

Application Operations Guide



Trigger-Based Data Replication Using SAP LT (Landscape Transformation) Replication Server for SAP HANA

SAP HANA Appliance Software SPS 04

Document Version 1.0 –April 2012



SAP AG

Neurottstraße 16
69190 Walldorf
Germany
T +49/18 05/34 34 24
F +49/18 05/34 34 20
www.sap.com

© Copyright 2012 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, Outlook, and PowerPoint are registered trademarks of Microsoft Corporation.

IBM, DB2, DB2 Universal Database, OS/2, Parallel Sysplex, MVS/ESA, AIX, S/390, AS/400, OS/390, OS/400, iSeries, pSeries, xSeries, zSeries, z/OS, AFP, Intelligent Miner, WebSphere, Netfinity, Tivoli, and Informix are trademarks or registered trademarks of IBM Corporation 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.

MaxDB is a trademark of MySQL AB, Sweden.

SAP, R/3, mySAP, mySAP.com, xApps, xApp, SAP NetWeaver, 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.

Disclaimer

Some components of this product are based on Java™. Any code change in these components may cause unpredictable and severe malfunctions and is therefore expressly prohibited, as is any decompilation of these components.

Any Java™ Source Code delivered with this product is only to be used by SAP's Support Services and may not be modified or altered in any way.






Documentation in the SAP Service Marketplace

You can find this documentation at the following Internet address:
service.sap.com/instguides

Typographic Conventions

Type Style	Represents
Example Text	Words or characters that appear on the screen. These include field names, screen titles, pushbuttons as well as menu names, paths and options.
Example text	Cross-references to other documentation Emphasized words or phrases in body text, titles of graphics and tables
EXAMPLE TEXT	Names of elements in the system. These include report names, program names, transaction codes, table names, and individual key words of a programming language, when surrounded by body text, for example, SELECT and INCLUDE.
Example text	Screen output. This includes file and directory names and their paths, messages, names of variables and parameters, source code as well as names of installation, upgrade and database tools.
Example text	Exact user entry. These are words or characters that you enter in the system exactly as they appear in the documentation.
<Example text>	Variable user entry. Pointed brackets indicate that you replace these words and characters with appropriate entries.
EXAMPLE TEXT	Keys on the keyboard, for example, function keys (such as F2) or the ENTER key.

Icons

Icon	Meaning
	Caution
	Example
	Note
	Recommendation
	Syntax

Contents

1	Getting Started.....	6
1.1	Important SAP Notes	6
1.2	History of Changes	7
2	Technical System Landscape.....	8
2.1	Landscape Strategy	8
2.1.1	SLT System Setup for SAP Sources.....	9
2.1.2	SLT System Setup for Non-SAP Sources.....	10
2.2	System Requirements	10
3	Configuration of SAP LT Replication Server	11
3.1	System Connections.....	11
3.1.1	User Creation and Connection for an SAP Source System..	11
3.1.2	User Creation and Connection for a non-SAP Source System	13
3.1.3	User Creation and Connection for an SAP HANA System....	15
3.2	Specifying Replication Configuration	16
3.3	Data Provisioning.....	19
4	Operations and Management of SAP LT Replication Server	21
4.1	Important Tables and Table Structures of Trigger-based Replication Using SAP LT Replication Server	21
4.1.1	Important Replication-Relevant Tables	21
4.1.2	Table Structure in SAP HANA Modeler.....	21
4.2	Administration of Load and Replication Process	22
4.2.1	Mass Transfer ID and Schema GUID	22
4.2.2	Important SAP LT Replication Server Jobs	22
4.2.3	Job Handling and Performance Improvements for Initial Load of Tables	23
4.2.3.1	Restarting jobs.....	23
4.2.3.2	Stopping Jobs	23
4.2.3.3	Changing the Number of Jobs	24
4.2.3.4	Performance Improvements for Initial Load of Tables	24
4.3	Monitoring of Load and Replication Process.....	25
4.3.1	Status of Jobs and Connections.....	25
4.3.2	Status of Triggers.....	26
4.3.3	Statistics	26
4.4	Backup and Recovery Aspects.....	27
5	Special Considerations for Source Systems	28
5.1	Impact of Software Maintenance Events in SAP Source Systems	28

5.2	Archiving Data in Source Systems.....	29
5.3	Performance Implications in Source System using Trigger-based Data Replication	29
5.4	Data Volume Management	29
6	Further Tips and Tricks, Frequently Asked Questions, and Troubleshooting Recommendations	30
7	Appendix	32
7.1	Related Guides	32
7.2	Related Information	32

1 Getting Started



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

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
- Support Specialist

1.1 Important SAP Notes



Check regularly for updates available for the Application Operations Guide.

SAP Note Number	Title	Comment
19466	Downloading SAP Kernel patches	Downloading a kernel patch in the Service Marketplace, Software Distribution Center.
517484	Inactive services in the Internet Communication Framework	The Internet Communication Framework Services are inactive when you install the SAP Web Application Server.
1468391	Installation and delta upgrade of DMIS 2010_1	The SAP Landscape Transformation (SLT) component part of DMIS.
1597627	HANA Connection	Activating a secondary connection to the SAP HANA In-Memory Database
1603660	Individual release 7.20 kernel on SAP MaxDB for HANA LT	Using 7.20 EXT kernel with SAP MaxDB
1605140	SAP HANA: Central Note - SAP LT Replication Server	Collective Note for all the relevant Notes for LT Replication Server for HANA
1646371	HANA replication fails for sender systems with release 4.6C	For 4.6C source systems
1649910	Applying SLT (DMIS_2010 SP05) with	This note describes the recommended order for upgrading

	SAP HANA SPS03	the landscape to SPS03 and DMIS SP05
1655246	HANA LTR SP5: General Corrections	Fixes after applying DMIS_2010 SP05
1656370	HANA LTR SP5: General Corrections 2	Fixes after applying DMIS_2010 SP05
1662438	HANA LTR SP5: General Corrections 3	Fixes after applying DMIS_2010 SP05
1664883	HANA LTR SP5: General Corrections 4	Fixes after applying DMIS_2010 SP05
1663138	HANA data replication fails due to incorrect DB triggers	Problems with DB trigger definitions fixed

1.2 History of Changes



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*.

The following table provides an overview of the most important changes in prior versions.

Version	Important Changes
V0.5	Initial draft
V0.8	Initial completion of guide
V1.0	First published version, including revisions from Info Development

2 Technical System Landscape

2.1 Landscape Strategy

The SAP LT Replication Server is a replication technology to provide data from SAP source systems and non-SAP source systems in a SAP HANA environment. It acts as a key enabler for SAP HANA customers to supply their HANA environment with relevant data.

The following components are used in the technical system landscape:

SAP source system(s)

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

Non-SAP source system(s)

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

SLT system

If the source is an SAP system, the SLT system polls the logging tables in the source system with a remote function call (RFC) connection. If the source is a non-SAP system, the SLT system polls the logging tables in the non-SAP source system through a database connection.

SAP HANA system

The SAP HANA system contains the SAP HANA database. It is used to store the replicated data. The SLT system and the SAP HANA system communicate via database connection.

The SAP LT Replication Server can be used for data replication from SAP sources and non-SAP sources to the SAP HANA system.

The relevant information required to create the connection between the source system(s), the SLT system, and SAP HANA system(s) is specified in the SLT system as a configuration schema. In the *Configuration & Monitoring Dashboard* (transaction *LTR*), you can define a new configuration.

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

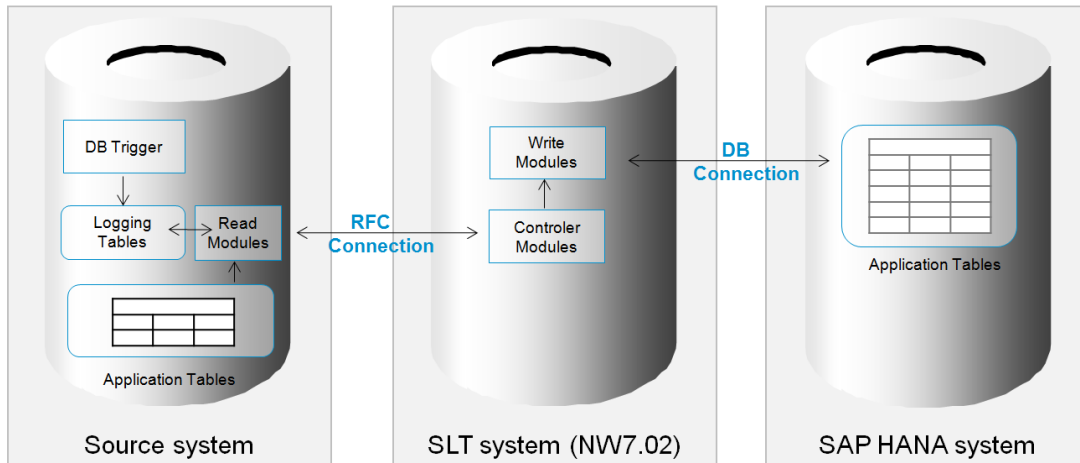
- For SAP sources, the SAP LT Replication Server can either be installed as part of the source system (if applicable to the customer's landscape strategy and technical prerequisites are met) or as a separate SAP system (recommended for productive use). However from a technical perspective, the SLT Replication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

In the following subsections, we will examine each option along with their technical prerequisites.

The SAP LT Replication Server uses background processing for data replication. This can be an important factor in deciding where to install SAP LT Replication Server since background processing affects performance. In addition, the flexibility to apply dedicated software changes without any impact on running applications might be another reason to keep SAP LT Replication Server as separate system. However, any additional SAP system also may require related system administration effort.

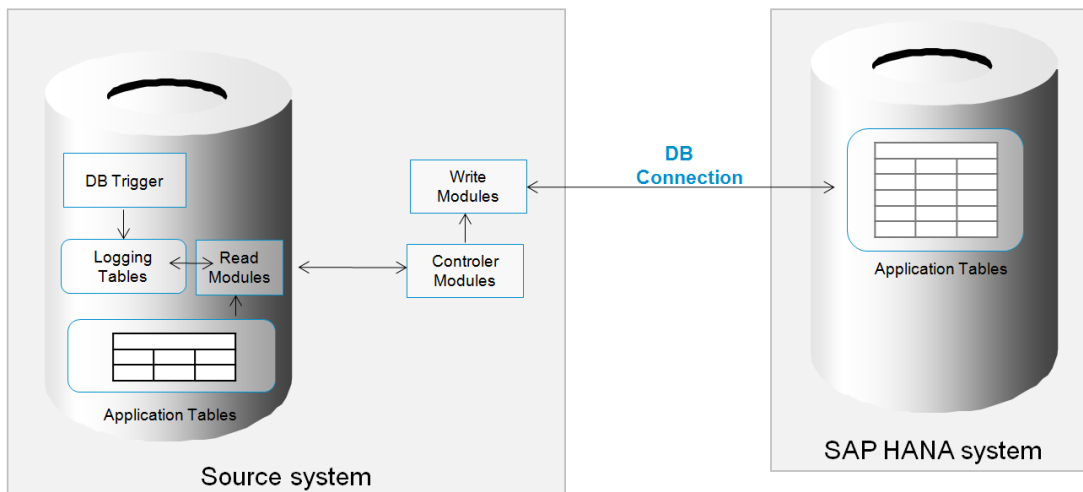
2.1.1 SLT System Setup for SAP Sources

Option 1 - Separate SLT system



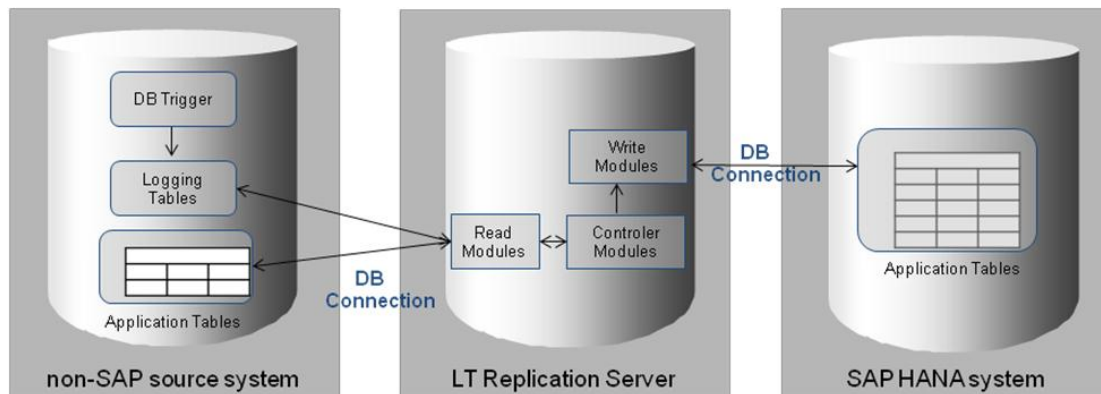
The diagram above shows the SLT component installed in a separate system. This three-tier approach is needed when the source system is not compliant with the required technical prerequisites of SLT. For data replication from SAP sources, we recommended that you keep the production SLT instance on a separate SAP system.

Option 2 - SLT installation in source system



The diagram above shows the SLT component installed in the SAP source system. If the source system supports the required SAP kernel and SAP NetWeaver versions, the SLT installation on top of the SAP source system leads to a simplified two-tier system landscape as shown. However, there is less flexibility and potential dependencies for software maintenance activities such as changing the SAP kernel or applying support packages to the SAP source system.

2.1.2 SLT System Setup for Non-SAP Sources



The diagram above shows SLT installed on a separate SAP system. Compared to a setup with an SAP source system, the read modules are created in the SLT system (for SAP sources the read modules are created in the source system). The connection from the SAP LT replication server to the non-SAP source system is established through a database connection.

Note: For database connections to non-SAP systems, the OS/DB restrictions of SAP NetWeaver 7.02 apply (see <http://service.sap.com/pam>). Furthermore, the customer database license has to cover a permanent database connection with third party products such as the SAP LT Replication Server.

Note the following when setting up connections between source system(s), an SLT system, and SAP HANA system(s):

- In the SLT system, each source system can be connected to only one SAP HANA system / database (1:1 relation)
- Each SLT system can connect multiple sources to one (or multiple) SAP HANA systems / databases - but only one source to one dedicated SAP HANA system is possible (multiple 1:1 relation)
- Multiple source systems can be connected to one SAP HANA system using the same database schema (N:1 relation)
- Replication of non-SAP sources requires that SLT is installed on a separate SAP system
- The SLT system must be a UNICODE system.

2.2 System Requirements

- SAP Basis: SAP NetWeaver 702, with Kernel 7.20EXT (64BIT Unicode)
- Filesystem: 100 GB
- RAM: 16-32 GB
- CPU: 2-4 cores
- Recommended to have 1 background work process for every 10-20 tables that have the status *In Replication*

The number of background jobs should be increased based on the following factors:

The number of tables to be replicated

The expected performance of the initial load (load time) as well as on the change rate of the tables that take part in the replication

In general, one job should be considered per 10-20 tables.

3 Configuration of SAP LT Replication Server

This section provides an overview of the transactions and tools you use to perform the basic configuration related to the trigger-based data replication using SAP LT Replication Server.

Initial Administrative Tasks

The following sections describe administrative tasks that should be performed initially after the SLT system is installed and should be checked in case of errors.

3.1 System Connections

3.1.1 User Creation and Connection for an SAP Source System

In order for the SLT system to operate, you must create an RFC connection to the source system(s). Also, a user is required in the SAP HANA database so that replication data can be stored.

Prerequisites

Refer to the SAP user administration guide for RFC user creation. Refer to the SAP HANA documentation for user creation information for the SAP HANA database.

Procedure

1. Create a user (of type *Dialog* or *System*) using transaction SU01 in your source system(s), generate and assign the following role to this user:
 - SAP_IUUC_REPL_REMOTE (available only if SP05 is applied in the SAP source system)

Note: Do not use user DDIC. The role SAP_IUUC_REPL_REMOTE is not generated by default. Generate and assign this role to the newly created user.
2. Create an RFC connection (type 3 – ABAP) from the SLT system to the source system with the user created above (if both systems are Unicode, specify this RFC connection as Unicode).

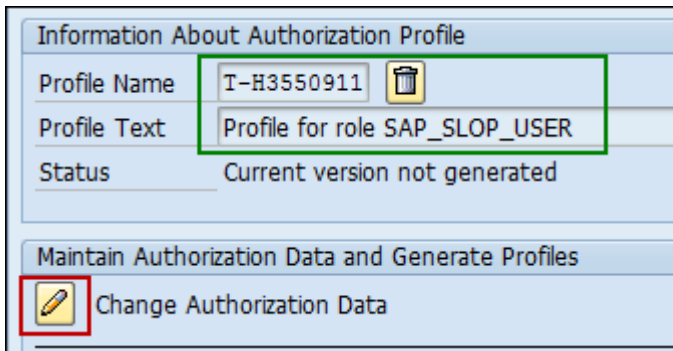
Note: Do not use DDIC for the RFC connection, as this will result in errors. If the source system and the SAP LT replication server are the same system, create an RFC connection and do not use the RFC connection *NONE*.
3. Create a user in the SAP HANA database. This user is necessary to setup the connection from the SLT system to the SAP HANA system during the schema creation (see chapter 5).

Details - Generation of user roles:

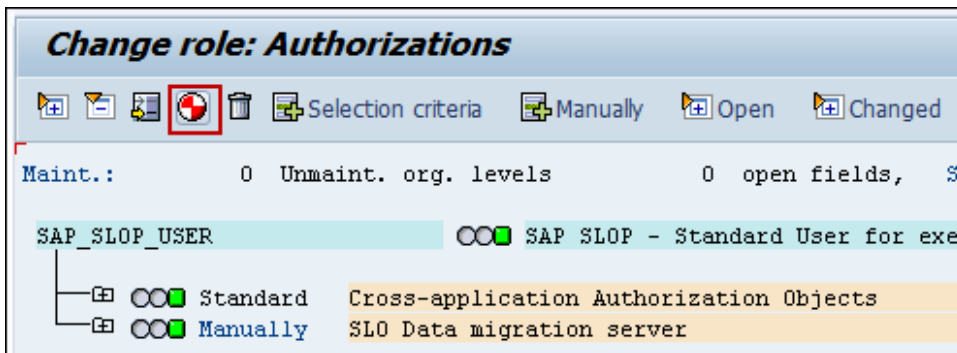
To generate a role (transaction PFCG) proceed as follows:

Note: In the following screenshots, the role SAP_SLOP_USER is used as an example. With SPS03 (DMIS_2010 SP05), the role SAP_IUUC_REPL_REMOTE is recommended to use.

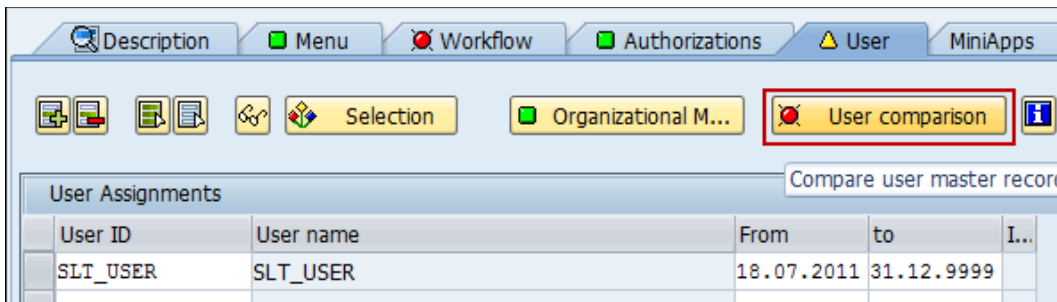
- a) Switch to change mode.
- b) Generate the profile for this role in the *Authorization* tab page.
- c) After the profile is generated, change authorization data as follows:



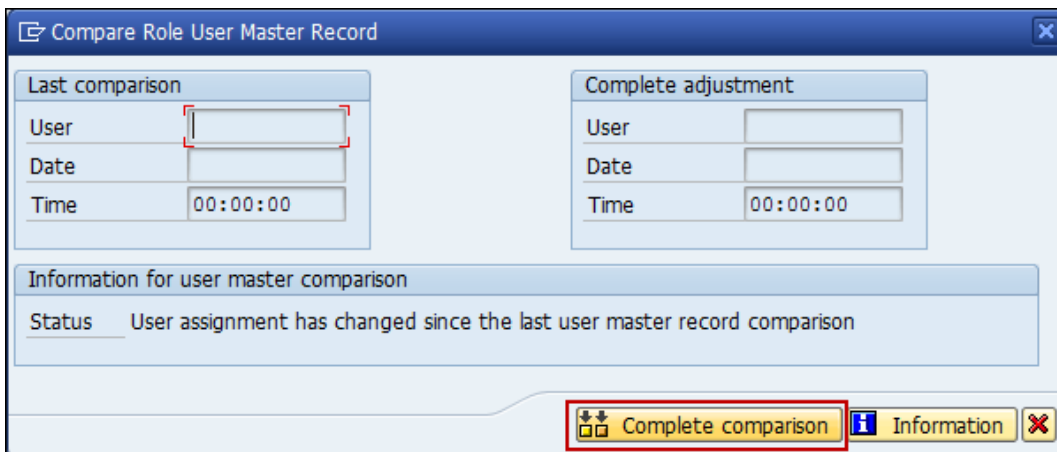
d) Click the *Generate* pushbutton:



e) Press F3 to go back, and you should see green light on the *Authorization* tab page. Switch to the *User* tab page, and press the pushbutton *User comparison*:



f) In the pop-up window, choose *Complete comparison*:



g) Now both *Authorization* and *User* tab pages have a green light, choose the *Save* pushbutton.

Note: Roles will be available after completion of the DMIS add-on and SP installation. If you use a new client after the DMIS add-on is applied, you must transport the necessary roles from client 000 into your target client.

For more details about the roles and authorization concept of SAP LT replication server, see the Security Guide for SAP LT Replication Server for SAP HANA.

3.1.2 User Creation and Connection for a non-SAP Source System

To establish a secondary database connection from an SAP system to an external database, the connection data and the user data of a user are required. This database user must be authorized to establish a connection to the external database. The SAP system connects to a specific schema from the database. To perform the replication and initially load a specific table from a given schema, the database user must have authorizations for the following actions:

- Selecting from the tables
- Creating tables in the given schema (for creating the logging table)
- Selecting from the logging tables
- Deleting the logging tables
- Creating database triggers for tables
- Deleting the triggers
- Creating synonyms for specific tables
- Deleting the synonyms

Depending on the specific external database system, the process of granting authorizations to a database user can vary.

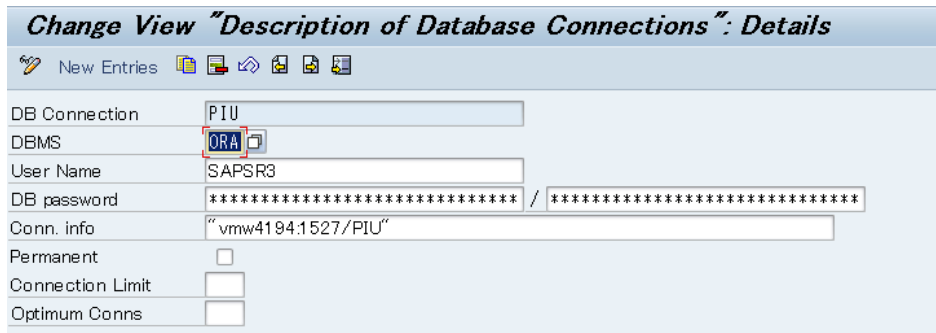
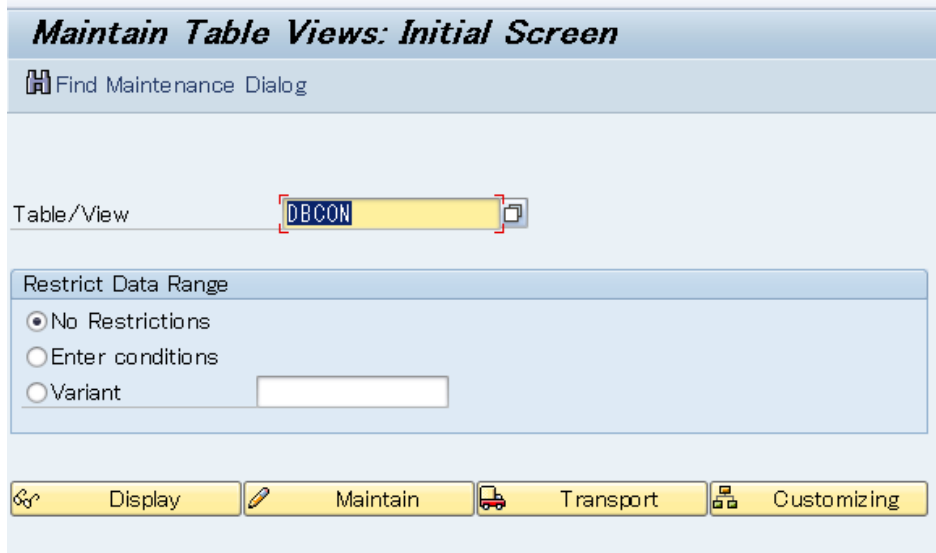
Details about non-SAP source connections:

Similar to SAP source system connections, you need:

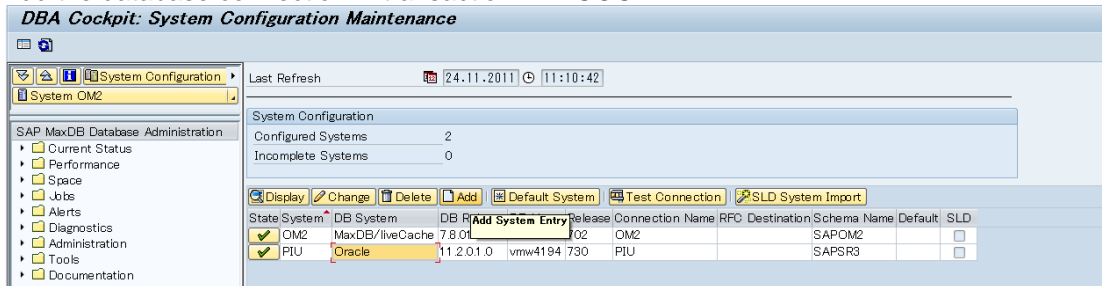
- A user in the source system with appropriate authorizations (see details in SAP LT Replication Server – Security Guide)
- A connection (in this case a database connection) to the source system. If you are replicating from non-SAP source systems, the user you specify needs authorizations as described in SAP LT Replication Server – Security Guide.

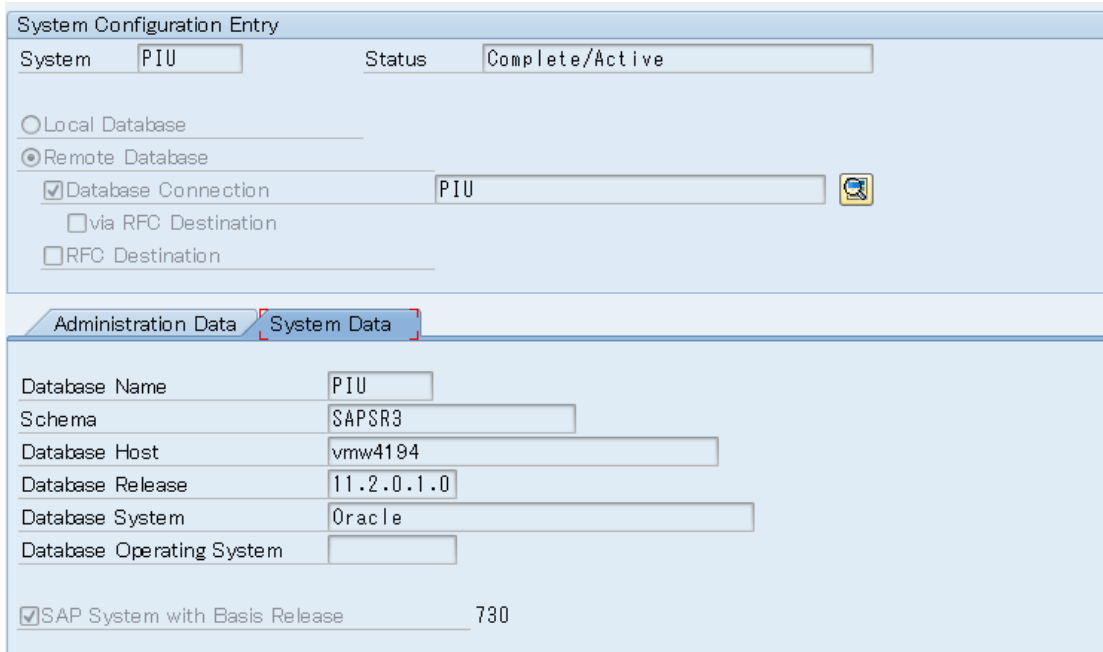
The following procedure outlines how to establish a database connection for an Oracle database as a non-SAP source system:

1. Install the Oracle instant client on the SLT system (if your SLT system is not based on Oracle).
2. Install the DBSL database dependent library for the SAP 7.20 EXT Kernel.
3. Create the database connection in table DBCON (via transaction SM30)

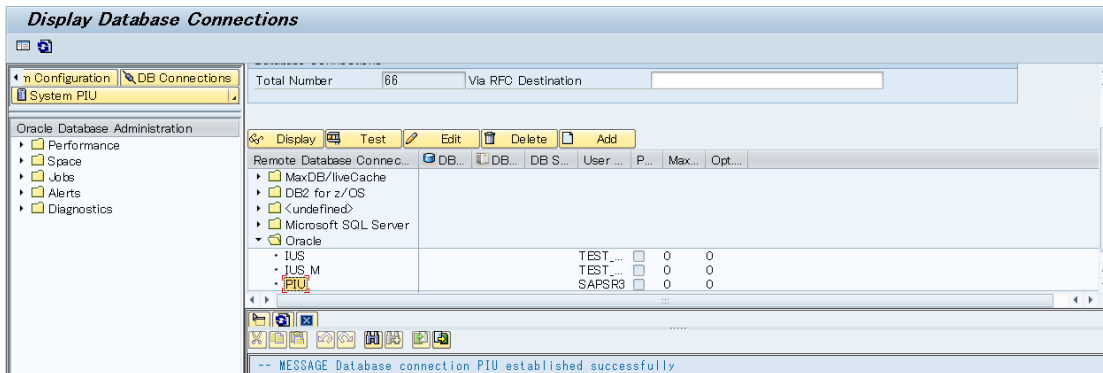


4. Add the database connection in transaction DBACOCKPIT:





5. Test the database connection:



You find more details about the roles and authorization concept of LT replication server in the Security Guide for SAP LT Replication Server for SAP HANA.

3.1.3 User Creation and Connection for an SAP HANA System

You can use a database user with the same authorizations as the user *SYSTEM* to establish a connection between the SLT system and the SAP HANA system.

While creating a configuration schema, a replication user (with name equal to schema name) is automatically generated for replication purposes only (with an unknown and hidden password). In addition, the following schema-specific roles are also generated. You explicitly need to assign one of the roles to HANA studio user(s) to grant access to content of the dedicated DB schema:

Name	Role	Purpose
Data Provisioning	<SCHEMA>_DATA_PROV	Data provisioning - deploy and control replication orders
Power User	<SCHEMA>_POWER_USER	Power user - allow modification within the replication schema

Specifying Replication Configuration

User Admin	<SCHEMA>_USER_ADMIN	Can call the RS_GRANT_* and RS_REVOKE_* procedures to assign specific SQL privileges of the respective schema
------------	---------------------	---

DATA PROV

- This role has to be assigned to users which have to configure and monitor the data provisioning process. This role has the right to SELECT data in the replication schema, and to INSERT values into the RS_ORDER table within the replication schema.

POWER USER

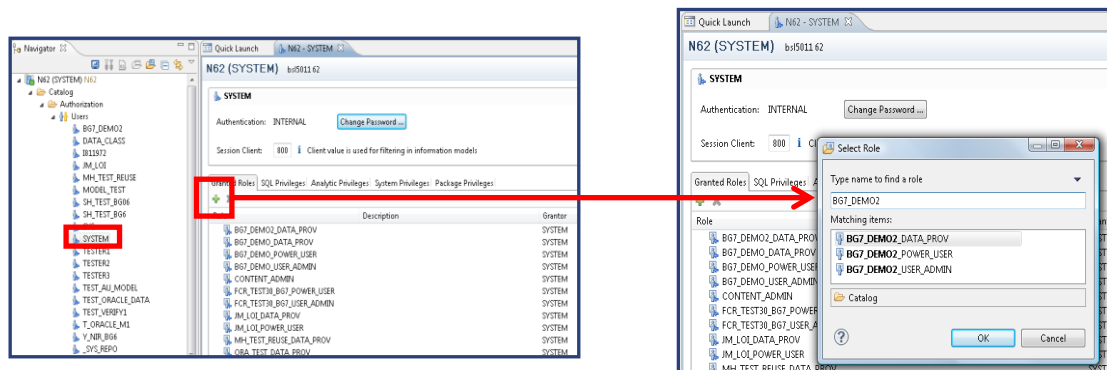
- This role provides full control on the contents of the replication schema.
- Note:** Only assign users to this role only for urgent maintenance operations, or similar operations. The rights granted by this role allow operations which might destroy the consistency of the replicated data.

USER ADMIN

- This role provides access to the database stored procedures RS_GRANT_ACCESS and RS_REVOKE_ACCESS which are used for fine-grained access control on the replication schema content.

Procedure

Select a user in the SAP HANA Studio and assign the relevant roles by choosing the green plus symbol:



After adding the relevant roles, the user has to be saved.

For more information about the roles and authorization concept of SAP LT replication server, see the Security Guide for SAP LT Replication Server for SAP HANA.

Note: in case you change the password of the generated replication user of a DB schema in the SAP HANA studio, you also need to adjust the password in the respective DB connection (transaction DBCO) for this replication user.

3.2 Specifying Replication Configuration

In the SLT system, you define the connection(s) between source system(s), the SLT system and the SAP HANA system(s) using transaction LTR. A dedicated connection is saved as a configuration.

Procedure

1. Choose *New* to create a new configuration.

SLT Replication Monitor

Available Configurations

New Delete Check Status

Schema Name	Target Parameters	Status
JM_LOI	bs5011.62	○○○
TEST_BGZ	bs5011.62	●○○

Configurations
 Schema Name: JM_LOI
 Target Parameter: bs5011.62

Connection to Source System
 RFC Connection: LOI

Connection to HANA System
 Host Name: bs5011
 Instance Number: 62

Last Refresh: 28.09.2011 15:44:15 Refresh

2. Specify the configuration

Create a new Configuration

General Data

Schema Name: *

Number of Replay Jobs: *

Connection to Source System

SAP Legacy

RFC Destination: *

Connection to HANA System

User Name: *

Password: *

Host Name: *

Instance Number: *

Tablespace Assignment

Data Class 1:

Data Class 2:

Data Class 3:

Replication

Real Time

Schedule by Interval

Schedule by Time

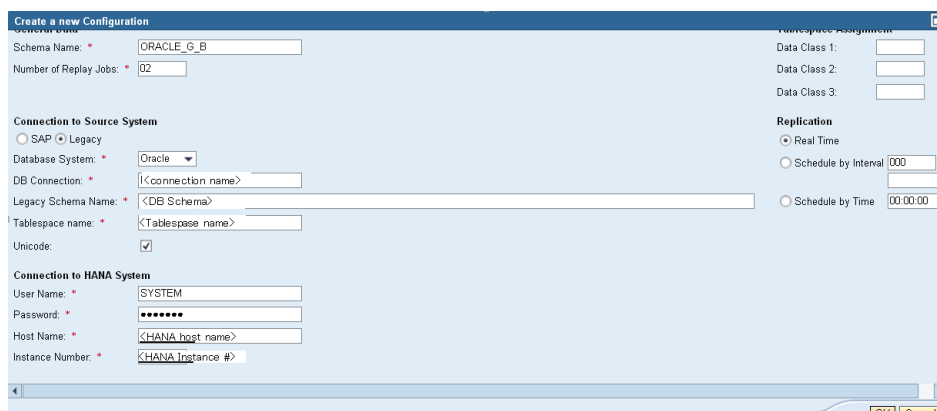
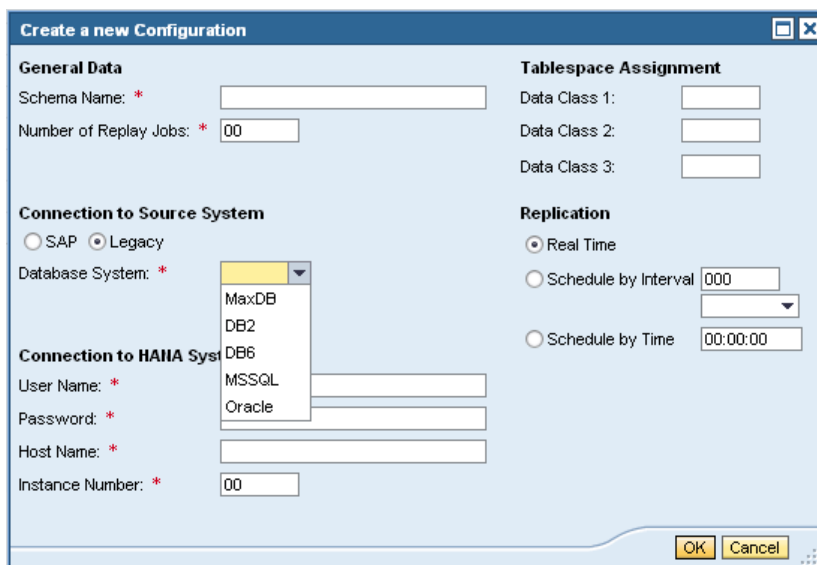
OK Cancel

a. General Data:

- **Schema Name:** Define a schema name
- **Number of Replay Jobs:** This value specifies the number of replay jobs (background jobs) which will run in the SLT system to replicate data from the connected source system to the SAP HANA system.

b. Connection to Source System:

- Choose **SAP** or **Legacy** (non-SAP) as your source system
- **RFC Destination** (if SAP source system): Enter the RFC connection you defined for the SAP source system.
- **Legacy-specific settings:** To replicate data from non-SAP source system select *Legacy* and the affected database system. Depending on the database system, additional information may be required.



- c. Connection to SAP HANA system: Supply the SAP HANA database connection information.
 - **User Name:** Specify the username that will be used to create the schema in the SAP HANA database. Usually this is the system user in the SAP HANA system.
 - **Host Name:** Specify the hostname of the SAP HANA database. This field is limited to 13 characters. A FQDN is not necessary. Ensure that the hostname corresponds to the correct IP address.
 - **Instance Number:** This is the instance number of the SAP HANA database.
 - d. Table space assignment: If you have chosen to use a separate table space for SLT related log tables in the source system (see section 3.2), you can optionally supply this information here.
 - e. Replication: you can select:
 - Real time - for real time, permanent replication
 - Scheduled – for scheduled replication by intervals
 - Scheduled by time – for example daily (at a fixed time) replication
3. To complete your schema configuration, choose *OK*.

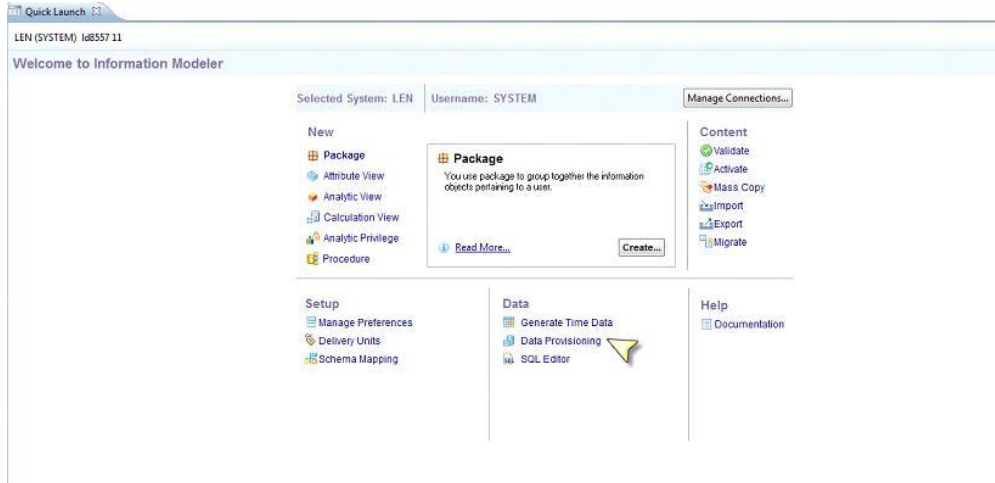
3.3 Data Provisioning

You use the SAP HANA Studio to initiate and control the replication process of the SLT replication system. There are some unique factors for the configuration using the trigger-based replication; they are covered in this section.

Note: The version of the SAP HANA studio has to at least match the revision level of the installed SAP HANA software.

Procedure

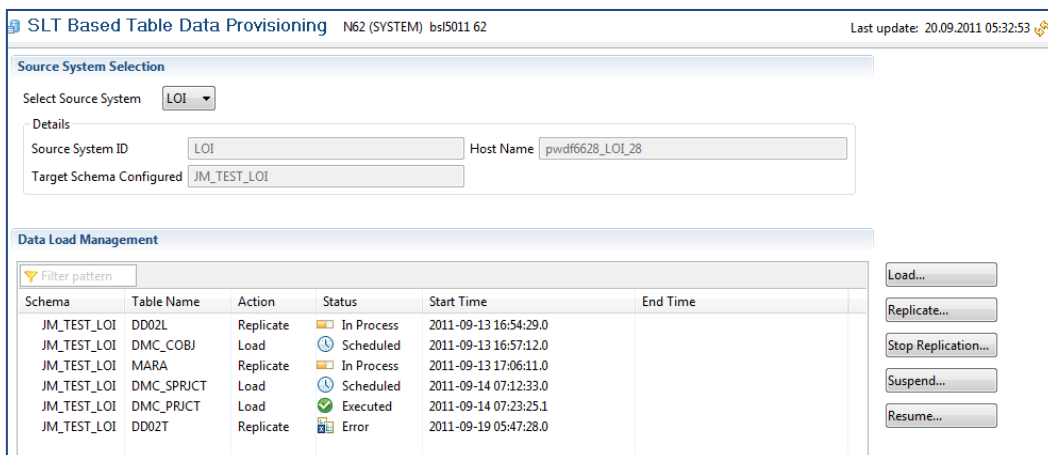
- Using the SAP HANA Studio, access the information modeler. Select *Data Provisioning*.



- In the *SLT Based Table Data Provisioning* screen, you can manage the SLT-based table replication.

Select source system: This is the system ID of the source system connected by a configuration in the SLT system as described in section 5.1.

Details: This information is automatically populated from the configuration.



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

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

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

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

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

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

4 Operations and Management of SAP LT Replication Server

SAP provides you with an infrastructure to help your technical support consultants and system administrators effectively manage all SAP components and complete all tasks related to technical administration and operation.

4.1 Important Tables and Table Structures of Trigger-based Replication Using SAP LT Replication Server

4.1.1 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.

SAP LT replication server as a trigger-based replication approach does not perform a mass import of all source tables to the SAP HANA system when you start to use it. SAP HANA studio users can select dedicated table(s) for initial load and/or for continuous replication. This is possible since the dictionary tables DD02L and DD02T are automatically replicated from the source system into SAP HANA when creating a schema. Once the dictionary tables are successfully replicated into SAP HANA, any application table from the source system can be selected for initial load and / or replication.

To ensure that new tables in the source system(s) are automatically reflected in SAP HANA, the dictionary tables DD02L and DD02T always remain in the status *Replicate, In process*.

4.1.2 Table Structure in SAP HANA Modeler

For each schema, related control tables are automatically generated while saving a configuration schema in the SLT system. 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.

Overview of important control tables:

Schema Name	Contains Tables	Owner
SYS_REPL	RS_REPLICATION_COMPONENTS	SYSTEM
<SCHEMA>	DD02L: see definition above DD02T: see definition above 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>

The control tables provide the flexibility to explicitly grant schema-specific authorizations to dedicated SAP HANA studio users. For more information about user creation and administration, see chapter 3.1.3.

4.2 Administration of Load and Replication Process

4.2.1 Mass Transfer ID and Schema GUID

When creating a new configuration in the SLT system, 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 SLT system.

4.2.2 Important SAP LT Replication Server Jobs

1. Monitoring Job

Naming Convention: IUUC_MONITOR_<MT_ID>

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 (under schema in SAP HANA).

2. Master Controller Job

Naming Convention: IUUC_REPLIC_CNTR_<MT_ID>_<4digits>

Maximal delay of this job is 59 seconds. This job is responsible for:

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

3. Data Load Job

Naming Convention: DTL_MT_DATA_LOAD_<MT_ID>_<2digits>

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 SLT server

4. Migration Object Definition Job

Naming Convention: IUUC_DEF_MIG_OBJ_<2digits>

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.

5. Access Plan Calculation Job

Naming Convention: ACC_PLAN_CALC_<MT_ID>_<2digits>

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 LT replication. For a normal sized table, access plan calculation

should finish quite quickly (less than 1 minute), while large tables might take more time to finish.

What is the relationship between “#Replay jobs” in the schema settings and the jobs?

Each job occupies 1 BGD work process in the SLT system. For each schema, the parameter *#Replay jobs* restricts the maximum number of Data load jobs the MT_ID can have. There is always one monitoring job (for the schema), one master controller job (for the MT_ID) and only one data load job (for the MT_ID).

For example, in the definition of “SCHEMA1”, parameter “#Replay Jobs” is set to 04, and mass transfer ID 001 is assigned. As a result, there should be the following jobs in the system:

- 1 Monitoring job: IUUC_MONITOR_SCHEMA1
- 1 Master controller job: IUUC_REPLIC_CNTR_001_0001
- At most 4 parallel jobs for MT_ID 001: DTL_MT_DATA_LOAD_001_01/~02/~03/~04

4.2.3 Job Handling and Performance Improvements for Initial Load of Tables

4.2.3.1 Restarting jobs

If the replication should be stopped on purpose or was aborted because an important job in the SLT system discontinued, it is possible to restart the monitoring job which initiates all subsequent jobs from the Configuration and Monitoring Dashboard (tab page *Jobs and Connections*):

The screenshot displays the 'Configuration JM_LOI' interface. At the top, it shows 'Schema Name JM_LOI' and 'Status OO' with a green indicator. Below this are 'Target Parameters bsI5011:62' and buttons for 'Edit', 'Close', and 'Refresh'. The main area has four tabs: 'Jobs and Connections' (selected), 'Triggers', 'Statistics', and 'Properties'. The 'Jobs and Connections' tab is divided into two panels. The 'Jobs' panel shows 'Master Job' with a status of 'Running' and a 'Restart' button, and 'Load Jobs' with a status of 'On'. It also displays 'Total Available: 10', 'Total Defined: 01', and 'Total Running: 01'. The 'Connections' panel shows 'Source System' and 'Target System' details, including 'RFC Connection: LOI', 'System User Name: SYSTEM', 'Host Name: bsI5011', and 'Instance Number: 62'. A footer note states: 'To change settings, enter Properties tab and switch to edit mode'.

4.2.3.2 Stopping Jobs

If you need to pause the replication from the SLT system, it is possible to stop schema-related jobs by using transactions SM37.

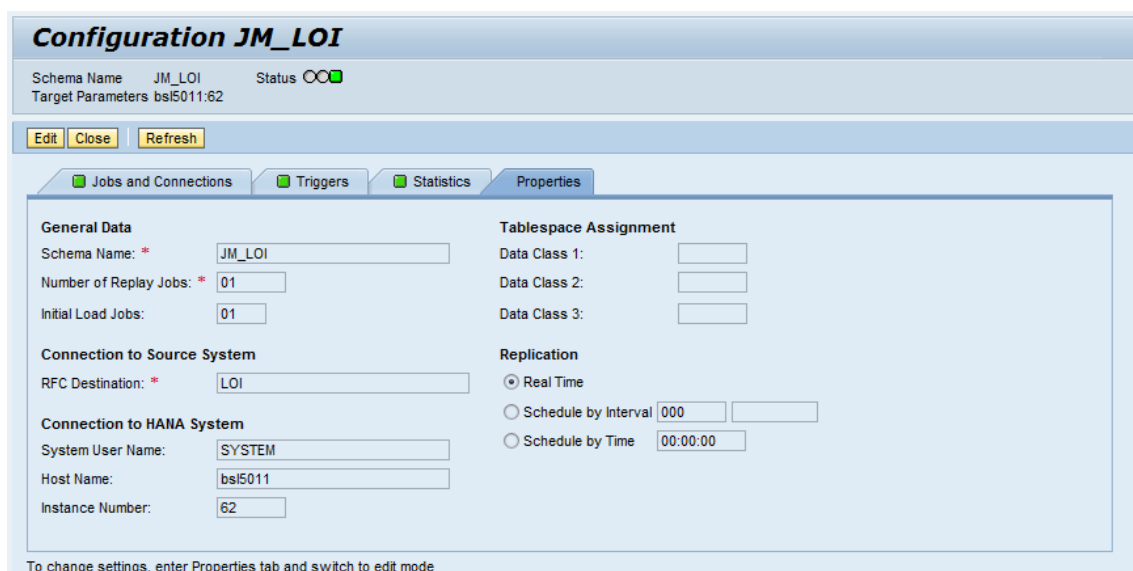
As an alternative, to temporarily stop and restart the replication by pausing schema-related jobs you can switch the replication mode in the tab page *Properties* from “Real time” to “Schedule by time”.

Note:

- You need to cancel all schema-related jobs – consider the naming convention with related mass transfer ID as explained above.
- When the data is being loaded initially, we recommend not stopping the schema-related jobs

4.2.3.3 Changing the Number of Jobs

Once a configuration and related schema is defined, it is still possible to adjust the number of jobs on the *Properties* tab page.



Configuration JM_LOI

Schema Name JM_LOI Status ○○
Target Parameters bsf5011:62

Edit Close Refresh

Jobs and Connections Triggers Statistics **Properties**

General Data

Schema Name: * JM_LOI
Number of Replay Jobs: * 01
Initial Load Jobs: 01

Connection to Source System

RFC Destination: * LOI

Connection to HANA System

System User Name: SYSTEM
Host Name: bsf5011
Instance Number: 62

Tablespace Assignment

Data Class 1:
Data Class 2:
Data Class 3:

Replication

Real Time
 Schedule by Interval 000
 Schedule by Time 00:00:00

To change settings, enter Properties tab and switch to edit mode

If you are not satisfied with the performance of the initial load and / or the replication time - or the SLT system does have more capacity available for the trigger- data replication - you can change the number of replay jobs and / or jobs for initial load for a configuration and related schema. Especially if the initial load of tables is completed, you may reduce the number of jobs for initial load.

4.2.3.4 Performance Improvements for Initial Load of Tables

Data replication using SAP LT replication server uses a dedicated technique of reading relevant data from the source system. This cluster-based approach, which considers fixed portions of data records has several advantages compared to other alternatives (reading types) – however it is currently not possible to speed-up the initial load especially for large tables by parallelizing the load using multi-threading technology (multiple jobs involved in data replication of one table).

Therefore, even changing the number of jobs for initial load as outlined in chapter 4.2.3.3 may not help to improve the performance of the initial load for a table. This is only possible by switching to a reading type, for example to “access plan based calculation” (reading type 1).

Note:

- With DMIS_2010 SP05, multi-threading technology cannot be used to accelerate the initial load for SAP cluster tables.

- We recommend that you engage the services of an SAP performance specialist before changing any performance settings.

4.3 Monitoring of Load and Replication Process

With SAP HANA SPS03 (SP05 of DMIS_2010_1_700), the *Configuration & Monitoring Dashboard* provides information you can use to monitor and identify potential replication issues.

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

SLT Replication Monitor

Available Configurations

New Delete Check Status

Schema Name	Target Parameters	Status
JM_LOI	bsl5011:62	○○○
TEST_BG7	bsl5011:62	●○○

Last Refresh: 28.09.2011 15:44:15 Refresh

Configurations
 Schema Name: JM_LOI
 Target Parameter: bsl5011:62

Connection to Source System
 RFC Connection: LOI

Connection to HANA System
 Host Name: bsl5011
 Instance Number: 62

From the overview screen, you can identify whether the overall status of the configuration and schema definition completed without errors.

4.3.1 Status of Jobs and Connections

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

Configuration JM_LOI

Schema Name JM_LOI Status ○○○
 Target Parameters bsl5011:62

Edit Close Refresh

Jobs and Connections Triggers Statistics Properties

Jobs

Master Job
 Status: Running Restart

Load Jobs
 Status: ■

Total Available: 10
 Total Defined: 01
 Total Running: 01

Connections

Source System
 Status: ■

RFC Connection: LOI

Target System
 Status: ■

System User Name: SYSTEM
 Host Name: bsl5011
 Instance Number: 62

To change settings, enter Properties tab and switch to edit mode

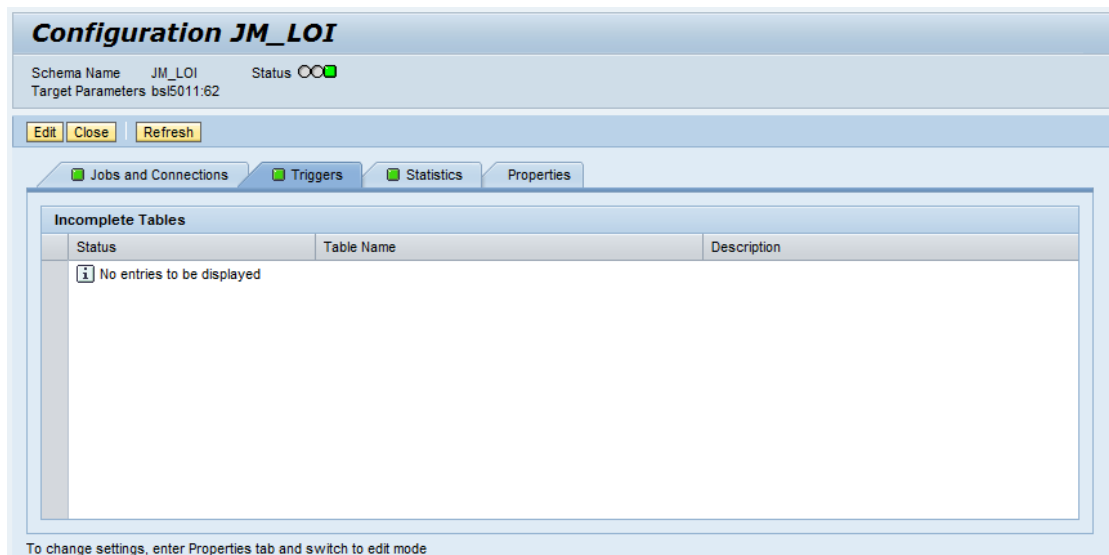
If the master job was aborted and is not running, you can try to identify the potential reasons by using transaction SM37 before restarting the job.

If status of jobs is not ok, you may check your settings compared to actual system performance.

If the connection to the source system (in case of SAP sources RFC connection) is not ok, you can verify if network and/or source system is still properly connected, for example using transaction SM59. Note that changes to the authorization setting of the RFC user in the source system could also cause this issue.

If the connection to the target (SAP HANA) system is not working properly, check system availability or other potential issues, such as the connectivity and configuration issues outlined in chapter 6.

4.3.2 Status of Triggers



Once a table is selected for replication in the *SLT Based Table Data Provisioning* screen, related database triggers are automatically created in the source system to record changes.

In the *Triggers* tab page, you can view all failures which appeared during the creation or the runtime of the triggers.

Note: In specific cases, you may need to delete existing database triggers directly in the source system. We recommend that you involve an SLT expert to support you with this task, or that you create a customer message using component BC-HAN-LTR.

4.3.3 Statistics

Similar to the *SLT Based Table Data Provisioning* screen, the *Statistics* tab page, provides detailed information about the tables in replication mode like the current action and status. In addition, you can view the replication latency of each table, calculated with the median, the minimum, and the maximum.

Configuration JM_LOI

Schema Name JM_LOI Status ○○○
Target Parameters bs15011:62

Edit Close Refresh

Jobs and Connections Triggers Statistics Properties

Replicated Tables									
Status	Table Name	Current Action	Current Status	Latency Med (all)	Latency Med (24h)	Latency Min (24h)	Latency Max (24h)	Last Replication	
■	DD02L	Replication	In Process	0.45 sec	0.30 sec	-0.00 sec	10.71 sec	16.59 min ago	
■	DD02T	Replication	In Process	0.41 sec	0.41 sec	-0.01 sec	8.68 hrs	16.59 min ago	
■	DMC_COBJ	Replication	In Process	0.36 sec	0.36 sec	-0.01 sec	3.46 sec	12.69 hrs ago	

To change settings, enter Properties tab and switch to edit mode

If the latency time does not meet your expectations, you can change the number of jobs as outlined in chapter 4.2.3.3.

4.4 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 SLT system 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	<ul style="list-style-type: none"> - Replication is stopped - SLT waits for source system to be available again 	<ul style="list-style-type: none"> - Restart the source system - SLT will continue from where it stopped
SLT system goes down (or source system and SLT, if in the same stack)	<ul style="list-style-type: none"> - Replication is stopped 	<ul style="list-style-type: none"> - Restart SLT jobs
SAP HANA system goes down	<ul style="list-style-type: none"> - Replication is stopped - SLT waits for HANA DB to be available again 	<ul style="list-style-type: none"> - Restart the HANA system - SLT will continue from where it stopped

If the source system or the SAP HANA system cannot be fully recovered, tables have to be dropped and loaded 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.

5 Special Considerations for Source Systems

5.1 Impact of Software Maintenance Events in SAP Source Systems

Software maintenance in the source system can affect tables that have the status *In Replication*. For tables that have the status *In Replication* (with database triggers activated), the ABAP data dictionary blocks any structural changes to the table.

As a consequence, the table activation will not be completed successfully if a trigger is active and the table structure was changed. Similar issues apply for the import of transports containing structural changes of tables – the ABAP dictionary activation step will end with an error.

Therefore, maintenance events that affect tables that have the status *In Replication* need to be identified in advance. There are no predictive or comparison tools that help to identify affected tables – so structure changes must be identified manually.

General procedure:

1. Request transport owner to identify any dictionary change of replicated tables
2. Stop or suspend tables that have the status *In Replication* prior to the maintenance event or prior to the import of transports
3. Take a full or incremental backup of the SAP source system prior the maintenance event or importing transports
4. Apply software maintenance or import transports in the test system first to identify the impact of replicated tables
5. Take any necessary actions depending on the dictionary changes of replicated tables
6. Restart replication or resume the replication after the maintenance event or after the import of transports.

Note:

- In general, we recommended that you stop and restart the replication in the SAP HANA studio to ensure data consistency
- If affected tables include a lot of data and the time for a new initial load is not acceptable, the following procedure might be an alternative:
 - Use the suspend function (in the HANA system) and delete the triggers directly in source system (transaction IUUC_REMOTE)
 - After the maintenance event, activate the delta recording (recreate the triggers) with a manual SQL command (insert in table RS_ORDER with action 'M') and start replication without initial load using a manual SQL command (insert in table RS_ORDER with action 'Q')

To use this procedure, the data replication of affected tables needs to be completed and a downtime of the source system is required (time from deletion of trigger to activation of the delta recording).

- If the structure changes affect key fields, you need to perform an initial load again
- If you expect data for new (non-key) fields also to be replicated or to drop data for removed fields in the SAP HANA database, you need to adjust the affected tables accordingly (alter tables) and start replication without initial load using a manual SQL command (insert in table RS_ORDER with action 'Q'). This may also require you to drop and to recreate views and indexes.

5.2 Archiving Data in Source Systems

The trigger-based approach replicates any deletion in source tables. Since it is currently not possible to distinguish on the database level between delete actions caused by archiving versus regular deletion of data records, SAP LT Replication Server will also replicate archiving activities in the source system as delete actions in the SAP HANA database.

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

Considering the relationship of jobs as outlined in chapter 4.2.3, the main system load of the trigger-based replication approach resides on the SLT system and related background jobs.

Detailed reviews and 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.

However, as a rule of thumb, SAP assumes a maximum of 2-3% performance impact on the source system. If the number of background jobs is significantly increased to speed up the initial load of tables or replication, the impact might be higher since a related number of dialog processes are consumed in the source system by the replication process.

In specific cases, it might be reasonable to leverage a more powerful application server especially during dedicated initial load activities.

5.4 Data Volume Management

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

1. Source system(s):
 - SLT related logging tables: The size of logging tables may increase in size, in case the replication is suspended for a longer time and/or there is a system outage of the system of the SLT system 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 SAP sources 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 – you may check if RFC related logging is critical.
2. SLT system: Since the data transfer – managed in dedicated portions of data records - during the replication process is handled by the memory of the SLT system, 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.
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 Further Tips and Tricks, Frequently Asked Questions, and Troubleshooting Recommendations

If SAP LT Replication Server for SAP HANA is already in use, what to consider when applying SP05 to leverage new features available HANA SPS03?

See SAP Note 1649910 for details.

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

SAP LT Replication Server includes various transformation capabilities. However - per 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 transformation rules for filtering of data in order to replicate only selected / relevant data into the SAP HANA database or for more advanced transformation requirements such as the scrambling of data. Transformation rules have to be implemented for respective tables on the SLT system 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 SLT expert in this matter.

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

The fundamental concept of the trigger-based replication is designed as for SAP source systems. The main differences are that the connection is a database connection (no RFC connection) and the read modules reside on the SLT system (as described in chapter 2.1.2).

Furthermore, consider the following:

- Due to the database trigger concept, tables in non-SAP source systems must have a primary key to be considered for replication.
- Tables DD002L and DD002T (for SAP source systems used to replicate the SAP dictionary information) include the metadata of the tables in the non-SAP 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 source systems.

Is it possible to replicate data from multiple source systems to one schema in the SAP HANA database?

Yes, SAP LT Replication Server supports N:1 replication (multiple source systems replicating into the same schema of the SAP HANA database). However, you need to consider specific preparation steps especially if the source systems include the same tables that should be considered for data replication. We recommend involving an SLT expert in this matter.

What are the potential issues if the creation of a configuration 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 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. Refer to 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.

How to ensure that data is consistent in the source system and 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. Since each entry of a logging table is deleted after a successful “commit” statement from the SAP HANA database, data consistency is guaranteed 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 the replication status to be verified from a business perspective are currently not available.

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.

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

If there are more than 2 billion records in a source table, it needs to be partitioned into multiple SAP HANA tables.

With SAP LT Replication Server for SAP HANA SP03, this type of partitioning is not automatically supported. However, it is planned to release soon an “active partition management” that allows you to add and to drop partitions later (for example, when doing a range partitioning for time periods).

If you face such a requirement today, we recommend you to proceed as follows:

1. Initiate table creation in SAP HANA database using the standard SLT approach (using load or replicate feature in the SAP HANA data provisioning UI)
2. Suspend data load and create partition(s) for tables manually as described in the Guide “SAP HANA database – Partitioning and Distribution of large tables”

Note: The decision on the most suitable partitioning type is very much application dependent (the way how to access the data).

3. Resume data load / replication

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 SAP source systems.

7 Appendix

7.1 Related Guides

You can find more information about installation and configuration in the Installation Guide – Trigger-based Data Replication Using SAP LT Replication Server.

7.2 Related Information

The following table contains links to information relating to the Application Operations Guide.

Content	Quick Link to the <i>SAP Service Marketplace</i> (service.sap.com)
Installation Guide - Trigger-based Data Replication Using SAP LT Replication Server	SAP HANA Installation Guide - Trigger-based Replication (SLT)
Security Guide – Trigger-based Data Replication Using SAP LT Replication Server	SAP HANA Security Guide - Trigger-based Replication (SLT)
Technical Operations Guide – SAP HANA Appliance Software	Technical Operations Guide – SAP HANA Appliance Software
SAP LT Replication Server for SAP HANA - Central SAP Note 1605140	1605140
Upgrade Considerations Using SAP LT Replication Server with SAP HANA SPS03	1649910