



One Identity Manager 8.1

Native Database Connector User Guide for Connecting SAP HANA Databases

Copyright 2019 One Identity LLC.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of One Identity LLC .

The information in this document is provided in connection with One Identity products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of One Identity LLC products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, ONE IDENTITY ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL ONE IDENTITY BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF ONE IDENTITY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. One Identity makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. One Identity does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

One Identity LLC.
Attn: LEGAL Dept
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our Web site (<http://www.OneIdentity.com>) for regional and international office information.

Patents

One Identity is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <http://www.OneIdentity.com/legal/patents.aspx>.

Trademarks

One Identity and the One Identity logo are trademarks and registered trademarks of One Identity LLC. in the U.S.A. and other countries. For a complete list of One Identity trademarks, please visit our website at www.OneIdentity.com/legal. All other trademarks are the property of their respective owners.

Legend

-  **WARNING:** A WARNING icon indicates a potential for property damage, personal injury, or death.
-  **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.
-  **IMPORTANT, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

Contents

Native database connector for SAP HANA databases	4
Users and permissions for synchronizing	5
Setting up a custom application role for synchronization	9
Setting up the synchronization server	10
Creating a synchronization project	13
How to set up a synchronization project	15
Connecting a system to an SAP HANA database	18
Updating schemas	23
Starting synchronization	24
Analyzing synchronization	25
Post-processing outstanding objects	25
Configuring target system synchronization	26
How to post-process outstanding objects	27
Configuring the provisioning of memberships	28
Error handling	30
Help for the analysis of synchronization issues	30
About us	31
Contacting us	31
Technical support resources	31
Index	32

Native database connector for SAP HANA databases

With this native database connector, you can synchronize external databases with the One Identity Manager database. One Identity Manager supports connecting to SAP HANA databases, amongst others.

The native database connector cannot load any random external database system data configuration. For example, custom data types and columns containing value list are not currently supported.

The native database connection does not provide a project template for setting up synchronization. You must create synchronization configuration components (mappings, workflows, start up configurations ...) manually after the synchronization project has been saved.

In the Synchronization Editor, external database tables and columns are referenced as schema types and schema properties.

To set up synchronization with a database

1. Install and configure a synchronization server and declare the server as Job server in One Identity Manager.
2. Provide One Identity Manager users with the required permissions for setting up synchronization and post-processing of synchronization objects.
3. Create a synchronization project with the Synchronization Editor.

Detailed information about this topic

- [Setting up the synchronization server](#) on page 10
- [Users and permissions for synchronizing](#) on page 5
- [Creating a synchronization project](#) on page 13

Users and permissions for synchronizing

In the synchronization with the database connectors, there are three use cases for mapping synchronization objects in the One Identity Manager data model.

1. Mapping custom target systems
2. Mapping default tables (for example Person, Department)
3. Mapping custom tables

In the case of non-role-based login to One Identity Manager tools, it is sufficient to add one system user in the **DPR_EditRights_Methods** permissions group. For detailed information about system users and permissions groups, see the *One Identity Manager Authorization and Authentication Guide*.

Table 1: Users and Permissions Groups for Non Role-Based Login

Users	Task
One Identity Manager administrators	<ul style="list-style-type: none">• Create customized permissions groups for application roles for role-based login to administration tools in Designer as required.• Create system users and permissions groups for non-role-based login to administration tools in Designer as required.• Enable or disable additional configuration parameters in Designer as required.• Create custom processes in Designer as required.• Create and configures schedules as required.• Create and configure password policies as required.
System users in the DPR_EditRights_Methods permissions group	<ul style="list-style-type: none">• Configure and start synchronization in the Synchronization Editor.• Edit the synchronization's target system types as well as outstanding objects in the Manager.

There are different steps required for role-based login, in order to equip One Identity Manager users with the required permissions for setting up synchronization and post-processing of synchronization objects.

Table 2: User and permissions groups for role-based login: Mapped as custom target system

Users	Task
One Identity Manager administrators	<ul style="list-style-type: none"> • Create customized permissions groups for application roles for role-based login to administration tools in Designer as required. • Create system users and permissions groups for non-role-based login to administration tools in Designer as required. • Enable or disable additional configuration parameters in Designer as required. • Create custom processes in Designer as required. • Create and configures schedules as required. • Create and configure password policies as required.
Target system administrators	<p>Target system administrators must be assigned to the Target systems Administrators application role.</p> <p>Users with this application role:</p> <ul style="list-style-type: none"> • Administrate application roles for individual target systems types. • Specify the target system manager. • Set up other application roles for target system managers if required. • Specify which application roles are conflicting for target system managers • Authorize other employee to be target system administrators. • Do not assume any administrative tasks within the target system.
Target system managers	<p>Target system managers must be assigned to the application role Target systems Custom target systems or a sub application role.</p> <p>Users with this application role:</p> <ul style="list-style-type: none"> • Assume administrative tasks for the target system. • Create, change or delete target system objects, like user accounts or groups. • Edit password policies for the target system. • Prepare groups for adding to the IT Shop.

Users	Task
	<ul style="list-style-type: none"> • Can create employees with an identity that differs from the Primary identity. • Configure synchronization in the Synchronization Editor and defines the mapping for comparing target systems and One Identity Manager. • Edit the synchronization's target system types and outstanding objects. • Authorize other employees within their area of responsibility as target system managers and create child application roles if required.

Table 3: User and permissions groups for role-based login: Mapped as default tables

Users	Task
One Identity Manager administrators	<ul style="list-style-type: none"> • Create customized permissions groups for application roles for role-based login to administration tools in Designer as required. • Create system users and permissions groups for non-role-based login to administration tools in Designer as required. • Enable or disable additional configuration parameters in Designer as required. • Create custom processes in Designer as required. • Create and configures schedules as required. • Create and configure password policies as required.
Custom application role	<p>Users with this application role:</p> <ul style="list-style-type: none"> • Configure and start synchronization in the Synchronization Editor. • Edit the synchronization's target system types as well as outstanding objects in the Manager. <p>This application role gets its write access through a custom permissions group and the permissions group vi_4_SYNCPROJECT_ADMIN.</p>

Table 4: Users and permissions groups for role-based login: Mapped in custom tables

Users	Task
One Identity Manager	<ul style="list-style-type: none"> • Create customized permissions groups for application

Users	Task
administrators	<p>roles for role-based login to administration tools in Designer as required.</p> <ul style="list-style-type: none"> • Create system users and permissions groups for non-role-based login to administration tools in Designer as required. • Enable or disable additional configuration parameters in Designer as required. • Create custom processes in Designer as required. • Create and configures schedules as required. • Create and configure password policies as required.
Application roles for custom tasks	<p>Administrators must be assigned to Custom Administrators.</p> <p>Users with this application role:</p> <ul style="list-style-type: none"> • Administrate custom application roles. • Set up other application roles for managers if required.
Manager for custom tasks	<p>Managers must be assigned to Custom Managers or a subordinate role.</p> <p>Users with this application role:</p> <ul style="list-style-type: none"> • Add custom task in One Identity Manager. • Configure and start synchronization in the Synchronization Editor. • Edit the synchronization's target system types as well as outstanding objects in the Manager. <p>You can use these application roles, for example, to guarantee One Identity Manager users write permissions on custom tables or columns. All application roles that you define here must obtain their write permissions through custom permissions groups.</p> <p>This application role gets its write access through a custom permissions group and the permissions group vi_4_SYNCPROJECT_ADMIN.</p>

To configure synchronization projects and target system synchronization (in the use cases 2 and 3)

1. Set up a custom permissions group with all permissions for configuring synchronization and editing synchronization objects.
2. Assign a custom application role to this permission group.

Detailed information about this topic

- [Setting up a custom application role for synchronization](#)

Setting up a custom application role for synchronization

For role-based login, create a custom application role to guarantee One Identity Manager users the necessary permissions for configuring synchronization and handling outstanding objects. This application role obtains the required permissions by using a custom permissions group.

To set up an application role for synchronization (use case 2):

1. Select the default application role to use to edit the objects you want to synchronize in the Manager.
 - Establish the application role's default permissions group.

If you want to import employee data, for example, select **Identity Management | Employees | Administrators**. The default permissions group of this application role is `vi_4_PERSONADMIN`.
2. Create a new permissions group in the Designer.
 - Set **Only use for role based authentication**.
3. Make the new permissions group dependent on the **vi_4_SYNCPROJECT_ADMIN** permission group.

The `vi_4_SYNCPROJECT_ADMIN` permissions groups must be assigned as the parent permissions group. This means that the new permissions group inherits the properties.
4. Make the new permissions group dependent on the default permission group of the selected default application role.

The default permissions group must be assigned as a subgroup. This means that the new permissions group inherits the properties.
5. Save the changes.
6. Create a new application role in the Manager.
 - a. Assign the selected application role to be the parent application role.
 - b. Assign the new permissions group.
7. Assign employees to this application role.
8. Save the changes.

To set up an application role for synchronization (use case 3):

1. Create a new permissions group for custom tables, which are populated through synchronization, in the Designer.
 - Set **Only use for role based authentication**.
2. Guarantee this permissions group all the required permissions to the custom tables.
3. Create another permissions group for synchronization.
 - Set **Only use for role based authentication**.
4. Make the permissions group for synchronization dependent on the permissions group for custom tables.

The permissions group for custom tables must be assigned as parent permissions group. This means the permissions groups for synchronization inherits its properties.
5. Make the permissions group for synchronization dependent on the **vi_4_SYNCPROJECT_ADMIN** permission group.

The vi_4_SYNCPROJECT_ADMIN permissions groups must be assigned as the parent permissions group. This means the permissions groups for synchronization inherits its properties.
6. Save the changes.
7. Create a new application role in the Manager.
 - a. Assign the application role **Custom | Managers** as the parent application role.
 - b. Assign the permissions group for the synchronization.
8. Assign employees to this application role.
9. Save the changes.

For detailed information about setting up application roles and permissions groups, see the *One Identity Manager Authorization and Authentication Guide*.

Setting up the synchronization server

A server with the following software must be available for setting up synchronization:

- SAP HANA Data Provider for ADO.NET
- One Identity Manager Service
 - Install One Identity Manager components with the installation wizard.
 1. Select **Select installation modules with existing database**.
 2. Select the **Server | Job server** machine role.

For more detailed information about system requirements for installing the One Identity Manager Service, see the *One Identity Manager Installation Guide*.

The synchronization server must be declared as a Job server in One Identity Manager.

Use the Server Installer to install the One Identity Manager Service. The program executes the following steps:

- Setting up a Job server.
- Specifying machine roles and server function for the Job server.
- Remote installation of One Identity Manager Service components corresponding to the machine roles.
- Configuration of One Identity Manager Service.
- Starts the One Identity Manager Service.

NOTE: The program executes remote installation of the One Identity Manager Service. Local installation of the service is not possible with this program. Remote installation is only supported within a domain or a trusted domain.

For remote installation of One Identity Manager Service, you require an administrative workstation on which the One Identity Manager components are installed. For detailed information about installing a workstation, see the *One Identity Manager Installation Guide*.

To install and configure One Identity Manager Service remotely on a server

1. Start the program Server Installer on your administrative workstation.
2. Enter the valid connection credentials for the One Identity Manager database on the **Database connection** page.
3. Specify the server on which you want to install One Identity Manager Service on the **Server properties** page.
 - a. Select a Job server from the **Server** menu.
- OR -
To create a new Job server, click **Add**.
 - b. Enter the following data for the Job server.

Table 5: Job Server Properties

Property	Description
Server	Job server name.
Queue	Name of the queue to handle the process steps. Each One Identity Manager Service within the network must have a unique queue identifier. The process steps are requested by the job queue using exactly this queue name. The queue identifier is entered in the One Identity Manager Service configuration file.
Full server name	Full server name in accordance with DNS syntax. Example: <Name of servers>.<Fully qualified domain name>

NOTE: You can use the **Extended** option to make changes to other properties for the Job server. You can also edit the properties later with Designer.

4. Select **Job server** on the **Machine roles** page.
5. On the **Server functions** page, select **Native database connector**.
6. Check the One Identity Manager Service configuration on the **Service settings** page.

NOTE: The initial service configuration is predefined already. If further changes need to be made to the configuration, you can do this later with the Designer. For detailed information about configuring the service, see the *One Identity Manager Configuration Guide*.

7. To configure remote installations, click **Next**.
8. Confirm the security prompt with **Yes**.
9. Select the directory with the install files on **Select installation source**.
10. Select the file with the private key on the page **Select private key file**.

NOTE: This page is only displayed when the database is encrypted.

11. Enter the service's installation data on the **Service access** page.

Table 6: Installation Data

Data	Description
Computer	<p>Server on which to install and start the service from.</p> <p>To select a server</p> <ul style="list-style-type: none"> • Enter a name for the server. - OR - • Select a entry from the list.
Service account	<p>User account data for the One Identity Manager Service.</p> <p>To enter a user account for the One Identity Manager Service</p> <ul style="list-style-type: none"> • Set the option Local system account. <p>This starts the One Identity Manager Service under the NT AUTHORITY\SYSTEM account.</p> - OR - <ul style="list-style-type: none"> • Enter user account, password and password confirmation.
Installation account	<p>Data for the administrative user account to install the service.</p> <p>To enter an administrative user account for installation</p>

Data	Description
	<ul style="list-style-type: none"> • Enable Advanced. • Enable Current user. This uses the user account of the current user. - OR - • Enter user account, password and password confirmation.

12. Click **Next** to start installing the service.

Installation of the service occurs automatically and may take some time.

13. Click **Finish** on the last page of Server Installer.

NOTE: The service is entered with the name **One Identity Manager Service** in the server service management.

Creating a synchronization project

A synchronization project collects all the information required for synchronizing the One Identity Manager database with a target system. Connection data for target systems, schema types and properties, mapping and synchronization workflows all belong to this.

Make the following information available for setting up a synchronization project for synchronizing with the native database connector.

Table 7: Information Required for Setting up a Synchronization Project

Data	Explanation
Synchronization server	<p>All One Identity Manager Service actions are executed against the target system environment on the synchronization server. Data entries required for synchronization and administration with the One Identity Manager database are processed by the synchronization server.</p> <p>Installed components:</p> <ul style="list-style-type: none"> • One Identity Manager Service (started) <p>The synchronization server must be declared as a Job server in One Identity Manager. The Job server name is required.</p> <p>For more information, see Setting up the synchronization server on page 10.</p>
Remote connection server	<p>To configure synchronization with a target system, One Identity Manager must load the data from the target system. One</p>

Data

Explanation

Identity Manager communicates directly with target system to do this. Sometimes direct access from the workstation on which the Synchronization Editor is installed is not possible, because of the firewall configuration, for example, or because the workstation does not fulfill the necessary hardware and software requirements. If direct access to the workstation is not possible, you can set up a remote connection.

The remote connection server and the workstation must be in the same Active Directory domain.

Remote connection server configuration:

- One Identity Manager Service is started
- RemoteConnectPlugin is installed
- SAP HANA data provider for ADO.NET is installed

The remote connection server must be declared as a Job server in One Identity Manager. The Job server name is required.

TIP: The remote connection server requires the same configuration as the synchronization server (with regard to the installed software and entitlements). Use the synchronization as remote connection server at the same time, by simply installing the RemoteConnectPlugin as well.

For more detailed information about setting up a remote connection, see the *One Identity Manager Target System Synchronization Reference Guide*.

Synchronization workflow

Set the option **Data import** in the synchronization step if synchronization data is imported from a secondary system.

For more detailed information about synchronizing user data with different systems, see the *One Identity Manager Target System Synchronization Reference Guide*.

Base object

You cannot normally specify a base object for synchronizing. In this case, assignment of one base table and the synchronization server is sufficient.

- Select the **Base table** from the menu in which to load the objects. The base table can be used to defined downstream processes for synchronization. For more information about downstream processes, see the *One Identity Manager Target System Synchronization Reference Guide*.
- The **Synchronization servers** menu displays all Job servers for which the **Native database connector**

Data	Explanation
Variable set	server function is activated. If you implement specialized variable sets, ensure that the start up configuration and the base object use the same variable set.

To configure synchronization with the native database connector

1. Create a new synchronization project.
2. Add mappings. Define property mapping rules and object matching rules.
3. Create synchronization workflows.
4. Create a start up configuration.
5. Define the synchronization scope.
6. Specify the base object of the synchronization.
7. Specify the extent of the synchronization log.
8. Run a consistency check.
9. Activate the synchronization project.
10. Save the new synchronization project in the database.

Detailed information about this topic

- [How to set up a synchronization project](#) on page 15

How to set up a synchronization project

There is an wizard to assist you with setting up a synchronization project. This wizard takes you all the steps you need to set up initial synchronization with a target system. Click **Next** once you have entered all the data for a step.

NOTE: The following sequence describes how you configure a synchronization project if Synchronization Editor is both:

- executed In default mode, and
- started from the launchpad

If you execute the project wizard in expert mode or directly from Synchronization Editor, additional configuration settings can be made. Follow the project wizard instructions through these steps.

Prerequisite


- The SAP HANA Data Provider for ADO.NET must be installed on the workstation, which is executing the Launchpad.

To set up a synchronization project

1. Start the Launchpad and log on to the One Identity Manager database.
 - ① **NOTE:** If synchronization is executed by an application server, connect the database through the application server.
2. Select **Native Database Connector** and click on **Run**.

This starts the Synchronization Editor's project wizard.
3. Specify how One Identity Manager can access the target system on the **System access** page.
 - If access is possible from the workstation on which you started Synchronization Editor, you do not need to make any settings.
 - If access is not possible from the workstation on which you started Synchronization Editor, you can set up a remote connection.

Enable the **Connect using remote connection server** option and select the server to be used for the connection under **Job server**.
 - Click **Next** to start the system connection wizard to create a connection to an external database.
4. Select the database system to which you want to connect on the **Select database system** page.
 - Select **SAP HANA**.
5. Configure the system connection.

For more information, see [Connecting a system to an SAP HANA database](#) on page 18.
6. You can save the current configuration as a template on the **Save configuration** page. When you reconnect to a database system of the same type, you can use this configuration as a template.
 - Click  and enter the name and repository of the configuration file.
7. You can save the connection data on the last page of the system connection wizard.
 - Set the **Save connection locally** option to save the connection data. This can be reused when you set up other synchronization projects.
 - Click **Finish**, to end the system connection wizard and return to the project wizard.
8. On the **One Identity Manager Connection** tab, test the data for connecting to the One Identity Manager database. The data is loaded from the connected database. Reenter the password.

NOTE: If you use an unencrypted One Identity Manager database and have not yet saved any synchronization projects to the database, you need to enter all connection data again. This page is not shown if a synchronization project already exists.

9. The wizard loads the target system schema. This may take a few minutes depending on the type of target system access and the size of the target system.
10. Select a project template on the **Select project template** page to use for setting up the synchronization configuration.

NOTE: The native database connector does not provide a default project template for setting up synchronization. If you have created your own project template, you can select it to configure the synchronization project. Otherwise, select **Create blank project**.

11. Enter the general setting for the synchronization project under **General**.

Table 8: General properties of the synchronization project

Property	Description
Display name	Display name for the synchronization project.
Script language	<p>Language in which the scripts for this synchronization project are written.</p> <p>Scripts are implemented at various points in the synchronization configuration. Specify the script language when you set up an empty project.</p> <p>IMPORTANT: You cannot change the script language once the synchronization project has been saved.</p> <p>If you use a project template, the template's script language is used.</p>
Description	Spare text box for additional explanation.

12. To close the project wizard, click **Finish**.
13. Save the synchronization project in the database.

Connecting a system to an SAP HANA database

Table 9: Required information for connecting the system

Data	Explanation
Server	Name of the server on which the database server is installed. The fully qualified server name or the IP address may be given.
User account and password	User account and password used by the native database connector to log in to the external database. Make a user account available with sufficient permissions.
Database	Name of the external database to be synchronized.
Installed SAP HANA provider	Provider that connects to the SAP HANA database.

To configure the connection to an SAP HANA database:

1. Enter the connection parameters on the **Database connection** page. Select the SAP HANA provider and enter all the parameters required by the database connector to create a connection with the selected database system.
 - To enter additional system-specific information about the system connection, click **Advanced**.

The database system connection is tested the moment you click **Next**.

2. Enter a display name and a unique identifier for the database connection on the **Describe the database** page.

Table 10: Name of the database

Property	Description
Display name of database	Display name of the database for display in the One Identity Manager tools.
System identifier	Unique identifier of the database.

- !** **IMPORTANT:** The system identifier of the database must be unique. These identifiers help to differentiate between the databases. To prevent incorrect behavior and loss of data ensure that the system identifiers are unique within the One Identity Manager environment.
- Identifiers may not be defined more than once.
 - Identifiers must not be changed after the connection is saved.

3. You can enter a file on the **Load configuration** page from which the connection configuration can be loaded. This data is used in subsequent steps in the connection wizard and can be modified there.
4. Select the time zone for the time zone data in the database on the page, **Time zone selection**. The time zone is required to convert the time saved in the database into the local time. The local time is displayed in the One Identity Manager tools.
5. You can specify additional connection settings on the **Initializing** page. Write a script in the database syntax to specify number and date formats, language and data sort order, for example. This script is then executed every time you connect the system.
6. You can write a script on the **Initialization** page that is executed each time a connection is established. For example, connection settings for data sort order, language or date format can be defined by the script.
7. On the **Select partial schemas** page, you can reduce the database schema by selecting partial schemas. If the database contains several schema, specify here, which schemas are loaded into the synchronization project.
 - Enable all the schemas to process in the **Partial schemas/owner** list.
8. The database schema is loaded on the **Schema detection** page during which One Identity Manager tries to identify a known schema.
 - If a One Identity Manager schema is detected, the **Fill in system description completely** option is displayed. If you only want allow read-only access to the database, you can deactivate this option.

If the schema is loaded successfully, the next step in the sequence can be carried out.

9. On the **Extend key information** page, specify columns for each table to be used as unique keys for identifying objects.

NOTE:

- This page is only displayed if the schema of the external database there are tables with no identifiable unique keys.
- Tables without unique keys are not used in the synchronization configuration.

Table 11: Defining unique keys

Property	Description
Hide unconfigured tables	Specifies whether table are hidden if no settings have been changed.
Schema	Tables without a unique key.
Column is	Specifies whether the column contains a unique key.

Property	Description
key	
Column group	<p>Button for editing column groups. Create a column group, if a unique key can only be made of a combination of more than one column.</p> <ul style="list-style-type: none"> To create a column group, click Add To edit or remove an existing column group, click Edit or remove

Table 12: Column group properties

Property	Description
Key name	Column group identifier. Permitted characters are letters and underscore. A virtual schema property is formed from the column group with the name <code>vrtColumnGroup<column group></code> .
Columns	Columns included in the column group. Mark all the columns that together make up the unique key.

10. You can enter information about object relations in the **Define data relations** page.

Table 13: Defining column relations

Property	Description
Hide unconfigured tables	Specifies whether table are hidden if no settings have been changed.
Schema	Database schema tables.
Target(s)	<p>Columns to which the reference refers. Enter table and column names in the following syntax: [<code><schema></code>].<code><table name></code>.<code><column name></code>. If a reference points to several column, enter the targets in a comma delimited list. The target columns must be labeled as key columns.</p> <p>TIP: You can copy the column name of a referenced column using the Copy fully qualified column names item in the context menu and add this as a target.</p>
Referential integrity enabled	Specifies whether the referential integrity of the data in the target table has been tested.

11. You can enter additional schema information on the **Complete schema** page.

Table 14: Additional schema information

Property	Description
Hide unconfigured tables	Specifies whether table are hidden if no settings have been changed.
Schema	Tables and schemas of the database schema.
Display value	Column used in the display pattern. <ul style="list-style-type: none"> To use the column in the display pattern, click Add.
Preferred key	Specifies whether the column is to be primarily used for object identification. A preferred key can defined, if a table has more than one unique key. Only columns with the String data type can be selected.
Contains sensitive data	Specifies whether the column contains sensitive data.
Revision counter	Specifies whether the column contains the revision counter. The data in this column form the comparison value for revision filtering.
Sort criteria for hierarchies	Specify whether the column maps the path in an object hierarchy. Synchronization objects are sorted by this order. This makes it possible to resolve object dependencies. Only one column per table can be used as a sort criterion.
Scope reference	Specifies whether the column can be used to form the reference scope.

Table 15: Table properties

Property	Description
Display template	Display pattern with which the objects in Synchronization Editor are displayed. The display pattern is, for example, used in error messages or test result from object matching rules. The display pattern is, for example, used in error messages or in the test results from object matching rules. Enter a display table for each display pattern. <ul style="list-style-type: none"> To use a column in the display pattern, select a column and click Add.

- You can specify special operations for changing data in the external database on the **Define data operations** page. This is only required, if the default operations INSERT, UPDATE and DELETE cannot be used in the external database system.

⚠ WARNING: A good knowledge of programming is required to implement data operations. Errors in this implementation can lead to loss of data.

To define a data operation

- a. Select a table and mark the operation you want to define.
- b. Select a strategy.
- c. Enter the data operation you want to run in the **Settings** input field.

Table 16: Defining data operations

Property	Description
Hide unconfigured tables	Specifies whether table are hidden if no settings have been changed.
Table/operation	Tables for which the data operations are to be defined.
Strategy	Strategy with which the data operation is created and run. A simple procedure can be called for a data operation or a script can be executed. Select the strategy you want use to define the data operation.

Table 17: Strategies for running data operations

Strategy	Description
Pattern-based	Simple procedure call that runs the operation.
Script-based	Script that runs a complex data operation. You can use custom code snippets in the script. The code snippets must contain a keyword element with the DML keyword. For more detailed information about support for writing scripts, see the <i>One Identity Manager Target System Synchronization Reference Guide</i> .

- Click  to delete a data operation.

Required columns	List of required key columns in a script-based data operation. The columns must be entered if they are
------------------	--

Property	Description
	not part of the display name.
Settings...	<p>Define the data operation that is to be run when objects are added, updated, or deleted. Enter the procedure call or create a script depending on the selected strategy.</p> <p>Example of a pattern-based data operation:</p> <pre>exec CreateUser ('%Uid%', '%FirstName%', '%LastName%')</pre> <p>It has an advanced edit mode which provides additional actions. For detailed information about support for creating scripts, see the <i>One Identity Manager Target System Synchronization Reference Guide</i>.</p>

Related Topics

- [How to set up a synchronization project](#) on page 15

Updating schemas

All the schema data (schema types and schema properties) of the target system schema and the One Identity Manager schema are available when you are editing a synchronization project. Only a part of this data is really needed for configuring synchronization. If a synchronization project is finished, the schema is compressed to remove unnecessary data from the synchronization project. This can speed up loading the synchronization project. Deleted schema data can be added to the synchronization configuration again at a later point.

If the target system schema or the One Identity Manager schema has changed, these changes must also be added to the synchronization configuration. Then the changes can be added to the schema property mapping.

To include schema data that have been deleted through compressing and schema modifications in the synchronization project, update each schema in the synchronization project. This may be necessary if:

- A schema was changed by:
 - Changes to a target system schema
 - Customizations to the One Identity Manager schema
 - A One Identity Manager update migration
- A schema in the synchronization project was shrunk by:

- enabling the synchronization project
- saving the synchronization project for the first time
- compressing a schema

To update a system connection schema

1. Select the category **Configuration | Target systems**.
- OR -
Select the category **Configuration | One Identity Manager connection**.
2. Select the view **General** and click **Update schema**.
3. Confirm the security prompt with **Yes**.
This reloads the schema data.

To edit a mapping

1. Select the category **Mappings**.
2. Select a mapping in the navigation view.
Opens the Mapping Editor. For more detailed information about mappings, see the *One Identity Manager Target System Synchronization Reference Guide*.

NOTE: The synchronization is deactivated if the schema of an activated synchronization project is updated. Reactivate the synchronization project to synchronize.

Starting synchronization

Synchronization is started using scheduled process plans. A scheduled process plan is added once a start up configuration is assigned to a schedule. Use schedules to define executing times for synchronization.

NOTE: Synchronization can only be started if the synchronization project is enabled.

To execute synchronization regularly, configure and activate the a schedule. You can also start synchronization manually if there is no active schedule.


IMPORTANT: As long as synchronization is running, you must not start another synchronization for the same target system. This applies especially, if the same synchronization objects would be processed.

- If another synchronization is started with the same start up configuration, this process is stop and is assigned the **Frozen** execution status. An error message is written to the One Identity Manager Service log file.
- If another synchronization is started with another start up configuration, that addresses same target system, it may lead to synchronization error or loss of data. Specify One Identity Manager behavior in this case, in the start up configuration. Group start up configurations with the same start up behavior.

Analyzing synchronization

Synchronization results are summarized in the synchronization log. You can specify the extent of the synchronization log for each system connection individually. One Identity Manager provides several reports in which the synchronization results are organized under different criteria.

To display a synchronization log

1. Open the synchronization project in the Synchronization Editor.
2. Select **Logs**.
3. Click  in the navigation view toolbar.
Logs for all completed synchronization runs are displayed in the navigation view.
4. Select a log by double-clicking on it.
An analysis of the synchronization is shown as a report. You can save the report.

Synchronization logs are stored for a fixed length of time.

To modify the retention period for synchronization logs

- In Designer, enable the **DPR | Journal | LifeTime** configuration parameter and enter the maximum retention period.

Post-processing outstanding objects

Objects, which do not exist in the target system, can be marked as outstanding in One Identity Manager by synchronizing. This prevents objects being deleted because of an incorrect data situation or an incorrect synchronization configuration.

Outstanding objects

- Cannot be edited in One Identity Manager.
- Are ignored by subsequent synchronization.
- Are ignored by inheritance calculations.

This means, all memberships and assignments remain intact until the outstanding objects have been processed.

Start target system synchronization to do this.

To allow post-processing of outstanding objects

- Configure target system synchronization.
For more information, see [Configuring target system synchronization](#) on page 26.


Related Topics

- [How to post-process outstanding objects](#) on page 27
- [Users and permissions for synchronizing](#) on page 5

Configuring target system synchronization


Create a target system for post-processing outstanding objects. Assign tables you want to be populated by synchronization, to this target system type. Specify the tables for which outstanding objects can be published in the target system during post-processing. Define a process for publishing the objects.

To create a target system type

1. In the Manager, select the category **Data Synchronization | Basic configuration data | Target system types**.
2. Click  in the result list toolbar.
3. Edit the target system type master data.
4. Save the changes.

Enter the following data for a target system type.

Table 18: Master Data for a Target System Type

Property	Description
Target system type	Target system type description.
Description	Spare text box for additional explanation.
Display name	Name of the target system type as displayed in One Identity Manager tools.
Cross-boundary inheritance	Specifies whether user accounts can be assigned to groups if they belong to different custom target systems.  NOTE: If this option is not set, the target system type is used to group the target systems.
Show in compliance rule wizard	Specifies whether the target system type for compliance rule wizard can be selected when rule conditions are being set up.
Text snippet	Text snippets used for linking text in the compliance rule wizard.

To add tables to the target system synchronization

1. In the result list, select the target system type.
2. Select **Assign synchronization tables**.

3. Assign tables whose outstanding objects you want to handle in **Add assignments**.
4. Save the changes.
5. Select **Configure tables for publishing**.
6. Select tables whose outstanding objects can be published in the target system and set **Publishable**.
7. Save the changes.

NOTE: The connector must have write access to the target system in order to publish outstanding objects that are being post-processed. That means, the option **Connection is read only** must not be set for the target system connection.

To publish outstanding objects

- For each table for which you want to publish outstanding objects, create a process, which is triggered by the event `HandleOutstanding` and which executes the provisioning of the objects. Use the `AdHocProjection` process function of the `ProjectorComponent` process component. For more detailed information about defining processes, see *One Identity Manager Configuration Guide*.

How to post-process outstanding objects

To post-process outstanding objects

1. In Manager, select **Data synchronization | Target system synchronization: <target system type>**.

All tables assigned to the target system type are displayed in the navigation view.

2. Select the table whose outstanding objects you want to edit in the navigation view. All objects marked as outstanding are shown on the form.


TIP:



To display object properties of an outstanding object

- a. Select the object on the target system synchronization form.
- b. Open the context menu and click **Show object**.

3. Select the objects you want to rework. Multi-select is possible.
4. Click one of the following icons in the form toolbar to execute the respective method.

Table 19: Methods for handling outstanding objects

Icon	Method	Description
	Delete	The object is immediately deleted in the One Identity Manager database. Deferred deletion is not taken into account. The

Icon	Method	Description
		Outstanding label is removed for the object. Indirect memberships cannot be deleted.
	Publish	The object is added in the target system. The Outstanding label is removed for the object. The method triggers the <code>HandleOutstanding</code> event. This runs a target system specific process that triggers the provisioning process for the object. Prerequisites: <ul style="list-style-type: none"> • The table containing the object can be published. • The target system connector has write access to the target system. • A custom process is set up for provisioning the object.
	Reset	The Outstanding label is removed for the object.

5. Confirm the security prompt with **Yes**.

NOTE: By default, the selected objects are processed in parallel, which speeds up execution of the selected method. If an error occurs during processing, the action is stopped and all changes are discarded.

Bulk processing of objects must be disabled if errors are to be localized, which means the objects are processed sequentially. Failed objects are named in the error message. All changes that were made up until the error occurred are saved.

To disable bulk processing

- Deactivate  in the form toolbar.

Related Topics

- [Configuring target system synchronization](#) on page 26
- [Users and permissions for synchronizing](#) on page 5

Configuring the provisioning of memberships

Memberships, for example, user accounts, are saved in assignment tables in the One Identity Manager database. During provisioning of modified memberships, changes made in the target system will probably be overwritten. This behavior can occur under the following conditions:

- Memberships are saved in the target system as an object property in list form (Example: List of user accounts in the `Members` property of an Active Directory group).
- Memberships can be modified in either of the connected systems.
- A provisioning workflow and provisioning processes are set up.

If a membership in One Identity Manager changes, the complete list of members is transferred to the target system by default. Memberships, previously added to the target system are removed by this; previously deleted memberships are added again.

To prevent this, provisioning can be configured such that only the modified membership is provisioned in the target system. The corresponding behavior is configured separately for each assignment table.

To allow separate provisioning of memberships

1. In the Manager, select the category **Data Synchronization | Basic configuration data | Target system types**.
2. Select **Configure tables for publishing**.
3. Select the assignment tables for which you want to allow separate provisioning. Multi-select is possible.
 - This option can only be enabled for assignment tables that have a base table with `XDateSubItem` or `CCC_XDateSubItem` column.
 - Assignment tables that are grouped together in a virtual schema property in the mapping must be marked identically (for example, `ADSAccountInADSGroup`, `ADSGroupInADSGroup` and `ADSMachineInADSGroup`).
4. Click **Enable merging**.
5. Save the changes.

For each assignment table labeled like this, the changes made in One Identity Manager are saved in a separate table. During modification provisioning, the members list in the target system is compared to the entries in this table. This means that only modified memberships are provisioned and the members list does not get entirely overwritten.

NOTE: The complete members list is updated by synchronization. During this process, objects with changes but incomplete provisioning are not handled. These objects are logged in the synchronization log.

For detailed information about provisioning memberships, see the *One Identity Manager Target System Synchronization Reference Guide*.

Error handling

For detailed information about correcting errors during synchronization of object hierarchies, see the One Identity Manager Target System Synchronization Reference Guide.

Help for the analysis of synchronization issues

You can generate a report for analyzing problems which occur during synchronization, for example, insufficient performance. The report contains information such as:

- Consistency check results
- Revision filter settings
- Scope applied
- Analysis of the synchronization buffer
- Object access times in the One Identity Manager database and in the target system

To generate a synchronization analysis report

1. Select the menu **Help | Generate synchronization analysis report** and answer the security prompt with **Yes**.
The report may take a few minutes to generate. It is displayed in a separate window.
2. Print the report or save it in one of the available output formats.

One Identity solutions eliminate the complexities and time-consuming processes often required to govern identities, manage privileged accounts and control access. Our solutions enhance business agility while addressing your IAM challenges with on-premises, cloud and hybrid environments.

Contacting us

For sales or other inquiries, visit <https://www.oneidentity.com/company/contact-us.aspx> or call +1-800-306-9329.

Technical support resources

Technical support is available to One Identity customers with a valid maintenance contract and customers who have trial versions. You can access the Support Portal at <https://support.oneidentity.com/>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to-videos at www.YouTube.com/OneIdentity
- Engage in community discussions
- Chat with support engineers online
- View services to assist you with your product

A

- application role 5
- application role for synchronization 9

B

- base object 13

D

- database connector
 - native 4

J

- Job server
 - process 10

M

- membership
 - change provisioning 28

O

- object
 - delete immediately 27
 - outstanding 25, 27
 - publish 27
- outstanding object 25

P

- provisioning
 - member list 28

R

- remote connection server 13

S

- SAP HANA provider 4
- schema
 - changes 23
 - compress 23
 - update 23
- synchronization
 - start 24
- synchronization analysis report 30
- synchronization configuration 13, 15
- synchronization log 25
- synchronization server 13
 - configuring 10
 - install 10
 - Job server 10

T

- target system synchronization
 - assign tables 26
- target system type 26

V

- variable set 13

W

workflow 13