

Application Operations Guide

SAP Landscape Transformation Replication Server
Document Version: 3.7 – 2018-05-25

CUSTOMER

SAP Landscape Transformation Replication Server

For SAP HANA Platform 2.0 SPS02



Typographic Conventions

Type Style	Description
<i>Example</i>	Words or characters quoted from the screen. These include field names, screen titles, pushbuttons labels, menu names, menu paths, and menu options. Textual cross-references to other documents.
Example	Emphasized words or expressions.
EXAMPLE	Technical names of system objects. These include report names, program names, transaction codes, table names, and key concepts of a programming language when they are surrounded by body text, for example, SELECT and INCLUDE.
Example	Output on the screen. This includes file and directory names and their paths, messages, names of variables and parameters, source text, and names of installation, upgrade and database tools.
Example	Exact user entry. These are words or characters that you enter in the system exactly as they appear in the documentation.
<Example>	Variable user entry. Angle brackets indicate that you replace these words and characters with appropriate entries to make entries in the system.
EXAMPLE	Keys on the keyboard, for example, F2 or ENTER.

Document History



Caution

Make sure you use the **current** version of the Application Operations Guide.

The current version of the Application Operations Guide is at service.sap.com/instguides on SAP Service Marketplace.

Version	Date	Change
3.4	2017-06-14	Revised published version for SAP HANA Platform 2.0 SPS01 and DMIS 2011 SP13.
3.5	2018-01-16	Revised published version for SAP HANA Platform 2.0 SPS02 and DMIS 2011 SP14.
3.6	2018-03-14	Revised published version for SAP HANA Platform 2.0 SPS02 and DMIS 2011 SP14. Update details a limitation regarding the prevention of archive deletes from being replicated to a target system.
3.7	2018-05-25	Revised published version for SAP HANA Platform 2.0 SPS02 and DMIS 2011 SP14. Update details additional information about using multiple 1:1 configurations for a source system.

Table of Contents

1	Getting Started.....	6
1.1	About this Guide	6
1.2	Global Definitions	6
1.3	Important SAP Notes	7
1.4	SAP Landscape Transformation Replication Server Guides	7
1.5	Key Terms.....	8
2	Basic Concepts and System Landscape Options.....	11
2.1	Basic Concepts of Trigger-based data Replication using SAP LT Replication Server for SAP HANA	11
2.2	System Landscape and Installation Options.....	12
3	Configuration Information and Replication Concepts	14
3.1	Technical Prerequisites and Authorization Aspects	14
3.1.1	Software Installation.....	14
3.1.2	System Connections and Authorizations.....	15
3.1.3	Creating a Configuration.....	16
3.1.4	Additional Information for Non- ABAP Source Systems.....	19
3.2	Managing the Replication Process Using the SAP HANA Studio	20
3.3	Accessing the Configuration and Monitoring Dashboard.....	21
3.4	Important Transactions and Control Tables.....	22
3.4.1	Mass Transfer ID and Schema GUID	22
3.4.2	Important Replication-Relevant Tables.....	23
3.4.3	Table Structure in SAP HANA Modeler	23
3.4.4	Data Transformation Capabilities within SAP Landscape Transformation Replication Server.....	24
3.5	Sizing and Basic Job Handling.....	24
3.5.1	Sizing of SAP LT Replication Server - Basic Rules and Influencing Factors	24
3.5.2	Important SAP LT Replication Server Jobs.....	26
3.5.3	Stopping Jobs	29
3.5.4	Restarting Jobs	29
3.6	Changing Load and Replication Procedures.....	30
3.6.1	Changing the number of Jobs.....	30
3.6.2	Embedded Health Checks	30
3.6.3	Performance Improvements for the Initial Load of Tables	31
3.7	Monitoring of Load and Replication Process	32
3.7.1	Status of Jobs and Connections.....	33
3.7.2	Status of Triggers	34
3.7.3	Statistics	34
3.7.4	Using the Monitoring Transaction LTRO	36
3.7.5	Monitoring using SAP Solution Manager.....	37
3.8	Backup and Recovery Aspects.....	38

4	SAP Landscape Transformation Replication Server Cockpit.....	39
4.1	Introduction	39
4.2	Starting and Stopping all Configurations	39
4.3	Prerequisites: Choosing a Mass Transfer ID	39
4.4	Navigating from the SAP Landscape Transformation Replication Server Cockpit	40
4.5	SAP Landscape Transformation Replication Server Cockpit Tab Pages	40
5	Special Considerations for Source Systems	53
5.1	Impact of Software Maintenance activities in ABAP Source Systems	53
5.2	Archiving Data in Source Systems	56
5.3	Preventing Archive Deletes from being Replicated to a Target System.....	57
5.4	Performance Implications in Source System using Trigger-based Data Replication	59
5.5	Improving Performance for Configurations with Many Tables and Low Change Rates.....	59
5.6	Data Volume Management.....	61
6	Replication Logging	62
6.1	Overview	62
6.2	Prerequisites.....	62
6.3	Changing the Replication Logging Settings	62
6.3.1	Operational Considerations.....	63
7	Frequently Asked Questions and Troubleshooting Recommendations.....	65
7.1	If SAP LT Replication Server for SAP HANA is already in use, what are the considerations when applying a new DMIS version and/or support package?	65
7.2	Does SAP LT Replication Server for SAP HANA only support 1:1 data replication or is it possible to do filtering (selective data replication) and / or transform data during the replication?	65
7.3	Are there any special considerations if the source system is a non-ABAP system?.....	66
7.4	Is it possible to use a configuration for multiple source and target SAP HANA systems?.....	66
7.5	What are the potential issues if the creation of a configuration and a related schema fails?	66
7.6	How can I ensure that data is consistent in the source system and SAP HANA system?	67
7.7	What happens in case of network failures?.....	67
7.8	Is the partitioning of SAP HANA tables supported by SAP LT Replication Server?	67
7.9	Does SAP LT Replication Server for SAP HANA support data compression like the SAP HANA database?	68
7.10	How to use Authorization Groups to control access on configuration level by LT Replication Server administration tasks?	68
7.11	How to define an alternative connection for the initial load?	69
7.12	How to change the settings of the DB connection to the target system?	70
7.13	The Notification Console	70
8	Appendix.....	72
8.1	SAP Landscape Transformation Replication Server on SAP Community Network (SCN).....	72
8.2	Related Guides.....	72
8.3	Related SAP Notes	72

1 Getting Started



Caution

This guide does not replace the daily operations handbook that we recommend customers to create for their specific production operations.

1.1 About this Guide

Designing, implementing, and running your SAP applications at peak performance 24 hours a day has never been more vital for your business success than now.

This guide provides a starting point for managing your SAP applications and maintaining and running them optimally. It contains specific information for various tasks and lists the tools that you can use to implement them. This guide also provides references to the documentation required for these tasks, so you will sometimes also need other Guides such as the Master Guide, Technical Infrastructure Guide, and SAP Library.

Target Groups

- Technical Consultants
- System Administrators
- Solution Consultants
- Business Process Owner
- Support Specialist

1.2 Global Definitions

SAP Application

A SAP application is an SAP software solution that serves a specific business area like ERP, CRM, PLM, SRM, and SCM.

Business Scenario

From a microeconomic perspective, a business scenario is a cycle, which consists of several different interconnected logical processes in time. Typically, a business scenario includes several company departments and involves with other business partners. From a technical point of view, a business scenario needs at least one SAP application (SAP ERP, SAP SCM, or others) for each cycle and possibly other third-party systems. A business scenario is a unit which can be implemented separately and reflects the customer's prospective course of business.

Component

A component is the smallest individual unit considered within the Solution Development Lifecycle; components are separately produced, delivered, installed and maintained.

1.3 Important SAP Notes



Caution

Check regularly for updates available for the Application Operations Guide.

Important SAP Notes

SAP Note Number	Title	Comment
1605140	Central Note - SAP LT Replication Server	Collective Note for all the relevant Notes for LT Replication Server for HANA
2570052	Installation/Upgrade SLT – DMIS 2011 SP14	This SAP Note describes the installation or upgrade of SAP Landscape Transformation Replication Server to the relevant DMIS SP.
1691975	HANA LTR – Clarification on DMIS releases	Decision criteria for choosing the correct software version
1768805	SAP LT Replication Server: Collective Note - Non-SAP Sources	Collective Note that describes details about using the SAP LT Replication Server for non-ABAP source systems
1963522	Limitation for SAP LT Replication Server on MaxDB	Due to issues with MaxDB database triggers, using SAP LT Replication Server to replicate data from source systems based on some versions of Max DB could lead to data inconsistencies in the target system. Therefore only MaxDB versions specified in this SAP Note can be used for SAP LT Replication Server scenarios.

1.4 SAP Landscape Transformation Replication Server Guides

The following table contains useful links to related guides:

Guide	Location
Security Guide - Replicating Data to SAP HANA	http://help.sap.com/sapslt

Guide	Location
Installation Guide – Replicating Data to SAP HANA	http://help.sap.com/sapslt
Sizing Guide	https://service.sap.com/sizing → Sizing Guidelines → Database and Technology → SAP In-Memory Computing → SAP Landscape Transformation Replication Server, SAP HANA

1.5 Key Terms

The following table contains key terms regarding the SAP Landscape Transformation Replication Server for SAP HANA:

Term	Definition
Advanced Replication Settings	A program (transaction code LTRS) that runs on the SAP LT Replication Server that you can use to specify advanced replication settings. For example, you can: <ul style="list-style-type: none"> • Modify target table structures • Specify performance optimization settings • Define transformation rules
Configuration	The definition of the parameters that the SAP LT Replication Server uses to replicate data from one or more source systems to one or more consumers, i.e. target systems and/or DB schemas. The configuration specifies the source system, the target system, and the relevant connections.
Configuration and Monitoring Dashboard	An application that runs on the SAP LT Replication Server that you use to specify configuration information (such as the source and target systems, and relevant connections) so that data can be replicated. You can also use it to monitor the replication status.
Database trigger	A database trigger is procedural code that is automatically executed in response to certain events to a particular database table or view.
Data transfer Job	A job that is used for the data transfer process in the SAP LT Replication Server.

Term	Definition
Initial load	A step within the trigger-based replication process that loads data from the source system to target system.
Initial load Job	A job that is used for the initial load process in the SAP LT Replication Server.
Latency	Latency is defined as the length of time it takes until a data change to a source table has been committed to the target system database.
Logging table	A table in the source system that records any changes to a table that is being replicated. This ensures that the SAP LT Replication Server can replicate these changes to the target system.
Master job	One job that exists for all active configurations that starts the initial load jobs and the data transfer jobs in the SAP LT Replication Server.
Maximum latency	The maximum latency which occurred within a selected time frame.
Median latency	The median latency which occurred within a selected time frame.
Minimum latency	The minimum latency which occurred within a selected time frame.
Reading type	A technique for reading data from tables in the source system during the initial load process.
Replication phase	A phase in the trigger-based replication process whereby only changes to the source database (recorded by databases triggers) are replicated to the target database, thereby facilitating real-time data replication.
SAP LT Replication Server	An SAP system that facilitates the replication of data from one or more source systems to one or more target systems. The source systems can be SAP or non-SAP systems.
Tables for replication settings	The set of tables in the SAP LT Replication Server that contain the replication settings for a specific configuration.
Trigger-based replication	A technique for replicating data where an initial load is first performed that loads data from the source to the target system, and a replication phase begins whereby only changes to the source database (recorded by databases triggers) are replicated to the target database, thereby facilitating data replication in real-

Term	Definition
	time.
Transformation rules	<p>A rule that you can specify in the Advanced Replication Settings transaction for source tables such that data is transformed during the replication process. The transformed data will then be present in the target table. For example, you can specify a rule to:</p> <ul style="list-style-type: none"> • Convert fields • Fill empty fields • Skip records

2 Basic Concepts and System Landscape Options

2.1 Basic Concepts of Trigger-based data Replication using SAP LT Replication Server for SAP HANA

The SAP Landscape Transformation (SAP LT) Replication Server is the SAP technology that allows you to load and replicate data in real-time from ABAP source systems and non- ABAP source systems to an SAP HANA environment.

The SAP LT Replication Server uses a trigger-based replication approach to pass data from the source system to the target system.

The SAP LT Replication Server can be installed either as a separate SAP system, or if the technical prerequisites permit, on an ABAP source system.

In order to replicate data, you must first define the parameters that the SAP LT Replication Server will use to replicate data from one or more source systems to one or more target systems. You specify this information in a Configuration. A configuration contains information about the source system, the target system, and the relevant connections. You define configurations in the SAP LT Replication Server.

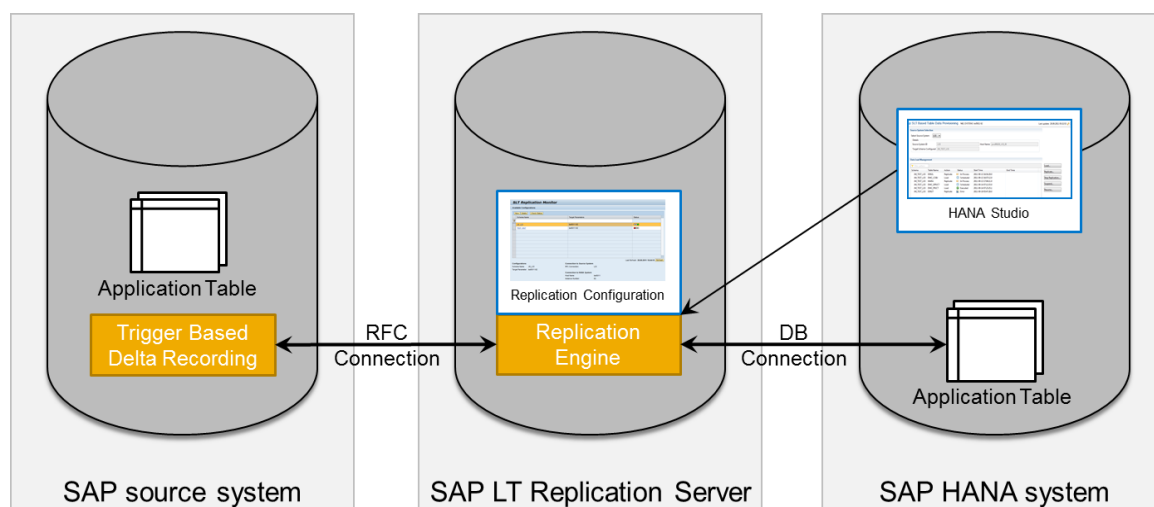
You use a configuration to load and replicate data from one source system to one target database schema of a HANA system (1:1), or from multiple source systems to one target database schema of an SAP HANA system (N:1). Furthermore, it is possible to load and replicate data from one source system to multiple (up to 4) target database schemas of one or more HANA systems (1:N). You can also specify the type of data load and replication - either in real-time, or scheduled by time or by interval.

Note

By default, a maximum of one 1:1 configuration is possible for a source system. If more 1:1 configurations are required for a source system, set the value of the field MULTI_1_1 to X in table IUUC_RT_PARAMS in the SAP LT Replication Server system.

The trigger-based data replication approach is a table-based concept that retrieves data from application tables in the source system (or source systems). A user uses the SAP HANA Studio to select the tables for the replication. If not all data records of a table should be transferred, you can create transformation rules to selectively filter the data (selective data replication) or to enable other transformations during the data replication process.

The following graphic outlines the basic concept and the typical landscape (for an ABAP source system) using the trigger-based data replication approach of the SAP LT Replication Server:



2.2 System Landscape and Installation Options

The SAP LT Replication Server can be used for data replication from ABAP source systems and non- ABAP source systems to the SAP HANA system. In the Configuration & Monitoring Dashboard (transaction LTR) of the SAP LT Replication Server, you can define a new configuration that contains the relevant information required to create the connection between the source systems and the target SAP HANA systems.

The technical landscape consists of following components:

- ABAP source system (or systems)

The source system tracks database changes by using database triggers. It records information about changes in the logging tables. Read modules (located on the ABAP source system) transfer the data from the source system to the SAP LT Replication Server. The relevant data is read from the application tables.

- Non-ABAP source system (or systems)

The non- ABAP source system tracks database changes by using database triggers. It records information about changes in the logging tables. Read modules (located at the SAP LT Replication Server) transfer the data from the non-ABAP source system to the SAP LT Replication Server. The relevant data is read from the application tables.

SAP LT Replication Server

An SAP system that facilitates the replication of data from one or more source systems to one or more target systems. The source systems can be SAP or non- ABAP systems.

- SAP HANA system

The SAP HANA system contains the SAP HANA database. It is used to store the replicated data. The SAP LT Replication Server and the SAP HANA system communicate by means of a database connection.

Prior to the installation, it is important to understand the various system landscape options available:

- For an ABAP source system, the SAP LT Replication Server can either be installed as part of the source system (depending on the customer's system landscape strategy, and only if certain technical prerequisites are fulfilled) or as a separate SAP system (recommended for productive use). However, from a technical perspective, the SAP LT Replication Server does not have to be a separate SAP system if the source system complies with the technical prerequisites.
- For non- ABAP source systems, the SAP LT Replication Server has to be installed as a separate SAP system.

The following table outlines in more detail the advantages and disadvantages of the different installation options:

	Source System (if SAP system)	SAP Solution Manager	Dedicated System
Advantages	Simplified landscape and administration	Re-use of existing NW instance	No software maintenance dependencies Flexibility
Disadvantages	Performance impact Potential software maintenance dependencies	Performance impact Potential software maintenance dependencies	Investment and maintenance effort for separate server / NW instance

Experiences show that HANA customers using the SAP LT Replication Server tend to use a dedicated SAP LT Replication Server for productive use. An SLT sandbox or quality assurance system is installed sometimes on top of an appropriate ABAP source system or on top of an SAP Solution Manager system.

All details about the preparation, the installation, and the configuration procedure of SAP LT Replication Server for SAP HANA are described in the Installation Guide. You can find the latest version of these guides on the SAP Help Portal at <http://help.sap.com/sapslt>. The sizing guide can be found in the SAP Service Marketplace under <http://service.sap.com/sizing> → Sizing Guidelines → Database & Technology → SAP In-Memory Computing.

For more information about sizing-related information, see chapter 3.5.

3 Configuration Information and Replication Concepts

This section provides an overview of the transactions and tools you use to define a configuration and related table settings. It also includes basics about the SLT-based replication concept which is key to understanding important SAP Landscape Transformation Replication Server-related operational aspects.

3.1 Technical Prerequisites and Authorization Aspects

Before you can create a configuration and start the data replication, you have to install the required software and establish related system connections with appropriate user authorizations.

For more information, see the relevant installation and security guides on the SAP Help Portal at <http://help.sap.com/sapslt>, as well as in the specific SAP Notes listed below.

3.1.1 Software Installation

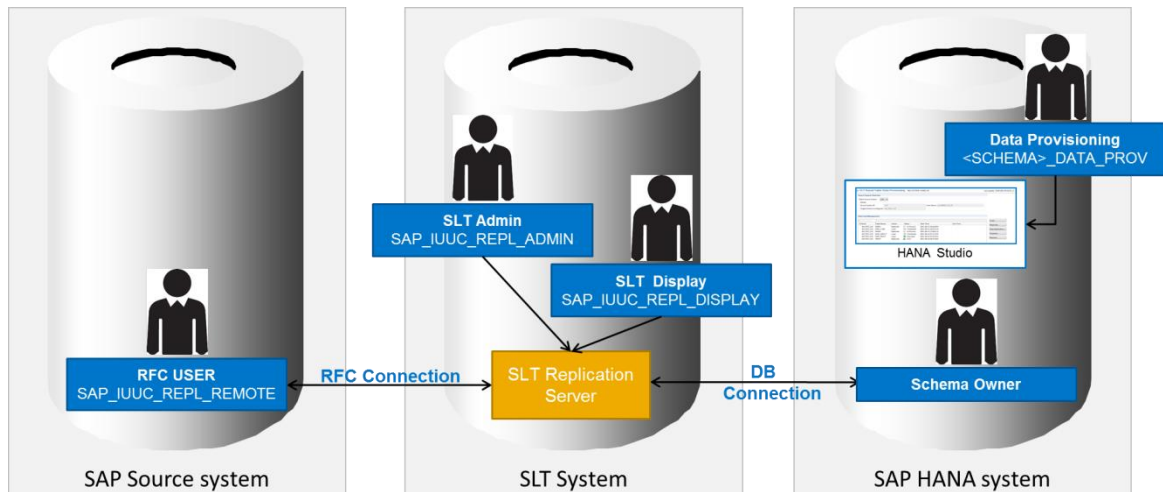
SAP LT Replication Server is shipped in a specific add-on (DMIS_2010* or DMIS_2011*). Depending on the installation option, you have to ensure that the respective DMIS add-on is installed in the ABAP source system(s) and in the SAP LT Replication Server. The installation of the DMIS add-on is not required (or possible) on non-ABAP source systems.

Note

- Ensure that the correct SAP kernel version is in use in the SAP LT Replication Server with the appropriate patch level - otherwise the connection to the HANA system will fail. For more information, see SAP Note [1597627](#)
- In the ABAP source system(s) and the SAP LT Replication Server, check if the minimum support package level for the DMIS add-on is applied to ensure compliance with the respective SAP HANA Studio version. For more information, see SAP Note [1759156](#)
- Ensure that the appropriate front-end requirements are met (SAP GUI version).
- Verify that the SLT-specific web services are activated before you launch the Configuration & Monitoring Dashboard (transaction LTR). For more information, see SAP Note [517484](#)

3.1.2 System Connections and Authorizations

Before you create a configuration in the SAP LT Replication Server, you need to define the related users with relevant authorizations in the source system(s), and establish the appropriate system connections. The following graphic illustrates the authorization concept:



For ABAP source systems, a predefined user role (SAP_IUUC_REPL_REMOTE) is available, and an RFC connection between the ABAP source system and the SAP LT Replication Server has to be established.

Since DMIS 2010 SP9 / DMIS 2011 SP4 two new roles for the user in the SLT System have been created. Role SAP_IUUC_REPL_ADMIN is required to create, change or display configurations. Role SAP_IUUC_REPL_DISPLAY can be used if a user should be able to see the existing configuration and their status but is not allowed to change any settings or create a new configuration.

For non- ABAP sources, you need to create database user(s) with appropriate authorizations in advance and establish the database connection by using transaction DBCO in the SAP LT Replication Server.

The connection to the SAP HANA system is established when you create a configuration in the SAP LT Replication Server. You need to use an SAP HANA Studio user with appropriate authorizations (such as user SYSTEM) to establish a database connection and to successfully create a database schema (for the data replication) in the SAP HANA system.

Note

- For ABAP source systems:
 - Ensure that the user for the RFC connection has the role IUUC_REPL_REMOTE assigned and that the user role is generated correctly.
 - Do not use a DDIC user for RFC connection.
 - If the SAP Landscape Transformation Replication Server system is connected to an ABAP-based SAP source system by means of an RFC connection, then there will be an increase in the number of log entries (for example ST03 traces) in the SAP Landscape Transformation Replication Server system.
- For non- ABAP source systems:

- Ensure a database user has the sufficient authorizations for data replication. For more information, see the Security Guide.
- Ensure the database-specific library components for the used kernel are installed in the SAP LT Replication Server.
- Review other database specific requirements and potential limitations by searching for SAP Notes that contain the key word "non-SAP", under application area HAN-DP-LTR.
- For the SAP HANA system:
 - In order to create a configuration, the user SYSTEM is required, or a user with similar authorizations with at least the authorizations for creating a HANA DB schema is required.
 - While creating a configuration, a replication user (with name equal to the schema name) is generated automatically.
 - The following authorizations are granted to the schema-specific roles for the replication user for the created database schema:
 - <SCHEMA>_DATA_PROV;
 - <SCHEMA>_POWER_USER;
 - <SCHEMA>_USER_ADMIN,
 - <SCHEMA>_DATASELECT)
- For a consolidation scenario (N data sources replicating to the same SAP HANA schema) the SLT Server connects to an already existing schema and the user has to enter the password of the schema owner which is normally not known. If SLT connects to an existing schema, it checks whether a database connection already exists that points to this schema. If such a connection exists, the logon data will be copied and the user does not need to enter the password any longer (just a confirmation popup that the schema exists must be confirmed in order to proceed).
 If you need to change the password of the database connection to the SAP HANA system, you need change the password of the HANA schema user in the SAP HANA studio and then specify the database connection password in table DBCON by using transaction SM30. Note that providing a new initial password to the SAP HANA schema user requires a dedicated dialog user log-on, and that the initial password is changed to a permanent one.

You can find more information about the roles and authorization concept of the SAP LT Replication Server in the Security Guide.

3.1.3 Creating a Configuration

In order to replicate data, you must first define the parameters that SAP Landscape Transformation Replication Server will use to replicate data from a source system to a target system. You specify this information in a configuration. A configuration contains information about the source system, the target system, and the relevant connections. You create a configuration in SAP Landscape Transformation Replication Server system.

To create a configuration, open the SAP Landscape Transformation Replication Server Cockpit (transaction **LTRC**) and choose *Create Configuration*. Alternatively, you can create a configuration by using the Configuration and Monitoring Dashboard (transaction **LTR**).

Creating a configuration involves the following steps:

1. Specify *General Data*

In this step, you specify the following information:

- Configuration Name
 - You must specify a name for the configuration. This name is used for the schema that is created automatically in the target system.
- Description (Optional)
 - You can specify a description for the configuration. .
- Authorization Group (Optional)
 - If you have sufficient authorizations, and require specific authorizations for the configuration, you can specify an authorization group here, and use this authorization group in the corresponding authorization object S_DMIS_SLT.

2. Specify Source System

For ABAP source systems, you specify the following information:

- System Data
 - You can use either an SAP system or a non-SAP system as a source system.
- RFC Destination
 - Specify the RFC destination to the source system.
- Allow Multiple Usage
 - Select this checkbox if you want to replicate data from an ABAP source system to multiple target systems.
- Read from Single Client
 - Select this checkbox if you only want to replicate data from the client that is specified in the RFC destination.

For non- ABAP source systems, you specify the relevant database and connection information.

3. Specify Target System

A) Connection to SAP HANA System.

To connect to the SAP HANA system, supply the relevant in-memory database connection information:

- **User Name:** Specify the username that will be used to create the schema in the HANA database. Usually this is the system user in the HANA system.
- **Host Name:** Specify the hostname of the HANA database. This field is limited to 32 characters. A FQDN (Fully Qualified Domain Name) is not necessary. Ensure that the hostname corresponds to the correct IP address.
- **Instance Number:** This is the instance number of the HANA database.
- **Tablespace Assignment:** If you chose to configure a separate log table space you can optionally supply this information.

It is possible (but not essential) to store the source system replication log tables in a separate table space. The decision to do this or not rests with the system administrator. One advantages of having the log tables in their own table space is that you can easily monitor the size of the log tables.

As each database system has its own method of providing this functionality, refer to your database documentation for this procedure.

If you use own data classes and tablespaces, see SAP Note [46272](#).

B) Connection to SAP System via RFC Connection

Enter the RFC connection you defined to the SAP target system. This connection option can be used for configuration of standard RFC Scenario for connecting a generic SAP NetWeaver ABAP-based System, the BW PSA replication scenario, or the ODP replication scenario.

4. Specify Transfer Settings

In this step, you specify the following information:

- **Initial Load Mode**
 - There are different options (reading types) available for the initial load. These reading types access the data in the source system table in different ways. In order to accelerate the initial load, you can change the data load behavior for the entire configuration. The default setting is *Resource Optimized* (reading type 3 for all tables), but you can accelerate the data load by choosing *Performance Optimized* (reading type 5 for transparent tables, and reading type 4 for cluster tables).
- **Dataclass of Tablespace**
 - You can specify a tablespace for the logging tables in the source system.
- **No. of Data Transfer Jobs**
 - You can specify the number of jobs that are used for the data transfer process in the SAP LT Replication Server.
- **No. of Initial Load Jobs**
 - You can specify the number of jobs that are used for the initial load in the SAP LT Replication Server.
- **No. of Calculation Jobs**
 - For reading types 1 and 3, this is the number of jobs that are used to calculate the data transfer portions that are used for the initial load. For reading types 4 and 5, it is the number of jobs that transfer the portions to table DMC_INDCL.

Replication Options

You can choose one of the following replication options:

- Real-time - The trigger-based data replication method that continuously replicates any data changes to the target system.
- Schedule by Interval - You can specify a time and frequency for the replication. For example, every 30 minutes or every 12 hours.
- Schedule by Time - You can specify a specific time for the replication, for example 23:00. The SAP LT Replication Server would then replicate any database changes to the target system every day at 23:00.

Activation of Replication Logging - If you choose this option, the replicated data entries will be saved in the SAP LT Replication Server system for all tables in the configuration. When a table is replicated to the target system, the replicated data will also be saved to database tables in the SAP LT Replication Server system for a configurable period of time. This means if any data is missing from the target system, you can view this data and replicate it again from the SAP LT Replication Server system to the target system (for instance in case of a point-in-time recovery of the target system).

Application (Optional)

Depending on your use case, you may be required to specify an application.

5. Review and Create

In this step, you can review your settings and create the configuration.

3.1.4 Additional Information for Non- ABAP Source Systems

If you are replicating from non- ABAP source systems, the user you specify need authorizations as described in SAP LT Replication Server – Security Guide.

The actual privilege to be granted to the database user depends on the database system (Oracle/DB2/MSSQL and so on). For example, if you want to configure an Oracle database as a non- ABAP source system, the following steps apply:

1. Install the Oracle instant client on the SAP LT Replication Server (if your SAP LT Replication Server is not based on Oracle).
2. Install the DBSL database dependent library for the correct kernel version (i.e. Kernel 7.20 for a SAP NetWeaver 7.0.2 SAP LT Replication Server system).
3. Create the database connection in table DBCON (by using transaction **SM30**)
4. Add the database connection in transaction DBACOCKPIT
5. Test the database connection
6. In the SAP LT Replication Server, use transaction **LTR** to complete the configuration for the non- ABAP source system.

Constraints

- Only tables with a primary key can be replicated.
- Details for tables DD02L and DD02T cannot be displayed within the Configuration and Monitoring Dashboard, because these tables do not exist physically in the non-ABAP source system.

- However, the metadata from the non-SAP system will be loaded into both SAP LT Replication Server and the target system as 'initial load' (this means that no automatic updates will occur for the metadata tables).
- The metadata information is represented as table DD02L and DD02T in the HANA system.

Note

For important considerations about non- ABAP source systems, see SAP Note [1768805](#).

3.2 Managing the Replication Process Using the SAP HANA Studio

Once a configuration is created in the SAP LT Replication Server, you use the SAP HANA Studio to initiate and control the table-based replication process of the SAP LT Replication Server.

Procedure

1. In the SAP In-Memory Studio, access the information modeler by selecting Data Provisioning.
2. In the SAP LT Replication Server Based Table Data Provisioning screen, you can control the SAP LT Replication Server.

Select Source System: This is the system ID of the source system retrieved by the RFC connection you configured in section 3.2.1.

Details: This information is automatically populated from the schema configuration.

3. You can use the Load, Replicate, Stop Replication, Suspend and Resume pushbuttons to control the replication for the selected source system.

Note

Before you can select application tables related to the selected source system and target schema, the initial load of the tables DD02L, DD02T and DD08L must be completed as they include important meta data information. For more information, see chapter 3.3.2.

Load: Starts an initial load of replication data from the source system. The procedure is a one-time event. After it is completed, further changes to the source system database will not be replicated.

Note

- For the initial load procedure, neither database triggers nor logging tables are created in the source system
- Default settings use reading type 3 (DB_SETGET) with up to three background jobs in parallel to load tables in parallel or subsequently into the HANA system. For very large tables, you may consider in advance to accelerate the initial load procedure as outlined in chapter 4.

Replicate: Combines an initial load procedure and the subsequent replication procedure (real-time or scheduled).

Note

Before the initial load procedure will start, database trigger and related logging table are created for each table in the source system (as well as additional replication-related objects in the SAP LT Replication Server).

Stop Replication: Stops any current load or replication process of a table.

Note

The stop function will remove the database trigger and related logging tables (and additional replication-related objects) completely. Only use this function if you do not want to continue a selected table otherwise you must initially load the table again to ensure data consistency.

Suspend: Pauses a table from a running replication. The database trigger will not be deleted from the source system. The recording of changes will continue, and related information is stored in the related logging tables in the source system.

Note

If you suspend tables for a long time, watch carefully the size of the logging table(s) and adjust the table space if required

Resume: Restarts the replication for a suspended table. The previous suspended replication will be resumed (no new initial load required).

3.3 Accessing the Configuration and Monitoring Dashboard

You use the Configuration and Monitoring Dashboard to create configurations, and to view status information for the replication.

You can access the Configuration and Monitoring Dashboard by using transaction LTR.

Note

In order to ensure compliance with security standards, the Configuration and Monitoring Dashboard (transaction LTR) requires the use of the SAP NetWeaver Business Client 3.5 or 4.0.

The reason for this is that the SAP NetWeaver Business Client supports a logout for all Web Dynpro windows. Web browsers do not support this logout feature. For example if you use a web browser to access the work center, there is no option to logout. Simply closing the web browser window does not log the user out of the system. The session runs on the server until it times out, and this is a potential security risk.

In order to use the SAP NetWeaver Business Client, the relevant user must have the role `SAP_IUUC_REPL_NWBC` assigned to them.

Alternatively, you can use the SAP LT Replication Server Cockpit (transaction `LTRC`) in order to work with configurations.

3.4 Important Transactions and Control Tables

The SAP LT Replication Server has two transactions for managing configuration-specific settings and table-specific settings:

- The Configuration & Monitoring Dashboard (transaction `LTR`) provides capabilities to initially create, manage, and monitor configuration and related settings. After the creation of a configuration, you can:
 - Use functions such as stop/restart the configuration master job.
 - Change configuration settings such as the number of total jobs (and number of total jobs for initial load).
 - Monitor the trigger and replication status, including detailed throughput information for the data replication.
- Advanced Replication Settings (transaction `LTRS`) allows you to define and change various table settings for a configuration such as:
 - Partitioning and structure changes for target tables in SAP HANA
 - Table-specific transformation and filter rules
 - Adjusting the number of jobs (and reading type) to accelerate the load/replication process

For a more detailed analysis, and for additional scenarios that are not yet supported by the Configuration & Monitoring Dashboard, SAP LT Replication Server provides an additional application, the SAP LT Replication Server Cockpit, which can be accessed by using transaction `LTRC`.

The following sections provide you with a summary of the most important control tables that are relevant for understanding in more detail the SLT-based replication process.

3.4.1 Mass Transfer ID and Schema GUID

When creating a new configuration in the SAP LT Replication Server, a schema GUID and a mass transfer ID (`MT_ID`) is automatically created and assigned to the configuration.

A schema GUID ensures that configurations with the same schema name can be created.

The mass transfer ID is used in the naming of SLT jobs. With the mass transfer ID, the system can uniquely identify a schema.

The mapping (and further details) of mass transfer IDs and related schema are stored in the SLT control table DMC_MT_HEADER in the SAP LT Replication Server.

3.4.2 Important Replication-Relevant Tables

Application tables in the source system can be selected for initial load or for replication.

In SAP systems, table definitions of the SAP dictionary are stored in table DD02L, related texts in DD02T. These dictionary tables as well as table DD08L (containing foreign key relationships of SAP application tables) are automatically replicated from the source system into the SAP HANA system when creating a schema. Once the replication is successfully completed, SAP HANA Studio users can select tables for initial load and/or for replication.

To ensure that new tables or structure changes of existing tables in the source system(s) are automatically reflected in SAP HANA, the dictionary tables DD02L, DD02T and DD08L always remain in the status *Replicate*, or *In Process*.

3.4.3 Table Structure in SAP HANA Modeler

While saving a configuration in the SAP LT Replication Server the system automatically generates related control tables in the HANA database for each schema.

The general definition and registration of all schemas are stored in the replication configuration table RS_REPLICATION_COMPONENTS which is included in schema SYS_REPL.

The following table provides an overview of the important control tables:

Schema Name	Contains Tables	Owner
SYS_REPL	RS_REPLICATION_COMPONENTS	SYSTEM
<SCHEMA>	RS_ORDER: includes basic information about table selected for initial load/replication RS_STATUS: includes information about replication status of each table RS_LOG_FILES <loaded_and_replicated_tables>: provides details and log information of replicated tables	<SCHEMA>

These control tables provide the flexibility to explicitly grant schema-specific authorizations to dedicated SAP HANA studio users. For more information, see the Security Guide – Trigger-Based Data Replication Using SAP LT Replication Server on the SAP Help Portal at http://help.sap.com/hana_platform.

3.4.4 Data Transformation Capabilities within SAP Landscape Transformation Replication Server

If you use SAP LT Replication Server (DIMS 2010 with SP07 or higher, or equivalent DMIS 2011 SP2 or higher) with SAP HANA 1.0 for trigger-based replication, you can use transaction LTRS to specify advanced replication settings. For example, you can:

- Modify target table structures
- Specify performance optimization settings
- Define transformation rules to transform data from the source table during the replication process. The transformed data will then be present in the target table. For example, you can specify rules to convert fields, fill empty fields, and skip records.

For more information about these advanced replication settings, see the documentation that is accessible from the UI of transaction LTRS.

3.5 Sizing and Basic Job Handling

3.5.1 Sizing of SAP LT Replication Server - Basic Rules and Influencing Factors

As a minimum requirement, an SAP LT Replication Server should provide the following hardware configuration:

- File system: 100 GB
RAM: 8-10 GB

- CPU: 2-4 cores
- Minimum number of background jobs: 10
- Network: 1GB/sec or better

In order to provide a rough order of magnitude sizing SAP offers a simple approach with three different SLT configurations:

	SMALL	MEDIUM	LARGE
Use Case	A small scenario with typically one LTR configuration with approx. up to 50 tables weighted table size category S-M an overall expected throughput of less than 1.000.000 records/hour	A moderate mid-range scenario with approx. ~ 3 different Source Systems (equivalent to 3 LTR Configurations), and/or up to 200 tables in total; weighted table size category M-L an overall expected throughput of less than 10.000.000 records/hour	A upper mid-range scenario with up to 10 different Source Systems (equivalent to 10 LTR Configurations), and/or up to 500 tables (in total); weighted table size category M-XL an overall expected throughput of up to 50.000.000 records/hour
SLT Server	1 LTR configuration with 2 Data Transfer Jobs Hardware: 2-4 CPU Cores, 8-10 GB Main Memory	10 Data Transfer Jobs in total (sum of all configurations) Hardware: 4-6 CPU Cores, 8-16 GB Main Memory	25 Data Transfer Jobs in total (sum of all configurations) Hardware: 8-10 CPU Cores, 16-32 GB Main Memory
Source System(s)	1:1 relation to data transfer jobs per source Reserve 2 BTC work processes for ACL (Access plan calculation), ensure 2 free Dialog work processes for data load/replication Additional Hardware required: ~ 1 CPU Core (0.5 CPU per data transfer job, APPL & DB)	1:1 relation to data transfer jobs per source sum over all source systems: Reserve 2-4 BTC work processes for ACL (Access plan calculation), ensure 10 free Dialog work processes for data load/replication Additional Hardware required: ~ 5 CPU Core in total (0.5 CPU per data transfer job, APPL & DB)	1:1 relation to data transfer jobs per source sum over all source systems: Reserve 4-8 BTC work processes for ACL (Access plan calculation), ensure in sum 25 free Dialog work processes for data load/replication Additional Hardware required: ~ 12 CPU Core in total (0.5 CPU per data

	SMALL	MEDIUM	LARGE
			transfer job, APPL & DB)
SAP HANA System	~ 1 additional CPU core	~ 3-4 additional CPU cores	~ 8 additional CPU cores

SAP provides general sizing information on the SAP Service Marketplace. You can find more information in the Sizing Guide for SAP Landscape Transformation (SLT) Replication Server for SAP HANA, SPS 05 at <http://service.sap.com/sizing> → Sizing Guidelines → Database and Technology → SAP In-Memory Computing.

However, for productive use and especially if you are dealing with large volume tables – the sizing may need to be adjusted. See also chapter 3.6, Changing Load and Replication Procedures.

In general, the load and replication procedure uses reading type 3 (DB_SETGET) with a maximum of three BGD jobs for each table for the load/replication procedure. Depending on the number of jobs available, the load/replication of tables is performed in parallel or sequentially.

The following criteria influence the requirement to change or increase the number of data transfer jobs:

- The number of configurations managed by the SAP LT Replication Server.
- The number of tables to be loaded/replicated for each configuration.
- The expected speed of the initial load (load time).
- The expected replication latency time (which mainly depends on the number and volume of changes of the tables in replication). As a simplified rule of thumb, one background job should be used for every 5 to 10 tables to achieve acceptable latency times ("real-time" data replication). However, the SAP LT Replication Server Sizing Guide provides more precise guidance about how to calculate the required number of data transfer jobs for a given scenario.

3.5.2 Important SAP LT Replication Server Jobs

3.5.2.1 Overview

To understand the concept of the load and replication procedures in more detail, the following section explains the major jobs involved in the replication process.

Master Job (Monitoring Job)

There is only one master job which handles all active configurations (a configuration is active if flag is set on IUUC_REPL_HDR).

An additional master job is started for system event SAP_SYSTEM_START which ensures that the master job is started automatically after a restart of the SLT system.

The decrease and increase of initial load and data transfer jobs is controlled by the composite master job.

Naming Convention: /1LT/IUC_REP_MSTR

Every 5 seconds, the monitoring job checks in the SAP HANA system whether there are new tasks and, if so, triggers the master control jobs. It also deletes the processed entries (tasks) from table RS_ORDER and writes statistics entries into table RS_STATUS (in the relevant schema in the SAP HANA system).

Note

Every day at midnight, the monitoring job and related data load jobs - but not the migration object definition or access plan calculation jobs – are automatically stopped and restarted immediately. This action has no negative impact on the ongoing data load and replication, it will simply resume immediately.

However, this procedure allows several activities on specific SAP LT Replication Server control tables that avoids potential long-term performance issues with the SAP LT Replication Server and simplifies a lot of log and job-related troubleshooting activities, for example by providing the option to filter by date in transaction SM37.

Master Controller Job

Naming Convention: /1LT/IUC_REP_CNTR_<mtid>

This job is scheduled on demand and is responsible for:

- Creating database triggers and logging table in the source system
- Creating synonyms
- Writing new entries in admin tables in SLT server when a new table is loaded/replicated

Data Load Job

Naming Convention: 1LT/IUC_LOAD_MT_<mtid>_nnn

This job should always be active. If the job does not complete successfully, the master controller job restarts it.

This job is responsible for:

- Loading data (load)
- Replicating data (replication)
- Changing status flag for entries in control tables in the SAP LT Replication Server

Migration Object Definition Job

Naming Convention: /1LT/IUC_DEF_COBJ_<mtid>

This job defines the migration object of a specific table (that you choose to load/replicate), which is the fundamental object for LT replication. The migration object definition should normally be quite fast for all tables.

Access Plan Calculation Job

Naming Convention: /1LT/IUC_CALC_ACP_<mtid>_n

This job calculates the access plan of a specific table (that you choose to load/replicate), and the access plan is used for data load or replication. The access plan is also a fundamental object for the replication. For a normal sized table, access plan calculation should finish quite quickly (less than 1 minute) while large tables might take up to several hours to finish.

3.5.2.2 Data Transfer Jobs

This section explains the relationship between the number of data transfer jobs and the number of available background work process.

Data transfer and data transformation processing on SLT server system is accomplished by the background work processes of the underlying SAP NetWeaver ABAP application server. Each job occupies 1 background work process in the SAP LT Replication Server system. For each configuration, the parameter [Data Transfer Jobs](#) restricts the maximum number of data load job for each mass transfer ID (MT_ID).

In total, a mass transfer ID (MT_ID) requires at least 4 background jobs to be available:

- One monitoring job (master job)
- One master controller job
- At least one data load job
- One additional job either for the migration objects definition, access plan calculation or to change configuration settings in the [Configuration & Monitoring Dashboard](#).

Example

If you set the parameter [Data Transfer Jobs](#) to 04 in a configuration "SCHEMA1", a mass transfer ID 001 is assigned. As a result, the following jobs should be in the system:

- 1 Master controller job: /1LT/IUC_REP_CNTR_001
- At most 4 parallel jobs for MT_ID 001: /1LT/IUC_LOAD_MT_001_001/~002/~003/~004

When configuring your data load or replication scenario, consider the following:

- Do not define more data transfer jobs than the number of available application server background work processes. If all available background work processes are already occupied by jobs, any other job will have to wait until a free work process becomes available. This can lead to long wait times until a new activity (for example creating triggers) can start, and can also result in significantly increased latency times for data replication.
- The number of dialog work processes in the source system corresponds 1:1 with the number of data transfer jobs in the SAP LT Replication Server system.
- Besides the work processes allocated by the data transfer jobs you need to provide additional available work processes for controller and monitoring jobs, the migration objects definition, access plan calculation or to perform configuration changes, and so on.

Sizing for SAP LT Replication Server involves determining how many work processes are required to perform the initial load of data into the target system within an acceptable timeframe, and accomplish the change capturing and the transfer of data changes to the target system within expected latency times.

Ensure that you add enough additional work processes to allow other required SAP LT Replication Server jobs to run.

Finally, you map the number of required application server work processes to their system resource consumption (CPU, memory, disc space) using the formulas provided by the SLT Sizing Guide.

With the simple formula below, you can calculate the number of required application server work processes (WPs) on the SLT Server for each active SLT configuration.

The number of required work processes can be determined by adding

- The Number of required data transfer jobs ,
- plus one background work process for Central Master (Monitoring) Job (only one per system!),
- plus one background work process for Master Controller Job,
- plus 3-5 additional empty background work processes (recommended per configuration),
- plus approx. 3 dialog work processes (recommended for each configuration).

Note: A lack of available free application server work processes can negatively affect the data load or data replication processes.

3.5.3 Stopping Jobs

In the Configuration & Monitoring Dashboard (tab page [Job and Connections](#)), you can suspend the load and / or replication for all tables of a configuration using the pushbutton Stop for the master job. The master job stops as well all related jobs of a configuration – initial load and / or replication will immediately discontinue – however, a database trigger in the source systems will continuously record changes in the log tables.

Note

As an alternative to automatically temporarily stopping and restarting the replication after a certain point in time, you can switch the replication mode in the tab page [Settings](#) from “Real time” to “Schedule by time”.

3.5.4 Restarting Jobs

If you stopped the master job of a configuration, or if it was aborted, you can restart the master job from the Configuration & Monitoring Dashboard (tab page [Jobs and Connections](#)). The master job resumes as well as all related jobs.

3.6 Changing Load and Replication Procedures

3.6.1 Changing the number of Jobs

Once a configuration is created, you can adjust the number of jobs in the Settings tab page by choosing the [Edit](#) pushbutton.

You may want to change the number of the jobs for the following reasons:

- If you are not satisfied with the speed of the initial load and / or the replication latency time
- If the SAP LT Replication Server has more resources than initially available, you can increase the number of data transfer and / or initial load jobs
- After completion of the initial load, you may want to reduce the number of initial load jobs
- You can customize the number of access plans. The default setting is one, but you can define a higher number of access plan calculation jobs that can run in parallel.

Note

There must be at least one free background job to be able to change any configuration settings in the Configuration & Monitoring Dashboard.

3.6.2 Embedded Health Checks

SAP Landscape Transformation Replication Server runs health checks. There are health checks that cover jobs, for example they check whether all relevant jobs are running for the current configuration. These health checks can support with regard to determining the root cause for an issue related to jobs that are not running. Other health checks cover the tables, for example they check whether all required steps are executed (for example the creation of triggers) for all tables of the current configuration. These checks can support with regard to determining the root cause if a table is stuck on a certain status, as it indicates which steps are not successfully executed yet.

You can access these health checks by opening the monitoring transaction (transaction [LTRO](#)), and navigating to the [Health Check Results](#) tab.

In addition, a health check for ensuring a consistent 1:N replication scenario setup is available.

- Health Check for 1:N Replication

This troubleshooting activity is relevant for 1:N replication scenarios. The system checks the relevant system tables for any inconsistencies that could result in errors or that could affect performance.

In this troubleshooting activity, the system checks the following:

- Whether multiple configuration registrations exist for the same logging table.
- Whether obsolete records exist in SAP LT Replication Server system tables.
- If configurations that are associated with application and logging tables are not registered in SAP LT Replication Server.
- If any tables are incorrectly registered for the 1:N replication process.
- If any configurations are incorrectly set up for the 1:N replication process.

In addition, you can view all the tables in the system that are being replicated, as well as their association with particular configurations and mass transfer IDs.

By default, the system uses the connection specified in configuration (the mass transfer ID) that is currently loaded in the LT Replication Server Cockpit. If required, you can specify another connection in the Source RFC Connection field.

You can view the result of the checks in the section Health Checks. A red traffic light indicates errors. You can navigate directly to detailed information by double-clicking an item in this section.

3.6.3 Performance Improvements for the Initial Load of Tables

The SAP LT Replication Server uses reading type 3 (DB_SETGET) as the default technique to retrieve fixed portions of data records from the source system(s). This procedure is optimized especially for cluster-tables and uses up to three background jobs in parallel for each table.

Note

If a lot of tables are selected for load / replication at the same time, it may happen that there are not enough background jobs available to start the load procedure for all tables immediately. In this case, you may increase the number of initial load loads in the tab page Settings – assuming sufficient system resources are available. Otherwise the initial load of tables will be handled sequentially.

For tables with a large volume of data (especially if they are cluster tables - where the read performance is about 70% compared to transparent tables), you can use the transaction Advanced Replication Settings (transaction LTRS) to further optimize the load and replication procedure for dedicated tables.

In the tab page [Performance Options](#), you can assign more (or less) jobs to run in parallel for a table or define a sequence that should be considered if load and replication of tables should run subsequently.

Furthermore, it is possible to switch the default reading type 3 to another more appropriate reading procedure. The following table compares the advantages and disadvantages of the three most important reading types:

Reading type	Advantages	Disadvantages
1 – Access Plan Calculation	<ul style="list-style-type: none"> • Fast data load if index exists • Parallel data load possible 	<ul style="list-style-type: none"> • Additional index may be required • Requires a key field which is selective enough to significantly speed-up the data access • Calculation time required before load
3 – DB_SETGET (Default)	<ul style="list-style-type: none"> • No separate index required • Parallel data load (multi-threading) with starting with DMIS_2010 SP07 	<ul style="list-style-type: none"> • Additional consumption of database buffer
4 & 5 – Index Cluster	<ul style="list-style-type: none"> • Very fast data load after data is extracted to table DMC_INDCL • Minimal usage of DB buffer 	<ul style="list-style-type: none"> • Additional table space temporarily required in the source system

Note

Other reading types such as “1 access plan calculation” may require the creating of a secondary index beforehand otherwise there will be no positive impact for switching the reading type.

Switching to a specific reading type requires additional preparation steps and expert knowledge – therefore we recommend that you engage the services of an SAP performance expert in this matter.

3.7 Monitoring of Load and Replication Process

The Configuration & Monitoring Dashboard includes several information you can use to monitor and identify potential replication issues.

You can access the Configuration & Monitoring Dashboard by using transaction LTR.

From the overview screen, you can identify the overall status of the configuration.

Details on Status Information:

Status Yellow: May occur if trigger status is set to yellow, for example because the triggers are not yet successfully created.

Status Red: May occur if master job aborted, for example because it was stopped manually in transaction SM37.

3.7.1 Status of Jobs and Connections

On the tab page Jobs and Connections, you can identify potential issues of the configuration-related jobs and / or the connections to the source system and target (SAP HANA) system.

If the job status is not ok, check your current settings with regard to the current system utilization, that is, with regard to the number of available work processes.

Details on status information:

- Master job - Can have status Red for example, for the following reasons:
 - The SAP LT Replication Server was restarted
 - The job was manually stopped in the Configuration and Monitoring Dashboard or by using transaction SM37
 - A support package or SAP Note was applied that changed program code and therefore the ABAP compiler aborted the job
- Data Transfer Jobs – Can have status Red if the number of defined jobs is greater than the number of available jobs.
- Connections:
 - The ABAP source system RFC connection can have the status Red if the source system or network is down (if ping to the ABAP source system exceeds time limit defined in the RFC settings of the SAP LT Replication Server) or authorization of RFC user has changed.
 - Connection to non- ABAP source system and HANA system: Can have the status Red if the open connection fails.

3.7.2 Status of Triggers

Once a table is selected for replication in the SAP HANA Studio, related database triggers are automatically created in the source system to record changes.

In the Triggers tab page, you can view the trigger status of each table in replication. The column Description provides with useful details about the trigger status – this is especially important if status indicator is not green.

Details about status information:

The trigger status may turn to yellow if the creation of the trigger fails for example for one of the following reasons:

- Missing authorization for source system to create trigger
- Related logging tables are for whatever reason corrupted
- If MaxDB or Sybase ASE is used, if other database triggers already exist



Note

If you delete trigger manually in source system, the trigger status will stay green.

3.7.3 Statistics

Similar to the SLT-based Table Data Provisioning screen in the SAP HANA Studio, the [Statistics](#) tab page provides key information about tables. It includes details about the current action and status, and replication latency time of each table (calculated with the median, the minimum, and the maximum).

You can customize the latency threshold for warnings and errors for each configuration at table level. To change the default settings, you can use the expert function [Settings for Notifications](#) in the SAP LT Replication Server Cockpit (under section [Status and Notifications](#) in the [Expert Functions](#) tab page). In this activity, you specify a mass transfer ID, and use the tab page [Latency Thresholds](#) to specify threshold latency values (in seconds) for warning and error conditions for each table. These settings are optional; if you do not make any settings here, tables use default system settings.



Note

For 1:1 replication scenarios, the maximum latency value is calculated based on the oldest timestamp in the logging table and not in current data package portion.

In addition, you can view [Replication Details](#) including throughput information of inserted, updated, and deleted data records.

In the [Statistics](#) tab page, you can view both latency and throughput information (you can switch between the two views).

Note that no statistical data can be displayed for logical tables that are part of a table cluster.

If the throughput or replication latency time does not meet your expectations, you can change the number of jobs as outlined in chapter 3.6.1.

Details on status information:

The statistic status for a table may change to yellow if the replication latency - median (24h) exceeds 5 seconds.

Note

Since calculation of replication latency also considered initial load time, the statistic status most likely will remain on status yellow at least in the first 24h after the successful completion of the initial load.

The statistic status for a table may change to red if the column *Current Status* shows “Error” for example because of the following reasons:

- Connection issues with the HANA system or source system
- Master job was not stopped before applying software maintenance activities in the ABAP source system
- Inconsistent mapping of data formats, such as NC → INT; DATS → DATE

Note

As long as the connection to the SAP HANA system exists, the column status Current Status reflects the same table status as control table RS_Status of the related schema in the SAP HANA system.

In order to reduce the amount of data in your system, the system summarizes statistical data based on both days and hours. By default, the system summarizes statistical data based on days for the previous 30 days, and summarizes statistical data based on hours for the previous 10 days. By default, the system deletes all statistical data that is older than 365 days. You can change these values by choosing *Other Action* → *Change Settings for Statistical Data*. You can change the following variables:

- No. of Days (Based on Hours)
Specify the number of days for which the system summarizes the relevant statistical data based on hours.
- No. of Days (Based on Days)
Specify the number of days for which the system summarizes the relevant statistical data based on days.
- No. of Days (Deletion)
The system deletes all statistical data that is older than the number of days that you specify here (the default value is 365 days).

3.7.4 Using the Monitoring Transaction LTRO

3.7.4.1 Overview

The new monitoring transaction, LTRO, is a tool which runs on the SAP LT Replication Server system and can assist system administrators with obtaining a faster overview about the participating systems and configurations. For example, you can use the monitoring transaction to get information about:

- The number of current running and free batch and dialog processes running on all application servers of the SAP LT Replication Server system
- The availability status for all connected target systems, and the corresponding log entries (warnings and errors)
- The status of multiple configurations (mass transfer IDs) within a single screen

Even though you can use the SAP LT Replication Server Cockpit (transaction LTRC) to view the status of a single configuration, system administrators managing multiple configurations have to switch from configuration to configuration in order to see the relevant status. The benefit of using the monitoring transaction is to have all relevant information in one screen.

3.7.4.2 Overview Tab Page

The [Overview](#) tab page is divided into two sections. The section on the left hand side contains SAP LT Replication Server system-specific information, and the right hand side you can view details about all connected systems.

On the left hand side, under Details, you can view the available application servers for the SAP LT Replication Server system. For each application server, you can view the number of free and available processes, and the maximum runtime (longest running process in seconds) for the processes (in seconds). The number of free process is an indicator for the load on the SAP LT Replication Server system. If a high maximum runtime value occurs, this could indicate an issue with a long running process.

At the top of the tab page, you can view aggregated values for the application servers. This is the total number of free and available processes, and the highest maximum runtime value (the longest running process across all application servers).

On the right hand side, you can view all the target systems that are connected to the SAP LT Replication Server system. For each target system, you can view the target system status (a connection status), as well as the configurations assigned to the target system. Under Target Alerts, the system displays up to 20 database log entries for the last 12 hours. You can customize these settings in table IUUC_RM_PARAMS (field PARAM, value

RCVR_ALERTS_INT for the time period, and value RCVR_ALERTS_NO for the maximum number of database log entries to be displayed).

3.7.4.3 Configuration View Tab Page

In the [Configuration View](#) tab page, you can view the status of the master job, and an overview of all configurations selected when the transaction was executed.

For each configuration, you can view information such as the name, status, technical ID, and so on.

The advantage of this view is that system administrators can see multiple configurations on a single screen without needing to switch from one configuration to another.

You can navigate to the details of an individual configuration by clicking an entry in the [Mass Transfer ID](#) field or by choosing the [Open Tables View](#) pushbutton. To view additional information for the configuration, you can navigate to the [SAP LT Replication Server Cockpit](#) (transaction LTRC) and use the tab pages [Table Overview](#), Data Transfer Monitor, and [Application Logs](#).

3.7.5 Monitoring using SAP Solution Manager

Starting with SAP Solution Manager 7.1 SP05 (and using at least DMIS_2010 SP07 on the SAP LT Replication Server), you can also use the system monitoring capabilities of SAP Solution Manager to monitor the status of a configuration and related schema.

From the SAP Solution Manager system, you can connect to the SAP LT Replication Server during the set-up steps using the SAP template “SAP SLT ABAP Add on” that is automatically proposed to be assigned, if the DMIS Add-on is detected.

After the completion of the set-up steps, you are able to monitor – aggregated for a schema - basic information on job, trigger and table status.

SAP Note [1558756](#) – ‘Solution Manager 7.1. – BI Monitoring: Prerequisites’ describes the prerequisites to be met and the necessary configuration steps in order to integrate the SAP LT Replication Server Monitoring Information into the SAP Solution Manager E2E Alerting Infrastructure. After successful configuration, you can monitor the information provided in the Configuration and Monitoring Dashboard (transaction LTR) in SAP Solution Manager too, and set customized alerts for system conditions you want to receive notifications for.

Note

The system monitoring capabilities of SAP Solution Manager do not allow the viewing of details for each table and related trigger – it contains only aggregated information for a schema.

3.8 Backup and Recovery Aspects

You need to back up your system landscape regularly to ensure that you can restore and recover it in case of system outages or other failures.

The backup and restore strategy for SAP LT Replication Server consists of two parts:

1. Backup and restore coverage for each component (see table below)
2. Cross-system data dependencies and handling

The backup and recovery strategy for your system landscape should not only consider SAP systems but should also be embedded in overall business requirements and incorporate your company's entire process flow.

In addition, the backup and recovery strategy must cover disaster recovery processes, such as the loss of a data center through fire.

Situation	Consequences	Actions
Source system goes down	Replication is stopped SLT waits for source system to be available again	Restart the source system SLT will continue from where it stopped
SAPLT Replication Server goes down (or source system and SLT, if in the same stack)	Replication is stopped	Restart SLT jobs
SAP HANA system goes down	Replication is stopped SLT waits for HANA DB to be available again	Restart the HANA system SLT will continue from where it stopped

If the source system or the SAP HANA system cannot be fully recovered to the same point in time, tables have to be dropped and reloaded into the SAP HANA system again to ensure data consistency between both systems. Therefore, the replication needs to be stopped and restarted for all tables of a schema from the SLT Based Table Data Provisioning screen in the SAP HANA Studio.

4 SAP Landscape Transformation Replication Server Cockpit

4.1 Introduction

The SAP LT Replication Server provides a Web Dynpro application, the Configuration and Monitoring Dashboard, which can be accessed by using transaction LTR. You can use the Configuration and Monitoring Dashboard to define new replication configurations, and view the status, as well as some statistical information, for existing configurations.

For a more detailed analysis, and for additional scenarios that are not yet supported by the Configuration and Monitoring Dashboard, SAP LT Replication Server provides an additional application, the SAP LT Replication Server Cockpit (the SLT Cockpit), which can be accessed by using transaction LTRC.

4.2 Starting and Stopping all Configurations

On the initial Screen of the SLT Cockpit, you can start or stop all replication configurations by using the [Start All Configurations](#) or [Stop All Configurations](#) pushbuttons. This may be convenient when you want to begin or return from a system maintenance activity which requires the stop of SAP LT Replication Server, or if you want to start or stop the SAP LT Replication Server system.

4.3 Prerequisites: Choosing a Mass Transfer ID

In order to use the SLT Cockpit, you must first identify the corresponding mass transfer ID for your replication configuration. There are different ways to get the corresponding mass transfer ID.

In the Configuration and Monitoring Dashboard you can see the corresponding mass transfer ID both on the overview screen or the screen that shows information about a configuration.

Choosing a Mass Transfer ID by using the Search Help

On the initial screen, you can use the input help on Mass Transfer ID input field to choose the relevant mass transfer ID. The input help only displays the mass transfer IDs for which a valid configuration is available and that are displayed in the Configuration and Monitoring Dashboard. Obsolete mass transfer IDs are not displayed in the search help.

Creating a New Configuration / Mass Transfer ID

For scenarios that are currently not supported in the Configuration and Monitoring Dashboard, you have to create the new configuration (and mass transfer ID) by using the SLT Cockpit. Choose the Create New Configuration pushbutton to do so.

4.4 Navigating from the SAP Landscape Transformation Replication Server Cockpit

From the SAP LT Replication Server Cockpit, you can use the [Goto](#) menu in the navigation bar to navigate to related applications. You can navigate directly to the:

- Configuration and Monitoring Dashboard (transaction LTR) - the system opens a new session in your Web browser.
- SAP LT Replication Server Monitoring transaction (transaction LTRO)
- Advanced Replication Settings (transaction LTRS)

4.5 SAP Landscape Transformation Replication Server Cockpit Tab Pages

The SAP LT Replication Server Cockpit contains the following tab pages:

Administration Data

On the [Administration Data](#) tab page, you can view the following sections:

General Data

You can view general information, such as the name of the configuration, the replication scenario, and details about the source and target systems. For more information about the configuration, choose the [Show Configuration Data](#) pushbutton.

Authorization

You can view the authorization group (if one exists), or change it if required by using the [Change Authorization Group](#) pushbutton.

If you require specific authorizations for the configuration, you can specify an authorization group, and use it in the corresponding authorization object S_DMIS_SLT.

Job Options

You can view the status of the configuration, for example [Running](#) or [Finished](#). You can start and stop the configuration as required. Note that all configurations are handled by one master job, and this master job only stops if no configurations are active.

You can view and change the following:

- No. of Data Transfer Jobs
 - You can specify the number of jobs that are used for the data transfer process in the SAP Landscape Transformation Replication Server.
- No. of Initial Load Jobs
 - You can specify the number of jobs that are used for the initial load in the SAP Landscape Transformation Replication Server.
- No. of Calculation Jobs

-
- For reading types 1 and 3, this is the number of jobs that are used to calculate the data transfer portions that are used for the initial load. For reading types 4 and 5, it is the number of jobs that transfer the portions to table DMC_INDCL.

For both the source system and the replication server the job class, user and execution server for the data load jobs can be individually defined for each configuration.

To display an overview about current configuration settings, choose the [Show Configuration Data](#) pushbutton.

Processing Steps

The [Processing Steps](#) tab page contains all the steps required to start or stop an initial load or a replication.

You can execute the preparatory steps for an initial load or for a replication. These steps are normally executed by the replication controller job automatically, but you can use this function to execute the steps if any issues have occurred.

To view the documentation for a step, select it, and choose the [Documentation](#) pushbutton. To execute a step, select it, and choose the [Execute](#) pushbutton.

There are two options:

- Execute all preparatory steps with one activity. If you execute this activity, the system only executes the steps that have not completed successfully.
- Execute a dedicated single preparatory step.

If there are inconsistent objects or settings for a particular table you can use the menu options to delete a load/replication object, the related generated runtime object, or even use the option to reset all objects and settings for a table.

Table Overview

In the [Table Overview](#) tab page, you can view all tables where an initial load or a replication has been initiated. You can use the table to identify what action was started for a table.

For a triggered replication, the field [Logging Table](#) indicates the logging table name, and the field [Logging Table Generated](#) indicates whether the logging table has already been created in the sender system. The field [Trigger State](#) indicates whether the trigger for a certain table is already active or still has to be created. In the columns [Local Table in Receiver](#), [Table in Receiver](#) and [Synonym in Receiver](#), you can view the status of the objects to be created in the target system.

If one of the following steps is in process, or fails, the status of the corresponding table will be updated accordingly (in columns [In Process](#) and [Failed](#)):

- Create logging table in source system
- Create trigger in source system
- Create table in target system
- Create synonym in target system
- Create load / replication object

Note

If any errors exist for a table, you can choose the [Show Error Log](#) pushbutton to view them.

Data Transfer Monitor

As soon as the data load or transfer object is created, you can view it in the Data Transfer Monitor tab page. If a table is displayed in the tab page [Table Overview](#) but not in the tab page [Data Transfer Monitor](#), then one of the following steps has not yet completed successfully:

- Create logging table in source system
- Create trigger in source system
- Create table in target system
- Create synonym in target system
- Create load / replication object

Note that for cluster tables, you can find the name of the logical table in the tab page [Table Overview](#), while in the tab page [Data Transfer Monitor](#) you can view the logical table name during the initial load and the cluster table name during the replication.

Note

If any errors exist for a table, you can choose the [Show Error Log](#) pushbutton to view them.

Application Logs

The system saves the logs of the individual work steps, and the error logs for the data load and the replication in the SAP Application Log. You can view the log by using transaction SLG1, or in the tab page [Application Logs](#). In order to display the logs, you must first define filter criteria. To find the root cause of an error it makes sense to limit the message output to those message logs which contain at least one error message (only logs with errors).

Load Statistics

The display of runtime information is integrated into the SLT Cockpit in the tab page [Load Statistics](#). You can view runtime information (for example the number of loaded portions, runtime statistics, and so on) in this tab page during the initial load.

The runtime information of the initial load is not deleted when the object switches to the delta mode and can be displayed, when switching the selection to [Load – Finished](#).

Expert Functions

In the Expert Functions tab page, you can find all required functions to reset a table from an error status, if the error cannot be resolved by the tool automatically. The new expert functions are combinations of the old existing ones, and are provided for the individual work steps. For a detailed explanation of prerequisites, available features, required inputs, and the output of these functions, refer to the documentation that is accessible from the UI. An overview of the individual expert functions is provided in the following tables:

Expert Functions for Reset Indicator / Status

Function	Description
Reset Status for Triggers and Logging Tables	<p>In this step, you can reset the relevant flags for database triggers and logging tables. You can view the flags on the tab page Table Overview.</p> <p>If an issue occurs during the creation of database triggers or logging tables, the system sets the status of the table to Failed. If the process that creates logging tables or triggers stopped unexpectedly, the table may remain at status In Process, but the table will not be processed any further. In such situations, you can reset the Failed or In Process flag so that the required processing is started again.</p> <p>You can also use this step to reset the creation status for triggers or logging tables. We recommend that you use this feature very carefully, as you should only reset these flags if you are sure that the logging tables or triggers do not exist in the source system. If the initial load or the replication is already running, do not reset the flags, as this would result in the loss of delta data. Instead, you can stop the load or the replication and</p>

Function	Description
	start the table again.
Reset Status of Tables and Synonyms	<p>In this step, you can reset the relevant flags for database tables and synonyms in the source or receiver systems. You can reset flags for the tables listed on the Table Overview tab page.</p> <p>If an issue occurs during the creation of the database table in the target system, or during the creation of the synonyms used to communicate with the source system (only relevant for direct database connections) or target system, the system sets the status of the table to Failed. If the process that creates the target table and synonyms stopped unexpectedly, the table may remain in the status In Process, but the table will not be processed any further. In such situations, you can reset the Failed or In Process flag so that the required processing is started again.</p> <p>You can also reset the creation status for the relevant tables and synonyms that have been created in the source or target system.</p>
Reset Replication Object Definition	<p>In this step, you can reset the relevant flags for the definition of the initial load or replication object.</p> <p>You can use this step if you changed the replication content definition, and you want to recreate the initial load or replication object. The system resets the relevant flags but does not reset the definition of the initial load or replication object. If you want to delete an existing initial load or replication object, and not just reset the flags, use the step Delete Load / Replication Objects (accessible from the tab page Processing Steps) instead. This step also resets the status.</p> <p>Only reset the flags if you are sure that the flags are not defined properly and therefore must be reset. If the load or replication object still exists, use the step Delete Load / Replication Objects instead.</p>
Reset Runtime Object Flags	<p>In this step, you can reset the relevant flags for the generated runtime objects (generated modules for data load or replication). You can view the flags on the tab page Data Transfer Monitor.</p> <p>If you reset the Generated flag, the system regenerates the object automatically if the master job is running. You can reset the flag at any time as the initial load can continue from where it stopped before (as long as you only reset the flags and does not change the reading type for the table). However, do not reset the flag if the object is still being processed by a job. We therefore</p>

Function	Description
	<p>recommend that you stop the jobs for the configuration when you reset the generated flags.</p>
Reset Load and Replication Status	<p>In this step, you can reset the flags that are relevant for the data load or replication. These flags are:</p> <ul style="list-style-type: none"> • Error • In Process • Loaded • Blocked Processing Step <p>You can view these flags on the tab page Data Transfer Monitor.</p> <p>As long as the load or the replication is running properly, we recommend that you do not reset any of the flags (the system resets the Error flag automatically, and retries the load or the replication periodically). We recommend that you exercise extreme caution when using this step, as resetting these flags may result in duplicate records or data inconsistencies in the target system.</p> <p>If a load or replication object terminates, a table may remain in the status In Process. If you observe that more tables are in process than there are load and replication jobs defined, this can indicate that some tables are stuck in the status In Process without making any progress. You should stop the jobs, and examine the tab page Data Transfer Monitor. If jobs are stopped, and some tables are still In Process, you can reset the In Process flag for those tables. If you reset the flag while load jobs are still running, the same portions of a table might be handled by different jobs, therefore resulting in duplicate key errors.</p> <p>The load jobs automatically reset the Failed flag, but you can also reset the flag at any time as a table is blocked as long as it remains in an error state.</p> <p>The system resets the Loaded flag for the replication objects automatically, but the flag will remain for finished initial loads. If you reset the Loaded flag for an initial load, the system resets the Loaded flag for the selected tables but not for the loaded portions. You can reset the Loaded flag for tables that are being replicated, as a table is blocked as long as it remains in a loaded state.</p>

Function	Description
	<p>You can also reset the Blocked Processing Step flag. This flag is used to ensure that the initial load will not start before delta recording is active in the source system. The system also sets this flag if a trigger is deleted in the source system while the initial load or the replication is running.</p> <p>Note that there is a high risk of data inconsistencies if you start the load while delta recording is not active, or if you just resume replication if triggers were not active for some time. Therefore, we recommend that you exercise extreme caution when resetting the Blocked Processing Step flag.</p>

Expert Functions for Information and Analyses - Display Table Information

Function	Description
Display Table Information	<p>In this step, you can view information about tables from the sender system. You can select tables that have been selected for data provisioning (load or replication) or a list of individual tables (also including tables that are not started yet).</p> <p>You can use this step to view the table size in the source system and also the information on created trigger, logging tables and registered replication consumers (in case of 1:N replication) that replicate the same table to different targets.</p>

Expert Functions for Information and Analyses - Triggers and Logging Tables

Function	Description
View Trigger Source Code	<p>In this step, you can view the trigger source code for a given database table and mass transfer ID. You can view either the active or generated trigger statements.</p> <p>An active trigger statement contains the current active SQL trigger statement, while a generated trigger statement displays the SQL trigger statements that would be generated with the current settings. These</p>

Function	Description
	statements are not yet active.
View Unprocessed Logging Table records	You can use this step to check the logging tables in the source system for any remaining unprocessed logging records. The system displays the oldest timestamp for an unprocessed record in the logging table, and also the number of unprocessed records.
Count Logging Table Records	In this step, the system counts the number of records in the logging table and also the number of processed records.

Expert Functions for Information and Analyses - Status and Notifications

Function	Description
Status Notification Settings	It is important that the system administrator responsible for SAP LT Replication Server is notified of any issues regarding the replication process as soon as possible. To avoid permanent monitoring of the SAP LT Replication Server system, you can activate notifications for a configuration. If notifications are active for a configuration, the system uses a background job to check the configuration periodically (the notification job).
Job Health Check	<p>In this troubleshooting activity, the system analyzes the jobs involved in data load and replication processes for the SAP LT Replication Server.</p> <p>The system checks the following:</p> <ul style="list-style-type: none"> • Whether the connections to the source and target system are working correctly. • Whether all data load jobs are running for the current configuration • Whether the load and replication processes are working correctly, and if jobs are waiting for any resources (for example table locks, the connection to the HANA system, or free batch processes). • Whether the job status is valid.
Table Health Check	In this troubleshooting activity, the system verifies the integrity of the SAP LT Replication Server configuration, as well as many operational aspects

Function	Description
	<p>involved in the table load and replication processes.</p> <p>This troubleshooting activity is relevant for tables that have the status <i>Error</i> or if a table remains on status <i>Scheduled</i>. You can execute this troubleshooting activity for all tables of the current schema or for selected tables.</p>
1:N Replication Health Check	<p>This troubleshooting activity is relevant for 1:N replication scenarios. The system checks the relevant system tables for any inconsistencies that could result in errors or that could affect performance.</p> <p>In this troubleshooting activity, the system checks the following:</p> <ul style="list-style-type: none"> • Whether multiple configuration registrations exist for the same logging table. • Whether obsolete records exist in SAP LT Replication Server system tables. • If configurations that are associated with application and logging tables are not registered in SAP LT Replication Server. • If any tables are incorrectly registered for the 1:N replication process. • If any configurations are incorrectly set up for the 1:N replication process. <p>In addition, you can view all the tables in the system that are being replicated, as well as their association with particular configurations and mass transfer IDs.</p>

Expert Functions for Information and Analyses - Display Replication Statistics

Function	Description
Display Replication Statistics	<p>In this step, you can view statistical information for tables. This statistical information includes, for example, the number of inserted, updated, and deleted records for each table, as well as latency information.</p>

Repair Replication Tables

Function	Description
Resolve Inconsistencies Identified by CDC Application	<p>When you use the Cross-Database Comparison (CDC) application to analyze the source and target systems with regard to data consistency, the application can identify the inconsistencies outlined below.</p> <p>In this expert function, you can use this information to resolve these inconsistencies.</p> <p>The Cross-Database Comparison (CDC) application can identify the following inconsistencies, and take appropriate action:</p> <ul style="list-style-type: none"> ○ Objects exist in the source system, but not in the target system. SAP LT Replication Server inserts the missing rows in the target system ○ Objects exist in the target system but not in the source system. SAP LT Replication Server removes the rows from the target system. ○ Objects exist in both systems, but there are differences with regard to data. SAP LT Replication Server adjusts the values in the target system so that they are consistent with the values in the sender system. <p>Prerequisites:</p> <p>You have successfully configured and used the Cross-Database Comparison (CDC) application on a connected SAP Solution Manager system to analyze the source and target systems.</p> <p>For more information, see the documentation that is available in the system for this expert function.</p>

Additional Functions:

Function	Description
Logged Replication Data	<p>When a table is replicated to the target system, the replicated data can also be saved to database tables in the SAP LT Replication Server system for a period of time. This means if any data is missing from the target system, you can view this data and replicate it again from the SAP LT Replication Server system to the target system.</p> <p>All tables having already transferred data in replication stage (no initial load stage) will be shown and selectable</p>

Function	Description
	<p>For all selectable tables it is possible to see when and how many Inserts/Deletes and Update-Statements have been performed.</p> <p>Furthermore, if you have appropriate authorization assigned (add Authority object S_DMIS with Activity 29, this is not part of any SLT standard role) you can display the data content of the replication packages when clicking on the Package ID. Otherwise an error message in the status bar will be provided.</p> <p>For more information see the embedded system help of this expert function.</p>
Define Connection to Alternative System for Initial Load	<p>In this step, you can define an alternative sender system that is used for the initial load. To do this, you specify an RFC destination or a database connection to the alternative system. You can also delete the connection information for an alternative sender system.</p> <p>After you have defined an alternative sender system, the system connects to the system as a data source during the initial load whenever you:</p> <ul style="list-style-type: none"> Add a new table to be replicated for the configuration (schema) in the SAP HANA Studio. Stop the replication of a table for this configuration and restart it in the SAP HANA Studio. Add a new table to be loaded only for the configuration in the SAP HANA Studio. Stop the replication of a table for the configuration, and add it to be loaded only (in the SAP HANA Studio). <p>The system uses an RFC destination or database connection to the alternative sender system during the Initial Load phase. For a replication, as soon as the system switches to Delta mode, the system also switches the sender destination and connects to the sender system specified in the configuration.</p> <p>If the replication sender system is connected by means of an RFC destination, you can choose whether you want to connect the alternative sender system by means of an RFC destination or a database connection. If the replication sender system is connected by means of a database connection, you</p>

Function	Description
	must also use a database connection to the alternative sender system.
Change Settings for Connection to Target System	<p>In this step, you can change the connection settings to a target system for a mass transfer ID.</p> <p>You can change the hostname and instance number, and set the password to a new one (by choosing the Change Password pushbutton).</p> <p>There are two connections, one for setting up the configuration, and the other for handling the actual data replication.</p> <p>Note for the Connection for the Data Replication:</p> <p>If you have a multi-node HANA system, enter all nodes in the Connection for Data Replication section; this ensures that if one node is not available, the database connection will use the next node.</p>
Download Load Statistics to File	This function allows for downloading the load statistics into a file.

5 Special Considerations for Source Systems

5.1 Impact of Software Maintenance activities in ABAP Source Systems

Software maintenance activities (such as applying support packages or using the SAP transport management system TMS) in the source system may affect tables that have the status In Replication using active database triggers.

What happens if there are structural changes on tables with active DB trigger?

When a table is registered for replication, a corresponding logging table and a database trigger are created in the source system to record any INSERT, UPDATE, or DELETE statements. The question is now how the SLT system and the source system behave regarding any structural changes to a table for which a DB trigger is active. As the behavior depends on the SAP Basis release on the source system the behavior is described for the respective Basis version. For the SLT System we strongly recommend to use DMIS 2011 SP2 or higher / DMIS 2010 SP7 or higher. If you use an older SLT version you have to stop and restart the affected tables in any case.

As the behavior is the same no matter if the structural change is done directly in the system or executed within a transport, we do not have to distinguish between both use cases.

Source System on SAP NetWeaver Basis 6.40 or lower / non-ABAP Source system:

There is no check for active trigger in the ABAP dictionary. I.e. any structural change can be executed on a table no matter if a trigger is active for that table.

If SLT is installed on DMIS 2011 SP2 or higher / DMIS 2010 SP7 or higher, SLT will recognize structural changes and will adapt the table on the HANA side accordingly. As the logging table stores the key of the original table, the trigger and logging table have to be dropped and recreated in case of changes of key definition. Please ensure that the trigger is dropped before the key definition is changed as the trigger will fail once the key definition is changed. Therefore in case of key changes you have to stop the replication for the affected table and start the replication (including a new load) after the applying the change to the table.

Source System on SAP NetWeaver Basis 7.00 - 7.30:

There is a check in the ABAP dictionary which prevents any structural change on a table where a database trigger is active. There are two notes available which downport the adaption which is part of the standard installation starting with SAP NetWeaver 7.31 SP5 and higher releases (enhancement packages). Note that both notes handle SLT triggers only. If customer-specific triggers are active on a table also, any structural change will still be prevented.

SAP Note 1655725 – ‘Check for DB trigger made less strict’ together with SAP note 1875769 - ‘Check for DB trigger made less strict (not null)’:

By implementing these two notes on the ABAP source system, simple table structure changes (adding a non-key field) are possible with active SLT triggers. The trigger will remain active and delta recording continues without any interruption. SLT checks before every replication if the activation timestamp of the source table changed to detect any structural change. When the table is changed on the source system, SLT pauses the replication of that table, calculates the ALTER statement for the HANA DB and does the same adjustment (including DEFAULT VALUE statements) on the HANA side. Afterwards the replication object is adjusted to the new structure and replication continues without the need of a new initial load. In this case the table is paused for a few minutes, therefore Latency might go up for the first portions after the table adjustment was processed automatically. For the automatic adaption by SLT, DMIS 2011 SP2 or higher / DMIS 2010 SP7 or higher is required on the SLT System.

SAP Note 1749824 - Treatment of SLT triggers during transport:

In case of complex table structure changes (change data type or length, add or delete key fields) the trigger and logging table have to be dropped and recreated. Also the table on HANA side has to be dropped and recreated to ensure data consistency. Therefore you normally have to stop the replication for the affected tables upfront before applying a complex structural change. By implementing this note, the data dictionary checks in case of complex structural changes if the active triggers are SLT triggers and will drop the SLT trigger to ensure that the structural change can be applied. SLT permanently checks if the created triggers still exist in the source system. In case a trigger is dropped manually or by the data dictionary, SLT suspends the replication for that table and sets it to error state with an according error message in the application log (that the trigger was deleted). So this note does not prevent that you have to reload the table after a complex structural change but you can apply complex structural changes with stopping the replication of the affected table upfront.

For additional information about simple and complex table structure changes, see SAP Knowledge Base Article [2254415](#) – ‘Source structure change(s) not moving over to HANA system’.

Source System on SAP NetWeaver Basis 7.31 SP5 or higher:

Simple Table Structure Changes

Simple table structure changes (adding a non-key field) are possible with active SLT triggers. The trigger will remain active and delta recording continues without any interruption. SLT checks before every replication if the activation timestamp of the source table changed to detect any structural change. If the table is changed in the source system, SLT will pause the replication of that table, calculate the ALTER statement for the HANA DB and do the same adjustment (including DEFAULT VALUE statements) on the SAP HANA side. Afterwards the replication object is adjusted to the new structure and replication continues without the need of a new initial load. In this case the table is paused for a few minutes, therefore Latency might go up for the first portions after the table adjustment was processed automatically. For the automatic adaption by SLT, DMIS 2011 SP2 or higher / DMIS 2010 SP7 or higher is required on the SLT System.

Complex Table Structure Changes

In case of complex table structure changes (change data type or length, add or delete key fields) the trigger and logging table have to be dropped and recreated. Also the table on HANA side has to be dropped and recreated to ensure data consistency. Therefore you normally have to stop the replication for the affected tables upfront before applying a complex structural change. The data dictionary checks in case of complex structural changes if the active triggers are SLT triggers and will drop the SLT trigger to ensure that the structural change can be applied. SLT permanently checks if the created triggers still exist in the source system. In case a trigger is dropped manually or by the data dictionary, SLT suspends the replication for that table and sets it to error state with an according error message in the application log (that the trigger was deleted). So the table has to be reloaded again after a complex structural change but you can apply complex structural changes with stopping the replication of the affected table upfront.

Transport Analyzer for Trigger Conflicts

Starting with DMIS 2011 SP3 / DMIS 2010 SP8, SLT provides a transport analyzer to identify if a table is affected by a structural change when certain transports are imported. The analysis program can be created in any system by using SAP Note 1728711. As the program does not use any other DMIS objects you can install the program also on a system where no DMIS add on is installed.

The transport analyzer evaluates the selected transport request and identifies whether they contain any structural changes for database tables. The transport requests are checked whether they contain any of the following objects:

- *Table
- *Table Definition
- *Table Type
- *Table Type Definition
- *Pool / Cluster Table
- *Text Component of Table Type D
- *Table Index
- *Selection View
- *View

As not every change for these objects might lead to a structural change to the database, the result list could be restricted to those tables which are activated within a certain time frame.

Note that the analysis may result in a few tables which are not really changed at database level, as the analysis works with the above objects and where-used lists. This means that if a change is done and discarded afterwards, the transport analyzer will still return it as a potential structural change.

Recommendation how to deal with SLT Replication during a Software Maintenance Event:

If you use the correct SAP NetWeaver Basis Release, or the related SAP Notes for the related SAP NetWeaver Basis release, the created SAP LT Replication Server DB triggers will remain active during potential Data Dictionary changes (adding of non-key fields) or will be automatically deleted by the Data Dictionary tools.

Knowing this, you can in principle continue to have the triggers active and only need to re-load those tables where the triggers have been deleted during a software maintenance activity.

The SAP LT Replication Server checks for missing DB triggers and reports this error condition.

Note

In case that XPRA's or After import Methods (AIMs) change data records on a larger table (usually a very rare case) you better stop the replication in order not to increase the upgrade processing time. Furthermore, in such a case a reload using parallel import processes is very likely much faster than continuing the replication after software maintenance.

Recommendation:

In general SAP recommends for any SAP software upgrade/update activity as well as for imports of structural changes into a SLT source system that you stop the data replication from the concerned source system and start new replication processes (incl. initial load) after successful end of the software upgrade/update/import process.

If this is not acceptable, because of a long runtime of initial load, a test upgrade should be performed on a comparable test system (i.e. a copy of production system) while keeping the DB triggers active. At the beginning of the downtime phases of the software maintenance process you should pause the replication by stopping the data transfer jobs (that is, suspend the replication) so that the DB triggers are still able to record changes. After the upgrade you need to analyze if triggers have been deleted and if there is a high volume of changes recorded in the remaining logging tables due to a XPRA/AIM run on the affected table. Only for those tables a new initial reload is necessary, for all other tables it is sufficient, that you restart the replication process so that the consuming system gets back in sync, then.

Note

If you cannot restart the replication, including a full initial load, contact an SLT expert or create a message using component HAN-DP-LTR.

5.2 Archiving Data in Source Systems

The trigger-based replication also considers the deletion in source tables by archive activities (since it is not possible to distinguish on the database level between delete actions caused by archiving versus regular deletion of data records). Consequently, SAP LT Replication Server will also replicate archiving activities as delete actions in the SAP HANA database.

If archived data of ABAP source systems should to be also available in the SAP HANA database, you can use report IUUC_CREATE_ARCHIVE_OBJECT.

The report creates a replication object, allows the selection of relevant tables of an archive object and the loading of the archived data (by the date of the archiving session) into the respective the schema.

As a technical prerequisite, related SAP ILM (Information Lifecycle Management) APIs need to be available in the ABAP source system. For more information, see SAP Note 1652039.

Note

Similar to the initial load procedure, for the archive load procedure, no trigger and no logging tables are created.

5.3 Preventing Archive Deletes from being Replicated to a Target System

Archiving activities can occur in the source system at any time. If data is being replicated from a source system, it is important to understand the impact that any archiving activities can have on the target system. When SAP LT Replication Server replicates a table to the target system, it creates a database trigger that records any changes in a logging table. If a table is being replicated, and a record in the table is archived, the database trigger interprets the archiving operation as a DELETE statement, and SAP Landscape Transformation Replication Server deletes the corresponding record from the target table. However, you may want to prevent the record from being deleted in the target system.

There are two options to prevent the record from being deleted in the target system:

1. Specify a User

If a user archives data from any table in the source system, you can prevent the data from being deleted from the corresponding table in the target system.

Note

This option is only supported for the following databases:

- Oracle
- DB2 z/OS
- DB2 LUW (DB6)
- SAP HANA
- Microsoft SQL Server

This option is not supported for the following databases:

- DB4
- Sybase ASE
- SAP MaxDB

For more information, see SAP Note 2585153.

If a user archives data from any table in the source system, you can prevent the data from being deleted from the corresponding table in the target system. To do this, proceed as follows:

1. In the SAP Landscape Transformation Replication Server system, open transaction LTRS (Advanced Replication Settings)
2. Under Trigger Options, select the option Do Not Delete Corresponding Target Table Record,
3. Specify the relevant user (or users).

2. Specify an Application Server

As a prerequisite the archive process needs to run on a dedicated, separate application server in the source system. No other data processing should take place on this server.

Prerequisites:

Some ADK (Archive Development Kit) prerequisite steps are required in order to ensure that all archive jobs run on a dedicated Application Server or dedicated server group. You need to create a server group by using transaction SM61. The server group contains the designated application server on which the archiving jobs should run. After you have created the server group, you specify the server group name in [Cross-Archiving Object Customizing](#) (transaction SARA → [Customizing](#) → [Cross-Archiving Object Customizing](#) → [Technical Settings](#)).

Once this is done, all archiving-related jobs (write jobs, delete jobs, and so on) will be scheduled and executed only on the assigned server or server group. The ADK-related documentation can be found in the SAP Help Portal (help.sap.com) for example for systems based on SAP NetWeaver 7.0 EHP2 Stack, see [Data Archiving in the ABAP Application System](#).

Configuration steps in the source system and in the SAP LT Replication Server system

- In the source system in table IUUC_RT_PARAMS, set the field IU_PARAMETER to **ARCHSERVER**, and the field IU_VALUE to <the Name of Application Server> (note that you can only enter one application server).
- In the SAP LT Server system in table IUUC_REPL_CONFIG, set the field IGNO_ARCHIVE_DEL to **x** for the relevant mass transfer ID (MT_ID). Note that the field IGNO_ARCHIVE_DEL can have the content 'x', '-', or ''.

Note

The system creates the trigger code according to settings in the tables above. Therefore, you should implement these configuration steps before the relevant tables are replicated (before the system has created triggers for these tables).

If you want to activate this feature in a configuration which is already replicating data, you must delete the triggers for the relevant tables and re-create them after implementing the configuration steps described above. In other words, if you already have a replication running, this would require that the replication is stopped, and that a new replication (including an initial load of the relevant tables) is started. To avoid repeating the initial load, you can apply SAP note [1898479](#). This SAP Note contains the report IUUC_REDEFINE_DB_TRIGGERS. In this report, you can specify the mass transfer ID (MT_ID), leave field [Lock Timeout Value](#) unchanged, and then specify the relevant tables in the [Table Name](#) field. This means that you do not need to stop and restart the replication (including dropping and recreate triggers and reloading the relevant tables).

5.4 Performance Implications in Source System using Trigger-based Data Replication

Considering the relationship of jobs as outlined before, the main system load of the trigger-based replication approach resides on the SAP LT Replication Server.

Detailed experience with the trigger-based technology has proven that the performance impact of establishing triggers (only for the limited number of tables that have the status In Replication) on the database of the source system is not significant and in many cases not measurable compared to the transactional load of the SAP application.

There is direct relation between the configuration of number SLT data load/replication jobs and the number of batch work process for SLT server system and dialog work processes in the source system.

The source system and SLT system should be configured to have appropriate number of background jobs and dialog work processes to ensure parallel execution and full allocation of the hardware configuration which is sized for the expected replications. An insufficient number of work processes results in a sequential execution of the replication and therefore in much longer execution times. On the other hand, an insufficient number of available work processes in the source system may impact the performance of business applications running there.

For each SLT configuration, the parameter Data Transfer Jobs restricts the maximum number of data load jobs, which can be started for one mass transfer ID (MT_ID).

In the source system, the number of available DIA work processes, which are reserved for SLT replication should be equal to the number of configured Data Transfer Jobs on SLT side. In addition free work processes are required when creating new logging tables and database triggers, as well as creating runtime objects and calculating access plans.

Note that these are the net requirements for SLT Replication only! The source system and SLT system should have enough additional DIA and BGD work processes to handle their own tasks.

The Sizing Guide for SAP Landscape Transformation (SLT) Replication Server for SAP HANA provides detailed information about how to calculate the necessary system resources for a given configuration. For more information, see <https://service.sap.com/sizing> → Sizing Guidelines → Database and Technology → SAP In-Memory Computing → SAP Landscape Transformation Replication Server, SAP HANA

5.5 Improving Performance for Configurations with Many Tables and Low Change Rates

By default, SAP LT Replication Server is configured to handle high change rates and high data transfer volumes. This involves processing all logging tables, regardless of whether they contain entries. However, if your configuration has many tables in replication mode, and changes are occurring to only some of these tables, there is alternative approach.

Background

A logging table is a table in the source system that records any changes to a table is being replicated. This ensures that SAP LT Replication Server can replicate these changes to the target system.

In general, for each application table being replicated, there exists a corresponding logging table. If changes occur to the application table, SAP LT Replication Server creates entries in the logging table to record the changes.

In order to replicate changes from the source system application tables to the target system, SAP LT Replication Server checks the logging tables to determine whether entries exist.

If entries do not exist, this indicates that no changes have occurred to the corresponding application table. SAP LT Replication Server proceeds to check the next logging table in the queue.

If entries exist, this indicates that changes have occurred to the corresponding application table. SAP LT Replication Server replicates these changes to the target system immediately, and proceeds to check the next logging table in the queue.

Note

You can view the progress of the process in the [Data Transfer Monitor](#) tab page of the SAP LT Replication Server Cockpit (transaction LTRC). The data transfer jobs check the tables sequentially. A free data transfer job will reset the [In Process](#) flag for a table, and check whether the corresponding logging table contains new entries. As a result, the [Loaded](#) flags on the Data Transfer Monitor tab page are constantly changing (an X in the column [Loaded](#) indicates that SAP LT Replication Server has processed the table).

Changing the Default Behavior for Configurations with Many Tables and Low Change Rates

If a configuration has a large number of tables, and changes are occurring to only some of these tables, there is no need for all work processes to permanently scan for new logging table entries. In order to improve performance for this scenario, you can use the option [Only Process Logging Tables with Records](#).

If you select this option, SAP LT Replication Server first performs a check to determine which logging tables contain entries, and then only processes these tables. This results in a significantly reduced number of active data transfer jobs and related background work processes.

You can find this option in the SAP LT Replication Server Cockpit (transaction LTRC) under Utilities -> Specify Options for Logging Tables.

You need to activate this feature by stopping and restarting all jobs for the configuration. You can deactivate and activate the configuration by using the corresponding pushbuttons on the [Administration Data](#) tab page of the SAP LT Replication Server Cockpit (transaction LTRC).

Note

In order to use this feature, DMIS 2011 SP6 or higher must be installed in the source system and DMIS 2011 SP8 must be installed in the SLT system.

Examples

You have tables being replicated, and frequent changes are occurring to most of these tables. The option [Process all Logging Tables](#) is the most suitable.

You have many tables being replicated (for example more than 100), and changes are occurring to only some of these tables. The option [Only Process Logging Tables with Records](#) is the most suitable option. This allows you to reduce the number of data transfer jobs that process the tables, and still replicate to the target system in real time.

Additional Information

In the SAP LT Replication Server Cockpit (transaction LTRC), in the tab page Data Transfer Monitor, you can view whether the tables have been processed. An X in the column *Loaded* indicates that SAP LT Replication Server has processed the table.

5.6 Data Volume Management

Data replication using SAP LT Replication Server transfers data from the source system(s) to the target (SAP HANA) system. The following considerations apply:

Source system(s):

- SLT related logging tables: The size of logging tables may increase in size, if the replication is suspended for a long period of time and/or there is a system outage of the system of the SAP LT Replication Server or the SAP HANA system. You should carefully monitor the table space size (assuming logging tables have been assigned to a separate table space) and take appropriate action early enough.
- RFC logging: Since data replication from ABAP source systems is managed by RFC connection, the size of related log files should be reviewed. Especially, if performance improvements are implemented that will increase the number of jobs for initial load / replication, check if RFC related logging is critical.

SAP LT Replication Server:

Since the data transfer – managed in dedicated portions of data records - during the replication process is handled by the memory of the SAP LT Replication Server, there are no specific data volume issues to be considered. In case of system (source and/or SAP HANA system) or network outages, the replication process will generate logging information for related failures. To avoid unnecessary logging information from being stored, you can pause the replication by stopping the schema-related jobs as outlined in chapter 4.2.3.

SAP HANA system:

Depending on the number of tables that have the status In Replication (or initially loaded only) and the growth rate of tables in the source system, the size of the SAP HANA database will increase accordingly.

Note

Consider that the SAP HANA database supports compression rates – therefore the table size in the SAP HANA database may be different to that of the source system.

If the table size (in the SAP HANA database) exceeds 2 billion records, you must split the table by using the available partitioning features.

6 Replication Logging

6.1 Overview

When a table is replicated to the target system, the replicated data can also be saved to database tables in the SAP LT Replication Server system for a configurable period of time. This means if any data is missing from the target system (for instance after a point-in-time recovery of the target system to a previous state), you can view this data and replicate it again from the SAP LT Replication Server system to the target system.

6.2 Prerequisites

This activity is relevant only for configurations for which the checkbox [Activate Replication Logging](#) was selected when the configuration was created in the [Configuration and Monitoring Dashboard](#) (for more information, see section 3.1.3, [Creating a Configuration](#) → [Replication Options](#)). If you choose this option, the replicated data packages from all tables which are in replication mode within this configuration will be saved in the SAP LT Replication Server system.

6.3 Changing the Replication Logging Settings

You can use transaction LTRS (Advanced Replication Settings) to activate replication logging for an existing configuration. You can also deactivate replication logging here.

Under [Replication Logging](#), you can specify settings for replication logging. If replication logging is active, the system saves the replicated data entries in the SAP LT Replication Server system. This means that if any data is missing from the target system, you can view this data and replicate it again from the SAP LT Replication Server system to the target system.

By default, replication logging applies to all tables in a configuration. Under [Table-Specific Settings](#), you can activate replication logging for specific tables in a configuration. You can also change the recording interval for the tables.

6.3.1 Operational Considerations

Replicated data packages are stored in database tables in the SAP LT Replication Server for a configurable period of time. The default retention time is 3 days. Data packages which have been read from SAP cluster tables are stored in table IUUC_RL_DATA_CL. All other data packages (originating from transparent and SAP pool tables) are stored in table IUUC_RL_DATA.

Note

Storing the replicated data packages on the SLT Replication server requires sufficient space in the underlying database. Data growth in the above mentioned tables needs to be carefully monitored and managed and is dependent from the number of tables that are in replication logging state, the chosen data retention period, and the volume of data changes on the concerned source tables.

6.3.1.1 Display and Retransfer Replicated Data Packages

You can display and, if necessary, retransfer already replicated data packages by the expert function [Logged Replication Data](#) in the SAP LT Replication Server Cockpit. To view the documentation for this activity, select it and choose the [Documentation](#) pushbutton.



7 Frequently Asked Questions and Troubleshooting Recommendations

7.1 If SAP LT Replication Server for SAP HANA is already in use, what are the considerations when applying a new DMIS version and/or support package?

If the SAP LT Replication Server is installed on a separate system, we recommend that the same support package levels is installed on this system and on the source system.

As it might not always be possible to install the new support package on all involved systems at the same time, several combinations are supported. If not all systems are installed on the newest support package, some new functionalities might not be available.

If you upgrade the SAP LT Replication Server system from DMIS 2010 to DMIS 2011 ensure that you install DMIS 2011 and all relevant support packages in one installation queue. This is important as some additional table fields (for example partitioning command) were introduced on a certain DMIS 2010 SP level and will get lost if DMIS 2011 basis package and support packages are installed sequentially.

If you upgrade from DMIS 2010 to DMIS 2011 you have to upgrade to at least the corresponding support package level. The corresponding level for DMIS 2010 SP9 is DMIS 2011 SP4

More details are provided in SAP Note [1824710 - Installation/Upgrade SLT - DMIS 2011 SP4 / 2010 SP9](#).

7.2 Does SAP LT Replication Server for SAP HANA only support 1:1 data replication or is it possible to do filtering (selective data replication) and / or transform data during the replication?

The SAP LT Replication Server includes various transformation capabilities. However, by default the replication process is a 1:1 table replication. Several adjustments of data from source systems to the SAP HANA database format such as conversion to UNICODE format happen automatically during the replication process.

You can use Advanced Replication Settings (transaction LTRS, tab page [Rule Assignment](#)) to define transformation rules for filtering of data or for more advanced transformation requirements such as the scrambling of data. Transformation rules have to be implemented for respective tables on the SAP LT Replication Server before you start with the data replication.

Note: Since applying transformation rules may lead to unexpected data inconsistencies between source and target (SAP HANA) system or may have a negative impact on the performance of the data replication process, we highly recommend involving an SAP LT Replication Server expert in this matter.

7.3 Are there any special considerations if the source system is a non-ABAP system?

The fundamental concept of the trigger-based replication is designed as for ABAP source systems. The main differences are that the connection is a database connection (no RFC connection) and the read modules reside on the SAP LT Replication Server.

Furthermore, consider the following:

- Due to the database trigger concept, tables in non-ABAP source systems must have a primary key to be considered for replication.
- Tables DD02L, DD02T (for ABAP source systems used to replicate the SAP dictionary information) include the metadata of the tables in the non-ABAP system – however these tables are just initially loaded and not automatically updated (replicated).
- Tables that have database-specific formats may need to apply dedicated mappings (transformation rules) before they can be properly replicated.
- Only SAP supported databases (with respective DBSL for SAP NetWeaver 7.02) are supported as non-SAP ABAP systems.

7.4 Is it possible to use a configuration for multiple source and target SAP HANA systems?

Yes, the SAP LT Replication Server supports both (N:1 replication and 1:N replication – limited to 1:4)

For N:1 replication: consider specific preparation steps especially if you want to replicate data from the same tables.

For 1:N replication: consider that you flag the related configuration during the initial creation already accordingly (transaction: LTR)

In any case, we recommend involving an SLT expert for the initial set-up of those scenarios.

7.5 What are the potential issues if the creation of a configuration and a related schema fails?

Missing add-on DMIS_2010 in your source system: Check and make sure that in your source system you have installed the required add-on DMIS_2010. See SAP Note 1468391 for further information.

Missing the proper roles for the RFC user: In the system log (SM21) and ABAP dump (ST22), both indicated that RFC_NO_AUTHORIZATION error occurred. Logon to the source system and check if role SAP_IUUC_REPL_REMOTE is assigned to the RFC user.

Roles are not generated properly: If the roles are assigned correctly to the RFC user, but still get the error you might not have generated the roles properly. The role should have a green traffic light on the tab pages

[Authorization](#) and [User](#). If there is a yellow light on these tabs, generate the roles and execute the user comparison (as outlined in chapter 3).

DDIC user is used for RFC communication: You cannot use user DDIC for RFC connections between the source system and SLT, as it is in conflict with coding in the SAP core function.

HANA DBSL is not installed or database client software path not specified in the OS environment: check the work process log and refer to SAP Note 1597627.

DB client software is outdated. Check the work process log and update the database client software. See SAP Note 1603671 for further information.

Logon credentials for SAP HANA system are not correct: Check the logon information when you create a new schema.

7.6 How can I ensure that data is consistent in the source system and SAP HANA system?

Since any change in the source system is tracked in dedicated logging tables, the replication status for each changed data record is transparent. An entry of a logging table is deleted after a successful “commit” statement from the SAP HANA database, this procedure guarantees data consistency between source and target (SAP HANA) system – even in the case of system outages or network failures. Advanced monitoring and further expert functions allow you to track the replication progress of each data portion in detail – however, dedicated reconciliation lists that allow reviewing the replication status from a business perspective are currently not available.

7.7 What happens in case of network failures?

As long as there was no successful “commit” statement from the SAP HANA database, respective information in the logging tables stay in place – therefore the replication of related changes will be repeated until the replication is successfully completed.

7.8 Is the partitioning of SAP HANA tables supported by SAP LT Replication Server?

There may be several reasons why partitioning of tables replicated into SAP HANA is needed - for example if there are more than 2 billion records in a source table, it needs to be partitioned into multiple SAP HANA tables.

You can use Advanced Replication Settings (transaction LTRS, tab page [Table Settings](#)) to define partitioning commands. You have to enter the SQL statements using the same syntax as in the SQL editor of the SAP HANA Studio (see also in the guide [SAP HANA Database – Partitioning and Distribution of large tables](#)).

Note: You have to complete the partitioning definition before you set the related table(s) into status Load or Replication from the SAP HANA Studio.

7.9 Does SAP LT Replication Server for SAP HANA support data compression like the SAP HANA database?

Yes, this is automatically covered by the RFC connection used for the data replication from the ABAP source systems.

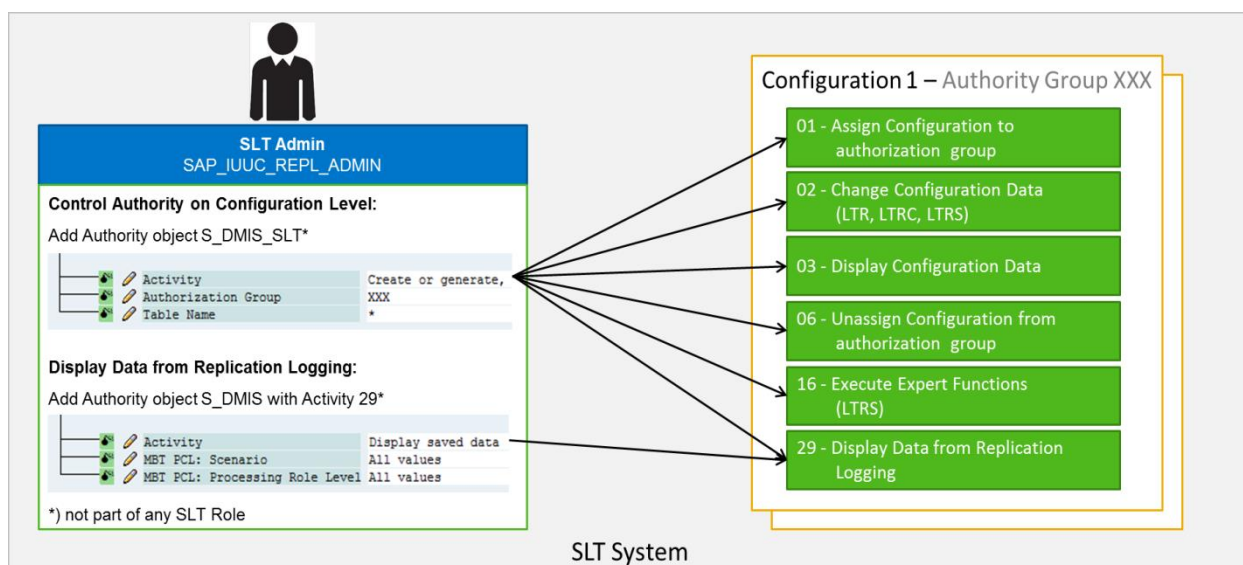
7.10 How to use Authorization Groups to control access on configuration level by LT Replication Server administration tasks?

You can view the authorization group (if one exists), or change it if required by using the [Change Authorization Group](#) pushbutton.

If you require specific authorizations for the configuration, you can specify an authorization group, and use it in the corresponding authorization object S_DMIS_SLT.

Authorization Object Definition

The authority object controls the authorizations for different activities for configurations and mass transfers.



Defined fields

MBT_AUTHGR: specifies the authorization group assigned to the configuration or mass transfer ID

ACTVT: The activity specifies the possible actions that can be performed:

- 01 Create or Generate: A new configuration or mass transfer can be created or assigned to this authorization group
- 02 Change: Settings for configurations or mass transfers assigned to this authorization group can be changed
- 03 Display: UI information for configurations or mass transfers assigned to this authorization group can be displayed
- 06 Delete: Configurations or mass transfers assigned to this authorization group can be deleted
- 16 Execute: Functions for configurations or mass transfers assigned to this authorization group can be executed

7.11 How to define an alternative connection for the initial load?

You can define a deviating connection for initial load for a dedicated mass transfer ID with transaction LTRC. Navigate to the tab page [Expert functions](#) in section [Additional Functions](#) and choose [Define Connection to Alternative System for Initial Load](#).

You can define an alternative sender system that is used for the initial load. To do this, you specify an RFC destination or a database connection to the alternative system. You can also delete the connection information for an alternative sender system.

After you have defined an alternative sender system, the system connects to the system as a data source during the initial load whenever you:

- Add a new table to be replicated for the configuration (schema) in the SAP HANA Studio.
- Stop the replication of a table for this configuration and restart it in the SAP HANA Studio.
- Add a new table to be loaded only for the configuration in the SAP HANA Studio.
- Stop the replication of a table for the configuration, and add it to be loaded only (in the SAP HANA Studio).

The system uses an RFC destination or database connection to the alternative sender system during the Initial Load phase. For a replication, as soon as the system switches to delta mode, the system also switches the sender destination and connects to the sender system specified in the configuration.

If the replication sender system is connected by means of an RFC destination, you can choose whether you want to connect the alternative sender system by means of an RFC destination or a database connection. If the replication sender system is connected by means of a database connection, you must also use a database connection to the alternative sender system.

Prerequisites

You must ensure that the sender system and the alternative initial load sender system are synchronized before the replication of a table is started. If the table content is not identical in both systems, then incorrect data will be transferred to the target system during the initial load.

7.12 How to change the settings of the DB connection to the target system?

You can change the connection settings to a target system for a dedicated mass transfer ID with transaction LTRC. Navigate to the tab page *Expert functions* in section *Additional Functions* and choose *Change Settings for Connection to Target System*.

You can change the hostname and instance number, and set the password to a new one (by choosing the *Change Password* pushbutton).

There are two connections, one for setting up the configuration, and the other for handling the actual data replication.

Note for the Connection for the Data Replication:

If you have a multi-node SAP HANA system, enter all nodes in the *Connection for Data Replication* section; this ensures that if one node is not available, the database connection will use the next node.

With DMIS 2011 SP8, you can create or connect to a SAP HANA Database schema with a case sensitive schema name and/or a schema name using lower case letters.

In addition you can optionally connect to an existing SAP HANA database schema with a different user than the schema owner. In this case you need to specify the user name and password during the process of creating an SAP LT connection to an already existing SAP HANA database schema.

7.13 The Notification Console

Purpose

It is important that the system administrator responsible for SAP LT Replication Server is notified of any issues regarding the replication process as soon as possible. To avoid permanent monitoring of the SAP LT Replication Server system, you can activate notifications for a configuration. If notifications are active for a configuration, the system uses a background job to check the configuration periodically (the notification job).

Integration

You can configure the notification settings in the SAP LT Replication Server Cockpit (transaction LTRC) or by using transaction IUUC_REPL_EN.

Prerequisites

SAPconnect must be configured for e-mail output. You can check this in transaction SCOT.

The UI comprises the following tab pages:

General Settings

If you want the system to check a configuration (mass transfer ID) periodically, you must activate notifications. To do this, select the *Notifications Active* checkbox, and choose *Save*.

You can change the frequency of the checks by specifying a value (in seconds) in the *Interval for Checks (Sec)* field, and change the frequency of the notification e-mail by specifying a value (in minutes) in the *Interval for Notifications (Mins)* field.

For each configuration, you can also activate and deactivate the individual checks made by the notification job.

The system can check the following for the SAP LT Replication Server system:

- Whether the connections to the source system and target system are working correctly
- Whether the master job is running
- Whether the data load jobs are running
- Whether acceptable latency times have been exceeded
- Whether any errors occurred during the load or replication

For the source system, the system can check the status of the triggers in the source system, and whether the logging tables have been created.

The system can also check whether the tables have been created in the target system.

Latency Thresholds

Latency is defined as the length of time it takes until a data change to a source table has been committed to the target system database. For each table, you can define individual thresholds in order to trigger error notifications and warning notifications. If the latency time for a table exceeds this threshold, the system sends a warning or error notification to the relevant recipients.

For example, you can trigger a warning e-mail if the latency time has exceeded 10 seconds, and an error e-mail if the latency time has exceeded 50 seconds. These settings are optional; if you do not make any settings here, the system uses default system settings for the tables.

E-Mail Recipients

In this tab page, you can specify e-mail addresses for the recipients of the notifications. You can specify multiple recipients for a configuration.

For each recipient, you specify the type of notification that they receive. There are three categories, Information Notifications, Warning Notifications, and Error Notifications.

Output

If the system detects any issues, an e-mail is sent to all registered users.

8 Appendix

8.1 SAP Landscape Transformation Replication Server on SAP Community Network (SCN)

The following location contains information such as latest available guides, how-to documents, as well as blogs and discussion threads about the SAP LT Replication Server:

<http://scn.sap.com/community/replication-server>

8.2 Related Guides

The following table contains useful links to related guides:

Content	Guide
Installation Guide - Trigger-based Data Replication Using SAP LT Replication Server	http://help.sap.com/sapslt
Security Guide – Trigger-based Data Replication Using SAP LT Replication Server	http://help.sap.com/sapslt
SAP HANA Administration Guide	http://help.sap.com/hana → Administration
Sizing Guide for SAP Landscape Transformation (SLT) Replication Server for SAP HANA, SPS 05	https://service.sap.com/sizing → Sizing Guidelines → Database and Technology → SAP In-Memory Computing → SAP Landscape Transformation Replication Server, SAP HANA

8.3 Related SAP Notes

The following table contains links to information related SAP Notes:

SAP Note Number	Title	Comment
1605140	Central Note - SAP LT Replication Server	Collective Note for all the relevant Notes for LT Replication Server for SAP HANA.

SAP Note Number	Title	Comment
2570052	Installation/Upgrade SLT – DMIS 2011 SP14	This SAP Note describes the installation or upgrade of SAP Landscape Transformation Replication Server to the relevant DMIS SP.
1768805	SAP LT Replication Server: Collective Note - non-SAP Sources	Collective Note that describes details about using the SAP LT Replication Server for non-ABAP source systems.
1691975	HANA LTR – Clarification on DMIS releases	Decision criteria for choosing the correct software version.
1963522	Limitation for SAP LT Replication Server on MaxDB	Due to issues with MaxDB database triggers, using SAP LT Replication Server to replicate data from source systems based on some versions of Max DB could lead to data inconsistencies in the target system. Therefore only MaxDB versions specified in this SAP Note can be used for SAP LT Replication Server scenarios.

Relevant Notes from related software components:

SAP Note Number	Title	Comment
1875338	SLT: Table structure changes are not replicated	Structural Changes are not handled automatically : Your SLT System is running on SAP_BASIS 710/711 or 720 and therefore structural changes are not done by SLT automatically
1860891	Transfer of new table fields from the sender system to SAP HANA fails	Structural Changes are not handled automatically: After applying note 1749824 structural changes are not done by SLT automatically
1762397	GEN_ALTER_TABLE cannot be used for secondary database	Structural changes are not done by SLT automatically
1875769	Check for DB trigger made less strict (not null)	For more details see section 5.1 - 'What happens if there are structural changes on tables with active DB trigger?'

SAP Note Number	Title	Comment
1749824	Treatment of SLT triggers during transport	For more details see section 5.1 - 'What happens if there are structural changes on tables with active DB trigger?'
1655725	Check for DB trigger made less strict	For more details see section 5.1 - 'What happens if there are structural changes on tables with active DB trigger?'

Material Number

© 2018 SAP AG. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.

Microsoft, Windows, Excel, Outlook, and PowerPoint are registered trademarks of Microsoft Corporation.

IBM, DB2, DB2 Universal Database, System ads, System i5, System p, System p5, System x, System z, System z10, System z9, z10, z9, iSeries, pSeries, xSeries, zSeries, eServer, z/VM, z/OS, i5/OS, S/390, OS/390, OS/400, AS/400, S/390 Parallel Enterprise Server, PowerVM, Power Architecture, POWER6+, POWER6, POWER5+, POWER5, POWER, OpenPower, PowerPC, BatchPipes, BladeCenter, System Storage, GPFS, HACMP, RETAIN, DB2 Connect, RACF, Redbooks, OS/2, Parallel Sysplex, MVS/ESA, AIX, Intelligent Miner, WebSphere, Netfinity, Tivoli and Informix are trademarks or registered trademarks of IBM Corporation.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Adobe, the Adobe logo, Acrobat, PostScript, and Reader are either trademarks or registered trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Oracle is a registered trademark of Oracle Corporation.

UNIX, X/Open, OSF/1, and Motif are registered trademarks of the Open Group.

Citrix, ICA, Program Neighborhood, MetaFrame, WinFrame, VideoFrame, and MultiWin are trademarks or registered trademarks of Citrix Systems, Inc.

HTML, XML, XHTML and W3C are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Massachusetts Institute of Technology.

Java is a registered trademark of Sun Microsystems, Inc.

JavaScript is a registered trademark of Sun Microsystems, Inc., used under license for technology invented and implemented by Netscape.

SAP, R/3, xApps, xApp, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP Business ByDesign, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world. All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

These materials are subject to change without notice. These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.