



One Identity Manager 8.0.5

Native Database Connector User
Guide for Connecting MySQL
Databases

Copyright 2020 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 MySQL Databases	4
Users and Permissions for Synchronizing	5
Setting Up a Custom Application Role for Synchronization	8
Setting Up the Synchronization Server	10
Creating a Synchronization Project	13
How to Set up a Synchronization Project	15
Connecting a System to a MySQL Database	17
Updating Schemas	23
Starting Synchronization	24
Analyzing Synchronization	24
Post-Processing Outstanding Objects	25
Configuring Target System Synchronization	25
How to Post-Process Outstanding Objects	27
Configuring Memberships Provisioning	28
Troubleshooting	30
Help for Analyzing Synchronization Issues	30
About us	31
Contacting us	31
Technical support resources	31
Index	32

Native Database Connector for MySQL Databases

With this native database connector, you can synchronize external databases with the One Identity Manager database. One Identity Manager supports connecting to MySQL 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 process, there are three use cases for mapping synchronization objects in the One Identity Manager data model with database connectors.

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

In the case of non role-based log in on One Identity Manager tools, it is sufficient to add one system user in the permissions group "DPR_EditRights_Methods". For more detailed information about system users and permissions groups, see the One Identity Manager Configuration Guide.

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

User	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 the Designer, as required.• Create system users and permissions groups for non-role based login to administration tools, as required.• Enable or disable additional configuration parameters in the Designer, as required.• Create custom processes in the Designer, as required.• Create and configures schedules, as required.• Create and configure password policies, as required.
System users in the permissions group "DPR_EditRights_Methods"	<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

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

User	Task
administrators	<p>roles for role-based login to administration tools in the Designer, as required.</p> <ul style="list-style-type: none"> • Create system users and permissions groups for non-role based login to administration tools, as required. • Enable or disable additional configuration parameters in the Designer, as required. • Create custom processes in the 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 application role Target system Administrators.</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. • 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

User	Task
	<p>outstanding objects.</p> <ul style="list-style-type: none"> • 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

User	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 the Designer, as required. • Create system users and permissions groups for non-role based login to administration tools, as required. • Enable or disable additional configuration parameters in the Designer, as required. • Create custom processes in the 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

User	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 the Designer, as required. • Create system users and permissions groups for non-role based login to administration tools, as required. • Enable or disable additional configuration parameters in the Designer, as required. • Create custom processes in the Designer, as required.

User	Task
	<ul style="list-style-type: none"> • Create and configures schedules, as required. • Create and configure password policies, as required.
Application roles for custom tasks	<p>Administrators must be assigned to the application role 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 the application role Custom Managers or a subordinate role.</p> <p>Users with this application role:</p> <ul style="list-style-type: none"> • Add custom task in the 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>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](#) on page 8

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 (use case 2)

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

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

The permissions "vi_4_SYNCPROJECT_ADMIN" must be assigned as 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 newly created permissions group to it
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 though synchronization, in the Designer.
 - Set the option **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 the option **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 permissions group "vi_4_SYNCPROJECT_ADMIN".

The permissions "vi_4_SYNCPROJECT_ADMIN" must be assigned as 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 parent application role.
 - b. Assign the permissions group for synchronization.
8. Assign employees to this application role.
9. Save the changes.

For more detailed information about setting up application roles, see the One Identity Manager Application Roles Administration Guide. For more detailed information about permissions groups, see the One Identity Manager Configuration Guide.

Setting Up the Synchronization Server

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

- One Identity Manager Service
 - Install One Identity Manager components with the installation wizard.
 1. Select the option **Select installation modules with existing database**.
 2. Select the machine role **Server | Job server**.

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

- ADO.NET provider for MySQL
 - Copy the provider DLL to the One Identity Manager Service's installation directory.

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. This 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.
- Configures the 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.

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

1. Start the program Server Installer on your administrative workstation.
2. Enter valid data for connecting to One Identity Manager on the **Database connection** page and click **Next**.
3. Specify on which server you want to install the One Identity Manager Service on the **Server properties** page.
 - a. Select a job server in the **Server** menu.
- OR -
Click **Add** to add a new job server.
 - b. Enter the following data for the Job server.

Table 5: Job Servers Properties

Property	Description
Server	Name of the Job servers.
Queue	Name of 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 name of the server in DNS syntax. Example: <name of server>.<fully qualified domain name>

NOTE: Use the **Advanced** option to edit other Job server properties. You can use the Designer to change properties at a later date.

4. Specify which job server roles to include in One Identity Manager on the **Machine role** page. Installation packages to be installed on the Job server are found depending on the selected machine role.
Select at least the following roles:
 - Job Server
5. Specify the server's functions in One Identity Manager on the **Server functions** page. One Identity Manager processes are handled depending on the server function. The server's functions depend on which machine roles you have selected. You can limit the server's functionality further here.
Select at least one of the following server functions:
 - Native database connector
6. Check the One Identity Manager Service configuration on the **Service settings** page.

NOTE: The initial service configuration is already predefined. If further changes need to be made to the configuration, you can do this later with the Designer. For more detailed information about configuring the service, see 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 the **Select installation source** page.
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 the server name. - OR - • Select a entry from the list.
Service account	<p>One Identity Manager Service user account data.</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. This starts the One Identity Manager Service under the account "NT AUTHORITY\SYSTEM". - OR - • Enter user account, password and password confirmation. The One Identity Manager Service farm's server farm account must be used as user account for SharePoint.
Installation account	<p>Data for the administrative user account to install the service.</p> <p>To enter an administrative user account for installation</p> <ul style="list-style-type: none"> • Enable Advanced. • Enable the option 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 the Server Installer.

NOTE: The is entered with the name "One Identity Manager Service" in the server's service administration.

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 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.</p> <p>The remote connection server and the workstation must be in the same Active Directory domain.</p> <p>Remote connection server configuration:</p>

Data	Explanation
	<ul style="list-style-type: none"> • One Identity Manager Service is started • RemoteConnectPlugin is installed • ADO.NET provider for MySQL <ul style="list-style-type: none"> • Copy the provider DLL to the One Identity Manager Service's installation directory. <p>The remote connection server must be declared as a Job server in One Identity Manager. The Job server name is required.</p> <p>TIP: The remote connection server requires the same configuration (with respect to the installed software) as the synchronization server. Use the synchronization as remote connection server at the same time, by simply installing the RemoteConnectPlugin as well.</p> <p>For more detailed information about setting up a remote connection, see the One Identity Manager Target System Synchronization Reference Guide.</p>
Synchronization workflow	<p>Set the option Data import in the synchronization step if synchronization data is imported from a secondary system.</p> <p>For more detailed information about synchronizing user data with different systems, see the <i>One Identity Manager Target System Synchronization Reference Guide</i>.</p>
Base object	<p>You cannot normally specify a base object for synchronizing with database connectors. In this case, assignment of one base table and the synchronization server is sufficient.</p> <ul style="list-style-type: none"> • 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 <i>.One Identity Manager Target System Synchronization Reference Guide</i> • All Job servers, which have the server function "native connector" enabled are displayed in the Synchronization servers menu.
Variable set	<p>If you implement specialized variable sets, ensure that the start up configuration and the base object use the same variable set.</p>

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 the Synchronization Editor is both:

- In default mode
- Started from the launchpad

Additional settings can be made if the project wizard is run in expert mode or is started directly from the Synchronization Editor. Follow the project wizard instructions through these steps.

To set up a synchronization project

1. Start the Launchpad and log on to the One Identity Manager database.

NOTE: If synchronization is executed by an application server, connect the database through the application server.

2. Select **Native Database Connector**. Click **Run**.


This starts the Synchronization Editor's project wizard.

3. Specify how the One Identity Manager can access the target system on the **System access** page.

- If you have access from the workstation from which you started the Synchronization Editor, do not set anything.
- If you do not have access from the workstation from which you started the Synchronization Editor, you can set up a remote connection.

In this case, set the option **Connect using remote connection server** and select, under **Job server**, the server you want to use for the connection.

- Click **Next** to start the system connection wizard to create a connection to an external database.
4. Click **Next** on the start page of system connection wizard.
 5. Select the database system to which you want to connect on the **Select database system** page.
 - Select **MySQL**.
 6. Configure the system connection.

For more information, see [Connecting a System to a MySQL Database](#) on page 17.
 7. 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.
 8. You can save the connection data on the last page of the system connection wizard.
 - Set the option **Save connection locally** 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.
 9. Verify the One Identity Manager database connection data on the **One Identity Manager connection** page. The data is loaded from the connected database. Reenter the password.
 - ① **NOTE:** Reenter all the connection data if you are not working with an encrypted One Identity Manager database and no synchronization project has been saved yet in the database. This page is not shown if a synchronization project already exists.
 10. 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.
 11. Select a project template on the **Select project template** page to use for setting up the synchronization configuration.
 - ① **NOTE:** The native database connection 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**.

- Enter the general setting for the synchronization project on the **General** page.

Table 8: General Synchronization Project Properties

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>i IMPORTANT: The script language cannot be changed after 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.

- Click **Finish** to complete the project wizard.
- Save the synchronization project in the database.

Connecting a System to a MySQL Database

Table 9: Information Required 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.

To configure the connection to a MySQL database

- Enter the connection parameters on the **Database connection** page. 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: Naming the Database

Property	Database display name
Database display name	Database display name used in the One Identity Manager tools.
System identifier	Unique identifier for the database.

- i** **IMPORTANT:** The system identifier must describe the database uniquely. 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 may not be changed after the connection has been 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. 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.

7. 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 key	Specifies whether the column contains a unique key.
Column group	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. <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 column is formed from the column group with the name "vrtColumnGroup<column group>".
Columns	Columns included in the column group. Mark all the column, which together make up the unique key.

8. 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 pointed to by the reference. 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 enter the column name of a reference column with the context menu item Copy fully qualified column names and enter the target.</p>
Referential integrity enabled	Specifies whether referential integrity of the target table data is ensured.

9. 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 columns in the database.
Display value	<p>Columns used in the display pattern.</p> <ul style="list-style-type: none"> To use a display pattern, click Add.
Preferred key	Specifies whether the column is primarily used for object identification. A preferred key can defined, if a table has more than one unique key. Only columns with the data type "String" can be selected.
Contains sensitive data	Specifies whether the column contains sensitive data.
Revision	Specifies whether columns have a revision counter. The data in this

Property	Description
counter	column form the comparison value for revision filtering.
Hierarchy sort criterion	Specifies whether the column maps the path in an object hierarchy. Synchronization objects are sorted by this order. This make 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 for setting the reference scope.

Table 15: Table Properties

Property	Description
Display pattern	<p>Pattern for displaying objects in the Synchronization Editor. 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.</p> <ul style="list-style-type: none"> To use a column in a display pattern, select a column and click Add.

10. 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: Well-founded programming knowledge is required for setting up data operations. Errors at this stage can lead to loss of data.

To define a data operation

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

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 data operation are being defined.

Property	Description
Strategy	Strategy for creating the data operation and executing it. 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 Executing Data Operations

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

- To delete a data operation, click .

Options	<p>Define the data operation to be executed when objects are inserted, updated or deleted. Enter the procedure call or create a script depending on the selected strategy.</p> <p>Example of pattern based data:</p> <pre>exec CreateUser ('%Uid%', '%FirstName%', '%LastName%')</pre> <p>It has an advanced edit mode which provides additional actions. For more detailed information about support for creating scripts, see the One Identity Manager Target System Synchronization Reference Guide.</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:
 - Activating the synchronization project
 - Synchronization project initial save
 - Compressing a schema

To update a system connection schema

1. Select the category **Configuration | Target system**.
- 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 editing mappings, see 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.

- The moment another synchronization is started with the same start up configuration, the running synchronization process is stopped and given the status, "Frozen". 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. Select the category **Logs**.
2. Click  in the navigation view toolbar.
Logs for all completed synchronization runs are displayed in the navigation view.
3. 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 the Designer, set the "DPR\Journal\LifeTime" configuration parameter and enter the maximum retention time.

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.

Objects marked as outstanding:

- Cannot be edited in One Identity Manager.
- Are ignored by subsequent synchronization.
- Must be post-processed separately in One Identity Manager.

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 25.


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. Start the Manager.
2. Select the category **Data Synchronization | Basic configuration data | Target system types**.
3. Click  in the result list toolbar.
4. Edit the target system type master data.
5. 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. i 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. Select the category **Data Synchronization | Basic configuration data | Target system types**.
2. Select the target system type in the result list.
3. Select **Assign synchronization tables** in the task view.
4. Assign tables whose outstanding objects you want to handle in **Add assignments**.
5. Save the changes.
6. Select **Configure tables for publishing**.
7. Select tables whose outstanding objects can be published in the target system and set the option **Publishable**.
8. Save the changes.

i **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

- Create a process for each table with outstanding objects you want to publish. The process is triggered by the event "HandleOutstanding" and provisions the objects. Use the process function "AdHocProjection" of the process component "ProjectorComponent". 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. Start the Manager.
2. Select the category **Data synchronization | Target system synchronization: <target system type>**.

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




3. 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**.
4. Select the objects you want to rework. Multi-select is possible.
 5. 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. Deferred deletion is not taken into account. The "outstanding" label is removed from the object. Indirect memberships cannot be deleted.
	Publish	The object is added in the target system. The "outstanding" label is removed from the object. The method triggers the event "HandleOutstanding". 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 from the object.

6. 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 25
- [Users and Permissions for Synchronizing](#) on page 5

Configuring Memberships Provisioning

Memberships, for example, user accounts in , 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 property Members 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. Start the Manager.
2. Select the category **Data Synchronization | Basic configuration data | Target system types**.
3. Select **Configure tables for publishing**.
4. Select the assignment tables for which you want to allow separate provisioning. Multi-select is possible.

- The option can only be set for assignment tables whose base table has a XDateSubItem or a CCC_XDateSubItem.
- Assignment tables, which are grouped together in a virtual schema property in the mapping, must be labeled identically (For example ADSAccountInADSGroup, ADSGroupInADSGroup and ADSMachineInADSGroup).

5. Click **Enable merging**.

6. Save the changes.

For each assignment table labeled like this, the changes made in the 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 more detailed information about provisioning memberships, see the One Identity Manager Target System Synchronization Reference Guide.

Troubleshooting

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

Help for Analyzing 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. Open the synchronization project in the Synchronization Editor.
2. 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.
3. 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, such as licensing, support, and renewals, 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

B

base object 13

D

database connector
native 4

J

Job server
edit 10

M

membership
modify provisioning 28

O

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

P

provisioning
members list 28

R

remote connection server 13

S

schema
changes 23
shrink 23
update 23
synchronization
start 24
synchronization analysis report 30
synchronization configuration 13, 15
synchronization log 24
synchronization server 13
configure 10
install 10
Job server 10

T

target system synchronization
table to assign 25
target system type 25

V

variable set 13

W

workflow 13