



One Identity Manager 9.2.1

Generic Database Connector User
Guide for Connecting DB2 (LUW)
Databases

Copyright 2024 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 website (<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/trademark-information.aspx. All other trademarks are the property of their respective owners.

Legend

 **WARNING:** A WARNING icon highlights a potential risk of bodily injury or property damage, for which industry-standard safety precautions are advised. This icon is often associated with electrical hazards related to hardware.

 **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

One Identity Manager Generic Database Connector User Guide for Connecting DB2 (LUW) Databases
Updated - 16 October 2024, 14:51

For the most recent documents and product information, see [Online product documentation](#).

Contents

Generic database connector for DB2 (LUW) databases	4
Users and permissions for synchronizing	5
Setting up custom application roles for synchronization	9
Setting up the synchronization server	10
Creating a synchronization project	13
Creating a synchronization project	16
Connecting a system to a DB2 (LUW) database	18
Updating schemas	25
Starting synchronization	26
Analyzing synchronization	27
Post-processing outstanding objects	27
Configuring target system synchronization	28
Post-processing outstanding objects	30
Configuring the provisioning of memberships	31
Configuring single object synchronization	33
Ignoring data error in synchronization	34
Troubleshooting	36
Help for analyzing synchronization issues	36
About us	37
Contacting us	37
Technical support resources	37
Index	38

Generic database connector for DB2 (LUW) databases

With this generic database connector, you can synchronize external databases with the One Identity Manager database. One Identity Manager supports connecting to DB2 (LUW) databases, amongst others. Only databases for the operating systems Linux, UNIX, and Windows can be synchronized.

The generic 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 generic database connection does not provide a project template for setting up synchronization. You must create synchronization configuration components (such as mappings, workflows or 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 a Job server in One Identity Manager.
2. Provide One Identity Manager users with the required permissions for setting up synchronization and post-processing 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 in default tables (such as Person, Department)
3. Mapping custom tables

In the case of One Identity Manager tools non role-based login, it is sufficient to add a system user in the **DPR_EditRights_Methods** and **QBM_LaunchPad** permissions groups. For more 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

User	Tasks
One Identity Manager administrators	<p>One Identity Manager administrator and administrative system users Administrative system users are not added to application roles.</p> <p>One Identity Manager administrators:</p> <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 in the Designer as required.• Enable or disable additional configuration parameters in the Designer as required.• Create custom processes in the Designer as required.• Create and configure schedules 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.
System users in the QBM_LaunchPad permissions group	<ul style="list-style-type: none">• Working with the Launchpad.

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	Tasks
One Identity Manager administrators	<p>One Identity Manager administrator and administrative system users Administrative system users are not added to application roles.</p> <p>One Identity Manager administrators:</p> <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 in the Designer as required. • Enable or disable additional configuration parameters in the Designer as required. • Create custom processes in the Designer as required. • Create and configure schedules 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"> • Administer application roles for individual target system types. • Specify the target system manager. • Set up other application roles for target system managers if required. • Specify which application roles for target system managers are mutually exclusive. • Authorize other identities 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 Target systems Custom target systems application role or a child application role.</p> <p>Users with this application role:</p> <ul style="list-style-type: none"> • Assume administrative tasks for the target system.

User	Tasks
	<ul style="list-style-type: none"> • Create, change, or delete target system objects. • Edit password policies for the target system. • Prepare groups to add to the IT Shop. • Can add identities that do not have the Primary identity identity type. • Configure synchronization in the Synchronization Editor and define the mapping for comparing target systems and One Identity Manager. • Edit the synchronization's target system types and outstanding objects. • Authorize other identities 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: Default table mapping

User	Tasks
One Identity Manager administrators	<p>One Identity Manager administrator and administrative system users Administrative system users are not added to application roles.</p> <p>One Identity Manager administrators:</p> <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 in the Designer as required. • Enable or disable additional configuration parameters in the Designer as required. • Create custom processes in the Designer as required. • Create and configure schedules 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>The application role gets its permissions through a custom permissions group and the vi_4_SYNCPROJECT_ADMIN permissions group.</p>

Table 4: Users and permissions groups for role-based login: Custom table mapping

User	Tasks
One Identity Manager administrators	<p>One Identity Manager administrator and administrative system users Administrative system users are not added to application roles.</p> <p>One Identity Manager administrators:</p> <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 in the Designer as required. • Enable or disable additional configuration parameters in the Designer as required. • Create custom processes in the Designer as required. • Create and configure schedules as required.
Application roles for custom tasks	<p>Administrators must be assigned to the Custom Administrators application role.</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 Custom Managers application role or a child 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 user permissions on custom tables or columns. All application roles that you define here must obtain their permissions through custom permissions groups.</p> <p>The application role gets its permissions through a custom permissions group and the vi_4_SYNCPROJECT_ADMIN permissions group.</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 permissions group.

Detailed information about this topic

- [Setting up custom application roles for synchronization](#) on page 9

Setting up custom application roles 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. In the Manager, select the default application role to use to edit the objects you want to synchronize.
 - Establish the application role's default permissions group.

If you want to import identity main data, for example, select the **Identity Management | Identities | Administrators** application role. The default permissions group of this application role is **vi_4_PERSONADMIN**.

2. In the Designer, create a new permissions group .
 - Set the **Only use for role based authentication** option.
3. Make the new permissions group dependent on the **vi_4_SYNCPROJECT_ADMIN** permissions group.

Then 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 permissions group of the selected default application role.

Then the default permissions groups must be assigned as the parent permissions group. This means that the new permissions group inherits the properties.

5. Save the changes.
6. In the Manager, create a new application role.

- a. Assign the selected application role to be the parent application role.
- b. Assign the newly created permissions group.
7. Assign identities to this application role.
8. Save the changes.

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

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

Then the permissions group for custom tables must be assigned as the 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** permissions group.

Then 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. In the Manager, create a new application role.
 - a. Assign the **Custom | Managers** application role as the parent application role.
 - b. Assign the permissions group for the synchronization.
8. Assign identities to this application role.
9. Save the changes.

For more 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:

- ADO.NET provider for DB2 (LUW)
- 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 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.

To set up a Job server, perform the following steps.

1. Create a Job server and install and configure the One Identity Manager Service. Use the One Identity Manager Service to install the Server Installer. The program runs the following steps:
 - Sets up a Job server.
 - Specifies machine roles and server function for the Job server.
 - Installs One Identity Manager Service components corresponding to the machine roles.
 - Configures the One Identity Manager Service.
 - Starts the One Identity Manager Service.

Use the Server Installer to install the One Identity Manager Service locally or remotely.

To remotely install the One Identity Manager Service, provide an administrative workstation on which the One Identity Manager components are installed. Ensure that the One Identity Manager components are installed on the server before installing locally. For more information about installing One Identity Manager components, see the *One Identity Manager Installation Guide*.

2. If you are working with an encrypted One Identity Manager database, declare the database key in the One Identity Manager Service. For more information about working with an encrypted One Identity Manager database, see the *One Identity Manager Installation Guide*.
3. To generate processes for the Job server, you need the provider, connection parameters and the authentication data. By default, this information is determined from the database connection data. If the Job server runs through an application server, you must configure extra connection data in the Designer. For more information about connection data, see the *One Identity Manager Configuration Guide*.

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

1. Start the Server Installer program.

NOTE: To install remotely, start the Server Installer program on your administrative workstation. To install locally, start the program on the server.

2. On the **Database connection** page, enter the valid connection credentials for the One Identity Manager database.

You can connect via the application server or directly to connect to the database.

3. On the **Server properties** page, specify the server on which you want to install the One Identity Manager Service.

- 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.

- **Server:** Name of the Job server.
- **Queue:** Name of the queue to handle the process steps. Each Job server within the network must have a unique queue identifier. The process steps are requested by the Job queue using this exact 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.

Syntax:

<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 the Designer.

4. On the **Machine roles** page, select **Job server**.
5. On the **Server functions** page, select **Generic database connector**.
6. On the **Service Settings** page, enter the connection data and check the One Identity Manager Service configuration.

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

For a direct connection to the database:

- a. In the module list, select **Process collection > sqlprovider**.
- b. Click the **Connection parameter** entry, then click the **Edit** button.
- c. Enter the connection data for the One Identity Manager database.
- d. Click **OK**.

For a connection to the application server:

- a. In the module list, select the **Process collection** entry and click the **Insert** button.
- b. Select **AppServerJobProvider** and click **OK**.

- c. In the module list, select **Process collection > AppServerJobProvider**.
 - d. Click the **Connection parameter** entry, then click the **Edit** button.
 - e. Enter the address (URL) for the application server and click **OK**.
 - f. Click the **Authentication data** entry and click the **Edit** button.
 - g. In the **Authentication method** dialog, select the authentication module for logging in. Depending on the authentication module, other data may be required, such as user and password. For more information about One Identity Manager authentication modules, see the *One Identity Manager Authorization and Authentication Guide*.
 - h. Click **OK**.
7. To configure the installation, click **Next**.
 8. Confirm the security prompt with **Yes**.
 9. On the **Select installation source** page, select the directory with the install files. Change the directory if necessary.
 10. On the **Service access** page, enter the service's installation data.
 - **Computer:** Select the server, on which you want to install and start the service, from the menu or enter the server's name or IP address.
To run the installation locally, select **Local installation** from the menu.
 - **Service account:** Enter the details of the user account that the One Identity Manager Service is running under. Enter the user account, the user account's password and password confirmation.

The service is installed using the user account with which you are logged in to the administrative workstation. If you want to use another user account for installing the service, you can enter it in the advanced options.

You can also change the One Identity Manager Service details, such as the installation directory, name, display name, and the One Identity Manager Service description, using the advanced options.

11. Click **Next** to start installing the service.
Installation of the service occurs automatically and may take some time.
12. Click **Finish** on the last page of the Server Installer.

NOTE: In a default installation, the service is entered in the server's service management with the name **One Identity Manager Service**.

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 generic database connector.

Table 5: Information required for setting up a synchronization project

Data	Explanation
Synchronization server	<p>All One Identity Manager Service actions are run 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 the target system to do this. Sometimes direct access from the workstation, on which the Synchronization Editor is installed, is not possible. For example, because of the firewall configuration or the workstation does not fulfill the necessary hardware and software requirements. If direct access is not possible from the workstation, you can set up a remote connection.</p> <p>Remote connection server configuration:</p> <ul style="list-style-type: none">• One Identity Manager Service is started• RemoteConnectPlugin is installed and an authentication method is set up• ADO.NET provider for DB2 (LUW) is installed <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 as the synchronization server (with regard to the installed software and entitlements). Use the synchronization server as remote connection server as well by installing the RemoteConnectPlugin.</p> <p>For more detailed information about setting up a remote connection, see the <i>One Identity Manager Target System Synchronization Reference Guide</i>.</p>
Synchronization	Set the Data import option in the synchronization step if

Data	Explanation
workflow	<p>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 table from the Base table 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>. • The Synchronization servers menu displays all Job servers for which the Generic database connector server function is set.
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 generic 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

- [Creating a synchronization project](#) on page 16

Creating a synchronization project

There is a wizard to assist you with setting up a synchronization project. This wizard takes you through 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 to configure a synchronization project if the Synchronization Editor is both:

- Run in default mode
- Started from the Launchpad

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

To set up a synchronization project

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

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

2. Select the **Generic Database Connector** and click **Run**.
This starts the Synchronization Editor's project wizard.
3. On the wizard's start page, click **Next**.
4. On the **System access** page, specify how One Identity Manager can access the target system.


- If access is possible from the workstation on which you started the Synchronization Editor, do not change any settings.
- If access is not possible from the workstation on which you started the Synchronization Editor, you can set up a remote connection.

Select the **Connect using remote connection server** and enter the remote connection properties.

Remote connection properties

- **Access parameters**

- **Server:** Full server name or IP address of the server.

To select an existing Job server as the remote connection server, click  and select the server from the menu. This displays all the Job servers that have the **One Identity Manager Service installed** server function selected.

- **Port:** Port that is configured for the RemoteConnectPlugin.

- **Authentication**


If **SecretAuthentication** is configured for the RemoteConnectPlugin:

- **Secret:** Secret used by the Synchronization Editor to authenticate on the RemoteConnectPlugin.

If **ADGroupAuthentication** is configured for the RemoteConnectPlugin, no data is required.

- **Options**

- **Request timeout:** Maximum time allowed for a server query in seconds. If the time is exceeded, the request is canceled.
 - **Accept self-signed certificates:** Specifies whether self-signed certificates can be accepted.
- Click **Next** to start the system connection wizard to create a connection to an external database.
5. On the **Select database system** page, select the database system to which you want to connect.
 - Select **DB2 (LUW)**.
 6. Configure the system connection.

For more information, see [Connecting a system to a DB2 \(LUW\) database](#) on page 18.
 7. On the **Save configuration** page, you can save the current configuration as a template. 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. On the last page of the system connection wizard, you can save the connection data.
 - 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.
 9. 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.
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. On the **Select project template** page, select a project template to use for setting up the synchronization configuration.

NOTE: The generic 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**.

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

Table 6: General properties of the synchronization project

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

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

Connecting a system to a DB2 (LUW) database

Table 7: 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 generic 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 provider	Provider used to connect to the external database.

To configure the connection to a DB2 (LUW) database:

1. On the **Database connection** page, enter the connection parameters. Enter all the parameters required by the generic 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. On the **Describe the database** page, enter a display name and a unique identifier for the database connection.
 - **Database display name:** Display name of the database for displaying in the One Identity Manager tools.
 - **System identifier:** Unique system identifier.

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.
- **System category:** Category for further differentiation of system types. A maximum of 16 characters, consisting of numbers and letters, are permitted for the identifier.

Connections with the generic database connector are usually stored with the **DB** system type. The system category allows one more classification into subtypes. This allows different database connections with differing content to be distinguished from one another.

Example:

Connection 1: type=**DB**, system category=**Rights**

Connection 2: type=**DB**, system category=**PersonImport**

3. On the **Load configuration** page, you can enter a file 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. On the **Time zone selection** page, select the time zone for the time zone data in the database. The time zone is required to convert the time saved in the database into the local time. The local time is displayed in One Identity Manager tools.
5. On the **Initializing** page, you can specify additional connection settings. Write a script in the database syntax to specify number and date formats, language, and data sort order, for example. This script is then run every time you connect the system.
6. 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.

- In the **Partial schemas/owner** list, enable the schemas you want to be processed.
7. 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.

8. 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 8: Defining unique keys

Property	Description
Hide unconfigured tables	Specifies whether tables 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	<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 9: Column group properties

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

| **NOTE:** Columns of type Char are not supported.

- On the **Define data relations** page, you can enter information about object relations.

Table 10: Defining column relations

Property	Description
Hide unconfigured tables	Specifies whether tables are hidden if no settings have been changed.
Schema	Database schema tables.
Target(s)	<p>Columns to which the reference refers. Enter the table and column name in the following syntax:</p> <pre><schema>.<table name>.<column name></pre> <p>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.

- On the **Complete schema** page, you can enter additional schema information.

Table 11: Additional schema information

Property	Description
Hide unconfigured tables	Specifies whether tables are hidden if no settings have been changed.
Schema	Tables and schemas of the database schema.
Display value	<p>Column used in the display pattern.</p> <ul style="list-style-type: none"> To use the column in the 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 String or Integer data type can be selected.
Contains sensitive data	Specifies whether the column contains sensitive data.

Property	Description
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	Specifies whether the value in this column maps the path in an object hierarchy. If this table's objects are sorted by this column, it results in a list sorted in hierarchical order. This makes it possible to resolve object dependencies. Only one column per table can be marked as a sort criterion. An example is the CanonicalName column.
Scope reference	Specifies whether the column can be used to form the reference scope. Only one column per schema type can be labeled as the reference scope.
Auto fill behavior	Specifies whether the values entered in the column are identified automatically.
Enable compatible null handling	Specifies whether a null value in the column is converted to an empty string.
Incompatible data handling	Specifies how to handle data that is not allowed in Microsoft .NET Framework and thus cannot be mapped in One Identity Manager. <ul style="list-style-type: none"> • Post error: Cancels schema type matching with an error message. • Use default value: Uses the One Identity Manager default value. • Use minimum value: Uses the minimum value of the data type. • Use maximum value: Uses the maximum value of the data type.

Table 12: Table properties

Property	Description
Display template	Display pattern with which the objects in the 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.


11. On the **Define data operations** page, you can specify special operations for changing data in the external database. 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 13: Defining data operations

Property	Description
Hide unconfigured tables	Specifies whether tables are hidden if no settings have been changed.
Table/operation	Tables for which the data operations are to be defined.
Strategy	<p>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 run. Select the strategy you want use to define the data operation.</p> <ul style="list-style-type: none"> • Pattern based: Simple procedure call that runs the operation. <ul style="list-style-type: none"> NOTE: If column names are referenced that contain special characters, replace the special characters with an underscore (_). Special characters are everything apart from letters, numbers and underscore. • Script based: Script that performs a complex data operation. <p>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>.</p> • Click  to delete a data operation.

Property	Description
Required columns	List of required key columns in a script-based data operation. The columns must be entered if they are 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 more information about support for creating scripts, see the <i>One Identity Manager Target System Synchronization Reference Guide</i>.</p>

12. The **Extend target system schema** page opens if you enable the **Fill in system description completely** option on the **Load schema** page or make settings on the **Define data operations** page. You can add virtual schema properties to the target system schema here. Use the virtual schema properties to provide additional data for your own DML handling.

Table 14: Virtual schema properties

Property	Description
Hide unconfigured tables	Specifies whether tables are hidden if no settings have been changed.
Schema	Tables in the target system schema for which virtual schema properties can be added or exist already.
Virtual schema properties	<p>Buttons for editing virtual schema properties</p> <ul style="list-style-type: none"> Click Add to add a virtual schema property. Click Edit or remove to edit or delete a virtual schema property.
Data type	Data type of the schema property
Multivalue	Specifies whether the schema property has multiple values
Required field	Specifies whether the schema property is a mandatory property
Secret	Specifies whether the schema property value may be displayed in

Property	Description
	the Synchronization Editor or in logs, reports, and messages. If this option is set, the value is not displayed.

To edit or delete a virtual schema property

- a. In the **Schema** column, open the node of the table with the schema properties that you want to edit or delete.
- b. Click **Edit or remove**.
- c. Edit the properties of the virtual schema property.
 - OR -
 - Click **Delete**.

Related topics

- [Creating a synchronization project](#) on page 16

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 the loading of 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 compression 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. In the Synchronization Editor, open the synchronization project.
2. Select the **Configuration > Target system** category.
- OR -
Select the **Configuration > One Identity Manager connection** category.
3. Select the **General** view and click **Update schema**.
4. Confirm the security prompt with **Yes**.
This reloads the schema data.

To edit a mapping

1. In the Synchronization Editor, open the synchronization project.
2. Select the **Mappings** category.
3. Select a mapping in the navigation view.
Opens the Mapping Editor. For more 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 running times for synchronization.

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

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

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

- If another synchronization process is started with the same start up configuration, the process is stopped and is assigned **Frozen** status. An error message is written to the One Identity Manager Service log file.
 - Ensure that start up configurations that are used in start up sequences are not started individually at the same time. Assign start up sequences and start

up configurations different schedules.

- Starting another synchronization process with different start up configuration that addresses same target system may lead to synchronization errors or loss of data. Specify One Identity Manager behavior in this case, in the start up configuration.
 - Use the schedule to ensure that the start up configurations are run in sequence.
 - Group start up configurations with the same start up behavior.

If you want to specify the order in which target systems are synchronized, use the start up sequence to run synchronization. In a start up sequence, you can combine start up configurations from different synchronization projects and specify the order in which they are run. For more information about start up sequences, see the *One Identity Manager Target System Synchronization Reference Guide*.

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. In the Synchronization Editor, open the synchronization project.
2. Select the **Logs** category.
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 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, 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 synchronizations.
- 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 28.


Related topics

- [Post-processing outstanding objects](#) on page 30
- [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 **Data Synchronization > Basic configuration data > Target system types** category.
2. Click  in the result list.
3. Edit the target system type main data.
4. Save the changes.

Enter the following data for a target system type.

Table 15: main data for a target system type

Property	Description
Target system type	Target system type description.
Description	Text field for additional explanation.
Display	Name of the target system type as displayed in One Identity Manager

Property	Description
name	tools.
Cross-boundary inheritance	<p>Specifies how user accounts are assigned to or inherit groups and system entitlements if they belong to different custom target systems.</p> <ul style="list-style-type: none"> If the option is set, groups and system entitlements can be assigned to user accounts that belong to the same target system or to different target systems. The target systems must have the same target system type. <p>For all target systems of a target system type, the settings for the User account has memberships column (UNSRootB.UserContainsGroupList) must be identical.</p> <ul style="list-style-type: none"> If the option is not set, groups and system entitlements can only be assigned to the same target system. <p>NOTE: If the option is not set, the target system type is used to simplify grouping of the target systems.</p>
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.
Alternative connectors	List of connector that can process this type of target system.

To add tables to target system synchronization

1. In the Manager, select the **Data Synchronization > Basic configuration data > Target system types** category.
2. In the result list, select the target system type.
3. Select the **Assign synchronization tables** task.
4. In the pane, assign **custom** tables to the outstanding objects you want to handle.
5. Save the changes.
6. Select the **Configure tables for publishing** task.
7. Select the tables that contain the outstanding objects that can be published in the target system and set the **Publishable** option.
8. 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 **Connection is read-only** option 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 runs the provisioning of the objects. Use the `AdHocProjection` process task of the `ProjectorComponent` process component.

For more information about defining processes, see the *One Identity Manager Configuration Guide*.

Post-processing outstanding objects

To post-process outstanding objects



1. In the Manager, select the **Data synchronization > Target system synchronization: <target system type>** category.
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 the properties of an outstanding object

1. Select the object on the target system synchronization form.
 2. Open the context menu and click **Show object**.
 3. For memberships, select the object whose properties you want to display.
3. Select the objects you want to rework. Multi-select is possible.
 4. Click on one of the following icons in the form toolbar to run the respective method.

Table 16: Methods for handling outstanding objects

Icon	Method	Description
	Delete	The object is immediately deleted from the One Identity Manager database. Deferred deletion is not taken into account. Indirect memberships cannot be deleted.
	Publish	The object is added to the target system. The Outstanding label is removed from the object. This runs a target system specific process that triggers the provisioning process for the object. Prerequisites:

Icon	Method	Description
		<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.

TIP: If a method cannot be run due to certain restrictions, the respective icon is disabled.

- To display the constraint's details, click the **Show** button in the **Constraints** column.

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

NOTE: By default, the selected objects are processed in parallel, which speeds up 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

- Disable the  icon in the form's toolbar.

Related topics

- [Configuring target system synchronization on page 28](#)
- [Users and permissions for synchronizing on page 5](#)

Configuring the provisioning of memberships

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

- Memberships are saved as an object property in list form in the target system.
Examples: List of user accounts in the Member property of a group - OR - List of profiles in the MemberOf property of a user account

- Memberships can be modified in either of the connected systems.
- A provisioning workflow and provisioning processes are set up.

If one membership in One Identity Manager changes, by default, the complete list of members is transferred to the target system. Therefore, memberships that were previously added to the target system are removed in the process and 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 **Data Synchronization > Basic configuration data > Target system types** category.
2. In the result list, select the target system type.
3. Select the **Configure tables for publishing** task.
4. Select the assignment tables that you want to set up for single provisioning. Multi-select is possible.
5. Click **Merge mode**.

NOTE:


- This option can only be enabled for assignment tables that have a base table with a `XDateSubItem` column.
- Assignment tables that are grouped together in a virtual schema property in the mapping must be marked identically.

Example: `ADSAccountInADSGroup`, `ADSGroupInADSGroup`, and `ADSMachineInADSGroup`

6. Save the changes.

For each assignment table labeled like this, the changes made in One Identity Manager are saved in a separate table. Therefore, only newly added and deleted assignments are processed. 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 not the entire members list.

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.

You can restrict single provisioning of memberships with a condition. Once merge mode has been disabled for a table, the condition is deleted. Tables that have had the condition deleted or edited are marked with the following icon: . You can restore the original condition at any time.

To restore the original condition

1. Select the auxiliary table for which you want to restore the condition.
2. Right-click on the selected row and select the **Restore original values** context menu item.
3. Save the changes.

NOTE: To create the reference to the added or deleted assignments in the condition, use the *i* table alias.

Example of a condition on the UNSAccountBInUNSGroupB assignment table:

```
exists (select top 1 1 from UNSGroupB g
        where g.UID_UNSGroupB = i.UID_UNSGroupB
        and <limiting condition>)
```

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

Configuring single object synchronization

Changes made to individual objects in the target system can be immediately applied in the One Identity Manager database without having to start a full synchronization of the target system environment. Individual objects can only be synchronized if the object is already present in the One Identity Manager database. The changes are applied to the mapped object properties. If the object is no longer present in the target system, then it is deleted from the One Identity Manager database.

Prerequisites

- A synchronization step exists that can import the changes to the changed object into One Identity Manager.
- The table that contains the changed object is assigned to a target system type.
- The path to the base object of the synchronization is defined for the table that contains the changed object.

Specify the tables that you want to synchronize using single object synchronization and configure single object synchronization for these tables. For more information, see the *One Identity Manager Target System Synchronization Reference Guide*, section *Include custom tables in the synchronization*.

To define the path to the base object for synchronization for a table

1. In the Manager, select the **Data Synchronization > Basic configuration data > Target system types** category.
2. In the result list, select the target system type.

3. Select the **Assign synchronization tables** task.
4. In the **Add assignments** pane, assign the table for which you want to use single object synchronization.
5. Save the changes.
6. Select the **Configure tables for publishing** task.
7. Select the table and enter the **Root object path**.
 - If a concrete base object is defined for the target system, enter the path to the base object in the ObjectWalker notation of the VI.DB.
Example: `FK(UID_GAPCustomer).XObjectKey`
 - If no concrete base object is defined for the target system, enter the XObjectKey of the base table.
Example: `<Key><T>DialogTable</T><P>RMB-T-Org</P></Key>`
8. Save the changes.

Ignoring data error in synchronization

By default, objects with incorrect data are not synchronized. For example, a user account is not loaded in the One Identity Manager database if, in the user account table, the formatting script of a column contains an email address detects invalid data. These objects can be synchronized once the data has been corrected. In certain situations, however, it might be necessary to synchronize objects like these and ignore the data properties that have errors. This synchronization behavior can be configured in One Identity Manager.

To ignoring data errors during synchronization in One Identity Manager

1. In the Synchronization Editor, open the synchronization project.
2. Select the **Configuration > One Identity Manager connection** category.
3. In the **General** view, click **Edit connection**.
This starts the system connection wizard.
4. On the **Additional options** page, enable **Try to ignore data errors**.
This option is only effective if **Continue on error** is set in the synchronization workflow.
Default columns, such as primary keys, UID columns, or mandatory input columns cannot be ignored.
5. Save the changes.

IMPORTANT: If this option is set, One Identity Manager tries to ignore commit errors that could be related to data errors in a single column. This causes the data changed in the affected column to be discarded and the object is subsequently saved again. This effects performance and leads to loss of data.

Only set this option in the exceptional circumstance of not being able to correct the data before synchronization.

Troubleshooting

For more 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 that arise during synchronization, inadequate performance for example. The report contains information such as:

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

To generate a synchronization analysis report

1. In the Synchronization Editor, open the synchronization project.
2. Select the **Help > Generate synchronization analysis report** menu item and click **Yes** in the security prompt.

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 and other inquiries, such as licensing, support, and renewals, visit <https://www.oneidentity.com/company/contact-us.aspx>.

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, 33

D

database connector
generic 4

J

Job server
edit 10

M

membership
modify provisioning 31

O

object
delete immediately 30
outstanding 27, 30
publish 30
outstanding object 27

P

provisioning
members list 31

R

remote connection server 13

S

schema
changes 25
shrink 25
update 25
single object synchronization 33
synchronization
start 26
synchronization analysis report 36
synchronization configuration 13, 16
synchronization log 27
synchronization server 13
configure 10
install 10
Job server 10

T

target system synchronization
table to assign 28
target system type 28

V

variable set 13

W

workflow 13