PUBLIC

SAP HANA Platform SPS 12
Document Version: 1.2 – 2018-01-24

SAP HANA Master Guide

Content

1 SAP HANA Master Guide

This guide is the entry point for planning the installation of your SAP HANA system landscape.

2 Getting Started with SAP HANA

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

2.1 About This Document

This Master Guide is the central starting point for the technical implementation of SAP HANA.

The Master Guide provides the following information about SAP HANA:

- Overview
- Architecture
- Software components
- Deployment scenarios

2.2 What Is SAP HANA?

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

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

3 SAP HANA Use Cases

SAP HANA can be used in a variety of use cases and scenarios.

3.1 SAP HANA as Primary Persistence for SAP NetWeaver-Based Applications

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

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

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

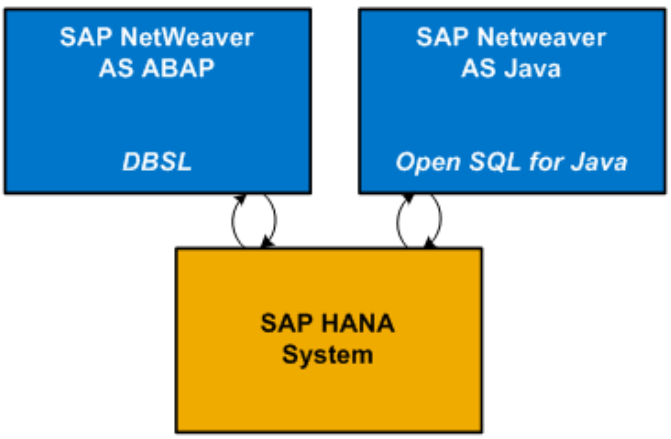

SAP HANA and SAP NetWeaver AS ABAP/Java on one Server

SAP HANA and SAP NetWeaver AS ABAP or SAP NetWeaver AS Java deployed on one server is a multi-component, resource-optimized and cost-optimized deployment approach.

For more information, see *SAP Note 1953429 - SAP HANA and SAP NetWeaver AS ABAP on one Server* and *SAP Note 2043509 - SAP HANA and SAP NetWeaver Java on a Single Host*.

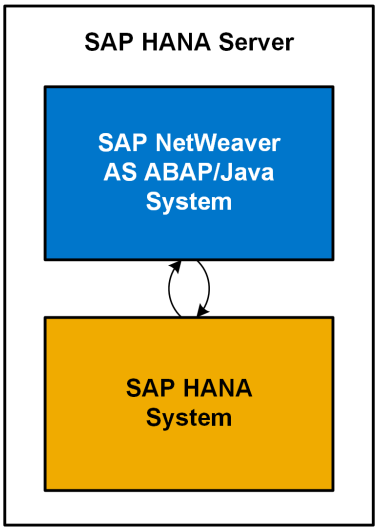

Related Information

SAP Note 1953429 - SAP HANA and SAP NetWeaver AS ABAP on one Server
SAP Note 2043509 - SAP HANA and SAP NetWeaver Java on a Single Host
End-to-End Implementation Roadmap for SAP NetWeaver AS ABAP on SAP HANA
End-to-End Implementation Roadmap for SAP NetWeaver AS Java on SAP HANA
End-to-End Implementation Roadmap for SAP BW powered by SAP HANA

3.2 SAP HANA as Data Mart

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

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

Real-time Operational Data Marts with an SAP Business Suite System

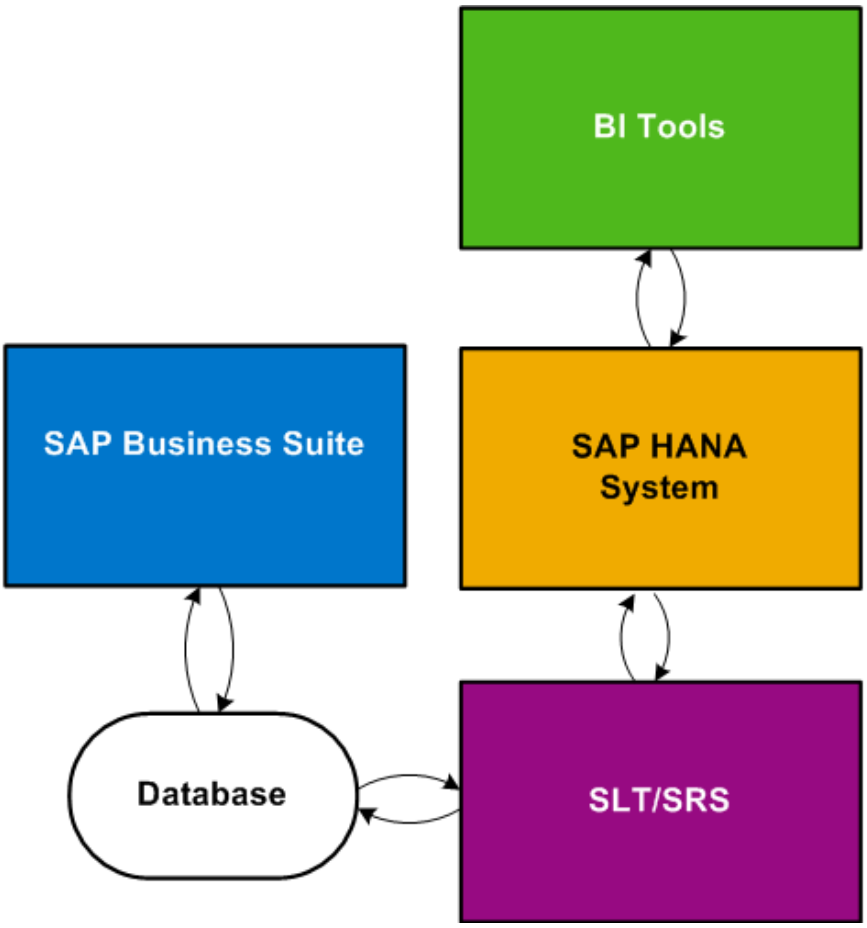

SAP Business Suite is a source system for operational data marts in SAP HANA. The SAP Landscape Transformation Replication Server (SLT) is an SAP NetWeaver ABAP-based tool that provides real-time data replication. In addition, a log-based SAP Replication Server (SRS) can also be used to provide real-time data replication for an SAP Business Suite system.

Real-time Operational Data Marts with a non-SAP OLTP System

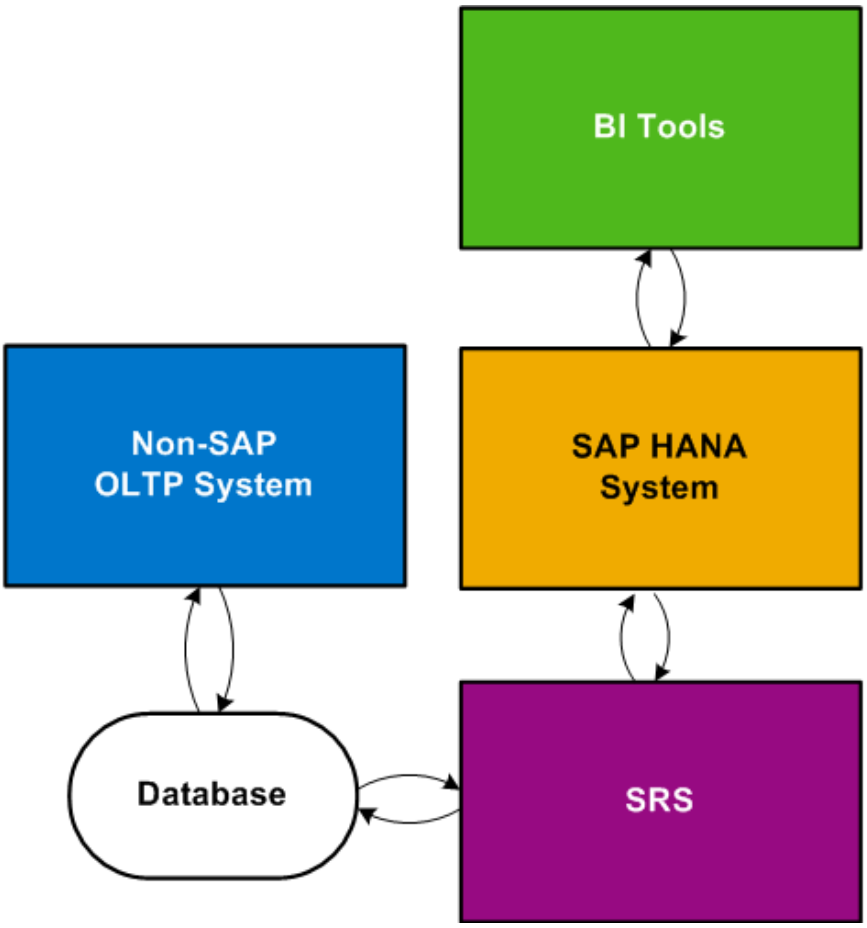

A non-SAP transactional source system is used as a basis for real-time operational data marts in SAP HANA. SAP Replication Server (SRS) refers to the SAP Replication Server application, which is a tool which provides real-time data replication.

Agile Data Marts

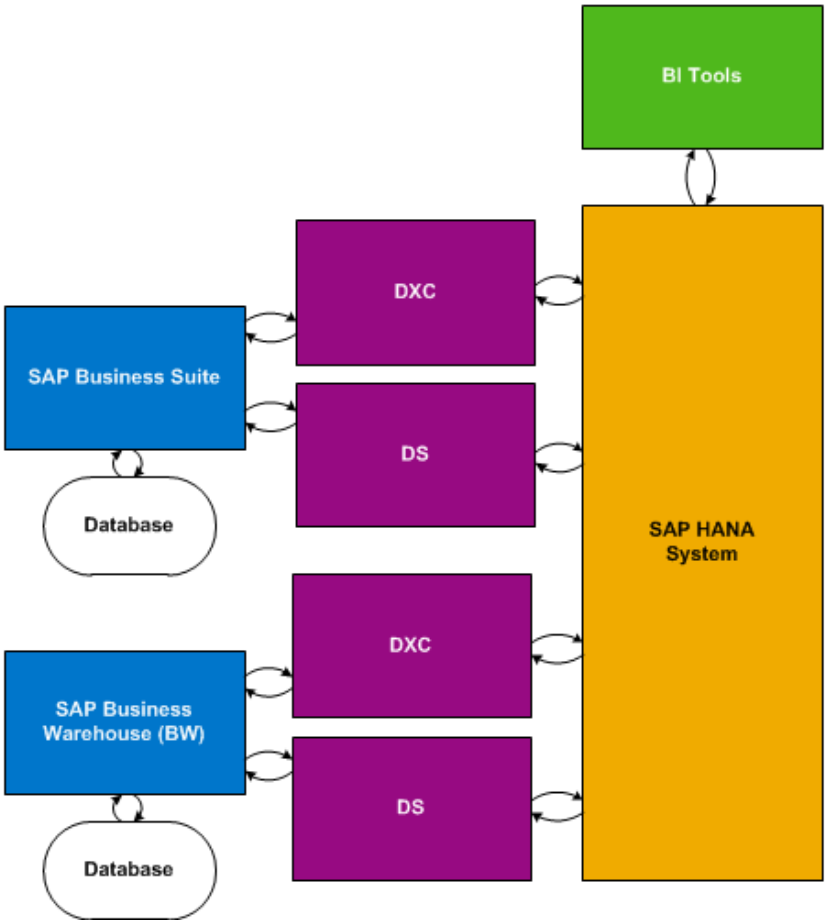

Agile data marts are a type of data mart that offer analytics and reporting on data acquired from a transactional system. When deployed in SAP HANA, they may offer advantages of flexibility when compared to taking a more comprehensive approach to organizational information management, such as deploying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

3.3 SAP HANA-Based Accelerators

SAP HANA-based accelerators are types of applications or scenarios that extend the capabilities of business processes in SAP Business Suite systems by leveraging the performance and scalability advantages that SAP HANA provides. This is implemented by performing operations for certain parts of computing operations of designated business processes or reports.

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

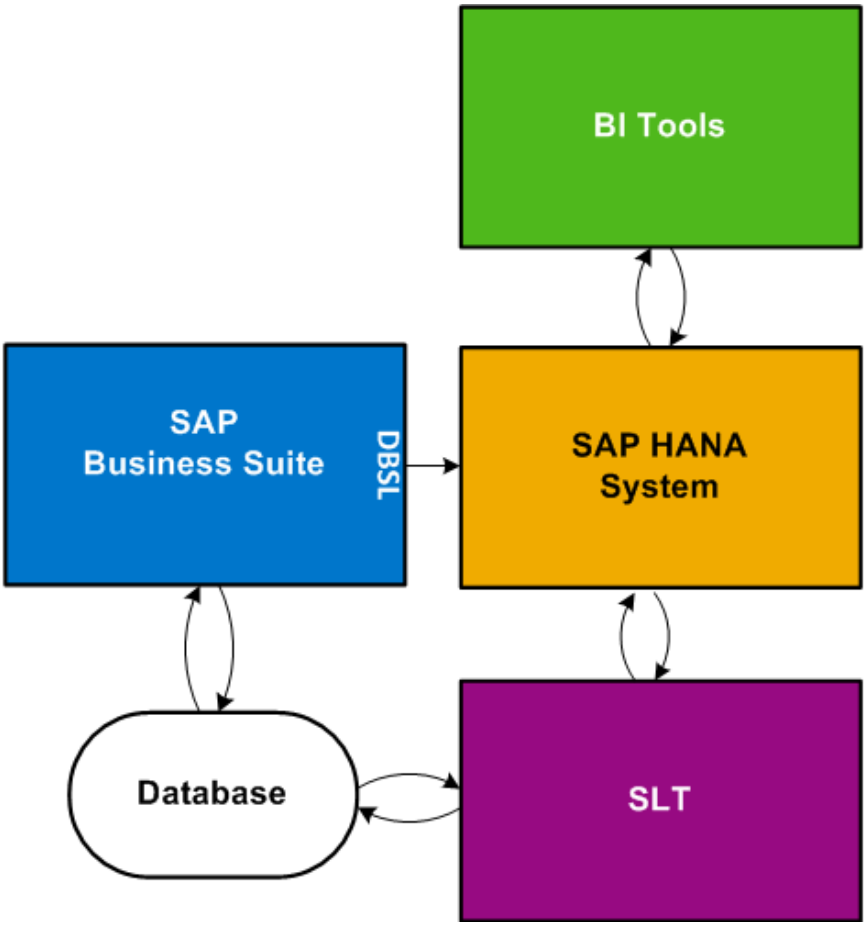
BI Tools
SAP
Business Suite
DBSL
SAP HANA
System
Database
SLT

3.4 SAP HANA Data Provisioning

In-memory reporting and analysis of business data require data provisioning from a source system to the SAP HANA database.

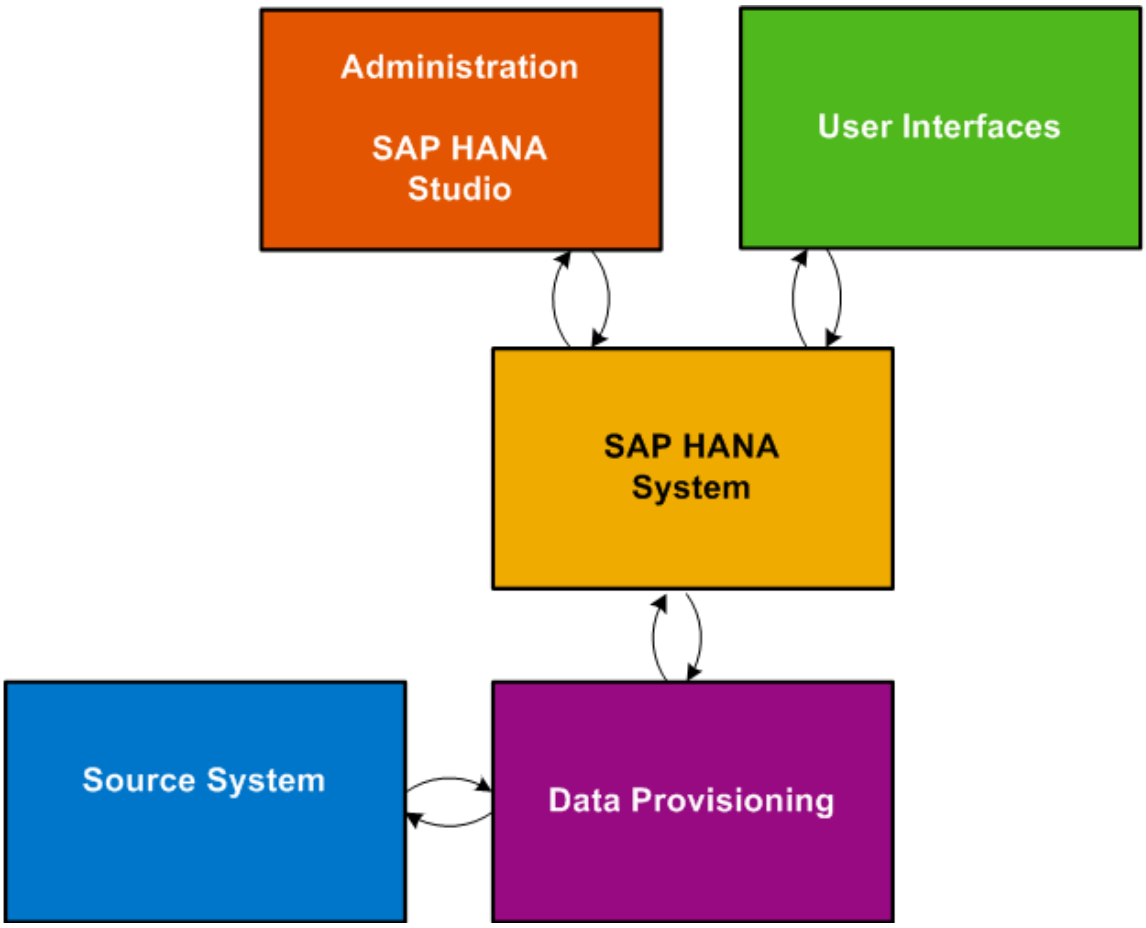

The figure above focuses on the task of loading business data from a source system to the SAP HANA database.

The methods for performing data replication are shown in the figure above. The main components involved in all replication scenarios are:

- SAP HANA, consisting of the SAP HANA database and SAP HANA studio, which is an administration tool. User interfaces, such as SAP BusinessObjects Dashboards or Web Intelligence, are not part of SAP HANA.
- Source system
- Software components supporting the data replication

> **Note**
>
> For more information about the availability of the data replication methods, see the *Product Availability Matrix* on *SAP Service Marketplace*.

Related Information

Product Availability Matrix (PAM)

3.4.1 Trigger-Based Replication (SAP Landscape Transformation Replication Server)

The Trigger-Based Replication method uses the SAP Landscape Transformation (LT) Replication Server component to pass data from the source system to the SAP HANA database target system.

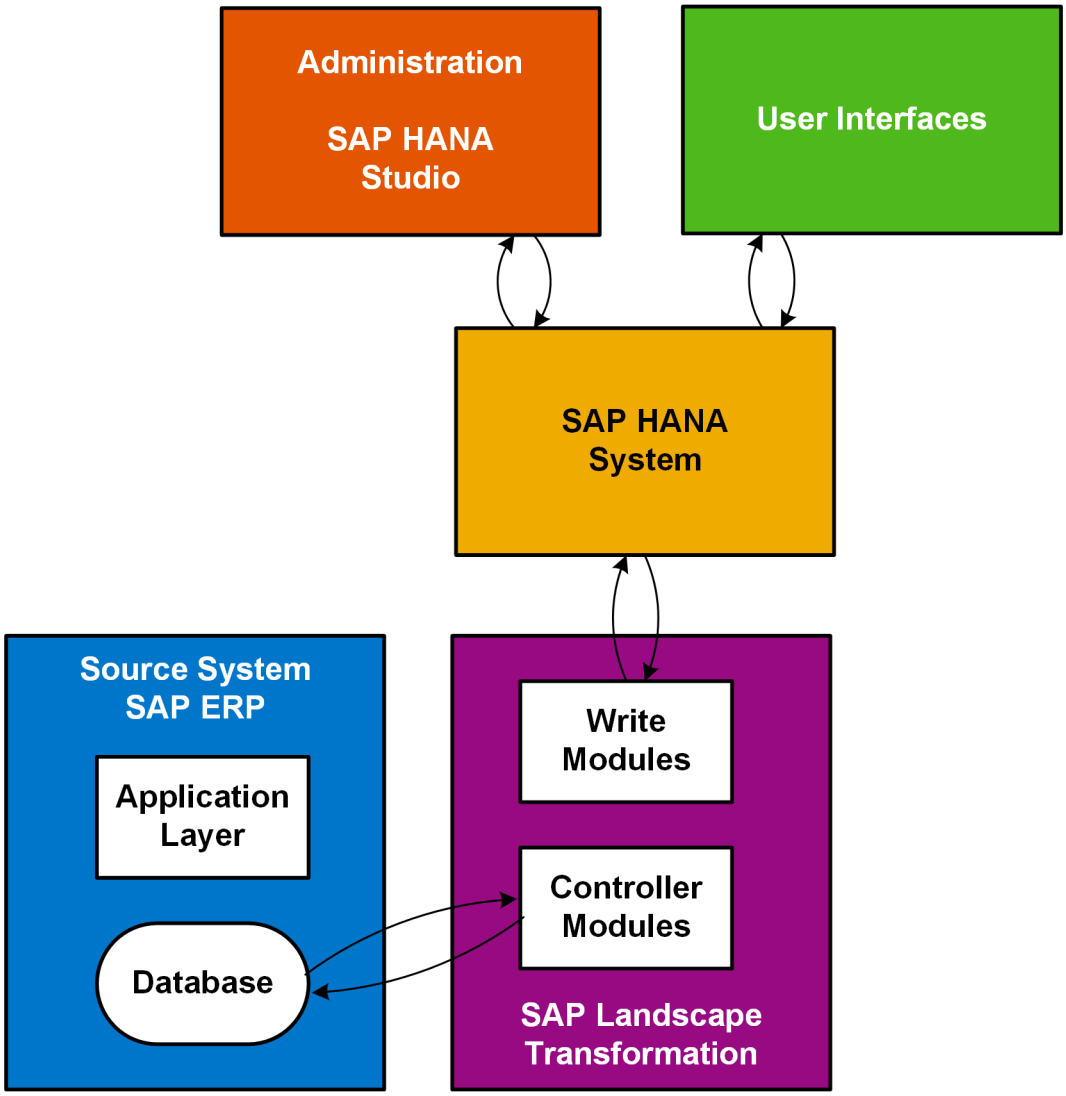

Initial Load and Simultaneous Delta Replication

The initial load of business data is initiated using the SAP HANA studio. The initial load message is sent from the SAP HANA system to the SLT system, which in turn passes the initialization message to the ERP system. Furthermore, the SLT system initiates the set-up of replication log tables in the database of the ERP system for

each table to be replicated. After the transaction tables are completed, the SLT system begins a multi-threaded replication of data to the target system, which enables high speed data transfer.

The initial load of data can be executed while the source system is active. The system load that this process causes can be controlled by adjusting the number of worker threads performing the initial replication.

In parallel to the initial load, by means of database-specific triggers, the SLT system begins detecting any data changes that occur while the initial load process is running. These changes are already recorded in logging tables during the initial load phase and are propagated during the replication phase to the target SAP HANA system after the initial load has been completed. The multi-version concurrency control (MVCC) of the SAP HANA database prevents issues that might be caused by the overlapping of the initial load process and new database transactions

Continuous Delta Replication After Initial Load

After the initial load process has completed, the SLT system continues to monitor the transaction tables in the ERP system, and replicates data changes in the source system to the SAP HANA system in near real time.

Required Software Components

This replication method requires the following component:

- SAP Landscape Transformation: this controls the entire replication process by triggering the initial load and coordinating the delta replication.

Installation considerations

The SLT system can be installed in the ways shown below. You can select between these options depending on your current system landscape and the software versions in your landscape:

- Installation on your ERP system
- Installation on a standalone SAP system (recommended setup)

3.4.2 SAP HANA Direct Extractor Connection (DXC)

The SAP HANA Direct Extractor Connection (DXC) provides SAP HANA with out-of-the-box foundational data models based on SAP Business Suite entities, and is a data acquisition method as well.

Customer projects may face significant complexity in modeling entities in SAP Business Suite systems. In many cases, data from different areas in SAP Business Suite systems requires application logic to appropriately represent the state of business documents. SAP Business Content DataSource Extractors have been available for many years as a basis for data modeling and data acquisition for SAP Business Warehouse; now with DXC, these SAP Business Content DataSource Extractors are available to deliver data directly to SAP HANA.

DXC is a batch-driven data acquisition technique; it should be considered as a form of extraction, transformation and load although its transformation capabilities are limited to user exit for extraction.

A key point about DXC is that in many use cases, batch-driven data acquisition at certain intervals is sufficient (for example, every 15 minutes).

Overview of the DXC Rationale

- Leverage pre-existing foundational data models of SAP Business Suite entities for use in SAP HANA data mart scenarios:
  - Significantly reduces complexity of data modeling tasks in SAP HANA
  - Speeds up timelines for SAP HANA implementation projects
- Provide semantically rich data from SAP Business Suite to SAP HANA:
  - Ensures that data appropriately represents the state of business documents from ERP
  - Application logic to give the data the appropriate contextual meaning is already built into many extractors
- Simplicity/Low TCO:
  - Re-uses existing proprietary extraction, transformation, and load mechanism built into SAP Business Suite systems over a simple http(s) connection to SAP HANA
  - No additional server or application needed in system landscape
- Change data capture (delta handling):
  - Efficient data acquisition – only bring new or changed data into SAP HANA
  - DXC provides a mechanism to properly handle data from all delta processing types

Default DXC Configuration for SAP Business Suite

DXC is available in different configurations based on the SAP Business Suite system:

- The default configuration is available for SAP Business Suite systems based on SAP NetWeaver 7.0 or higher – such as ECC 6.0.
- The alternative configuration is available for SAP Business Suite systems based on releases lower than SAP NetWeaver 7.0 – such as SAP ERP 4.6, for example.

An SAP Business Suite system is based on SAP NetWeaver. As of SAP NetWeaver version 7.0, SAP Business Warehouse (BW) is part of SAP NetWeaver itself, which means a BW system exists inside SAP Business Suite systems such as ERP (ECC 6.0 or higher). This BW system is referred to as an "embedded BW system". Typically, this embedded BW system inside SAP Business Suite systems is actually not utilized, since most customers who run BW have it installed on a separate server, and they rely on that one. With the default DXC configuration, we utilize the scheduling and monitoring features of this embedded BW system, but do not utilize its other aspects such as storing data, data warehousing, or reporting / BI. DXC extraction processing essentially bypasses the normal dataflow, and instead sends data to SAP HANA. The following illustration depicts the default configuration of DXC.

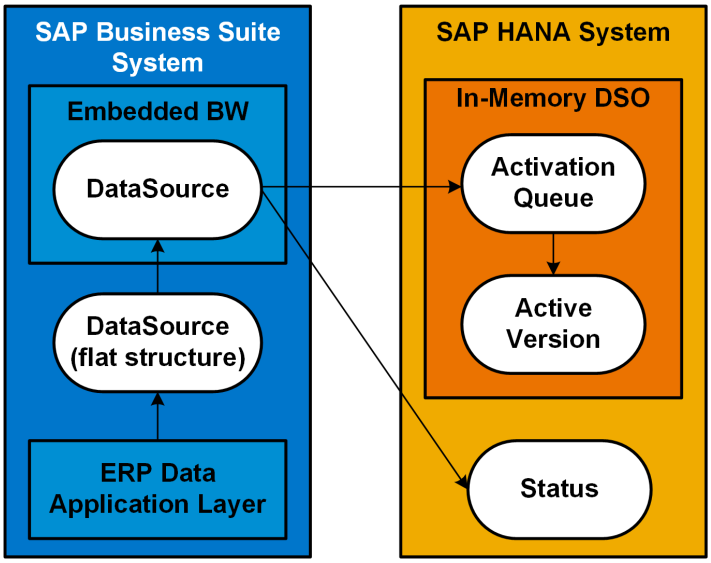

An In-Memory DataStore Object (IMDSO) is generated in SAP HANA, which directly corresponds to the structure of the DataSource you are working with. This IMDSO consists of several tables and an activation mechanism. The active data table of the IMDSO can be utilized as a basis for building data models in SAP HANA (attribute views, analytical views, and calculation views).

Data is transferred from the source SAP Business Suite system using an HTTP connection. Generally, the extraction and load process is virtually the same as when extracting and loading SAP Business Warehouse – you rely on InfoPackage scheduling, the data load monitor, process chains, etc. – which are all well-known from operating SAP Business Warehouse.

> **Note**
>
> DXC does not require BW on SAP HANA. Also with DXC, data is not loaded into the embedded BW system. Instead, data is redirected into SAP HANA.

Related Information

Editing DataSources and Application Component Hierarchies
Enhancing DataSources

3.4.3 ETL-Based Replication (SAP Data Services)

Extraction-Transformation-Load (ETL) based data replication uses SAP Data Services (also called Data Services) to load relevant business data from SAP ERP to the SAP HANA database. This lets you read the business data on the application layer level.

> **Note**
>
> For more information about the availability of the data replication methods, see the *Product Availability Matrix (PAM)* on *SAP Service Marketplace*.

You deploy this method by defining data flows in Data Services and scheduling the replication jobs. Since this method uses batch processing, it also permits data checks, transformations, synchronizing with additional data providers, and the merging of data streams.

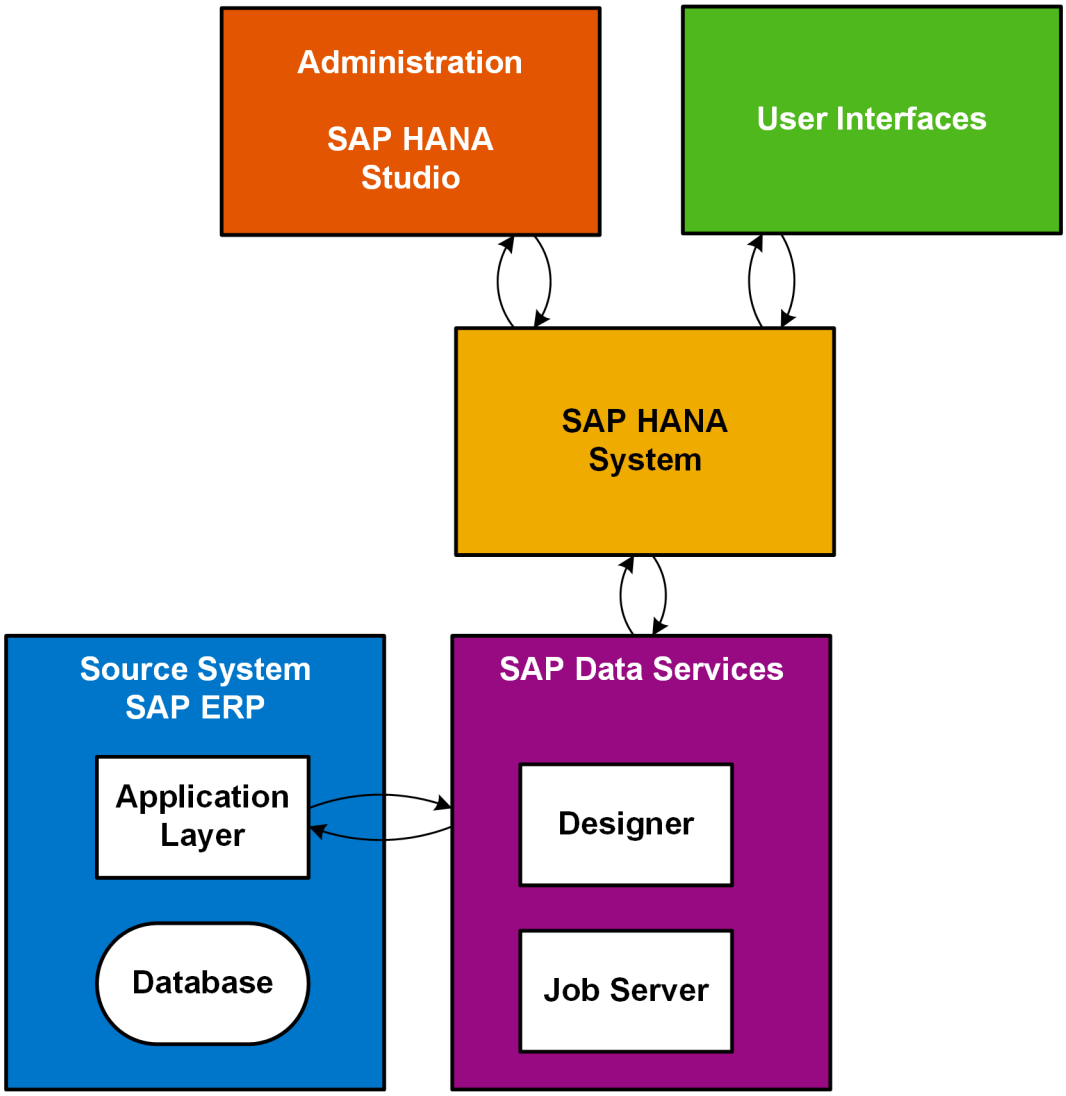

The figure above gives an overview of the ETL-based replication method. Here, data replication is operated by Data Services. Its main components are the Data Services Designer, where you model the data flow, and the Data Services Job Server for the execution of the replication jobs. An additional repository is used to store the metadata and the job definitions.

Data Flow

As for any replication scenario you have to define a series of parameters for the two systems involved. Utilizing Data Services you have to set up datastores to define such parameters. You use the Designer to set up datastores.

Datastore Setup

Setting up a datastore for the source system SAP ERP, choose SAP Applications for the type of datastore, and specify the address of the system, the user name and password allowing Data Services to access the system. Additional settings depend on the type of SAP ERP objects to be read.

For the target system of the replication, the SAP HANA database, you have to set up a separate datastore as done for the source system.

Data Flow Modeling

Once datastores are set up, Data Services can connect to the source system by RFC. Based on the metadata imported from the ABAP Data Dictionary to Data Services, you can determine the business data to be replicated. Data Services offers replication functions for a variety of data source types. However, for the replication of SAP ERP data to SAP HANA database, we recommend you to use extractors.

> **Note**
>
> - You must apply *SAP Note 1522554 - NetWeaver Support Package requirement for Data Services SAP Extractor support* to fully benefit from the extractor support.
> - In the source system, the extractors must be released for the replication access by Data Services. In addition, you have to indicate the primary key, such as the GUID, to enable the correct replication.
> - The extractors must support delta handling.

Choose the extractors that are relevant for the replication job.

Model the data flow for each extractor you have selected: indicate the source for the data flow, which is the extractor. For the target of the replication, choose a template table, which is then used in the SAP HANA database to store the replaced data.

Data Flow for Initial Load and Update

Both the initial load of business data from the source system into SAP HANA database as well as updating the replicated data (delta handling) is done using SAP Data Services. The initial load can be set up modeling a simple data flow from source to target. For the update, in most cases, the data flow is enhanced by a delta handling element, such as `Map_CDC_Operation` or `Table_Comparison Transform`. It depends on the

environment and the requested setup of the target tables which data flow design best serves your requirements.

Although we recommend you to use delta supporting extractors, you can also use SAP ABAP tables.

Replication Job Schedule

Since you can schedule the replication jobs when using Data Services, this method is suitable where the source system must be protected from additional load during the main business hours. In this way, you can shift the replication workload, for example, to the night. As a result, the data that is available for reporting always represents the state reached by the time when the latest replication job was started.

Use the Management Console, which comes with Data Services, to schedule replication jobs. You can choose from different tools and methods for the scheduling. You can also use the Management Console to monitor the replication process.

Required Software Components

This replication method requires the following main components:

- SAP HANA database
- SAP Data Services (contains Information Platform Services)

Related Information

Product Availability Matrix (PAM)
SAP Data Services on SAP Help Portal
SAP Note 1522554 - NetWeaver Support Package requirement for Data Services SAP Extractor support

3.4.4 Log-Based Replication (SAP Replication Server)

SAP Replication Server (SRS) moves and synchronizes transactional data including DML and DDL across the enterprise, providing low impact, guaranteed data delivery, real-time business intelligence, and zero operational downtime.

> **Note**
>
> For more information about the availability of the data replication methods, see the *Product Availability Matrix (PAM)* on *SAP Service Marketplace*.

SRS supports log-based replication from and to heterogeneous databases, except for the homogeneous SAP HANA to SAP HANA replication, which is trigger-based. You can use SAP Replication Server to do an initial load

as well as replication in real time to SAP HANA, at both table and database level, from these primary databases:

- SAP Adaptive Server Enterprise
- Oracle
- Microsoft SQL Server
- DB2 UDB LUW
- SAP Business Suite running on SAP ASE, Oracle, MSSQL Server or DB2 UDB LUW
- SAP HANA database

Also, SRS includes Replication Agent for SAP HANA (RAH), a light-weight server that replicates data from a primary SAP HANA to a replicate SAP HANA database using trigger-based replication, primarily used for real time data distribution and real time reporting.

For all of the above primary databases, initial load materialization of data as well as continuous real-time transactional replication are supported. The initial load materialization feature allows you to set up replication without any downtime of the primary data server and offers high performance.

You can set up the replication environment for replication into the SAP HANA database using the Replication Management Agent (RMA). Additionally, SRS offers the Data Assurance that compares row data and schema between two or more databases, reports and rectifies discrepancies. You can compare row data between any combinations of SAP® Adaptive Server® Enterprise (SAP® ASE), SAP HANA®, IBM DB2 Universal Database (UDB), Microsoft SQL Server, or Oracle databases in a heterogeneous comparison environment.

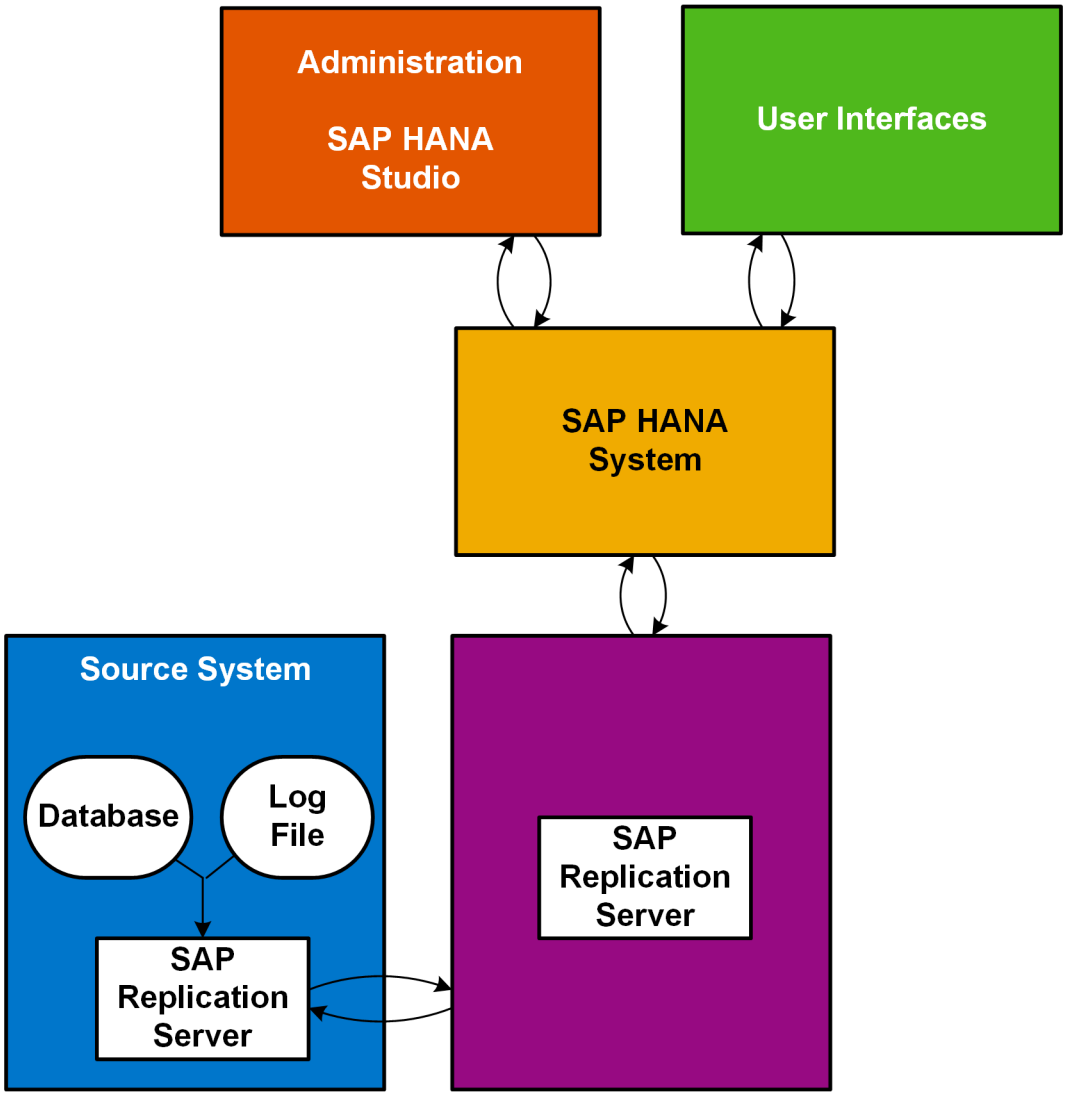

These are the components required to implement a Primary DB-to-SAP-HANA database replication system:

- A primary data server
- A replicate SAP HANA database data server
- A Replication Server® (with ExpressConnect for SAP HANA database)
- The Replication Server Options component (this component is not required for ASE Primary DB to SAP HANA): Replication Agent for MSSQL or DB2 or Oracle
- Replication Agent for SAP HANA (RAH)

> **Note**
>
> *License Requirements:*
>
> You can use these licenses for replication into the SAP HANA database:
>
> - Separately purchase SAP Replication Server, SAP HANA edition, for replicating from multiple data sources. This license includes all components required to provision data from UDB DB2, Oracle, or MSSQL into the SAP HANA database.

- Alternatively, when you buy the Real-time-data edition for SAP HANA, you automatically receive a license for SAP Replication Server, SAP HANA edition.

Individual components can also be licensed separately for replication from UDB DB2, Oracle, or MSSQL into the SAP HANA database through the SAP Replication Server, option for SAP HANA. A prerequisite for this is that you must have SAP Replication Server, enterprise edition.

Generate license keys for various components in SAP Replication Server, SAP HANA edition at https://support.sap.com/licensekey. For information on generating license keys and how to generate sub-capacity license keys, see http://infocenter.sybase.com/help/topic/com.sybase.infocenter.dc32237.1571200/doc/html/ang1278662839532.html. The license for ExpressConnect for SAP HANA Database (ECH) is distributed with Replication Server and is enabled when you purchase SAP Replication Server, SAP HANA edition.

Related Information

Product Availability Matrix (PAM)
Documentation on SAP Help Portal

3.4.5 SAP HANA Smart Data Integration

SAP HANA smart data integration and SAP HANA smart data quality load data, in batch or real-time, into HANA (on premise or in the cloud) from a variety of sources using pre-built and custom adapters.

You deploy this method by installing a Data Provisioning Agent to house adapters and connect the source system with the Data Provisioning server, housed in the HANA system. You then create replication tasks, using WebIDE, to replicate data, or flowgraphs, using Application Function Modeler nodes, to transform and cleanse the data on its way to HANA. For more information about deploying the smart data integration, see the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality Master Guide* on the SAP Help Portal.

> **Note**
>
> SAP HANA with the data provisioning server can run on IBM Power. However, the data provisioning agent needs to be hosted on an Intel machine. It is possible to connect between the two.

Related Information

SAP HANA Smart Data Integration

3.5 SAP HANA as Application and Development Platform

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

> **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 XS Classic Model

Within this category, there are two different types of applications that can be designed in this manner: native SAP HANA applications, and applications with another application server that connects to SAP HANA:

- Native SAP HANA Applications
  While SAP HANA is a database comprised of innovative technology, it is also much more than that. SAP HANA includes a small-footprint application server, a web server, and a repository for content, which provides lifecycle management functionality for development artifacts. Together with development tools, these components form an application development platform and runtime that can be used to build, deploy, and operate SAP HANA-based software applications of all kinds. These applications normally have an HTML or mobile app user interface that connects to SAP HANA using HTTP. The name for these described capabilities is SAP HANA Extended Application Services, or simply XS.
- SAP HANA-based Applications with Another Type of Application Server (for example, .NET or Java)
  Various types of applications can be built on, and run on, SAP HANA utilizing the architecture of other widely-known application servers and languages. Applications written using .NET are integrated with SAP HANA using Open Database Connectivity (ODBC), which is a standard, implementation-agnostic C-based API for accessing a database. Applications written using Java integrate using Java Database Connectivity (JDBC), which functions similarly to ODBC in principle. These interface types provide methods for creating and maintaining connections, transactions, and other mechanisms for create, read, update, and delete operations in SAP HANA; these methods map directly to the underlying SQL semantics, hiding the actual communication details. Essentially, any application that can utilize ODBC, ODBO, or JDBC can integrate with SAP HANA.

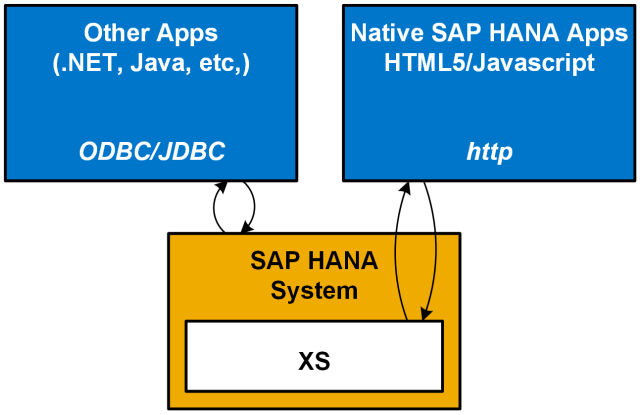

3.6 SAP HANA Smart Data Access

SAP HANA smart data access enables remote data to be accessed as if they were local tables in SAP HANA, without copying the data into SAP HANA.

Not only does this capability provide operational and cost benefits, but most importantly it supports the development and deployment of the next generation of analytical applications which require the ability to access, synthesize and integrate data from multiple systems in real-time regardless of where the data is located or what systems are generating it.

Specifically, in SAP HANA, you can create virtual tables which point to remote tables in different data sources. Customers can then write SQL queries in SAP HANA, which could operate on virtual tables. The SAP HANA query processor optimizes these queries, and executes the relevant part of the query in the target database, returns the results of the query to SAP HANA, and completes the operation.

> **Note**
>
> For IBM Power Systems, only the following data sources are supported:
>
> - SAP HANA
> - SAP IQ
> - SAP Adaptive Service Enterprise
> - Oracle Database 12C
>
> For more information, see *SAP HANA Smart Data Access* in the *SAP HANA Administration Guide*.

4 SAP HANA Architecture

Various aspects are relevant to the SAP HANA architecture.

4.1 SAP HANA Platform Software Components

SAP HANA platform is composed of several components.

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 features, **SAP HANA capabilities**, **SAP HANA options** provide additional functions. To use the *SAP HANA options* and *SAP HANA capabilities* you need a dedicated license for the options or capabilities you want to use (see disclaimer below). *SAP HANA options* and *SAP HANA capabilities* are among others:

- SAP HANA Accelerator for SAP ASE
- SAP HANA Dynamic Tiering
- SAP HANA Remote Data Sync
- SAP Landscape Transformation Replication Server
- SAP HANA Smart Data Streaming

> **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*.

> **Caution**
>
> SAP HANA server software and tools can be used for several SAP HANA platform and options scenarios as well as the respective capabilities used in these scenarios. The availability of these is based on the available SAP HANA licenses and the SAP HANA landscape, including the type and version of the back-end systems the SAP HANA administration and development tools are connected to. There are several types of licenses available for SAP HANA. Depending on your SAP HANA installation license type, some of the features and tools described in the SAP HANA platform documentation may only be available in the SAP HANA options and capabilities, which may be released independently of an SAP HANA Platform Support Package Stack

(SPS). Although various features included in SAP HANA options and capabilities are cited in the SAP HANA platform documentation, each SAP HANA edition governs the options and capabilities available. Based on this, customers do not necessarily have the right to use features included in SAP HANA options and capabilities. For customers to whom these license restrictions apply, the use of features included in SAP HANA options and capabilities in a production system requires purchasing the corresponding software license(s) from SAP. The documentation for the SAP HANA options is available in SAP Help Portal. If you have additional questions about what your particular license provides, or wish to discuss licensing features available in SAP HANA options, please contact your SAP account team representative.

Related Information

SAP HANA Hardware and Software Requirements [page 91]

4.2 SAP HANA Technical Deployment Options

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

4.2.1 SAP HANA Tenant Databases

It is now possible to install to SAP HANA to support multitenant database containers.

An SAP HANA system installed in multiple-container mode is capable of containing more than one multitenant database containers. Otherwise, it is a single-container system. Single-container systems can be converted to multiple-container systems.

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

All the databases in a multiple-container system share the same installation of database system software, the same computing resources, and the same system administration. However, each database is self-contained and fully isolated with its own:

- 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.

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

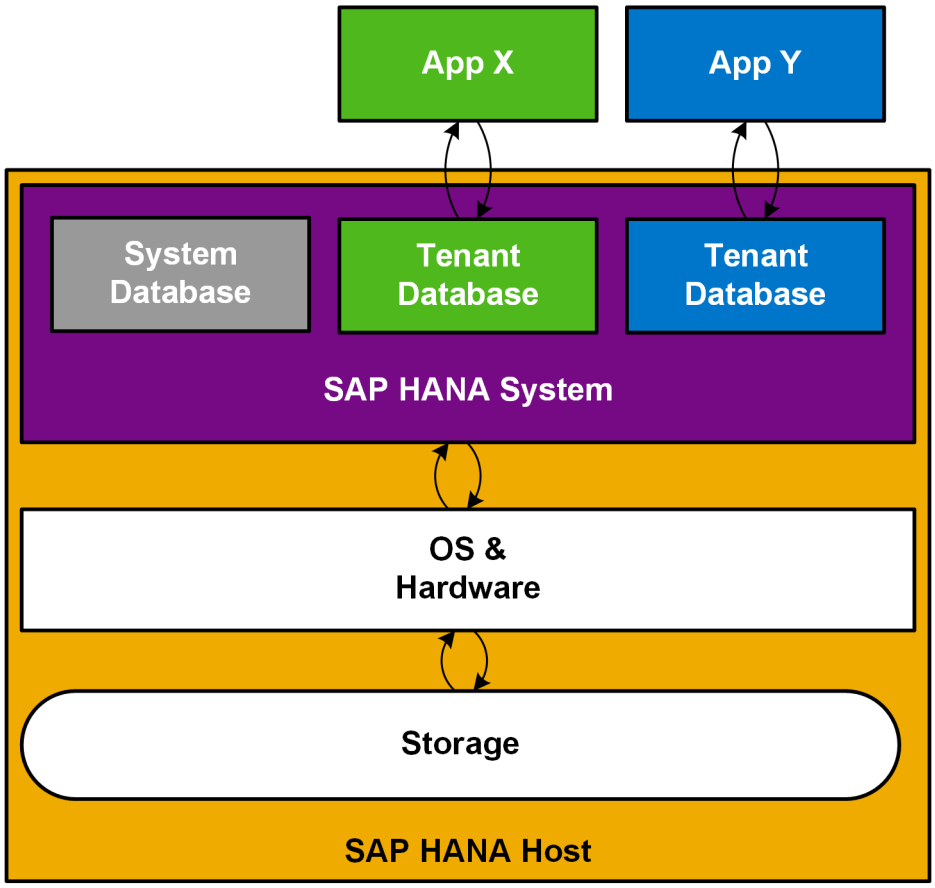

> **Note**
>
> For more information, see *Installing a Multitenant Database Container Enabled SAP HANA System* in *SAP HANA Server Installation and Update Guide* and *Creating and Configuring Tenant Databases* in the *SAP HANA Administration Guide*.

Related Information

SAP Note 2096000 - SAP HANA multitenant database containers - Additional Information

4.2.2 Single Application on One SAP HANA System (SCOS)

A single application on one SAP HANA system is also known as Single Component on One System (SCOS).

To more readily describe the various other options for technical deployment, it is useful to first illustrate the simple, straightforward approach to deploying an application on an SAP HANA system. This will be useful for comparison purposes.

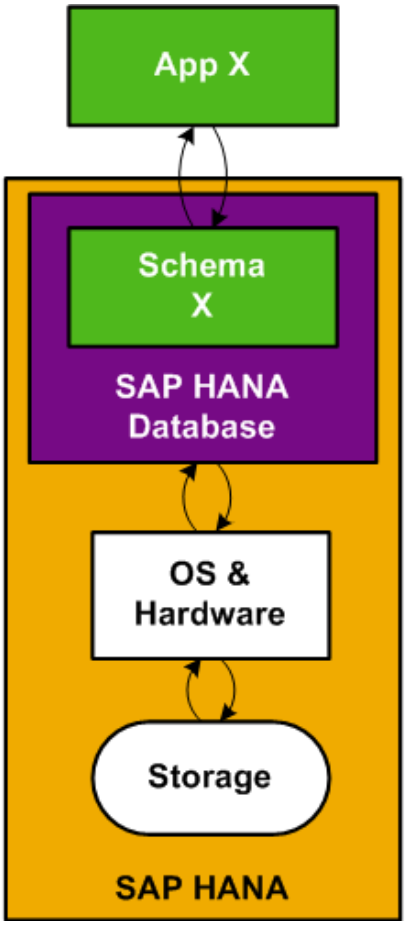

In this configuration, a single application runs in a single schema, in a single SAP HANA database as part of an SAP HANA system. This is a simple, straightforward scenario that is supported for all scenarios without restriction.

4.2.3 Multiple Applications on One SAP HANA System (MCOD)

Multiple applications on one SAP HANA system is also known as Multiple Components on One Database (MCOD).

The technical deployment type MCOD refers to the scenario where more than one application, scenario, or component runs on one SAP HANA system. This deployment type is available, with restrictions, for production SAP HANA systems.

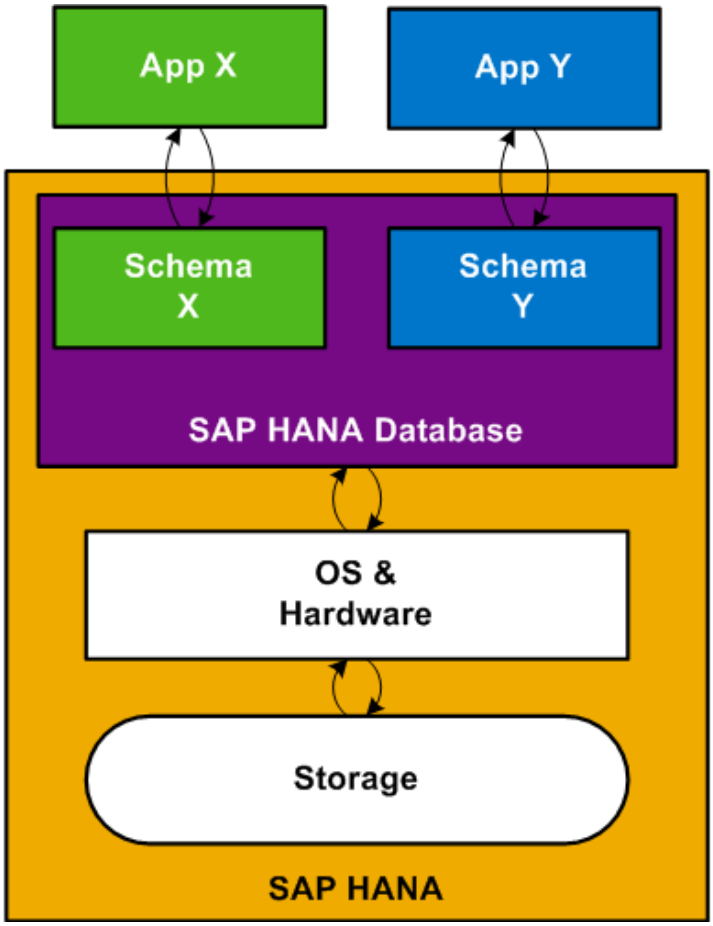

Related Information

SAP Note 1661202 - Support for multiple applications on SAP HANA
SAP Note 1826100 - Multiple applications SAP Business Suite powered by SAP HANA
SAP HANA Tenant Databases [page 26]

4.2.4 Multiple SAP HANA Systems on One Host (MCOS)

Multiple SAP HANA systems on one host are also known as Multiple Components on One System (MCOS).

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

Production support is restricted to SAP HANA SPS 09 or higher due to the availability of some resource management parameters (for example affinity). Be aware that running multi-SID on one SAP HANA host may impact performance of various types of operations, as contention for computing resources may occur (memory, cpu, i/o, ...).

SAP strongly recommends performing requisite testing in any project before going live; in general, stress/volume testing is recommended in order to provide good indicators of expected performance. When operating a system that features a multi-SID deployment, SAP recommends to actively make use of the resource management features of SAP HANA (e.g. parameters controlling memory limits, and influencing utilization of CPU cores, etc.) in order to optimize performance.

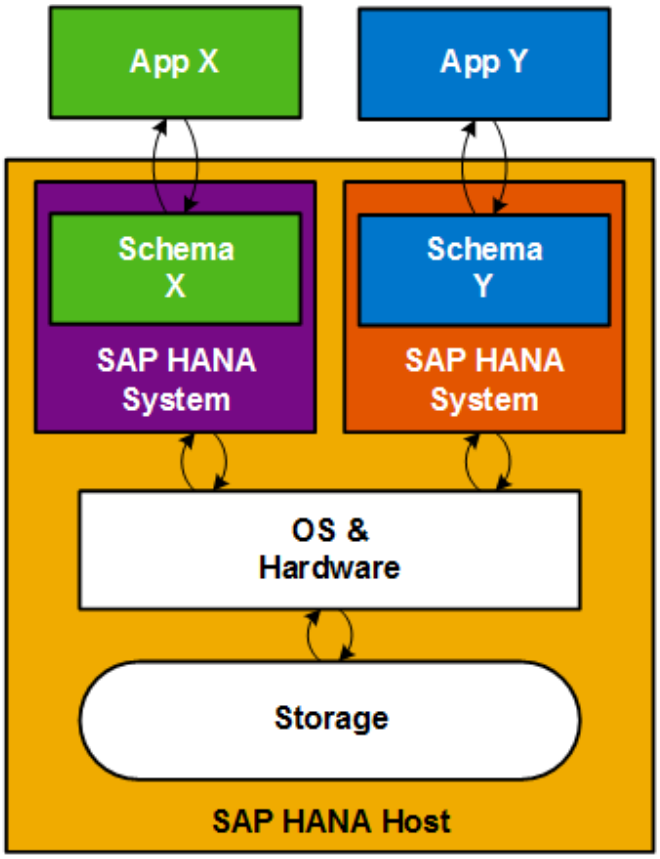

Related Information

SAP Note 1681092 - Multiple SAP HANA databases on one SAP HANA system
SAP Note 1666670 - BW on SAP HANA - landscape deployment planning

4.2.5 SAP HANA System Types

The number of hosts in a SAP HANA system landscape determines the SAP HANA system type.

The 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 host provides links to the installation directory, data directory, and log directory or to the storage itself. The storage needed for an installation does not have to be on the host. In particular, shared data storage is required for distributed systems.

An SAP HANA system can be configured as one of the following types:

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

For more information about SAP HANA system types, see the *SAP HANA Server Installation and Update Guide*.

4.2.5.1 Single-Host System

If the system consists of only one host, it is called a single-host system.

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

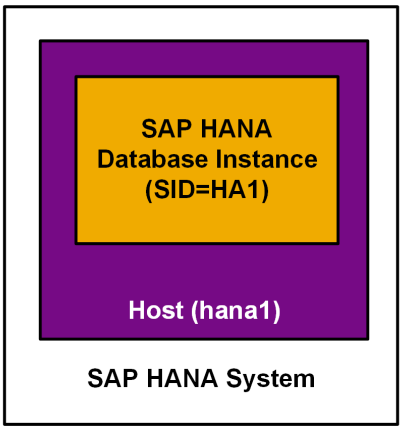

4.2.5.2 Distributed System (Multiple-Host System)

If the system consists of multiple connected hosts, it is called a distributed system.

The following graphic shows the file system for a distributed installation:

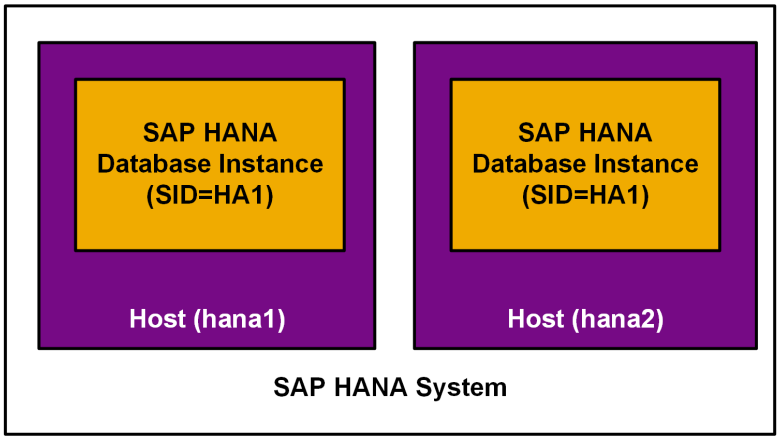

A distributed system might be necessary in the following cases:

- You can scale SAP HANA either by increasing RAM for a single server, or by adding hosts to the system to deal with larger workloads. This allows you to go beyond the limits of a single physical server.
- Distributed systems can be used for failover scenarios and to implement high availability. Individual hosts in a distributed system have different roles (master, worker, slave, and standby) depending on the task.

For more information about scaling, failover scenarios, and high availability, see the *SAP HANA Administration Guide*.

4.2.6 SAP HANA with Virtualization

The technical deployment type SAP HANA with Virtualization refers to the scenario where one or more SAP HANA database SIDs are deployed on one or more virtual machines running on SAP HANA server hardware.

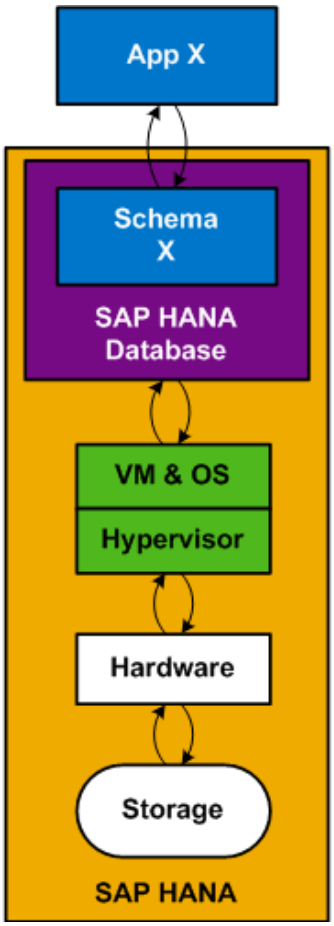

For information about SAP HANA with virtualization, see:

- Intel-Based Hardware Platforms - VMware
  - SAP Note 1788665 - SAP HANA Support for VMware vSphere Environments
  - SAP Note 2024433 - Multiple SAP HANA VMs on VMware vSphere in production
  - SAP Note 2157587 - SAP Business Warehouse, powered by SAP HANA on VMware vSphere in scale-out and production
- IBM Power Systems - IBM PowerVM
  SAP Note 2055470 - SAP HANA on POWER Planning and Installation Specifics - Central Note

For more information, see *SAP HANA and Virtualization* in *SAP HANA Server Installation and Update Guide*.

Related Information

SAP and VMware Announce SAP HANA for Production Use on VMware vSphere 5.5
SAP on VMware
SAP HANA virtualized - Overview
SAP HANA Guidelines for being virtualized with VMware vSphere
SAP Note 1788665 - SAP HANA Support for VMware vSphere Environments
SAP on VMware
SAP Note 2055470 - SAP HANA on POWER Planning and Installation Specifics - Central Note
SAP Note 2024433 - Multiple SAP HANA VMs on VMware vSphere in production
SAP Note 2157587 - SAP Business Warehouse, powered by SAP HANA on VMware vSphere in scale-out and production

4.3 The SAP HANA Network

An SAP HANA data center deployment can range from a database running on a single host to a complex distributed system with multiple hosts located at a primary and one or more secondary sites, and supporting a distributed multi-terabyte database with full high availability and disaster recovery.

In terms of network connectivity, SAP HANA supports traditional database client connections and, with SAP HANA Extended Application Services (SAP HANA XS), Web-based clients. SAP HANA can be integrated with transaction-oriented databases using replication services, as well as with high-speed event sources. SAP HANA-based applications can be integrated with external services such as e-mail, Web, and R-code execution.

The setup of an SAP HANA system, and the corresponding data center and network configurations, depends on your company's environment and implementation considerations. Some of these considerations are:

- Support for traditional database clients, Web-based clients, and administrative connections
- The number of hosts used for the SAP HANA system, ranging from a single-host system to a complex distributed system with multiple hosts
- Support for high availability through the use of standby hosts, and support for disaster recovery through the use of multiple datacenters
- Security and performance

SAP HANA has different types of network communication channels to support the different SAP HANA scenarios and setups:

- Channels used for external access to SAP HANA functionality by end-user clients, administration clients, application servers, and for data provisioning via SQL or HTTP
- Channels used for SAP HANA internal communication within the database or, in a distributed scenario, for communication between hosts

SAP HANA supports the isolation of internal communication from outside access. To separate external and internal communication, SAP HANA hosts use a separate network adapter with a separate IP address for each

of the different networks. For IBM Power systems, this might be different. In addition, SAP HANA can be configured to use SSL for secure communication.

For information about troubleshooting the SAP HANA network, see the section "Network Performance and Connectivity Problems" in the *SAP HANA Troubleshooting and Performance Analysis Guide*.

Related Information

Network Zones [page 34]
Ports and Connections [page 35]
Default Host Names and Virtual Host Names [page 59]

4.3.1 Network Zones

Separate network zones, each with its own configuration, allow you to control and limit network access to SAP HANA to only those channels required for your scenarios, while ensuring the required communication between all components in the SAP HANA network.

These network zones can be basically described as follows:

- Client zone
  The network in this zone is used by SAP application servers, by clients such as the SAP HANA studio or Web applications running against the SAP HANA XS server, and by other data sources such as SAP Business Warehouse.
- Internal zone
  This zone covers the interhost network between hosts in a distributed system as well as the SAP HANA system replication network.
- Storage zone
  This zone refers to the network connections for backup storage and enterprise storage.
  In most cases, the preferred storage solution involves separate, externally attached storage subsystem devices that are capable of providing dynamic mount-points for the different hosts, according to the overall landscape. A storage area network (SAN) can also be used for storage connectivity – for example, when running SAP HANA on IBM Power.
  See also Related Information below.

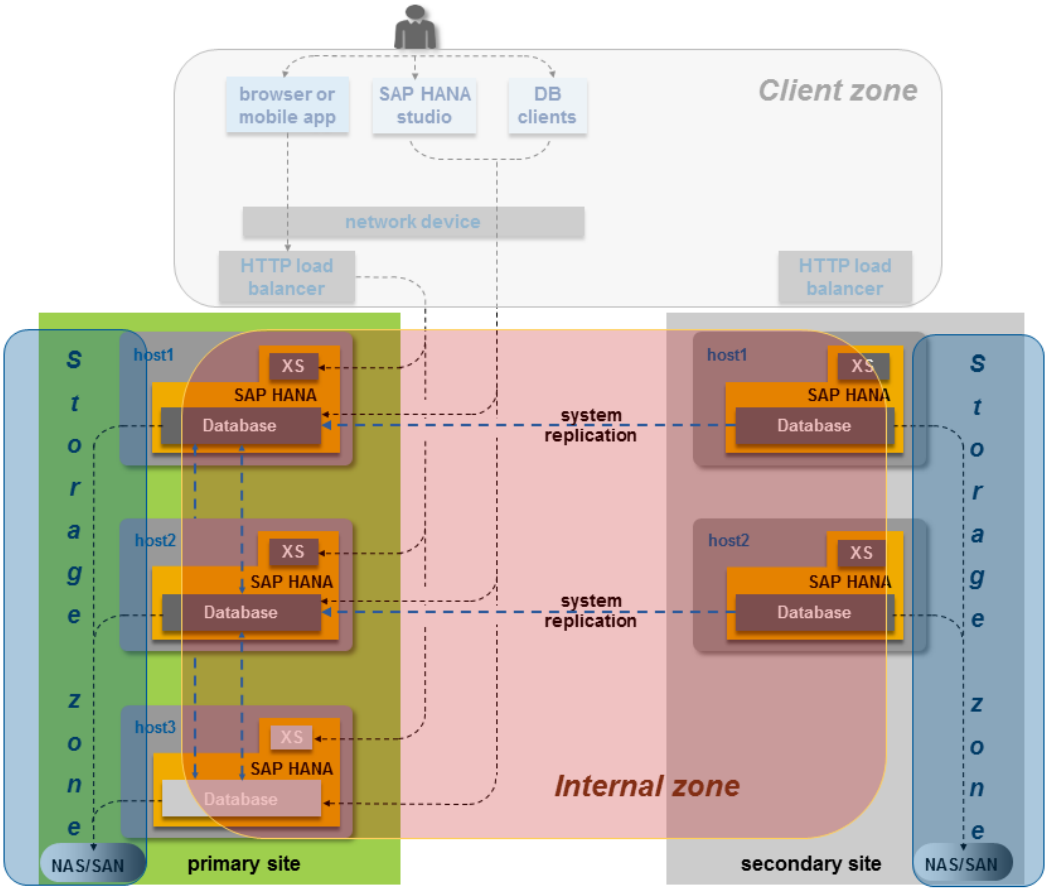

Related Information

Client Zone

Connections from Database Clients and Web Clients to SAP HANA [page 36]

Internal Zone

Connections for Distributed SAP HANA Systems [page 42]

Internal Host Name Resolution [page 62]

Host Name Resolution for System Replication [page 64]

Storage Zone

SAP HANA - Storage Requirements

FAQ - SAP HANA Tailored Data Center Integration FAQ

4.3.1.1 Ports and Connections

Before you start configuring the network for SAP HANA, you'll want to get an overview of the different types of connections to, from, and within SAP HANA and which ports to configure for them.

Related Information

Connections from Database Clients and Web Clients to SAP HANA [page 36]
Connections for Distributed SAP HANA Systems [page 42]

4.3.1.1.1 Connections from Database Clients and Web Clients to SAP HANA

Before you start configuring the network for SAP HANA, you'll want to get an overview of the external client connections to SAP HANA .

The connections between SAP HANA and external components and applications can be classified as follows:

- Connections that are used for administrative purposes
- Connections that are used for data provisioning
- Connections from database clients that access the SQL/MDX interface of the SAP HANA database
- Connections from HTTP/S clients
- Outbound connections

You can see an example of what these connections look like in the diagram below. Network connections are depicted by dashed arrows. The direction of each arrow indicates which component is the initiator (start of arrow) and which component is the listener (end point of arrow). Administrative access to and from SAP HANA is depicted by the blue dashed arrows. Port numbers are shown with a pink background. The xx in the port numbers stands for your SAP HANA instance number.

The diagram shows all the network channels used by the SAP HANA software. For purposes of illustration, the diagram shows a single host installation of SAP HANA. However, the connections shown apply equally to a distributed scenario.

> **Note**
>
> In distributed scenarios, you must also ensure that every database client can connect to every host (not shown in the diagram). Moreover, additional network channels are required in distributed scenarios for communication between the different hosts of a HANA system or between the different sites.

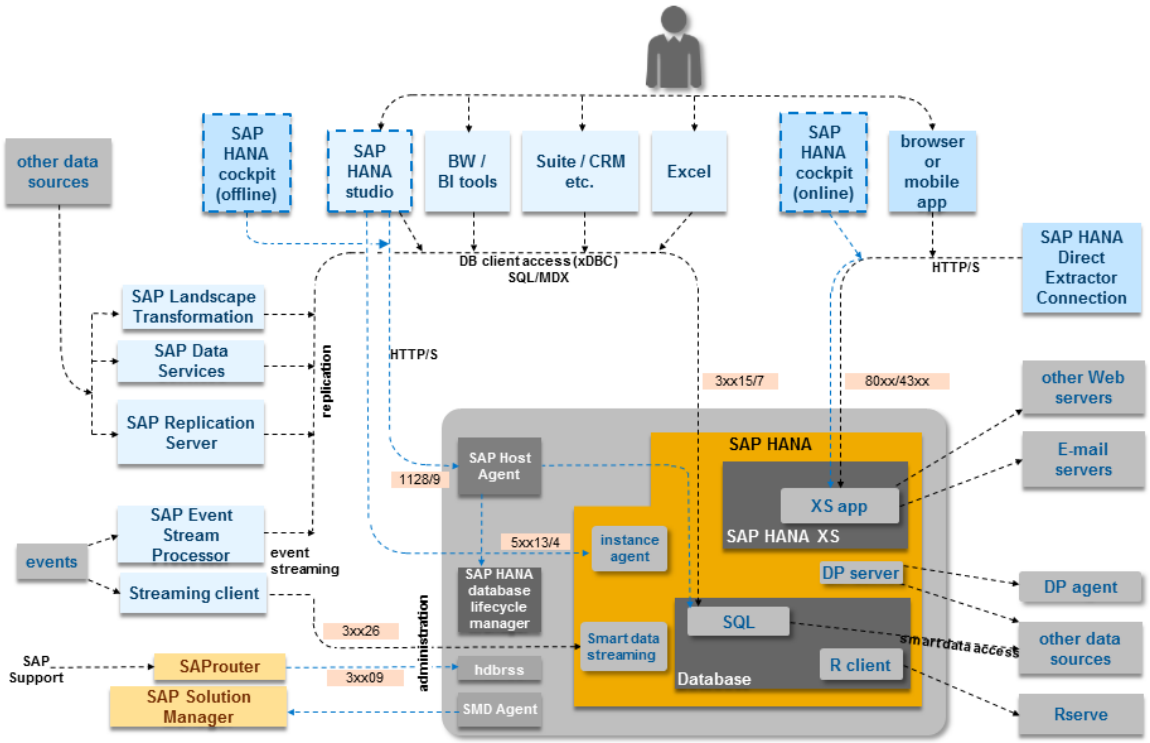

The following tables explain the diagram and the different categories described above in more detail.

Database Client Access

Client	Protocol and additional information	TCP port
Application servers that use SAP HANA as a database Examples: SAP Business Warehouse and one or more components of SAP Business Suite	You must enable SQL/MDX access for all database clients. External and internal host names are mapped for the purposes of database client access. You can change the default mapping. For more information, see the *SAP HANA Administration Guide*. The protocol used for database client access is SQLDBC (ODBC/JDBC).	3xx15 3xx17
End-user clients that access the SAP HANA database directly Example: Microsoft Excel		

Client	Protocol and additional information	TCP port
SAP HANA studio This connection is used for administrative purposes (for example, to access user data, configuration data or trace files) or for modeling purposes (to access data models).		

HTTP/S Client Access

Client	Additional information	TCP port
Examples: a Web browser or a mobile device	Access for applications based on SAP HANA Extended Application Services (SAP HANA XS). For more information, see the *SAP HANA Developer Guide* . The SAP HANA platform itself has a number of Web applications that run on SAP HANA Extended Application Services: for example, the SAP HANA cockpit, the SAP HANA Web-based Development Workbench, SAP HANA Application Lifecycle Management, and the SAP DB Control Center.	80xx/43xx
SAP HANA cockpit	This connection is required for the online administration of SAP HANA by the Web-based applications that comprise the SAP HANA cockpit.	
SAP HANA Direct Extractor Connection (DXC)	This connection is used for ETL-based data acquisition. For more information, see the *SAP HANA Direct Extractor Connection Implementation Guide*.	

Client	Additional information	TCP port
UI toolkit for SAP HANA Info Access	For more information, see the *SAP HANA Search Developer Guide*. **Caution** The toolkit is part of the SAP HANA Advanced Data Processing option. 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 [page 94].	
SAP HANA studio	This is the connection to the SAP HANA database lifecycle manager via SAP Host Agent. For more information about the SAP HANA database lifecycle manager, see the *SAP HANA Administration Guide*.	1128 1129 (SSL)
SAP HANA cockpit for offline administration	The offline version of the SAP HANA cockpit communicates with SAP HANA using SAP Host Agent. For more information about the SAP HANA cockpit for offline administration, see the *SAP HANA Administration Guide*.	

Each SAP HANA host comes with an SAP HANA XS engine and a native SAP HANA service for Web dispatcher administration (`webdispatcher` with operating system process `hdbwebdispatcher`). In addition, the clients in distributed or system replication landscapes typically connect through a web dispatcher that is external to SAP HANA. This web dispatcher is used for load balancing. It can be either SAP Web Dispatcher (`sapwebdisp`) or a third-party tool. Similar to the client libraries, the HTTP load balancer has its own `ini` file in which the available SAP HANA XS engines are configured. Typically, the high availability of the load balancer is supported by the use of an external cluster manager as well as by virtual host names.

Administrative Tasks

Client	Protocol and additional information	TCP port
SAP support	The connection is not active by default because it is required only in certain support cases. To find out how to open a support connection, see the *SAP HANA Administration Guide*. An internal SAP protocol is used for this connection.	3xx09

Client	Protocol and additional information	TCP port
SAP HANA studio	The connection to the instance agent acts as an administrative channel for low-level access to the SAP HANA instance to allow features such as starting or stopping of the SAP HANA database. The protocol used for this connection is SQLDBC (ODBC/JDBC).	5xx13 5xx14 (SSL)

Other administrative tasks, mainly database administration, use the SQL/MDX channel of the database.

Data Provisioning

Client	Protocol and additional information	TCP port
Replication systems for external data sources	• The following replication technologies may be used: • SAP Landscape Transformation (SLT) The protocol is SQLDBC (ODBC/JDBC). • SAP Data Services (DS) The protocol is SQLDBC (ODBC/JDBC). • SAP Replication Server (not included with all licensed editions of SAP HANA) The protocol is SQLDBC (ODBC/JDBC). • SAP HANA Direct Extractor Connection (DXC). This technology uses HTTP/S access.	3xx15 3xx17
Streaming client	This connection is used for SAP HANA Smart Data Streaming (supported on Intel-based platforms only). **⚠ Caution** SAP HANA Smart Data Streaming is an SAP HANA option. 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 [page 94].	3xx26

Outbound Connections

Connection	Additional information
From the SAP Solution Manager diagnostics (SMD) agent to SAP Solution Manager	For information about how to install the SAP Solution Manager diagnostics agent, see SAP Note 1858920.
Calls from SAP HANA Extended Application Services to external servers	Examples: a Web server or an e-mail server (depends on what applications your company has deployed)

Connection	Additional information
Smart data access from SAP HANA to external data sources for data federation purposes	Smart data access for SAP HANA is described elsewhere in this document.
From SAP HANA to the R environment	Only required for scenarios which use the **R** integration supported by SAP HANA. For more information, see the *SAP HANA R Integration Guide*.
From the data provisioning server of the SAP HANA database to the data provisioning agent and, depending on the type of adapter used, to the external data source(s)	This connection is used for SAP HANA smart data integration in scenarios where SAP HANA is deployed on premise. For more information, see the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality Master Guide*. SAP HANA with the data provisioning server can run on IBM Power. However, the data provisioning agent needs to be hosted on an Intel machine. It is possible to connect between the two. **⚠ Caution** SAP HANA smart data integration is an SAP HANA option. 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 [page 94].

Related Information

SAP HANA platform documentation on SAP Help Portal

Client connections

Host Name Resolution for SQL Client Communication [page 68]

Mapping Host Names for Database Client Access [page 70]

SAP HANA smart data access

SAP HANA Smart Data Access [page 24]

SAP Solution Manager diagnostics agent

SAP Note 1858920

SAP HANA options

Connections for SAP HANA Smart Data Streaming [page 49]

SAP HANA Smart Data Streaming documentation on SAP Help Portal

SAP HANA Smart Data Integration documentation on SAP Help Portal

4.3.1.1.2 Connections for Distributed SAP HANA Systems

Before you start configuring the network for SAP HANA, you'll want to get an overview of the internal connections between server components. In SAP HANA, server components are distributed across multiple hosts for the purposes of scalability and availability.

An installed SAP HANA system is identified by a system ID (SID). It is perceived as one unit from the perspective of the administrator, who can install, update, start up, shut down, or back up the system as a whole. The different services of the system share the same metadata, and requests from client applications are transparently dispatched to the different services in the system. A **distributed SAP HANA system** is a system that is installed on more than one host. Otherwise it is a single-host system. An **SAP HANA instance** is a set of components of a distributed system that are installed on one host.

In addition to external network connections, SAP HANA uses separate, dedicated connections exclusively for internal communication. There are two types of internal communication:

- Distributed scenarios
  Internal network communication takes place between the hosts of a distributed system on one site. SAP HANA hosts contain a separate network interface card that is configured as part of a private network, using separate IP addresses and ports. For IBM Power systems, this might be different.

  > **Note**
  >
  > In single-host scenarios, the same communication channels are used for communication between the different processes on a single host and the internal IP addresses/ports are by default bound to the `localhost` interface. Before SPS 06, these ports were by default bound to all network interfaces.

  There are a number of ways to isolate internal network ports from the client network. The preferred method depends on the data center configuration, on hardware vendor delivered options, and on the high availability implementation. Applying network separation for the internal communication prevents unauthorized access from outside networks. For additional security it is possible to encrypt the internal communication using SSL. For more information about security, see the *SAP HANA Security Guide*.

  Ports for Distributed Scenarios

  The xx in the port numbers stands for your SAP HANA instance number.

Client	TCP port	Service	Use
Hosts of a distributed system on one site	3xx00	daemon	
	3xx01	nameserver	
	3xx02	preprocessor	
	3xx03	indexserver	
	3xx04	scriptserver	Optional
	3xx05	statisticsserver	Applicable only if you do **not** use the embedded statistics server. For more information, see SAP Note 1917938.

Client	TCP port	Service	Use
	3xx07	xsengine	
	3xx10	compileserver	
	3xx40-3xx99	indexservers	Optional, after n->1 recovery

- System replication
  Internal network communication for system replication takes place between a primary site and a secondary site. In a multitier setup, this communication takes place between the tier-1 primary system and tier-2 secondary system as well as, asynchronously, between the tier-2 and tier-3 secondary systems. For more information about system replication and multitier setups, see the *SAP HANA Administration Guide*. You can secure system replication connections using the Secure Sockets Layer (SSL) protocol. In this case, landscape topology communication on the one hand, and data replication and log replication channels on the other, must be secured in separate steps. By default, SSL is turned off. For more information about configuring SSL for internal communication as well as securing communication between sites in system replication scenarios, see the *SAP HANA Security Guide*.

  Ports for System Replication
  The xy in the port numbers stands for your SAP HANA instance number plus 1.

Client	TCP port	Service	Used for...
Hosts on primary and secondary sites	3xy01	nameserver	Log and data shipping
	3xy02	nameserver	Metadata communication
	3xy03	indexserver	Log and data shipping
	3xy04	scriptserver	Log and data shipping (optional)
	3xy05	statisticsserver	Log and data shipping Applicable only if you do **not** use the embedded statistics server. For more information, see SAP Note 1917938.
	3xy07	xsengine	Log and data shipping
	3xy40-3xy99	indexservers	Log and data shipping (optional, after n->1 recovery)

> **Note**
>
> SAP HANA internal communication has sometimes been unofficially referred to as TREXNet communication. However, the term TREXNet is not valid in the context of SAP HANA.

Example 1

The following diagram shows a distributed SAP HANA system with two active hosts and an extra standby host, fully system-replicated to a secondary site to provide full disaster recovery support.

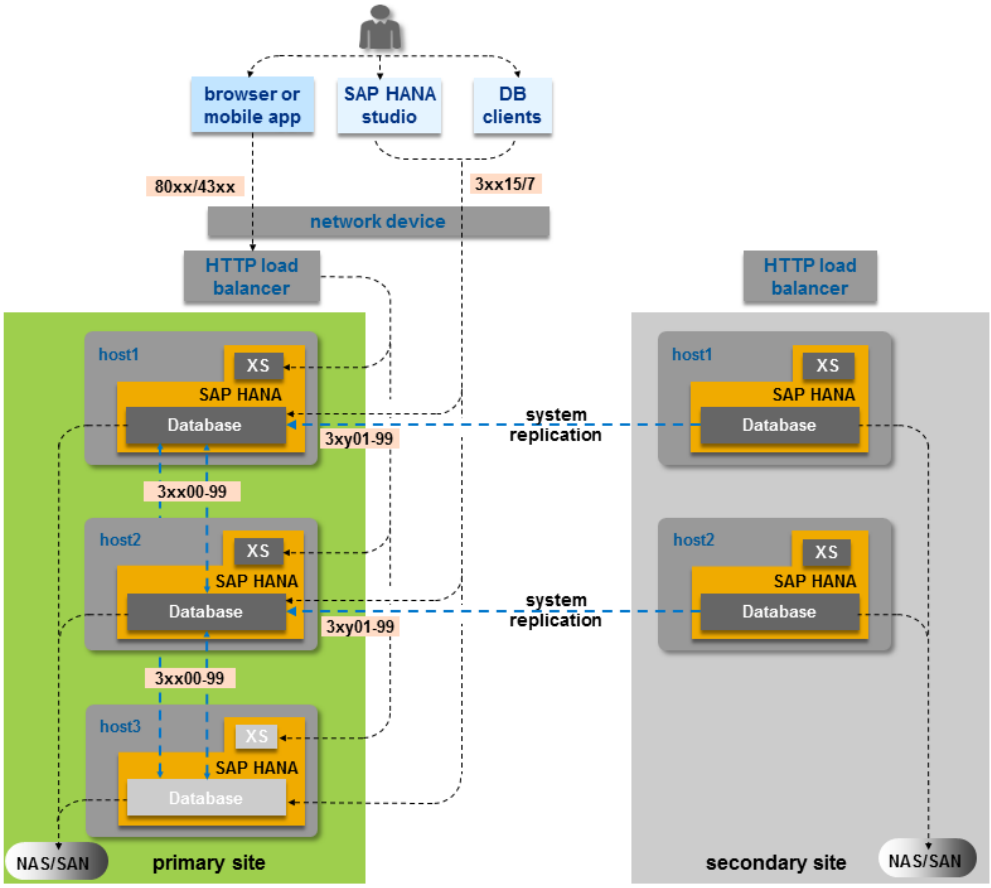

The vertical blue dashed lines show the communications between the services of the system; all instances communicate with all other instances of a distributed system on one site. The horizontal blue dashed lines show the initial connection for system replication communication between services on hosts on the primary site and the corresponding services on hosts of the secondary site (typically over a high-performance fiber network). The details of system replication configuration depend on the specific network setup of your company.

Also shown is the connection to a storage subsystem, which can be either a network attached storage device (NAS) or a storage area network (SAN).

One of the most critical aspects of the network design of a highly available distributed system is the question of how the different clients manage to reconnect to the system when its topology changes due to the recovery operations following a failure or disaster. The diagram shows two additional components that can be used to handle client reconnection:

- An **HTTP load balancer** (such as SAP Web Dispatcher) acts as a reverse proxy for HTTP connections and exposes a consistent external network address to the client network. The HTTP load balancer can also be used to provide load-balanced access to multiple distributed SAP HANA Extended Application Services (XS) servers.

- A **network device** (router and/or switch), which can be used in conjunction with DNS or virtual IP redirection

For information about how to set up multiple XS instances as well as client connection recovery, see the *SAP HANA Administration Guide*.

Example 2

The following diagram shows an example of multitier system replication:

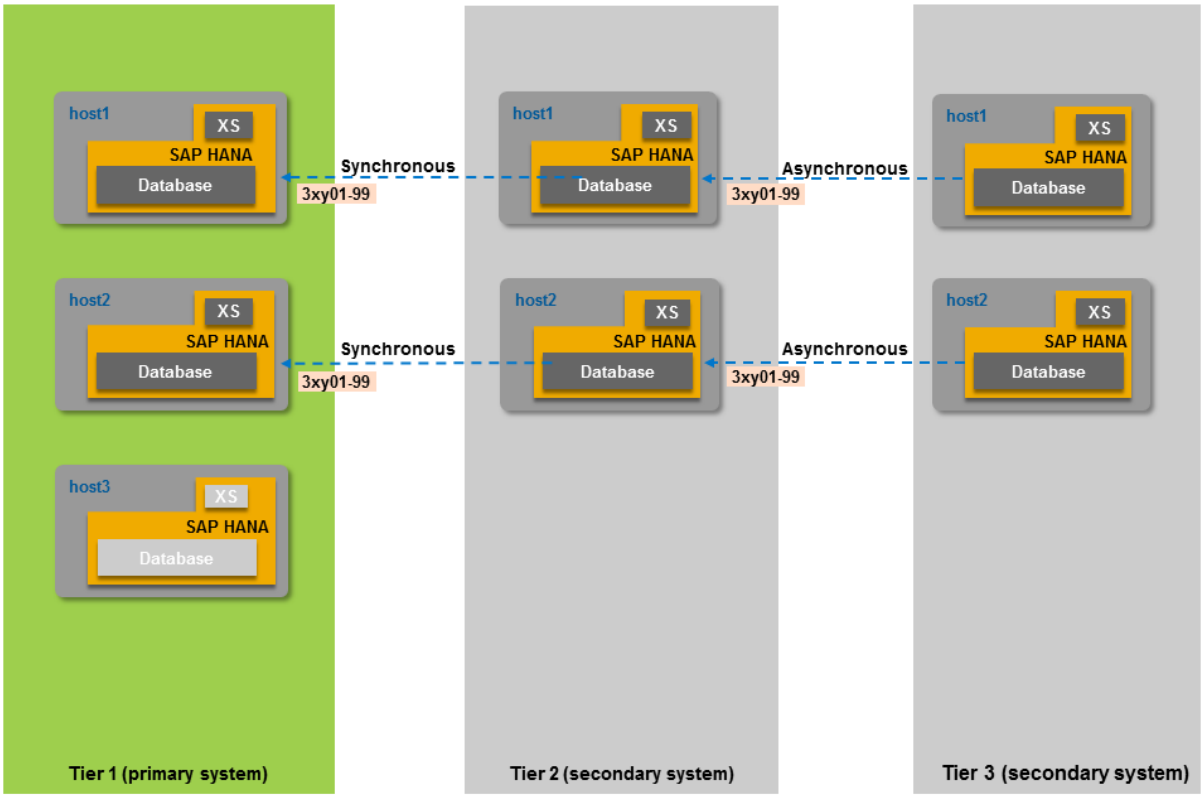

Related Information

Internal Host Name Resolution [page 62]
Host Name Resolution for System Replication [page 64]
Connections from Database Clients and Web Clients to SAP HANA [page 36]
SAP Note 1917938

4.3.1.1.2.1 High Availability for SAP HANA

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

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

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

1. Hardware Redundancy – SAP HANA appliance and tailored data center integration vendors offer multiple layers of redundant hardware, software and network components, such as redundant power supplies and fans, enterprise grade error-correcting memories, fully redundant network switches and routers, and uninterrupted power supply (UPS). Disk storage systems use batteries to guarantee writing even in the presence of power failure, and use striping and mirroring to provide redundancy for automatic recovery from disk failures. Generally speaking, all these redundancy solutions are transparent to SAP HANA's operation, but they form part of the defense against system outage due to single component failures.
2. Software – SAP HANA is based on Linux (see *SAP Note 2235581 - SAP HANA: Supported Operating Systems*) and includes security pre-configurations (for example, minimal network services). Additionally, the SAP HANA system software also includes a watchdog function, which automatically restarts configured services (index server, name server, and so on), in case of detected stoppage (killed or crashed).
3. Persistence – SAP HANA persists transaction logs, savepoints and snapshots to support system restart and recovery from host failures, with minimal delay and without loss of data.
4. Standby and Failover – Separate, dedicated standby hosts are used for failover, in case of failure of the primary, active hosts. This improves the availability by significantly reducing the recovery time from an outage.

SAP HANA High Availability Support

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

SAP HANA supports the following recovery measures from failures:

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

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

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

SAP HANA supports system replication for multitenant database containers on the system database level, this means the multitenant database system as a whole including all tenant databases. An SAP HANA system installed in multiple-container mode always has exactly one system database and any number of multitenant database containers (including zero), also called tenant databases. For more information on multitenant database containers see *Creating and Configuring Tenant Databases* in the *SAP HANA Administration Guide*.

Using Secondary Servers for Non-Production systems

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

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

Related Information

Network Required for SAP HANA System Replication
How to Perform System Replication for SAP HANA
SAP Note 2235581 - SAP HANA: Supported Operating Systems

4.3.1.1.3 Connections for SAP HANA Options

Dedicated ports are used to connect SAP HANA options.

> **Caution**
>
> SAP HANA server software and tools can be used for several SAP HANA platform and options scenarios as well as the respective capabilities used in these scenarios. The availability of these is based on the available SAP HANA licenses and the SAP HANA landscape, including the type and version of the back-end systems the SAP HANA administration and development tools are connected to. There are several types of licenses available for SAP HANA. Depending on your SAP HANA installation license type, some of the features and

tools described in the SAP HANA platform documentation may only be available in the SAP HANA options and capabilities, which may be released independently of an SAP HANA Platform Support Package Stack (SPS). Although various features included in SAP HANA options and capabilities are cited in the SAP HANA platform documentation, each SAP HANA edition governs the options and capabilities available. Based on this, customers do not necessarily have the right to use features included in SAP HANA options and capabilities. For customers to whom these license restrictions apply, the use of features included in SAP HANA options and capabilities in a production system requires purchasing the corresponding software license(s) from SAP. The documentation for the SAP HANA options is available in SAP Help Portal. If you have additional questions about what your particular license provides, or wish to discuss licensing features available in SAP HANA options, please contact your SAP account team representative.

4.3.1.1.3.1 Connections for SAP HANA Dynamic Tiering

No additional manual configuration of connections and ports is required in the SAP HANA software for the SAP HANA dynamic tiering option.

When an external client sends a request for warm data, it connects to the SAP HANA host which passes the request to the dynamic tiering host. The dynamic tiering host listens on internal port 3xx12. There is no direct connection between external components and the dynamic tiering host. The connection back from the dynamic tiering host to the SAP HANA host is through the SQL port 3xx15 of the SAP HANA host.

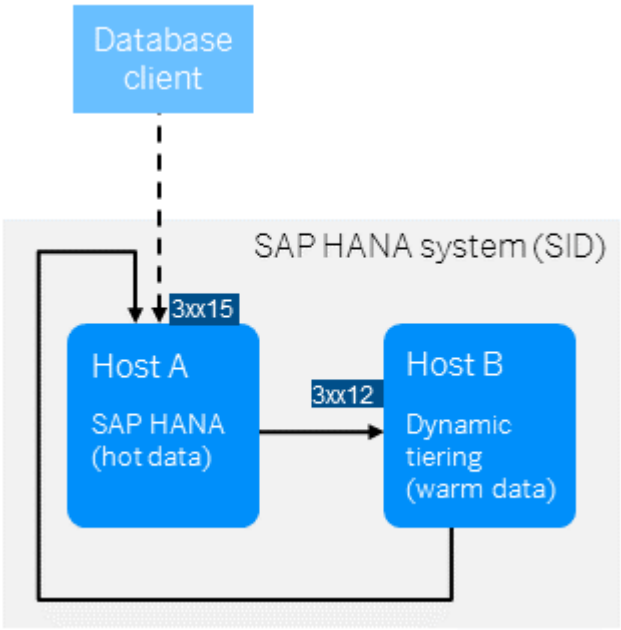

Related Information

SAP HANA dynamic tiering on SAP Help Portal
Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities [page 94]

4.3.1.1.3.2 Connections for SAP HANA Smart Data Integration

The connections between the components for SAP HANA smart data integration may differ depending on whether SAP HANA is deployed on premise, in the cloud, or behind a firewall.

For more information, see the *SAP HANA Smart Data Integration and SAP HANA Smart Data Quality Master Guide*.

> ⚠ Caution
>
> SAP HANA smart data integration is an SAP HANA option. 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 [page 94].

Related Information

SAP HANA Smart Data Integration documentation on SAP Help Portal

4.3.1.1.3.3 Connections for SAP HANA Smart Data Streaming

The internal connections and ports for the SAP HANA smart data streaming option are set up automatically. None of the ports are configurable.

> ℹ Note
>
> SAP HANA smart data streaming is supported on Intel-based platforms only.

The SAP HANA server connects to one or more smart data streaming servers on internal port 3xx16. Through this connection, SAP HANA gathers smart data streaming statistics. The connection is triggered by the SAP HANA cockpit monitoring views.

The smart data streaming hosts connect to the SAP HANA server on port 3xx15. The streaming hosts retrieve the streaming license information and the streaming cluster configuration (which is stored on the SAP HANA database). If the smart data streaming project has an SAP HANA adapter or a generic database adapter that connects to SAP HANA, it would also use the 3xx15 port connection. The connection on port 3xx15 is initiated upon startup of the streaming host.

Any streaming clients that run outside the SAP HANA system (such as custom-built external adapters) connect to a streaming node via the XML/RPC protocol on port 3xx26.

In a multinode setup, the 193xx port is used for interserver communication between streaming hosts. This port is for internal use, but you may want to make a note of it for firewall settings.

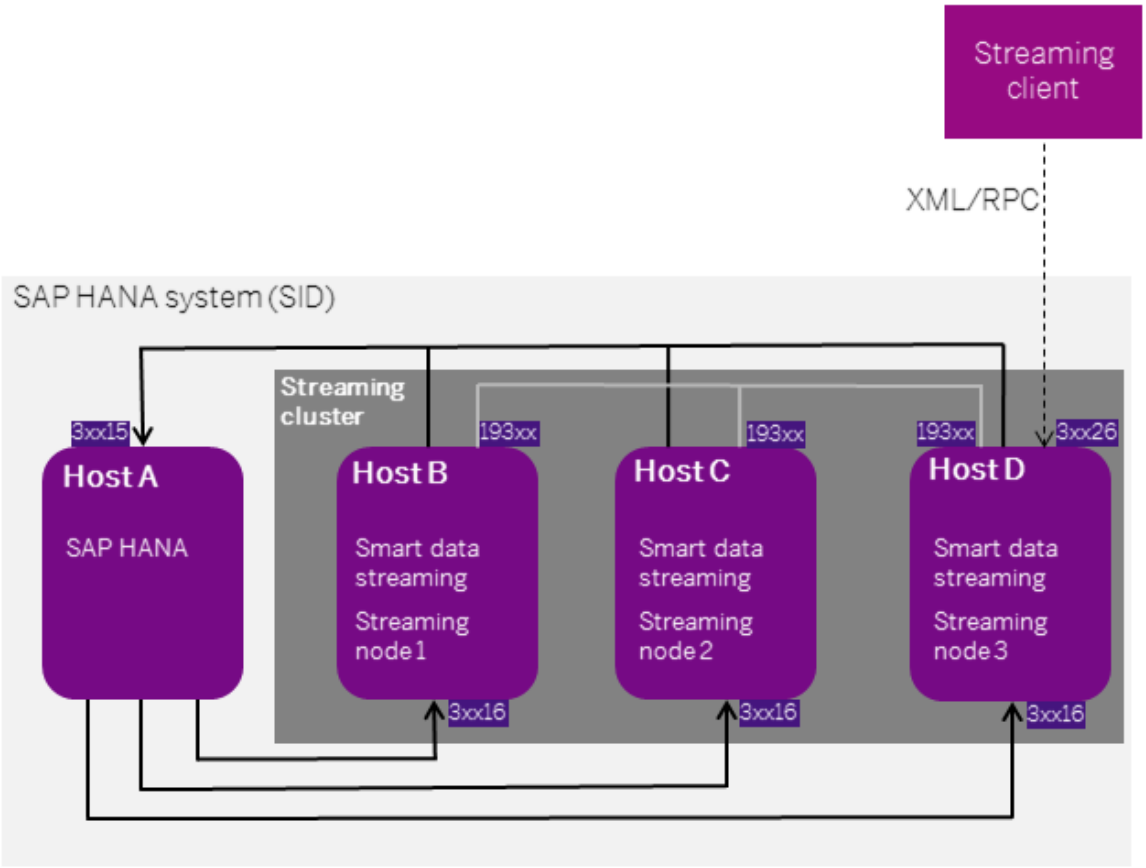

Related Information

SAP HANA smart data streaming on SAP Help Portal
Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities [page 94]

4.3.1.1.3.4 Connections for SAP HANA Accelerator for SAP ASE

The internal connections and ports for the SAP HANA accelerator for SAP ASE option are set up automatically.

When an external client sends a request for warm data, it connects to the SAP HANA host which passes the request to the accelerator for SAP ASE host. The accelerator for SAP ASE host listens on internal port 3xx21. The connection back from the accelerator for SAP ASE host to the SAP HANA host is through the SQL port 3xx15 of the SAP HANA host. Any SAP ASE clients that run outside the SAP HANA system can connect to an accelerator for SAP ASE node on port 3xx21 directly.

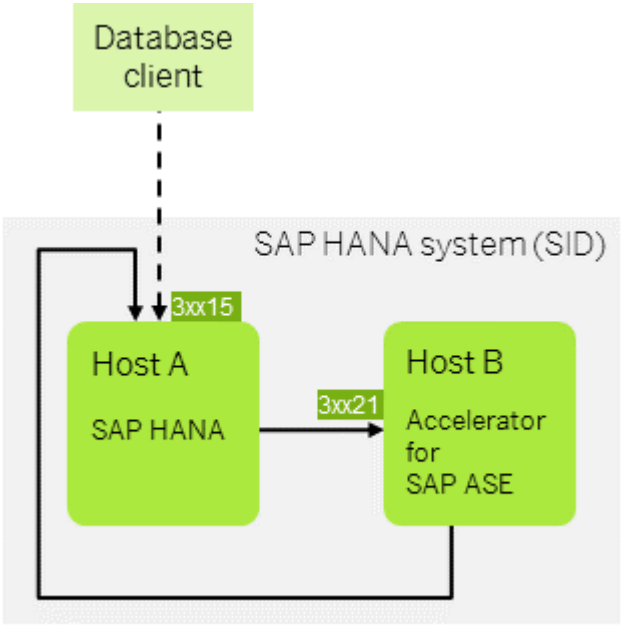

Related Information

SAP HANA accelerator for SAP ASE on SAP Help Portal
Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities [page 94]

4.3.1.1.4 Connections for Multitenant Database Containers

Additional ports and connections are required to run SAP HANA with multitenant database containers.

Port Assignment in Tenant Databases

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

The default port number range for tenant databases is 3`<instance>`40—3`<instance>`99. This means that the maximum number of tenant databases that can be created per instance is 20. However, you can increase this by reserving the port numbers of further instances. You do this by configuring the property `[multidb]`

`reserved_instance_numbers` in the `global.ini` file. The default value of this property is 0. If you change the value to 1, the port numbers of one further instance are available (for example, 30040—30199 if the first instance is 00). If you change it to 2, the port numbers of two further instances are available (for example, 30040—30299 if the first instance is 00). And so on.

Let's look at some simple examples.

Example

Example 1:

You install a new SAP HANA system in multiple-container mode. Then, you create three tenant databases. Each of these tenant databases is automatically assigned three port numbers, one for each of the following connection types:

- Internal communication
- SQL
- HTTP (This is the port of the XS server embedded in the index server.)

The first tenant database is assigned port numbers 3<instance>40—42, the second ports 3<instance>43—45, and the third 3<instance>46—48.

Example 2:

You install a new SAP HANA system in multiple-container mode. Then, you create a tenant database. The same three port numbers as above are assigned: 3<instance>40 (internal communication), 3<instance>41 (SQL), and 3<instance>42 (HTTP). Next, you add a separate xsengine service to the first database. This service is automatically assigned the next three available port numbers: 3<instance>43—45. Finally, you create a second tenant database. This tenant database is automatically assigned the next three available port numbers: 3<instance>46—48.

Example 3:

You convert a single-container system to a multiple-container system. This results in the automatic creation of one tenant database. This tenant database has the same port numbers as the original single-container system: 3<instance>03 (internal communication), 3<instance>15 (SQL), 3<instance>08 (HTTP). Then, you add a second indexserver to the tenant database. It is automatically assigned port numbers 3<instance>40—42. Finally, you create a second tenant database. It is automatically assigned ports the next three available port numbers: 3<instance>43—45.

Note

All of the above examples refer to single-host systems and are based on automatic port number assignment.

Note

The port number of the system database are fixed: 3<instance>01 (internal), 3<instance>13 (SQL), and 3<instance>14 (HTTP).

You can determine the ports used by a particular tenant database by querying the M_SERVICES system view, either from the tenant database itself or from the system database.

- From the tenant database: `SELECT SERVICE_NAME, PORT, SQL_PORT, (PORT + 2) HTTP_PORT FROM SYS.M_SERVICES WHERE ((SERVICE_NAME='indexserver' and COORDINATOR_TYPE= 'MASTER') or (SERVICE_NAME='xsengine'))`
- From the system database: `SELECT DATABASE_NAME, SERVICE_NAME, PORT, SQL_PORT, (PORT + 2) HTTP_PORT FROM SYS_DATABASES.M_SERVICES WHERE DATABASE_NAME='<DBNAME>' and ((SERVICE_NAME='indexserver' and COORDINATOR_TYPE= 'MASTER') or (SERVICE_NAME='xsengine'))`

Remember

If your system was converted from single-container mode to multiple-container mode, the HTTP port number of the first tenant database is **always** 3`<instance>`08 and not the port number returned using the above queries.

Note

System privilege DATABASE ADMIN or CATALOG READ is required to read the M_SERVICES system view.

The following diagram shows an example of the connections and ports used in a multiple-container system with two tenant databases, installed on a single host. It is a new SAP HANA system that was installed in multiple-container mode, to which two tenant databases have been added.

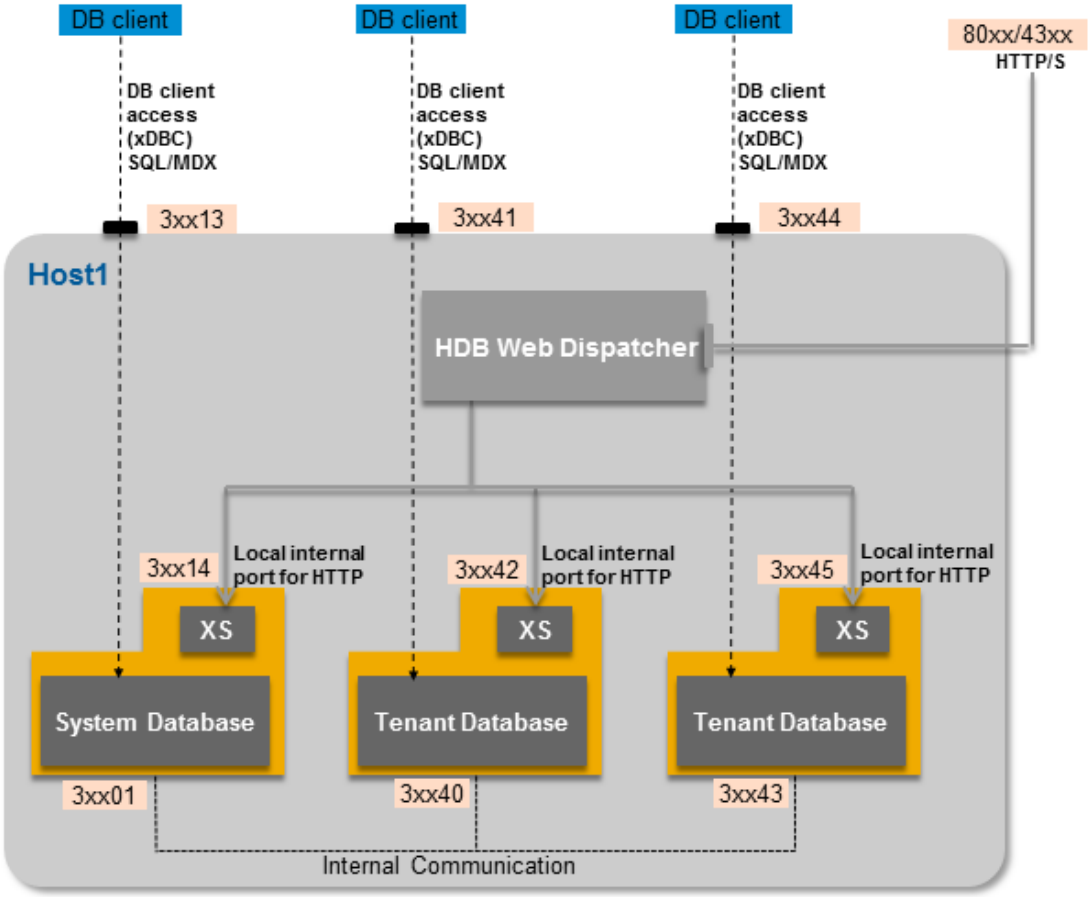

HTTP(S) Client Access

The XS server allows Web-based applications to access SAP HANA via HTTP(S). The internal Web Dispatcher of the SAP HANA system manages these incoming HTTP(S) requests. To allow applications to send requests to specific databases in a multiple-container system, every tenant database needs an alias host name. Requests to the alias host name can then be forwarded to the XS server of the corresponding tenant database. Requests with the physical host name in the HTTP host header are forwarded to the XS server running on the system database.

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

You configure HTTP(S) access to tenant databases by specifying in the `xsengine.ini` file the URLs by which each tenant database is publicly accessible. The system then automatically configures the Web Dispatcher by generating the required profile entries in the `webdispatcher.ini` configuration file. It is not necessary to specify the URL of the system database, this is done automatically.

> **Note**
>
> This automatic configuration of the Web Dispatcher is controlled by the parameter `[profile] wdisp/system_auto_configuration` in the `webdispatcher.ini` configuration file. If this parameter is set to **`false`** or is not available (revisions earlier than SPS 10), you need to configure the `webdispatcher.ini` file manually.

For more information, see *Configure HTTP Access to Multitenant Database Containers* in the *SAP HANA Administration Guide*.

> **Caution**
>
> Do not use a wildcard server certificate if strict isolation between tenant databases is required. If authentication relies on a wildcard certificate and a shared trust store, users of one tenant database will be able to log on to other databases in the system.

4.3.1.1.5 Connections for SAP HANA Extended Application Services, Advanced Model

Additional ports and connections are required if you are using SAP HANA extended application services, advanced model.

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.

Client(s)	Service	TCP Port	Use
Application UI (browser, mobile, and so on)	xsuaaserver	3xx32	Client HTTP(S) This port is used for the connection from the client to the xscontroller-managed Web Dispatcher (platform router) for purposes of user authentication.
Port 3xx32		3xx31	Internal HTTP(S) This port is used for the connection from the xscontroller-managed Web Dispatcher (platform router) to the xsuaaserver for purposes of user authentication.
• Command line client • Client library (Java) • One or more SAP HANA XS advanced model applications used, for example, for administrative and/or monitoring purposes	xscontroller	3xx30	Client HTTP(S) This port is used for the connection to the xscontroller-managed Web Dispatcher for purposes of data access.
Port 3xx30		dynamic, in the range 51000-51500	Internal HTTP(S) This port range is used for the connection from the xscontroller-managed Web Dispatcher (platform router) to the xscontroller for purposes of data access.
Application UI (browser, mobile, and so on)	Instances	dynamic, in the range 51000-51500	Client HTTP(S) This port range is used for the connection from the client to the .xscontroller-managed Web Dispatcher (platform router) for access to the application instance.

Client(s)	Service	TCP Port	Use
Application ports 51000-51500	Instances	dynamic, in the range 50000-50999	Internal HTTP(S) This port range is used in **single-host scenarios** for the connection from the xscontroller-managed Web Dispatcher (platform router) to the application instances.
Application ports 51000-51500	Host-internal Web Dispatcher	dynamic, in the range 50500-50999	Internal HTTP(S) This port range is used in **multihost scenarios** for the connection from the xscontroller-managed Web Dispatcher (platform router) to the host-internal platform router (host-specific Web Dispatcher).
Host-internal Web Dispatcher	Instances	dynamic, in the range 50000-50499	Internal HTTP(S) This port range is used in **multihost scenarios** for the connection from the host-internal platform router (host-specific Web Dispatcher) to the application instance.
xsexecagent	xscontroller	3xx29	Internal HTTP(S) These ports are used for the connection between the xs execution agent and the xscontroller.
xscontroller	xsexecagent	system	

Client(s)	Service	TCP Port	Use
Application UI (browser, mobile, and so on)	Application instances	3xx33	Web Dispatcher HTTP(S) This port is used for the xscontroller-managed Web Dispatcher where routing is done by host names instead of ports. In this case, the xscontroller is available with URL `https://api.<example.com>:3xx33` and the xsuaaserver is available with URL `https://uaa-server.<example.com>:3xx33`. You specify the routing method - ports or host names - during installation. You can subsequently change the routing method in the `communication` section of the `xscontroller.ini` file. For more information, see SAP Note 2245631.

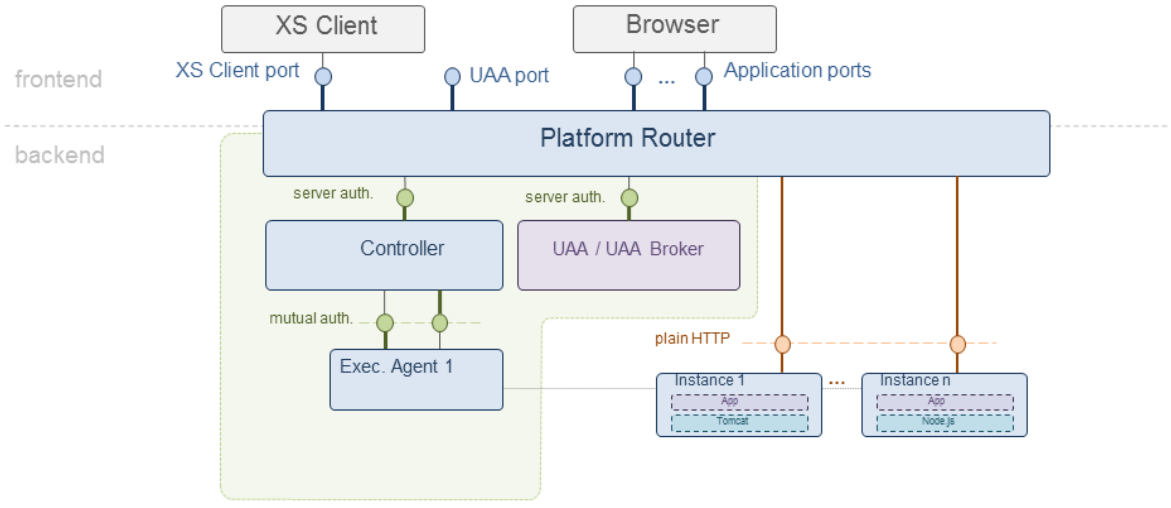

XSA Ports in a Single-Host Scenario

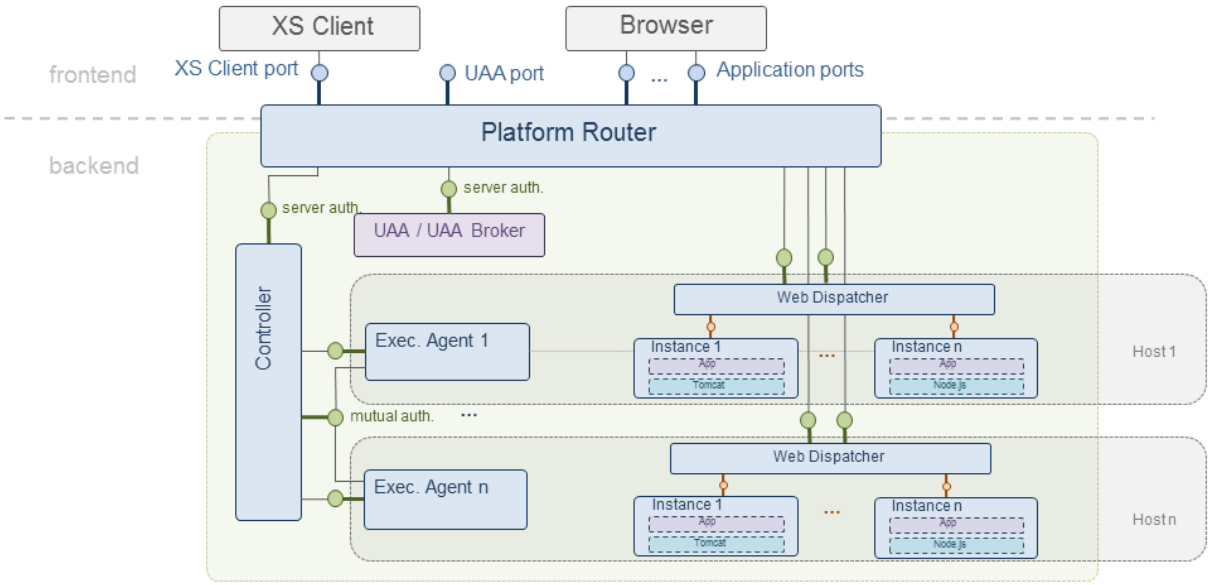

XSA Ports in a Multihost Scenario

For more information, see the section "Network and Communication Security with SAP HANA XS Advanced" in the *SAP HANA Security Guide*.

Related Information

SAP Note 2245631

4.3.1.2 Host Name Resolution

4.3.1.2.1 Default Host Names and Virtual Host Names

The assignment of multiple host names to the same host supports performance optimization as well as the security of your SAP HANA system. Moreover, some cluster managers and third-party backup tools as well as SAP Landscape Virtualization Management work on the basis of virtual host names or IP aliases.

Default Host Names

The default host names if nothing else is configured during the installation of SAP HANA are the host names defined at operating system level. The installation extracts the host names known to the operating system (that is, the names of the SAP HANA instances) and stores them in the sapstart service profiles, that is, in the following files:

```
/usr/sap/sapservices
/usr/sap/<SID>/HDB<instance_number>/<hostname>/sapprofile.ini
```

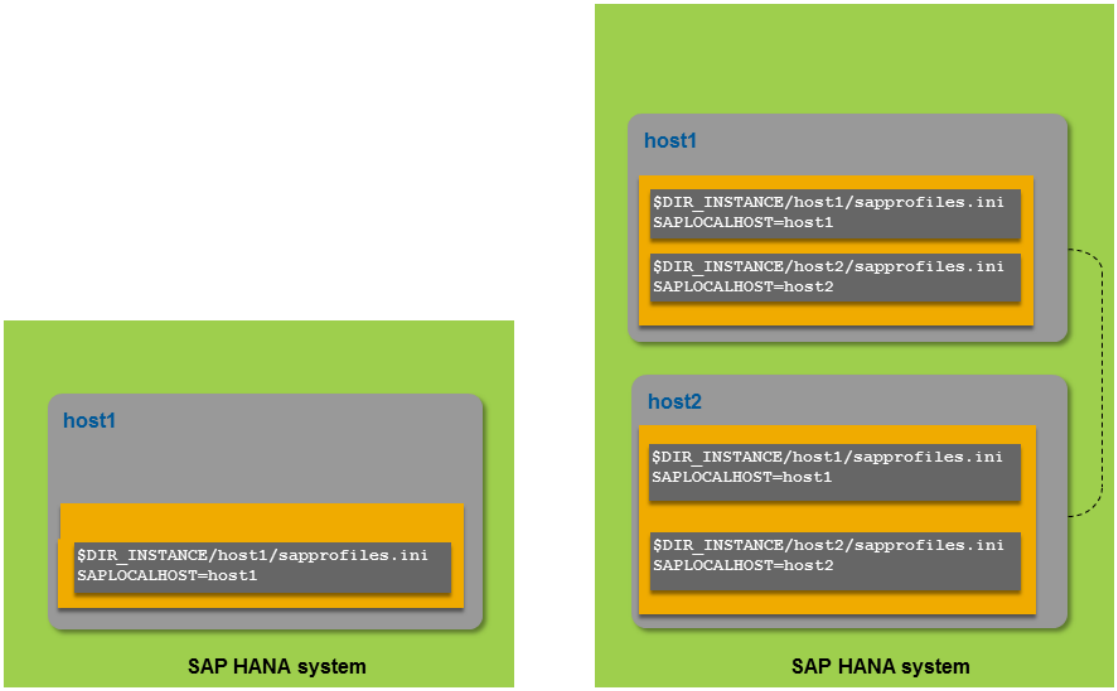

Example of Default Host Names for SAP HANA

These host names are then used for all internal communications between the SAP HANA services (nameserver, indexserver, and so on) and the SAP start service (`sapstartsrv`). In addition, SAP HANA system views with a HOST column show these host names.

Virtual Host Names

Another approach is to specify alternative host names during installation. These are referred to as virtual host names. Virtual host names must also be unique across multiple SAP HANA systems if more than one data center or site is used.

Host names specified in this manner must be resolvable during installation time as well as when SAP HANA is in operation. This is achieved, for example, by adding an `<ip> <hostname>` line to the operating system file `/etc/hosts` that contains the hostname-to-IP address mappings for the TCP/IP subsystem. Here is an example of what this might look like at operating system level for one host:

```
127.0.0.1       localhost
10.68.91.226    virtualhost1.wdf.sap.corp virtualhost1
```

Virtual host names are assigned as part of the installation process with the platform LCM command-line tool `hdblcm` using the `hostname` parameter. For more information about using the command-line tool or the `hostname` parameter, see the topics "Use the Command-Line Interface to Perform Platform LCM Tasks" and "hostname" in the *SAP HANA Server Installation and Update Guide*.

The `<virtualhostname>` is then stored as the internal host name in the sapstart service profiles and shows up in the HOST column of any system view.

It is also possible to assign virtual host names once the system is up and running, by using the platform LCM action `system_rename` with the `hostmap` parameter. For more information about mapping hosts, see the topics "Rename an SAP HANA System Host" and "Parameter Reference: Register and Rename" in the *SAP HANA Administration Guide*.

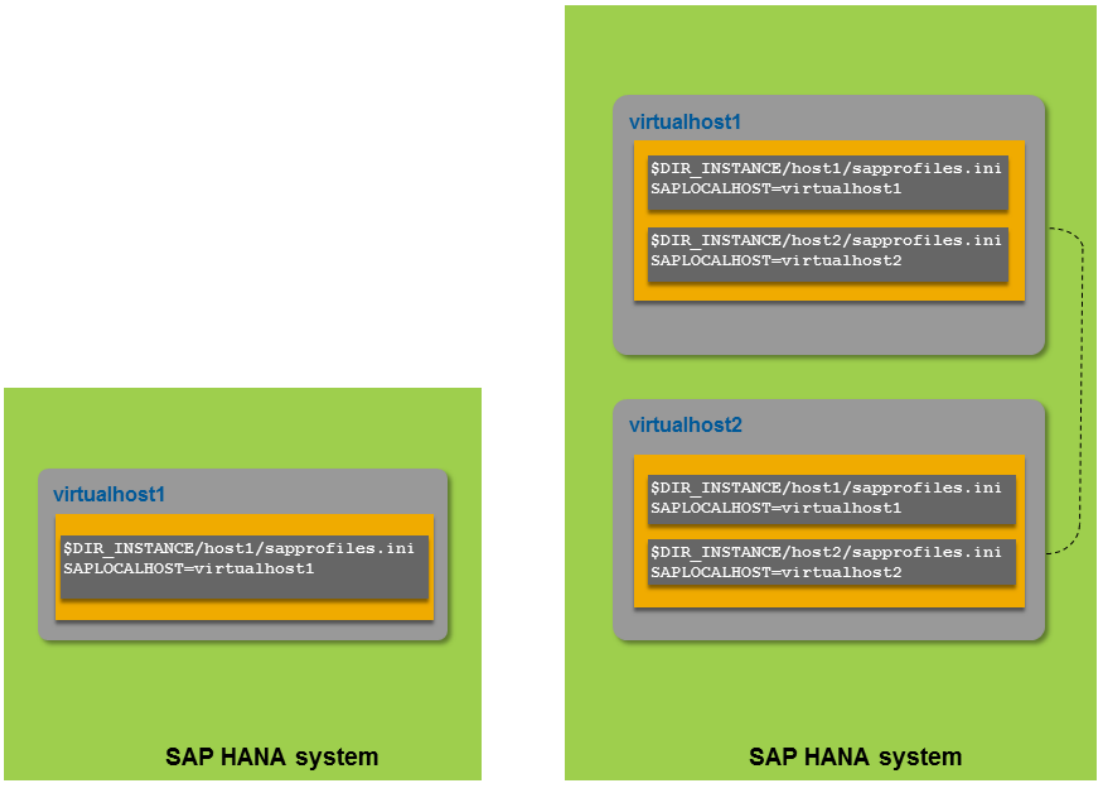

Example of Virtual (Internal) Host Names for SAP HANA

Distributed Landscapes

In multiple-host systems used for scale-out, the host names of all hosts must be known to each SAP HANA host. The `/etc/hosts` file for each host must include the corresponding lines:

```
host1
127.0.0.1       localhost
10.68.91.226    virtualhost1.wdf.sap.corp virtualhost1
10.68.91.227    virtualhost2.wdf.sap.corp virtualhost2
```

```
host2
127.0.0.1       localhost
10.68.91.226    virtualhost1.wdf.sap.corp virtualhost1
10.68.91.227    virtualhost2.wdf.sap.corp virtualhost2
```

Related Information

Internal Host Name Resolution [page 62]
Host Name Resolution for System Replication [page 64]

4.3.1.2.2 Internal Host Name Resolution

The SAP HANA services use IP addresses to communicate with each other. Host names are mapped to these IP addresses through internal host name resolution, a technique by which the use of specific and/or fast networks can be enforced and communication restricted to a specific network.

Single Host Versus Multiple Hosts

For single-host systems, no additional configuration is required. The services listen on the loopback interface only (IP address 127.0.0.1). In the `global.ini` files, the `[communication] listeninterface` is set to **`.local`**:

```
global.ini
[communication]
listeninterface=.local
```

In a distributed scenario with multiple hosts, the network needs to be configured so that interservice communication is operational throughout the entire landscape. In this setup, the host names (these could be virtual host names) of all hosts must be known to each other and thus to the SAP HANA system. This can be achieved by manually adding all hosts to each `/etc/hosts` file on the operating system of each host.

A distributed system can run with or without a separate network definition for interservice communication.

Distributed System Without a Separate Internal Network

If no separate network is defined for internal communication, the SAP HANA services listen on all available network interfaces. In the `global.ini` file, the listening interface is set to **`.global`**:

```
global.ini
[communication]
listeninterface=.global
```

> ⚠ Caution
>
> If the `listeninterface` parameter is set to **`.global`**, we strongly recommend that you secure the SAP HANA servers with additional measures such as a firewall and/or SSL. Otherwise, the internal service ports of the system are exposed and can be used to attack SAP HANA.

Distributed System with a Separate Internal Network

A distributed system can be configured with a dedicated internal network in either of the following ways:

- at installation time, using the `HDBLCM` **command line option** as in the following example:

  ```
  /<installation medium>/DATA_UNITS/HDB_LCM_LINUX_X86_64/hdblcm --
  internal_network=10.66.128.0/20
  ```

- post installation, using the **resident HDBLCM** from the GUI, command-line, or Web user interface. The following example, in command-line mode, binds the processes to this address only and to all local host interfaces. This option requires an internal network address entry:

  ```
  /<sapmnt>/<SID>/hdblcm/hdblcm --action=configure_internal_network --
  listen_interface=internal --internal_address=10.66.8/21
  ```

  For more information, see "Configuring SAP HANA Inter-Service Communication" in the Lifecycle Management chapter of the *SAP HANA Administration Guide*.
  Post-installation configuration as described here is to be done by SAP HANA system administrators who have root credentials or on SAP HANA systems where SSH is configured. If root privileges or SSH are not available, you can still perform network configuration but you will need to use a host-by-host approach, also known as decentralized execution. In this case, see SAP Note 2048681.

SAP HANA automatically chooses on each host a network interface within the allowed network mask. If the network interface is defined as `.internal` in the `global.ini` file as described above, the SAP HANA services listen on this interface only:

```
global.ini
[communication]
listeninterface=.internal
```

Only the SAP start service (`sapstartsrv`) still listens on all interfaces, to accept `start/stop/...` commands from outside the SAP HANA system.

This illustration shows a simple example of how a separate internal network might be configured for an SAP HANA database with three hosts:

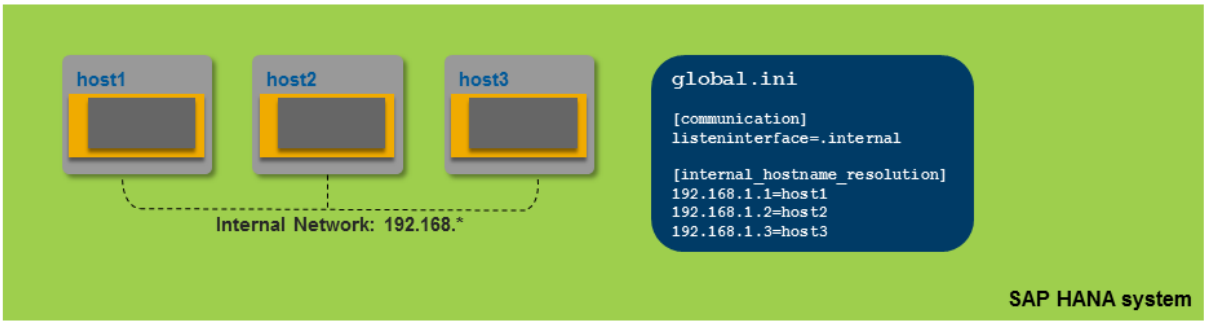

Simple Example of a Separate Internal Network for a Distributed SAP HANA System

For a more complex example, see "Host Name Resolution for System Replication".

For more information about configuring the network for multiple hosts, see the *SAP HANA Administration Guide*.

For information about the security of internal networks, see "Secure Internal Communication" and the "Security Configuration Checklist" for networks in the *SAP HANA Security Guide*.

> **Note**
>
> SAP HANA internal communication has sometimes been unofficially referred to as TREXNet communication. However, the term TREXNet is not valid in the context of SAP HANA.

Related Information

Host Name Resolution for System Replication [page 64]
SAP Note 2048681

4.3.1.2.3 Host Name Resolution for System Replication

The correct mapping of internal host names between primary and secondary systems is required for system replication.

With SAP HANA system replication, each SAP HANA instance communicates on the service level with a corresponding peer in the secondary system to persist the same data and logs as in the primary system. The replication of the transactional load can be configured to work in synchronous or asynchronous mode, depending mainly on the distance between the two sites. For a full description of system replication, see the *SAP HANA Administration Guide* and the white paper *Introduction to High Availability for SAP HANA*.

Communication between the primary and the secondary system is based on internal host names. The host names of the other site must always be resolvable, for example, through configuration in SAP HANA or corresponding entries in the `/etc/hosts` file.

To enforce specific networks and to avoid issues with host name resolution (for example, because SAP HANA was installed with short names) without the need to adapt the entries in the `/etc/hosts` file, internal host names can be mapped to IP addresses in the `global.ini` file as follows:

```
global.ini
[system_replication_hostname_resolution]
<ip-address_site>=<internal-host-name_site>
<...>
<...>
<...>
```

The inclusion of the internal host names ensures that each site can resolve the host name of the other site. Listing the hosts from all sites, also in multitier setups, ensures that the replication chain can switch seamlessly in the event of a takeover.

The entries in the `[system_replication_hostname_resolution]` section are used in combination with the `listeninterface` parameter in the `[system_replication_communication]` section. The following combinations are possible:

`[system_replication_communication] listeninterface`	`[system_replication_hostname_resolution]`	Additional Information
`.global`	No mappings specified	This is the default if nothing is specified. The default network route is used for system replication communication. This is normally the public network. **⚠ Caution** If you use a public network instead of a separate network, you **must** secure this connection with additional measures such as a firewall or a virtual private network and/or SSL.
`.global`	Entries for all hosts of neighboring sites (minimum) or for all hosts of own site as well as for all hosts of neighboring sites	A separate network is used for system replication communication. **➔ Tip** For three-tier setups, this is how you can use a dedicated network for system replication communication.
`.internal`	Entries for all hosts of own site as well as for all hosts of neighboring sites	A separate network is used for system replication communication. The primary hosts listen on the dedicated ports of the separate network only, and incoming requests on the public interfaces are rejected. **⚠ Caution** As of SAP HANA SPS 11, network communication for system replication with `listeninterface=.internal` is supported for two-tier replication but **not for three-tier setups**.

The parameters in the `global.ini` file must be set prior to registering the secondary system, because the `–sr_register` command uses this mapping. Registration is one step in the process of configuring the

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

> **Note**
>
> The `listeninterface=.global/.internal` parameter in the `[communication]` section is required for the communication between SAP HANA services (name server, index server, and so on) in a distributed system; it has no impact on system replication.

Examples

The following examples shows the host name resolution configuration for system replication to a secondary site using a dedicated network, with internal networks for the communication between hosts at each site:

- Public network 10.0.1.*
- Network for internal SAP HANA communication: 192.168.1.*
- Dedicated network for system replication: 10.5.1.*

In the first example, the `[system_replication_communication]listeninterface` parameter has been set to **`.global`** and the neighboring hosts are specified.

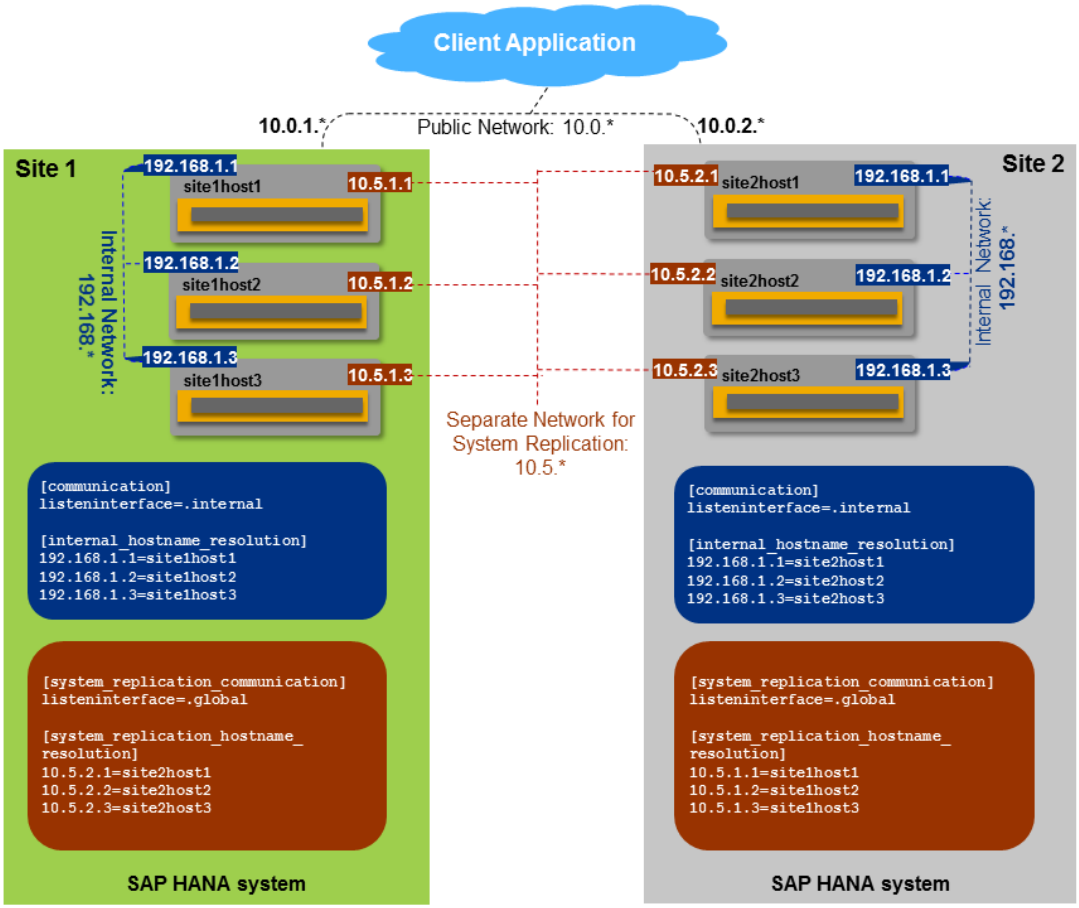

In the next example, the `[system_replication_communication]listeninterface` parameter has been set to `.internal` and all hosts are specified.

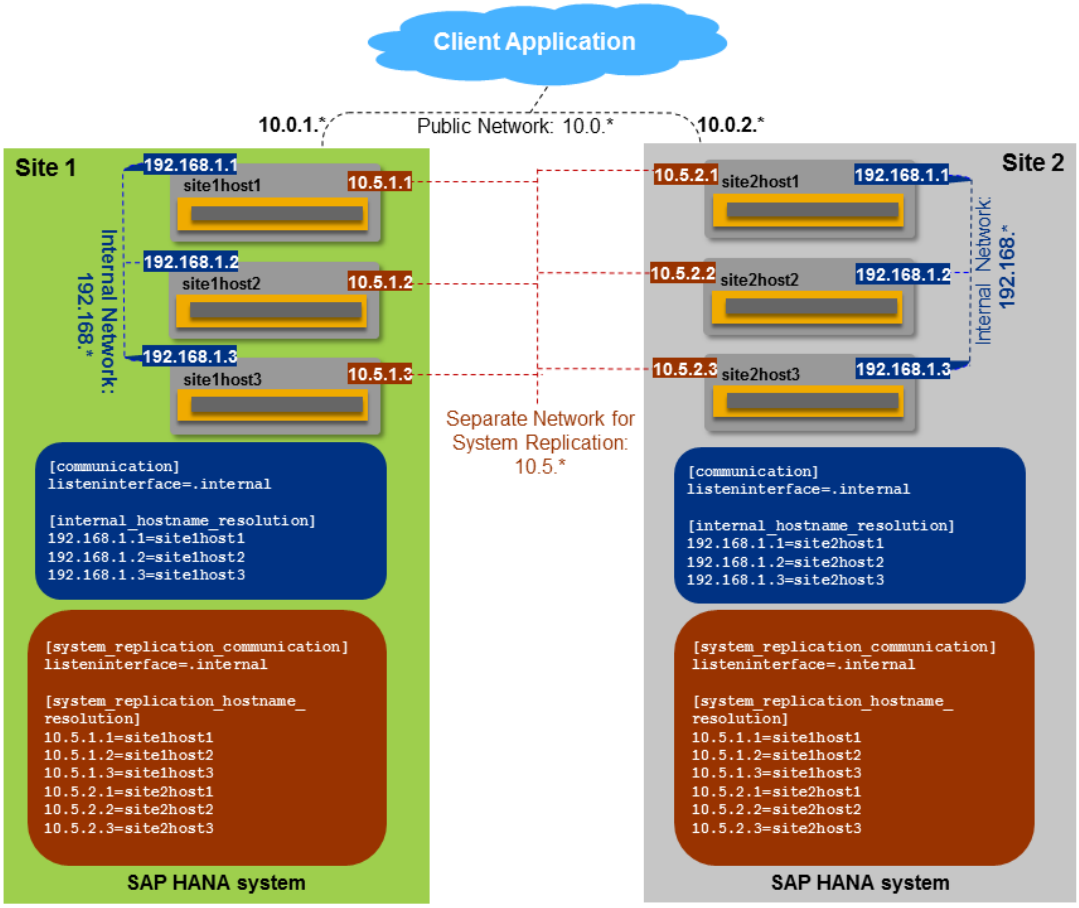

Important Security Information

- If you use a public network instead of a separate internal network and/or the parameter `listeninterface=.global`, you **must** secure this connection with additional measures such as a firewall or a virtual private network and/or SSL.
- If no separate internal network channel is configured for SAP HANA system replication, the `allowed_sender` parameter can be used to restrict communication between the primary and secondary sites to certain hosts. For this purpose, the following settings are made in the `global.ini` file on the primary site:

  ```
  global.ini
  [system_replication_communication]
  allowed_sender=<list of IP-addresses of secondary or CIDR-netmasks>
  ```

  An example of this parameter value would be `10.0.1.0/30`. The default is no restriction.
- For more security-related information, see "Secure Internal Communication Between Sites in System Replication Scenarios" and the "Security Configuration Checklist" for networks in the *SAP HANA Security Guide*.

Related Information

White paper "Introduction to High Availability for SAP HANA"
SAP Note 2036111: Configuration parameters for SAP HANA (including system replication)
Network Configuration for SAP HANA System Replication

4.3.1.2.4 Host Name Resolution for SQL Client Communication

Client applications communicate with SAP HANA servers from different platforms and types of clients via a client library (such as SQLDBC, JDBC, ODBC, DBSL, ODBO or ADO.NET) for SQL or MDX access.

In distributed systems, the application has a **logical connection** to the SAP HANA system: that is, the client library may in fact use multiple connections to different servers or change to a different underlying connection. The client library supports load balancing and minimizes communication overhead by:

- Selecting connections based on load data
- Routing statements based on information about the location of data

> **Note**
>
> Communication with SAP HANA hosts from a Web browser or a mobile application is requested using the HTTP protocol, which enables access to SAP HANA Extended Application Services (SAP HANA XS).

Public Host Name Resolution

An SQL client library always connects to the first available host specified in the connect string. From this host, the client library then receives a list of all the hosts. During operations, statements may be sent to any of these hosts.

By default, the IP address of the primary network interface is returned to the clients, as configured in the following parameter:

```
global.ini
[public_hostname_resolution]
use_default_route=ip
```

This works as long as there is only one external network. If a hostname or IP address is unresolvable, the client library falls back on the host names in the connect string:

- In single-host systems, the user doesn't normally notice this. In rare cases, the connection attempt does not fail immediately but waits for a tcp timeout, making the first statement run very slowly.
- In distributed systems, performance is impaired because statements must first be sent to the initial host and then forwarded on the server side to the right host.

For more information, see Related Information below.

Connect String with Multiple Hostnames

In a distributed SAP HANA system consisting of more than one host, a list of hosts (host:port) is specified in the SQL client library connect string.

The connect string for JDBC, for example, could look like this:

```
jdbc:sap://host1:30015;host2:30015;host3:30015/
```

All hosts that could become the active master, because they are one of the three configured master candidates, must be listed in the connect string to allow an initial connection to any of them in the event of a host auto-failover. A host auto-failover is an automatic switch from a crashed host to a standby host in the same system. One (or more) standby hosts are added to a SAP HANA system and configured to work in standby mode. As long as they are in standby mode, these hosts do not contain any data and do not accept requests or queries. When an active (worker) host fails, a standby host automatically takes its place.

Inclusion of the standby hosts in the connect string is mandatory if they are master candidates, otherwise optional.

The client connection code (ODBC/JDBC) uses a "round-robin" approach to reconnection, ensuring that the clients can always access the SAP HANA database, even after failover.

The following diagram illustrates how host auto-failover works. An active host fails (in this example, Host 2), and the standby host takes over its role by starting its database instance using the persisted data and log files of the failed host.

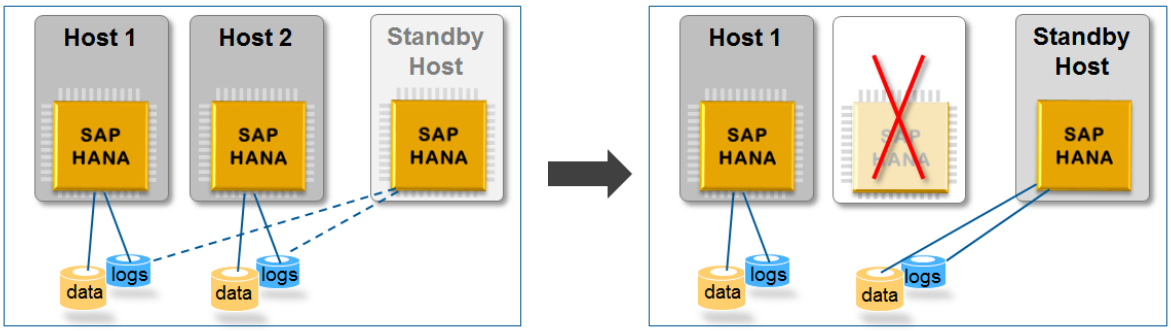

Example of Auto Host-Failover

One way to look up the master candidates in your distributed SAP HANA database is to use the following SQL statement:

```
select HOST
from SYS.M_LANDSCAPE_HOST_CONFIGURATION
where NAMESERVER_CONFIG_ROLE like 'MASTER%'
order by NAMESERVER_CONFIG_ROLE
```

For more information about configuring clients for failover, see the *SAP HANA Administration Guide*.

Connect String for SAP HANA System Replication

If system replication is used, we recommend that you do **not** specify physical host names in the SQL client connect string. Otherwise, you would have to reconfigure all of your applications after a takeover. Instead, use a

virtual host name or **virtual IP address**, and manage it using an external cluster manager. This virtual host name or IP address must point to the active master host on the active primary site.

System replication takeover hooks can be implemented to provide notification about the takeover. For more information about takeover hooks and client connection recovery, see the *SAP HANA Administration Guide*.

Related Information

Mapping Host Names for Database Client Access [page 70]
SAP Note 1780950
SAP Note 1876398

4.3.1.2.4.1 Mapping Host Names for Database Client Access

Clients communicate with the database through external hostnames or external IP addresses. A default mapping of external hostnames to internal hostnames enables statement routing and automatic reconnection in the event of a failover.

By default, the IP address of the primary network interface is used but there may be situations where you need to change this configuration, such as for certain firewall configurations, network address translation (NAT) types, or multiple external networks. For this purpose, a `[public_hostname_resolution]` section in the `global.ini` file is used with:

```
use_default_route = ip # values: no,ip,name,fqdn
optional pattern mapping: map_<internal-prefix>* = <public-prefix>*<public-
suffix>
optional exact mapping: map_<internal-name> = <public-name>
```

If optional mappings exist, they are always considered regardless of the `use_default_route` parameter value. Exact mappings have higher priority than pattern mappings.

Each host identifies the network interface and thus the default route for the connection:

Description	Parameter	Example
IP address of the interface	*use_default_route* = **ip**	10.4.2.71
Host name of the interface	*use_default_route* = **name**	lnd8520
Fully qualified name of the interface	*use_default_route* = **fqdn**	lnd8520.lnd.abc.corp
Disable feature and use internal host name	*use_default_route* = **no**	hananode01

In most cases, you do not need to configure anything. If you do need to configure something, see if you can use one of the default route mechanisms. You need to specify your own mapping only if the default route mechanisms do not fit your network requirements.

Example

Here are some examples of how you might customize this parameter:

```
[public_hostname_resolution]
map_hananode* = myservername*
```

```
[public_hostname_resolution]
map_hananode* = hananode*.lnd.abc.corp
```

```
[public_hostname_resolution]
map_hananode01 = 10.4.2.71
map_hananode02 = 10.4.2.72
map_hananode03 = 10.4.2.73
map_hananode04 = 10.4.2.74
```

```
[public_hostname_resolution]
map_hananode0* = 10.4.2.7*
map_hananode1* = 10.4.2.8*
```

Changes to configuration and default routes are checked once a minute and become effective within a minute after the SQL system management statement `ALTER SYSTEM ALTER CONFIGURATION ... WITH RECONFIGURE`.

4.3.1.2.4.2 SQL Connection Information for New Clients

It can be convenient for new SQL clients to be able to query the connectivity information of an existing client.

The connect string of the existing client was stored in the secure store and cannot be accessed. However, you can use the `global.ini/[communication]/sql_connect_hosts` parameter to record the connectivity information in the SAP HANA server so that it is available for the database connection from new clients. This information is a list of host names or IP addresses, which could be virtual host name or IP addresses, separated by commas.

The SAP HANA server does not use this parameter. It is used by applications and components that connect to SAP HANA. If the parameter is not filled, the application needs to consume the host values as follows:

```
select HOST
from SYS.M_LANDSCAPE_HOST_CONFIGURATION
where NAMESERVER_CONFIG_ROLE like 'MASTER%'
order by NAMESERVER_CONFIG_ROLE
```

In the following example, an ABAP system is installed on SAP HANA and the connection information is stored on the client side in the connect string including the logon data. This is the standard case. The ABAP client then sets the `sql_connect_hosts` parameter on the SAP HANA server. A System Landscape Directory (SLD) is subsequently installed on the same SAP HANA system. The SLD agent is able to look up the parameter in SAP HANA to find out the connection information. If the parameter values are missing, SLD uses the above SQL statement. If more clients are added, they follow the same procedure.

The example shows a single host but the parameter can also be useful in scenarios with multiple hosts.

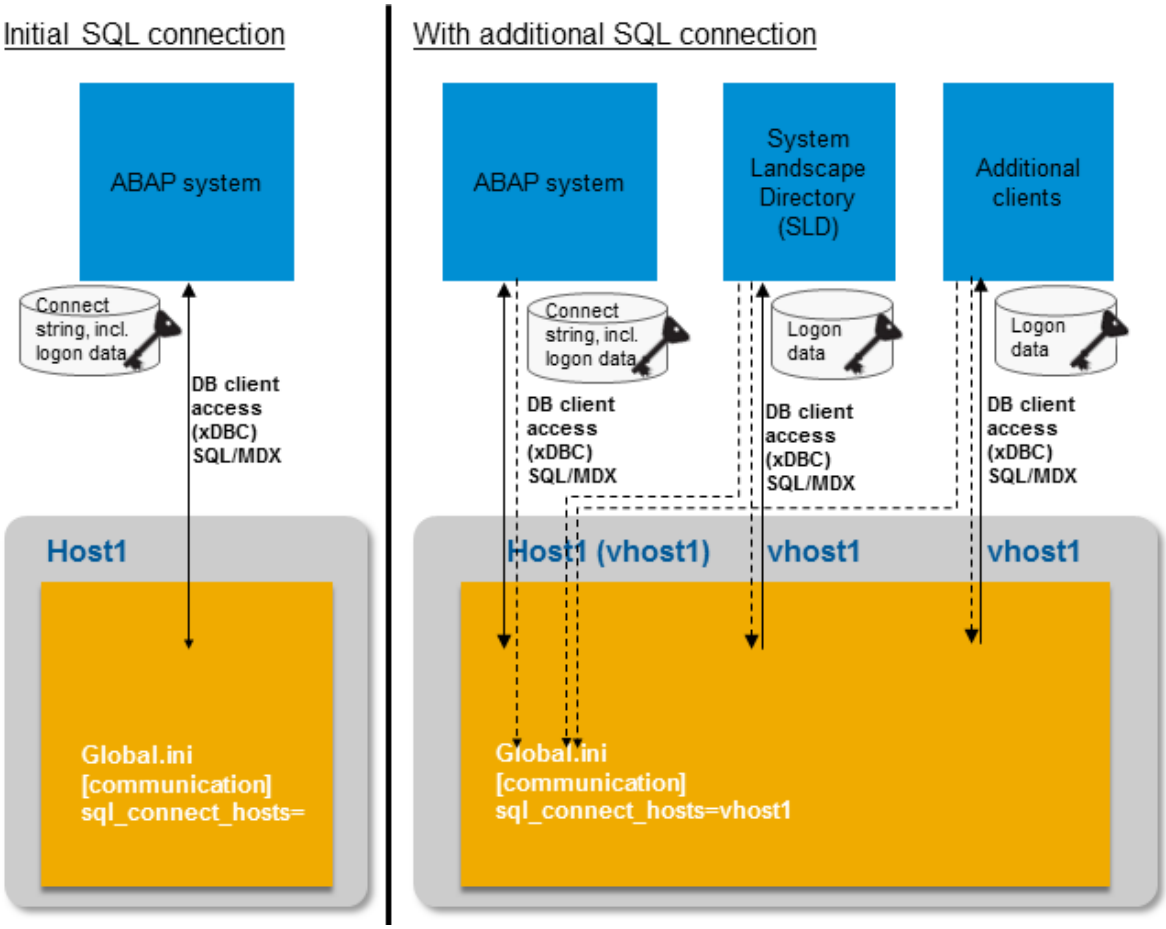
Initial SQL connection
ABAP system
Connect string, incl. logon data
DB client access (xDBC) SQL/MDX
Host1
Global.ini [communication] sql_connect_hosts=
With additional SQL connection
ABAP system
System Landscape Directory (SLD)
Additional clients
Connect string, incl. logon data
Logon data
Logon data
DB client access (xDBC) SQL/MDX
DB client access (xDBC) SQL/MDX
DB client access (xDBC) SQL/MDX
Host1 (vhost1)
vhost1
vhost1
Global.ini [communication] sql_connect_hosts=vhost1

5 SAP HANA Deployment Options

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

5.1 On-Premise

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

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

- As an **appliance**, SAP HANA combines software components from SAP-optimized on proven hardware provided by SAP's hardware partners. This approach is valid for Intel-based hardware platforms only.
- Compared with the appliance delivery approach, **SAP HANA Tailored Datacenter Integration** is a more open and flexible approach, designed to serve your needs regarding the integration of SAP HANA in the data center. The requirements for this deployment option are as follows:
  - The **server** is certified and belongs to the allowed hardware.
  - The **storage solution** has successfully passed SAP HANA hardware certification.
  - The components of SAP HANA can only be **installed** by certified hardware partners, or any person holding **certification** on certified hardware running an approved operating system. See the SAP Education resources for information about the SAP HANA certification exams.

> **Note**
>
> The supported hardware for SAP HANA depends on the deployment method used (appliance or TDI). For more information, see the *Related Information* in this section and under *SAP HANA Hardware and Software Requirements* in the *SAP HANA Master Guide*.

Related Information

Certified Appliance Hardware (also applicable for compute servers in TDI environments)
SAP Certified and Supported SAP HANA Hardware
SAP HANA Tailored Datacenter Integration (TDI)
SAP HANA TDI – FAQ
SAP HANA Tailored Data Center Integration 2017 - Overview
SAP Certified Enterprise Storage Hardware for SAP HANA
SAP Approved Hardware for Entry-level Compute Servers for SAP HANA
SAP Note 1943937 - Hardware Configuration Check Tool - Central Note
SAP Note 1900823 - SAP HANA Storage Requirements and SAP HANA Storage Connector API
SAP HANA Network Requirements

SAP HANA on IBM Power Systems

SAP Note 2055470 - HANA on POWER Planning and Installation Specifics - Central Note

SAP Note 2218464 - Supported products when running SAP HANA on IBM Power Systems

General Links

SAP Training & Education

5.2 In the Cloud

In the **cloud**, SAP HANA is offered as a comprehensive infrastructure combined with managed services.

SAP HANA is deployed through the following cloud offerings:

- SAP HANA Cloud Platform
  - SAP HANA Infrastructure Services
    High-performance cloud infrastructure to quickly deploy existing SAP HANA licenses
    Infrastructure-only on a monthly subscription (license bought separately)
  - SAP HANA DB Services
    Fully-featured SAP HANA hosted in the public cloud
    Infrastructure and license on a monthly subscription
  - SAP HANA App Services
    SAP HANA Platform-as-a-Service (PaaS) in a cloud environment
    Infrastructure and license on a monthly subscription
- SAP HANA Enterprise Cloud
  Enterprise-class SAP HANA managed cloud offering
  Infrastructure and managed services on a monthly subscription (license bought separately)
- SAP HANA One
  Fully-featured SAP HANA hosted in the public cloud
  Infrastructure and license on an hourly subscription

6 SAP HANA Implementation and Operation

The implementation and operation of SAP HANA depends on the chosen use case and the technical deployment.

6.1 Sizing SAP HANA

The SAP HANA database can be deployed as an SAP In-Memory Appliance (SAP HANA) or deployed following the SAP HANA Tailored Datacenter Integration (TDI) approach.

Memory Sizing

Every SAP HANA customer must perform memory sizing as the first step to sizing an SAP HANA deployment.

- For new SAP HANA implementations, it is necessary to size the memory for an SAP HANA system using the SAP Quick Sizer in Related Information.
- For systems that are migrating to SAP HANA we recommend
  - Using a Sizing Report on the source database if the migration is from a SAP NetWeaver based system.
  - Applying a sizing SAP Note, if the migration is from a non-SAP NetWeaver data source.

Any system that is very large or complex requires sizing from an SAP sizing expert.

For more information about memory sizing, we recommend the following Related Information:

Starting Point	Sizing for BW on HANA	Sizing for Suite on HANA	Sizing non-NetWeaver
Greenfield / new implementation	Quick Sizer Questionnaire for BW on HANA	Quick Sizer Questionnaire plus additional formula	Sizing guide as available
Migration	SAP Note 1736976 - Sizing Report for BW-on-HANA	• SAP Note 1872170 - Suite on HANA memory sizing • SAP Note 1793345 - SAP BW on HANA: Sizing SAP HANA Database	SAP Note 1514966 - SAP HANA 1.0: Sizing SAP In-Memory Database

The result of the memory sizing is the basis for the hardware recommendation for an SAP HANA system. If you decide to buy the In-Memory Appliance (HANA), you have a selection of certified appliances from certified hardware partners. You should check the *SAP Certified and Supported SAP HANA Hardware* for hardware that matches your memory sizing results. Ivy Bridge customers should check the SAP Community Network (SCN). For an In-Memory Appliance, you don't need to consider storage and CPU sizing, because they are inclusive in the certified appliance offering.

SAP HANA Tailored Datacenter Integration (TDI) Approach

TDI follows the approach of SAP sizing and mapping to authorized servers, network and storage. Depending on the chosen vendor, various configurations can be employed to ensure valid mapping. Contact your vendor for information about the options available. If you decide to build the SAP HANA system based on the SAP HANA TDI approach, you must become TDI certified.

For storage sizing recommendations, see the *SAP HANA Storage Requirements* whitepaper in Related Information.

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

Related Information

SAP Quick Sizer
SAP Note 1736976 - Sizing Report for BW-on-HANA
SAP Note 1872170 - Suite on HANA memory sizing
SAP Note 1793345 - Sizing for SAP Suite on HANA
SAP Note 1514966 - SAP HANA: Sizing SAP HANA
SAP Note 2055470 - SAP HANA on POWER Planning and Installation Specifics - Central Note
SAP Certified and Supported SAP HANA Hardware
SAP HANA TDI - Storage Requirements
SAP Certified Enterprise Storage Hardware for SAP HANA
Sizing Approaches for SAP HANA

6.2 Installing SAP HANA

The installation of SAP HANA comprises several steps. The central part is the installation of the SAP HANA Platform Edition. You need to check specific installation information, depending on the data provisioning technology you use or other components you want to add to your SAP HANA landscape.

SAP HANA Platform Edition

For information about installing SAP HANA, see the SAP HANA documentation on SAP Help Portal:

- *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 tools.

> **Note**
>
> SAP HANA installations are performed using the SAP HANA database lifecycle manager (HDBLCM). SAP HANA installations cannot be performed using the Software Provisioning Manager (SWPM).

- *SAP HANA Client Installation and Update Guide*
- *SAP HANA Studio Installation and Update Guide*

Data Provisioning Technologies

You can find the documentation for the data provisioning technologies on different publication channels:

- *SAP HANA Installation Guide – Trigger-Based Replication (SLT)*
- *SAP HANA Direct Extractor Connection Implementation Guide*
- *SAP HANA smart data access*
  SAP HANA smart data access is part of SAP HANA. However, it is not installed during the installation of the SAP HANA Platform Edition. For more information about installing SAP HANA smart data access, see *SAP HANA Smart Data Access* in the *SAP HANA Administration Guide*.
  For IBM Power Systems, only the following data sources are supported: SAP HANA, SAP IQ, SAP Adaptive Service Enterprise, and Oracle Database 12C. For more information, see *SAP HANA Smart Data Access* in the *SAP HANA Administration Guide*.
- ETL-Based Replication (SAP Data Services)
  See the *SAP Data Services* page on SAP Help Portal.
- Log-Based Replication (SAP Replication Server)
  See the *SAP Replication Server* page on SAP Help Portal.
- SAP HANA smart data integration
  See the *SAP HANA Smart Data Integration* page on SAP Help Portal.

Related Information

SAP HANA Smart Data Access [page 24]
SAP Data Services
SAP Replication Server
SAP HANA Smart Data Integration and SAP HANA Smart Data Quality
SAP HANA Platform
SAP HANA Real-Time Replication

6.3 Administrating SAP HANA

The SAP HANA Technical Operations Manual and the SAP HANA Administration Guide provide information about administering and operating your SAP HANA system landscape.

6.4 SAP HANA Lifecycle Management

SAP HANA lifecycle management covers two aspects: platform lifecycle management for customizing and updating your SAP HANA platform and application lifecycle management for managing SAP HANA content products and transports.

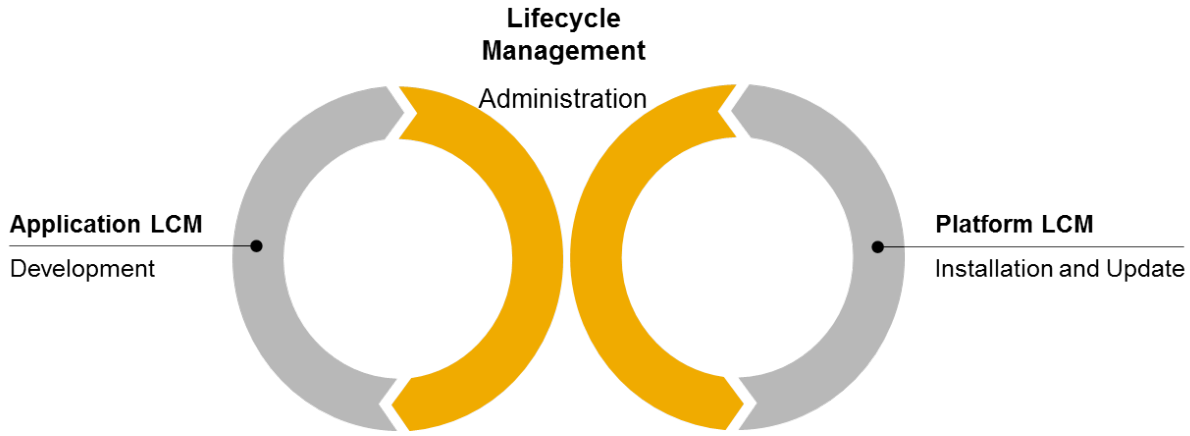

Platform Lifecycle Management Aspects

You can customize platform lifecycle management aspects of your SAP HANA system by accessing the SAP HANA database lifecycle manager from three user interfaces: the graphical user interface, the command-line interface, or the Web user interface in a stand-alone Web browser, in the SAP HANA studio, or via the SAP HANA cockpit.

SAP HANA platform lifecycle management encompasses the installation and update of an SAP HANA server, mandatory components, and additional components, as well as the post-installation configuration. The concepts and procedures for SAP HANA platform installation and update are described in the *SAP HANA Server Installation and Update Guide* on SAP Help Portal.

A number of system configuration features are integrated into the SAP HANA database lifecycle manager, such as:

- The initial configuration of your SAP HANA platform to integrate it into your landscape. For example, by registering it in a system landscape directory, or configuring the inter-service communication.
- Adapting the topology of your SAP HANA platform by adding or removing additional SAP HANA hosts.
- Reconfiguring the system. For example, by renaming your SAP HANA system, relocating the system to different hardware, or converting the system to a multiple-container enabled system.

System configuration as it pertains to SAP HANA lifecycle management is described in the *SAP HANA Platform Lifecycle Management* section of this *SAP HANA Administration Guide*.

Application Lifecycle Management Aspects

SAP HANA application lifecycle management aspects can be accessed in different user interfaces: an interface that runs as an SAP HANA XS application in a web browser, a command-line tool hdbalm, integrated in SAP HANA studio, or via the SAP HANA cockpit.

SAP HANA application lifecycle management supports you in all phases of the lifecycle of an SAP HANA application or add-on product, from modelling your product structure, through application development, transport, assembly, to installing and updating products that you have downloaded from SAP Support Portal or which you have assembled yourself.

All application lifecycle management tasks are documented in the guide *SAP HANA Application Lifecycle Management* on SAP Help Portal.

System administrators use SAP HANA application lifecycle management mainly to install and update SAP HANA applications or add-on products. Therefore, these tasks are documented in this *SAP HANA Administration Guide*. Tasks related to SAP HANA development are documented in the *SAP HANA Developer Guide - For SAP HANA Studio* (on SAP Help Portal) under *SAP HANA Application Lifecycle Management*.

6.5 SAP HANA Content

SAP HANA content is structured in the way that delivery units (DUs) are used to group SAP HANA content artifacts (such as analytic, attribute or calculation views, and SQLScript procedures).

DUs are grouped to SAP HANA products in order to ship and install SAP HANA applications with all dependent artifacts (grouped in DUs). To distribute SAP HANA content, a product archive (*.ZIP file) or a delivery unit archive (*.tgz file) is used. There are various ways of acquiring and deploying these archive types.

SAP HANA content, which is developed on SAP HANA Extended Application Services (SAP HANA XS), can also be grouped in a DU.

For more information about SAP HANA content, see *Components Delivered as SAP HANA Content* in the *SAP HANA Security Guide*.

.

6.5.1 SAP HANA Archive Types

The difference between the various archive types is their method of deployment, and when the content is deployed.

The following archive types are available:

- **Product archive file** (`*.ZIP`)
  A product version archive is a `*.ZIP` file containing 1-n software component archive files and the following metadata files: `stack.xml`, `pd.xml`. A software component archive file is created for each DU containing its archive file (`*.tgz`).
  A product is usually the entity that delivers SAP HANA applications, but it can also be used for transports. SAP HANA content that can be downloaded independently is shipped as SAP HANA products in SAP HANA product archives. SAP HANA content that is not part of the SAP HANA database is called SAP HANA content add-on (or SAP HANA product). SAP HANA content add-ons are developed as part of the SAP HANA platform or as part of an application that runs on top of SAP HANA.
  For information about how to deploy a product archive, see *Deploy a Product Archive (*.ZIP)*.
- **Software Component Archive** (`*.ZIP`)
  A software component archive is a `*.ZIP` file (in previous versions also `*.SAR` files were delivered as software component archives) containing one delivery unit archive file (`*.tgz`) and (optionally) a corresponding translation DU and the metadata file `SL_MANIFEST.XML`. A software component archive can be deployed with the same tool as product archives.
  For information about how to deploy a software component archive, see *Deploy a Product Archive (*.ZIP)*.
- **Delivery unit archive file** (`*.tgz`)
  A delivery unit archive is a `*.tgz` file containing the SAP HANA content artifacts that are created in the SAP HANA repository. A DU is used to deliver one or more software components from SAP (or a partner) to a customer.
  For distribution using export/import and deployment, a DU is contained in a delivery unit archive (`*.tgz` file). It contains the objects and packages of a DU together with the metadata file `manifest.txt`. The transport is also offered at DU level.
  The following types of delivery unit archive files are available:
  - **Delivery unit archives as part of the SAP HANA database**
    The following types of delivery unit archive files that are part of the SAP HANA database are available:
    - **Automated content** is installed together with SAP HANA and imported into the SAP HANA repository during installation. This is an integral part of the SAP HANA database and is used by every SAP HANA database customer.
      Automated content is located on the SAP HANA system in the following folder: `/usr/sap/<SID>/SYS/global/hdb/auto_content`.
    - **Non-automated content** is installed with SAP HANA, but needs to be imported into the SAP HANA repository manually by the system administrator. It is used for integral parts of the SAP HANA database, but is only used by a small number of customers.
      Non-automated content is located on the SAP HANA system in the following folder: `/usr/sap/<SID>/SYS/global/hdb/content`.
      Delivery unit archives that are non-automated content of the SAP HANA database need to be deployed manually.
  - **Independent delivery unit archives that are not part of the SAP HANA database**
    Delivery unit archives that are not installed together with the SAP HANA database and are not part of the SAP HANA database need to be deployed manually.

For information about how to deploy or activate a delivery unit archive, see *Deploy a Delivery Unit Archive (*.tgz)*.

Related Information

Deploy a Product Archive (*.ZIP) [page 81]
Deploy a Delivery Unit Archive (*.tgz) [page 81]

6.5.2 Deploy a Product Archive (*.ZIP)

SAP HANA application lifecycle management provides a method of deploying a product archive file (`*.ZIP` file containing a product) or software component archive files (`*.ZIP`).

For more information, see *Installing and Updating SAP HANA Products and Software Components* in the *SAP HANA Application Lifecycle Management Guide*.

6.5.3 Deploy a Delivery Unit Archive (*.tgz)

The following deployment methods for deploying a delivery unit archive file (*.tgz file containing a DU) are provided:

- SAP HANA Application Lifecycle Management
  Choose *Products* > *Delivery Units* > *Import*.
  This tool runs on the SAP HANA XS Web server.
  For more information, see *Import a Delivery Unit* in the *SAP HANA Developer Guide (For SAP HANA Studio)*.
- SAP HANA Application Lifecycle Management
  SAP HANA application lifecycle management provides functions for installing and updating SAP HANA products:
  - SAP Fiori application integrated in the SAP HANA Application Lifecycle Management XS application
  - `hdbalm` command line tool

  For more information, see *Installing and Updating SAP HANA Products and Software Components* in the *SAP HANA Application Lifecycle Management Guide*.
- SAP HANA studio
  Import function of the SAP HANA Modeler
  Choose *File* > *Import* > *SAP HANA Content* > *Delivery Unit*.

For more information, see *SAP HANA Modeling Guide*.

7 Appendix

The appendix provides additional information.

7.1 Related Information

Links are provided to documentation on planning your deployment that is useful to know but not necessarily directly connected to SAP HANA.

Content	Location
SAP HANA 1.0 SPS 12 documentation	https://help.sap.com/viewer/product/SAP_HANA_PLATFORM/1.0.12/en-US
Previous documentation sets for SAP HANA	https://help.sap.com/viewer/product/SAP_HANA_PLATFORM/1.0/en-US
Sizing, calculation of hardware requirements, such as CPU, disk, and memory resources	https://service.sap.com/sizing
Sizing, calculation of hardware requirements - such as CPU, disk and memory resources - with the QuickSizer tool	https://service.sap.com/quicksizer
Released platforms and technology-related topics such as maintenance strategies and language support – Platform Availability Matrix (PAM)	https://support.sap.com/pam
Security	https://service.sap.com/security
Unicode SAP systems and their availability	https://service.sap.com/unicode

To get the full access to the SAP Service Marketplace or the SAP Support Portal you need an authorized user ID to access this information. Follow the instructions on the corresponding page how to register for a login.

The following table lists further useful links:

Content	Location
SAP incident wizard	https://support.sap.com/incident
SAP Notes search	https://support.sap.com/notes
SAP Software Download Center – software download and ordering of software	https://support.sap.com/swdc

7.2 Important SAP Notes

Read the following SAP Notes before you start the installation. These SAP Notes contain the latest information about the installation, as well as corrections to the installation documentation.

Make sure that you have the most up-to-date version of each SAP Note, which you can find on SAP Service Marketplace at https://service.sap.com/notes.

SAP Note Number	Title
1514967	SAP HANA: Central Note
2298750	SAP HANA Platform SPS 12 Release Note
1523337	SAP HANA Database: Central Note
2000003	FAQ: SAP HANA
1944799	SAP HANA Guidelines for SLES Operating System
1824819	SAP HANA DB: Recommended OS settings for SLES 11 / SLES for SAP Applications 11 SP2
1954788	SAP HANA DB: Recommended OS settings for SLES 11 / SLES for SAP Applications 11 SP3
2240716	SAP HANA DB: Recommended OS settings for SLES 11 / SLES for SAP Applications 11 SP4
2205917	SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12
2009879	SAP HANA Guidelines for Red Hat Enterprise Linux (RHEL)
2013638	SAP HANA DB: Recommended OS settings for RHEL 6.5
2136965	SAP HANA DB: Recommended OS settings for RHEL 6.6
2247020	SAP HANA DB: Recommended OS settings for RHEL 6.7
2292690	SAP HANA DB: Recommended OS settings for RHEL 7.2
2055470	HANA on POWER Planning and Installation Specifics - Central Note
2218464	Supported products when running SAP HANA on IBM Power Systems
2235581	SAP HANA: Supported Operating Systems
52505	Support after end of mainstream/extended maintenance

SAP Note Number	Title
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
1828400	SAPUI5 tools disappeared after updating HANA Studio
1917938	Migrating the Statistic Server During Update to SPS 07
1927949	Standard Behavior for SAP Logon Tickets
1577128	Supported clients for SAP HANA
1514966	SAP HANA: Sizing SAP HANA Database
1637145	SAP BW on HANA: Sizing SAP HANA Database
1793345	Sizing for Suite on HANA
1824819	Optimal settings for SLES 11 SP2 and SLES 11 for SAP SP2
1597355	Swap space recommendation for Linux For the SAP HANA system there are no special requirements regarding swap space.

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* SAP HANA Data Center Intelligence
- *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-IOA* SAP IT Operations Analytics

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-FLY* SAP HANA InA Firefly
- *HAN-AS-INA-SVC* SAP HANA InA Service
- *HAN-AS-INA-UI* SAP HANA InA Toolkit, Fiori Search UI
- *HAN-AS-MDS* SAP HANA Multidimensional Service
- *HAN-AS-RPO* SAP HANA Repository

- *HAN-AS-RST* SAP HANA Development Environment REST API
- *HAN-AS-RUL* SAP HANA Rules Framework
- *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
- *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-TM* SAP HANA Task management
- *HAN-AS-XSA-WF* SAP HANA Workflow

SAP HANA Accelerator for SAP ASE

- *HAN-ASE* SAP HANA Accelerator for SAP ASE

SAP HANA Cockpit

- *HAN-CPT* SAP HANA Cockpit
- *HAN-CPT-ADM* SAP HANA Administration Core
- *HAN-CPT-ASE* SAP HANA Accelerator for SAP ASE Administration
- *HAN-CPT-BAC* SAP HANA Backup and Recovery
- *HAN-CPT-CNR* SAP HANA Workload Capture and Replay
- *HAN-CPT-DCC* SAP DB Control Center
- *HAN-CPT-DP* Please use HAN-DP-SDI
- *HAN-CPT-DYT* SAP HANA Dynamic Tiering Administration
- *HAN-CPT-SDS* SAP HANA Smart Data Streaming Administration
- *HAN-CPT-SYN* SAP HANA Remote Data Sync Cockpit
- *HAN-CPT-UM* SAP HANA User Management
- *HAN-CPT-XS* Please use HAN-AS-XS-ADM

SAP HANA Database

- *HAN-DB* SAP HANA Database
- *HAN-DB-AFL* Please use subcomponents, see SAP Note 2198403
- *HAN-DB-AFL-DQ* SAP HANA Data Quality 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-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-DI* HANA Deployment Infrastructure (HDI)

- *HAN-DB-ENG* SAP HANA DB Engines
- *HAN-DB-ENG-BW* SAP HANA BW Engine
- *HAN-DB-ENG-GPH* SAP HANA Graph Engine
- *HAN-DB-ENG-GPH-API* SAP HANA Graph Engine API
- *HAN-DB-ENG-GPH-GEM* Graph Exploration and Manipulation (GEM) Language
- *HAN-DB-ENG-IM* Please use HAN-DB-SDQ
- *HAN-DB-ENG-PLE* SAP HANA Planning Engine
- *HAN-DB-ENG-SPA* SAP HANA Spatial Engine
- *HAN-DB-ENG-TXT* SAP HANA Text Engine
- *HAN-DB-EPM* SAP HANA Enterprise Performance Management 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
- *HAN-DB-LVC* SAP HANA integrated liveCache
- *HAN-DB-MDX* SAP HANA MDX Engine/Excel Client
- *HAN-DB-MON* SAP HANA Monitoring
- *HAN-DB-PER* SAP HANA Database Persistence
- *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

SAP HANA Data Provisioning Services

- *HAN-DP* SAP HANA Data Provisioning Services
- *HAN-DP-DS* SAP Data Services
- *HAN-DP-DXC* SAP HANA Direct Extractor Connector
- *HAN-DP-ESS* SAP HANA Enterprise Semantic Services (ESS)
- *HAN-DP-LTR* SAP Landscape Transformation Replication Server
- *HAN-DP-REP* SAP Sybase Replication Server
- *HAN-DP-SDI* SAP HANA smart data integration

SAP HANA Dynamic Tiering

- *HAN-DYT* SAP HANA Dynamic Tiering

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

SAP HANA Remote Data Sync

- *HAN-SYN* SAP HANA Remote Data Sync

SAP HANA Smart Data Streaming

- *HAN-SDS* SAP HANA Smart Data Streaming

SAP HANA Studio (Eclipse)

- *HAN-STD* SAP HANA Studio (Eclipse)
- *HAN-STD-ADM* SAP HANA Administration
- *HAN-STD-ADM-BAC* SAP HANA Backup and Recovery (Studio)
- *HAN-STD-ADM-DBA* SAP HANA Database Administration and Monitoring
- *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-CDS* SAP HANA Core Data Services Tools
- *HAN-STD-DEV-CDS-GRA* Please use HAN-STD-DEV-CDS
- *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 Component
- *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 Tools Platform / 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 Fiori Launchpad
- *HAN-STD-DEV-XS* SAP HANA XS Editors and Wizards

SAP HANA remote data sync

- *HAN-SYN* SAP HANA remote data sync

SAP HANA Vora

- *HAN-VO* SAP HANA Vora
- *HAN-VO-EN* SAP HANA Vora Engine
- *HAN-VO-SE* SAP HANA Vora Spark Extension Library

SAP HANA Web IDE

- *HAN-WDE* SAP HANA Web IDE
- *HAN-WDE-BLD* SAP Web IDE for Hana building applications
- *HAN-WDE-CHE* SAP Web IDE for Hana CHE
- *HAN-WDE-DBG* SAP Web IDE for Hana debugging applications
- *HAN-WDE-DOC* SAP Web IDE for Hana documentation
- *HAN-WDE-EDT* SAP Web IDE for Hana editor

- *HAN-WDE-EDT-CDS* SAP Web IDE for Hana editor for Core Data Services
- *HAN-WDE-EDT-MOD* SAP Web IDE editor for HANA Analytical Modeling
- *HAN-WDE-EDT-NJS* SAP Web IDE for Hana Node.js Tools
- *HAN-WDE-EDT-UI5* SAP Web IDE for Hana editor for UI5 applications
- *HAN-WDE-EIM* Flowgraph, RepTasks and other SDA Tools
- *HAN-WDE-GIT* SAP Web IDE for Hana GIT
- *HAN-WDE-INS* Installation SAP Web IDE for HANA
- *HAN-WDE-MTA* Multi Targeted Application in Web IDE
- *HAN-WDE-PREF* SAP Web IDE for Hana user and project settings
- *HAN-WDE-RTT* SAP Web IDE for Hana Runtime and SQL Tools
- *HAN-WDE-RUN* SAP Web IDE for Hana running applications
- *HAN-WDE-SDS* Smart Data Streaming Tools
- *HAN-WDE-SRC* Search
- *HAN-WDE-TPL* Project creation, template and wizards
- *HAN-WDE-XSC* Old SAP HANA Web IDE

SAP HANA XS Advanced

- *BC-XS* XS Engine (XS Advanced)
- *BC-XS-JAS* Java Runtime
- *BC-XS-JS* Javascript runtime
- *BC-XS-SEC* UAA and Security for XS engine
- *BC-XS-SRV* XS Engine Services and Administration

SAP HANA Database (CCMS, Porting and DB Interface)

- *BC-DB-HDB-CCM* CCMS for SAP HANA
- *BC-DB-HDB-POR* DB Porting for SAP HANA
- *BC-DB-HDB-SYS* SAP HANA database interface/DBMS

End User Clients

- *BI-BIP, BI-BIP-CMC* Business intelligence platform (formerly known as BOE)
- *BI-RA-EXP* SAP BusinessObjects Explorer
- *BI-RA-CR, BI-BIP-CRS* SAP Crystal Reports
- *BI-RA-XL* Dashboard Designer
- *BI-BIP-IDT* Information design tool
- *BI-RA-WBI* Web Intelligence
- *BI-RA-MULTIDRAGAO-XLA* MS Excel Add-In

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 sub-components.

Reporting Incidents

If you encounter any problems with the software, report an incident on the SAP Service Marketplace 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.

SAP HANA server software and tools can be used for several SAP HANA platform and options scenarios as well as the respective capabilities used in these scenarios. The availability of these is based on the available SAP HANA licenses and the SAP HANA landscape, including the type and version of the back-end systems the SAP HANA administration and development tools are connected to. There are several types of licenses available for SAP HANA. Depending on your SAP HANA installation license type, some of the features and tools described in the SAP HANA platform documentation may only be available in the SAP HANA options and capabilities, which may be released independently of an SAP HANA Platform Support Package Stack (SPS). Although various features included in SAP HANA options and capabilities are cited in the SAP HANA platform documentation, each SAP HANA edition governs the options and capabilities available. Based on this, customers do not necessarily have the right to use features included in SAP HANA options and capabilities. For customers to whom these license restrictions apply, the use of features included in SAP HANA options and capabilities in a production system requires purchasing the corresponding software license(s) from SAP. The documentation for the SAP HANA options is available in SAP Help Portal. If you have additional questions about what your particular license provides, or wish to discuss licensing features available in SAP HANA options, please contact your SAP account team representative.

7.3 SAP License Key

You must install a **permanent** SAP license for the SAP HANA database.

When you install your SAP system, a **temporary** license (90 days) is automatically installed. Before the temporary license expires, you must apply for a permanent license key from SAP.

We recommend that you apply for a permanent license key as soon as possible after installing your system.

> **Note**
>
> For more information, see section *Managing SAP HANA Licenses* in the *SAP HANA Administration Guide*.

7.4 Software Download

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

Installation Media and Components for SAP HANA

Installation Media for an SAP HANA SPS

1. Open the *SAP Support Portal Home*.
2. Choose *Download Software*.
3. Go to *INSTALLATIONS & UPGRADES*, if not already chosen.
4. Open *> By Alphabetical Index (A-Z)*.
5. Choose *H*.
6. Choose *SAP HANA PLATFORM EDITION*.
7. Go to *DOWNLOADS*, if not already opened.
8. Choose *SAP HANA PLATFORM EDITION 2.0*.
9. Open *DOWNLOADS*, if not already opened.
10. Choose *INSTALLATION*.
11. 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 Support Portal Home*.
2. Choose *Download Software*.
3. Choose *SUPPORT PACKAGES & PATCHES*, if not already chosen.
4. Open *> By Alphabetical Index (A-Z)*.
5. Choose *H*.
6. Choose *SAP HANA PLATFORM EDITION*.
7. Choose *DOWNLOADS*, if not already chosen.
8. Choose *SAP HANA PLATFORM EDITION 2.0*.
9. Choose *DOWNLOADS*, if not already chosen.
10. 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**, 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 of the product version. Support Package Stacks (SPS) can be downloaded and applied to appliances in accordance with agreements with the respective hardware partner.

Related Information

SAP Support Portal Home → Download Software

7.5 SAP HANA Hardware and Software Requirements

For SAP HANA, a number of hardware and software requirements apply.

> **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) on the SAP Service Marketplace, 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 the blog *Recent changes in the SAP HANA Technology certification program 2016* in the *Related Information* section.

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* for your SAP NetWeaver release under *Installation* > *System Copy*.

The following **SAP HANA features**, **SAP HANA capabilities**, and **SAP HANA options** are **supported on Intel-based hardware platforms only**:

- SAP HANA Accelerator for SAP ASE
- SAP HANA Hadoop Controller
- SAP HANA Remote Data Sync
- SAP HANA Smart Data Streaming
- SAP HANA Data Provisioning Agent
- Hive ODBC Driver

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 HANA Guidelines for SLES Operating System

SAP Note 2205917 - SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12

SAP Note 1984787 - SUSE LINUX Enterprise Server 12: Installation notes

Red Hat Enterprise Linux (RHEL)

SAP Note 2009879 - SAP HANA Guidelines for Red Hat Enterprise Linux (RHEL) Operating System

SAP Note 2292690 - SAP HANA DB: Recommended OS settings for RHEL 7.2

Supported Hardware Platforms

SAP Certified and Supported SAP HANA Hardware

On-Premise [page 73]

SAP Note 1943937 - Hardware Configuration Check Tool - Central Note

SAP Note 2055470 - HANA on POWER Planning and Installation Specifics - Central Note

SAP Note 2218464 - Supported products when running SAP HANA on IBM Power Systems

SAP Note 2188482 - SAP HANA on IBM Power Systems: Allowed Hardware

General Links

SAP HANA Tailored Data Center Integration 2017 - Overview

Recent changes in the SAP HANA Technology certification program 2016

SAP Note 52505 - Support after end of mainstream/extended maintenance

SAP Note 2235581 - SAP HANA: Supported Operating Systems

Product Availability Matrix

SAP HANA Network Requirements

SAP Note 1900823 - SAP HANA Storage Requirements and SAP HANA Storage Connector API

Important Disclaimer for Features in SAP HANA Platform, Options and Capabilities

SAP HANA server software and tools can be used for several SAP HANA platform and options scenarios as well as the respective capabilities used in these scenarios. The availability of these is based on the available SAP HANA licenses and the SAP HANA landscape, including the type and version of the back-end systems the SAP HANA administration and development tools are connected to. There are several types of licenses available for SAP HANA. Depending on your SAP HANA installation license type, some of the features and tools described in the SAP HANA platform documentation may only be available in the SAP HANA options and capabilities, which may be released independently of an SAP HANA Platform Support Package Stack (SPS). Although various features included in SAP HANA options and capabilities are cited in the SAP HANA platform documentation, each SAP HANA edition governs the options and capabilities available. Based on this, customers do not necessarily have the right to use features included in SAP HANA options and capabilities. For customers to whom these license restrictions apply, the use of features included in SAP HANA options and capabilities in a production system requires purchasing the corresponding software license(s) from SAP. The documentation for the SAP HANA options is available in SAP Help Portal. If you have additional questions about what your particular license provides, or wish to discuss licensing features available in SAP HANA options, please contact your SAP account team representative.

Important Disclaimers and Legal Information

Coding Samples

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

Gender-Neutral Language

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

Internet Hyperlinks

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

go.sap.com/registration/contact.html