Configuration and administration

ELO Replication

Table of contents

ELO Replication

Getting started

You will find an overview of configuring ELO Replication at two sites here. Follow the links for more detailed information about the individual steps.

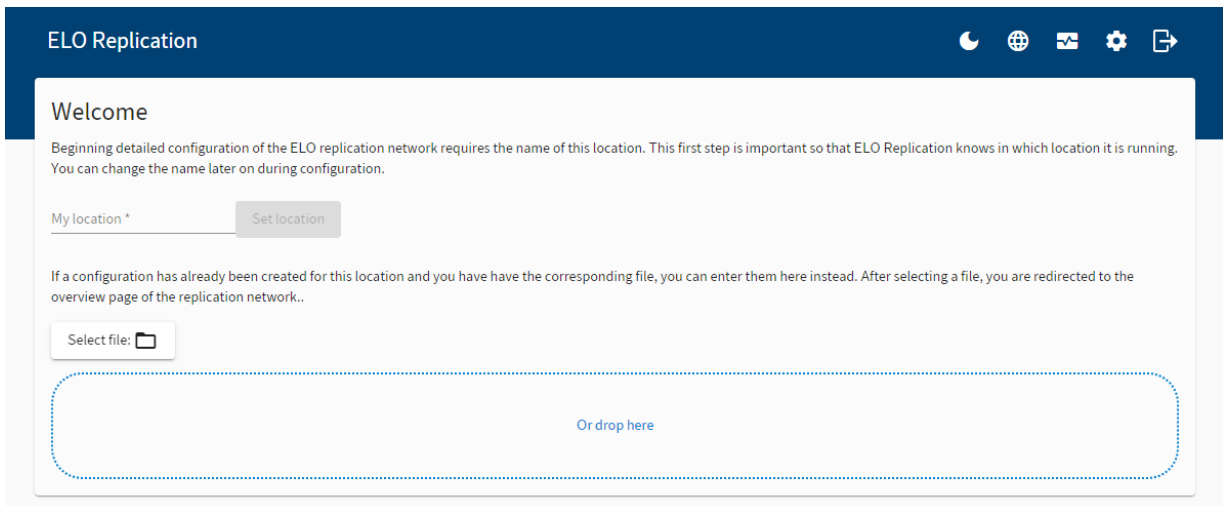### Overview: Configuring replication at two sites

To configure ELO Replication at at least two sites, perform the following steps:

1. Install at site 1 and site 2.

Install ELO Replication at both sites using the ELO Server Setup on the *Applications* tab.

2. Configure site 1.

    1. Open your browser and go to the ELO Application Server overview page.

    2. To open the configuration, select *rp*.

    3. Log on to ELO Replication with your Apache Tomcat administrator account.

    4. Choose a site name.

       Once you have set your site, draft mode opens.

You can enter the settings for all participating sites from your site.

5. To add a new site, select the plus icon and then click anywhere in the viewer pane.

6. To add a new connection between two sites, place the mouse on a site. The connection icon appears. Press and hold down the mouse button and drag the connection to another site.

7. To show the setting options, select a site.

   The site settings appear on the left.

Site-Name

Stuttgart

Interfaces ⌃

Scheme

SSH ▾   🖊

IP address or server name

localhost

Enter the public address for this branch

Port
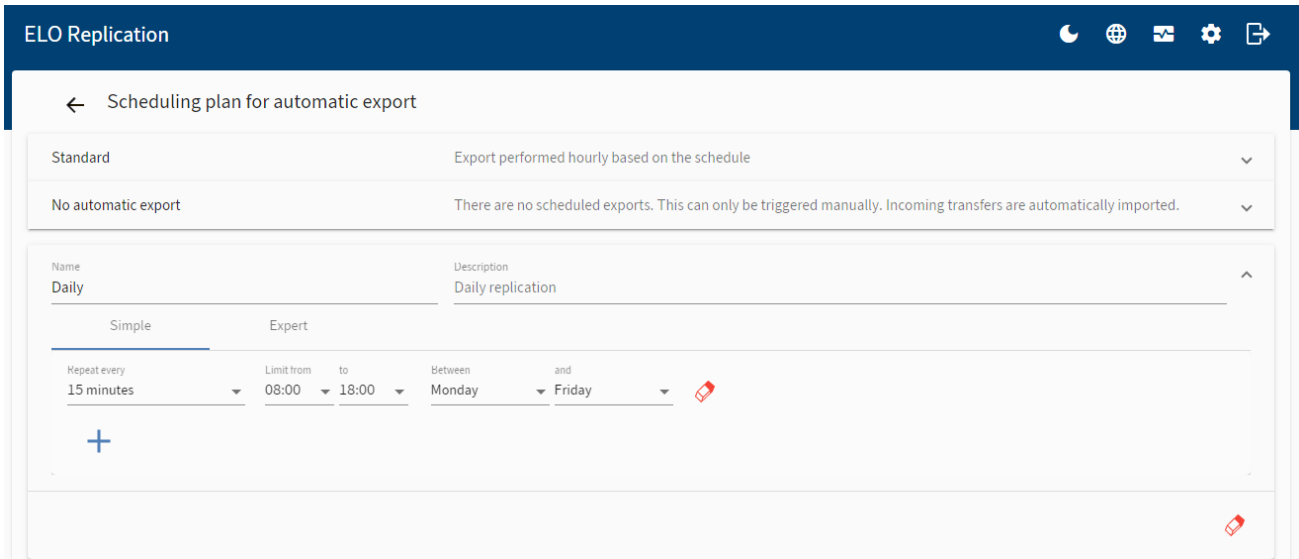
9096

Enter the address port

＋

Repositories ⌄

8. Fill out the fields accordingly.

9. Click the plus icon in the *Repositories* area to add a repository. A replication set is automatically created for each repository, which you can assign to an entry in the ELO Java Client or the ELO Web Client.

10. Under Scheduler for automatic transfer, define when to export data from the repository.

> **Please note**
>
> The specified interval time refers to the local time of the respective server. All servers hosting repositories involved in replication must be in sync in terms of time for the comparison via timestamp to work.

> **Please note: The scheduler calculates in hours. For example, if you enter a time from 8:00 a.m. to 6:00 p.m., replication begins at 8:00 a.m. and ends at 6:59 p.m.**

In the drop-down menu you will find the preconfigured intervals *Default* and *Import only*, as well as all intervals you have created. To create new intervals or change existing ones, select the pencil icon.

The *Scheduler for automatic transfer* menu opens. Clicking the plus icon adds a new schedule. Clicking the eraser icon deletes a schedule.

The input options *Simple* and *Expert* are available for the *scheduler for automatic transfer*. With the input option *Simple*, you can choose from default intervals and customize them. With the input option *Expert*, you can control the scheduler for automatic transfer more precisely using variables.

3. Publish replication at site 1.

Once you are finished making settings, you can enable replication. In draft mode, select *Publish*.

Please note

During initial configuration, you will have to publish the configuration at all participating sites.

4. Send configuration file to site 2.

A configuration file with your settings is created automatically for each site. Download the configuration file by selecting *Download configuration* on the overview page. Send the configuration file to the relevant site. There, the administrator logs on to ELO Replication and uploads the configuration file.
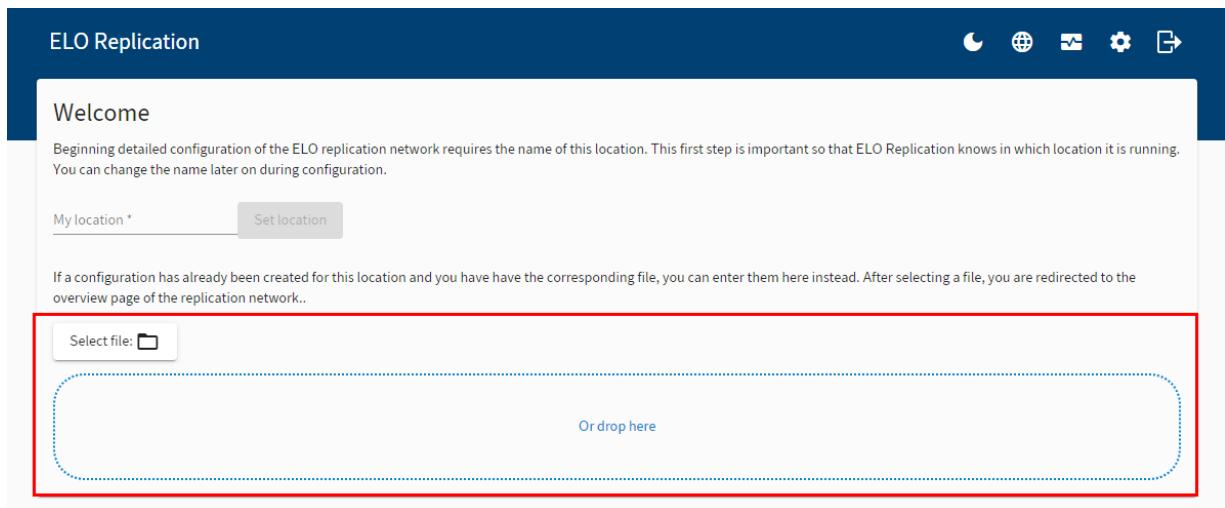
Site: Hamburg
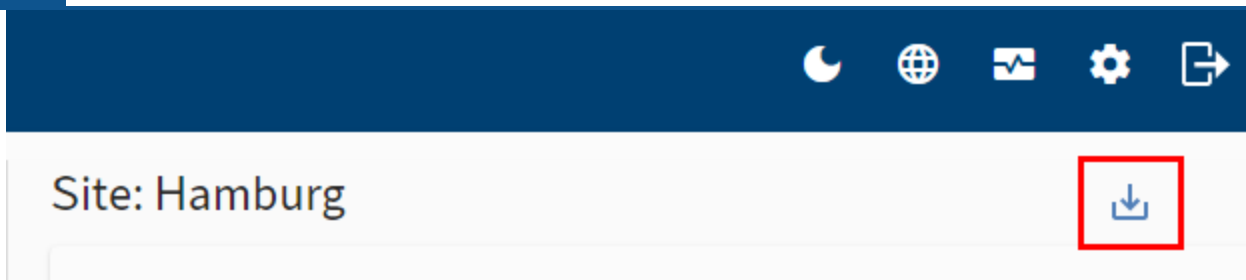
5. Insert configuration file at site 2.

    1. Open the configuration via the ELO Application Server overview page and log on to ELO Replication with the Apache Tomcat administrator account.
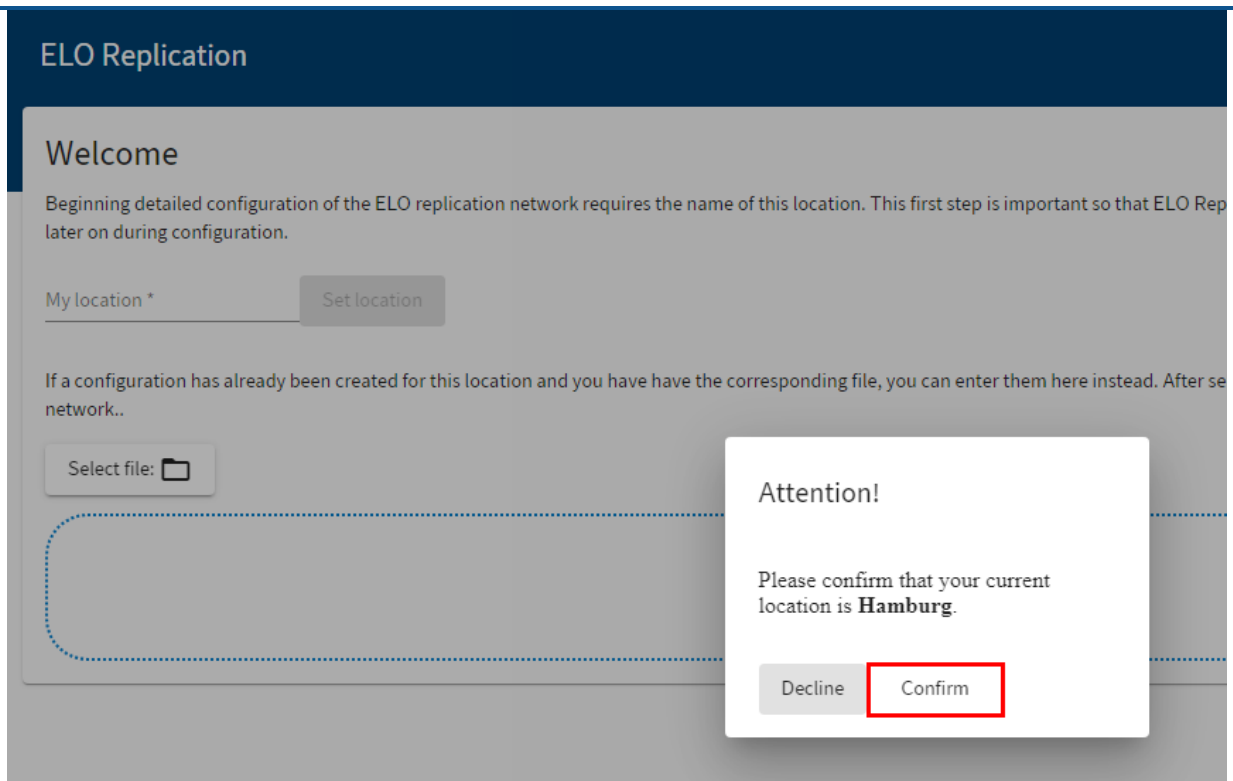
> **Information**
>
> You do not have to enter a site name. The site name from the configuration file is automatically applied.

ELO Replication

Welcome

Beginning detailed configuration of the ELO replication network requires the name of this location. This first step is important so that ELO Replication knows in which location it is running. You can change the name later on during configuration.

My location *   Set location

If a configuration has already been created for this location and you have have the corresponding file, you can enter them here instead. After selecting a file, you are redirected to the overview page of the replication network..

Select file:

Or drop here

   2. Insert the configuration file.

    3. Confirm the site.

The settings are applied.

6. Publish replication at site 2.

To enable replication, in draft mode, select *Publish*.

7. Assign replication sets at site 1 and/or site 2.

Once you have published the replication at the participating sites, you have to assign the individual entries replication sets. Each replication set stands for one repository. The replication set is a characteristic that you assign to individual repositories in order to replicate them in another repository. The individual entries to be replicated are selected in the ELO Java Client or ELO Web Client using the *Assign replication sets* function.
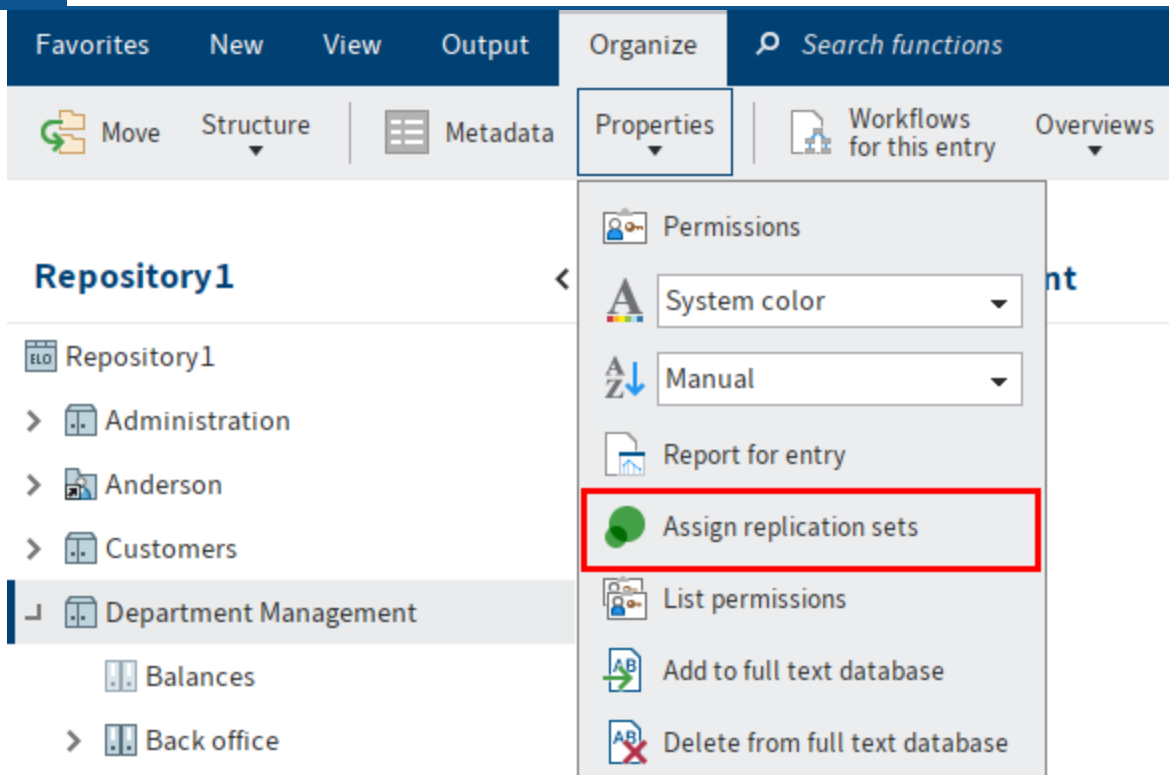
You will find the *Assign replication sets* functions in *Ribbon > Organize > Properties*.

> **Please note You should log which entries you assign which replication sets.**

> **Please note**
>
> Replication sets are inherited upward to parent folders in the repository structure. To prevent inheritance, select the *Start point for replication* option in the *Metadata* dialog box of the entry you want to replicate. In the ELO Java Client, this function is only available for folders. You will find more information in the chapter Replication set inheritance.

Result

Once you have performed the steps above, the entries are replicated according to your settings in the configuration.

Basics

In the following chapter, you will learn how data synchronization and data transfer work.

Functionality

ELO Replication compares entries between multiple repositories. The relevant repositories can be installed at different sites. ELO Replication transfers the data to the relevant repositories. This means the repositories do not necessarily have to be available to one another.

ELO Replication is a web app based on Java that is installed on an Apache Tomcat.

Replication set

The replication set is a characteristic that you assign to individual repositories in order to replicate them in another repository. Create sites and add repositories in the web-based configuration. Using the ELO Indexserver URL, specify which repositories data is replicated from. A replication set is created automatically in the configuration when you add a new repository. Each replication set stands for one repository. The individual entries (folders, documents) to be replicated are selected in the ELO Java Client or ELO Web Client using the "Assign replication sets" function. With this function, you select the repository you want to export the entries to. This means you select a repository the entries are replicated to.

Data synchronization

ELO Replication captures, distributes, and transfers changes to the relevant repositories. An extension to the ELO Indexserver creates a synchronization data set with the changes to a repository. The format of this data set is a compressed stream of JSON objects from the ELO Indexserver API. The data in this stream is selected in the ELO Indexserver based on its synchronization status. The following options are available:

Entries without a replication set: Entries that you do not want to synchronize are not assigned a replication set. A replication set determines which other repositories the entry is synchronized with. Entries without a replication set are not added to the synchronization data set.

Entries with a new replication set: Entries that you want to synchronize are assigned new replication sets. The information from the entries is added to the synchronization data set. All the entry information is only sent to the repositories assigned the new replication set. If entries have already been assigned replication sets, their repositories are simply informed that the entries are now replicated with additional repositories.

The ELO Indexserver assigns replicated entries the *tstampsync* field when creating the synchronization data set in the database. The *tstampsync* field contains the value from the corresponding *tstamp* field at the time the data was read from the database. In the synchronization data set, the *tstampsync* contains the value it had when the database was read. This difference plays an important role when importing the data set.

> **Information**
>
> The name of the *tstampsync* field varies depending on the table.

Entries changed since the last synchronization: The entries were changed since the last synchronization. To recognize a change, the *tstamp* and *tstampsync* fields are compared. In case of changes, the ELO Indexserver automatically sets the *tstamp* field to the current time in UTC (Coordinated Universal Time).

Entries unchanged since the last synchronization: The entries were not changed since the last synchronization. For unchanged entries, the values in the *tstamp* and *tstampsync* fields are identical. Unchanged entries are not added to the synchronization data set.

Data transfer

The synchronization data set is created by the ELO Indexserver. ELO Replication initiates its creation based on a configured schedule. The synchronization data set is streamed by the ELO Replication to the other sites, and from there it is streamed to the ELO Replication of the other repositories so it can be imported. During streaming, the data set is processed between the ELO Replication instances. Only the data required is sent at the target is sent.

To compensate for instabilities during transmission, ELO Replication caches the data sets. If disconnected, ELO Replication re-attempts to send the data set once a minute. The data set is transferred all over again, regardless of when the previous transfer was interrupted.

The SSHD library from the Apache MINA project is used for data transfer. The method with public and private keys is used exclusively for authentication. The keys are automatically generated for each site when configuring ELO Replication.

> **Information**
>
> You can add additional options for transferring data by creating and integrating plug-ins.

The ELO Indexserver imports the synchronization data set into the target repository. The following options are available:

Entry does not exist in repository: The GUID is used to check whether an entry already exists in the repository. If the entry does not exist in the repository, it is imported.

Entry already exists in the repository and has not changed since the last synchronization: The values in the *tstampsync* fields of the synchronization data set and repository entries are identical. The entry is imported into the repository. The entry in the repository is overwritten with the values from the synchronization data set.

Entry already exists in the repository and has competing changes: The values in the *tstamp* fields of the synchronization data set and repository entries do not match. If an entry has been changed in multiple repositories, the latest change is applied. If the latest change comes from the target

repository, the entry in the synchronization data set is ignored. Otherwise, the values of the entry in the repository would be replaced with the values in the synchronization data set.

In rare cases, the values in the *tstampsync* fields may differ. This occurs when a synchronization data set is created in the local repository before the synchronization data set of the other repository has been read and the entry is changed in both repositories. The latest change is applied in this case as well.

What data is replicated?

The following data is synchronized during replication:

- Folders
- Documents
- Sticky notes
- Relations
- Keyword lists
- Workflows
- Workflow templates
- Map data
- Feed
- Translations
- Master data: users and groups (via owners and ACLs), metadata forms, index field templates, aspects, colors, replication sets

Master data is resolved recursively. For example, if a user is listed in the ACL of a folder, this user's groups are also included in the synchronization data set.

> **Please note**
>
> File version histories are not replicated.

How is master data replicated?

During data export, the ELO Indexserver checks which master data belongs to a SORD.

Only users that a SORD explicitly refers to are replicated. If you want to replicate a specific user, this user has to be referenced by a SORD, e.g. via permissions (*Metadata > Permissions*) or owner rights (e.g. Create folder, File document, Apply stamp/annotation, or Start workflow).

> **Please note**
>
> If a user is replicated, the groups the user belongs to are also replicated (without the individual members).
>
> If a group is replicated, for example because it has permission to a SORD, the individual group members are not replicated.

If a metadata form is replicated, a time stamp is set in the *masktstampsync* field of the *docmasks* table in the database. Metadata forms that have already been replicated are added to the synchronization data set if changes have been made to the forms.

Templates for index fields are first replicated based on their use in replicating metadata forms. After this, changes to the index field templates are captured and replicated via the values of the *tstamp* and *tstampsync* fields.

Aspects are replicated based on their use in replicated metadata forms. Changes are captured via the values of the *tstamp* and *tstampsync* fields. In the event of a conflict, the "winning" aspect overwrites the "losing" aspect. In the future, after initial replication due to metadata form use, aspects will be replicated after every change.

How are workflows replicated?

The following section describes the behavior starting with ELO Indexserver version 23.

With the default settings, a workflow can only run in one repository. During data export, the workflow is assigned a flag indicating which repository it is running in. In the target repository, the workflow is displayed after replication but it will not continue. You cannot start, edit, or delete the workflow in the target repository.

To use the workflow in the target repository as well, the flag needs to be changed during export. This is done with a server transfer node, which is added in the workflow designer. If a server transfer node is set, the workflows stops at this node. Once the data has been transferred through replication, the workflow continues in the target repository. The entire workflow including all subworkflows is always replicated.

> **Please note**
>
> The server transfer can only take place in a main workflow, not in a subworkflow. If you want to start a subworkflow at site B, the server transfer must take place in the main workflow at site A. If a group is replicated, for example because it has permission to a SORD, the individual group members are not replicated.
>
> A subworkflow should only run at one site. You should not start a subworkflow at site A and continue it at site B, as conflicts that cannot be resolved automatically can occur if the main workflow and subworkflow are running in different repositories at the same time.

How are workflow templates replicated?

Workflow templates are first replicated based on their use in replicated workflows. After this, changes to workflow templates are also tracked with the values of the *tstamp* and *tstampsync* fields and replicated after every change. In the event of a conflict, the "winning" template overwrites the "losing" template.

How is map data replicated?

Map data is replicated based on its association with the SORD/document of the respective assigned replication sets. Map data is arranged in different map domains. These domains are also replicated. In the map domain, a flag can be used to indicate whether they should also be replicated. If this flag is set to *false*, the entire domain and all map data in it are ignored by replication.

Changes are tracked per map and SORD and captured via the values of the *tstamp* and *tstampsync* fields. Conflicts during import are therefore handled at the map level, and not for individual map fields. The "winning" map therefore overwrites the "losing" map.

How are keyword lists replicated?

Keyword lists are first replicated based on their use in metadata forms and aspects to be replicated. After this, changes to the keyword lists are captured and replicated via the values of the *tstamp* and *tstampsync* fields.

How are translations replicated?

Translations that are stored in the translation table are first replicated when the translation key is used in objects to be replicated. After this, changes to translations are captured and replicated via the values of the *tstamp* and *tstampsync* fields.

Installation

This chapter explains how to install ELO Replication.

Requirements

All servers hosting repositories involved in replication must be in sync in terms of time for the comparison via timestamp to work.

> **Please note**
>
> The same version of ELO Replication and the ELO Indexserver should be installed at all sites.

Method

> **Important**
>
> Replication sets and their assignments in the repository that were created in a version older than ELO Replication 12 will be deleted when migrating to ELO Replication 23. Before performing a migration, create an overview of replication sets and their assignments so you can reconfigure them in version 23.

1. Start the ELO Server Setup.
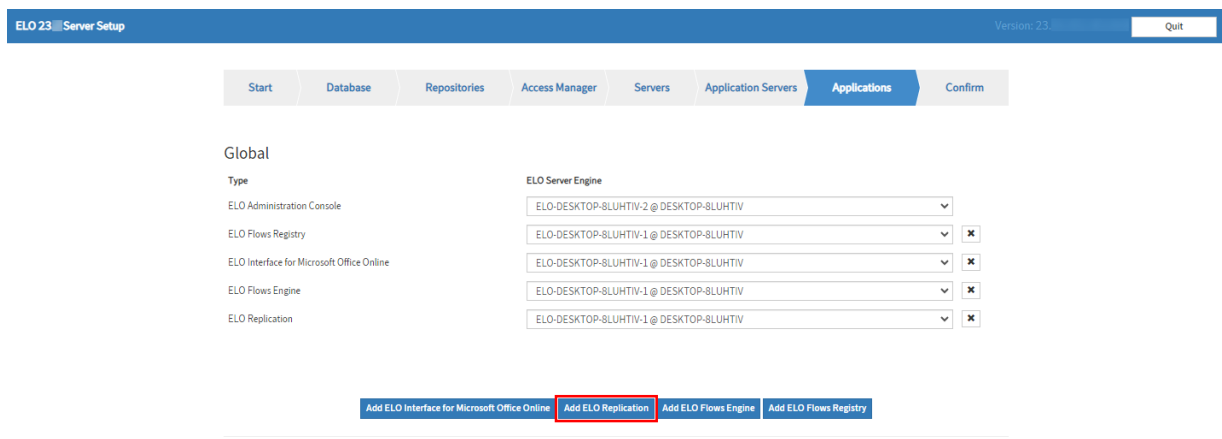
2. On the *Applications* tab, select *Add ELO Replication*.

3. Select an ELO server you want to install the module to.

4. Follow the instructions of the ELO Server Setup to complete installation.

Result

You have installed ELO Replication.

To open the ELO Replication configuration, open your browser and go to the ELO Application Server overview page. For more information, refer to the chapter Configuration.

Moving the data directory

If you would like to create a separate data directory for ELO Replication for space reasons, you will have to change the file path in the ELO Replication configuration file.
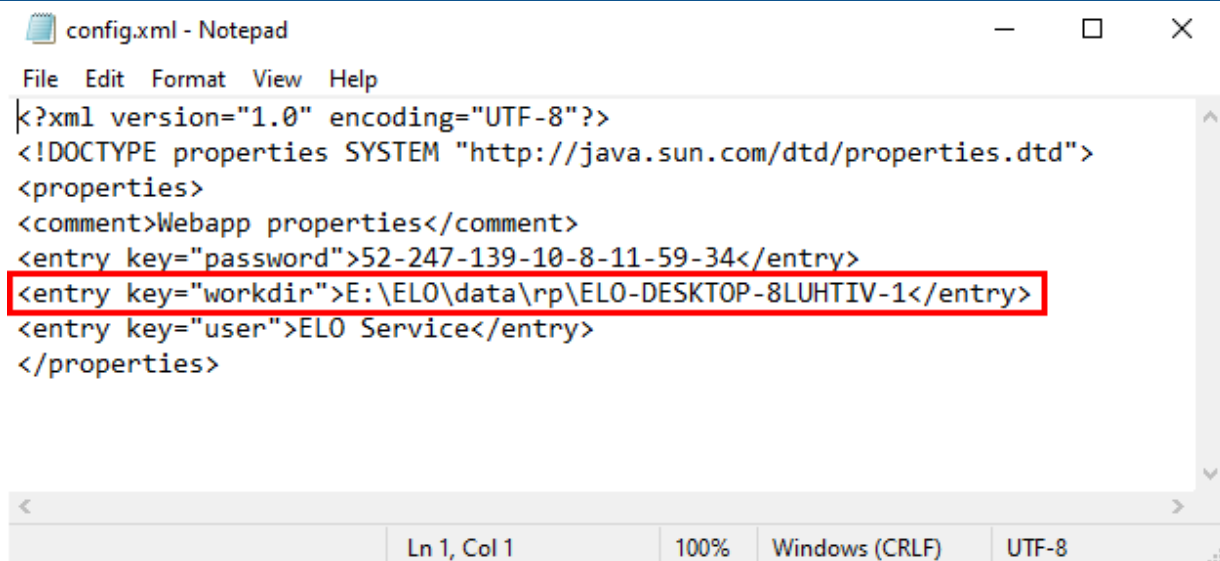
Method

1. In the *config* folder of your ELO directory, navigate to the folder for the Tomcat that ELO Replication is installed on:

   <ELO>\config\rp\<Tomcat of replication site>
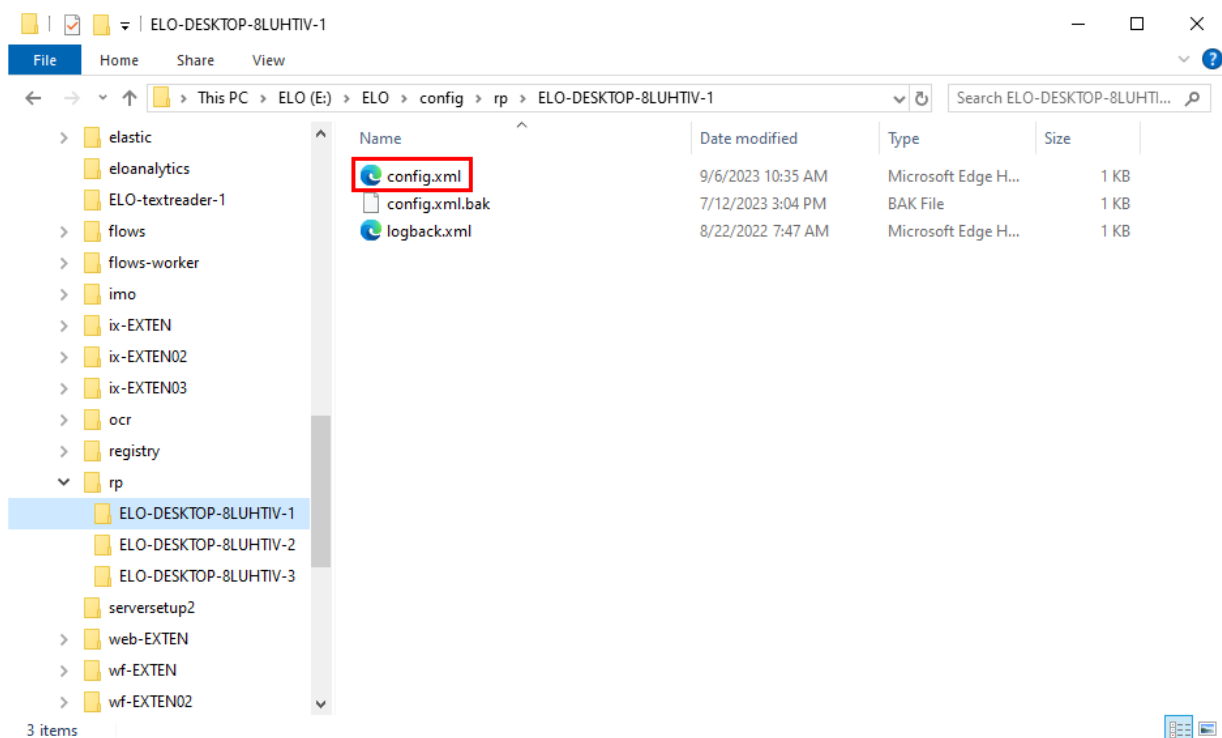
2. Open the configuration file *config.xml* with an editor.

```
config.xml - Notepad                                              —    □    ✕

File  Edit  Format  View  Help
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE properties SYSTEM "http://java.sun.com/dtd/properties.dtd">
<properties>
<comment>Webapp properties</comment>
<entry key="password">52-247-139-10-8-11-59-34</entry>
<entry key="workdir">E:\ELO\data\rp\ELO-DESKTOP-8LUHTIV-1</entry>
<entry key="user">ELO Service</entry>
</properties>




                      Ln 1, Col 1              100%   Windows (CRLF)     UTF-8
```

3. In the line *<entry key="workdir">*, change the file path to your new directory.

4. Save the changes to the file.

Result

The data for ELO Replication is now stored in the new directory.

Please note

If you perform an update of the ELO Server Setup to move the ELO Replication module to another Apache Tomcat server, files from the *data* folder in the directory *<ELO>\data\rp\<Tomcat of replication site>* are not transferred. To ensure that there are no files missing in the target repository, all files from the old *data* folder must be moved to the *data* folder in the new Tomcat directory.

You need to stop the Apache Tomcat server before you move the files. After restarting the Tomcat, the configuration is loaded.

If you make changes to the Tomcat, such as adding a new host or changing the ports, you need to publish the respective configurations in the ELO Replication menu again.

Configuration

Create sites and add repositories in the web-based configuration. Using the ELO Indexserver URL, specify which repositories data is replicated from. A replication set is created automatically in the configuration when you add a new repository. Each replication set stands for one repository. The replication set is a characteristic that you assign to individual repositories in order to replicate them in another repository. The individual entries (folders, documents) to be replicated are selected in the ELO Java Client or ELO Web Client using the "Assign replication sets" function. With this function, you select the repository you want to export the entries to. This means you select a repository the entries are replicated to.

> **Important**
>
> Replication sets and their assignments in the repository that were created in a version older than ELO Replication 12 will be deleted when migrating to ELO Replication 23. Before performing a migration, create an overview of replication sets and their assignments so you can reconfigure them in version 23.

To configure ELO Replication for the first time, log on at any site and configure the settings as desired. These settings are saved in a configuration file. You can send the configuration file to other sites. This means you only have to configure replication manually at one site.

Overview
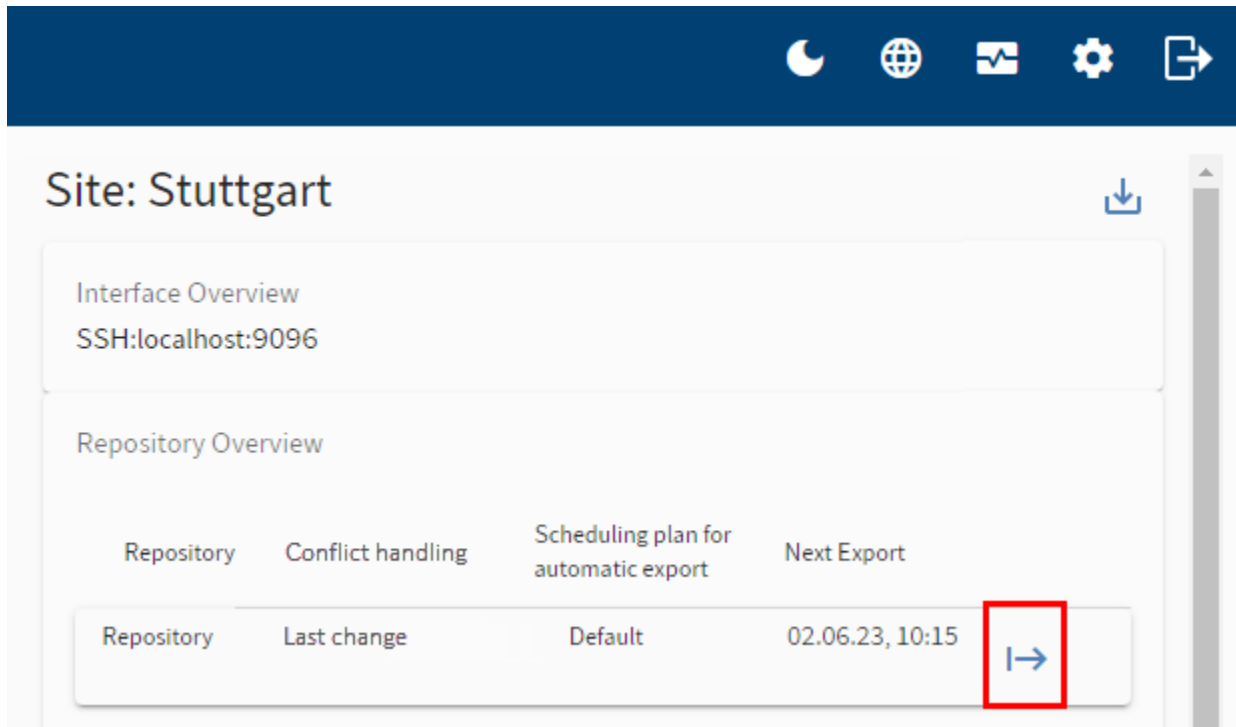
To configure ELO Replication at at least two sites, perform the following steps. For more information on the individual steps, refer to the chapters Installation and the following sections.

1. Install at site 1 and site 2.

2. Configure site 1.

3. Publish replication at site 1.

4. Send configuration file to site 2.

5. Insert configuration file at site 2.

6. Publish replication at site 2.

7. Assign replication sets at site 1 and/or site 2.

> **Information**
>
> Once you have performed steps 1 - 7, you can manually start an export from the local repository for test purposes. The export starts immediately, regardless of the configured schedule. Select *Start export* on the ELO Replication overview page. Once you have performed steps 1 - 7, you can manually start an export from the local repository for test

purposes. The export starts immediately, regardless of the configured schedule. Select *Start export* on the ELO Replication overview page.
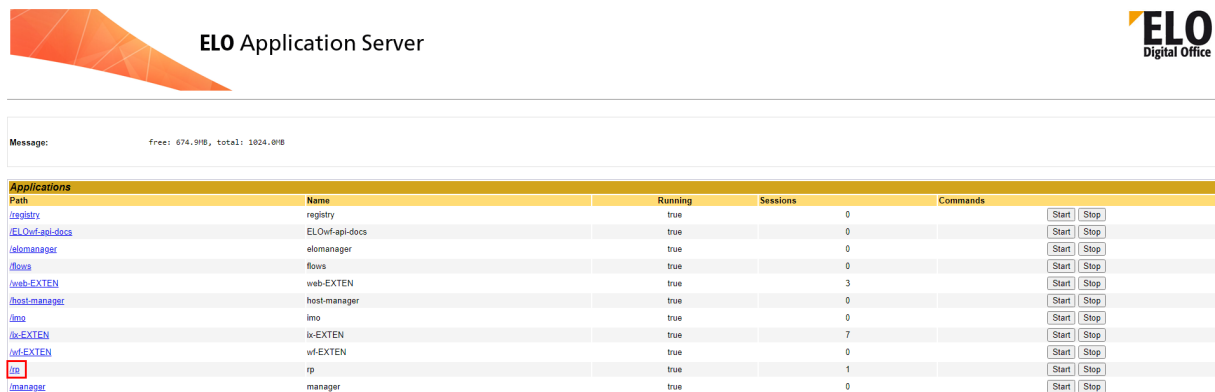
Initial logon

Information

Once you have installed ELO Replication, you can set up a replication network in the configuration. The first time you log on to ELO Replication, proceed as follows:

Step by step

1. Open your browser and go to the ELO Application Server overview page.

2. To open the configuration, select */rp*.

3.

Log on to ELO Replication with your Apache Tomcat administrator account.

4. The first time you log on, you will have to set a site name. If you already have a configuration file for ELO Replication, you can select it here.

> **Information**
>
> For more detailed information on inserting a configuration file, refer to the chapter Create/ insert a configuration for other sites.

Result

Once you have set your site, draft mode opens. Here, you can configure and edit a new replication as well as publish drafts.

Next step

To show the setting options, select the site.

> **Information**

> For more detailed information on setting up a replication network, refer to the additional sections in the chapter *Configuration*.

Draft mode (overview)

In draft mode, you can configure and edit a new replication as well as publish drafts.

> **Please note**
>
> Bi-directional replication is possible between two sites. Serial replication between multiple sites in a loop is not possible.

1 Viewer pane: Here, you will see the sites and connections in the replication network. Your local site is marked with a pin icon. The site currently selected that you are changing the settings for is shown in yellow.

2 Back: If you don't want to publish a draft yet or want to continue editing it at a later point in time, exit draft mode by clicking the back arrow. The draft is saved. You return to the overview page.

3 Selection: If this button is enabled, you can select a site with the mouse in the viewer pane and move it. A dialog box appears on the right for each selected site where you can enter the data for the site.

4 New site/new connection: If this button is enabled, you can add new sites and connections.

To add a new site, click anywhere in the viewer pane.
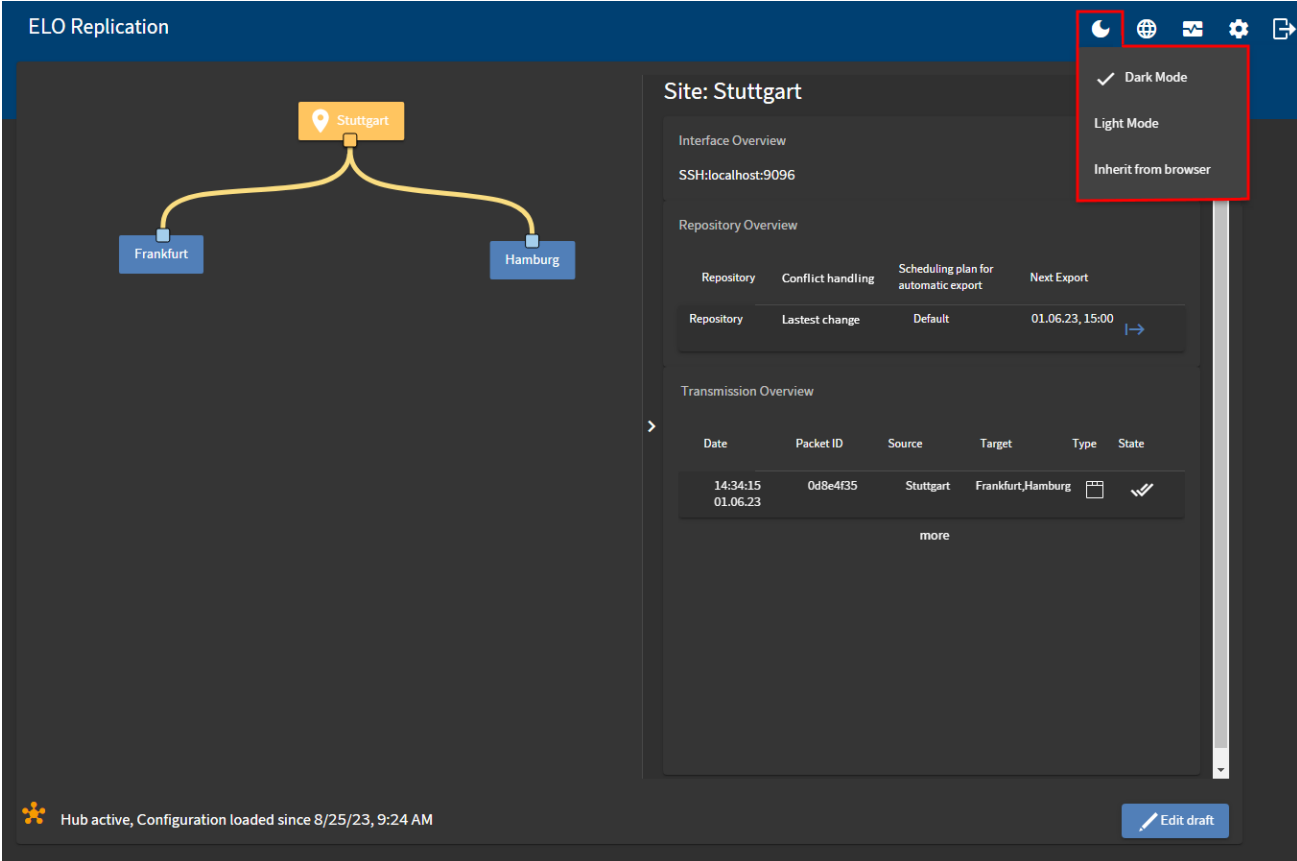
Before you add a connection, you need to set up an interface in the site settings. For more information, refer to the Site settings chapter. To add a new connection, place the mouse on a site. The connection icon appears. Press and hold down the left mouse button and drag the connection to another site.

5 Delete: If this button is enabled, you can delete sites and connections.

6 Scheduler for automatic transfer: This button opens the *scheduler for automatic transfer*. For more information, refer to the Scheduler for automatic transfer chapter.

7 Site settings: If you select a site in the viewer pane, the setting options for the site appear here. Your entries are saved automatically as soon as you leave an input field. For more information on these settings, refer to the chapters Site settings and Scheduler for automatic transfer.

8 Dark mode/light mode: You can choose between a dark or light screen background or use the browser default.

9 Language: Select the display language here You can choose between German and English.

10 Status information: This area contains information about monitoring the replication process. The package processing status is displayed as well as the number of data packages currently in the processing queue. The status when forwarding to the ELO Indexserver and to other sites is also displayed. This includes the number of data packages in the transmission queue.

ELO Replication

← Information   ⧉ http monitoring endpoint

Hub Status: RUNNING

Package processing

Status	Queued	Last Update at	Last result
OK	0	15:12 12.01.2024	OK

Redircrects
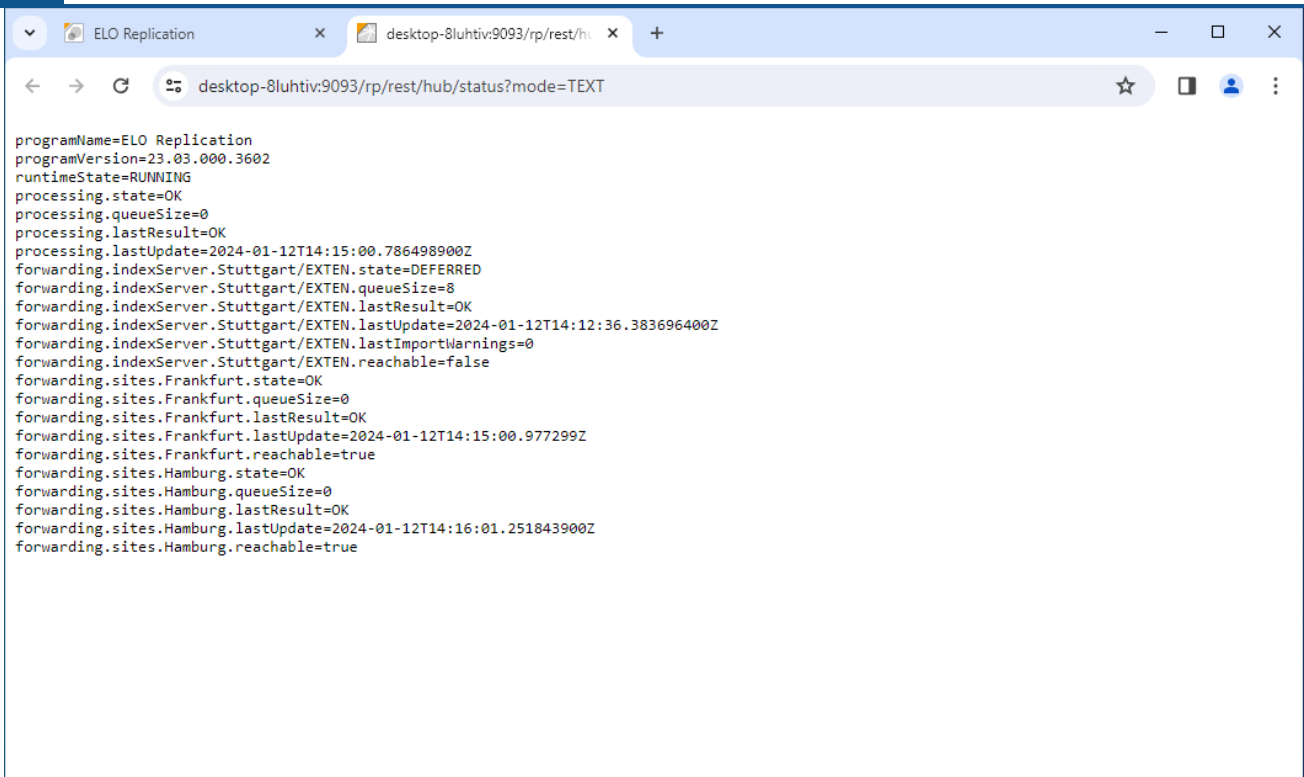
Indexserver

Repository	Status	Queued	Last Update at	Last result	Available	Import warnings
Stuttgart/EXTEN	DEFERRED	6	15:12 12.01.2024	OK	false	0

To other sites

To	Status	Queued	Last Update at	Last result	Available
Frankfurt	OK	0	15:12 12.01.2024	OK	true
Hamburg	OK	0	15:12 12.01.2024	OK	true

The link to *http monitoring endpoint* contains machine-readable information about the status page. This can be used for automatic monitoring.

```
programName=ELO Replication
programVersion=23.03.000.3602
runtimeState=RUNNING
processing.state=OK
processing.queueSize=0
processing.lastResult=OK
processing.lastUpdate=2024-01-12T14:15:00.786498900Z
forwarding.indexServer.Stuttgart/EXTEN.state=DEFERRED
forwarding.indexServer.Stuttgart/EXTEN.queueSize=8
forwarding.indexServer.Stuttgart/EXTEN.lastResult=OK
forwarding.indexServer.Stuttgart/EXTEN.lastUpdate=2024-01-12T14:12:36.383696400Z
forwarding.indexServer.Stuttgart/EXTEN.lastImportWarnings=0
forwarding.indexServer.Stuttgart/EXTEN.reachable=false
forwarding.sites.Frankfurt.state=OK
forwarding.sites.Frankfurt.queueSize=0
forwarding.sites.Frankfurt.lastResult=OK
forwarding.sites.Frankfurt.lastUpdate=2024-01-12T14:15:00.977299Z
forwarding.sites.Frankfurt.reachable=true
forwarding.sites.Hamburg.state=OK
forwarding.sites.Hamburg.queueSize=0
forwarding.sites.Hamburg.lastResult=OK
forwarding.sites.Hamburg.lastUpdate=2024-01-12T14:16:01.251843900Z
forwarding.sites.Hamburg.reachable=true
```

11 Settings: This contains the version number and runtime details for ELO Replication. You can also select a period of time after which the saved transfer protocol files from the *TransmissionDetailsHistory* folder in the ELO directory.
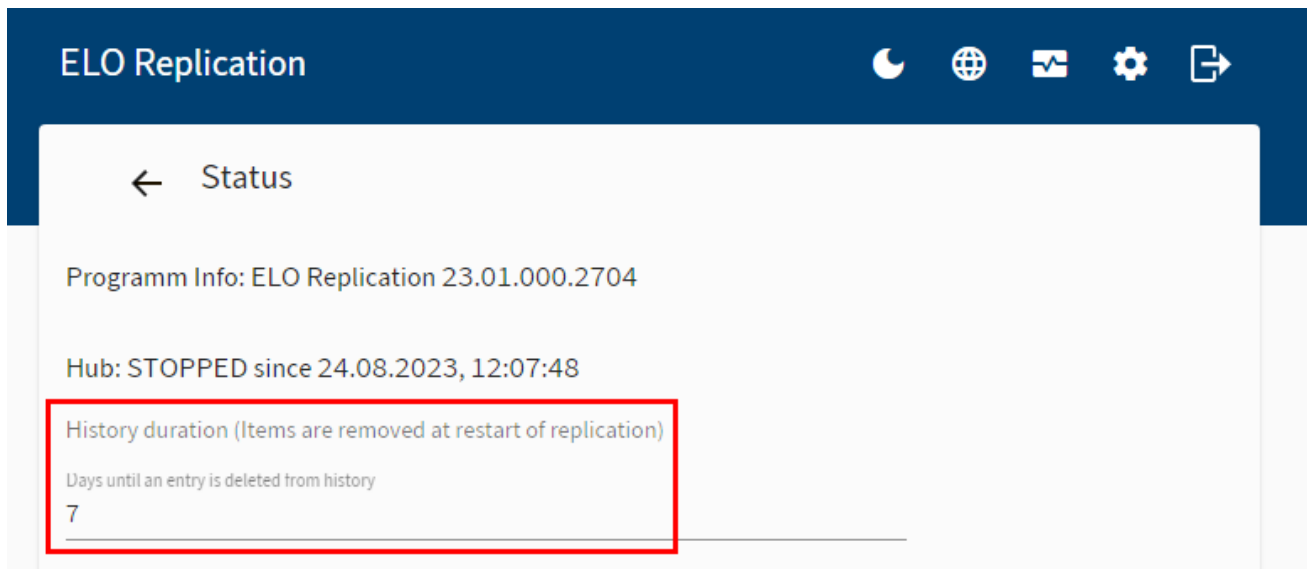
12 Log off: Select this button to log off of ELO Replication.

13 Publish: To enable a replication, select *Publish*. For more information, refer to the chapter Publish replication.

14 Discard: To delete a draft or discard your changes, select *Discard*.

Site settings

Overview

In the configuration, you can add sites. A dialog box appears for each site where you can enter the data for the site. Your entries are saved automatically as soon as you leave an input field.

You can enter the configuration settings for all participating sites from your site. A configuration file with your settings is created automatically for each site. Download the configuration file by selecting *Download configuration* on the overview page. Send the configuration file to the relevant site. There, the administrator logs on to ELO Replication and uploads the configuration file. You can find more information in the section Create/insert a configuration for other sites.

Name: Enter the name of the site.

Interfaces

Click the plus icon to add an interface.

Schema: Select SSH or BundleExchange.

Once you have made a selection, a rectangle appears over the site in the viewer pane. You can move the rectangle with the mouse to connect two sites. To add a new connection, enable the *New connection* button and place the mouse on a site. The connection icon appears. Press and hold down the left mouse button and drag the connection to a rectangle of another site. You can only connect interfaces of the same schema. You can only place one rectangle per schema for a single site. This rectangle can be used to add connections to multiple sites. To delete a rectangle, enable the *Delete* button and click on the rectangle.

SSH schema

IP address or server name: Enter the IP address or the name of the server ELO Replication is installed on.

Port: Enter the port for establishing the connection.

> **Please note**
>
> Do not change the firewall rules created by the ELO Server Setup in Microsoft Windows by adding the ELO Replication port. The ELO Server Setup may overwrite the rule in the future. You have to create a new rule for ELO Replication in the firewall settings.
>
> The port must be enabled for ELO Replication.

BundleExchange schema
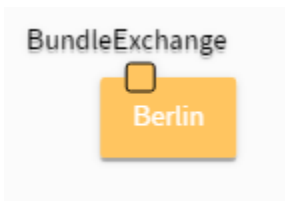
BundleExchange enables data transmission between sites that do not share a network connection. In the configuration, specify a source folder and a target folder for file transfer. Enter the full folder path, e.g. *C:\Replication\Target*. ELO Replication writes files to the target folder in rBundle format for transferring to other sites. The target folder is monitored by ELO Replication. All rBundle files in this folder are processed as incoming file transfers. Data is transferred between the source and target folders using a third-party application that you must install yourself.

> **Information**
>
> You can add additional options for transferring data by creating and integrating plug-ins and by selecting them from the drop-down menu.

Repositories

Click the plus icon to add repositories. A replication set is automatically created for each repository, which you can assign to an entry in the ELO Java Client or the ELO Web Client.

Stuttgart

Interfaces

Repositorys

Repository   Indexserver URL
https://desktop-8luhtiv:9093/ix-

Default

Conflict handling
Latest change

Name: Enter the name of the repository.

Indexserver URL: Enter the URL to the ELO Indexserver.

Scheduler for automatic transfer: Define when to export data from the repository. In the drop-down menu you will find the preconfigured intervals *Default* and *No automatic export*, as well as all intervals you have created. To create new intervals or change existing ones, select the pencil icon. For detailed information, refer to the Scheduler for automatic transfer chapter.

> **Information**
>
> The *scheduler for automatic data transfer* allows you to control the export of data. If you do not want to export data from the selected repository, select the *No automatic import* schedule. In this case, there is no scheduled data export from the repository; instead, data will only be automatically imported into the repository. If necessary, you can trigger data export manually.

Conflict handling: To prevent conflicts during the replication process, you can choose from three different options:

- *Retain local*: If changes are made at both sites, the local changes are retained and changes imported during the replication process are discarded.
-

Overwrite local: If changes are made at both sites, the local changes are discarded and changes imported during the replication process are applied.

- *Latest change*: This is the default setting. If changes are made at both sites, the most recent changes are always applied.
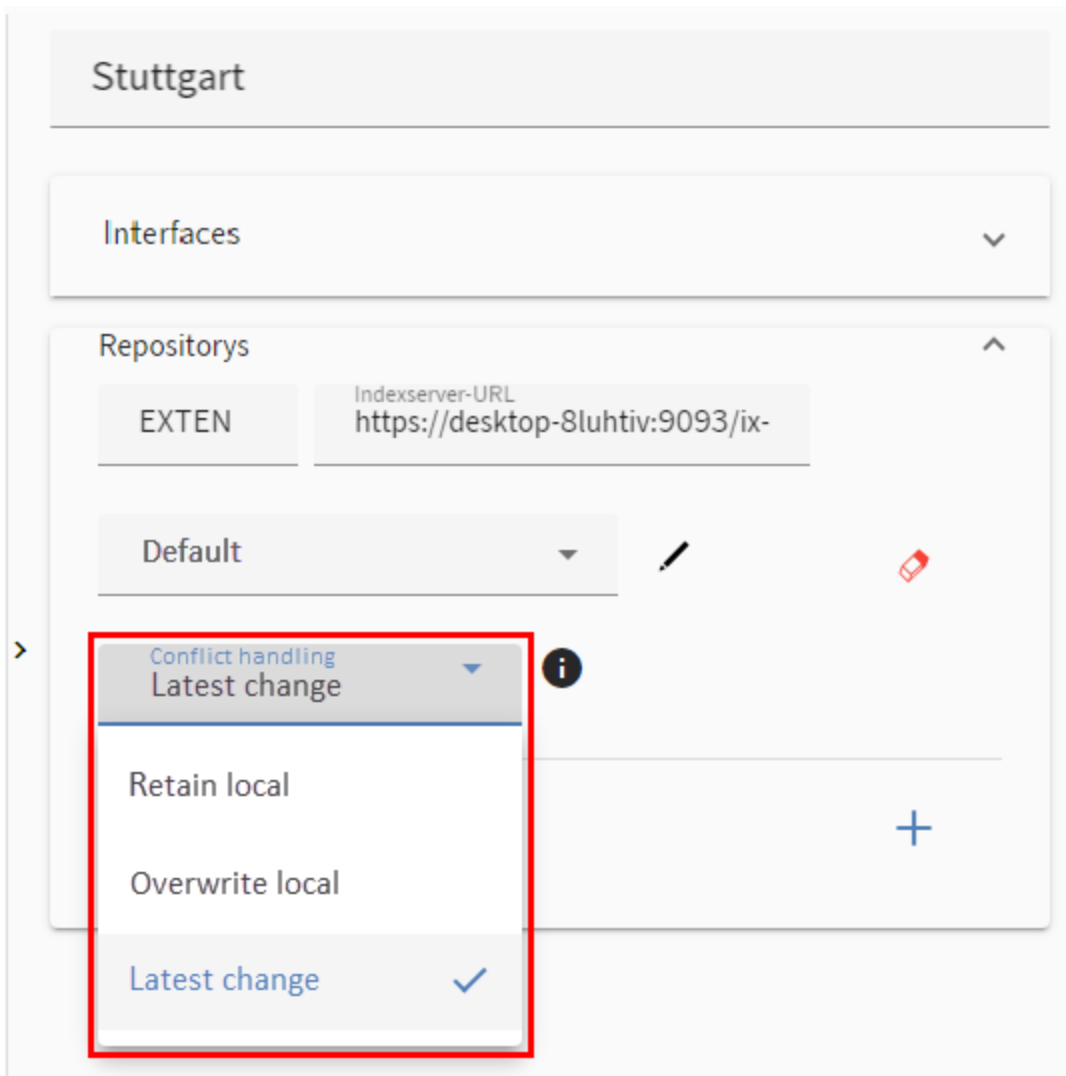
Scheduler for automatic transfer

The *scheduler for automatic transfer* allows you to define when to export data from the repository. In the drop-down menu you will find the preconfigured intervals *Default* and *No automatic export*, as well as all intervals you have created.

> **Please note**
>
> The specified interval time refers to the local time of the respective server. All servers hosting repositories involved in replication must be in sync in terms of time for the comparison via timestamp to work.

> **Please note**
>
> The scheduler calculates in hours. For example, if you enter a time from 8:00 a.m. to 6:00 p.m., replication begins at 8:00 a.m. and ends at 6:59 p.m.

'Default' interval

In the *Default* interval, data is exported from the repository: every 60 minutes, from 8 a.m. to 6 p.m., from Monday to Friday.

'No automatic export' interval

The *scheduler for automatic data transfer* allows you to control the export of data. If you do not want to export data from the selected repository, select the *No automatic import* schedule. In this case, there is no scheduled data export from the repository; instead, data will only be automatically imported into the repository. If necessary, you can trigger data export manually.

'Scheduler for automatic transfer' menu

To create new intervals or change existing ones, select the pencil icon.

The *Scheduler for automatic transfer* menu opens. Clicking the plus icon adds a new schedule. Clicking the eraser icon deletes a schedule.

The input options *Simple* and *Expert* are available for the *scheduler for automatic transfer*. With the input option *Simple*, you can choose from default intervals and customize them. With the input option *Expert*, you can control the *scheduler for automatic transfer* more precisely using variables.

Only one of the two options can be enabled. To use the expert input option, go to the *Expert* tab and check the *Use expert mode* box.

'Simple' input option

Add a new schedule with the plus icon.

Simple	Expert

Repeat every	Limit from	to	Between	and
15 minutes ▼	08:00 ▼	18:00 ▼	Monday ▼	Friday ▼

＋

In the *Repeat every* field, set an interval for exporting data.

In the fields *Limit from ... to ...*, you specify during which hours replication should take place. In the example above, replication takes place in the morning at 8:00, 8:15, 8:30, and 8:45.

In the fields *Between ... and ...*, set the days of the week to repeat replication.

'Expert' input option

With the input option *Expert*, you can control the *scheduler for automatic transfer* more precisely using variables.

Add a new schedule with the plus icon.

Simple	Expert

☑ Use expert mode

Minute	Hour	Day of the month	Weekday	Month
*	8-18	?	2-6	*

＋

Syntax

Field	Input
Minute	0-59
Hour	0-23
Day of the month	1-31
Weekday	1-7 (1 is Sunday)

Field	Input
Month	1-12

> **Please note**
>
> You can set the days in two fields: *Day of the month* or *Weekday*. To avoid competing entries, set the days in either the *Day of the month* or *Weekday* field. In the other field, enter the variable *?*. The field with the variable *?* is not run.

Variables

Variable	Meaning
*	Always run (every...)
*/n	Run every n
n,x,y	Run at/on n, x, and y
n-x	Run from n to x (x included)
?	Field is not run; only applies to *Day of the month* and *Weekday*

The following presents two examples of expert scheduling:

Example 1

In this example, replication takes place:

- Every 15 minutes
- From 6:00 a.m. to 8:45 p.m.
- From Sunday to Thursday
- Every month

The field *Day of the month* is not evaluated as it contains the variable *?*.

Example 2

In this example, replication takes place:

- In minute 0 and 30
- From 8 a.m.
- From the 1st to the 10th day of the month
- In February

The field *Weekday* is not evaluated as it contains the variable *?*.

Time display of the next data export

After creating the schedule and publishing the replication, you can see when the time of the next data export is scheduled in the overview page of each site.

Publish replication

To enable a replication, you have to publish your draft. In draft mode, select *Publish*.

Please note

During initial configuration, you will have to publish the replication at all participating sites. In case of later changes, you only have to publish the new draft at one site. The changes are automatically applied to all sites.

If you add a new site for an active replication, you will have to upload the configuration file at the new site and publish the replication there.

Please note

You cannot publish a draft until you have completed all entries. You can check this with the status message that is displayed at the bottom of the window.

Changes: Correct: Your entries are complete. You can publish the draft.

Changes: Incomplete: You still need to make entries. An error message is displayed in the corresponding section and you can select *Navigate to error* to go straight to the incomplete input field.

Create/insert a configuration for other sites

Information

You can enter the configuration settings for all participating sites from your site. A configuration file with your settings is created automatically for each site. Download the configuration file by selecting *Download configuration* on the overview page. Send the configuration file to the relevant site. There, the administrator logs on to ELO Replication and uploads the configuration file.

> **Please note**
>
> During initial configuration, you will have to upload a configuration file at all sites. In case of later changes, you only have to publish the new draft at one site. The changes are automatically applied to all sites. For more information, refer to the chapter Edit configuration.

If you add a new site for an active replication, you will have to upload the configuration file at the new site and publish the replication there.

Site 1: Download configuration file

Requirement: You have completed site configuration and published the replication at site 1.

1. On the overview page, select an external site you have already configured.
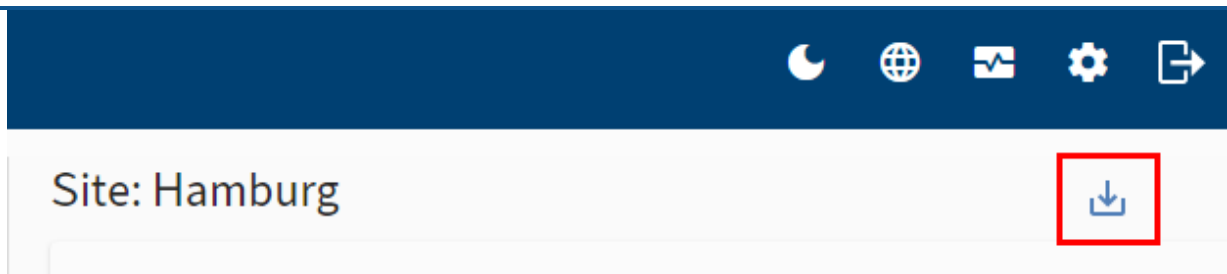
2. Select the *Load configuration* button.

The configuration is downloaded as a JSON file. The site name is used as the file name.

3. Send the configuration file to the administrator at the other site.

Site 2: Upload configuration file

1. Log on to ELO Replication with your Apache Tomcat administrator account.

> **Information**
>
> You do not have to enter a site name. The site name from the configuration file is automatically applied.
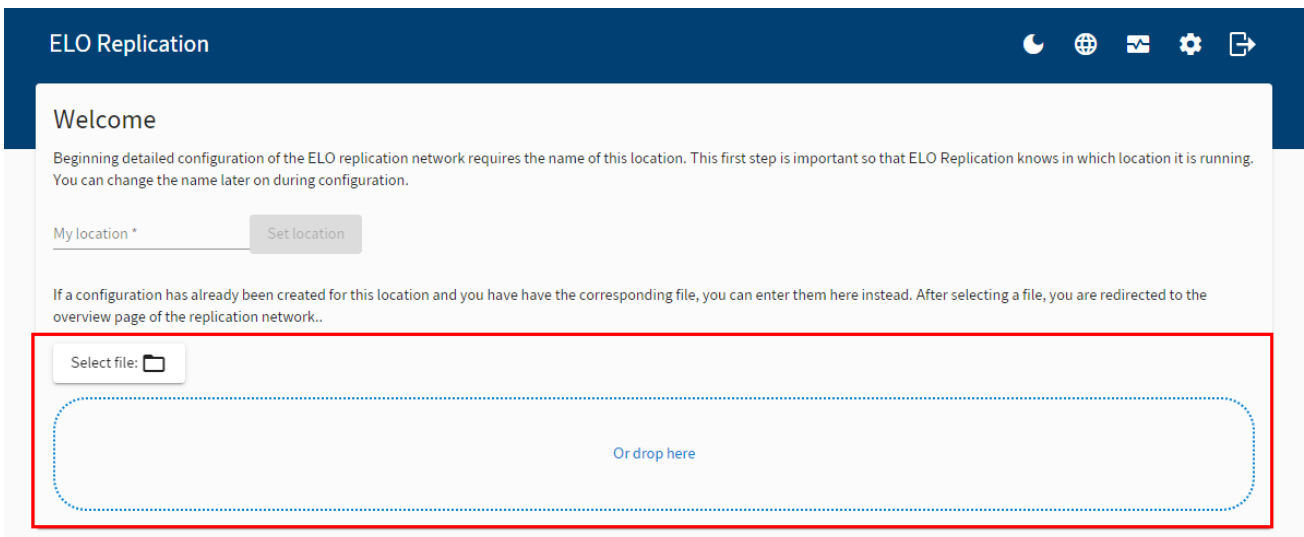
1. Insert the configuration file.

2. Confirm the site.
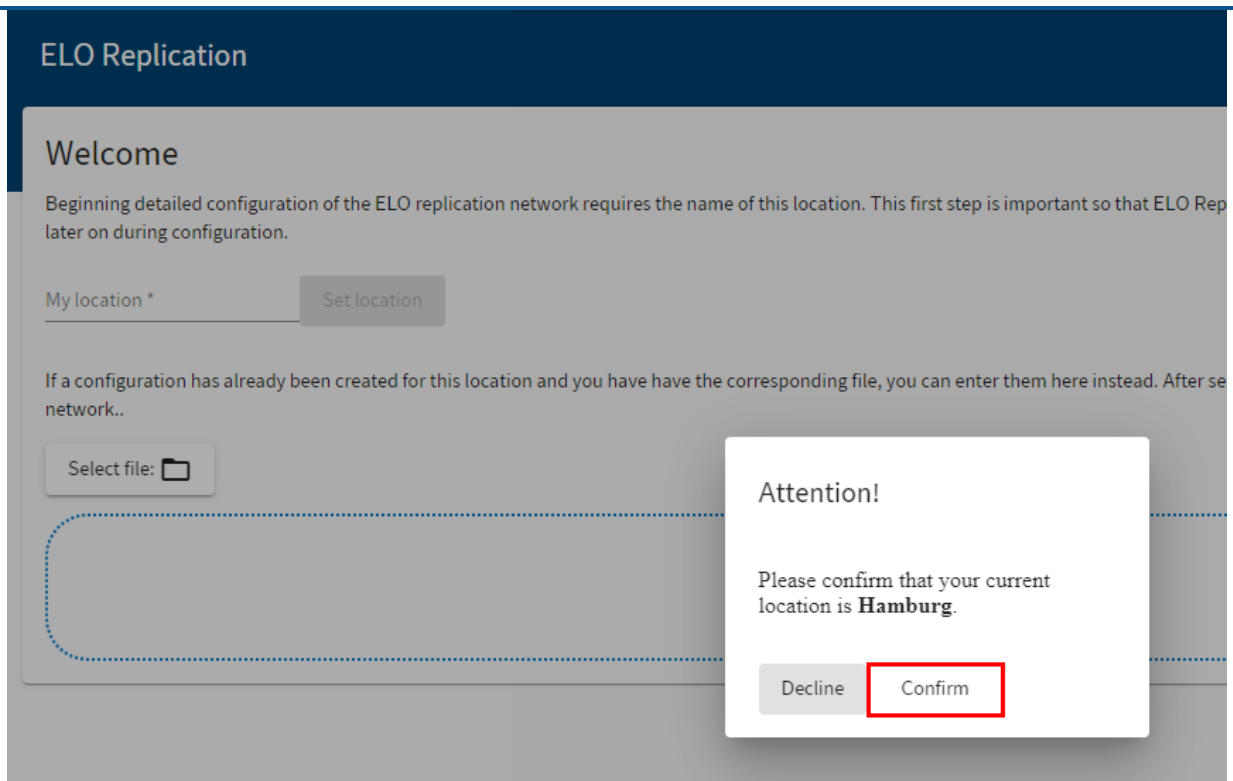
Once you have confirmed the configuration file, draft mode opens. You will see the sites in the replication network. The settings from the configuration file were applied.

3. To enable the replication for your site, select *Publish*.

Result

The replication sets are created. The SSH port is started.

Next step

Once you have published the replication at the participating sites, you have to assign the individual entries replication sets. This means you determine what entries are replicated and where. For more detailed information, refer to the chapters Assign replication sets and Replication set inheritance.

Replication set inheritance

Replication sets are automatically created in the configuration when you add a new repository. Each replication set stands for one repository. The replication set is a characteristic that you assign to individual repositories in order to replicate them in another repository. The individual entries (folders, documents) to be replicated are selected in the ELO Java Client or ELO Web Client using the *Assign replication sets* function.

Information

You only have to assign replication sets at one site. The assigned replication sets are also configured at the other sites through replication.

Default: Inheritance to parent folder

Replication sets are inherited upward to parent folders in the repository structure.

Example: If you assign the folder *C* including its child entries a replication set, the entries marked in yellow are replicated. The replication set is inherited to the parent folders.
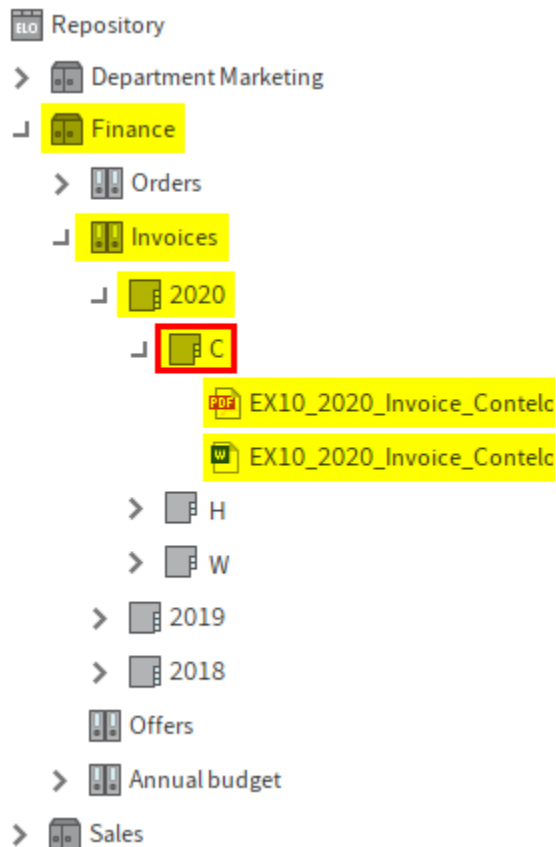
'Start point for replication' option

To prevent the replication sets from being inherited to the parent folders, select the *Start point for replication* option in the *Metadata* dialog box of the entry you want to replicate. The *Start point for replication* option allows you to replicate parts of repositories that are not filed to identical repository structures.

> **Please note**
>
> In the ELO Java Client, this function is only available for folders.

> **Please note**

Enable the *Start point for replication* option before assigning the entry a replication set.

Example: You have enabled the *Start point for replication* option for the *C* folder. If you then assign the *C* folder including its child entries a replicated sets, only the folder *C* and its child entries are replicated (marked yellow). The replication set is not inherited to the parent folders.
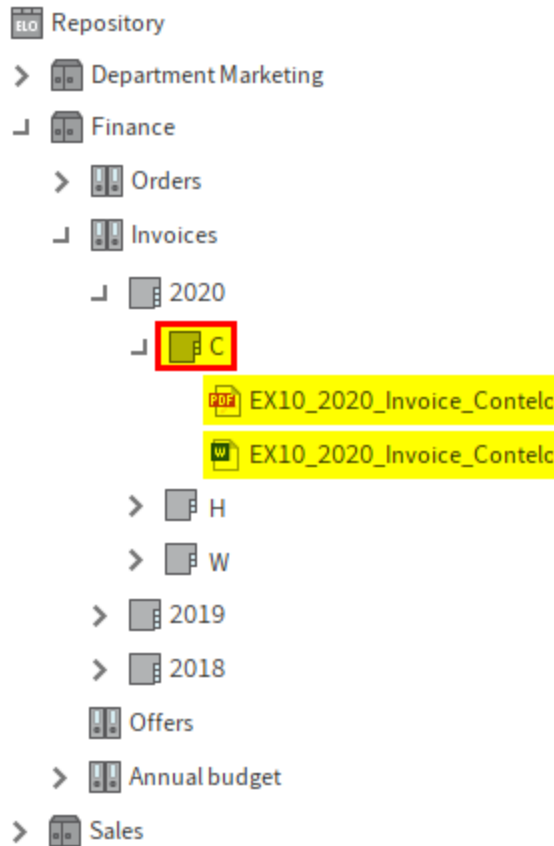
During the next replication process, the *C* folder is filed to the following path in the target repository: *// Administration // Replication Base*. From there, the administrator moves the folder to the desired site in the target repository. This only has to be performed during the initial replication process. The replication remembers the repository structure.

Assign replication sets

Information

Replication sets are automatically created in the configuration when you add a new repository. Each replication set stands for one repository. The replication set is a characteristic that you assign to individual repositories in order to replicate them in another repository. The individual entries (folders, documents) to be replicated are selected in the ELO Java Client or ELO Web Client using the *Assign replication sets* function. This means you select a repository the entries are replicated to. The name of a replication set in the ELO client is made up of the site name and the name of the repository.

> **Information**
>
> The following uses screenshots of the ELO Java Client.

> **Information**

You only have to assign replication sets at one site. The assigned replication sets are configured at the other sites through replication.
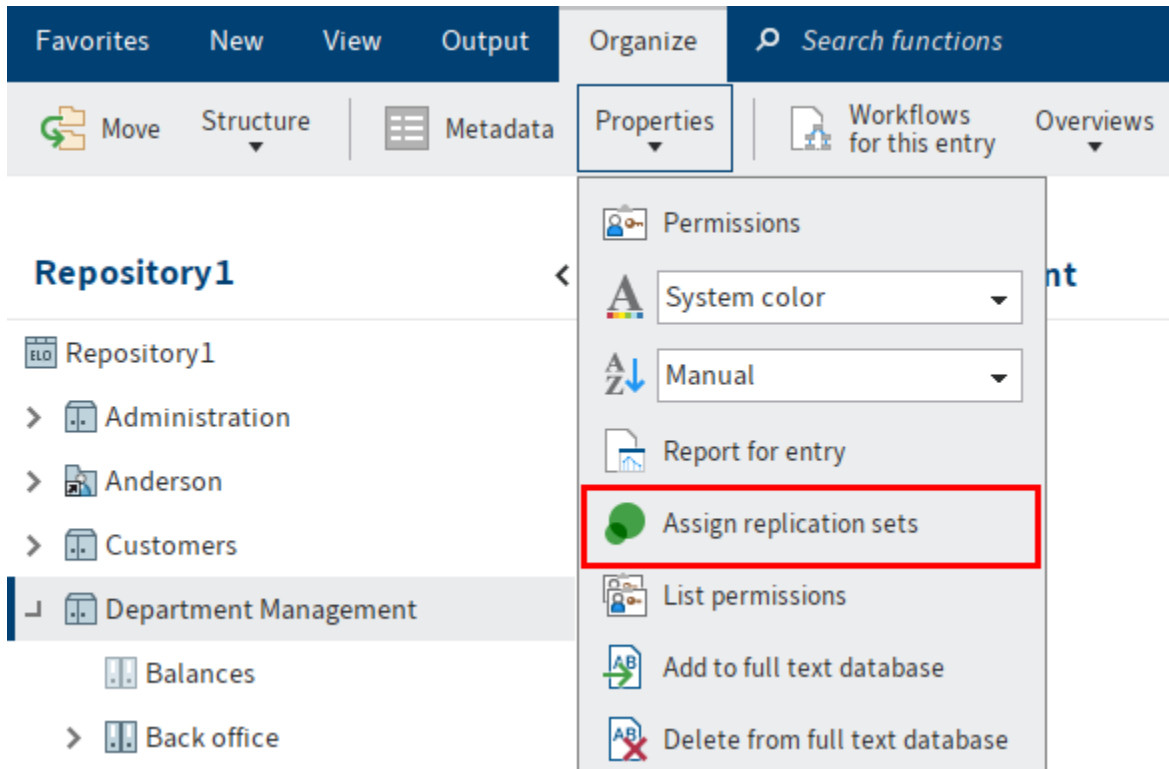
You will find the *Assign replication sets* functions in *Ribbon > Organize > Properties*.

Requirements

You need the right *Assign replication sets*. This right is enabled in the ELO Administration Console.

Please note

You should log which entries you assign which replication sets.

Please note

Documents with the status *Version control disabled* cannot be replicated. You will find the document status at *Ribbon > Organize > Metadata > 'Options' tab*.

Step by step

1. In the ELO client, select the entry you want to assign a replication set to.

2. On the ribbon, select *Organize > Properties > Assign replication sets*.

   The *Assign replication sets* dialog box opens.

3.

Select *Add*.
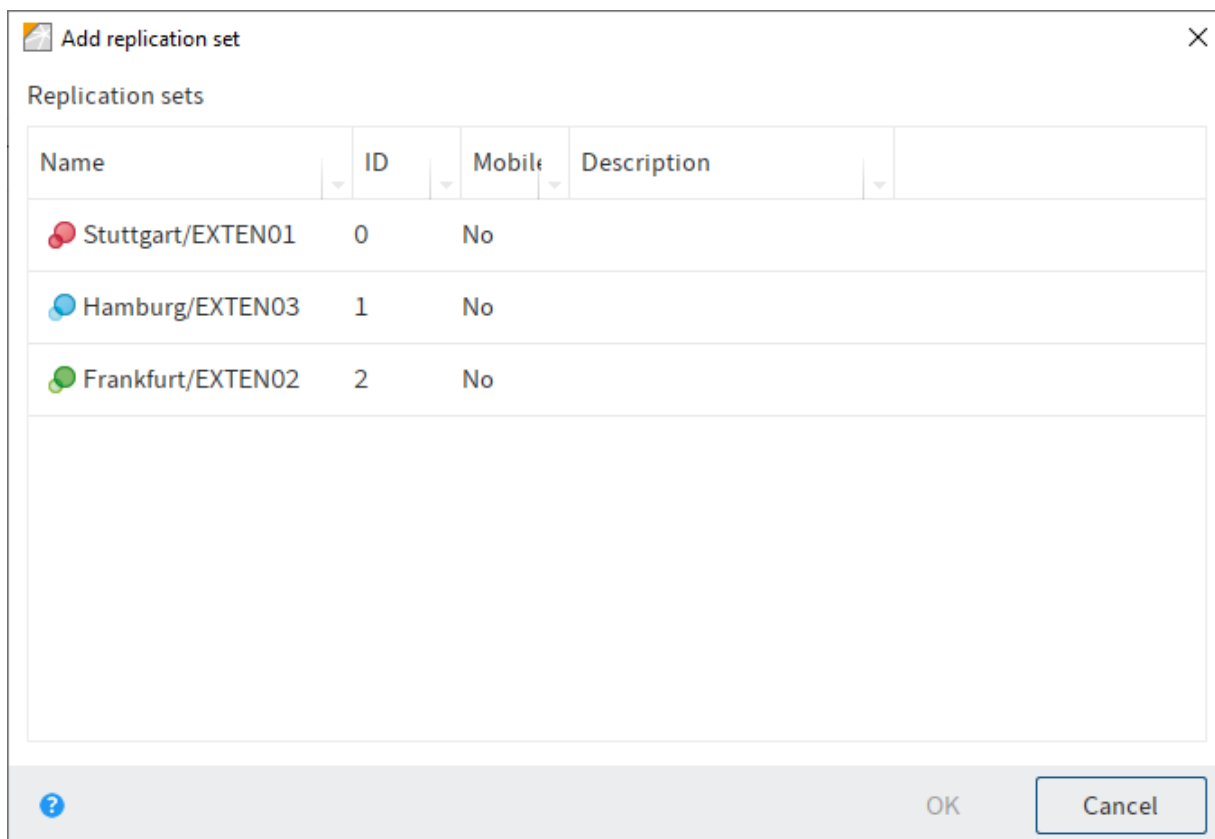
The *Add replication set* dialog box opens.

You will see all the replication sets defined in the configuration. The name of the replication set is made up of the site name and the name of the repository.

> **Information**
>
> The local replication set, that is the replication set where the entries are located, is assigned automatically. You do not have to add it manually.

4. Select one or more replication sets and confirm with *OK*.
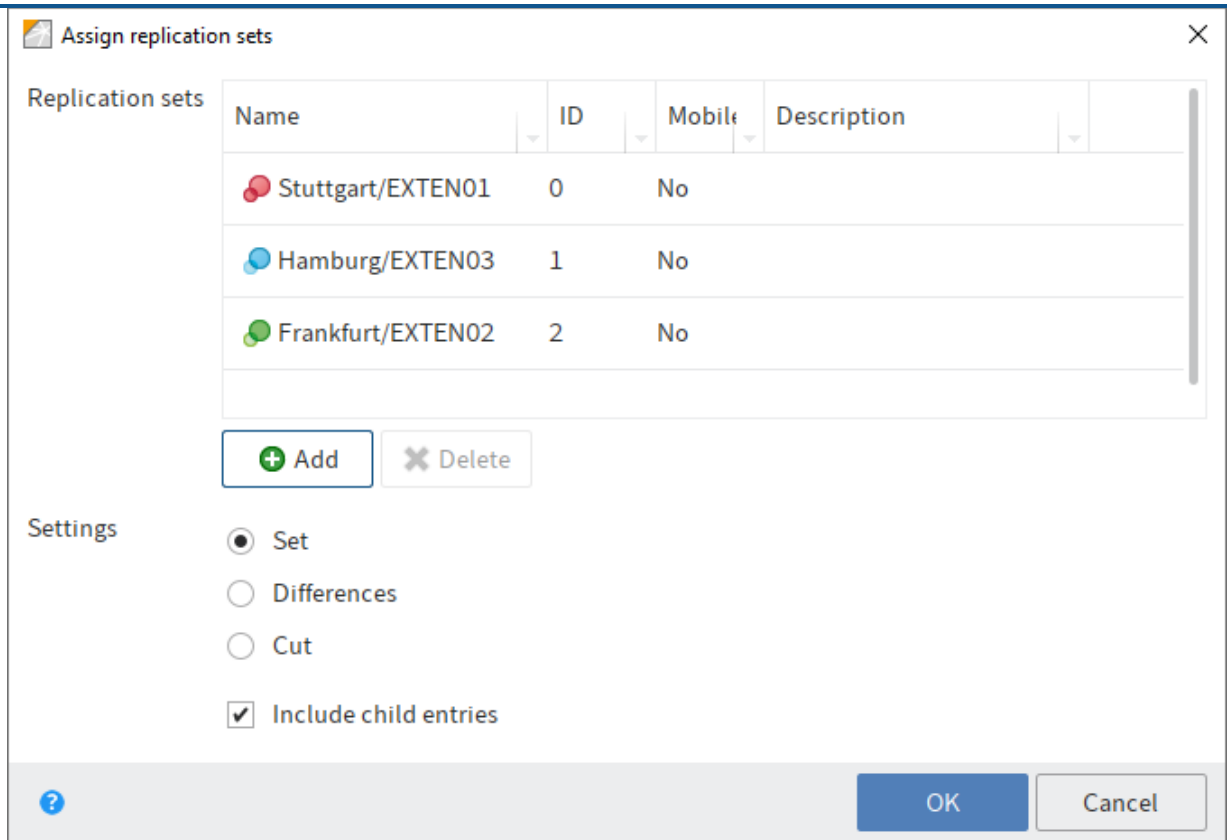
The replication sets appear in the *Assign replication sets* dialog box.

> **Please note**
>
> Replication sets are inherited upward to parent folders in the repository structure. To prevent inheritance, select the *Start point for replication* option in the *Metadata* dialog box of the entry you want to replicate. In the ELO Java Client, this function is only available for folders. You will find more information in the chapter Replication set inheritance.

The following options are available:

Including child entries: If this option is enabled, the replication set is also assigned to the child entries of the selected entry.

Set: All the listed replication sets are assigned to the entry.

Differences: This option should only be used in combination with the *Include child entries* option. Differences: If you removed replication sets from the list and/or added replication sets to the list, this change is passed on to the entries. Replication sets from the list that were not changed are not passed on. The replication sets of the entries that were not changed are retained.
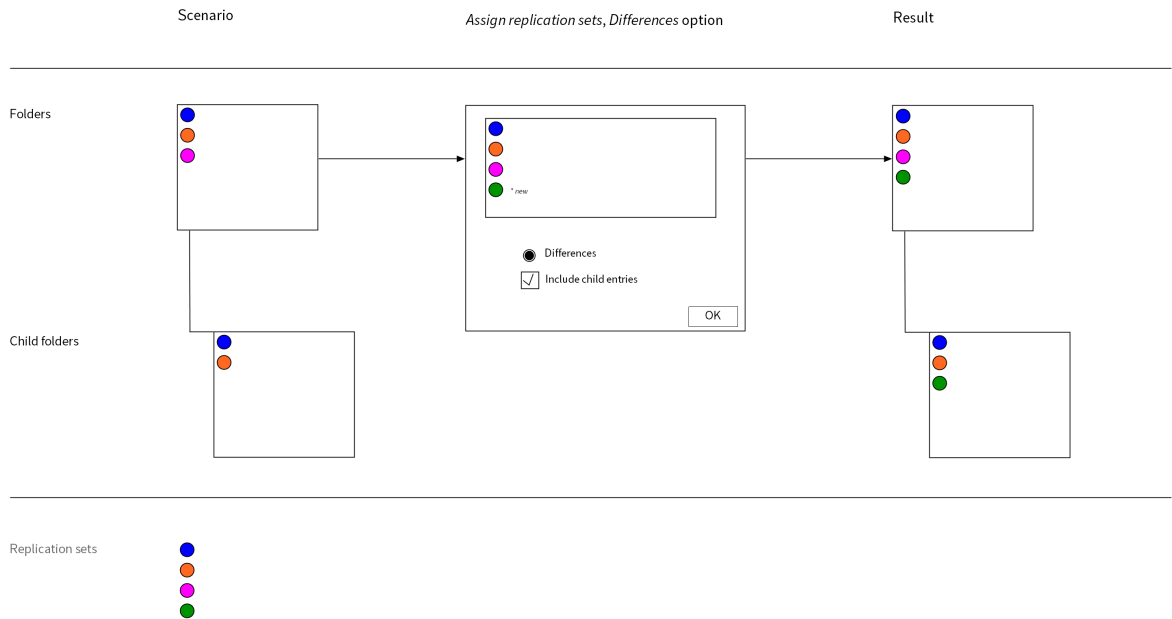
Scenario   *Assign replication sets*, *Differences* option   Result

Folders

Child folders

Replication sets

In the example above, a folder and its child folders have different replication sets. The *Assign replication sets* dialog box opens for the folder. The *Include child entries* option is enabled. The *Blue*, *Orange*, and *Pink* replication sets are already assigned to the folder. The *Green* replication set is added. The *Differences* option is selected. The *Green* replication set is added to the list again. This change is passed on to the folder and its child folders. Replication sets from the list that were not changed are not passed on. The *Blue*, *Orange*, and *Pink* replication sets are not assigned again. This means the child folder is not assigned the *Pink* replication set. The replication sets of the entries that were not changed are retained. Folders and child folders retain their original replication sets and are additionally assigned the *Green* replication set.

Cut: The replication sets are assigned to the entry. The replication sets that are not within this intersection are removed. This option should only be used in combination with the *Include child entries* option. Use the *Cut* option if, for example, it is unclear due to the set *Start point for replication* which replication sets are assigned to the child entries and you only want to inherit the replication sets selected in the dialog box. In contrast to the *Set* option, with the *Cut* option you cannot assign any new replication sets. With the *Cut* function, you can delete replication sets from child entries that are not contained in the parent entry.
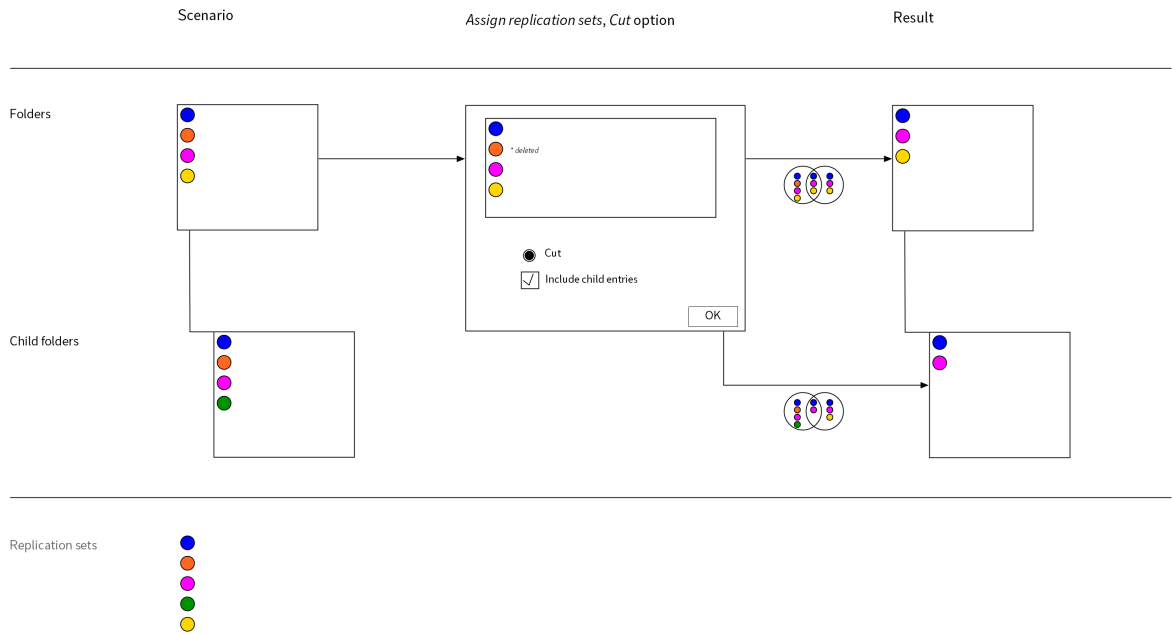
In the example above, a folder and its child folders have different replication sets. The *Assign replication sets* dialog box opens for the folder. The *Include child entries* option is enabled. The *Blue*, *Orange*, *Pink*, and *Yellow* replication sets are already assigned to the folder. The *Orange* replication set is deleted. The *Cut* option is selected. The replication sets are assigned. The replication sets that do not overlap between the initial status and the reassignment are removed. Only the *Orange* replication set is removed from the folder. With the child folder, the *Orange*, *Green*, and *Yellow* replication sets do not overlap between the initial status and the reassignment. For this reason, the child folder is not assigned the *Yellow* replication set. The *Orange* and *Green* replication sets are removed from the child folder.

5. Make the desired changes to the settings and confirm with *OK*.

Result

You have assigned the selected replication sets. During the next replication process, the selected entry will be replicated.

Outlook

You can check the transfer status on the ELO Replication overview page.

On the overview page, select a site. The transfer overview opens on the right. You will recognize a replication based on the repository icon in the *Type* column. Successful transfer is indicated by two check marks in the *Status* column. For more information, refer to the chapter [Monitoring](#).

Edit configuration

Information

You can edit the configuration of a published replication. During editing, the replication remains active. To apply your changes to the published replication, select *Publish* in draft mode.

> **Information**
>
> If you publish a new draft, the changes are automatically applied to all sites. You do not have to update the configuration manually at the different sites.

If you add a new site, you will have to upload the configuration file at the new site and publish the replication there. You can find more information in the section [Create/insert a configuration for other sites](#).

Step by step

1. On the overview page, select *Edit draft*.

   Draft mode opens.

   For more detailed information on the configuration settings, refer to the [Draft mode (overview)](#), [Site settings](#), and [Scheduler for automatic transfer](#) chapters.
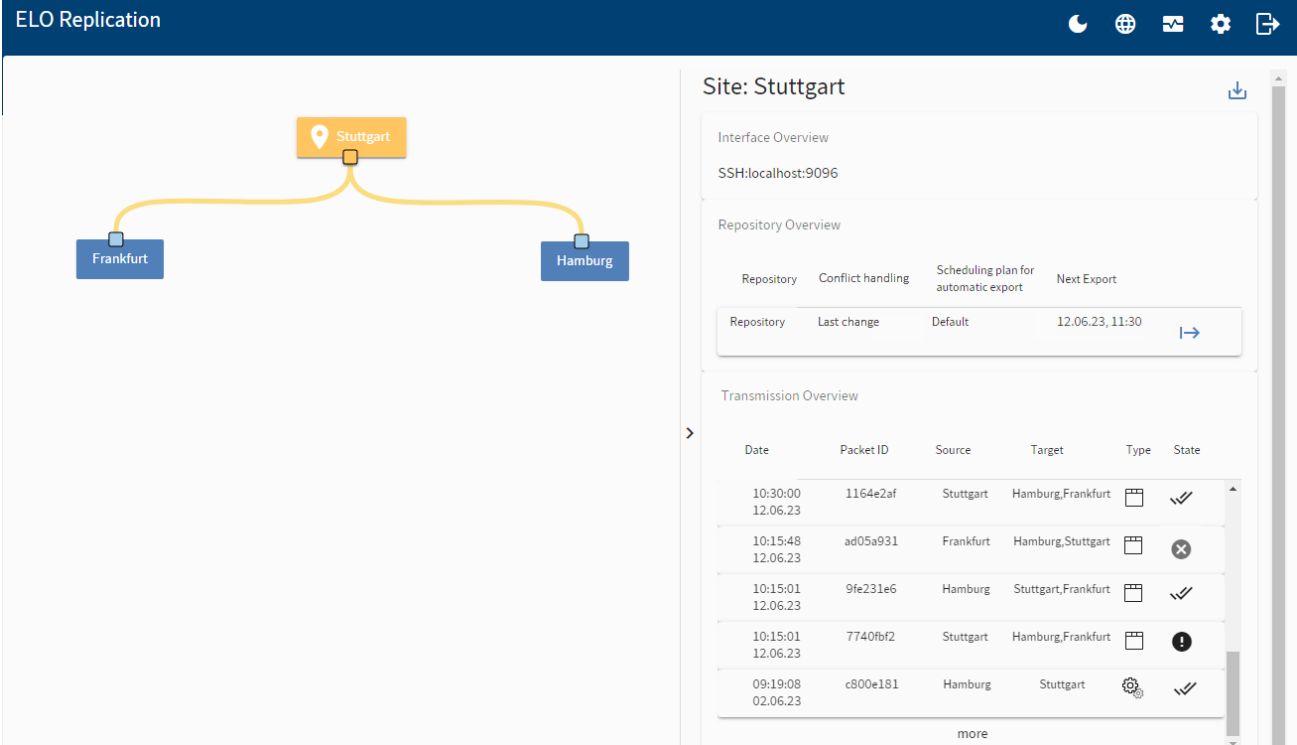
2. Make the desired changes. Your changes are saved automatically as soon as you leave an input field.

3. To end editing, you have the following options:

   - Publish changes: To apply your changes to the published replication, select *Publish*.
   - Discard changes: To delete your changes, select *Discard*. The published replication is not changed.
   - Save draft: If you don't want to publish the new draft yet or want to continue editing it at a later point in time, exit draft mode by clicking the back arrow or log off. Your draft is saved automatically. The published replication is not changed. To continue editing, select *Edit draft* on the overview page.

Result

You have edited the replication.

Outlook

If you've published a new draft, the changes are automatically transferred to all sites. You can check the transfer status on the ELO Replication overview page.

On the overview page, select a site. The transfer overview opens on the right. You will recognize that a new draft is being transferred based on the gear icon in the *Type* column. Successful transfer is indicated by two check marks in the *Status* column. For more information, refer to the chapter Monitoring.

Pause transmission

Information

You can pause transmission during a replication process. In this case, the Scheduler for automatic transfer is temporarily disabled. During this time, you can perform actions such as maintenance work on the ELO system without having to change the scheduler configuration in ELO Replication.

Step by step

1. Select the *Pause export* button. This pauses the export of data from the repository.

2. Select the *Continue export* button.

Result

The scheduler for automatic transfer is enabled again and the export of data is resumed.

Monitoring

You can monitor an active replication using the transfer overview. To open the transfer overview, select a site on the overview page.

You can manually start an export from the local repository for test purposes. The export starts immediately, regardless of the configured schedule. Select *Start export* (1).

In the transfer overview, you see a list of the jobs. Selecting the entry opens the detailed view.

The following job types are shown in the transfer overview:

- Replication: You will recognize a replication job based on the repository icon (2).
- Configuration: If you have changed the configuration of an active replication and want to publish your changes, the changes are automatically transferred to all sites. You will recognize the transfer of a new configuration based on the gear icons (3).

The following statuses are shown in the transfer overview:

Status	Description
NEW	New transfer, not started
↦	Transfer started, export running
✓ ✓	Transfer successful: • One check mark per location in the detailed view • Two check marks in the *Status* column when all locations successful
⊙	Transfer failed: Network or software error, e.g. network disconnected or insufficient memory
⊗	Transfer canceled: Caused by user, e.g. server booted down

Handling errors

The following section describes the behavior starting with ELO Indexserver version 23.

Error when exporting from the Indexserver: If an error occurs while exporting an item, e.g. because a document was deleted in the document path, a SORD linking to the item that caused the error is created in the repository under *Administration > Replication Base*. The error message is displayed under *Metadata > Extra text*.

Errors that prevent an import (to hub or Indexserver): These are errors such as a corrupt replication package, an unavailable site/Indexserver, or the transfer was aborted. You can recognize errors like these by the exclamation mark in the *Status* column of the transfer overview. Selecting an entry opens the detailed view. A question mark is displayed next to the status. Select the question mark icon to download an error message in JSON format. ELO Replication attempts to send the data set once every minute. The data set is transferred all over again, regardless of when the previous transfer was interrupted.

Transmission Overview

Date	Packet ID	Source	Target	Type	State
09:19:08 02.06.23	c800e181	Hamburg	Stuttgart		!

Received by Site Hamburg:

Step	Started at	Finished at	Status	Details
Receive	09:19:08 02.06.23	09:19:08 02.06.23	✓✓	

Transfer to Stuttgart:

Step	Started at	Finished at	Status	Details
Send	09:19:08 02.06.23		NEW	
Processing	11:31:32 12.06.23		!	
Processing	11:30:32 12.06.23	11:31:32 12.06.23	!	?
Processing	11:29:32 12.06.23	11:30:32 12.06.23	!	?

Errors that prevent an import of individual items (to Indexserver): These are errors such as if the item dependency is not found, or an item violates rulesets (name too long, user GUID already assigned). In this case, ELO Replication does not attempt to transfer the data again. The import is completed and the error is logged. A log icon appears in the *Details* column of the transfer overview. Use this icon to open the transfer protocol that lists the errors.

After selecting the protocol icon, the information of the transfer protocol is displayed in the program interface.