



## SAP HANA Security Guide (Including SAP HANA Database Security)

- SAP HANA Appliance Software SPS 04

2012-04-24

## Copyright

© 2012 SAP AG. All rights reserved. SAP, R/3, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP BusinessObjects Explorer, StreamWork, SAP HANA and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and other countries. Business Objects and the Business Objects logo, BusinessObjects, Crystal Reports, Crystal Decisions, Web Intelligence, Xcelsius, and other Business Objects products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Business Objects Software Ltd. Business Objects is an SAP company. Sybase and Adaptive Server, iAnywhere, Sybase 365, SQL Anywhere, and other Sybase products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Sybase, Inc. Sybase is an SAP company. Crossgate, m@gic EDDY, B2B 360°, B2B 360° Services are registered trademarks of Crossgate AG in Germany and other countries. Crossgate is an SAP company. All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary. These materials are subject to change without notice. These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

2012-04-24

# Contents

<b>Chapter 1</b>	<b>Introduction.....</b>	<b>5</b>
1.1	Target Audience.....	5
1.2	Why Is Security Necessary?.....	5
1.3	About this Document.....	5
<b>Chapter 2</b>	<b>Before You Start.....</b>	<b>9</b>
2.1	SAP HANA Guides.....	9
2.2	Important SAP Notes.....	10
2.3	Additional Information.....	10
<b>Chapter 3</b>	<b>Technical System Landscape.....</b>	<b>13</b>
<b>Chapter 4</b>	<b>SAP HANA Database.....</b>	<b>15</b>
4.1	User Administration and Authentication.....	15
4.1.1	User Management.....	15
4.1.2	Integration into Single Sign-On Environments.....	25
4.1.3	Authentication Using SAML Bearer Token.....	26
4.1.4	Authorizations.....	28
4.1.5	Network and Communication Security.....	40
4.1.6	Communication Channel Security.....	41
4.1.7	Network Security.....	42
4.1.8	Secure Communication.....	45
4.1.9	Data Storage Security.....	51
4.1.10	Security Logging and Tracing.....	55
<b>Chapter 5</b>	<b>SAP HANA - Additional Components.....</b>	<b>61</b>
5.1	SAP HANA Information Composer.....	61
5.2	Lifecycle Management Tools.....	62
5.3	Unified Installer.....	62
5.4	SAP HANA UI Toolkit for INA.....	63

<b>Chapter 6</b>	<b>SAP HANA - Replication Technologies.....</b>	<b>65</b>
6.1	SAP HANA Security Extraction-Transformation-Load-(ETL-)Based Data Replication.....	66
6.1.1	Data Flow.....	67
6.1.2	Datastore Setup.....	68
6.2	SAP HANA Security ETL-based Data Acquisition by SAP HANA Direct Extractor Connection.....	68
<b>Chapter 7</b>	<b>Appendix.....</b>	<b>69</b>
7.1	Password Policy Parameters.....	69
7.2	How-To for Configuration of SAML Support.....	75

# Introduction

## **Caution:**

This guide does not replace the administration or operation guides that are available for productive operations.

## **1.1 Target Audience**

- Technology consultants
- System administrators

This document is not included as part of the installation guides, configuration guides, technical operation manuals, or upgrade guides. Such guides are only relevant for a certain phase of the software lifecycle, whereas security guides provide information that is relevant for all lifecycle phases.

## **1.2 Why Is Security Necessary?**

With the increasing use of distributed systems and the Internet for managing business data, the demands on security are also on the rise. When using a distributed system, you need to be sure that your data and processes support your business needs without allowing unauthorized access to critical information. User errors, negligence, or attempted manipulation on your system should not result in the loss of information or processing time. These demands on security also apply to the SAP HANA appliance software. This Security Guide is provided to assist you in securing SAP HANA appliance software.

## **1.3 About this Document**

The SAP HANA Security Guide provides an overview of the security-relevant information that applies to the SAP HANA appliance software including the SAP HANA database.

The SAP HANA Security Guide comprises the following main sections:

- Before You Start

This section contains references to the most important SAP Notes that apply to the security of the SAP HANA appliance software and further helpful resources.

- Technical System Landscape

This section provides an overview of the technical components and communication paths that are used by the SAP HANA appliance software.

- SAP HANA Database

This section provides an overview of the security aspects of the SAP HANA database. The SAP HANA database is used in different scenarios and solutions. Those solutions provide additional security guides, which cover aspects specific to the particular scenario.

- User Administration and Authentication

This section provides an overview of the following user administration and authentication aspects:

- Recommended tools for user management
    - User types that are required by the SAP HANA database
    - Standard users and roles that are delivered with the SAP HANA database
    - Overview of the password policy
    - Overview of the authentication mechanisms including the integration into Single Sign-On environments

- Authorizations

This section provides an overview of the SAP HANA database authorization concept.

- Network and Communication Security

This section provides an overview of the applicable communication paths used by the SAP HANA database and the security mechanisms. To restrict access at the network level, it also includes our recommendations for the network topology.

- Data Storage Security

This section provides an overview of any applicable critical data that is used by the SAP HANA database and the security mechanisms.

- Security Logging and Tracing

This section provides an overview of the trace and log files that contain security-relevant information, for example, so that you can reproduce activities if a security breach occurs.

- SAP HANA - Additional Components

- SAP HANA Information Composer

This section provides security-relevant information about the SAP HANA information composer, which is a Web application that allows you to upload data to and manipulate data on the SAP HANA database.

- Lifecycle Management Tools

This section provides security-relevant information about Lifecycle Management Tools such as the Software Update Manager (SUM).

- Unified Installer

This section provides security-relevant information about the unified installer.

- SAP HANA UI Toolkit for INA

This section provides security-relevant information about the SAP HANA UI toolkit for INA.

- SAP HANA - Replication Technologies

- Trigger-Based Data Replication

Trigger-Based Data Replication using SAP LT (Landscape Transformation) Replication Server

- Extraction-Transformation-Load -Based Replication

Extraction-Transformation-Load-(ETL-)Based Data Replication using SAP BusinessObjects Data Services Extraction-Transformation-Load (ETL) based data replication

- Log-Based Replication

Transaction Log-Based Data Replication using Sybase Replication

- ETL-based Data Acquisition by SAP HANA Direct Extractor Connection

- Appendix

This section provides information about further details, such as security-relevant parameter settings.



## Before You Start

### 2.1 SAP HANA Guides

For more information about SAP HANA landscape, deployment, installation, administration, and security, see the resources listed in the table below.

Topic	Guide/Tool	Quick Link
SAP HANA Landscape, Deployment & Installation	SAP HANA Appliance Software Knowledge Center on SAP Service Marketplace	<a href="https://service.sap.com/hana">https://service.sap.com/hana</a> -> <a href="#">SAP HANA Master Guide</a> -> <a href="#">SAP HANA Overall Installation Guide</a> -> <a href="#">SAP HANA Automated Update Guide</a> -> <a href="#">SAP HANA Database - SQL Reference Guide</a> -> <a href="#">SAP HANA Database - SQL Script Guide</a>
SAP HANA Administration & Security	SAP HANA Appliance Software Knowledge Center on SAP Help Portal	<a href="http://help.sap.com/hana">http://help.sap.com/hana</a> -> <a href="http://help.sap.com/hana_appliance">http://help.sap.com/hana_appliance</a> -> <a href="#">SAP HANA Technical Operations Manual</a> -> <a href="#">SAP HANA Security Guide</a>

For a complete list of the available SAP Security Guides, see <https://service.sap.com/securityguide> on the SAP Service Marketplace.

## 2.2 Important SAP Notes

The most important SAP Notes that apply to SAP HANA appliance software and SAP HANA database security are shown in the table below.

SAP Note	Title
<a href="#">1598623</a>	SAP HANA appliance: Security
<a href="#">1514967</a>	SAP HANA appliance: Central Note

In addition, you can find a list of security-relevant SAP Hot News and SAP Notes on the SAP Service Marketplace at <https://service.sap.com/securitynotes>.

## 2.3 Additional Information

For more information about specific topics, see the Quick Links in the table below.

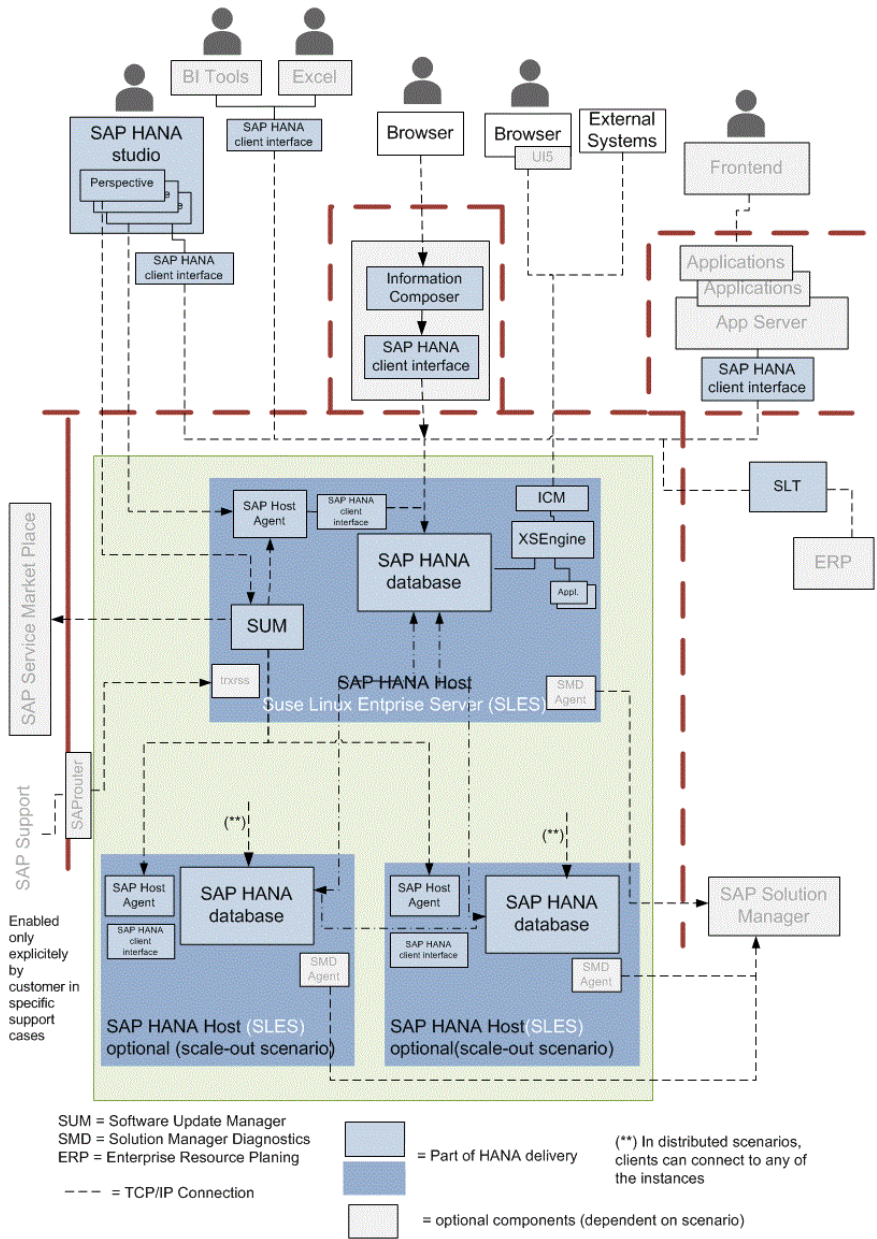
Content	Quick Link on the SAP Service Marketplace or SDN
Security	<a href="https://sdn.sap.com/irj/sdn/security">https://sdn.sap.com/irj/sdn/security</a>
Security Guides	<a href="https://service.sap.com/securityguide">https://service.sap.com/securityguide</a>
Related SAP Notes	<a href="https://service.sap.com/notes">https://service.sap.com/notes</a> <a href="https://service.sap.com/securitynotes">https://service.sap.com/securitynotes</a>
Released platforms	<a href="https://service.sap.com/pam">https://service.sap.com/pam</a>
Network security	<a href="https://service.sap.com/securityguide">https://service.sap.com/securityguide</a>
SAP Solution Manager	<a href="https://service.sap.com/solutionmanager">https://service.sap.com/solutionmanager</a>

<b>Content</b>	<b>Quick Link on the SAP Service Marketplace or SDN</b>
SAP NetWeaver	<a href="http://sdn.sap.com/irj/sdn/netweaver">http://sdn.sap.com/irj/sdn/netweaver</a>
In-Memory Computing	<a href="http://www.sdn.sap.com/irj/sdn/in-memory">http://www.sdn.sap.com/irj/sdn/in-memory</a>



## Technical System Landscape

The figure below shows an overview of the technical system landscape for the SAP HANA appliance software and its related components. The related components include the SAP HANA studio and other applications, such as the SAP HANA information composer. Note that the figure below shows a sample configuration with three SAP HANA appliances and three SAP HANA databases. The figure also shows some optional components that must be purchased separately.



# SAP HANA Database

## 4.1 User Administration and Authentication

The following sections contain information about user administration and authentication that specifically applies to the SAP HANA database:

- [User Management](#)

This section lists the tools to use for user management, the types of users required, and the standard users that are delivered with the SAP HANA database.

- [Integration into Single Sign-On Environments](#)

This section describes how the SAP HANA database supports single sign-on mechanisms.

- [Authentication Using SAML Bearer Token](#)

This section describes the SAP HANA database authentication using the Security Assertion Markup Language bearer token.

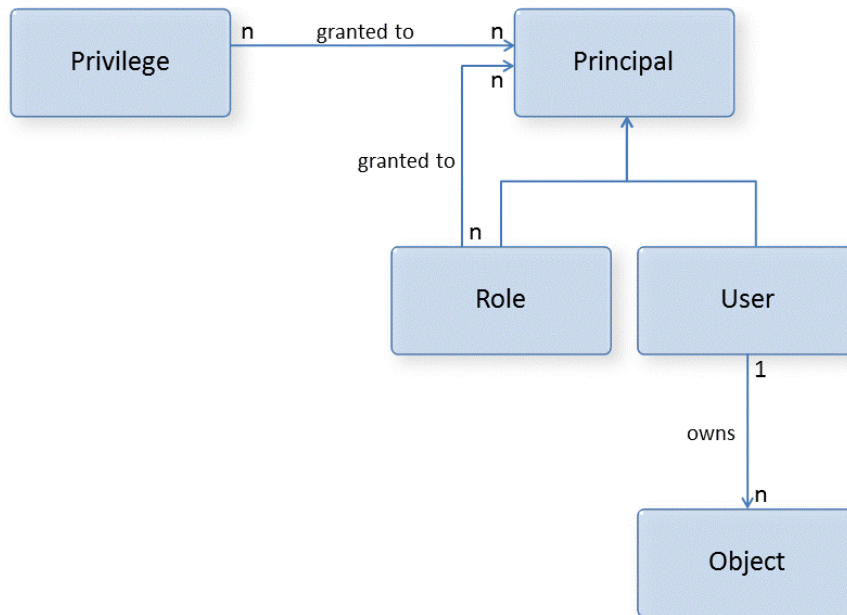
### 4.1.1 User Management

All users who want to access the SAP HANA database must be explicitly created in the database beforehand. A user who is connected using the external authentication provider but is not known to the database cannot access the database.

The relevant entities mentioned below relate to each other in the following way:

- A principal is either a role or a user.
- A known user can log on to the database. A user can be the owner of database objects.
- A role is a collection of privileges and can be granted to either a user or another role (nesting).
- A privilege is used to impose restrictions on operations carried out on certain objects. For more information, see [Privileges](#).

The entity relationships described above are depicted in the following figure:



Privileges can be assigned to users directly or indirectly using roles. Privileges are required to model access control. Roles can be used to structure the access control scheme and model reusable business roles.

**Tip:**

We recommend to manage authorization for users by using roles. Roles can be nested so that role hierarchies can be implemented. This makes them very flexible, allowing very fine- and coarse-grained authorization management for individual users.

All the privileges granted directly or indirectly to a user are combined. This means whenever a user tries to access an object, the system performs an authorization check using the user, the user's roles, and directly allocated privileges.

It is not possible to explicitly deny privileges. This means that the system does not need to check all the user's roles. As soon as all requested privileges have been found, the system aborts the check and grants access.

Several predefined roles exist in the database. Some of them are templates that need to be customized; others can be used as they are. For more information, see [Standard Roles](#).

User management is configured using SAP HANA studio. For an overview of how these mechanisms apply to the SAP HANA database, see the sections below. In addition, a list of the standard users required for operating the SAP HANA database is provided.

### 4.1.1.1 User Administration Tools

The table below shows the tools used with the SAP HANA database for user management and administration. To change passwords, and create or delete users, the USER ADMIN system privilege is required.

**Note:**

For more information about changing passwords with the On-Site Configuration tool, see the *SAP HANA Installation Guide with Unified Installer* at <http://help.sap.com/hana>.

Table 4-1: User Management Tools

Tool	Detailed Description
User and role maintenance with SAP HANA studio	For more information about SAP HANA administration and security, see the SAP HANA Knowledge Center on SAP Help Portal at <a href="http://help.sap.com/hana">http://help.sap.com/hana</a> -> <b>System Administration</b> .
User and role maintenance using the command line interface (hdbsql or other SQL tool)	By using SQL requests, for example, all the user management functions can also be executed from the command line. This is useful when using scripts for automated processing.  For more information about the command line tool hdbsql, see the <i>SAP HANA Database Administration Guide</i> at <a href="http://help.sap.com/hana">http://help.sap.com/hana</a> .

Tool	Detailed Description
SAP NetWeaver Identity Management	<p>The SAP NetWeaver Identity Management 7.2 SP 3 contains a connector to the SAP HANA database (IDM connector). With The SAP NetWeaver Identity Management you can perform the following actions in the SAP HANA database:</p> <ul style="list-style-type: none"> <li>• Creating and deleting user accounts</li> <li>• Assigning roles</li> <li>• Setting passwords for users</li> </ul> <p>For more information about the SAP NetWeaver Identity Management and the IDM connector, see the SAP Community Network at <a href="http://www.sdn.sap.com">http://www.sdn.sap.com</a> -&gt; <b>SAP NetWeaver Releases</b>.</p>

#### 4.1.1.1.1 User and Role Maintenance with SAP HANA Studio

The user and role concept of the SAP HANA database allows for a fine granularity of access control based on the user's tasks, for example:

- Business end users reading reports using client tools, for example, Microsoft Excel.
- Modelers creating models and reports using SAP HANA studio.
- Database administrators operating and maintaining the database and users using SAP HANA studio.

SAP HANA studio is used for the following security-relevant tasks:

- Creating users
- Creating roles and the role hierarchy

**Note:**

Only the following characters are allowed in user and role names:

A-Z, 0-9, \_

- Assigning privileges to roles and users
- Assigning roles to users
- Modeling and activating Analytic Privileges

The process flow for user management is as follows:

1. Define and create privileges.

This only applies to Analytic Privileges. All other privileges are built-in.

2. Define and create roles.
3. Assign privileges to roles.
4. Create users.
5. Assign roles to users.

We recommend using roles for assigning privileges to individual users.

For more information about SAP HANA administration and security as well as about how to create and modify users and roles, see the SAP HANA Appliance Software Knowledge Center on SAP Help Portal at <http://help.sap.com/hana> -> **System Administration**.

#### 4.1.1.1.2 Authentication

The SAP HANA database provides the following options for authentication:

- Direct logon to the SAP HANA database with user name and password

The SAP HANA database authenticates the user.

**Note:**

For some administrative operations, such as database recovery, the credentials of the SAP operating system user (<sapsid>adm) are also required.

- Authentication using Kerberos (third-party authentication provider)

For more information, see [Integration into Single Sign-On Environments](#).

- Authentication using Security Assertion Markup Language (SAML) bearer token

For more information, see [Authentication Using SAML Bearer Token](#).

#### 4.1.1.1.3 User Types

It is often necessary to specify different security policies for different types of users.

The user types that are required for the SAP HANA database include:

- Users

Named users represent real persons and are used for daily working with the SAP HANA database. These users are created by the user administrator.

- SYSTEM user

This user is the built-in overall system administrator. For more information, see the “Standard Users” table below.

- Technical users

- SYS, \_SYS\_STATISTICS, and \_SYS\_REPO

These are internal users within the SAP HANA database. The `_SYS_STATISTICS` user is used by the statistics server, and the `_SYS_REPO` by the repository. These users cannot log on from outside.

**Caution:**

These users must not be used to log on to the database and to perform day-to-day activities.

- Application-specific technical users

For example, an application server may log on to the SAP HANA database using a dedicated technical user.

Only a conceptual separation, but no technical differences, exists between these user types. The SAP HANA database authentication is the same for both users and technical users.

#### 4.1.1.1.4 Standard Users

The table below shows the standard users that are necessary for installing, upgrading, and operating the SAP HANA database.

Table 4-2: Standard Users

User ID	Type	Password	Description
SYSTEM	Overall system administrator	You specify the initial password during installation.	<p>The SYSTEM user is the initial bootstrap user that is created during the installation of the SAP HANA database.</p> <p><b>Tip:</b> Do not use the SYSTEM user for day-to-day activities.</p> <p>Use this user to create dedicated administrator users and to assign privileges to the administrator users.</p>

User ID	Type	Password	Description
<sid>adm <sid> = system ID	SAP system user (operating system user)	You specify the initial password during installation.	The SAP system administrator has unlimited access to all local resources related to SAP systems.  This user is not a database user but a user on the operating system level.
ROOT	User for installation and upgrade	You specify the initial password during installation.	The ROOT user is used for installation and upgrade, only. <b>Tip:</b> Do not use the ROOT user for day-to-day activities.

#### 4.1.1.1.5 Standard Roles

The table below shows the standard roles that are delivered with the SAP HANA database.

Table 4-3: Standard Roles

Role	Description
MODELING	<p>Contains all privileges required for using the information modeler in the SAP HANA studio.</p> <p>Contains the database authorization for a modeler to create all kinds of views and Analytic Privileges.</p> <p>Allows access to all data in activated views without any filter (<code>_SYS_BI_CP_ALL</code> Analytic Privilege). However, this is restricted by missing SQL Privileges on those activated objects.</p> <p><b>Note:</b></p> <p>Use caution when using the <code>_SYS_BI_CP_ALL</code> Analytic Privilege.</p> <p>Use this predefined role as a template.</p>
MONITORING	<p>Contains privileges for full read-only access to all meta data, the current system status in system and monitoring views, and the data of the statistics server.</p>
PUBLIC	<p>Contains privileges for filtered read-only access to the system views. Only objects for which the users have access rights are visible. By default, this role is assigned to each user.</p>
CONTENT_ADMIN	<p>Contains the same privileges as the MODELING role, but with the extension that users allocated this role are allowed to grant these privileges to other users.</p> <p>In addition, it contains Repository Privileges for working with imported objects.</p> <p>Use this role as a template for what content administrators might need as privileges.</p>

#### 4.1.1.1.6 Password Policy

Passwords for database users are subject to certain security rules, which are configured using parameters. You can find these parameters in the `indexserver.ini` file in the password policy section.

**Note:**

To view the contents of the INI file, use the `M_INIFILE_CONTENTS` view.

The password policy parameters can be found in the `M_PASSWORD_POLICY` view. For more information about the system tables and monitoring views, see *System Tables and Monitoring Views* at <http://help.sap.com/hana>.

To change the password policy rules, you need the `INIFILE ADMIN` system privilege.

To change the parameter values in the "Password policy" section, you have the following options:

- Using SAP HANA studio, follow these steps:
  1. Open the "Administration editor" and go to the **Configuration** tab.
  2. Expand the **indexserver.ini** section.
  3. In the "Password policy" section, change the required parameters.
- Using an SQL statement

```
Alter system alter configuration ('indexserver.ini' , 'SYSTEM')
set ('password policy' , '<parameter_name>') = '<new_value>' with reconfigure
```

If a parameter is set to a value outside the value range, either the minimum value or the maximum value of the value range, whichever is appropriate, is used instead.

**Note:**

No warning is given if parameters are set to a value outside the permitted range.

The current parameter values (or the corrected internal values) are available for all users in the `M_PASSWORD_POLICY` view.

To display information about the current state of users, use the `USERS` system view. You can do this to, for example, find out when the password of a user expires or why the user account is locked.

For more information about the parameter values, see the [Password Policy Parameters](#) in the "Appendix" of this document.

#### 4.1.1.1.7 Deactivation of Users

The administrator can deactivate a user account in SAP HANA studio or with the following SQL command:

```
ALTER USER <user_name> DEACTIVATE [USER NOW]
```

After the user account is deactivated, the user cannot log on to the SAP HANA database until the administrator reactivates that user again.

For more information about user account deactivation in SAP HANA studio, see the *SAP HANA Database Administration Guide* at <http://help.sap.com/hana>.

#### 4.1.1.1.8 Reactivation of Users

The administrator can reactivate a user account. A user account can be locked because of the following reasons:

- The user was deactivated explicitly.
- The user's password has expired.
- The user has made too many invalid logon attempts.

If the user was deactivated explicitly, an administrator has to execute the following SQL command:

```
ALTER USER <user_name> ACTIVATE [USER NOW]
```

If the deactivated user uses a password to connect to the SAP HANA database, the administrator can reactivate the user by changing their password to a new value.

If the user's password has expired, the user has to change the password to a new value. If the user has made too many invalid logon attempts, the administrator can unlock the user account in SAP HANA studio or with an SQL command. For more information, see the `password_lock_time` parameter in the [Password Policy Parameters](#) in the "Appendix" of this document.

For more information about user account reactivation in SAP HANA studio, see the *SAP HANA Database Administration Guide* at <http://help.sap.com/hana>.

#### 4.1.1.1.9 Emergency User

If the SYSTEM user's password is lost, you can use the SAP system user to reset the password. To recover a SAP HANA instance where SYSTEM user's password is lost, you need to have <sid>adm access to the instance where SAP HANA's master index server is running.

1. Open a command line interface, and log on to the server on which the instance of the SAP HANA master index server is running.
2. Shut down the instance.
3. Start the name server:

```
./usr/sap/<SID>/HDB<instance>/hdbenv.sh  
/usr/sap/<SID>/HDB<instance>/exe/hdbnameserver
```

4. Start an index server in console mode:

```
./usr/sap/<SID>/HDB<instance>/hdbenv.sh  
/usr/sap/<SID>/HDB<instance>/exe/hdbindexserver -console
```

You see the output of a starting index server. When the service has started, you have a console to the SAP HANA instance where you are logged on as a SYSTEM user.

5. You can reset the SYSTEM user's password and store the new password in a secure location with the following SQL command:

```
ALTER USER SYSTEM password <new password>
```

**Note:**

As you are logged on as a SYSTEM user in this console, you do not have to change this password after the next logon, regardless of what your password policy setting is.

## 4.1.2 Integration into Single Sign-On Environments

SAP HANA supports the Kerberos protocol for single sign-on. It has been tested with Windows Active Directory Domain Kerberos implementation and MIT Kerberos network authentication protocol. The ODBC database client and the JDBC database client support Kerberos.

To implement this, you need to install the MIT Kerberos client software on the host of the SAP HANA database. Once installation is complete, configure Kerberos as follows:

1. Create a user that serves as the service user for the SAP HANA database.
2. For this user, create one service principal name (SPN) for each host of the system.

The SPN needs to have the following syntax:

```
hdb/<domain_name>@Kerberos realm name
```

where *<domain\_name>* is the fully qualified domain name of the host.

3. Export each of the SPNs created into a separate file.
4. Import the SPNs in each of the files to the respective host.

For more information, see the Kerberos product documentation.

The users stored in the Microsoft Active Directory or the MIT Kerberos Key Distribution Center can be mapped to database users in the SAP HANA database. For this purpose, specify the user principal name (UPN) as the external ID when creating the database user.

To map the Kerberos-enabled provisioning mechanism to the SAP HANA database, do the following:

1. Launch SAP HANA studio.
2. In the navigator, select **Catalog > Authorization**.
3. From the context menu, select **New > User**.
4. In the "User Name" field, enter the user name.
5. In the "Authentication" section, choose **External** and enter the external ID specified in the Microsoft Active Directory or the MIT Kerberos Key Distribution Center for the user.

Alternatively, you can create the user with the following SQL command:

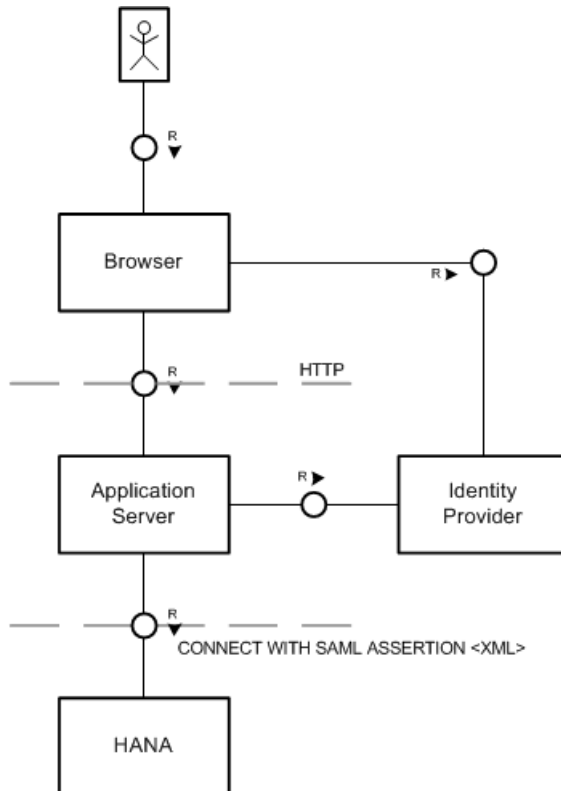
```
CREATE USER <username> IDENTIFIED EXTERNALLY AS '<external ID>'
```

For more information about SAP HANA administration, see the SAP HANA Appliance Software Knowledge Center on SAP Help Portal at <http://help.sap.com/hana> -> **System Administration**.

### 4.1.3 Authentication Using SAML Bearer Token

SAP HANA supports the Security Assertion Markup Language (SAML) as an additional authentication mechanism besides username/password and Kerberos. SAML is only used for authentication purposes and not for authorization.

The main purpose of SAML for SAP HANA is to support scenarios where clients are not directly connected to SAP HANA but to a middle tier application server (like the XS Engine) that runs an HTTP server. Whenever the application server needs to connect to the database on behalf of the user, it requests a SAML assertion from the client. The assertion is issued by an identity provider after the client was successfully authenticated. The assertion is then forwarded to the SAP HANA database, which will grant access based on the previously established trust to the identity provider. This scenario is depicted in the following figure:



For more information about the configuration, see [How-To for Configuration of SAML Support](#) in the "Appendix" of this document.

### 4.1.3.1 Supported SAML Features

SAP HANA supports plain SAML 2.0 assertions as well as unsolicited SAML responses that include an unencrypted SAML assertion. SAML assertions and responses have to be signed using XML signatures.

The following features of XML signatures are supported:

- SHA1 and MD5 for hash algorithms
- RSA-SHA1 as signature algorithm
- <X509Certificate> elements

**Note:**

For SPS 04, the XML signature must contain the X.509 certificate of the identity provider within the <X509Certificate> element.

The following SAML assertion features are supported:

- Assertion Subject with NameID
- Qualified NameID with SPProvidedID and SPNameQualifier
- Validity conditions (NotBefore, NotOnOrAfter)
- Audience restrictions

**Note:**

For SPS 04:

- SAML is not supported in distributed environments.
- Automatic client reconnect is not supported for SAML authenticated connections.

### 4.1.3.2 Assertion Properties Checked

The following properties of a SAML assertion are currently evaluated:

- saml:Assertion/@Version: Must be 2.0.
- saml:Subject/saml:NameID: Must exist.
- saml:Subject/saml:NameID/@Format: Must be "urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified".

- saml:Subject/saml:NameID/@SPProvidedID: Must either match an explicit mapping in the SAP HANA database or a wildcard mapping must have been set for the user.
- saml:Subject/saml:SubjectConfirmation: Must be `{{"urn:oasis:names:tc:SAML:2.0:cm:bearer"}}` if it exists.
- saml:Conditions: The following conditions are currently evaluated:
  - @NotBefore
  - @NotOnOrAfter: Must be set.
  - AudienceRestriction

### 4.1.3.3 User Mapping

Identity providers must be configured per user as a logon option. The following types of user mappings are supported:

- SAP HANA-based user mappings: The mapping to a SAP HANA database user is explicitly configured within SAP HANA per identity provider. The corresponding assertion subject looks like this:

```
<NameID Format="urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified">zgc2VLavgYy4hsohfYPM21</NameID>
```

- Identity provider-based user mappings: The identity provider maps its users to SAP HANA database users and provides this information via the SPProvidedID attribute. The corresponding assertion subject looks like this:

```
<NameID Format="urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified" SPProvidedID="BILLG">zgc2VLavgYy4hsohfYPM21</NameID>
```

If a SAP HANA-based user mapping exists for a given identity provider and a conflicting SPProvidedID is sent from the identity provider, an error is returned.

### 4.1.4 Authorizations

Depending on the customer landscape, preferred reporting tools, and the authorization requirements you have to decide whether to apply the authorization concept of SAP HANA, the authorization concept of the end-user tool (for example, SAP BusinessObjects Enterprise), or a hybrid of both.

### 4.1.4.1 Authorization Concept of the SAP HANA Database

When accessing the SAP HANA database using a client interface (such as ODBC, JDBC, MDX), any access to data must be backed by corresponding privileges. Different schemes are implemented. On a higher level, this concept provides authorization for the data contained in the database when it is accessed using client interfaces. In the SAP HANA database system, the regular SQL authorization concept is implemented.

For each SQL statement type (for example, SELECT, UPDATE, and CALL), a corresponding privilege exists that the executing user needs to have. Additionally, objects in the database (such as tables, views, or stored procedures) have an owner who can access the objects and grant privileges for them. No user, besides the owner of an object and users that the owner has provided with a privilege, can access this particular object. This authorization functions on the object level, whereby the smallest entities that can be privileged are, for example, a table or a view.

**Caution:**

The database “owner” concept stipulates that when a user (owner) is deleted, all objects created by or privileges granted to others by this user are deleted.

In addition, Analytic Privileges are used to provide row-level authorization on certain kinds of database objects, such as Analytic Views. Analytic Privileges can only be used for read operations and not for write operations. The Analytic Privileges provide the possibility of granting privileges to only view parts of the data in a view. An Analytic Privilege enables the grantee to see certain view rows that are identified by one or more column values. For example, an Analytic Privilege could enable the grantee to see only those entries in the SALES view for the years with the values 2006 to 2008.

With these two higher-level concepts, it is possible for an application using the database to control which user has access to what data. Package Privileges provide a further means of restricting access to different design time objects that are bundled in packages.

### 4.1.4.2 Privileges

The table below describes the types of privileges used by the SAP HANA database.

Table 4-4: Types of Privileges

Type of Privilege	Description
System Privileges	<p>These system-wide SQL Privileges control general system activities mainly for administrative purposes, such as creating schemas, creating and changing users and roles.</p> <p>System Privileges are assigned to users and roles. Following the principle of least privilege, users should only be given the smallest set of privileges required for their role.</p> <p>For SQL Privileges in the SAP HANA database, the SQL standard behavior is applied.</p>
Object Privileges	<p>These SQL Privileges are used to restrict access to and modification of database objects, such as tables. Depending on the object type (for example, table, view), actions (for example, CREATE ANY, ALTER, DROP) can be authorized per object.</p> <p>Object Privileges are assigned to users and roles. Following the principle of least privilege, users should only be given the smallest set of privileges required for their role.</p> <p>For SQL Privileges in the SAP HANA database, the SQL standard behavior is applied.</p>
Analytic Privileges	<p>Analytic Privileges are used to restrict the access for read operations to certain data in Analytic, Attribute, and Calculation Views by filtering the attribute values.</p> <p>Only applied at the processing time of the user query.</p> <p>Analytic Privileges need to be defined and activated before they can be granted to users and roles.</p> <p>For more information, see <a href="#">Analytic Privileges</a>.</p>

Type of Privilege	Description
Package Privileges	<p>Package Privileges are used to restrict the access to and the use of packages in the repository of the SAP HANA database.</p> <p>Packages contain design time versions of various objects, such as Analytic, Attribute, and Calculation Views, as well as Analytic Privileges, and functions. To be able to work with packages, the respective Package Privileges must be granted.</p>

For more information about all the supported System and Object Privileges, see the *SAP HANA Database - SQL Reference Guide* at <http://help.sap.com/hana>.

All privileges can be granted or revoked using the GRANT or REVOKE command in the SQL interface. Object Privileges need to be granted on individual objects. For more information, see the *SAP HANA Database - SQL Reference Guide* at <http://help.sap.com/hana>.

### 4.1.4.3 Analytic Privileges

Analytic Privileges are used in the SAP HANA database to provide fine-grained control of what data particular users can see for analytic use. They provide the ability for row-level authorization, based on filtering the values in one or more columns. All Attribute Views, Analytic Views, and Calculation Views, which have been designed in the information modeler and have been activated from the information modeler of the SAP HANA studio, are automatically supported by the Analytic Privilege mechanism.

If you are already familiar with the authorization model of SAP NetWeaver Business Warehouse (SAP NetWeaver BW), you will see many similarities between the two models. The main idea behind Analytic Privileges is to allow access of different users to different portions of data of the same view. For example, different regional sales managers, who are only allowed to see sales data for their regions, could access the same Analytic View. They would get the Analytic Privilege to see only data for their region, and their queries on the same view would return the corresponding data. This is a major difference to the SAP NetWeaver BW model. While the concept itself is very similar, SAP NetWeaver BW would forward an error message if you executed a query that would return values you are not authorized to see. With the SAP HANA database, the query would be executed and, corresponding to your authorization, only values you are entitled to see returned.

#### 4.1.4.3.1 Structure of Analytic Privileges

An Analytic Privilege consists of several restrictions. Three of these restrictions are always present and have the following special meanings:

- The Cube restriction determines for which column views (Attribute, Analytic, or Calculation Views) the privilege is used. This may involve a single view, a list of views or by means of a wildcard, all applicable views.
- The Activity restriction determines the affected activity, for example, READ. This means that the activity READ is restricted and not available for use.
- The Validity restriction determines at what times the privilege is valid.

In addition to these three restrictions, many additional Dimension restrictions are used. These are applied to the actual attributes of a view. Each Dimension restriction is relevant for one dimension attribute, which can contain multiple value filters. Each value filter is a tuple of an operator and its operands, which is used to represent the logical filter condition. For example, a value filter (EQUAL 2006) can be defined for a dimension attribute YEAR in a Dimension restriction to filter accessible data using the condition YEAR=2006 for potential users.

**Note:**

No measures or key figures in Analytic Views can be employed in Dimension restrictions.

#### 4.1.4.3.2 Use in Information Modeler

##### **Analytic Privilege-Capable Views**

The Analytic Privilege mechanism is automatically enforced for all three kinds of column views that can be defined using the information modeler, namely Attribute, Analytic, and Calculation Views:

- Attribute Views

These views are built on joins of existing column tables and views. Attribute Views cannot be nested in other Attribute Views.

- Analytic Views

These views are multidimensional cubes with a fact table joined with multiple dimension tables. The information modeler allows Analytic Views to be associated with Attribute Views to reuse the specified join paths. However, it is not possible to use existing Attribute or Analytic Views as base views (join candidates) and use these as the basis for defining new Analytic Views.

- Calculation Views

These views are defined directly using SQL Script or data flow operators in the graphical editor for Calculation Views. Any existing views, including Attribute, Analytic, and Calculation Views, can be used, for example, in a SELECT statement, to build new Calculation Views. This introduces interdependencies between the views.

##### **Definition of Analytic Privileges**

The information modeler provides a convenient way of defining and enforcing Analytic Privileges on column views. However, it also enforces several constraints on the behavior of Analytic Privileges on the views.

To define an Analytic Privilege, the user should first select all relevant views (reference models) to which this Analytic Privilege is applicable. The selectable views can be previously defined Attribute,

Analytic, or Calculation Views. Instead of choosing single relevant views, you can also select the **Applicable to all Information Models** checkbox to indicate that the Analytic Privilege is applicable to all possible views.

In addition, the different restrictions in the Analytic Privileges are set by the information modeler as follows:

1. For the Cube restriction:

In general, an Analytic Privilege is only valid for those views specified in the cube restriction.

For Analytic Views, the modeler automatically adds an association to further automatically generated views, such as views with the "/olap" suffix.

Access to a view is not granted if a valid Dimension restriction is found for the particular view (with the corresponding dimension attribute).

2. For the Activity restriction:

All Analytic Privileges are set to check for READ activity only. This is due to the fact that the target Attribute, Analytic, and Calculation Views are read-only views.

3. For the Validity restriction:

After their activation, all Analytic Privileges become immediately valid and have unlimited validity.

4. For the Dimension restriction:

The following operators can be employed for value filters in Dimension restrictions due to the query processing capabilities of the calculation database:

- = Equal
- <> Between
- > GreaterThan
- < LessThan
- >= GreaterEqual
- <= LessEqual
- ContainsPattern (with simple patterns, like AB\*, A\*B)

#### 4.1.4.3.3 Authorization Check

##### **General Behavior on Modeled Content**

In general, the user has access to an individual, independent view (Attribute, Analytic, or Calculation View) if the following prerequisites are met:

- The user was granted the SELECT privilege on the view or the containing schema.
- The user was granted an Analytic Privilege that is applicable to the view. An Analytic Privilege is applicable to a view if it contains the view in the Cube restriction and contains at least one filter on one attribute of this view.

**Note:**

No SELECT privilege on the underlying base tables or views of this view is required.

When access is possible, all relevant Analytic Privileges need to be evaluated so that view data can be filtered accordingly for user queries. The first step in the evaluation is to determine relevant Analytic Privileges for the current user and view. This is carried out in accordance with the following criteria:

- Analytic Privileges previously granted to this user are considered.
- Analytic Privileges with the Cube restriction covering this view are considered.
- Analytic Privileges with a currently valid Validity restriction are considered (only internal, cannot be set in the information modeler).
- Analytic Privileges with an Activity restriction covering the activity requested by the query are considered (only internal, because currently only the READ activity is supported).
- Analytic Privileges with Dimension restrictions covering some attributes of this view are considered.

If no relevant Analytic Privileges can be found, user queries are rejected with a Not authorized error. That is, despite the SELECT privilege on a view, access is not possible.

When relevant Analytic Privileges are found for the current user and the query directed to the particular view, the evaluation process ensures that, according to the value filters specified in the Dimension restrictions, the appropriate view data is presented to the user. In particular:

- Within one Dimension restriction, all value filters on the corresponding dimension attribute are combined with logical OR.
- Within one Analytic Privilege, all Dimension restrictions are combined with logical AND.
- Multiple Analytic Privileges are combined with logical OR.

For example, if there is only one Analytic Privilege found with two Dimension restrictions, YEAR=2008 and COUNTRY=US, the user is only allowed to see data fulfilling the condition YEAR=2008 AND COUNTRY=US. However, if these two conditions were put in two different Analytic Privileges found for this user and this view, the user is allowed to see more data, namely the OR combination of the filters of the individual Analytic Privileges: YEAR=2008 OR COUNTRY=US.

### **Behavior in Nested or Dependent Views**

The behavior of Analytic Privileges created and activated by the information modeler is as follows:

- Analytic Views associated with Attribute Views

Access to an Analytic View is granted as soon as the current user was granted an Analytic Privilege on the view itself, or an Analytic Privilege on an underlying Attribute View.

Result filtering is performed by a logical OR combination of all Analytic Privileges found on the Analytic View as well as on its underlying Attribute Views for the current user.

An Analytic Privilege on an Analytic View does not restrict data access on the view, but rather increases its scope, as the privilege is combined with the privileges defined on the associated Attribute Views of this Analytic View using logical OR.

- Calculation Views based on other views

Access to a Calculation View is only possible if the user was granted corresponding Analytic Privileges on this view, and also on all underlying views, if they are Attribute, Analytic, or Calculation Views.

Result filtering is performed for the individual views. In particular, the underlying views (Attribute, Analytic, Calculation Views) are authorized and filtered using the Analytic Privileges relevant for these views. Finally, the result of the top-level Calculation View is filtered using the Analytic Privileges defined on the Calculation View itself.

An Analytic Privilege on a Calculation View further restricts data access on the view, as the privilege is evaluated over and above (corresponding to logical AND combination) with the result filtered by the Analytic Privileges defined on the underlying views.

#### 4.1.4.3.4 Administration of Analytic Privileges

The CREATE STRUCTURED PRIVILEGE System Privilege is required to create Analytic Privileges in the database.

The STRUCTUREDPRIVILEGE ADMIN System Privilege allows users to drop existing Analytic Privileges.

In the information modeler, Analytic Privileges as well as the views are created (that is, activated) and dropped and recreated (that is, redeployed) by the SYS\_REPO technical user. By default, this user already has both privileges for the administration of Analytic Privileges at the database level.

A database user requires corresponding Repository Privileges, namely REPO.EDIT\_NATIVE\_OBJECTS and REPO.ACTIVATE\_NATIVE\_OBJECTS to activate and redeploy Analytic Privileges in the information modeler. For more information, see [Package Privileges](#).

### 4.1.4.4 Behavior of SQL and Analytic Privileges

SQL and Analytic Privileges show the following behavior:

- SQL Privileges can be defined for views and tables (including Attribute, Analytic, and Calculation Views). Currently, SELECT and DROP privileges can be granted for activated views.
- If a view contains objects of another view (for example, Attribute View in Analytic View), the view owner needs privileges for all objects within this view. Users accessing the containing view only need privileges for the actual view they want to access.
- The view owner needs at least the SELECT privilege on all underlying objects.
- Without SQL Privilege on the individual views, no access to any Attribute, Analytic, and Calculation Views is possible.
- Analytic Privileges can be defined for Attribute, Analytic, and Calculation Views.
- Analytic Privileges are applicable to Attribute, Analytic, and Calculation Views, and not to database tables. Access to database tables is governed entirely by SQL Privileges, whereas access to those views is controlled by SELECT privilege and a suitable Analytic Privilege.
- Different Analytic Privileges in Attribute and Analytic Views are combined by logical OR. For example, if a user has two Analytic Privileges, one restricting the product group to “123” and another restricting the customer to “456”, the user can access all data on products of group “123” (regardless of the

customer) as well as all data on customer “456” (regardless of the product). If an AND combination is desired, it must be defined within a single Analytic Privilege.

- Regarding Calculation Views, the behavior differs. The privileges of Attribute and Analytic Views are combined as described and provide a subset of data. The Analytic Privileges of the Calculation View are applied on this subset.
- If an Analytic View uses Attribute Views, it is sufficient to define at least one Analytic Privilege on the Analytic or any of the Attribute Views. If at least one Analytic Privilege is defined, then views without Analytic Privilege do not introduce further restrictions.

#### 4.1.4.5 Repository Privileges

The repository of the SAP HANA database consists of packages that contain design time versions of various objects, such as Attribute, Analytic, and Calculation Views, procedures, and Analytic Privileges. All repository methods that provide read or write access to content are secured with authorization checks. To allow users to work with packages in the repository, you must grant Repository Privileges, which include Package and System Privileges. All Repository Privileges can be granted or revoked using GRANT or REVOKE statements.

All users who want to access the repository from Eclipse or other clients need the EXECUTE permission for the database procedure SYS.REPOSITORY\_REST:

```
GRANT EXECUTE on SYS.REPOSITORY_REST to <USER/ROLE> [with GRANT OPTION];
```

##### 4.1.4.5.1 Package Privileges

In the SAP HANA studio, you can manage the Package Privileges on the **Package Privileges** tab.

The SAP HANA database repository is structured hierarchically with packages assigned to other packages as subpackages. If you grant privileges to a user for a package, the user is automatically also authorized for all corresponding subpackages.

In the SAP HANA database repository a distinction is made between native and imported packages. Native packages are packages that were created in the current system and should therefore, be edited in the current system. Imported packages from another system should not be edited, except by newly imported updates. An imported package should only be manually edited in exceptional cases.

Developers should be granted the following privileges for native packages:

- REPO.READ

This privilege authorizes read access to packages and design time objects, including both native and imported objects.

- REPO.EDIT\_NATIVE\_OBJECTS

This privilege authorizes all kinds of inactive changes to design time objects in native packages.

- REPO.ACTIVATE\_NATIVE\_OBJECTS

This privilege authorizes the user to activate or reactivate design time objects in native packages.

- REPO.MAINTAIN\_NATIVE\_PACKAGES

This privilege authorizes the user to update or delete native packages, or create subpackages of native packages.

Developers should only be granted the following privileges for imported packages in exceptional cases:

- REPO.EDIT\_IMPORTED\_OBJECTS

This privilege authorizes all kinds of inactive changes to design time objects in imported packages.

- REPO.ACTIVATE\_IMPORTED\_OBJECTS

This privilege authorizes the user to activate or reactivate design time objects in imported packages.

- REPO.MAINTAIN\_IMPORTED\_PACKAGES

This privilege authorizes the user to update or delete imported packages, or create subpackages of imported packages.

#### 4.1.4.5.2 System Privileges

In the SAP HANA studio, you can manage the repository System Privileges together with the other System Privileges on the **System Privileges** tab.

- REPO.EXPORT

This privilege authorizes the user to export, for example, delivery units.

- REPO.IMPORT

This privilege authorizes the user to import transport archives.

- REPO.MAINTAIN\_DELIVERY\_UNITS

This privilege authorizes the user to maintain delivery units (DU, DU vendor and system vendor must be the same).

- REPO.WORK\_IN\_FOREIGN\_WORKSPACE

This privilege authorizes the user to work in a foreign inactive workspace.

#### 4.1.4.6 Authorization for Activated Objects

The repository is part of the SAP HANA database. It stores both runtime data, such as calculation scenarios, and design time data, such as models used in analytic scenarios (Analytic, Calculation, and Attribute Views, as well as Analytic Privileges). Such design time objects must be activated to become runtime objects so that they can be used by regular users of SAP HANA and SAP HANA database.

The following section describes how the authorization for such objects is handled during and after activation. Note that the description does not cover details about the authorization during their lifetime as design time objects within the repository.

**Note:**

If you update from SAP HANA database Revision 10 or 11 to SAP HANA database Revision 12, different privileges for activated objects apply. For more information, see SAP Note [1605168](#).

#### 4.1.4.6.1 During Design Time and Activation

Inside the repository, only the `_SYS_REPO` user is used. This user is the owner of:

- All tables used in the repository
- All activated objects such as procedures, views, and Analytic Privileges

Therefore, only the `_SYS_REPO` user initially has privileges on those objects. However, `_SYS_REPO` is a technical user that does not log on and is only used internally.

#### 4.1.4.6.2 Privileges for `_SYS_REPO`

To perform the tasks described here, `_SYS_REPO` requires the following privileges:

- CREATE SCHEMA
- SCENARIO ADMIN
- CREATE STRUCTURED PRIVILEGE
- STRUCTUREDPRIVILEGE ADMIN

These privileges have already been granted internally to the user.

In addition, `_SYS_REPO` also requires the SELECT privilege on the data behind the modeled content with the option to grant this access to other users. This can be either on a per-table basis or via the schema.

**Note:**

If this privilege is missing, the activated views will be invalidated.

#### 4.1.4.6.3 Privilege Management

Only the `_SYS_REPO` user has any privileges on the created objects. Therefore, only this user can grant privileges on them. Since no user can log on as `_SYS_REPO`, another means of granting privileges is used.

This is provided by stored procedures in the `_SYS_REPO` schema. These procedures can be used to grant and revoke privileges on activated objects or schemas or grant and revoke Analytic Privileges. Stored procedures are beneficial because a user is not required to have a privilege in order to grant it.

The following procedures exist:

- GRANT\_PRIVILEGE\_ON\_ACTIVATED\_CONTENT(privilege IN VARCHAR(256), objectName IN VARCHAR(256), grantee IN VARCHAR(256))
- REVOKE\_PRIVILEGE\_ON\_ACTIVATED\_CONTENT(privilege IN VARCHAR(256), objectName IN VARCHAR(256), revokee IN VARCHAR(256))
- GRANT\_ACTIVATED\_ANALYTICAL\_PRIVILEGE(objectName IN VARCHAR(256), grantee IN VARCHAR(256))
- REVOKE\_ACTIVATED\_ANALYTICAL\_PRIVILEGE(objectName IN VARCHAR(256), revokee IN VARCHAR(256))
- GRANT\_SCHEMA\_PRIVILEGE\_ON\_ACTIVATED\_CONTENT (privilege IN VARCHAR(256), objectName IN VARCHAR(256), grantee IN VARCHAR(256))
- REVOKE\_SCHEMA\_PRIVILEGE\_ON\_ACTIVATED\_CONTENT (privilege IN VARCHAR(256), objectName IN VARCHAR(256), revokee IN VARCHAR(256))

There are public synonyms for those procedures with the same names. Therefore, these procedures can be used without specifying schema `_SYS_REPO` in front.

Having the EXECUTE privilege on any of the procedures enables a user to grant or revoke privileges. Using stored procedures and a technical user for privilege management also changes the behavior with regards to the revocation of privileges:

With regular SQL, privileges that were granted by a user are revoked when this user is dropped or loses the privilege that was granted. Also, only the granter can revoke privileges with SQL. Both details are not true with this approach. Any user with EXECUTE privilege on the revoke privilege procedure can revoke any privilege that was granted, regardless of the granter. Also, if a user that has granted privileges is dropped, none of the privileges that the user granted is revoked as part of dropping the user.

When using the SAP HANA studio for privilege management, this behavior is hidden. If privileges on activated objects or schemas are granted or revoked, the procedures are used automatically. Analytic Privileges are always granted and revoked using the procedures.

#### **4.1.4.7 Critical Combinations**

If users can change and activate objects and are also allowed to grant privileges on activated objects, they have access to all the SAP HANA content.

#### **4.1.4.8 Check of Privileges and Roles for Users and Roles**

To verify which privileges and roles have been granted to a user or role or have been granted by a user, you can check the GRANTED\_PRIVILEGES system view.

If you need to find out which privileges one single user has, you can check the `EFFECTIVE_PRIVILEGES` system view with the restriction for the user wanted (`WHERE user_name = '<user_name>'`).

The difference between those two system views is the fact, that in `EFFECTIVE_PRIVILEGES` all privileges granted directly or indirectly via a role are shown as single privileges, whereas `GRANTED_PRIVILEGES` only shows the roles directly granted, no matter, how many roles and privileges belong to these roles granted to this user indirectly.

`EFFECTIVE_PRIVILEGES` does not contain any information about privileges by the user specified, only those granted to the user.

For more information about these system views, see the *SAP HANA Database Administration Guide* at <http://help.sap.com/hana>.

## 4.1.5 Network and Communication Security

Your network infrastructure is extremely important in protecting your system. Your network needs to support the communication necessary for your business needs without allowing unauthorized access. A well-defined network topology can eliminate many security threats based on software flaws (at both the operating system and application level) or network attacks such as eavesdropping. If users are not able to connect to the server LAN (local area network), they cannot exploit well-known bugs and security holes in network services on the server machines.

The security guidelines and recommendations that apply to the SAP HANA database are described in the following sections:

- [Communication Channel Security](#)

This section describes the communication paths and protocols used by the SAP HANA database.

- [Network Security](#)

This section describes the recommended network topology for the SAP HANA database. It shows the appropriate network segments for the various client and server components and where to use firewalls for access protection. It also includes a list of the ports needed to operate the SAP HANA database.

For more information, see the SAP Library on SAP Help Portal at <http://help.sap.com> under **SAP NetWeaver > SAP NetWeaver 7.3 > System Administration > Security Guide > SAP NetWeaver Security Guide**:

- [Network and Communication Security](#)
- [Security Guides for Connectivity and Interoperability Technologies](#)

## 4.1.6 Communication Channel Security

The table below shows the communication channels used by SAP HANA, the protocol used for the connection and the type of data transferred.

Table 4-5: Communication Paths

Communication Path	Protocol Used	Typ of Data Transferred	Data Requiring Special Protection
Client Access (for example, replication, application server, end-user client, modeling, SAP HANA studio)			
SAP HANA database to data providers	ODBC/JDBC over TCP (SSL supported)	All application data	All application data
SAP HANA database to admin client	ODBC/JDBC over TCP (SSL supported)	User data, configuration data, trace files For modeling: Data models	User data, configuration data, trace files For modeling: Data models
SAP HANA database to end-user clients	ODBC/JDBC over TCP (SSL supported)	All application data	All application data
Administrative Access			
SAP Start Service	HTTP/HTTPS	Configuration data, trace files	Configuration data, trace files
Software Update Manager (SUM) with SAP HANA studio	HTTP/HTTPS	Configuration data	
SUM with SAP Host Agent	HTTPS	Configuration data	
SUM with Service Marketplace	HTTPS	Configuration data	
Database Internal Communication			

Communication Path	Protocol Used	Typ of Data Transferred	Data Requiring Special Protection
SAP HANA database internal communication and communication between SAP HANA database instances in distributed installations	TCP	All application data Configuration data	All application data Configuration data

### 4.1.6.1 Securing Data Communication

As shown in the table above, SAP HANA supports encrypted communication for the client to server communication.

We recommend to use encrypted channels in all cases where network attacks like eavesdropping are not protected by other network security measures (for example, access from end-user networks). For more information about encrypted communication, see [Secure Communication](#).

For communication within the SAP HANA database, explicit security measures are recommended, see [Network Security](#). As an alternative, VPN tunnels can be used for the transfer of encrypted information.

### 4.1.7 Network Security

When using SAP HANA appliance software, we recommend to operate different components of the solution in separate network segments. In order to prevent any unauthorized access to the SAP HANA appliance and the SAP HANA database via the network, we recommend controlling the network traffic between the different network segments by using a firewall or a packet filter. For more information about additional security mechanisms using encrypted communication, see [Secure Communication](#).

The system landscape pictured in the [Technical System Landscape](#) section gives an overview of the different network segments that, depending on the individual configuration, are available. The detailed setup is dependent on the specific application scenario and customer network infrastructure.

The SAP HANA appliance should be operated in a protected data center environment. Only dedicated authorized network traffic should be allowed from other network zones (for example, user access from client network zone):

- Client access (that is, all access to external standard database functionality, for example, SQL) only requires access to the client access port.

**Note:**

In distributed scenarios, clients must be able to access every node of the distributed SAP HANA appliance.

- Client HTTP access (for example, browser) in scenarios that use the HTTP access feature of SAP HANA (XSEngine), for example, ETL-based Data Acquisition by SAP HANA Direct Extractor Connection and SAP HANA UI Toolkit for INA.
- For some administrative functions (for example, starting and stopping the SAP HANA instance), access to the administrative ports is additionally required.
- Database internal communication is only used for communication within the database or in a distributed scenario, for the communication between hosts.
  - In a single blade scenario (one instance of SAP HANA on one blade), access to those ports from other network hosts must be blocked.
  - In a distributed scenario of SAP HANA (one instance of SAP HANA on multiple blades), we recommend to operate all blades in a dedicated subnet. We further recommend to ensure that communication on the internal communication channels is restricted to communication between authorized hosts of an instance.

**Caution:**

The internal communication must be strictly separated from the external or client communication paths. Access from hosts that are not part of an instance of the SAP HANA appliance should be blocked.

If your setup does not allow having the internal communication in a dedicated subnet, we recommend protecting the internal communication using encryption. For more information, see [Secure Communication](#).

Additional network configurations may be required for specific replication scenarios. For more information about the SAP HANA replication technologies, see the *SAP HANA Technical Operations Manual* at <http://help.sap.com/hana>.

### 4.1.7.1 Communication Ports

The table below lists the ports that are used by SAP HANA. We recommend controlling the network traffic between the different network segments by using a firewall or a packet filter.

**Tip:**

Block all access to other ports in the firewall that are not used by the SAP HANA database.

**Note:**

In certain scenarios, additional communication channels (for example, for remote operating system access) may be required.

The notation of the ports is as follows: n <instance> xy, where n is either 3 or 5 (see table below), <instance> is a two-digit number representing the instance number of the SAP HANA appliance, and xy represents a consecutive number.

### Communication Ports for Inbound Communication

Port Number	Used for
Client Access	
3<instance>15	Standard SQL communication for client access. This is the only port required for client access.
80<instance>/43<instance>	XSEngine (HTTP/HTTPS). Only enabled in scenarios that use XSEngine (for example, ETL-based Data Acquisition by SAP HANA Direct Extractor Connection).
Administrative Access	
5<instance>13 5<instance>14 (SSL)	System administration (for example, startup and shutdown) and communication between SUM and SAP Start Service on different hosts.  For more information about the SAP Start Service, see the SAP Library on SAP Help Portal at <a href="http://help.sap.com">http://help.sap.com</a> under <b>SAP NetWeaver &gt; SAP NetWeaver 7.3 &gt; Functional View &gt; SAP NetWeaver by Functional Areas &gt; Application Server &gt; Application Server Infrastructure &gt; Architecture of the SAP NetWeaver Application Server &gt; SAP Start Service</b> .  <b>Note:</b>  For SAP HANA appliance software SPS 04, the SAP Start Service is only used to start and stop an instance of the SAP HANA database and to monitor an instance of the SAP HANA database.
8080/8443	Software Update Manager (SUM) access (HTTP/HTTPS)
Database Internal Communication	

Port Number	Used for
3<instance>00	Used for database internal communication only. These ports should only be accessible from other hosts of the SAP HANA appliance.
3<instance>01	
3<instance>02	
3<instance>03	
3<instance>05	
3<instance>07	

### Communication Ports for Outbound Communication

The SUM connects to the SAP Service Marketplace to check if new updates for the SAP HANA software are available. In order to do so, the outbound communication channel from the SUM to the SAP Service Marketplace's address <https://service.sap.com> must be enabled by the customer's network setup.

## 4.1.8 Secure Communication

The SAP HANA appliance uses the secure sockets layer (SSL) protocol to provide secure communication between the individual components and client connections. Authentication is ensured by using certificates.

The communication between the following components can be secured by using SSL:

- Any ODBC- or JDBC-based connection
- SAP HANA studio – SAP HANA database (server authentication)

For more information, see [Configuring HTTPS Between SAP HANA Database and SAP HANA Studio](#).

- SAP HANA studio – Software Update Manager for SAP HANA

For more information about how to configure HTTPS for SAP HANA studio, see the *SAP HANA Automated Update Guide* at <https://service.sap.com/hana>.

- Software Update Manager for SAP HANA – SAP Service Marketplace

SAP HANA needs an SAP Service Marketplace user (S-user) to access the SAP Service Marketplace. These credentials are sent only via encrypted communication channels using an HTTPS connection. For more information about how to configure access to the SAP Service Marketplace, see the *SAP HANA Automated Update Guide* at <https://service.sap.com/hana>.

- Software Update Manager for SAP HANA – SAP Host Agent

For more information about how to configure HTTPS for the SAP Host Agent, see the *SAP HANA Automated Update Guide* at <https://service.sap.com/hana>.

- SAP HANA information composer – Browser  
For more information, see [SAP HANA Information Composer](#).
- Internal communication among the different components of a running SAP HANA database system  
For more information, see [Configuring SSL for SAP HANA Database Internal Communication](#).

### 4.1.8.1 Configuring HTTPS Between SAP HANA Database and SAP HANA Studio

Depending on the operating system used, the SAP HANA appliance software (SPS 03 and higher) supports different cryptographic library providers:

- For Microsoft Windows-based installation:
  - SChannel (default)
  - SAP Crypto
- For Linux-based installation:
  - OpenSSL (default)
  - SAP Crypto

#### 4.1.8.1.1 Setup on Server-Side

To protect your data during network transmission, only secure connections should be used. We recommend using the tools provided with OpenSSL to create the certificates required for SSL configuration.

##### **Prerequisites**

- The server possesses a public and private key pair and public-key certificate.

The SSL protocol uses public-key technology to provide its protection. Therefore, the server must possess a public and private key pair and a corresponding public-key certificate. It must possess one key pair and certificate to identify itself as the server component and another key pair. The key pair and certificate are stored in the server's own personal security environments (PSE), the SSL server PSE, and the SSL client PSE, respectively.

##### **Note:**

In case, your server keys are compromised, replace the certificate.

- You have installed a cryptographic provider such as OpenSSL or the SAP Cryptographic Library.

##### **Caution:**

The distribution of the SAP Cryptographic Library is subject to and controlled by German export regulations and is not available to all customers. In addition, usage of the SAP Cryptographic Library

or OpenSSL library may be subject to local regulations of your own country that may further restrict the import, use, and (re-)export of cryptographic software. If you have any further questions about this issue, contact your local SAP office.

## Features

By supporting SSL, the SAP HANA appliance software can provide the following:

- Server-side authentication

With server-side authentication, the server identifies itself to the client when the connection is established. This reduces the risk of using "fake" servers to gain information from clients.

- Data encryption

In addition to authenticating the communication partners, the data being transferred between the client and server is encrypted which provides for integrity and privacy protection. An eavesdropper cannot access or manipulate the data.

Client-side authentication and mutual authentication are not currently supported.

The following parameters can be used to configure the server connectivity. They are located in the `indexserver.ini` file, in the communication section.

### Note:

Configuration of cryptographic library providers is optional.

The parameters in the following table can be configured for the setup of secure connections.

Table 4-7: Configuration Parameters on Server-Side

Property Name	Property Value	Default	Description
sslCryptoProvider	{sapcrypto   openssl}	1. sapcrypto (if installed) 2. openssl	Cryptographic library provider to use for SSL connectivity.
sslKeyStore	<file>	\$HOME/.ssl/key.pem	Path to keystore file.
sslTrustStore	<file>	\$HOME/.ssl/trust.pem	Path to trust store file.
sslValidateCertificate	<bool value>	false	If set to true, validate the certificate of the communication partner.
sslCreateSelfSignedCertificate	<bool value>	false	If set to true, create a self-signed certificate if the keystore cannot be found.

**No Configuration Provided**

If no configuration for secure connections has been provided, the system determines which cryptographic library provider should be used as follows:

1. Checks whether the environment variable `SECUDIR` is set.
  - a. If the environment variable `SECUDIR` is set, it tries to load the `sapcrypto` library using first `SNCPATH` from the environment and then the regular paths for library lookup.
  - b. If `sapcrypto` cannot be loaded, it proceeds with the next cryptographic library provider.
  - c. If `sapcrypto` was loaded, it uses the path names given in `sslKeyStore` and `sslTrustStore` to check for a `*.pse` store.
  - d. If a PSE store could be found, the system verifies its integrity.
  - e. If no PSE store could be found or the PSE store's integrity could not be verified, SSL initialization fails and SSL is not available.
2. Checks whether OpenSSL is available.
  - a. If OpenSSL is available, it checks for key certificates at the path given in `sslKeyStore` and trusted certificates at the path given in `sslTrustStore`.
  - b. If any certificates were found, it checks for the integrity of the certificates.
  - c. If any of the above fails, SSL initialization fails and SSL is not available.

**Configuration Provided**

- If the value of the `sslCryptoProvider` parameter is set, the system tries to initialize the given cryptographic library provider. Any other installed cryptographic library providers are ignored.
- If the value of the `sslCryptoProvider` parameter is set but no paths are given for the `sslKeyStore` and `sslTrustStore` parameters, the system uses the default paths for initialization as if no configuration were provided.
- If the value of the `sslKeyStore` parameter or the `sslTrustStore` parameter is set, the system does not check the default paths. In this case, the `sslCryptoProvider` parameter must be set.
- If the values of both the `sslKeyStore` parameter and the `sslTrustStore` parameter are set, a value for the `sslCryptoProvider` parameter also has to be set; otherwise SSL initialization fails and SSL is not available.

**4.1.8.1.2 Setup on Client-Side (SQLDBC-Based Connections)**

Set the parameter values according to the operating system installed on the clients. For SQLDBC-based connectivity (for example ODBC), the parameters and their names are the same as for the server. Additionally, the `encrypt` parameter is available to initiate an SSL-secured connection.

Table 4-8: Configuration Parameters on Client-Side for SQLDBC-Based Connections

Property Name	Property Value	Default	Description
encrypt	<bool value>	False	Enables or disables SSL encryption.
sslCryptoProvider	{sapcrypto   openssl   mscrypto}	1. sapcrypto (if installed) 2. openssl/mscrypto	Cryptographic library provider to use for SSL connectivity.
sslKeyStore	<file>	\$HOME/.ssl/key.pem	Path to keystore file. Leave empty when using mscrypto.
sslTrustStore	<file>	\$HOME/.ssl/trust.pem	Path to trust store file. Leave empty when using mscrypto.
sslValidateCertificate	<bool value>	true	If set to true, validate the certificate of the communication partner.
sslHostNameInCertificate	<string value>	<empty>	Use the given host name for validation. <b>Tip:</b> Use this host name when validating the communication partner's certificate. Wildcards are not allowed. If the given host name is "*" then host name validation is disabled.
sslCreateSelfSignedCertificate	<bool value>	false	If set to true, create a self-signed certificate if the keystore cannot be found.

#### 4.1.8.1.3 Setup on Client-Side (JDBC-Based Connections)

For JDBC connections, the parameter names are the same as those for SQLDBC-based connections except for the missing prefix SSL. Additionally, some additional parameters to further characterize the

(Java-based) keystore and its password are used. If you use JDBC connections, deploy the certificates to the Java keystore.

For JDBC connections, the automatic creation of a self-signed certificate is currently not supported. Therefore, the `createSelfSignedCertificate` parameter is not available.

Table 4-9: Configuration Parameters on Client-Side for JDBC-Based Connections

Property Name	Property Value	Default	Description
<code>encrypt</code>	<bool value>	false	Enables or disables SSL encryption.
<code>validateCertificate</code>	<bool value>	true	If set to true, validate the certificate of the communication partner.
<code>hostNameInCertificate</code>	<string value>	<empty>	Use the given host name for validation. <b>Tip:</b> Use this host name when validating the communication partner's certificate. Wildcards are not allowed. If the given host name is "*" then host name validation is disabled.
<code>keyStore</code>	<file   store name>	<VM default>	
<code>keyStoreType</code>	<JKS   PKCS12>	<VM default>	
<code>keyStorePassword</code>	<password>	<VM default>	Password used to access the keystore.
<code>trustStore</code>	<file   store name>	<VM default>	
<code>trustStoreType</code>	<JKS>	<VM default>	
<code>trustStorePassword</code>	<password>	<VM default>	Password used to access the trust store.

If you do not specify any values for the \*Store\* parameters, the system uses the default values.

#### 4.1.8.1.4 Setup of SAP HANA Studio Connections (JDBC-Based-Connections)

As a prerequisite for SSL-secured connections to and from SAP HANA studio, the root certificate that was used to sign the server certificate must be available in the Java trust store. SAP HANA studio allows you to use either the system-wide trust store or the default user trust store for certificate validation. For more information about how to import certificates into trust stores, see the Java documentation.

### 4.1.8.2 Configuring SSL for SAP HANA Database Internal Communication

The certificates for internal network communication in the SAP HANA appliance software SPS 04 are specific for each host and different for the client and server side. This is necessary as every host shall be verified with its fully qualified domain name (FQDN). As the SAP HANA database deals with a set of certificates, we recommend using a dedicated certificate authority (CA) to sign these.

1. Download the SAP Cryptographic Library:

The standard installer does not provide the required binaries. You have to download them separately. The SAP Cryptographic Library is available at the SAP Service Marketplace.

2. Create a certificate authority (CA) designated to this installation using external tools, for example, the OpenSSL command line tool.

We recommend to store your CA certificate in `$DIR_INSTANCE/ca`.

3. Create certificates:

On every host you have to create the client-side and the server-side certificate. You have to sign these at the CA just created. The common name (CN) has to be the FQDN of the host you get by reverse DNS lookup. The other fields describe your organization. Make sure that the client-side certificate is created without a password. Create a local keystore named `SAPSSLC.pse` in directory `$SECUDIR` on every host and import the host's client certificate into `SAPSSLC.pse`.

4. Activate secure sockets:

Add the section [communication] to the custom layer of the file `global.ini`. Set the key `ssl = on`.

### 4.1.9 Data Storage Security

The SAP HANA database is stored in the file system (including configuration data). You can configure the base path during installation. If you change the base path, make sure that all instances of the distributed system still have access to the file system. For more information about how to create a distributed system, see the *SAP HANA Database - Server Installation and Update Guide* at <http://help.sap.com/hana>.

### 4.1.9.1 Data Protection on File System

The file permissions of the operating system are strictly configured. Therefore, we recommend that you do not change them after the installation of the SAP HANA database.

### 4.1.9.2 Data Encryption

For SAP HANA appliance software SPS 04, no data encryption is available. If data encryption is required in a specific scenario, we suggest you use the encryption features of the respective operating system or storage provider.

### 4.1.9.3 Secure Data Storage (SAP HANA)

The following properties apply to the secure data storage:

- System passwords are protected by the methods of the respective operating systems (for example, `/etc/passwd` in UNIX).
- All database user passwords on the SAP HANA database server are hashed with the secure hash algorithm SHA-256.
- Passwords used for authentication between components (for example, for scripting) are stored in the SAP secure password store in an encrypted form.
- Eclipse secure store (when using SAP HANA studio) is used to store saved passwords.

### 4.1.9.4 Secure User Store

In the secure user store of the SAP HANA client (`hdbuserstore`), you can securely store login information for the users, including passwords. This allows client programs to connect to the database without explicitly logging in. The secure user store is installed with the SAP HANA client package. After installation, it is located in the `/usr/sap/hdbclient` directory. The secure user store runs on all platforms supported by SAP HANA client interfaces and SAP BASIS 7.20 EXT.

The login information is stored in one of the following directories. If the path does not already exist, it is created by the hdbuserstore command.

- For systems using Microsoft Windows, the path is defined by <PROGRAMDATA>\.hdb\<COMPUTERNAME><SID>.

Where PROGRAMDATA is the path defined by CSIDL\_COMMON\_APPDATA resp. FOLDERID\_ProgramData and SID is the system ID of the user that uses the stored login information.

- For systems using other operating systems, the path is defined by <HOME>/ .hdb/<COMPUTERNAME>.
- HOME is the home directory of the user that uses the login information.

When executing the hdbuserstore script (in the context of the correct operating system user), the user store can be opened using a user key. Only the operating system user who writes the login information to the secure password store can access the corresponding files.

To edit the stored login information, you can use the following hdbuserstore commands:

Command	Parameter	Description
HELP		Displays a help message.
LIST	<user_key>	Lists entries with the key. Passwords are not displayed.
DELETE	<user_key>	Deletes entries with the key.
SET	<user_key>	Sets the entry key.
	<env>	Sets the connection environment (host and port).
	<user_name>	Sets the user name for the profile.
	<password>	Sets the password for the profile.

#### Example:

- Create a user key in the user store and store the password under this user key:

```
hdbuserstore SET <user_key> <env> <user_name> <password>
```

For example:

```
hdbuserstore SET millerj localhost:30115 JohnMiller 2wsx$RFV
```

- List all available user keys (passwords are not displayed):

```
hdbuserstore LIST <user_key>
```

For example:

```
hdbuserstore LIST millerj
```

The following information is displayed:

KEY: millerj

ENV: localhost:30115

USER: JohnMiller

- Call `hdbsql` with the user key:

```
hdbsql -U <user_key>
```

For example:

```
hdbsql -U millerj
```

---

### 4.1.9.5 Encryption Keys

All stored password information is encrypted using an encryption key. The system is provided with a default encryption key. If the encryption key is compromised, you can change the key.

#### **Caution:**

If the user forgets the stored password, you cannot recover that password because the system does not display passwords in a human-readable form. We recommend changing the encryption key.

#### **To change encryption keys**

1. Get the RSECSSFX command from SAP BASIS 7.20 EXT.
2. Specify the path based on the platform, as described above. The key path is the same as the data path.
3. Define the SAP system name as HDB.
4. Use the CHANGEKEY command to change the key.

### 4.1.9.6 History Tables

If you want to delete selected rows in the history table, follow the description in the *SAP HANA Database - SQL Reference Guide* at <http://help.sap.com/hana>.

## 4.1.10 Security Logging and Tracing

The main requirement for auditing a system is traceability of actions performed in that system. The main question is: Who did or tried to do what when?

Auditing does not directly increase the security of the system but if wisely designed, it can help:

- Uncover security holes
- Show security breaches and privilege misuses
- Protect the system owner against accusations of security violation and data misuse
- System owners to meet their security standards

### **Note:**

In the current version of the SAP HANA database, security logging and tracing is supported using the operating system log files (syslog). By default, the log files are written to the following location: `/var/log/messages`.

### 4.1.10.1 Auditing in a Running System

#### 4.1.10.1.1 Audit Policies

Audit policies define which events to audit. Each policy has a name and can be enabled or disabled by an administrator having the `AUDIT_ADMIN` privilege. Audit policies are owned by the system; they are not dropped when the creating user is removed. The policy has several further attributes, which are used to narrow the number of events that are audited.

#### **Policy Actions**

The action list describes the list of database actions triggering this particular audit policy. Actions are organized in different groups, some of them needing a target object, while others do not need any target object. Actions belonging to both these categories must never be combined in the same audit policy. Some actions can combine different database actions to one single entity, meaning that a policy having only one single audit action can fire on very different SQL statements. The action list can contain an arbitrary number of audit actions, but must have at least one action.

For more information about the possible actions and their meaning, the action statuses and audit levels, see the *SAP HANA Database – SQL Reference Guide* at <http://help.sap.com/hana>.

#### 4.1.10.1.2 Audit Entries

The table below shows the fields of audit entries and their meaning.

Table 4-11: Fields of Audit Entries

Field	Description	Example Value
Event Timestamp	When did the event occur? This field is in coordinated universal time (UTC).	
Service Name	Which service did the event occur in?	Indexserver
Hostname	Hostname of the system on which instance is running	myhanablade23.customer.corp
SID	Auditing system identifier	HAN
Instance Number	Instance number of the service in which the event occurred	23
Port Number	Port number of the service in which event occurred	32315
Policy Name	Name of the audit policy that was triggered	AUDIT_GRANT
Audit Level	Level of the audit policy that was triggered	WARNING
Audit Action	Action that triggered the event	GRANT PRIVILEGE
Active User	User that executed the statement	MYADMIN
Target Schema	On which schema was the privilege granted or in which schema is the target object?	PRIVATE
Target Object	On which object was the privilege granted?	
Privilege Name	Which privilege was granted or revoked?	SELECT

Field	Description	Example Value
Grantable	Was the privilege or role granted with or without GRANT/ADMIN OPTION?	NON GRANTABLE
Role Name	Which role was granted or revoked?	
Target Principal	Who was the target of the action? Useful for grant or revoke statements.	ATTACKER
Action Status	Was the statement or operation successful?	SUCCESSFUL
Component	Component affected by configuration change	
Section	Section affected by configuration change	
Parameter	Parameter affected by configuration change	
Old Value	Old value before configuration change	
New Value	New value after configuration change	
Comment	Additional information on current event	
Executed Statement	Statement that was executed	GRANT SELECT ON SCHEMA PRIVATE TO ATTACKER
Session Id	ID of the session the statement was executed in	400006

#### 4.1.10.1.3 Audit Trail

Currently, only the syslog protocol is a supported audit trail target. The syslog provides a means of safely storing the audit trail that even the database administrator cannot access or change. Numerous storage possibilities exist for the syslog, including storing it on other systems.

In addition, syslog is the default log daemon in \*nix systems and therefore it is assumed that many customers already have a strategy in place to deal with syslog entries. This provides a high degree of flexibility and security, as well as many integration possibilities into a larger system landscape.

The message written to the syslog is in comma-separated values (CSV) format so that it can be easily parsed and imported into other systems. The CSV format is as follows:

```
<Event Timestamp>;<Service Name>;<Hostname>;<SID>;<Instance Number>;<Port Number>;<Policy Name>;
<Audit Level>;<Audit Action>;<Active User>;<Target Schema>;<Target Object>;<Privilege Name>;<Grantable>;
<Role Name>;<Target Principal>;<Action Status>;<Component>;<Section>;<Parameter>;<Old Value>;<New Value>;
<Comment>;<Executed Statement>;<SESSION Id>;
```

**Caution:**

You can alter the audit configuration so that the audit trail is written to a text file. This must not be used on production systems. The text file writer has severe limitations. For example, it is not written in a thread-safe manner so that multiple, synchronously written entries can yield unexpected results. However, this can be very useful during the testing of audit policies because the results of policy actions can be more easily determined.

## 4.1.10.2 Global Audit Configuration

To maintain the global audit configuration, users need the INIFILE ADMIN System Privilege. Then, they can use the following statements to alter the configuration of the auditing system.

### 4.1.10.2.1 Activate or Deactivate Global Auditing

Activate global auditing:

```
alter system alter configuration ('global.ini','SYSTEM')
set ('auditing configuration'
,'global_auditing_state' ) = 'true'
with reconfigure;
```

Deactivate global auditing:

```
alter system alter configuration ('global.ini','SYSTEM')
set ('auditing configuration'
,'global_auditing_state' ) = 'false'
with reconfigure;
```

### 4.1.10.2.2 Set Audit Trail Type

Set the audit trail type CSVTEXTFILE:

**Note:**

This is for testing purposes only; do not use in productive environments.

```
alter system alter configuration ('global.ini','SYSTEM')
set ('auditing configuration'
,'default_audit_trail_type' ) = 'CSVTEXTFILE'
with reconfigure;
```

Set the audit trail type SYSLOGPROTOCOL:

```
alter system alter configuration ('global.ini','SYSTEM')
set ('auditing configuration'
```

```
, 'default_audit_trail_type' ) = 'SYSLOGPROTOCOL'
with reconfigure;
```

#### 4.1.10.2.3 Set Audit Target Path

To set the audit target path, the audit trail type `CSVTEXTFILE` must be set:

**Note:**

This is for testing purposes only; do not use in productive environments.

```
alter system alter configuration ('global.ini','SYSTEM')
set ('auditing configuration'
, 'default_audit_trail_path' ) = '<path>'
with reconfigure;
```

### 4.1.10.3 Manage Audit Policies

To manage audit policies, the `AUDIT ADMIN` System Privilege is required.

#### 4.1.10.3.1 Create Audit Policy

```
CREATE AUDIT POLICY <policy_name> AUDITING <ACTIONSTATUS> <ACTIONLIST> LEVEL <AUDITLEVEL>;
```

For more information about the possible values for `<ACTIONSTATUS>`, `<ACTIONLIST>`, and `<AUDITLEVEL>`, see the *SAP HANA Database – SQL Reference Guide* at <http://help.sap.com/hana>.

**Note:**

By default, all audit policies are created in the inactive state.

#### 4.1.10.3.2 Enable or Disable Audit Policy

Enable the audit policy:

```
ALTER AUDIT POLICY <policy_name> ENABLE;
```

Disable the audit policy:

```
ALTER AUDIT POLICY <policy_name> DISABLE;
```

#### 4.1.10.3.3 Drop Audit Policy

```
DROP AUDIT POLICY <policy_name>;
```

#### 4.1.10.3.4 Show Audit Policies

```
SELECT * FROM "PUBLIC"."AUDIT_POLICIES"
```

#### 4.1.10.3.5 Example

Audit logging for creating or dropping of roles and users, write audit messages with level "critical":

**1. Create audit policy:**

```
CREATE AUDIT POLICY policyAdministrators AUDITING ALL  
CREATE ROLE, DROP ROLE, CREATE USER, DROP USER LEVEL Critical;
```

**2. Activate audit policy:**

```
ALTER AUDIT POLICY policyAdministrators ENABLE;
```

# SAP HANA - Additional Components

In addition to the SAP HANA database, the following components are part of the SAP HANA landscape:

- SAP HANA information composer
- Lifecycle Management Tools
- Unified Installer
- SAP HANA UI toolkit for INA

## 5.1 SAP HANA Information Composer

The SAP HANA information composer is a Web application that allows you to upload and manipulate data on the SAP HANA database. The SAP HANA information composer uses a Java server which interacts with the SAP HANA database.

The Java server communicates with the SAP HANA information composer client via HTTP or HTTPS. The following ports are used by default:

- HTTP port 8080
- HTTPS port 8443

If HTTPS is used, the SSL certification must be configured by the administrator.

### **Note:**

The SAP HANA information composer can be configured to use antivirus software.

The SAP HANA information composer client is accessible to users who are assigned the IC\_MODELER role. This role allows users to upload new content into the SAP HANA database and to create physical tables and calculation views.

When content is marked as shared, it is accessible from users who are assigned the IC\_PUBLIC role. By default, the physical tables and calculation views are marked as private. This means that they are only visible to the user who created them. Calculation views are created by user `_SYS_REPO` in schema `_SYS_BIC` within the column Views.

The physical tables and calculation views can be shared with users who are assigned the IC\_PUBLIC role. The IC\_PUBLIC role is included in the IC\_MODELER role.

The created calculation view inherits the analytical privileges of the source data that is being used. Objects that are based on user data (spreadsheets) have no analytical privileges.

The technical user SAP\_IC is created during installation. After completing the installation, SAP\_IC is locked.

**Note:**

As long as the SAP HANA information composer is in use, the SAP\_IC user must not be deleted because otherwise, the IC\_MODELER and IC\_PUBLIC roles will also be deleted.

For more information about how to install and configure the SAP HANA information composer, see the *SAP HANA Information Composer – Installation and Configuration Guide* at <https://service.sap.com/hana>.

## 5.2 Lifecycle Management Tools

You can access the Lifecycle Management Tools from the "Lifecycle Management" perspective of SAP HANA studio. The Software Update Manager (SUM), which is part of the Lifecycle Management Tools, can be used to update the components of your SAP HANA installation.

To work properly, the SUM needs credentials for the following users:

- sapadm – used to authenticate to SAP Host Agent
- <sid>adm – required by SAP HANA database server update
- SAP Service Marketplace user – used to authenticate to SAP Service Marketplace

The SUM for SAP HANA communicates with the following components:

- SAP HANA studio
- SAP Service Marketplace
- SAP Host Agent

All these channels use encryption via HTTPS. For communication with the SAP HANA studio, the SUM for SAP HANA opens the server ports 8080 and 8443.

See the *SAP HANA Automated Update Guide* at <https://service.sap.com/hana> for more information about:

- How to set up and update the SUM (section "Configuring HTTPS for SAP HANA Automated Update").
- How to set up and update the Lifecycle Management Perspective (section "Setting Up the SAP HANA Studio").

## 5.3 Unified Installer

The SAP HANA Unified Installer is a tool to install the SAP HANA appliance software in a single, unified and predefined way. It is designed to be used by the SAP HANA hardware partners within their factory process. For more information about running the Unified Installer, see the *SAP HANA Installation Guide with SAP HANA Unified Installer* at <https://service.sap.com/hana>.

## 5.4 SAP HANA UI Toolkit for INA

The SAP HANA UI toolkit provides UI building blocks for developing search-based applications on SAP HANA. Such applications provide real-time information access (INA) and faceted search features on huge volumes of structured and unstructured text data.

By default, the SAP HANA UI toolkit is switched off. For more information, see *SAP HANA UI Toolkit for INA* at <http://help.sap.com/hana>.

**Note:**

Only activate technologies that you explicitly use.



## SAP HANA - Replication Technologies

SAP HANA supports the following replication technologies:

- Trigger-Based Data Replication using SAP LT (Landscape Transformation) Replication Server

SAP Landscape Transformation Replication Server is a replication technology to provide data from SAP systems in a SAP HANA environment. It acts as a key enabler for SAP HANA customers to supply their HANA environment with relevant data. For more information, see [SAP HANA Security Guide - Trigger-Based Data Replication](#).

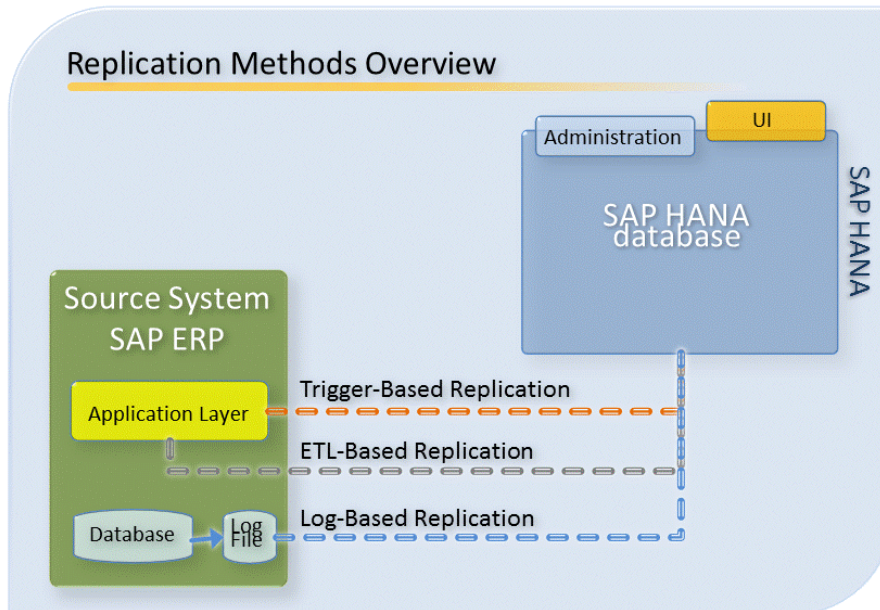
- Extraction-Transformation-Load-(ETL-)Based Replication

For more information, see [SAP HANA Security Extraction-Transformation-Load-\(ETL-\)Based Data Replication](#).

- Log-Based Replication

Transaction Log-Based Data Replication using Sybase Replication

For more information, see [SAP HANA Security Guide - Log-Based Replication](#).



The figure above gives an overview of the three methods for data replication from a source system to SAP HANA appliance software. Each method handles the required data replication differently, and consequently each method has varying security considerations.

For more information about SAP HANA replication methods and technologies, see:

- [SAP HANA Master Guide](https://service.sap.com/hana), section “SAP HANA Replication Technologies” on the SAP HANA Installation & Implementation Knowledge Center at <https://service.sap.com/hana>.
- [SAP HANA Technical Operations Manual](http://help.sap.com/hana), section “SAP HANA Replication Technologies” on the SAP HANA Installation & Implementation Knowledge Center at <http://help.sap.com/hana>.

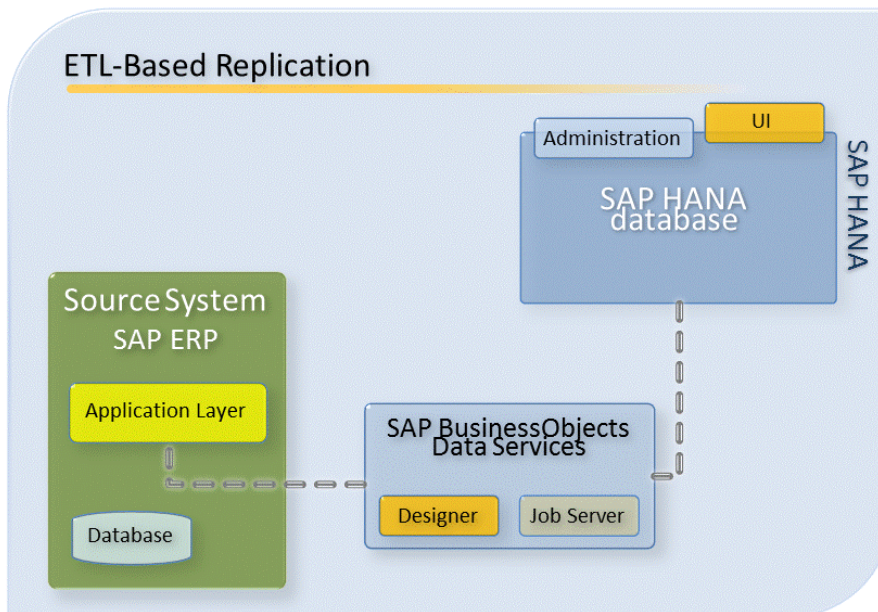
Furthermore, SAP HANA supports ETL-based Data Acquisition by SAP HANA Direct Extractor Connection based on the XSEngine. For more information about security aspects, see [SAP HANA Security ETL-based Data Acquisition by SAP HANA Direct Extractor Connection](#).

## 6.1 SAP HANA Security Extraction-Transformation-Load-(ETL-)Based Data Replication

The ETL-based data replication uses SAP BusinessObjects Data Services (hereafter referred to as Data Services) to load the relevant business data from the source system SAP ERP and replicate it to the target SAP HANA database. This method allows you to read the required business data on the

application layer level. You deploy this method by defining data flows in Data Services and scheduling the replication jobs.

Since this method uses batch processing, it also enables data checks, transformations, synchronizing with additional data providers, and the merging of data streams.



The figure above gives an overview of the ETL-based data replication method. Here, the 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.

For more information about ETL-based data replication using the SAP BusinessObjects Data Services database, see the sections "Security and User" and "Rights Management" in the [SAP BusinessObjects Data Services Administrator's Guide](#).

## 6.1.1 Data Flow

For any replication scenario, you have to define a series of parameters for the two systems involved. You define such parameters in Datastores by using the Designer for setup purposes.

For more information about the use of the Designer, see the SAP HANA Installation & Implementation Knowledge Center at <https://service.sap.com/hana>. Here, you can also find more information about installing and administering Data Services.

## 6.1.2 Datastore Setup

When 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 the password that allows Data Services to access the system. Any 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 in a similar way to that for the source system.

## 6.2 SAP HANA Security ETL-based Data Acquisition by SAP HANA Direct Extractor Connection

By default, the SAP HANA Direct Extractor Connection technology is switched off. For more information about how to switch it on, see the *ETL-based Data Acquisition by SAP HANA Direct Extractor Connection Installation and Configuration Guide* at <http://help.sap.com/hana>.

### **Note:**

Only activate technologies that you explicitly use.

For secure communication, the SAP HANA Direct Extractor Connection technology uses the SSL protocol (HTTPS) based on the Internet Communication Manager (ICM). For more information about ICM and SSL configuration, see the SAP Library on SAP Help Portal at <http://help.sap.com> under **SAP NetWeaver > SAP NetWeaver 7.3 > SAP NetWeaver Library: Function-Oriented View > Application Server > Application Server Infrastructure > Internet Communication Manager (ICM)**.

# Appendix

## 7.1 Password Policy Parameters

The table below contains the password policy parameters and their default values, and explains the function of each parameter.

Table 7-1: Password Policy Parameters

Parameter	Default Value	Description
<code>minimal_password_length</code>	8	Defines the minimum password length. The accepted value range is 6 to 64 characters. The allowed character classes are described directly below in the following table row.
<code>password_layout</code>	A1a	

Parameter	Default Value	Description
		<p>Defines the character classes which can be used within a password:</p> <ul style="list-style-type: none"> <li>• Upper-case ASCII: A-Z</li> <li>• Lower-case ASCII: a-z</li> <li>• Digits: 0-9</li> <li>• Special characters: _</li> </ul> <p>By default, a password must contain one or more upper-case characters, one or more lower-case characters and one or more digits. The use of special characters is optional.</p> <p>A password can contain any of the characters belonging to the character classes indicated above in any order. Each character that is not an upper-case letter, a lower-case letter, or a digit is interpreted as a special character.</p> <p>When the password is enclosed in double quotes during user creation, all Unicode characters may be used in the password.</p> <p><b>Note:</b></p> <p>The use of passwords enclosed in double quotes may cause logon issues, depending on the client used. SAP HANA studio, for example, supports passwords enclosed in double quotes, while the <code>hdbsql</code> command line tool does not.</p> <p>If no value is set for this parameter, the default value is used.</p>
<code>force_first_password_change</code>	true	

Parameter	Default Value	Description
		<p>Defines whether users have to change their initial passwords at first logon.</p> <p>Logging on with the initial password is still possible but only the <code>ALTER USER &lt;current_user&gt; PASSWORD &lt;password&gt;</code> command can be executed. All other statements give the error message user is forced to change password.</p> <p>Administrators can force a user to change the password at any time with the following SQL command:</p> <pre>ALTER USER &lt;user_name&gt; FORCE PASSWORD CHANGE</pre>
<code>maximum_invalid_connect_attempts</code>	6	

Parameter	Default Value	Description
		<p>Defines how many invalid logon attempts are allowed before the user account is locked.</p> <p>Administrators can reset the number of invalid logon attempts with the following SQL command:</p> <pre>ALTER USER &lt;user_name&gt; RESET CONNECT ATTEMPTS</pre> <p>With the first successful logon after an invalid logon attempt, an entry is made into the INVALID_CONNECT_ATTEMPTS view showing:</p> <ul style="list-style-type: none"> <li>• The number of invalid logon attempts since the last successful logon</li> <li>• The time of the last successful logon</li> </ul> <p>Administrators and users can delete the information of invalid logon attempts with the following SQL command:</p> <pre>ALTER USER &lt;user_name&gt; DROP CONNECT ATTEMPTS</pre> <p>If the value of this parameter is set to 0, no check is performed.</p>
password_lock_time	1440	<p>Defines the duration in minutes that a user account is locked after a defined number of failed logon attempts.</p> <p>The default value is set to 1,440 minutes (= 24 hours).</p> <p>Administrators can reset the number of invalid logon attempts and unlock the user account with the following SQL command:</p> <pre>ALTER USER &lt;user_name&gt; RESET CONNECT ATTEMPTS</pre>

Parameter	Default Value	Description
<code>last_used_passwords</code>	5	Defines the number of last used passwords that the user is not allowed to use when changing the current password.
<code>maximum_password_lifetime</code>	182	<p>Defines the duration in days that a password is valid.</p> <p>After the expiry of this validity period, users have to change their password at the next logon.</p> <p>Administrators can exclude users from this password lifetime check with the following SQL command:</p> <pre>ALTER USER &lt;user_name&gt; DISABLE PASSWORD LIFETIME</pre> <p><b>Note:</b></p> <p>This may be performed for technical users only, not for standard database users.</p>
<code>password_expire_warning_time</code>	14	<p>Defines a number of days before password expiration.</p> <p>Starting at the given period before the expiration date, users receive notification when logging on that their password will soon expire.</p>
<code>maximum_unused_initial_password_lifetime</code>	28	<p>Defines the duration in days that an initial password for a user account is valid.</p> <p>If an initial password has not been used for the first time within the given period of time, the password becomes invalid and the password must be reset.</p> <p>If the value of this parameter is set to 0, no check is performed.</p>
	365	

Parameter	Default Value	Description
<code>maximum_unused_productive_password_lifetime</code>		<p>Defines the duration in days that a user-defined password is valid.</p> <p>If a user-defined password has not been reused within the given period of time, the password becomes invalid and the password must be reset.</p> <p>If the value of this parameter is set to 0, no check is performed.</p>
<code>minimum_password_lifetime</code>	1	<p>Defines the minimum duration in days that a newly entered user-defined password remains valid before the user can change it again.</p> <p>If the value of this parameter is set to 0, no check is performed.</p>

## 7.2 How-To for Configuration of SAML Support

To use SAML on your SAP HANA instance, you have to configure SSL.

The following brief instruction shows the necessary steps to configure SAML support on an SAP HANA database system. For more information about individual SQL commands, see the *SAP HANA Database - SQL Reference Guide* at <http://help.sap.com/hana>.

1. Configure a new SAML identity provider:
  - a. From the X.509 certificate of the SAML identity provider that will be used for signing the SAML assertions, check the subject- and issuer-distinguished names.
  - b. Make sure that the entire certificate chain of the X.509 certificate is in the trust store of your SAP HANA instance.
  - c. Create a new SAML provider for SAP HANA:

```
CREATE SAML PROVIDER SAP_HANA_TEST_PROVIDER WITH
SUBJECT 'CN=SAP AG HANA SSL-Test-Server, OU=SAP HANA, O=SAP AG, L=Walldorf, ST=Germany, C=DE'
ISSUER 'CN=SAP AG HANA Test Root CA, OU=SAP HANA, O=SAP AG, L=Walldorf, ST=Germany, C=DE'
```

2. Enable a user to use SAML:

- For existing users:

```
ALTER USER SAP_HANA_TEST_USER ENABLE SAML
ALTER USER SAP_HANA_TEST_USER ADD IDENTITY 'EXTERNAL_USER_NAME' FOR SAML PROVIDER SAP_HANA_TEST_PROVIDER
```

**Note:**

EXTERNAL\_USER\_NAME has to match the NameID contents in your SAML assertion.

- For new users:

```
CREATE USER SAP_HANA_TEST_USER WITH IDENTITY
'EXTERNAL_USER_NAME' FOR SAML PROVIDER SAP_HANA_TEST_PROVIDER
```

**Note:**

As only SAML credentials are specified here, all other authentication mechanisms (including password) are disabled for this user.

### 3. Connect with a SAML assertion:

Now, you are able to log on using a SAML assertion. On the client side (JDBC/ODBC), in the connection properties, leave the user name empty and put the whole SAML assertion into the password property. Then connect as usual. As an alternative, you may also use the internal SQL connect:

```
CONNECT WITH SAML ASSERTION '<assertion>'
```

### 4. Configure identity provider-based user name mapping:

To allow your identity provider to map users to SAP HANA database users (and providing them in the SPProvidedID attribute of the NameID element in the assertions), you can also configure a user with a wildcard mapping:

```
ALTER USER SAP_HANA_TEST_USER ADD IDENTITY ANY FOR SAML PROVIDER SAP_HANA_TEST_PROVIDER
```

**Note:**

A user can either have a wildcard mapping or an explicit mapping for a particular identity provider, but not both.

### 5. Configure a SAML service provider (SAML audience restrictions):

If the SAML assertions that will be sent to your SAP HANA instance contain audience restrictions, you have to configure the SAML service provider name under which your SAP HANA instance acts as SAML service provider.

```
ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini','SYSTEM')
SET ('authentication','saml_service_provider_name') = 'sp.hana.com' WITH RECONFIGURE;
```