - SAP HANA Team Provider
- *HAN-STD-DEV-TP-CM* - SAP HANA Development Change Management
- *HAN-STD-DEV-UIS* - SAP HANA UI Integration Services
- *HAN-STD-DEV-UIS-FLP* - SAP HANA UI Integration Services

- *HAN-STD-DEV-XS* - SAP HANA XS Editors and Wizards

HAN-WDE - SAP Web IDE for Hana

- *HAN-WDE* - SAP Web IDE for Hana
- *HAN-WDE* - SAP Web IDE for Hana
- *HAN-WDE-BLD* - SAP Web IDE for Hana building applications
- *HAN-WDE-BLD-HDB* - SAP Web IDE for Hana HDB Build
- *HAN-WDE-DBX* - Database Explorer in Web IDE for SAP HANA
- *HAN-WDE-DOC* - SAP Web IDE for Hana documentation
- *HAN-WDE-EDT* - SAP Web IDE for Hana text editors
- *HAN-WDE-EDT-CDS* - SAP Web IDE for Hana editor for Core Data Services
- *HAN-WDE-EDT-GCDS* - Graphical Editor for HANA CDS
- *HAN-WDE-EDT-JAVA* - Java Support
- *HAN-WDE-EDT-MOD* - SAP Web IDE editor for HANA Analytical Modeling
- *HAN-WDE-EDT-NJS* - Node.js Tools
- *HAN-WDE-EIM* - Flowgraph, RepTasks and other SDA Tools
- *HAN-WDE-FPM* - SAP Web IDE for Hana feature management
- *HAN-WDE-GIT* - SAP Web IDE for Hana GIT
- *HAN-WDE-INS* - SAP Web IDE for Hana Installation
- *HAN-WDE-MTA* - SAP Web IDE for HANA Multi Targeted Application
- *HAN-WDE-PLF* - SAP Web IDE for Hana platform
- *HAN-WDE-RTT* - SAP Web IDE for Hana Runtime and SQL Tools
- *HAN-WDE-RUN* - SAP Web IDE for Hana running applications
- *HAN-WDE-RUN-UI* - SAP Web IDE for Hana - Run web and SAP Fiori applications
- *HAN-WDE-SA* - SAP HANA SQL Analyzer
- *HAN-WDE-SDS* - Smart Data Streaming Tools
- *HAN-WDE-TPL* - SAP Web IDE for Hana Project creation, template and wizards
- *HAN-WDE-TXT* - Text Analysis Web IDE extensions
- *HAN-WDE-XSC* - Old SAP HANA Web IDE
- *HAN-WDE-XSC-EIM* - Flowgraph, RepTasks and other SDA Tools
- *HAN-WDE-XSC-MOD* - Modeling
- *HAN-WDE-XSC-PVZ* - Performance Visualization Plugin

SAP HANA XS Advanced

- *BC* - Basis Components
- *BC-XS* - SAP HANA XS Advanced
- *BC-XS-BLDP* - Buildpacks (Java, Node.js, Python)
  Replaces the deprecated components `BC-XS-JAV` (Java), `BC-XS-JS` (Node.js), and `BC-XS-PY` (Python) in the XS advanced environment.

  > **Note**
  >
  > In the Cloud Foundry environment, the deprecated components `BC-XS-JS`, `BC-XS-PY`, and `BC-XS-JAV` are replaced by `BC-CP-CF-BLDP`.

- *BC-XS-JS* - Javascript runtime

  > **Note**
  >
  > Deprecated; replaced by `BC-XS-BLDP` in the XS advanced environment and by `BC-CP-CF-BLDP` in the Cloud Foundry environment.

- *BC-XS-PY* - Python Runtime

  > **Note**
  >
  > Deprecated; replaced by `BC-XS-BLDP` in the XS advanced environment and by `BC-CP-CF-BLDP` in the Cloud Foundry environment.

- *BC-XS-JAV* - Java Runtime

  > **Note**
  >
  > Deprecated; replaced by `BC-XS-BLDP` in the XS advanced environment and by `BC-CP-CF-BLDP` in the Cloud Foundry environment.

- *BC-XS-ADM* - SAP HANA XS Advanced Administration Tools
- *BC-XS-APR* - XS Advanced Application Router
- *BC-XS-ASYNCJS* - Successor for XSJS with node.js 16 or higher is async-xsjs
- *BC-XS-CDX* - SAP Cloud application programming model (CAP)
- *BC-XS-CDX-COR* - SAP CAP– Compiler and CDS language
- *BC-XS-CDX-JAV* - SAP CAP – Java runtime
- *BC-XS-CDX-JAV-V1* - Java Runtime Version 1
- *BC-XS-CDX-JAV-V2* - Java Runtime Version 2
- *BC-XS-CDX-NJS* - SAP CAP – Node.js runtime
- *BC-XS-CDX-TLS* - SAP CAP – Tools, IDEs, build, deployment
- *BC-XS-RT* - OP Runtime / XS Advanced Controller
- *BC-XS-SEC* - UAA and Security for the SAP HANA XS Advanced Engine
- *BC-XS-SL* - Software Logistics
- *BC-XS-SL-DS* - SAP HANA XS Advanced MTA Deployment Service
- *BC-XS-SL-PI* - SAP HANA XS Advanced Product Installer
- *BC-XS-SRV* - Services
- *BC-XS-SRV-ADT* - Audit Log Service for XS advanced
- *BC-XS-SRV-HSB* - SAP HANA Service Broker
- *BC-XS-SRV-JBS* - Job Scheduler
- *BC-XS-SRV-MESS* - Component for XS advanced Messaging Service
- *BC-XS-SRV-ODT* - OData Service
- *BC-XS-SRV-ODT-JS* - OData Node.js (for XSOData XS Classic, use HAN-AS-XS)
- *BC-XS-SRV-PTL* - SAP HANA XS Advanced Portal Services (for Fiori Launchpad)
- *BC-XS-TLS* - Tools
- *BC-XS-TLS-MIG* - XS Classic to XS Advanced Migration Assistant tooling

SAP HANA Database (CCMS, Porting and DB Interface)

- *BC-DB-HDB* Use HAN-DB*. Here CCMS, Porting, DB Interface issues only
- *BC-DB-HDB-AFL* Please use HAN-DB-AFL and its subcomponents
- *BC-DB-HDB-BAC* Please use HAN-DB-BAC - SAP HANA Backup and Recovery
- *BC-DB-HDB-CCM* CCMS / Database Monitors for SAP HANA
- *BC-DB-HDB-CLI* Please use HAN-DB-CLI - SAP HANA Clients (JDBC, ODBC)
- *BC-DB-HDB-DBA* Please use HAN-STD-ADM-DBA
- *BC-DB-HDB-DXC* Please use HAN-DP-DXC
- *BC-DB-HDB-ENG* Please use HAN-DB-ENG
- *BC-DB-HDB-PFW* SAP HANA Parallelization Framework
- *BC-DB-HDB-POR* Dictionary and Consistency checks for SAP HANA
- *BC-DB-HDB-SYS* Database Interface/DBMS for SAP HANA

End User Clients

- *BI-BIP* - Business intelligence platform (formerly known as BOE)
- *BI-BIP-CMC* - Central Management Console (CMC)
- *BI-BIP-CRS* - SAP Crystal Reports Server
- *BI-BIP-IDT* - Information design tool
- *BI-RA-AO-XLA* - MS Excel Add-In
- *BI-RA-CR* - SAP Crystal Reports
- *BI-RA-EXP* - SAP BusinessObjects Explorer
- *BI-RA-WBI* - Web Intelligence
- *BI-RA-XL* - Dashboard Designer

The search also supports using the wildcard asterisk (*), so you can, for example, also search for BC-DB-HDB* or similar and you will get results for all subcomponents.

Reporting Incidents

If you encounter any problems with the software, report an incident at http://support.sap.com/incident.

In addition, the Customer Interaction Center (CIC) is available 24 x 7 in every region to help you resolve any issues you may run into (https://support.sap.com/contactus).

The CIC requires a valid S-user number.

When reporting an incident, you can choose from the above list of components for the relevant software part.

Important Disclaimer for Features in SAP HANA

For information about the capabilities available for your license and installation scenario, refer to the Feature Scope Description for SAP HANA.

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.
About the icons:

- Links with the icon : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
  - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
  - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon : You are leaving the documentation for that particular SAP product or service and are entering an SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.
The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Bias-Free Language

SAP supports a culture of diversity and inclusion. Whenever possible, we use unbiased language in our documentation to refer to people of all cultures, ethnicities, genders, and abilities.

www.sap.com/contactsap