

Replication

Replicating data between locations and laptops





CAS Software AG

CAS-Weg 1 – 5

76131 Karlsruhe, Germany

+ 49 (0)721 9638 - 0

info@cas.de

www.cas-crm.com

Copyright

The information contained herein may be altered without prior notice. The names and data used in the examples are fictitious, except where otherwise stated. No part of this document may be reproduced or transmitted for any purposes whatsoever without the express written consent of CAS Software AG, irrespective of the manner or the means, electronic or mechanical, by which this occurs.

© 2009 - 2019 CAS Software AG. All rights reserved.

All trademarks are the property of their respective owners.

Disclaimer

The accuracy of the content is not guaranteed. Notification of errors would be appreciated.

February 2019

Contents

1	Introduction	5
1.1	Performance.....	6
1.2	Replication process.....	6
1.3	Main installation, subsidiary and laptops.....	8
1.4	Technical realization and connection costs.....	10
1.5	Replication licensing.....	11
1.5.1	Example for a cycle.....	12
1.5.2	Creating laptop licenses.....	13
1.6	Protecting data from outside the company network.....	13
1.6.1	Protecting data during transfer.....	13
1.6.2	Protecting data on laptops.....	14
1.6.3	Rights in CAS genesisWorld are replicated.....	15
2	Setting up the replication	15
2.1	Requirements and first steps.....	15
2.2	Special requirements for laptops.....	16
2.3	Steps and setup order.....	17
2.4	Administering application server profiles.....	17
2.5	Creating/editing application server profiles.....	19
2.6	Replication functions.....	21
3	Domain	24
3.1	Setting up the domain.....	24
3.1.1	New domain: Database and Document Archive window.....	26
3.1.2	New domain: Tasks window.....	29
3.2	Functions for domains.....	31
3.2.1	General tab.....	31
3.2.2	Settings tab.....	32
3.2.3	Maintenance tab.....	32
3.3	Receiving agent and schedule agent.....	33
3.3.1	Functions for agents.....	33
3.3.2	Defining settings for agents for laptops.....	34
3.4	Laptop automatics.....	34
3.5	Domain log window.....	35
3.6	Moving a domain to a new location.....	35
3.7	Fusing domains.....	37
3.8	User Management for more domains.....	38
3.8.1	Assigning a domain.....	39
3.8.2	Participants are assigned to domains.....	39
3.8.3	Users of other domains as participants.....	40

3.8.4	Participants in a domain connection and partial replication tasks	40
4	Domain connection	41
4.1	Setting up domain connections	41
4.2	General tab	42
4.3	Receive tab	42
4.4	Send tab	43
4.5	Settings tab	46
4.6	Administration tab	48
4.7	Authorization tab	51
4.8	Connection status window	52
5	Replication task	55
5.1	Creating replication tasks	57
5.2	User-defined task	59
5.3	Special tasks	60
5.4	Full replication task	61
5.5	Replication task properties	61
5.6	Which replication task is the right one?	63
5.7	Special considerations for full replication	64
5.8	Replicating data of a number of users	64
5.9	Deleting the domain connection or replication task	66
6	Replicating data	66
6.1	Concurrent replication	67
6.2	Replicating data from navigator views	68
6.3	Replicating certain data records	69
6.4	Replicating documents	69
6.4.1	Replicating the Document Archive	69
6.4.2	Replicating archive files	70
6.5	Resolving conflicts	71
7	Important information for administrators	72
7.1	Software updates	72
7.2	Code page and Unicode	72
7.3	Application server performance	72
7.4	Solving problems when replicating for the first time	73
7.5	Customizing the database	73
7.6	Service jobs in the Database area	74
7.7	Notification and action service	75
7.8	Checking mandatory fields	75
8	Important to know for certain functions	75

1 Introduction

Data replication enables you to synchronize CAS genesisWorld data between several locations, departments or laptops which you deploy outside the network.

You can also replicate individual changes to the database model made, for example, by CAS partners. This results, for example, in the creation of identical fields in the corresponding domain(s) and prevents data loss.

Data must thus be maintained only once, the connections between databases across different locations are kept and colleagues work with the same data stock.

You will need a license for replication.

Licenses for CAS genesisWorld are entered in the database. Other locations, departments or laptops outside the network access other databases. Licenses must also be entered in these databases. The administrator can create two laptop licenses for each CAS genesisWorld license in the Management Console, for example, a laptop license and a license for a home office workstation.

The administrator defines the replication rights and the rights to use certain functions. The right for the replication feature must also be enabled for the administrator.

If problems occur on the replication connection and need to be solved, the administrator has the option to be automatically notified via e-mail. This notification is sent by the notification service.

Integrated logs with status, warning and error messages notify users about the status of the connection. This makes it possible for the administrator to check activities after the replication has been finished.

Important information

If you set up the replication as administrator, please refer to our online help for administrators. The online help for the Desktop Client does not contain all relevant information.

Possible error messages and warnings in the log of the connection or the local domain as well as possible causes and solutions can be found in the online help pages of the Management Console and the Desktop Client, for example on the [Errors and warnings in the log](#) page.

When deploying the replication feature, other modules and integrations may be affected. You can look up possible effects at hilfe.cas.de at **Working with different modules** on the [Replication](#) page.

1.1 Performance

With regard to performance, you have to take into account several factors that affect each other. Please monitor and, if necessary, change the following influencing factors:

- Internet connection load
- Bandwidth of the Internet connection
- Deployed hardware for computers and network
- Number of users accessing the Web server
- Size of the replication packages

When deploying the Microsoft SQL Server Express Edition database system, you can specifically increase the performance.

- Deactivate the AUTO_CLOSE option for the database. This increases the database performance because the files of the database are not closed.

This option is only set to OFF by default for the Express Edition. The SQL command is ALTER DATABASE <name of database> SET AUTO_CLOSE OFF.

1.2 Replication process

The replication is illustrated in the figure below.

To start replicating data, two complete CAS genesisWorld installations including a database must exist – the so-called **Domains**. One domain can be, for example, the installation in the head office, the second domain can be a subsidiary or a laptop installation. Data which is relevant for the system already exists in a database after a subsidiary or a laptop have been set up. Alternative options are described in a separate section further below.

For each domain, **Domain connections** for the domain are set up. Data is exchanged via these domain connections.

A replication includes 2 processes: sending and receiving. Data is synchronized in both domains.

The figure below illustrates the start of the sending process in domain A and the process includes the steps 1 - 3:

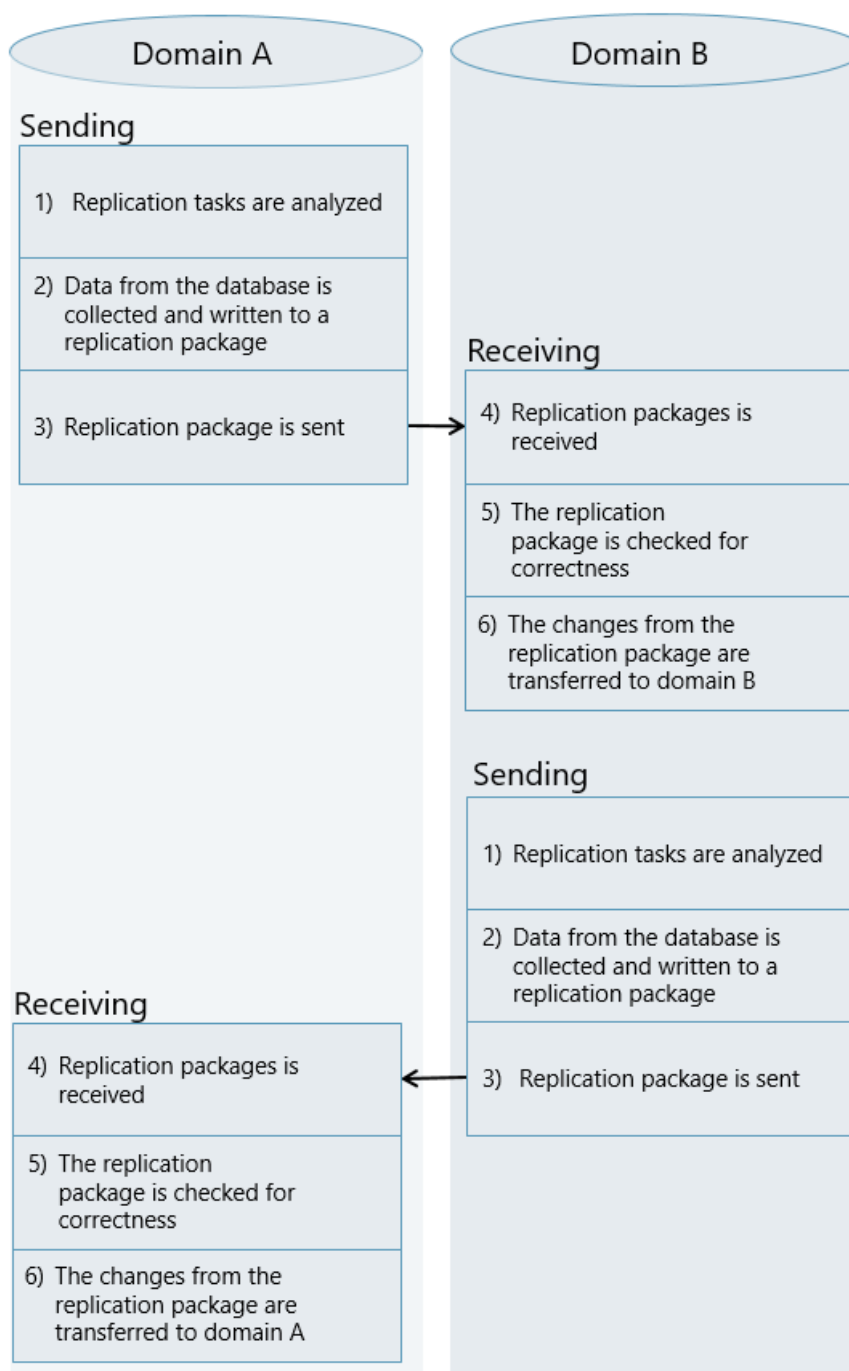
- ① A user defines which data he or she wants to synchronize with **Replication tasks**. This can be, for example, appointments or tasks within a certain time period, or certain documents.

Additionally, changes to system data are monitored and prepared to be transferred to the corresponding domain. System data contains all information about users, groups, rights settings, and so on.

With replication tasks, CAS genesisWorld checks which data is to be replicated. This is done via the **Change log** tab in data records: Which data selected in the replication tasks has been newly created or changed in domain A since the last time data has been successfully replicated?

Changes that were not documented in the **Change log** are not replicated. If you, for example, directly make changes in the database, these changes will not be replicated.

- 2 The new or changed data in domain A is collected in a replication package.
- 3 The replication package is sent to domain B.



Receiving in domain B consists of the steps 4 - 6:

- ④ The replication package is received in domain B.
- ⑤ The replication package is checked: is this the replication package for domain B and can be processed?
- ⑥ The changes to data records included in the replication package are transferred to domain B. Afterwards, a confirmation is sent to domain A. If no confirmation is sent, domain A sends the content of this replication package again.

Then, the sending process as described above in steps 1 - 3 for domain A is started.

Then, the receiving process is executed in domain A; it corresponds to steps 4 - 6 in domain B.

A replication is completed after the sending and receiving processes have been executed once for each of the two domains. As described in this document, this data replication is a bidirectional transfer. After the replication, the data which relates to the data records defined in the replication task is identical on both domains. Replication packages are deleted after the transfer.

If a replication package cannot be immediately received and transferred, no confirmation is sent to the corresponding domain. If no confirmation has been sent and a more up-to-date replication package is received, the previous replication package is replaced.

If replication packages exist which are not transferred yet, for example, due to a new start of the application server, then you can manually delete these packages. This can be done for replication packages which are dated before the application server has been newly started.

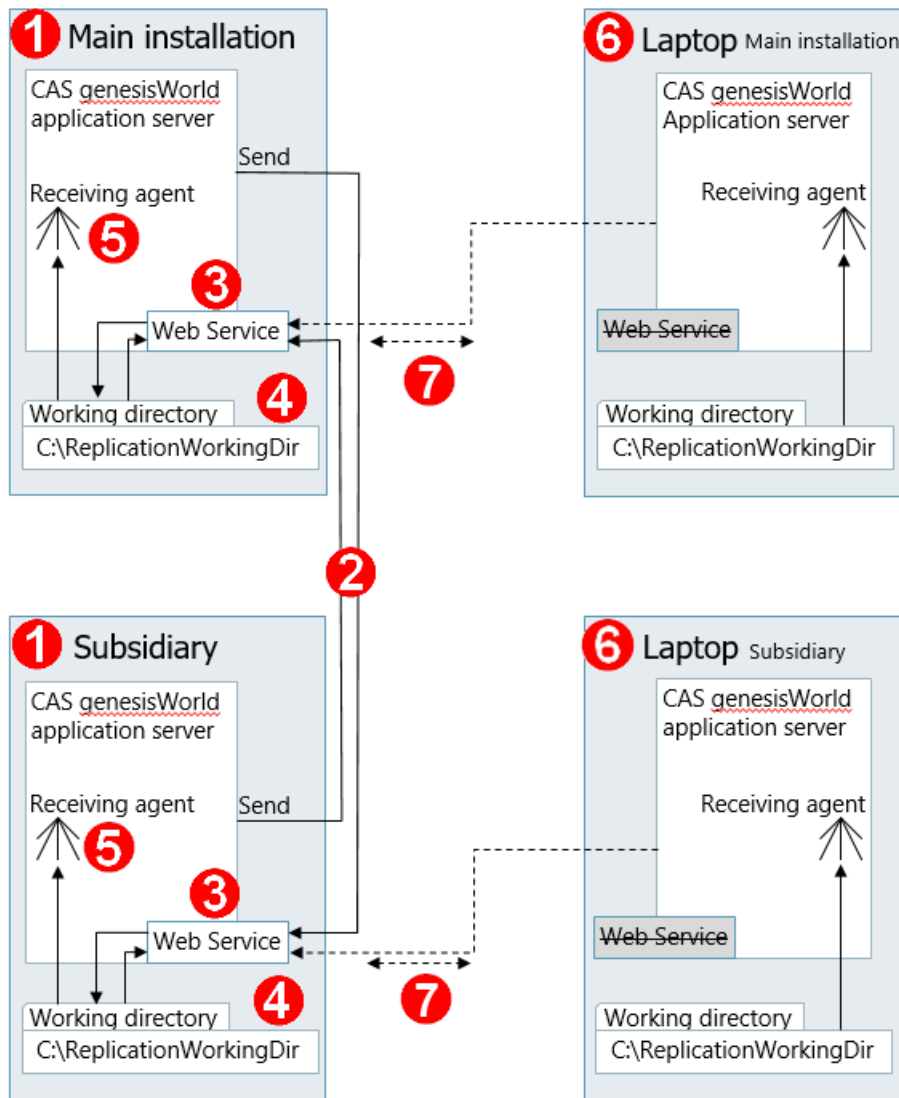
1.3 Main installation, subsidiary and laptops

The replication is executed via a Web Service with REST (Representational State Transfer). The operations for domains differ depending on the domain, that is, whether the domains are a main installation or other locations such as subsidiaries or laptops. The following figure describes a replication scenario between a main installation, a subsidiary and a laptop for each domain that replicates with a main installation and a subsidiary. The basic process is identical also for a number of subsidiaries or laptops.

- ① Data is replicated between the main installation and a subsidiary.
- ② The replication task with data is configured and sent by the application server for each replication.
- ③ The Web Service of the corresponding domain is the receiver of the replication task.

The Web Service for the replication is part of the integrated CAS genesisWorld Web server. This Web Service is, just like the Web Service for the SmartSearch, administered

via the CAS genesisWorld Server Manager, see the [Services](#) online help page of the Server Manager.



- 4** A received replication task is copied to the local C:\ReplicationWorkingDir by the Web Service.

The C:\ReplicationWorkingDir must be available locally on the computer of the CAS genesisWorld application server and/or meet the same requirements as a local folder. Only then can the folder be used for replication as described in this document.

- 5** The receiving agent is a replication service which is available for every domain and which is responsible for 2 operations: first, a receiving agent checks whether and which replication packages are available in the working directory. If replication packages are available for the domain with the receiving agent, they are transferred to the own domain.

The receiving agent prepares the receiving of replication packages. It is an integral part of the process for the steps 4 - 6 in the replication process, see "Replication process" on page 6.

The receiving agents from the main installation and subsidiary operate technically on the same level: they check via the Web Service whether or not a replication package is available for the "own" domain.

The receiving agents of laptops operate technically on the same level than the domain with which they are connected: via a laptop, the Web Service of the main installation or subsidiary checks whether or not a replication package for the laptop is available. The respective Web Service checks the information in the working directory of the respective domains.

- 6 Laptops just like the main installation and the subsidiary replicate independently and asynchronous because laptops may not always be connected online.

The Web Service on laptops is deactivated.

- 7 If a connection from the laptop to the main installation exists, the receiving agent of the laptop checks, via the Web Service of the main installation, whether or not a replication task for the laptop is available.

A list of available replication packages is sent to the receiving agent of the laptop via the Web Service. If a matching replication package is available, the Web Service of the main installation is informed by the receiving agent of the laptop.

The Web Service of the main installation then copies the matching replication package from the local working directory and sends it to the laptop. Transferred replication packages are then saved by the receiving agent of the laptop in the local working directory of the laptop.

The receiving agent always requests the replication task. This applies both to the laptop and the subsidiary.

1.4 Technical realization and connection costs

Replication data can be transferred via network connections, via LAN or worldwide via WAN, for example, by using a Virtual Private Network. This does not require both sides to be active.

Data is replicated through successive operations (threads) which work independently from each other. This enables parallel sending and receiving on a number of connections; but for each connection, only one of these operations can be executed.

A replication can always be executed with the corresponding function if the logged-on user has the required authorization. Additionally, a time-triggered execution can be performed: overnight or on the weekend. A time schedule can be created for each domain connection.

The data is always synchronized compressed and encrypted. The decompression is then made on the local system.

Possible connection costs

Two domains can be connected using a connection which is operated demand-driven via a dial-up line. In specific configurations it then occurs that the replication process keeps the connection open as the receiving agents regularly check whether a replication task is available.

Connection costs arise using a dial-up connection.

We recommend monitoring these connection costs during test operation and select the best configuration depending on the usage scenario or network environment. Contact your CAS Partner or CAS Software AG.

1.5 Replication licensing

Each domain with a corresponding database needs an extra license, excluding laptops and home office installations. If there is more than one user, licenses are required for each user.

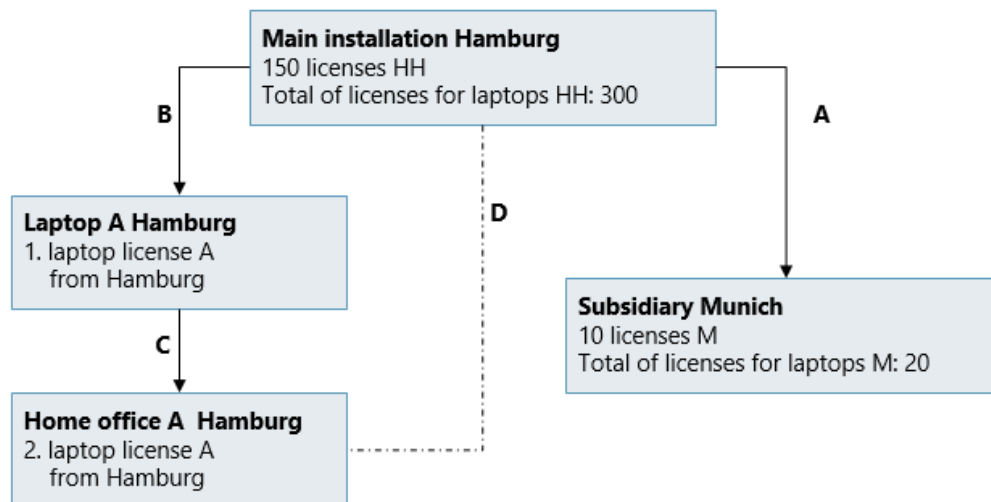
Replication can be performed between two database systems if the license numbers on both sides of the connection differ from each other.

For each Management Console license, you can create two laptop licenses. A separate license number is required for all laptop and home office installations if replications are to be performed. The license of the main installation cannot be used. These laptop licenses are free and are derived from the license of the main installation. Further licenses cannot be added to the laptop license. A laptop license is a single user license.

If users use a license for a module in the main installation, it can also be used by a user who works with a laptop license.

All license settings for users are replicated on all domains.

The number of laptop licenses is always twice the number of existing licenses. If, for example, ten licenses exist, users can use up to 20 free laptop licenses, that is, two laptop licenses per user.



- Data which is synchronized with laptop licenses can be synchronized between a license of the main installation and the corresponding laptop license, for example, when performing a replication between a laptop and the main installation.
- Data can also be synchronized between all laptop licenses which were derived from one identical license, for example, between laptop and home office.

We urgently advise you to not use cycles in the replication. This does not only apply to permanent but also to temporary cycles in the replication. If you do so, you will risk irreversible data loss.

The figure shows different scenarios.


- The replication between two licenses, for example, between two company sites: before the replication is started, the system checks if enough licenses on both sides have been installed and if the serial numbers differ from each other.
- The replication between license and corresponding laptop license, for example, a laptop that belongs to the subsidiary of a company: before replication is started, the program checks if the laptop license is derived from the appropriate group of licenses, that is, if the laptop license has been derived from the license of the main installation.
- The replication between two laptop licenses, for example, laptop and home office: before replication is started, the program checks if the two laptop licenses are derived from the appropriate group of licenses, that is, if the laptop licenses have been derived from the same license.
- Cycle: permanent and temporary cycles have to be avoided to prevent data loss.

1.5.1 Example for a cycle

Domain A can replicate directly with domain B that means that in both domains, a domain connection to each other domain has been created. If domain B replicates

directly with domain C, domain A replicates indirectly with domain C as data is transferred via B. Data is thus replicated from A via B to C.

A cycle is created if domain C replicates directly with domain A. A cycle is not permitted and must not be created.



When setting up a replication, please note that the replication can only be performed in one direction of a domain connection if two domains replicate with each other directly or indirectly via one or more domains.

1.5.2 Creating laptop licenses

Two laptop licenses can be created for each CAS genesisWorld license, for example, for a laptop and for working at home.

✓ Open the **Licenses** area in the Management Console.

 **Create laptop licenses** opens the respective window.

 **License list** opens the **Generated laptop licenses** window.

Print prints the list.

Copy copies the list to the clipboard.

Save as file saves the list as .txt file.

1.6 Protecting data from outside the company network

Data is outside of the company network if data is replicated via the Internet and if users work on laptops locally and do not depend on the company network.

1.6.1 Protecting data during transfer

Data is always transferred compressed and encrypted during replication. After the data has been transferred, the system checks whether or not data has been changed.

The first time data is replicated, a key is exchanged between both domains.

For each replication that follows, both domains must be identified through the exchanged keys via the Web Service.

If you protect the data transfer via the Internet by using the SSL protocol, you reduce the options for wiretapping or at least make it harder to wiretap.

The SSL protocol (Secure Socket Layer) secures the transfer of data via server authentication, data encryption and data integrity. The replication function uses a Web Service of the CAS genesisWorld application server. This means that the application server is used as a Web Client and must identify itself before each data transfer. This SSL technology provides far more security than just logging on with the CAS genesisWorld user name and password.

To use the SSL protocol, server certificates are installed on the computer on which the CAS genesisWorld application server and thus the Web server and/or the Web service is installed.

The application server uses HTTP port and thus also works if a firewall is set up between two computers. Thus, for replication purposes, you are able to work with HTTPS and you do not need a Virtual Private Network (VPN). The used ports must be included in the firewall.

1.6.2 Protecting data on laptops

Users working with laptops and the replication function should not be permitted to directly access the CAS genesisWorld database, the Document Archive or the SmartSearch index. You can define the respective settings for permissions in Windows and in the settings for the Microsoft SQL Server. Please note:

- A password to access the Microsoft SQL server should be provided. We strongly recommend entering a password during installation.
- The folder containing the Document Archive should be protected via NTFS (New Technology File System) so that users do not have direct access. Only local administrators should have access. For example, only allow local administrators access and assign main user rights under Windows for users of CAS genesisWorld.
- A user with local administration rights must be entered in the component services of the CAS genesisWorld package to make sure the application server runs on a laptop and users can access the Document Archive. The required changes can be made with the Server Manager.

In the user name field, you must enter the name of the domain and/or computer name in the **Server registration** tab of the Server Manager. The following syntax applies: Domain name\User name.

After you have made the change in the **Server registration** tab, click **Apply** to ensure that the user is entered with the next registration. If the logon is requested again, it means that the logon is invalid.

- As administrator, maintain the databases on laptops in order to ensure performance and security.

1.6.3 Rights in CAS genesisWorld are replicated

In addition to CAS genesisWorld data, such as addresses, documents, tasks, opportunities, and so on, the data for rights is also replicated.

Rights that enable the display of data, for example, in navigators or team calendars are also replicated; furthermore, rights for functions, for example, for actions or administration rights and rights through activated licenses are replicated.

To ensure that a user has consistent rights to his or her data across platforms, the corresponding data is synchronized each time a connection is established. This applies to both the direct settings for a user account and to inherited settings, for example, via group membership.

In this way, all replication data and functions are included in the extensive CAS genesisWorld user and access rights system.

2 Setting up the replication

Usually, the administrator defines the requirements for the replication in CAS genesisWorld and sets up the domains.

Both administrators and users can set up domain connections and replication tasks as well as execute replications.

2.1 Requirements and first steps

- The Desktop Client, the CAS genesisWorld application server and a database system must be installed on each computer on which you want to replicate.
- If you deploy replication and CAS genesisWorld must be completely installed or updated on a computer, you cannot do this via MSI.
- The database system must be installed separately on a computer. The database system can, as usually in CAS genesisWorld, be installed without setting up a database.
- In principle, you can use different releases of Microsoft SQL Servers on different domains. However, there are some technical restrictions if you set up new domains. Thus, we recommend that you use the same Microsoft SQL Server version on all replicating domains.

You will find more details on approved database systems and the respective versions in the current release notes at [CAS Product Documentation](#).

- The application server on a computer will be registered for a Windows user during installation. This Windows user requires the permission to access the

C:\ReplicationsWorkingDir and ..\Temp folder, see "Main installation, subsidiary and laptops" on page 8.

- An administrator requires rights to the respective folders if he or she wants to check these folders.

Entering licenses and defining rights

- ✓ As usual, the replication license has to be entered in the Management Console in the **Licenses** area and must then be activated for the corresponding user.
- ✓ As administrator, you must also assign the following right: open the **User management** area, go to **Properties** of the respective user and select the **Replication** option.
- ✓ The **Replication** right must also be set for the **Administrator** user.
- ✓ You can only set up a domain with the **Administrator** user account of CAS genesisWorld.

2.2 Special requirements for laptops

- Laptops require special laptop licenses; these licenses can be created by the administrator in the Management Console, see "Creating laptop licenses" on page 13.
- If there are no available laptop licenses in the main installation, the database is created on your laptop with a test license. You must then enter an additional license for a user in the Management Console, generate a new laptop license and register it on the laptop.
- A domain must already be set up on the application server of the main installation.
- Before setting up a domain on a laptop, a database system must be installed.

Depending on your requirements, you can also use the Express Edition of the Microsoft SQL Server.

The maximum data volume for the Microsoft SQL Server Express Edition is restricted. Please refer to the manufacturer's specifications.

You will find more details on approved database systems and the respective versions in the current release notes at [CAS Product Documentation](#).

For security reasons, we strongly recommend to assign a password for the database to prevent unauthorized access to your data.

- When working with a laptop, you can both work on your PC or within a company network. The domain, that is, the application server to which a user is logged on, is displayed on the title bar of the program window. If you log on locally, **CAS genesisWorld (localhost)** is displayed.

2.3 Steps and setup order

- 1 First, you define the settings in the Server Manager.

You must start the Web Service for the replication in the [Services](#) tab.

If you deploy a reverse proxy and you replicate data, the same rules apply as the ones for the settings for Mobile Apps, see Server Manager on [Configuring the reverse proxy](#).

- 2 The first domain is set up. This is usually done while installing CAS genesisWorld; this installation is then also regarded as the main installation.
- 3 One or more additional domains are then created either with a copy of the main installation database, or with an empty database. You require additional domains for subsidiaries or laptops, which have to be replicated with a main installation or subsidiary.


If you are using different Microsoft SQL Server versions, then you cannot import the database backup of a newer version of Microsoft SQL Server into a previous version.

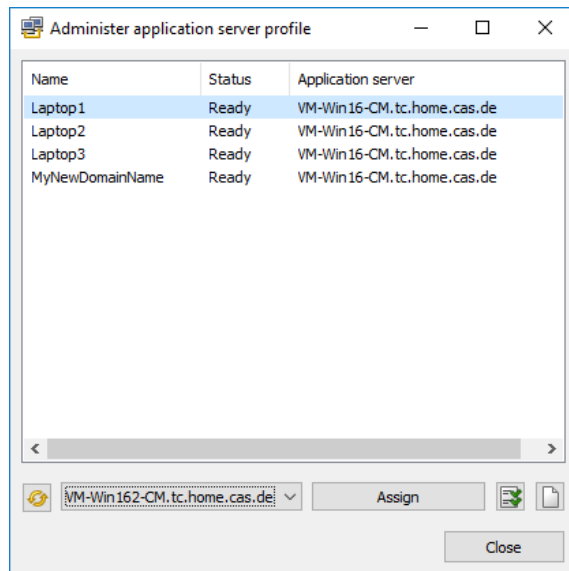
This is a significant disadvantage when creating new domains.

If the versions of the Microsoft SQL Server differ when you create them, the system displays a corresponding message.

- 4 On each domain, an application server profile must be set up in which the URL for the Web Service as well as directories to exchange data are defined, see "Main installation, subsidiary and laptops" on page 8.
- 5 Domain connections are set up for each domain with which data is to be replicated.
- 6 Replication tasks are created on a domain and contain the data that is to be replicated.
- 7 For laptops we recommend that the administrator defines settings for agents, see "Defining settings for agents for laptops" on page 34. This also applies to settings for automatic operations, see "Laptop automatics" on page 34.
- 8 Application server and Desktop Client must also be always updated on a laptop. This cannot be done via MSI, see "Software updates" on page 72.

2.4 Administering application server profiles

 **Administer application server profiles** in the **Replication management** window of the Desktop Client or the **Replication** area of the Management Console opens the respective window.




A list with all domain connections of the current domain is displayed. The service for the replication can be active on several application servers. The application server profile defines which application server or Web Service receives replication packages for which domain connection. The window displays which application server is assigned to which domain connection.

Only if you log on as administrator, you can select and configure an application server for a domain connection.


Selecting and configuring an application server

The information for the configuration of an application server for the replication is defined in the application server profile.

 **New** opens the **Application server profile** window. A new profile is created for the application server which is set in the drop-down list.

 **Properties** opens the **Application server profile** for the set application server.

✓ To define an application server for replication, select one or several domain connections in the list above.

 Update the drop-down list if necessary and select an application server.

A profile is always assigned in the same way. But only one application server is displayed in the drop-down list.

✓ Click **Assign**. The application server and the application server profile are used for the selected domain connections.

2.5 Creating/editing application server profiles


Application server profiles are created or edited in the **Administer application server profiles** window or in the **Receive** tab of a domain connection.

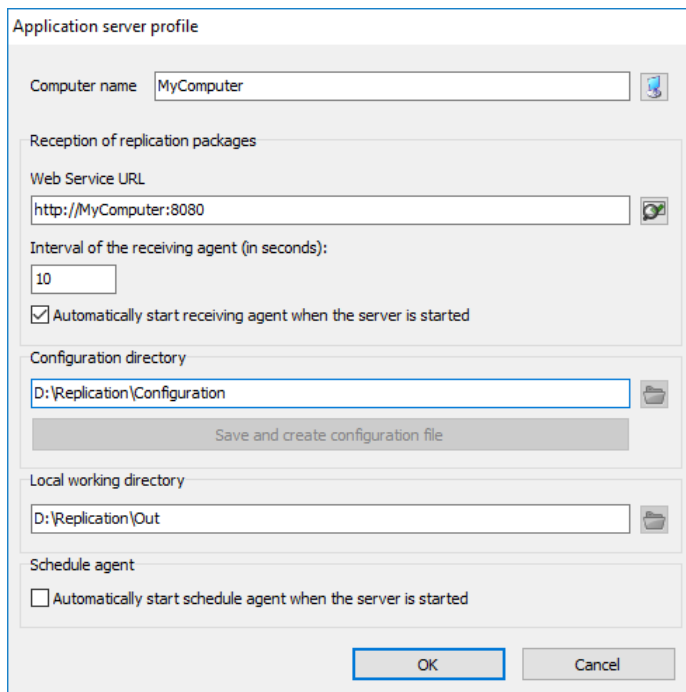
You can create and edit application server profiles only if you are logged on as administrator.


A separate application server profile must be available for each application server if the replication is to be used on this application server.

- In the application server profile, you define settings on how the application server is being used by the corresponding domain, that is, the main installation, subsidiary or laptop.
- The database must be available on this application server, for example, for the main installation, the subsidiary or the laptop.

All settings in the window are also saved in the database to which you are currently logged on.

  The **New** and/or **Properties** buttons open the **Application server profile** window.



 At **Computer name**, you enter the name of the application server of the local domain; alternatively, you can click on the button to the right to transfer the current computer name.

✓ Enter the **URL of the Web Service** of the application server via which you receive the replication packages.

For a main installation or a subsidiary, this would correspond to the Web Service of the computer which is also entered in **Computer Name**.

For a laptop, this corresponds to the Web Service of the main installation or subsidiary, see "Main installation, subsidiary and laptops" on page 8.

The URL can be, for example, the IP address or the name of the computer. Use https and/or http and enter the port.

Example: http://MyGWServerName:8080


By default, port 8080 is assigned if no other settings were made in the Server Manager for the Web Service, see "Setting up the replication" on page 15.

 With this button, you can test whether the Web Service can be accessed.

The system checks whether a domain is set up on the corresponding application server and a partial replication can be made.

- ✓ Define the **Interval of the receiving agent**.
- ✓ Define whether the **Receiving agent** is also to be started when starting the application server.

You can change this default value in the **Replication administration** window and/or the **Replication** area of the Management Console, see "Receiving agent and schedule agent" on page 33.

 The **Configuration directory** offers support when setting up new domains on laptops. The configuration directory saves a file containing the information for the configuration.


Use a local folder on the application server which can be accessed by the laptop via a UNC path.

- ✓ Click **Save and create configuration file** to save the settings made in the **Application server profile** window of the main installation, or subsidiary, in a file of the configuration directory.

This includes in particular the URL of the Web Services of the main installation and/or subsidiary.

The one-time password or the key are also saved in the configuration file to protect data during replication, see "Protecting data during transfer" on page 13.

You can only use configuration files for domains which use a laptop license.

 You create a folder on the application server of the local domain as **Local working directory**.

The replication packages of the local domain are created in this folder. Furthermore, the replication packages are saved in this folder after they have been received.


The folder must be available locally on both the main installation, the subsidiary and the laptop; furthermore, the application server must have full access rights to this folder. The Web Service runs via Windows services and also has access.

The receiving agent checks in this folder, whether replication tasks of other domains are copied via the Web Service.

- ✓ Define whether the **Schedule agent** should also be started when starting the application server.

In the **Replication administration** window and/or the **Replication** area of the Management Console, this default value can be changed, see "Receiving agent and schedule agent" on page 33.

Folders and paths

 With the buttons, you can only select the respective folders if the **Application server profile** window has been opened on the computer with the respective application server. As administrator, you must be logged on to the Desktop Client and/or the Management Console of the application server for which you want to set up the application server profile.

- The folder for the local working directory must have sufficient storage space. If a daily replication with the main installation is scheduled for 20 laptops and the replication packages are not transferred, they are stored in the local working directory. Replication packages are only deleted after the data has been transferred, see "Replication process" on page 6. If storage space is running out, programs may no longer be correctly executed.
- The application server of CAS genesisWorld is registered with a Windows user. This user must have read and write rights to the configuration directory and the local working directory.

This applies to each Windows user of an application server with which data is replicated.

- If a network path is mapped to a drive, it is identified as local path.

If folders do not have a computer name, the server accesses its local drive mappings, for example, **C:\Inbox** as local C hard drive.

When working with folders with a computer name, the application server accesses the network services of Windows, for example, **\\gWServer\Inbox**.

When sharing a folder, the **Full access right** must be entered for the Windows user.

2.6 Replication functions

All replication functions are both available for users in the Windows client and for administrators in the Management Console.

You will only have access to the respective replication functions, if you (as user) have received the respective rights. Only the administrator user account can access functions for the domains.


 Go to the **Tools** menu of the Desktop Client and click **Replication/Management**.

 In the Management Console, go to the **Replication** area.

The left pane contains a navigator with several levels. If you click an entry in the navigator, the subordinate levels are displayed in the right pane of the window.

- The highest level is the **Domain**. Only an administrator can set up a domain.

The display always shows the local domain. This is the domain to which you are logged on.



- The next level displays the **Domain connections**. Domain connections are created between domains. A domain connection is thus required for all domains with which replications are executed.
- The last level shows the **Replication tasks**. Replication tasks are matched to domains. A replication task defines which data is to be replicated. Several replication tasks can be created for different data.

With which application server do you work?

If you work with the replication function, you can either work on the local laptop or in the company network, for example, out of office or at the working station in the company.

When logging on to CAS genesisWorld, you enter your name and your password.

Depending on your replication scenario, you log on with the same client credentials, but to another CAS genesisWorld application server and therefore also to another database.

- ✓ When logging on to the company network, you enter the name of the company network in **CAS genesisWorld logon > Server name**.
- ✓ When logging on to the laptop locally, you do not fill in the **Server name** field in the **CAS genesisWorld logon** window.

The title line of your CAS genesisWorld program window displays the application server to which you have logged on. If you are logged on locally, **CAS genesisWorld (localhost)** is displayed.

On which database do you work?

- ✓ If you log on to different databases, the **Replication management** window can be opened for different databases.
- 🔄 When opening several **Replication management** windows with the same database, the modifications in a window can only be viewed after clicking the **Update** button in the other windows.
- 📄 **New** opens a menu with which you can create, depending on the selected entry, a **Domain connection**, a **Replication task** or - as administrator - a **Domain**.
- 👉 **Replicate now** starts the replication on the selected domain connection.

In the Desktop Client, the button is also displayed in the toolbar of the program window.



The settings for the domain connection, the option for this function and the properties of the **Replicate now** button must be defined. You can then start a replication with just a click and the **Replication management** window must not be opened, see "Settings tab" on page 46.

- 🏠 **Properties** opens the **Local domain** window for the domain to which you are logged on, see "Functions for domains" on page 31.

As user, you can see all properties of the local domain and you can inform yourself.

As administrator, you can change all properties of the local domain.

- 🔍 **Connection status** opens the respective window, with the log file of the domain connection and the corresponding functions, see "Connection status window" on page 52.

-  **Domain log** opens the respective window in which status, warning and error messages for the current domain are displayed.
-  **Administer application server profiles** opens a window with the domain connections of the current domain, see "Administering application server profiles" on page 17.
- ✓ The **Receiving agent** and **Schedule agent** options apply to all domain connections of a domain and display the current setting.

The administrator defines defaults and possible settings, see "Receiving agent and schedule agent" on page 33.

3 Domain

Each computer in a company can act as a domain if a database and an application server of CAS genesisWorld is installed.

Each domain must have a unique name. The name of the computer to which you are logged on is entered as default. Do not name the domain LOCAL.

Domains of a main installation and a subsidiary do not differ with regard to their technical characteristics. Laptop domains differ: laptops use a laptop license and the Web Service of the main installation or the subsidiary.

Database independence and multitenancy

Data is replicated database-independent and on an application server level. You can synchronize heterogeneous CAS genesisWorld installations with different databases and database systems. Replication is therefore also possible if you use the CAS genesisWorld multitenancy function.

The maximum data volume for the Microsoft SQL Server Express Edition is restricted. Please refer to the manufacturer's specifications.

3.1 Setting up the domain

At first, setting up a domain requires a full CAS genesisWorld installation and the corresponding rights, see "Requirements and first steps" on page 15.

You can set up a domain in the **Replication** area of the Management Console or in the **Replication management** window of the Desktop Client.

The new database and the way how you want to create this database determine to which application server and to which database you log on.

We recommend to log on to the Management Console if you, for example, also need to set rights.

Required rights

Only users with the CAS genesisWorld **Administrator** account can create and set up domains.

An exception is made for settings of the receiving and schedule agent of the domain: the administrator can individually define whether users can start or stop them.

Additionally, administrator rights to all folders in Windows which are used during the replication are necessary.

Setting up the replication function for the first time and creating the first domain

If you want to use the replication function for the first time, you first define the domain for the main installation.

- ✓ Log on to an application server which is linked to the database of the main installation and log on to the respective database.

You probably work with several application servers. It is then important that you log on to the database of the main installation and to one of the connected application servers.

- ✓ Open the **Replication** area in the Management Console or the **Replication administration** window in the Desktop Client.

The **Define domain name** window opens.

Creating a second or additional domains

Before creating a second domain or additional ones, you first decide whether you want to create an additional domain with an empty database or whether you want to use a copy of an existing database.

Depending on your decision, you log on to a specific application server and database.

Creating additional domains with empty database

- ✓ Log on to the application server on which you want to set up a new domain.
- ✓ Log on to any database which is connected to the application server used above.

- ✓ In the **Replication** area of the Management Console or in the **Replication administration** window of the Desktop Client, go to the **File** menu and click **New domain**.

The **New domain: Database and Document Archive** window opens.

Creating an additional domain with a database copy

The replication works efficiently: only changes made since the last replication are transferred for a replication task between two domains. This does not apply to the first time a replication is executed: in this case, the complete data of both sides must be transferred since the two domains do not have sufficient information from each other. This complete exchange of data usually takes very long, even if the data volume is in the medium range. You may also need more time if you want to set up a number of laptops.

You shorten the replication process with a copied database.

- ✓ First, create a database copy of an existing database with the Server Manager or by using the functions of the database system.

The database backup must include a CAS genesisWorld domain. This domain cannot be a laptop domain.

- ✓ Log on with CAS genesisWorld to an application server which is connected to the database that has been copied.
- ✓ Log on to the database from which you want to take the copy.
- ✓ In the **Replication** area of the Management Console or in the **Replication administration** window of the Desktop Client, go to the **File** menu and click **New domain**.

Next step

If you have selected the **New domain** function and if the name of a domain is defined, the **New domain: Database and Document Archive** window opens.

3.1.1 New domain: Database and Document Archive window

The window opens if the name of the new domain has already been defined. You have also decided whether you want to create a new domain with an empty database or a copy of an existing database. Depending on your choice, you will be logged on to a certain database.

In the **New domain: Database and Document Archive** window, you define how the database and the Document Archive for a new domain must be created.

The setup log is saved and can be opened after the creation of the domain, see "Requirements and first steps" on page 15:

%TEMP%\genesisWorld\CloneProtocol\ReplicationClone.log

Before setting up a domain, the system checks whether the CAS genesisWorld application server is installed on the computer.

- ✓ At **Name of database computer**, enter the computer on which the database system for the new domain is installed. In this database system, a new database is created for the new domain.

You can set up a domain for the replication with an empty database or with a copy of an already existing database. You will always need a database system which is already installed before you set up the replication, see "Special requirements for laptops" on page 16.

- ✓ You log on to the database system with the name and password of the **Database user** and you therefore use the credentials of the administrator for a Microsoft SQL Server.

With the existing logon information, you log on to the computer with the database system with which you want to create the new database and/or domain. With the following entries, you define how a new database is set up on the computer with the database system.

- ✓ If you want to create a new and empty database, select **Create new database with the following name** and enter a new name for the database of the new domain.

The application server of the new domain must not be active during this process.

If the computer with the database system already has a database with the entered name, a message will appear.

- ✓ With **Use copy of the current database**, you can create a new database from a backup of an existing database.

In the main installation, a new domain can be created and then, for example, be copied to and used on a laptop.

- ✓ Select a copy and/or the backup of the current database with **Select database**.
- ✓ With **Create full replication task**, you already create a replication task in this window. You can also do this in the following window.

The option is not displayed if you create a new database. A full replication task for an empty database is taking very long and we do not recommend it.

If you wish to perform replications with the new domain using a full replication task, then you have to ensure that a current copy of the database is available.

- ✓ If the **Copy archive files if replication tasks with documents exist** option is not active, you are not allowed to create replication tasks for document data records with archive files.

If you still create replication tasks for document data records with archive files you may create inconsistent data.

This option only affects the **New domain: Tasks** window. After the new domain has been set up, replication tasks for document data records with archive files can be created at any time.

- ✓ If the **Copy archive files if replication tasks with documents exist** option is active, archive files can be replicated. You can transfer an already existing Document Archive for the new domain.

- 📁 Enter the folder with the **Document Archive of the source**. The content of this folder is then copied to the **Document Archive of the new domain** folder.

When creating the domain, the folder specified for the **Document Archive of new domain** is used for copying files from the source. This can be the final folder, for example, for a laptop domain.

The folder for the Document Archive of the new domain is defined in the Database Wizard with **Connect database**.

As logged-on Windows user, you must have read and/or write rights to all used folders in the Document Archive.

- ✓ If you go to the **New domain: Database and Document Archive** window, and click **Next**, the **New domain: Tasks** window opens.

Notes on the Microsoft SQL Server rights system

Due to the superior rights system of the database system, problems with external database users (who are not CAS genesisWorld users) may occur when importing a backup to an external Microsoft SQL Server.

On the computer with the first database system, you have created the backup and you have given access to the corresponding database via a database logon. This database logon does not yet exist on the second computer.

When creating a logon on the second database system and assigning access permission to the restored database, a message from the Microsoft SQL Server informs you that the corresponding database logon already exists in the database.

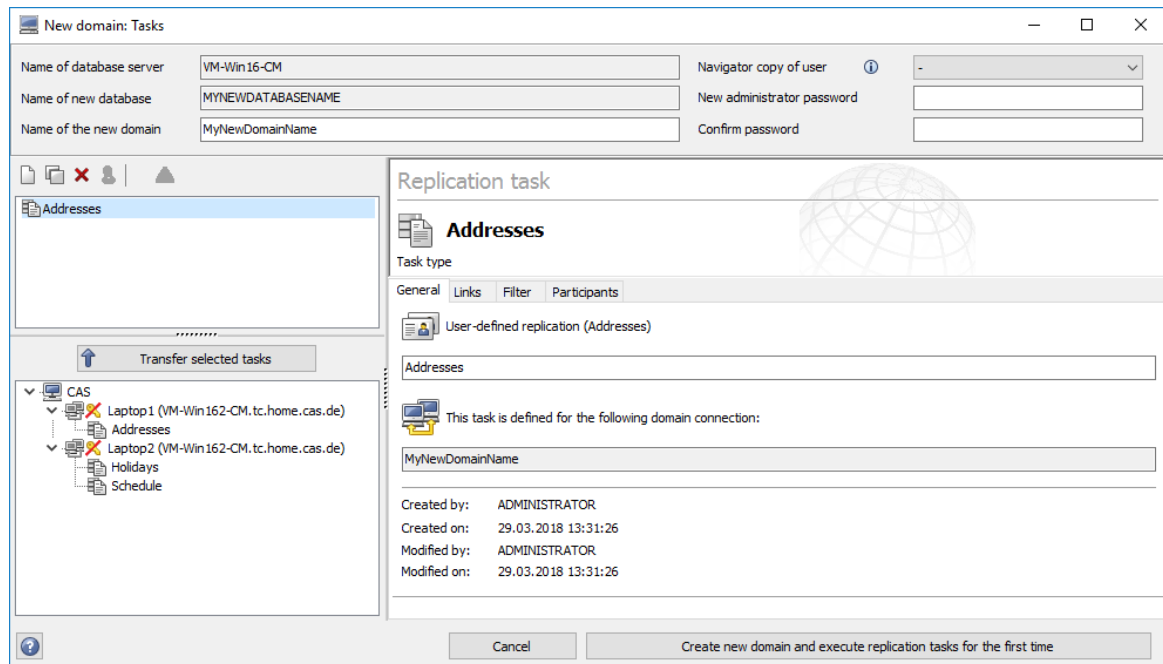
You can solve this problem with the following steps:

- ✓ Create a database logon on the second database system without giving this logon access rights to the restored database.
- ✓ Open the Query Analyzer and select the corresponding database from the drop-down list.
- ✓ Execute the following command: `sp_change_users_login 'Auto_Fix', '<login>'`
- ✓ `<login>` is the placeholder for your logon name.
- ✓ Check the access rights of the database logon to the restored database. The access right should now be assigned and the user should also have DBO (database owner) rights to the database. If this is not the case, you can make changes now.

3.1.2 New domain: Tasks window

In the **New domain: Tasks** window, you define the password for the administrator of the new domain, transfer navigators of a certain user and create new replication tasks for the new domain if required.

The **Name of the database computer**, the **Name of the database** and the **Name of the new domain** are adopted from the **New domain: Database and Document Archive** window.



Just like in the **New domain: Database and Document Archive** window, these entries refer to the database which will be newly created.

- ✓ In the **Navigator copy for** drop-down list, all user accounts of the domain and/or the database to which you have logged on are displayed.

When creating the new domain, the navigators are transferred for the user account that has been selected from the drop-down list.

- ✓ In **New administrator password** and **Confirm password**, you enter a password for a new domain. With this password, you then open the Management Console on a laptop and check or change the settings if required.

Replication tasks of all domain connections of the domain and/or the database to which you have logged on are displayed in the lower left-hand side corner of the window.

- ✓ The settings of a selected replication task in the lower left corner are displayed in tabs on the right-hand side. These settings cannot be changed.
- ✓ The properties of a selected replication task on the upper left-hand side can be immediately adjusted in the tabs on the right-hand side.
- ✓ Select the replication tasks and click **Transfer selected tasks** to transfer the tasks to the new domain.
- ✓ You select several replication tasks by clicking on them and pressing Ctrl at the same time.

📄 With **New**, you create a new replication task in the domain.

📄 **Duplicate** duplicates the selected replication task.

- ✗ Deletes the selected replication task.
- 👤 Select one or several replication tasks in the upper part and click the button. The user who is set as user from which navigators are copied in the upper right-hand window is added as participant to the selected replication tasks.
- ⚠️ The button opens a window with messages.

If the **Create full replication task** option is active in the **New domain: Database and Document Archive** window, then any replication tasks which have been transferred to the **New domain: Tasks** window are ignored.

- ✓ Click **Create new domain and execute replication tasks for the first time**.

The progress is displayed in a window.

In a copied database, the system data of the copied domain is transferred to the second domain. System data can include, for example, users or database extensions.

- ✓ If the new database has been created, select the **Open database wizard** button and connect the new domain with the respective application server.

When using a laptop with a copied database of the main installation, you can execute, for example, the first steps for setting up the new domain without having to work on the laptop.

- ✓ A laptop license is entered automatically in the new database. You can change this for a subsidiary if necessary.
- ✓ Then, you log on to the new domain (database). You can do this via the Management Console of the laptop or via the Desktop Client with the respective application server.

Define now, for example, the settings for replicating data with this domain.

3.2 Functions for domains

After you have set up a domain, you define the settings you require for replicating data with this domain. .

3.2.1 General tab

- ✓ The name of the domain can be changed.
- ✓ When activating the **Laptop automatic** option, you can transfer data records which were newly created on laptops automatically to replication tasks, see "Laptop automatics" on page 34.
- ✓ In the **E-mail recipients if error occurs** field, you enter one or several e-mail addresses.

To be able to send the e-mails, a valid e-mail account must exist for the notification and action service. This e-mail account is set up in the Management Console in the **Rules** area, **Server settings** tab. The notification and action service must not completely be set up on a laptop, see "Notification and action service" on page 75.

When errors occur and especially if the replication is stopped, e-mail notifications are sent; e-mail notifications are not sent for warnings. We recommend to check the log files for the domain and the connection.

3.2.2 Settings tab

- ✓ If you deactivate the **Only the administrator can start and stop agents** option, users are allowed to activate and/or deactivate the receiving and schedule agent, see "Receiving agent and schedule agent" on page 33.

The system then uses the settings for the agents which are defined in the application server profile for the used application server, see "Administering application server profiles" on page 17.

- ✓ If you activate the **The full replication task can only be set up by the administrator** option, users are not allowed to create full replication tasks.
- ✓ When activating the **Accept changes to database model** setting, this setting is adopted as default value for the new connection.

These settings only apply to the domain connection which is created after the option has been activated, see "Administration tab" on page 48.

3.2.3 Maintenance tab

The log may be very large. As administrator, you define if and how you want to automatically shorten the log.

The setting you select will be applied the first time you log on to a domain and then afterwards on a daily basis.

After the installation, the first option is set as the default so that initially no entries are deleted.

- ✓ **Do not automatically delete entries from log file** means that data is constantly logged. The newest entries are displayed at the top in the **Log** tab.
- ✓ **Automatically delete entries after a certain time period:** In the drop-down list, choose the time from which log file entries are to be deleted.
- ✓ **Automatically delete entries if the log file has reached a certain size:** Specifies how many entries should be kept. Entries exceeding this number are automatically deleted.

3.3 Receiving agent and schedule agent

These agents are services for the replication. The settings defined for the receiving and schedule agent apply to all domain connections of a domain.

Receiving agent

The **Receiving agent** checks in pre-defined time intervals whether replication packages are available in the working directory, that is, whether they were sent from another domain, see "Main installation, subsidiary and laptops" on page 8.

The receiving agent must be active if the replication is to be executed according to a schedule. Only then are replication packages automatically transferred.

As a rule, all receiving agents of the replicating domains should be activated if the schedule agent is activated. If, for example, two domains replicate hourly, the schedule agent must start the replication. At the same time, the receiving agent must also be activated to receive the sent data.

We only recommend to continuously activate the receiving agent if a permanent connection to the Internet is possible and can be maintained.

Schedule agent

The **Schedule agent** monitors all domain connections for which schedule agents have been defined. The replication is then started on the defined times. The schedule agent must be activated if you want to replicate data according to a time schedule.

A schedule agent does not require access via a Web Service and can always be activated.

3.3.1 Functions for agents

- ✓ In the application server profile, you define which agents are active when starting the application server of a domain, see "Administering application server profiles" on page 17.
- ✓ In the **Settings** tab of the **Domain**, the administrator defines whether users are allowed to change this option, see "Settings tab" on page 32.
- ✓ You can change the current status for the **Receiving agent** and the **Schedule agent** in the bottom left-hand corner of the **Replication management** window or in the **Replication** area of the Management Console.

Users are only allowed to change the status if the administrator has permitted the corresponding right in the domain settings.

If users are not allowed to make changes, only the **Administrator** user is allowed to change the status in the **Replication management** window or the **Replication** area of the Management Console.

3.3.2 Defining settings for agents for laptops

If no permanent connection between domains is available, receiving agents and schedule agents can always be activated manually. We recommend the manual activation for laptops.

Users can define settings for agents if the administrator has granted the respective permission.

The receiving agent permanently accesses the Web Service of the main installation. The schedule agent requires access to the connection when transferring data. This access is not always possible.

- ✓ When creating a domain on a laptop, deactivate the **Receiving agent** and **Schedule agent** options.
- ✓ Users then activate these two options in a domain connection before they start the replication.

3.4 Laptop automatics

Two automatics exist for laptops:

- When using the **Laptop automatic**, new data records on the laptop are transferred to the main installation without great efforts.

If new addresses, documents, tasks, opportunities, and so on, are created on the laptop, these new data records are automatically transferred to system tasks. This only happens if the new data records are not replicated via an existing task.

- The **Automatic clean up function** deletes those data records on a laptop for which no replication task has been defined.

Usually, you do not save all data of the main installation on the laptop, but only the data required, for example, for appointments, open tasks or documents with a customer dossier for a presentation. Replication tasks are created for these data records.

For example, you have given a presentation for customer A and have created a corresponding replication task. You now want to give a presentation at customer B's location. You delete the replication task for presentation A and create a new replication task for customer B. All data of the replication task A is deleted on your laptop.

If data is no longer included in the replication task a comparable process is triggered. For example, there is a replication task for addresses with postal code 7*. A customer moved and lives now in an area which starts with the postal code "8". This means that this address is also deleted from the laptop.

Functions for automatic actions

The automatics are activated or deactivated by the administrator in the domain settings.

- ✓ As administrator, activate the **Laptop automatic** option in the **General** tab.

The option both activates the task and clean-up automatic.

Notes




The laptop automatic function should only be used on laptops. Otherwise, you may delete data in a main installation or subsidiary which is not replicated.

The laptop automatic function must not be activated on both domains. This is checked when receiving replication packages before replicating the data.


3.5 Domain log window

 The **Domain log** window shows messages and functions for the displayed domain.

The **Domain log** window can be open while you are working with CAS genesisWorld.

   With these icons, you can activate or deactivate the display of status, warning and error messages.

✗ Use this button to delete one or more selected messages. If no message is selected, all messages are deleted after a security prompt was displayed.

 **Copy selected rows of the log into the clipboard** copies the rows, for example, for support requests.

- ✓ In the **Properties** window of a domain, you define whether and how you want to automatically cut the log file, see "Maintenance tab" on page 32. This setting applies to the log of the local domain and to all connection logs, see "Connection status window" on page 52.

3.6 Moving a domain to a new location

You may want to work with an existing domain on a different computer or a computer may not be fully set up yet. This may occur if

- a user receives a new laptop, or

- a laptop has already been set up and the Web Service is not yet configured.

The following steps can only be executed by the CAS genesisWorld administrator.

How to

- ✓ Stop the application server on the old computer.
- ✓ Disconnect the database from the application server on the old computer with the Database Wizard.
- ✓ CAS genesisWorld should be already installed on the new computer.
- ✓ Create a database backup on the old computer.
- ✓ Import the backup file to the new computer.
- ✓ Copy the Document Archive from the old computer to the new one.
- ✓ Connect the database and the Document Archive with the application server on the new computer.
- ✓ Start CAS genesisWorld on the new computer and log on as administrator of the local domain.

In the logon window, enter the name of the main installation as name of the application server.

- ✓ Open the **Replication management** window and the **Application server profile** window, see "Creating/editing application server profiles" on page 19.
- ✓ Select the **Configuration directory** on the application server on the main installation or enter the URL parameter and the one-time password for the communication.

The configuration file of the main installation including the relevant Web Service settings is adopted by the laptop. This is done as soon as the laptop has established a connection to the main installation.

New computer for a main installation

The configuration directory is also useful when the main installation is relocated to a new computer and updating laptops.

- ✓ The new domain of the main installation is set up.
- ✓ The UNC path for the configuration directory is identical for the old and the new computer.
- ✓ The configuration file for the main installation has been created.

The Web Service of the old application server is configured in the configuration file of the laptop. As the old Web Service can no longer be accessed, the laptop cannot replicate.

The configuration file is opened on the application server of the main installation via the laptop. This new folder can be accessed via the same UNC path as the laptop.

The new configuration with the new Web Service is transferred via the laptop and data can be replicated.

3.7 Fusing domains

Two companies at different sites merge. Both companies work with CAS genesisWorld and both work with a separate database each. Both companies want to work with the same data and to replicate their data in the future.

To be able to replicate between the two sites, the two independent databases have to be fused.

The following explanation describes the general basic steps. It is important to plan well and run extensive tests beforehand. Contact your CAS Partner to find the best solution for your company.

How to

- ✓ Make sure to create a complete and up-to-date backup of all important data for both databases before the fusion: this should include the database, the Document Archive, eventually portals, e-mails, and so on.
- ✓ First, the extensions of each database, like new data record types or fields, must be transferred to the respectively other database.
- ✓ This can be done by
 - importing script files with the Database Wizard, or
 - using the export/import function for XML files in the **Database** area of the Management Console.
- ✓ Adaptations and changes of the GUI for data record windows which were made with the Form & Database Designer must also be synchronized in both databases. You can use the export and import function of the Form & Database Designer.
- ✓ Export and import also the input assistance options of each database so they are identical in both domains.
- ✓ Set up a domain in each database if no domain exists.

If data has already been replicated in a domain, for example, with laptops, no new domain must be set up.

You must then only set up a new domain connection to a location.

- ✓ Set up a domain connection from each domain to the new corresponding domain.


Both domains and the domain connections in each domain to their corresponding domain are necessary to correctly execute the fusion.

- ✓ Each domain must be known to the corresponding domain: in the first domain, go to the **Settings** tab of the domain connection and click **Allow**. In the corresponding domain, click **Execute** in the same tab.
- ✓ A full replication task is created on a domain so that all user data can be synchronized between two databases.

In doing so, system data is also automatically replicated, just as with any replication task. System data contains all information about users, groups, rights settings, and so on.

A full replication task and a first replication of system data may take quite long.

3.8 User Management for more domains

 If a replication was made at least once with another domain, both domains are displayed as separate folders in the **User management** area of the Management Console. Each folder of a domain then contains a **User** folder.

When several domains with their own users and groups of users replicate with each other, you may have to adjust rights. When, for example, users and/or groups of users from domain A are to access data of users and/or groups of users from domain B, you must change external access rights.

- ✓ By clicking the respective folder of one domain, you switch to this domain.
- ✗ If you click directly on the domain folder, the **Delete** button is displayed on the right side. Clicking this button removes the displayed domain from the network. This is only recommended if you do not want to replicate with a domain anymore or want to newly set up the domain and use the same name again.

By deleting a domain from the network, the domain connection to this domain and existing replication tasks will also be deleted. Internal system data is removed from the domain. Only then it will be possible to use the name of a domain again.


- ✓ The administrator of domain A can also change the properties of users of domain B if he or she knows the user name and password of the administrator of this domain.

This does not apply to the user account of the administrator of a laptop. You cannot change rights or settings of users with a laptop administrator account. This also applies if the password of the main installation is known.

- ✓ Users are always assigned to one specific domain. This assignment can be changed, see "Assigning a domain" on page 41.

3.8.1 Assigning a domain

User accounts can be assigned to a different domain. The rights are adopted.

 If at least one replication has been performed with another domain, the **Assign domain** button is displayed in the **User management** module.

This function can only be used if a second domain is available, a replication was already performed and the connection is still active.

To move user accounts to another domain, you must be logged on as administrator to the Management Console. User name and password are thus already known.

- ✓ Clicking the button opens the **Assign domain to users** window.

In the upper pane of the window, you will find information about the domain to which you are currently logged on.

The current domain is entered as domain name.

In the lower pane, you enter information about the domain to which you want to newly assign the selected user accounts.

- ✓ First, select the domain
- ✓ and then enter the administrator password of the other domain.

Each domain has a separate Management Console. Thus, in each Management Console, an **administrator** user with the respective rights exists. If users are assigned to another domain, the user managers of the second domain are responsible for these users.

- ✓ It is therefore important that both administrators enter their password in this window or that each administrator knows the password of the other one.

3.8.2 Participants are assigned to domains

The replication process does not change rights or external access rights to data records. If you want to enter a user from a different domain as participant in a data record, the corresponding external access rights must exist for this user.

Usually, the administrator defines external access rights of users among each other. Users can define external access rights themselves in the **Tools** window of the Desktop Client.

In the **Miscellaneous** area of the Management Console, at **Participants/Resources**, the administrator defines whether the domain is displayed for users in the drop-down list.


When logging on to an external domain, you enter your default name and password. If the user Schmitt from the Karlsruhe domain wants to log on to the Munich domain and enters a name and password, the logon dialog is automatically expanded and he or she can select a domain.

Users of a domain cannot view or edit data of another domain which they cannot access. All public data can be viewed in all domains. If you create a public data record, this data record can be viewed, changed, and so on, by all users of all domains.

3.8.3 Users of other domains as participants

If the replication is deployed in a company, the domain is often displayed when selecting users.

To display users, resources and groups of another domain, you need the required external access rights. The administrator defines external access rights to users and/or groups of other domains in the Management Console, see "User Management for more domains" on page 41.



- ✓ The **Select participants** window includes the **Domain** drop-down list. When you select an entry from the drop-down list, all users, resources and groups of this domain are displayed.

In drop-down lists for the selection of users, entries are displayed like that: Domain\User name, for example, in the **Processed by** field of tasks, in the **Person responsible** field in addresses, and so on.

In CAS genesisWorld Web, the Domain\User name is displayed when selecting participants.

The entries are only displayed in the drop-down lists, if the administrator has permitted the display of domains in the selection lists in the Management Console.

3.8.4 Participants in a domain connection and partial replication tasks

Go to the **Authorization** tab in the domain properties to define which participants can view or change this domain connection. Rights can only be assigned for your own domain. These rights are the same as elsewhere in CAS genesisWorld.

If you click **Filter properties** in a partial replication task, the **Properties** window opens. Go to the **Team** tab to define the users whose data records you want to replicate.

You can only replicate data records to which you have access rights.

4 Domain connection

A domain connection has to be set up for two domains to be able to exchange data. A domain can be connected with various other domains. The main installation, that is, the installation in the head office of a company, can be connected to the employees' laptops or to subsidiaries via several domain connections.

A domain connection can connect two domains so they can exchange data.


A domain connection always refers to one direction: if you want to replicate data between two domains bidirectionally, you must set up a domain connection to each corresponding domain on each of the two domains. Domain A needs to be connected to the second domain B. And domain B needs to be connected to domain A to the other direction.

A connection between 2 domains is a logical connection, but it must be set up as separate domain connection in each domain. This is done because the properties in both domains differ partly. You only see the properties of the domain connection to which you are currently logged on. For example, the Web Service for sending data is configured differently for the main installation and the subsidiary. Other properties should be identical in two domains: for example, the function to replicate the Document Archive on two domain connections with their corresponding domains should be either activated or deactivated.

Direct or indirect cycles by connecting several domains to a network must be avoided, see "Replication licensing" on page 11.

If you want to replicate data, you exchange it using domain connections. The Web Service of an application server is used for the exchange. In the domain connection, you define which application server and/or which Web Service of which application server is used for a domain connection.

4.1 Setting up domain connections

 In the Desktop Client, go to the **Replication administration** window or the **Replication** area of the Management Console and click **File > New domain connection**.

The **Create new connection** window opens.

- ✓ Enter the name of the corresponding domain. This name was assigned when the corresponding domain was created.

Afterwards, the automatically entered connection properties are displayed and you can check and change them if required.

As administrator, you can already set up a domain connection in a main installation for the corresponding domain even if the corresponding domain has not been set up yet. This is helpful if you set up one or more laptops for replicating data. If you then create a domain on a laptop, you enter the name of this new domain in exactly the same way as you did when setting up the connection of the main installation, see "Special requirements for laptops" on page 16.

4.2 General tab

This tab displays general information about the domain connection which you cannot change.

Receipts are internal replication IDs which help the Support team when analyzing requests or errors.

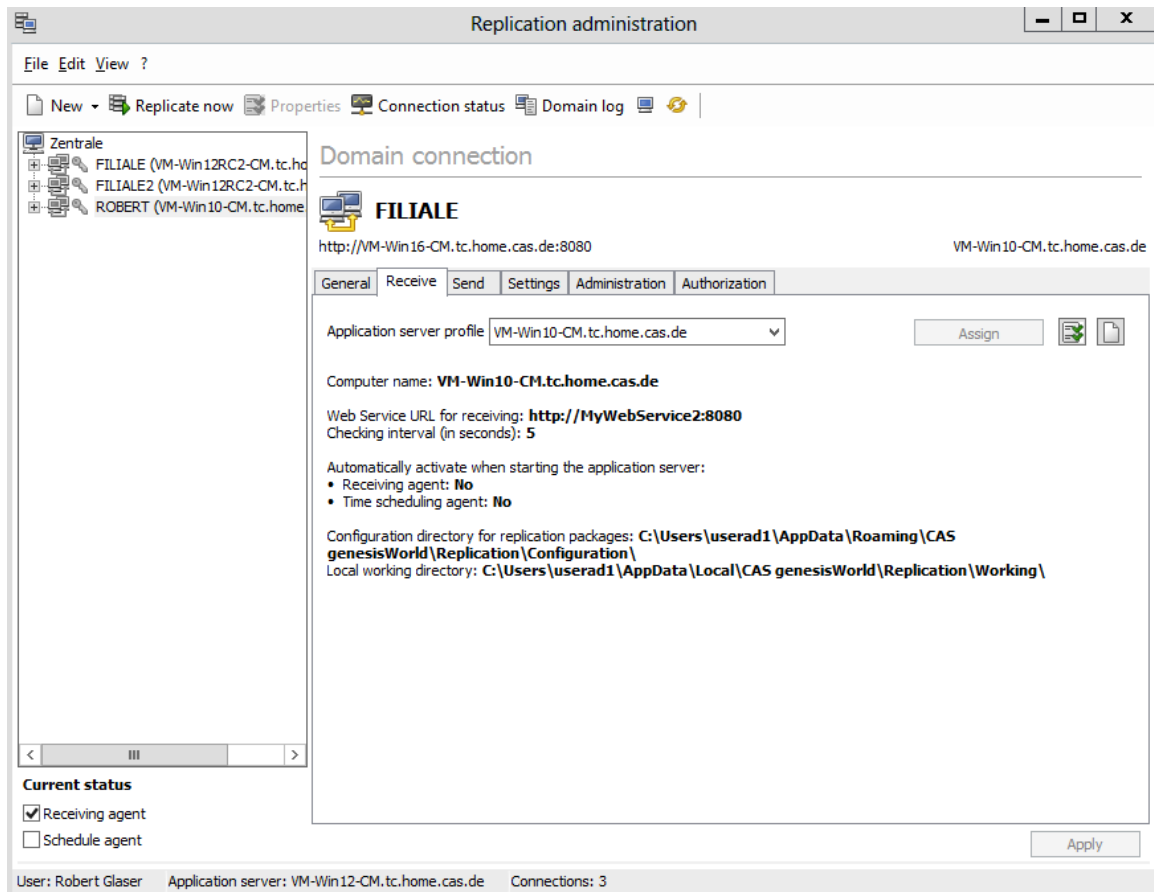
4.3 Receive tab



An application server profile is assigned to each domain connection.

By doing so, the application server and/or the corresponding Web Service of the application server as well as the folder via which replication packages are received is defined, see "Creating/editing application server profiles" on page 19.

If an application server profile has been assigned, all profile settings are displayed in the **Receive** tab.

You can only change the settings of the application server profile or assign another profile if you are logged on as administrator.



- ✓ In the **Application server profile** drop-down list, each profile which has been set up can be selected.
- ✓ With **Assign**, you define the selected application server profile from the drop-down list as profile to be used.
-  **Properties** opens the **Application server** window for the set application server.
-  **New** opens the **Application server profile** window. A new profile is created for the application server which is set in the drop-down list.

4.4 Send tab

A set up domain connection is always configured for one direction of a domain. A respective connection is also available for the corresponding domain of each domain.

In this tab, you define the settings for the sending of replication packages, check a domain connection and authenticate the domains after you have created a connection.

The application server profile contains the Web Service information for the application server via which replication packages are received, see "Creating/editing application server profiles" on page 19.

In the **Receive** tab, an application server profile is assigned to a domain connection. With this assignment, the Web Service for the domain connection via which replication packages are received is also defined, see "Receive tab" on page 42.

In the **Communication** tab of a domain connection, you define to which Web Service replication packages are sent.


- ✓ Activate the **Use the Web Service setting of the assigned application server** option.

If this option is ticked, the same Web Service which has been set for receiving in the application server profile is now also used for sending.

This is done for laptops, as they receive and send data via the Web Service of a main installation or subsidiary, see "Main installation, subsidiary and laptops" on page 8.

- ✓ Deactivate the **Use the Web Service setting of the assigned application server** option for the main installation or the subsidiary.
- ✓ Enter then the Web Service of the corresponding domain in the **Web Service URL for sending** field.

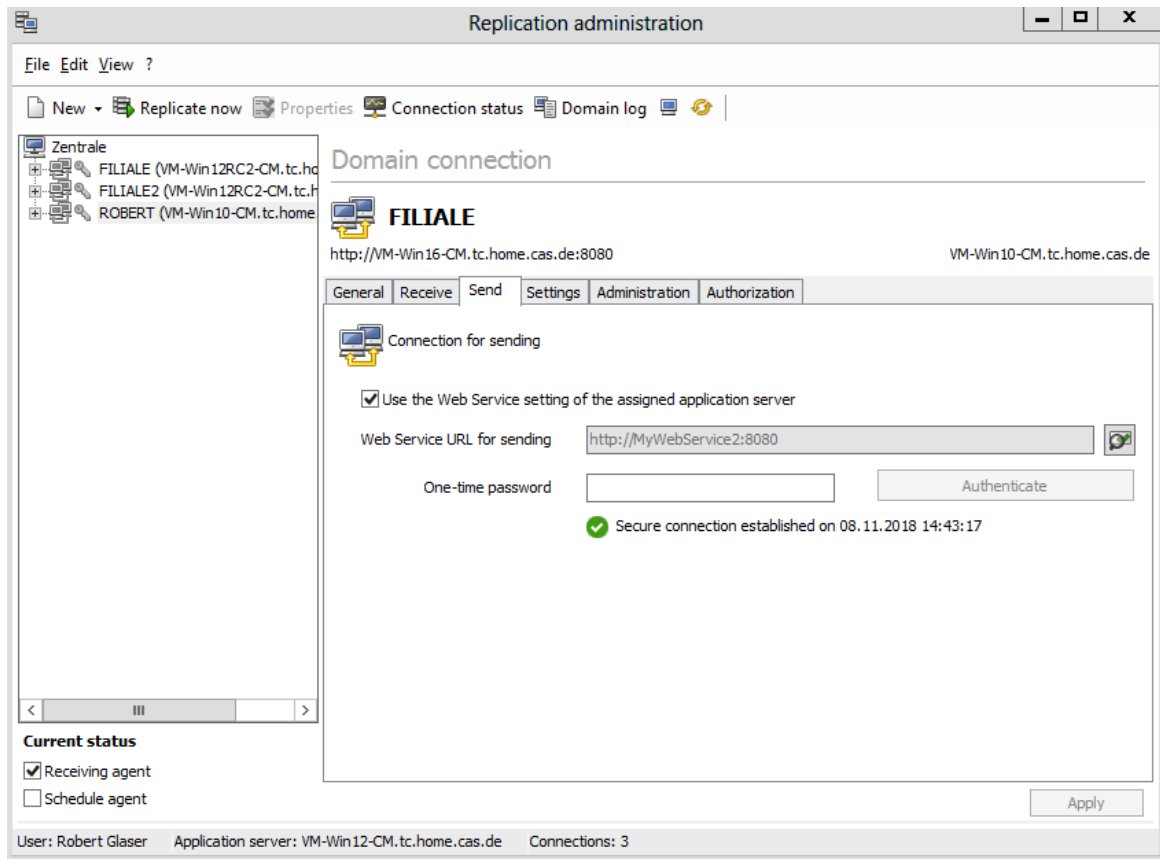
Data is then sent to the Web Service of the corresponding domain via a main installation or subsidiary. The receiving of data in a main installation or subsidiary is then monitored via their respective own Web Service. The respective Web Service has been defined in the application server profile of the main installation or subsidiary and has been assigned to the current connection.

-  With this button, you check whether a domain connection between the two domains can be created.

We also recommend to check the unidirectional connection as the corresponding domain sends a confirmation.

If a domain connection cannot be established, please check and compare the settings in the Web Service, for example, the setting for ports.

Even if the Web Service in the application server profile of the main domain can be accessed, the sending or receiving of replication packages may not be possible.



Initial mutual authentication of domains

To be able to replicate on an authorized domain connection, the matching corresponding domain must be correctly identified by the main installation. This first authentication is manually enabled by entering a one-time password when setting up the connection.

Keys will then be exchanged between the domains and the validity of a connection is checked each time before data is to be replicated. This authentication is made manually.

The authentication is triggered from the corresponding domain for the main installation.

Steps for the main installation

- ✓ Log on to the central domain and open the **Administration** tab of the domain connection for the corresponding domain.

Domain connections are displayed for all domains involved. If you are, for example, logged on to the domain of a main installation, you see all connections to all subsidiaries and all laptops. If you are logged on to the domain of a laptop, you see the connection to the main installation.

In the domain of the main installation, you see the properties of the domain connection to the subsidiary or the laptop, see "Domain connection" on page 41.

- ✓ Click **Generate one-time password**.

The one-time password is only valid until the next time the application server of the central domain is started.

- ✓ Click **Copy password** and save the password to the clipboard or copy it, for example, to a text editor.

Steps for the corresponding domain

- ✓ Log on to the corresponding domain and open the **Send** tab of the domain connection of the main installation.

In the domain of the subsidiary or the laptop, you see the properties of the domain connection to the main installation.

You are, for example, logged on to a subsidiary or a laptop: in the **Send** tab, you manually enter the information with which the subsidiary or the laptop are authenticated for the main installation.

- ✓ Enter the **One-time password** generated in the main installation from the clipboard or the text editor and click **Authenticate**.

With that, the corresponding domain (subsidiary or laptop) is initially authenticated in the central domain.

Afterwards, the **Secure connection established** notification is displayed.

Starting with the initial replication, keys will be used for replicating data, see "Creating/editing application server profiles" on page 19.

4.5 Settings tab

In this tab, you define how and which data you want to replicate on which connection.

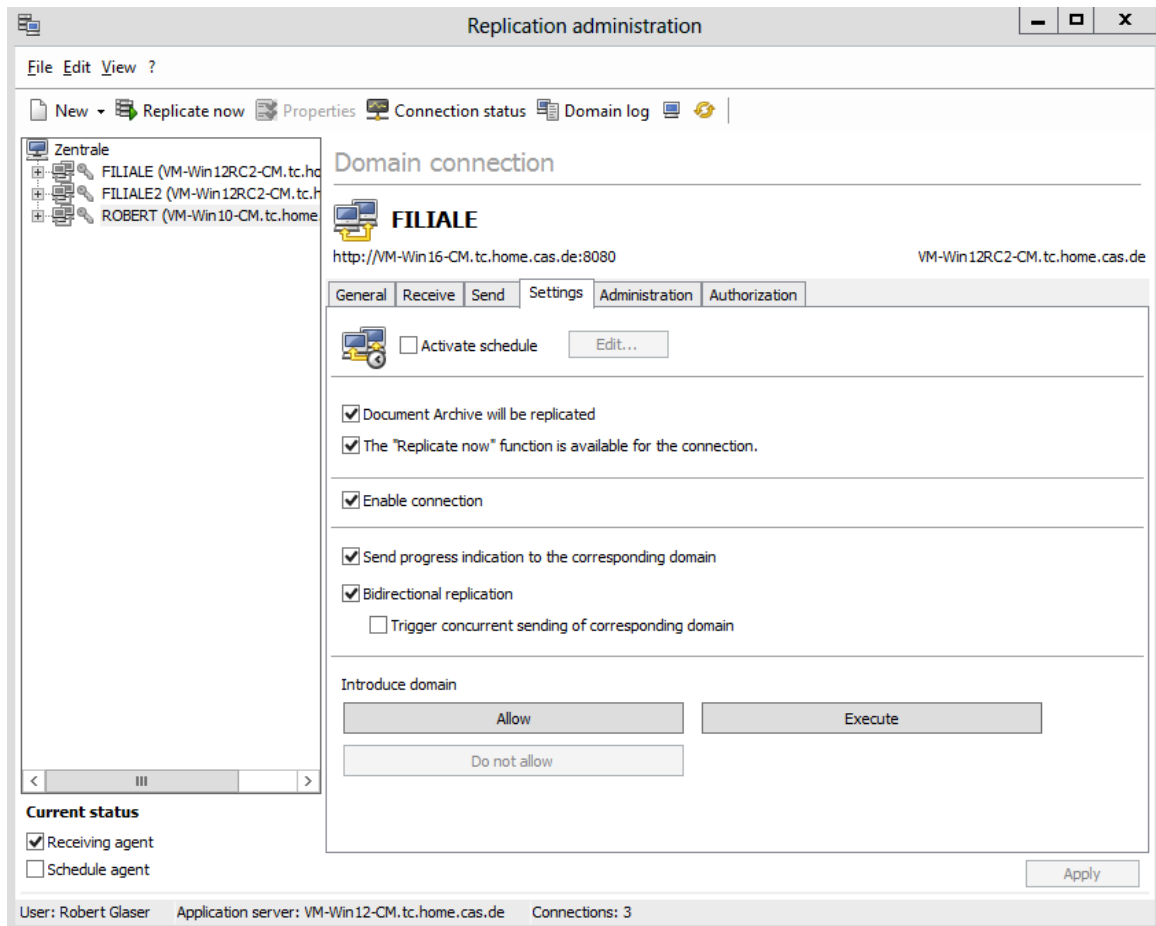
- ✓ With **Activate schedule**, the replication is executed according to the predefined schedule; you make settings for the schedule via the **Edit** button.

If the **Replicate now** function is available for a connection, you can always start a replication. It is not important whether or not a schedule has been defined.

- ✓ If the **Document Archive will be replicated** function is not activated, you can open documents, but you cannot open nor edit the archive files of these documents.

If you activate this option, the replication of data will take longer.

This setting applies to all tasks of the connection.



You can select the setting for replicating the Document Archive in both domains. A different setting for the two domains is not recommendable, as the archive files will be replicated, but the contents cannot be kept synchronous.

- ✓ **The "Replicate now" function for the connection is available** means that this connection function can be called via a button displayed in the toolbar.

Each domain connection is assigned to an application server via an application server profile, see "Administering application server profiles" on page 17. The **Replicate now** function can only be executed if you are also logged on to the application server for the domain connection.

- ✓ The domain connection is activated with **Activate connection**.
- ✓ With **Send progress indication to the corresponding domain**, you can, for example, see if data has been sent or received or if an error has occurred.
- ✓ **Bidirectional replication** defines how you want to replicate data via this connection.

Between two domains, a connection usually exists in one direction and a second connection in the other direction, see "Domain connection" on page 41.

A replication task is replicated bidirectionally: both domains send the changed data records to the other domain via the respective connection. Once the data has been replicated, the data of the replication task is identical on both domains.

With a unidirectional replication, the changed data of a replication task is only sent from one domain to another. With this option, edited data can for example be transferred from a laptop to the main installation. But the changes made on the main installation are not transferred back to the laptop.

When replicating data unidirectionally, replication tasks are sent only via one domain connection and only in one direction. The second connection still needs to be active in order to send messages.

When creating unidirectional replication tasks, different replication tasks can be created on two domains as they are only sent in one direction.

- ✓ **The Activate concurrent sending of the corresponding domain** option can only be set for a bidirectional replication. If this option is activated, a message is sent to the corresponding domain before replicating. When the corresponding domain has received this message, the corresponding domain starts sending the data.

This means that a concurrent sending process is triggered for both domains so that the replication is performed faster.

This setting is not considered if a connection has not replicated yet. The concurrent sending of a corresponding domain can only be made if at least one replication has been performed.

- ✓ Each domain must have been introduced to its corresponding domain: please always remember the following step when you start with an empty database or when fusing databases.

Click **Allow** in the first domain and then **Execute** in the corresponding domain.

The corresponding information already exists in a copied database.

- ✓ You enter the name of the database for the corresponding domain in the following window. When entering this information, the Web Service of the corresponding domain is also identified.
- ✓ Afterwards, you execute the initial authentication of the domains for each other, see "Send tab" on page 43.

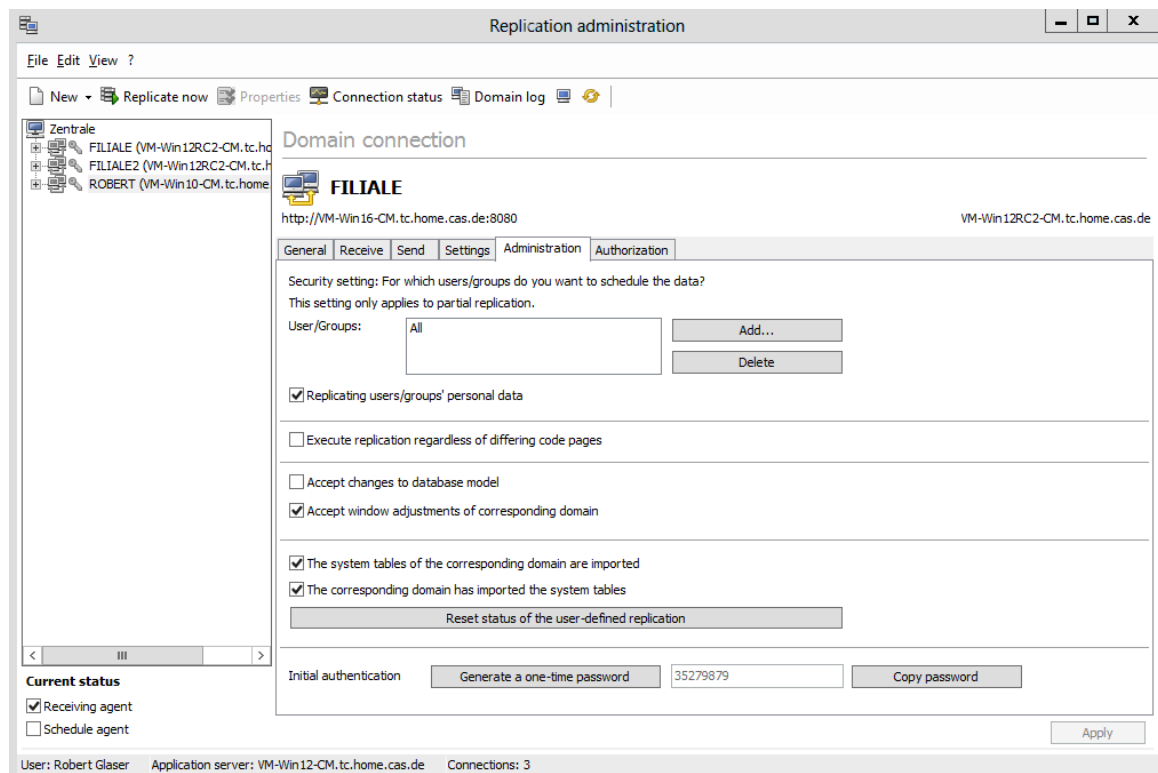
4.6 Administration tab

Only administrators can configure this tab's settings.

If a replication task is created by a user, specific data records are selected. Only those data records to which access rights exist can be selected.

With a replication task, not only the selected data records are replicated but also the linked data records (first level) if the corresponding settings were made in the replication task. Rights can differ for linked data records.

Here, you as administrator define which data of users or groups can be selected for a replication task and if you want to add personal data. With these settings you prevent that data to which users have no rights is replicated. This also applies to archived documents.



Examples

To connect to the laptop of the Head of Marketing, you select the Head of Marketing and the marketing group. The Head of Marketing has full rights to this group. The Head of Marketing creates a **My projects** replication task for this connection. Her projects and all linked data records are then replicated, that is, the data records of the Head of Marketing and the marketing group.

For the connection with a sales employee's laptop, select this employee. Only his data records and the linked data records to which the employee has external access to are replicated in his replication tasks for this connection.

Functions

- ✓ With the **Add/Remove** buttons at **Users/Groups**, you define users and/or groups for this connection.

When calculating and displaying the data volume for the replication, these users and/or groups are included.

- ✓ If you deactivate the **Replicating personal data from users/groups** option, the personal data of the selected users and/or groups are not replicated.
- ✓ The **Execute replication regardless of differing code pages** function may lead to problems when different code pages are used on the domains.

For additional information, please go to the online help pages of the Database Wizard, to the [Customizing the collation](#) page.

On computers with which data is replicated, a database system, a CAS genesisWorld application server and a Desktop Client are installed. The code page which is used for the database system is set in the Database Wizard. Users can furthermore work in different domains with different regional and language settings which means that different code pages are used when entering data. It may therefore occur that some characters in fields are displayed differently on different domains.

If characters are displayed incorrectly and replication packages with a different code page are received, changes may not be applied correctly or corrupted data may be created.

If this option is deactivated, replication packages with different code pages are not accepted. Activate this option if you still want to replicate data.

- ✓ With the **Accept changes to the database model** option, you can define that changes to the database model are accepted by the corresponding domain. You specify the default value for this setting in the properties of the local domain, see "Settings tab" on page 32.

Changes to the database may be additional fields or columns for data record types.

It is possible that different customizations have been made to the database model in two domains before changing to a new CAS genesisWorld version. If the **Accept changes to the database model** option is activated, these changes are not synchronized retroactively. In this case, the administrator must use the domain replication wizard. If two sides of a connection are based on the same database model, changes to the database model can also be replicated.

Deleting entire fields of a data record type is a special case: This change to a domain is not transferred to one or more domains via replication. The automatic deletion in another domain can result in unwanted data loss, because data deleted in this way cannot be restored. This is why you should delete fields separately on all domains.

- ✓ With **Accept window adjustments of the corresponding domain**, customizations to data record windows made with the Form & Database Designer are transferred from the corresponding domain.

Changes made with the Form & Database Designer must only be made on one domain. The **Accept window adjustments of the corresponding domain** option is then activated on the other domains to make sure that all changes are transferred on all domains.

The system tables of the corresponding domain are imported indicates that the system tables have been transferred on your domain. If you deactivate this option, the system tables of the corresponding domain are requested again and this may take some time. System tables contain among other all information about users, groups, rights settings, and so on.

The corresponding domain has imported the system tables indicates that the other domain has transferred the system tables. If you deactivate this option, this process will be repeated.

The **Reset status of the user-defined replication** function must only be executed if instructed to do so by the CAS genesisWorld support.

Domains are mutually introduced and authenticated with a one-time password in the **Administration** and **Send** tabs of the domain and the corresponding domain, see "Initial mutual authentication of domains" on page 45.

4.7 Authorization tab

In this tab, you define as administrator which participants are allowed to see and change the connection. These rights are the same as elsewhere in CAS genesisWorld.

The screenshot shows the 'Replication administration' window with the 'Authorization' tab selected for the 'FILIALE' domain connection. The interface includes a search bar for participants and a list of participants with their roles and descriptions.

Name	Description
ADMINISTRATC	
All (public)	
Antonio Matara	Produktmanager
Azubis	Auszubildende, Diplome
Besprechungsra	15 Plätze
Britta Glatt	Sekretärin
Buchhaltung	Buchhaltung
David Gruber	Vertriebsassistent
Entwicklung	Entwicklung und Konstr
GF	Geschäfts- und Abteilur
Konferenzraum	30 Plätze, Beamer, Konf
Marketing	Marketing

Current status:
 Receiving agent
 Schedule agent

User: Robert Glaser Application server: VM-Win12-CM.tc.home.cas.de Connections: 3

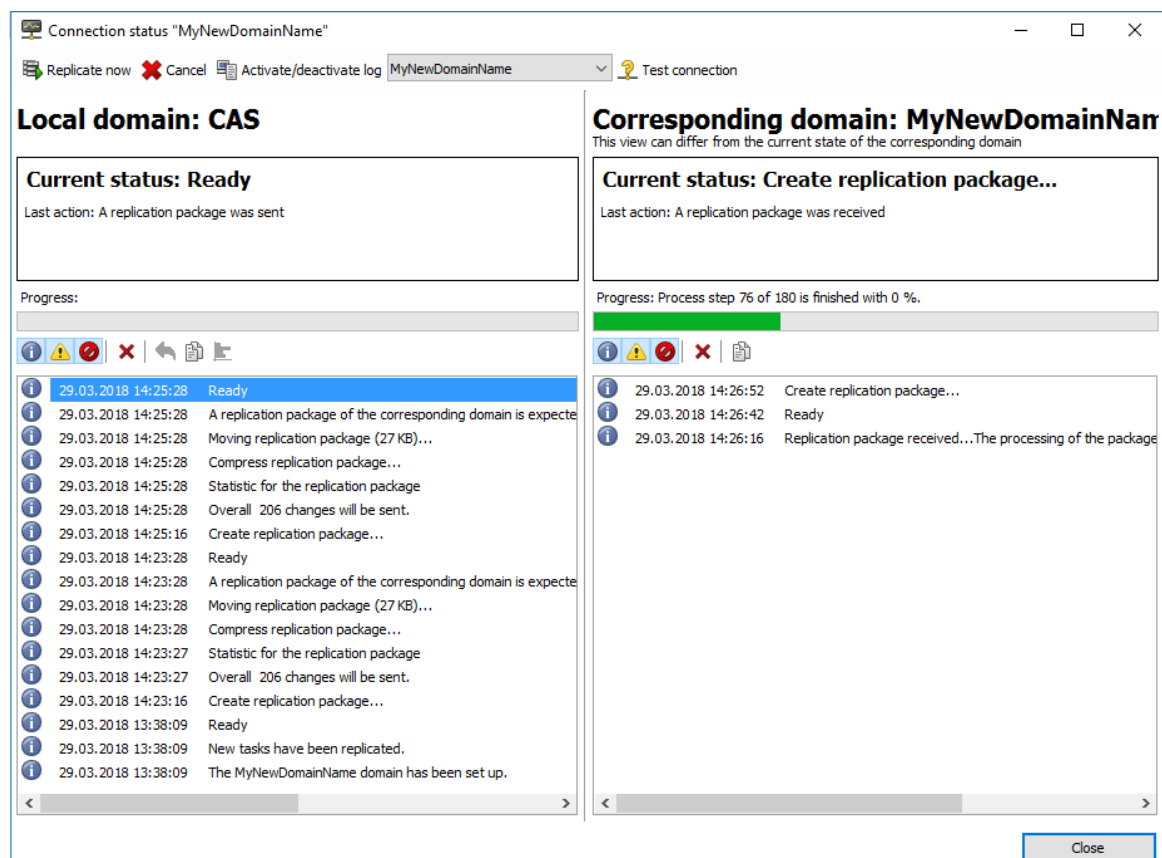
You can only assign rights for domain connections of the domain for which you have the **Administrator** right.

4.8 Connection status window

 In the **View** menu, click **Connection status** or the corresponding icon in the toolbar. The corresponding window opens.

The status of a connection is displayed in the Desktop Client.

This window shows all information about the domain, the corresponding domain, the connection as well as already ongoing or finished replications.



For information of occurring errors, we recommend you to have a look at this log and the **Domain log** window, see "Domain log window" on page 35.

In the upper part of the window, the connection status is displayed on the left-hand and, additionally, the last completed domain action may be displayed. To the right-hand side, the status, the last action and some messages of the corresponding domain are displayed.

In the lower pane, the messages of each domain are displayed. This data is only displayed if the **Send progress indication to the corresponding domain** option is activated; you activate this setting in the domain connection settings, see "Settings tab" on page 46.

The log does not inform about the connection status, but only displays the changes. The current status is displayed in the upper pane.


Notes

The information shown in the right pane may differ from the actual status of the corresponding domain depending on when the messages are transferred from the corresponding domain. If the connection is interrupted or not all messages are processed, the information on status and messages might not be complete. If the sending of messages is deactivated in the corresponding domain, no information is displayed in the right pane.

Progress bars show the progress of the replication process for each domain. Some steps of the replication are very complex and thus require some time. The progress bar may not move and will increase as each phase completes.

Buttons in the upper part of the window

 **Replicate now** starts the replication for the selected domain connection.


 **Cancel** ends an ongoing replication, for example, if errors occur which cannot be fixed. The status will then be reset to **Ready** and you can restart the replication.

If the status is set to **Ready**, the data to be replicated has been compiled and sent.


If you shut down the computer during an ongoing replication, the sending or receiving process will not be activated again when you restart the computer.

When starting the computer, the application server is also started and the status is set to **Ready**. Now you must start the replication again with **Replicate now**.


The **Ready** status does not mean that the corresponding domain has answered.


 **Activate/Deactivate log** activates or deactivates the update of the log. If the log is deactivated, messages are not lost. As soon as you enable it again, all messages that have not been displayed before are included in the log.


✓ With the **drop-down list of the connection**, you switch between each connection of the domain.

 **Test connection** tests the displayed domain connection.

Buttons in the lower part of the window

 With these icons, you can activate or deactivate the display of status, warning and error messages.


 Use this icon to delete one or several selected messages. If no message is selected all messages are deleted after a security prompt was displayed.


 **Reset replication** is only activated if the status of the domain connection is set to **Ready** after a replication has been finished.

Use this function, for example, after you have imported a backup to a domain.

Select a row in the log which is dated prior to the creation of the backup and where the status was set to **Ready**.

Click then **Reset replication**. All changes are now replicated again, starting from the selected time.

 **Copy selected rows of the log into the clipboard** copies the rows, for example, for support requests.

 **Display replication package statistics** opens a window displaying the tables, the type and the number of transferred entries of a table in a replication package.


Log of corresponding domain


These messages on the left-hand side of the pane can differ from the current state of the corresponding domain depending on when the messages are transferred from the corresponding domain. If the connection is interrupted or not all messages are processed, the information on status and messages might not be complete. If the sending of messages is deactivated in the corresponding domain, no information is displayed in the left pane.

The **Send progress indication to the corresponding domain** option must be activated in the connection properties of the corresponding domain, see "Settings tab" on page 46.

Displaying the connection status


The replication status is represented by colored icons which are displayed on the status bar in the Desktop Client.

 Green: Ready, no errors or warnings since the penultimate **Ready** status.

 Yellow: A warning has occurred since the penultimate **Ready** status.

 Red: An error has occurred since the penultimate **Ready** status.

 The right arrow indicates that data is being sent.

 The left arrow indicates that data is being received.

✓ Clicking an icon opens the **Connection status** window.

5 Replication task

A CAS genesisWorld data record does not only consist of the existing fields in the mask. An address does not only include name, street, town/city, phone number, and so on. But it also includes additional data such as links or rights settings which together can, for example, form a CAS genesisWorld dossier. For the replication function, this means that all data belonging to the address is consistently synchronized so that the address and the dossier have the same content and the same access rights apply at each location.

In a replication task, you define whether the complete data or a subset of data is to be synchronized. A replication task can include the following data:

- Addresses, documents, tasks, opportunities, and so on,
- Document data records including archive files from the Document Archive,
- Templates for data records without links,
- Settings of the administrator such as navigators, input assistance options or view formats,
- User settings, such as navigators, view formats, saved filters, holidays and colors for calendar views, and
- data and settings of the Management Console, such as data record types and fields including groups and the corresponding rights.

Replication tasks are synchronized by CAS genesisWorld between two domains via a domain connection.

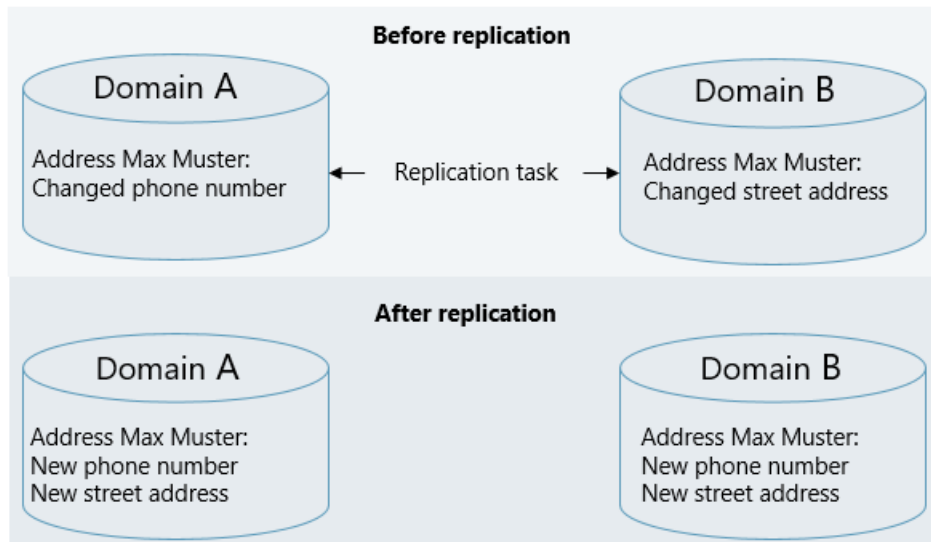
So-called replication packages are sent and received between two domains. A replication package always contains all created replication tasks. Data is usually sent and received bidirectionally, see "Replication process" on page 6.

The replication task for a bidirectional replication is always created on one single domain. After the first replication, an identical task also exists on the second domain.

- Changes to data in the replication task of domain A are available in domain B after the replication, and
- changes to data in the replication task of domain B are available in domain A after the replication.

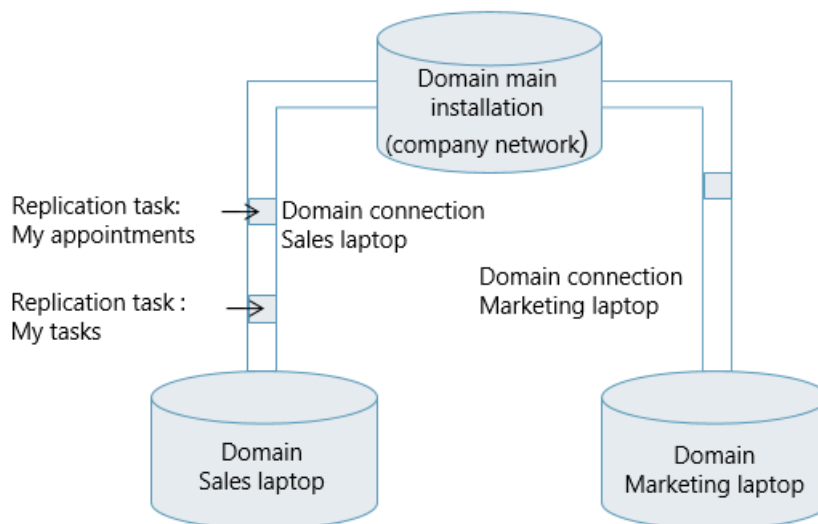
After a replication, the data in the databases of both domains is identical, for example, on a main installation in the office and on a laptop.

- In addition to data records such as addresses, documents, tasks, opportunities, and so on, the corresponding rights are also always replicated. That means that each user has the same access rights in each domain.
- The data is synchronized on field level. Document data records with archive files are treated as special case. In this case, the complete file needs to be transferred.



- **Conflicts** which occur when making changes in two domains at the same time only occur in one or several fields and not in the complete data record.

You can create one or several replication tasks for each domain. You define certain settings when creating a replication task.



You can create three different types of replication tasks: Full replication, partial replication or special replication tasks for navigators, archive documents and system tasks.

- When creating a **Full replication task**, all data is replicated. A full replication can be created, for example, if two departments work with the same data and access shared information, including rights and external access rights, complete dossiers, and more. Settings for this type of replication tasks cannot be changed.
- A **Partial replication task** technically corresponds to a **User-defined** task. In a **Partial replication task**, the selected data for special requirements is replicated. For example, you want to transfer the appointments of the next four weeks or several projects to a laptop. The data record type and a selection of data records for this type are defined.

Furthermore, you decide whether you want to also replicate linked data records and if so, which ones you want to replicate. If you replicate links all link types are automatically replicated.

Other replication task types may be templates based on the user-defined task, for example, **My current appointments** or **My open tasks** which you can choose when creating a replication task.

For a **User-defined task**, you have to edit all pages of the **Create a new replication task** wizard, but templates are easy to customize. Wizard pages that are not required are skipped. Custom templates can always be used again.

- The following tasks are **Special tasks: Navigator task, Archive documents, Archive documents (templates)** and **System tasks**.

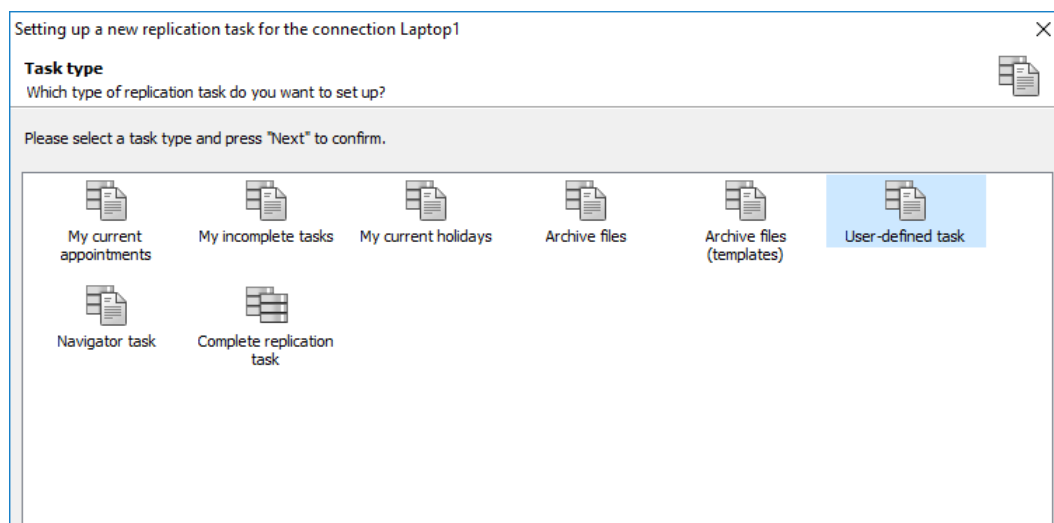
5.1 Creating replication tasks

You set up a replication task with a wizard and select one of several templates.

- ✓ Select the domain replication for which you want to set up a replication task.
- ✓ In the **Replication administration** window, open the **File** menu or the shortcut menu and click **New replication task**.

The **Create a new replication task wizard** opens.

The wizard consists of several dialogs. If you have defined all necessary settings and entered all information on all pages, the replication task is set up.



Afterwards, the properties of the replication task are displayed in several tabs to the right so you can change the properties.

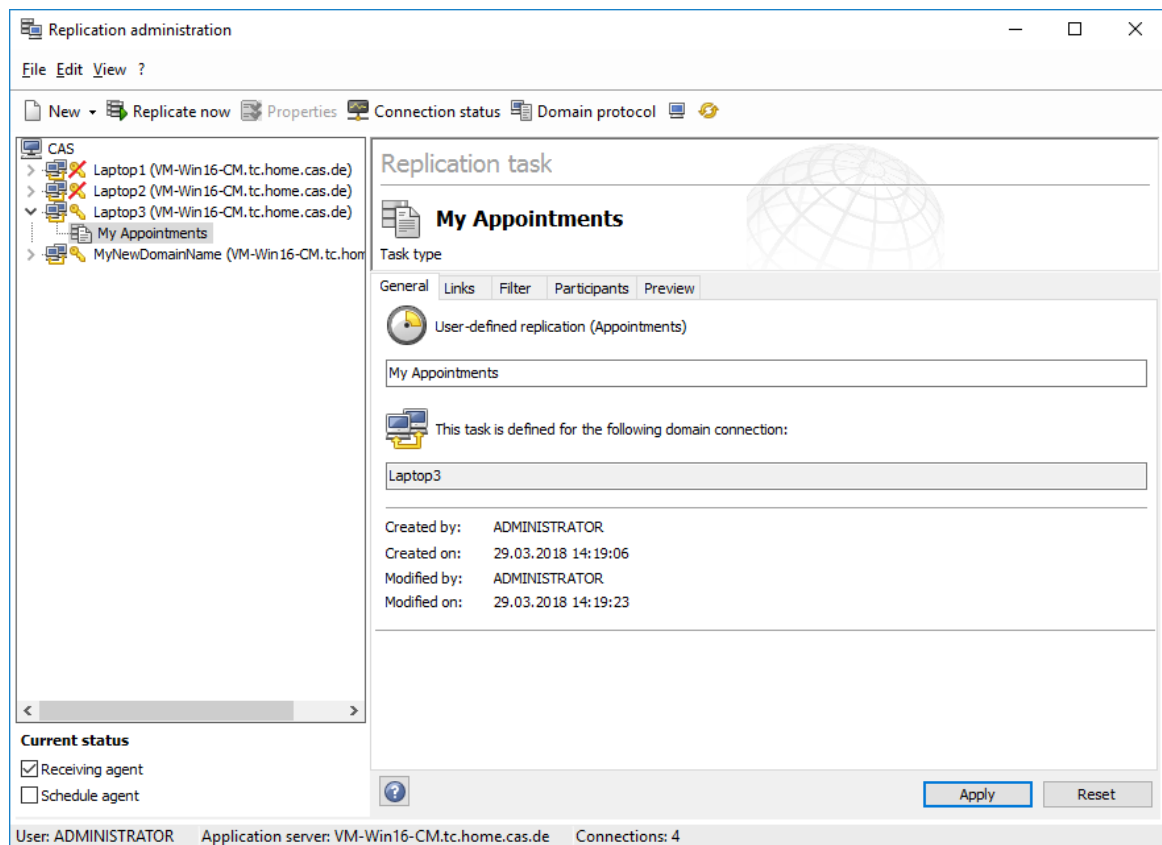
- ✓ On the **Task type** page, you select an entry.

With a **User-defined task**, all settings are individually defined.

My current appointments, My open tasks, and so on, are templates based on the user-defined task for a specific data record type. These tasks can be quickly adjusted to your requirements because you can skip those pages of the wizard that you do not require.

With a full replication task, you replicate all data of a domain with the exception of navigators.

Special tasks exist for navigators, archive documents and system tasks.




Replication task rules

The following rules apply to replication:

- If only user-defined tasks are replicated, all tasks of the connection are included.
- If user-defined tasks and a full replication task exist, the full replication task is executed and user-defined tasks are ignored.
- Special tasks like **Navigator tasks** or tasks for **Archive documents** are always replicated.
- If you use primary link functions and replicate links, the primary links between two linked data records are also replicated.

5.2 User-defined task

 Select this task type if you want to individually select the data you want to replicate. For example, if you want to transfer your appointments of the next four weeks or some documents for a product presentation to your laptop.

Once you have created the replication task, the data records of this task are displayed if the task is selected in the navigator. You can change the selected data records using the **Filter** settings later.

- ✓ On the **Data record type** page, you can select data records or templates of a data record type. If you replicate templates, the links of the templates are not replicated.
- ✓ On the **Links** page, you select whether and, if yes, which linked data records you want to replicate. If you replicate links all link types are automatically replicated.
 - The first hierarchy level of the links is replicated, further links to linked data records are no longer replicated.
- ✓ On the **Filter** page, you can define dynamic and static filter conditions to select the data records of the data record type that you want to add to the task.
- ✓ On the **Participants** page, you are displayed in the **Participants** list by default. If you add additional users, the data of these users is added to the task.
- ✓ The **Tasks, Appointments, Addresses, ...** page displays the data that is replicated. You see a list of data records by clicking **Display task preview**.
- ✓ On the **Description** page, you enter a name.
- ✓ The last page of the wizard contains a summary of your settings. Clicking **Display task preview** displays an overview of the replicated data.
- ✓ Once the replication task is created, the data records of this task are displayed in the **Replication management** window, if the task is selected in the navigator. You can change the **Filter properties** of the task anytime.

My ...

My ... are templates for partial replication tasks. An adjusted template can be reused anytime and only a few settings have to be made.

- ✓ If you have selected one of the templates, the **Links** page appears in the wizard directly after the **Task type** page. These templates are based on the user-defined task.

Inserting data records with drag and drop

You can easily add individual data records from CAS genesisWorld to the replication task with drag and drop.

Please note that the data records you want to insert with drag and drop must match the corresponding data record type. Example: You can only add tasks to the **My open tasks** replication task.

- ✓ Open the corresponding list view in CAS genesisWorld.
- ✓ Click on the desired row and drag the data record with the pressed mouse key in the list of the **Replication management** window.

This data record is now added to the replication task.

5.3 Special tasks

With a **Navigator task**, you can replicate own or public navigators. For the replication of navigators, a separate replication task must be created. Even when replicating data with full replication tasks, navigators are not included in these replication tasks.

Navigators are always replicated entirely.

The list displays all navigators to which you have access and which are saved in the database. If you are logged on as administrator, all navigators are displayed.

- ✓ By clicking the checkbox or using the buttons below the list, you select the navigators that you want to replicate.

The document template navigator is always replicated as only one document template navigator per user is available.

For program and link navigators, you either select all or certain navigators.

Tab groups you have saved in a navigator are also replicated.

- ✓ The settings for the navigators of a user are replicated with the task, for example, the start navigator and additional navigators.

With tasks for **Archive documents** and **Archive documents (templates)**, you can replicate document data records with archive files.

- ✓ One option to replicate archive documents is the **Document Archive will be replicated** option for domain connections. If you activate this option, all archive documents will be replicated.
- ✓ Alternatively, you create special tasks if not all archive documents are to be replicated.

▮ The links of data records in templates are usually not replicated.

If archive files are replicated, the replicated versions of the archive file are added to the existing versions and all versions may be newly numbered. If the archive file of the document data record has been changed on both domains or if the archive file has been checked out at the time of replication, a backup version of the document is created. To view a version, go to the **Edit document** window, **Versions** tab.

System tasks are automatically created if the laptop automatic is activated and new data records are not included in a replication task on a laptop. These new data records are then automatically replicated with each future replication.

Users are not permitted to change or delete these tasks.

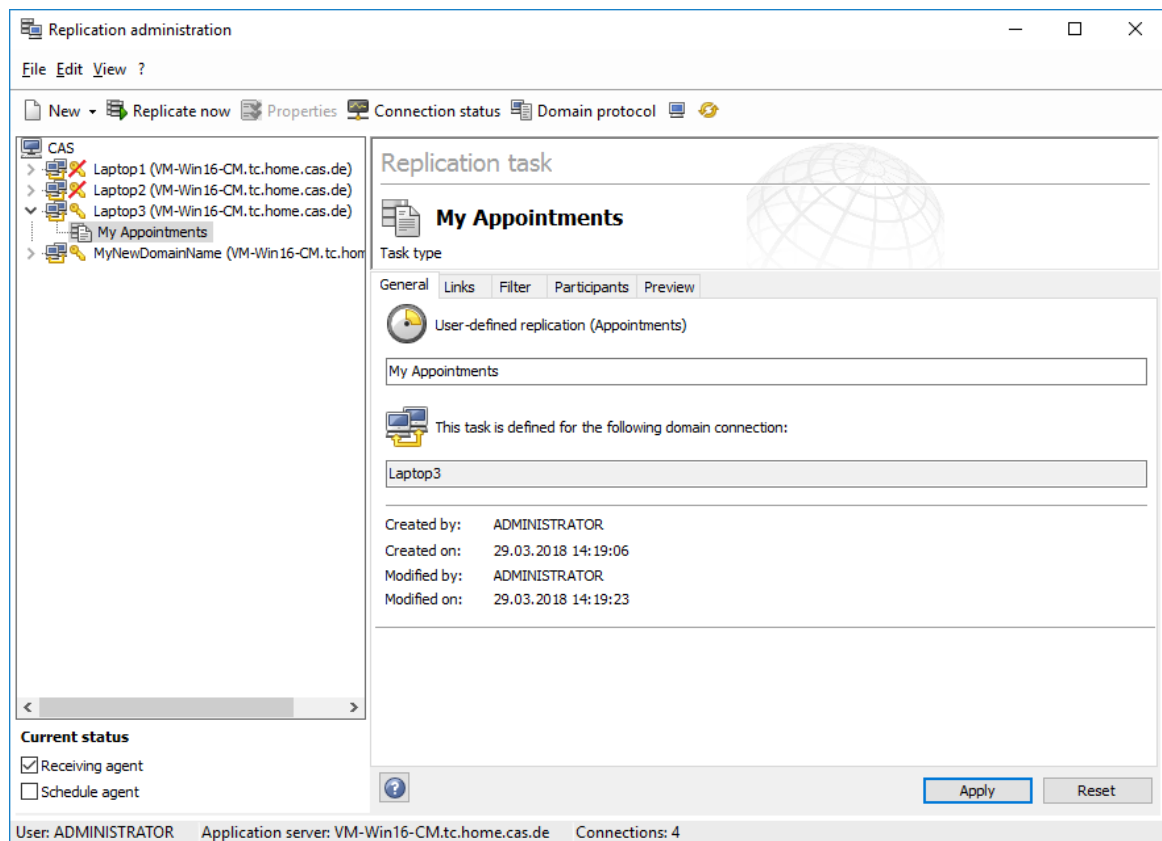
5.4 Full replication task

With a full replication task, you replicate all data of a domain with the exception of navigators.

- ✓ Enter a name for the replication task in the **Name** dialog.

5.5 Replication task properties

To view or change the properties of the replication task, select the task in the navigator of the **Replication administration** window. The properties are displayed in the right pane.



You can change all entries which you have defined when creating a task with the wizard in this window, see "Creating replication tasks" on page 57.

When replicating between two domains, a domain connection must have been set up on each domain. The replication task is only created on the domain which serves as master domain. After the first replication, the task also exists on the corresponding domain.

- ✓ Changed settings in the tabs are immediately displayed in the **Preview** tab. The changes are only saved in the task if you click on the **Apply** button.
- ✓ Use the **Reset** button to display the saved settings applied to the task before.

General tab

The **General** tab displays the type and the name of the replication task, the corresponding domain connection and the creation and change date of the task.

In a full replication task, only the **General** tab exists.

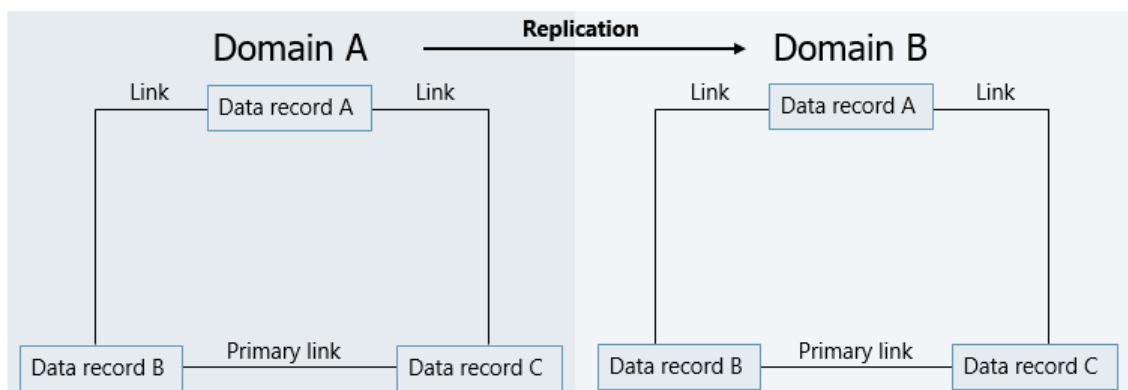
- ✓ You can change the name of the replication task by overwriting it.

Links tab

Replicated data records also include information about other data records that are linked to these replicated data records. Linked data records are only replicated if you define the corresponding option in the links tab of the replication task.

- ✓ Select whether you want to also replicate linked data records and if so, which ones you want to replicate.

Link types and primary links are also automatically replicated if links are replicated.



Filter tab

With filters, you define which data you want to add to the task.

- ✓ As with the creation of a replication task, you can add certain data records to the task with drag and drop.

Participants tab

If required, you can add additional users and the data of these users is added to the task, see "Replicating data of a number of users" on page 64.

Navigators tab

This tab is displayed for navigator tasks.

- ✓ By clicking the buttons or ticking the corresponding checkboxes, you define which navigator is to be replicated.
- ✓ An administrator can replicate all navigators in one task.
- ✓ Users can only replicate navigators which they are allowed to access.

Preview tab

An overview informs users of the data that has been replicated for user-defined replication tasks and tasks for archive documents.

- ✓ If you change a task in the other tabs, you can see the results in this tab.
- ✓ Users only see data in the preview to which they have access rights.
- ✓ The rights of users and groups in the **Administration** tab of the domain connection restrict the number of data records which may be included in the replication tasks of the connection.

5.6 Which replication task is the right one?

To make personal data of a user available on the laptop, a user-defined task is usually created. Depending on the requirements, the matching data records are then available:

- for the sales team,
- for customer presentations with appointments,
- open tasks,
- presentation material, or
- links to addresses of customers.

The same applies if you want to replicate, for example, addresses of a certain postal code area (7*) or specific projects with links.

If a full replication task is created, all data between two domains is replicated and synchronized with the initial replication; in further full replications, all changes are replicated.

The change volume and thus the transfer time of full replications is relatively high. But the full replication is also created and processed quickly: the system does not check which data records are replicated.

The full replication can also be used on a laptop if a LAN, RAS or VPN connection is available. The shorter the time between replications is, the less data changes have to be replicated.

5.7 Special considerations for full replication

When performing a full replication, you can only create a replication task on one domain. After the first replication is performed, the task also exists in the second domain. If you start with the smaller domain the extent of data to be replicated is less.

If a full replication of two domains is performed, the input assistance options of one domain are completely transferred when replicating for the first time. Input assistance options of the second domain are overwritten.

There is no logic which defines from which domain the input assistance options are transferred. The input assistance options of the domain with the smaller domain SysDomain.GGUID ID are transferred.

We recommend to ensure that input assistance options are identical in both domains before replicating data. In the **Database** area of the Management Console, you can export and import input assistance options.

If a new domain is created from a copied database, input assistance options are not overwritten.

5.8 Replicating data of a number of users

If participants are entered in a replication task, all data records of these users are replicated.

External access rights are not considered for replication tasks.

This also applies if groups are entered as participants in data records:

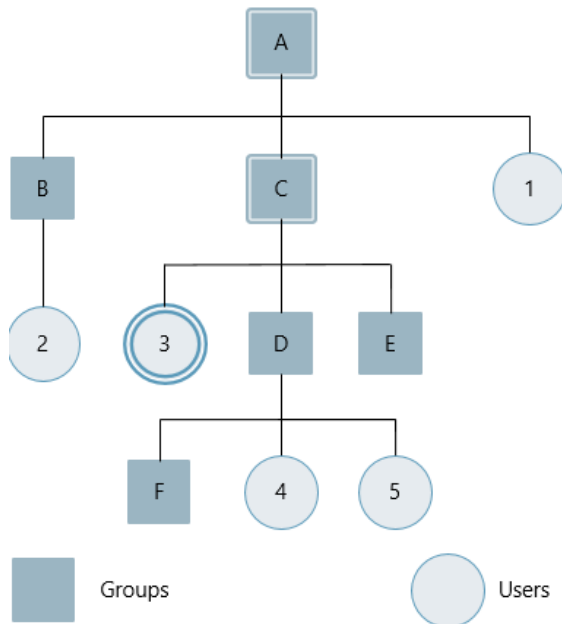
- A user is entered as participant for a replication task.
- The user is member of a group.
- With the replication task, all data records in which this group is entered as participant are replicated.
- This also applies to hierarchical groups, that is, if a group is member of another group.

The same also applies if you add a group as participant to a replication task:

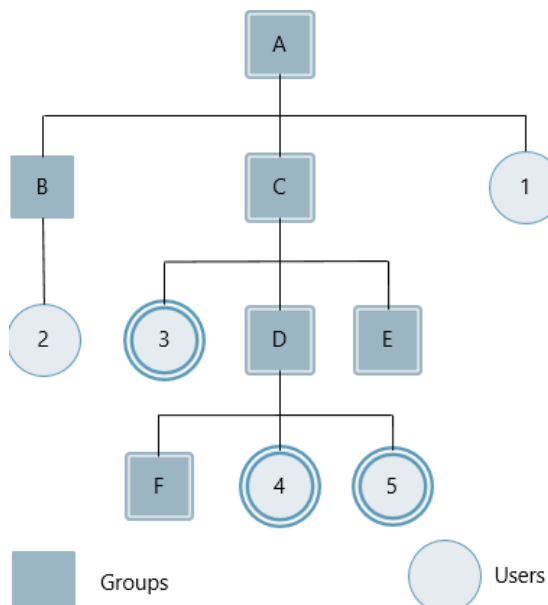
- A group is entered as participant for a replication task.

- With the replication task, all data records are replicated in which the group is entered as participant.
- Furthermore, all data records in which users of the group are entered as participants are also replicated.
- This also applies to hierarchical groups.

In the following example, user 3 is entered as participant in the replication task. All data records which have a user or a group with a double border as participant are replicated: this applies to groups A and C, as well as to user 3.



In the following example, group C is entered as participant in the replication task. All data records which have a user or a group with a double border as participant are replicated: this applies to all groups with exception of B and the users 3, 4 and 5.



5.9 Deleting the domain connection or replication task

We recommend to not delete the domain connections and replication tasks anymore. When setting up a new connection and a new replication task, a full synchronization of data is executed when users replicate for the first time. With each further replication, only the changes are transferred.

We recommend to delete a replication task only in the following cases:

- When the first replication is performed, the log of the local domain saves the information that different replication tasks have been used, that is, that a task has been created on both sides of the connection. One of the tasks must be deleted so that the other side of the connection can try to perform a new replication. The logical connection then refers to the same task.
- User data must again be replicated if, for example, data could not be transferred. Replication tasks which – after the first replication – exist on both sides of the connection must be deleted. A new replication task is then created on **one** side and will be replicated.

6 Replicating data

Usually, each user can set up a domain connection or create a replication task and also make replications. You need the **Replication** right which is set by the administrator in the Management Console in the user account properties. The domain is always configured by the administrator.

As a replication affects the company's data, sometimes also the data of a number of subsidiaries, contact your administrator to find out if you have the corresponding rights, who is in charge of executing replications and at what time.

For example, each user could be responsible for replicating his laptop but replications with other locations are only run by the IT department at certain intervals. Exactly how you decide to do things very much depends on the processes and regulations being used in any given company.

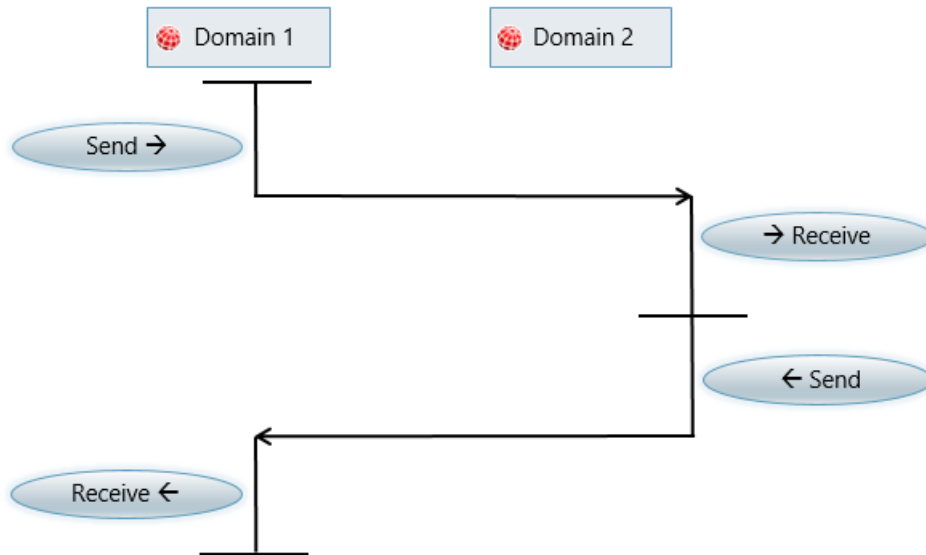
If you use the replication function, you can work with your laptop both locally and in the company network. The title row of the program window displays the domain to which you are logged on.

All functions to configure and execute a replication can be found in the **Replication management** window or in the **Replication** module of the Management Console, see "Replication functions" on page 21.

6.1 Concurrent replication

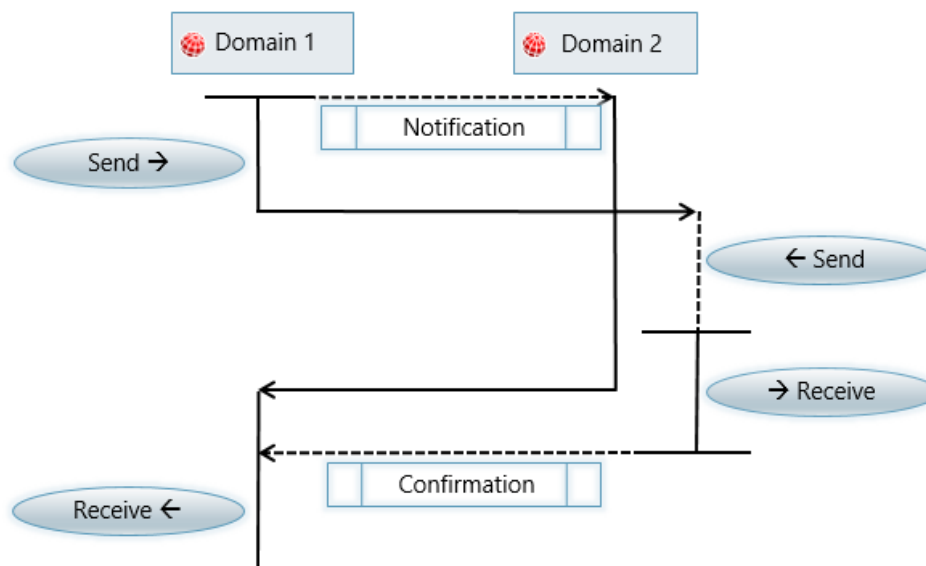
Partial concurrent sending can be applied to shorten the time it takes to complete a replication task.

As a rule, the replication process is run in sequences as follows: Sending – Receiving by the corresponding domain – Sending by the corresponding domain – Receiving.



In the **Settings** tab of the domain connection, the **Activate concurrent sending of the corresponding domain** option can be activated, see "Settings tab" on page 46.

Both domains will then send and receive data at the same time and the replication will then be executed as explained in the following:



- The first domain sends a message to the corresponding domain to trigger the sending process

- The first domain sends data and the corresponding domain also sends data at the same time (concurrent sending)
- The first domain and the corresponding domain both receive data at the same time

While the first domain packs the replication task, the corresponding domain can pack its own replication task. This means that the corresponding domain can start earlier with the transfer of the replication task. And the replication tasks received on each side can be processed simultaneously.

If you have not replicated on a connection yet, the **Activate concurrent sending of the corresponding domain** option is ignored. The concurrent sending of a corresponding domain can only be made if at least one replication has been performed.

6.2 Replicating data from navigator views

Data records of a program navigator view can be replicated. This applies to lists and dossiers as well as to e-mail lists. A special partial replication task for the selected connection is created if you mark the view for the replication. Using this function, you always replicate data records of a single view.

Notes

For this function, you require the necessary rights from the administrator: the right to replicate data and the rights to the navigator.

Users must be entered as participants in the view properties of the navigator.

The "**Current user**" variable must not be used in the filter settings for the view.

- ✓ Load a non-public navigator. If you have loaded a public navigator, go to **File > Navigator > Save as** and save the navigator as private navigator.
- ✓ Select the view that you want to replicate. Click **Properties** in the short cut menu.
- ✓ In the **General** tab of the **Properties** window, you select a domain connection from the **Replication** drop-down list.
- ✓ Make sure that the filter condition "**Current user**" is not selected in the **Filter** tab.
- ✓ Go to the **Team** tab and make sure the **Display data records of selected users if access rights exist** option is activated. Furthermore, one or several users must be entered as participants in the right list.
- ✓ Go to **File > Navigator > Save**.

A partial replication task is entered on the selected domain connection for this view.

The view is marked in the navigator and you will recognize the replication mark.

✘ With this icon you can remove the mark for replication in the **Properties** window. The partial replication task on the domain connection is automatically deleted.

You can also delete such a partial replication task with the replication functions.

This also deactivates the selection in a navigator view. The deactivated option is displayed after the navigator is loaded, for example, after having started the Desktop Client.

✓ You change the properties of the partial replication task by changing the settings for the view.

6.3 Replicating certain data records

You can replicate certain data records with a partial replication task, see "Creating replication tasks" on page 57.

If you, for example, update certain data records of a partial replication task, two special functions are available in the Desktop Client:

- ✓ To replicate data records again, select them in a list view.
- ✓ Go to the **Tools** menu and click **Replicate data records again**. With this function, these data records are completely replicated if they are included in a replication task.
- ✓ To newly replicate a single data record, you go to the **Edit** menu in a data record and click **Replicate data record again**.

6.4 Replicating documents

Special factors have to be considered:

- Document data records can contain archive files.
- Archive files are not saved in the database but in the document archive.
- Templates are available for document data records and archive files.

Attachments of document data records including file references are not replicated.

This results in special cases that are relevant for document replication.

6.4.1 Replicating the Document Archive

If archive files are replicated from the Document Archive, these Microsoft Word documents, Microsoft Excel files or PDF files, and others, can be opened and edited. If the archive files are not found, only the document data record is replicated. If users then try to open the archive file on the corresponding domain by clicking **Open archive file** in

the data record window of the document, the **The file could not be found in the archive** message is displayed.

When replicating the Document Archive, the archive files are handled like fields of the document, that is, each change is logged and transferred to the corresponding domain when replicating the next time. As CAS genesisWorld does not analyze the content of the archive files, the archive file is always transferred completely.

If the Document Archive is also replicated, the data volume that is transferred can become very large, depending on the operating scenario of CAS genesisWorld and the number, type and size of the managed archive files; the replication will then take longer.

The Document Archive can be replicated optionally and this option can be activated when setting up a domain connection, see "Settings tab" on page 46.

Notes

You can select the setting for replicating the Document Archive in both domains. However, we do not recommend to choose a different setting for the two domains, as the archive files will then be replicated, but the content cannot be kept synchronous.

If the administrator has finished the initial setup of a domain, see "Special requirements for laptops" on page 16 and if access to the archive files of the respective corresponding domains is to be available, then the folder of the Document Archive must also be copied.

6.4.2 Replicating archive files

In a document data record, both the document data record itself and the archive file may have been changed since the last time you have replicated. A document data record may include several versions of archive files. You can see the versions in the **Edit document** window, in the **Versions** tab.

Changes in an archive file are replicated when creating a new version of the archive file.

If only the archive file was changed since the last replication, a check sum is first calculated. If an archived file was changed, a new version of the archive file is created with the replication.

If versions of the archive file already exist, all versions of the archive file on both domains are merged by date and newly numbered during the replication.

Replicating specific documents

Archive files are usually replicated with the **Archive documents** replication task. If such a replication task does not exist, it will be created with the following functions.

- ✓ Go to a list view and select one or several document data records with archive files.

- ✓ Go to the **Tools** menu and click **Add to archive replication task**. An **Archive documents** replication task is automatically created and the archive files will be added to this task.
- ✓ The name of the replication task can be changed via the replication management.
- ✓ You add document data records including archive files to the replication task using a static filter. You can also remove them in the **Filter** tab.

6.5 Resolving conflicts

Data is synchronized when replicating on the field level. If, for example, the subject of an appointment is changed in a domain, the connected domains will only be informed about this change, but they will not receive the complete appointment with all further fields and related data records.

Usually, conflicts do not occur for entire data records when changes are made simultaneously in two domains. If, for example, the phone number of an address is changed in one domain, and the salutation of the address in another domain, then the two changes are merged and as a result, a complete address with up-to-date phone number and salutation is then available.

If the same field was changed in both domains, then the time stamp is used to compare the times at which both changes were made. The last change made is applied.

If a data record has been changed and the same data record has been deleted on the other domain, the deleting process is treated with higher priority. The deleted data record is stored in the recycle bin and can be restored. The recycle bin is not replicated and therefore this data record has not been replicated either.

The archive file in both domains can be changed for document data records. When making changes to the archive file in both domains, a new version is created. When replicating, all versions of the archive file are merged with existing document versions and new numbers are assigned referring to the date. Thus, conflicts only occur with the meta data of a document record.

If a version of the document data record is created each time an archive file is saved, the archive file itself is also saved in different versions. This means that the document data records including their related archive files of each version are replicated if the archive file of the corresponding document is replicated.

All the changes made to a data record, including on other domains, are logged in a list on the **Change log** tab of a data record. This enables you to track changes and undo them if necessary.

7 Important information for administrators

7.1 Software updates

Software updates on computers that replicate are executed in mostly the same way as in CAS genesisWorld.

You cannot update the application server of CAS genesisWorld via MSI. Thus, we recommend to use the update service for software updates on computers that replicate.

You have to enter a Windows user on the **Registering the application server** page in the Installation Wizard or the Server Manager.

This user must also have access to the folder in the network on which the CAB files for CAS genesisWorld software updates are saved.

Laptops are updated the next time they are connected to the network.

For further information, please refer to the **CAS genesisWorld Installation and Setup Manual** at [CAS Product Documentation](#).

7.2 Code page and Unicode

If you use the replication function and activate Unicode, the setting for Unicode must be activated in all domains.

In the Database Wizard, the administrator defines the settings for the code page and Unicode.

7.3 Application server performance

The replication's performance depends on several factors that influence each other.

The system environment, the computers that are deployed and your network are factors that have to be taken into account.

Some factors that influence performance are described in the following and you can adapt your settings to improve performance.

- The number of concurrent replication processes
 - You can execute replications with domains at different times.
- Partial replication tasks with complex filter settings
 - You can create a number of partial replication tasks.
- In the connection properties, in the **Administration** tab, the **All (public)** user is set.

This setting is not used very frequently. You can restrict the data records you want to replicate by creating a respective replication task.

- Linked data records are replicated

You can create a number of partial replication tasks.

- Other services are running on the Application Server

Use a separate Application Server for the replication.

- Large replication volume

You can replicate several times on these connections.

You can specify the filter criteria for the replication tasks more precisely. Also include the archive files of document data records and e-mail attachments.

- Many partial replication tasks with almost the volume of a full replication task

You can create a full replication instead. It will be executed faster than a partial replication of the same size. Also create corresponding tasks for archive documents and archive documents templates.

7.4 Solving problems when replicating for the first time

You can track the replication process in the connection log that can be found in the **Connection status** window.

If the corresponding domain does not respond, check the following:

- Check first the log files of the corresponding domain in the **Connection status** window and in the local domain log.
- Has the application server of the other domain been started?
- If the application server runs, is the receiving agent activated?
- We recommend to use the **Test connection** function: go to the **Replication management** window and open the **Edit** menu. In this case, the system tries to replicate completely without exchanging user data.

7.5 Customizing the database

You or your CAS genesisWorld partner can customize and extend the database or database model according to your requirements. For example, you can create new fields or edit existing ones and create new data record types.

If you use the Database Designer and create new data record types, the replication must be activated with the corresponding option: **Allow replication** means that data records of this type can be replicated.

Deleting from the whole field of a data record type is a special case: This change to a domain through replication is not transferred to one or more other domains. The so-called automatic deleting in another domain can lead to unwanted data loss, because data deleted in this way cannot be restored. This is why you should delete fields separately on all domains.

There are two ways of transferring customizations from one database to another:

- Whenever a new field or data record type is created, a script file is generated for it. This file can then be imported into another database using the Database Wizard.
- In the **Database** module in the Management Console, administrators can export data record types as XML files, which can then be imported into another database.

Customizations are not replicated to a new, empty database.

You do not have to export and import customizations if you set up a new domain with a copied database, see "New domain: Database and Document Archive" on page 26.

All changes made to a database after the first replication are replicated the second time. This is defined in the connection's properties, see "Administration tab" on page 48.

7.6 Service jobs in the Database area

Service jobs which are created to clear data from a database always relate to a specific database. When replicating, several databases are connected to each other and all service jobs of all domains are displayed in the Management Console.

You can create service jobs for each domain. Service jobs are included when replicating data.

Service jobs are only displayed with read rights on a laptop and cannot be changed.

Service jobs are executed on the domain that is configured for a task.

If a schedule for a service job exists, a service job is executed at this or the next possible time. To execute the service job, a laptop must not be connected to the main installation but only be turned on.

Please note that executed service jobs also affect the main installation. For example, data records which were deleted in the main installation through a service job are also deleted in other databases the next time a replication is executed. A special service job on the laptop is not required. Service jobs for databases on laptops may only be useful in special circumstances.

If you replicate data, only system service jobs on the own domain can be reactivated. The service jobs of all domains can be activated in the central domain.

7.7 Notification and action service

If a problem occurs on the replication connection and needs to be solved, the administrator has the option to be automatically notified via e-mail.

The e-mail address must be entered in the properties of the local domain, in the **General** tab, see "General tab" on page 31. Furthermore, an e-mail account must be specified and set up in the **Rules** area, **Server settings** tab.

Rules are not replicated: One rule is always created and implemented for each database. Thus an action cannot be performed on two databases at the same time. Rules are included in the new database if you create a new domain with the copy of a database, see "Setting up the domain" on page 24.

In the new domain and/or the new database, these rules are not executed, as server settings for the notification and action service in the Management Console are not replicated. This service is furthermore not started and set for the new domain. The notification and action service must be set up separately for a main installation.

7.8 Checking mandatory fields

Mandatory fields are not checked when replicating. Identical mandatory fields should exist in the respective databases.

8 Important to know for certain functions

When working with specific functions, there are some key points you need to be aware of.

Please note, that for these functions only an administrator can check or define the corresponding settings on a laptop. Contact your administrator for more information if required.

E-mail accounts

- ✓ Open the **Tools** menu of the Desktop Client, select **Settings** and then open the **E-mail** tab.
- ✓ The information about all existing e-mail accounts is replicated with the **Replicate e-mail accounts** option in the **E-mail accounts** tab.

To replicate e-mail accounts, you must either create a full replication task or a navigator replication task, see "Creating replication tasks" on page 57.

If this option is deactivated and you change the parameters for the e-mail account of a domain, you must change the settings also in the second domain. Make sure that

parameters for the account are identical after the changes have been applied. Otherwise, errors may occur, for example, if you create e-mails in another domain and send these e-mails in the other domain.

The out-of-office service cannot be started on a laptop without a laptop license.

Only archived e-mails are full CAS genesisWorld data records and only those are replicated.

E-mail campaigns

If you create an e-mail campaign in a main installation, the name of the domain is implicitly stored in this e-mail campaign. The e-mail campaign is sent from an application server to the specified e-mail account. Only the application servers which belong to the same domain are used for sending.

If the e-mail campaign is transferred to a laptop during replication, then the e-mail campaign will be sent via the laptop domain instead of the application server. This also applies when the e-mail campaign is sent with a time delay. If you edit and then save the e-mail campaign on your laptop, you will receive a message. Now you must decide whether the e-mail campaign should be sent via the laptop domain, or via the central domain.

When you create and save a new e-mail campaign on a laptop for replication, you will be prompted whether the e-mail campaign is to be sent with the laptop or the corresponding domain.

If you want to send an e-mail campaign on a laptop, you should ensure that the laptop is functioning at the time of sending.

Distribution list

Distribution lists can be replicated using a user-defined order. The addresses of the distribution list are also automatically replicated. When replicating a distribution list, the "distribution list-address" matches are merged in two domains. A duplicate check will not be made. If in both domains identical addresses with different GGUIDs exist, unwanted side effects may occur: For example, different duplicates can occur in two domains.

Duplicate check

If you replicate data on a laptop, you must only execute a duplicate check on the main installation, not on laptops. If the duplicate check is performed on several laptops simultaneously, different addresses may be deleted. If you then perform a replication with the main installation, both addresses will be deleted on both sides.

SmartSearch

If data is replicated on a laptop and you want to work with the SmartSearch, the corresponding settings must also be defined on the laptop, in the **Management Console > Database** area. Furthermore, the corresponding service for the application server must be started in the Server Manager on the laptop.

Invitation management

If you work with the replication function, the settings for the invitation management must be the same for all replicating domains. Otherwise the replication cannot be started. If you create domains on laptops as administrator, the domain setting of the main installation is applied. The setting for the invitation management cannot be changed on a laptop. For other domains, such as subsidiaries, the setting for the invitation management must be entered manually in each domain.

If you use the replication function, then e-mails may only be sent once the replication has been finished. This is always the case when changes are made through a domain which does not correspond to the domain of the organizing participant. E-mails are then only sent once the replication to the organizing participant's domain has been finished.