Sizing Guide
DOCUMENT VERSION: 3.1 – 2022-11-08

INTERNAL – AUTHORIZED FOR SAP CUSTOMERS AND PARTNERS

Sizing SAP Landscape Transformation Replication Server

THE BEST RUN SAP

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.

Document History

Version	Date	Change
1.0	2010-12-06	Initial Version
1.1	2014-01-29	Revised published version for SAP HANA Platform SPS07 and DMIS 2011 SP06
1.2	2014-03-10	Updates to reflect application component change from BC-HAN-LTR to HANDP-LTR.
1.21	2015-01-21	Revised published version for SAP HANA Platform SPS09 and DMIS 2011 SP08
1.22	2015-12-28	Revised published version for SAP HANA Platform SPS11 and DMIS 2011 SP10
1.23	2016-06-10	Revised published version for SAP HANA Platform SPS11 and DMIS 2011 SP11
1.24	2017-06-10	Revised published version for SAP HANA Platform 2.0 SPS00 and DMIS 2011 SP12
1.25	2018-01-16	Revised published version for SAP HANA Platform 2.0 SPS02 and DMIS 2011 SP14
1.26	2018-09-10	Revised published version for SAP HANA Platform 2.0 SPS03 and DMIS 2011 SP15
2.0	2019-01-28	Revised published version for DMIS 2011 SP16 and DMIS 2018 SP01
2.1	2019-09-09	Revised published version for DMIS 2011 SP17 and DMIS 2018 SP02
2.2	2020-02-26	Revised published version for DMIS 2011 SP18 and DMIS 2018 SP03
3.0	2022-08-03	Extended with guidance on how to size SAP Landscape Transformation Replication Server with SAP Quick Sizer. Resource consumption update. Examples improved.
3.1	2022-11-08	Additional information added about supported scenarios for SAP Quick Sizer.

Table of Contents

1. INTRODUCTION

The main function of SAP LT Replication Server is to replicate data from selected tables from SAP systems or non-SAP systems to SAP HANA systems. Some adjustments of data from source systems to the SAP HANA database format such as conversion to UNICODE format happen automatically during the replication process. It is possible to define transformation rules for filtering of data or for more advanced transformation requirements such as the scrambling of data. Transformation rules must be defined for the relevant tables in the SAP LT Replication Server system before the data replication starts.

1.1 Architecture of SAP Landscape Transformation Replication Server

SAP LT Replication Server is flexible with regards to landscape installation and configuration. It can be installed as an add-on on any SAP source system, or on a separate host. SAP LT Replication Server can also be used as a central component which serves multiple source and multiple target systems.

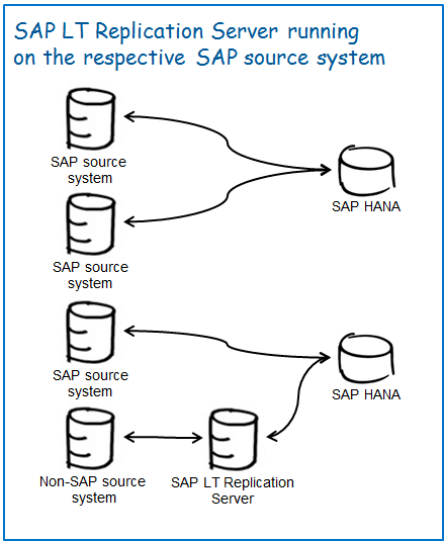

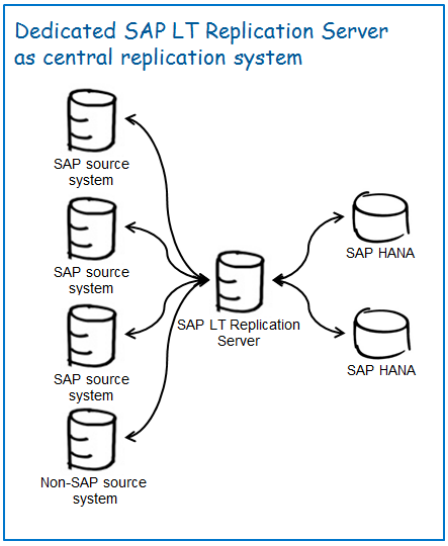

For SAP source systems, the data is extracted by means of remote function calls and pushed to HANA by means of direct SQL calls.

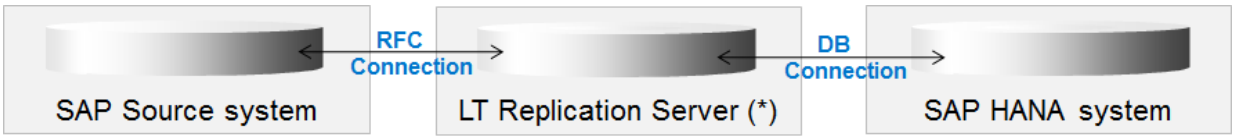

For non-SAP systems, the data is extracted with direct SQL calls through a database connection.

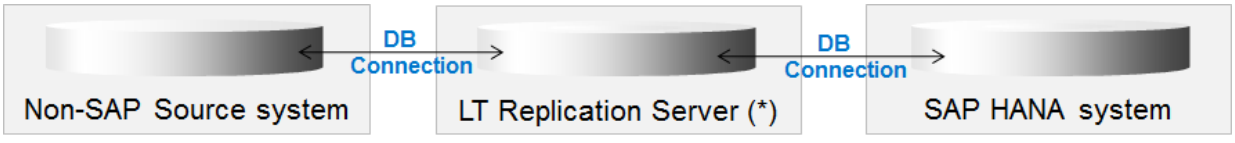

a) Multiple source systems are replicated to a single HANA schema

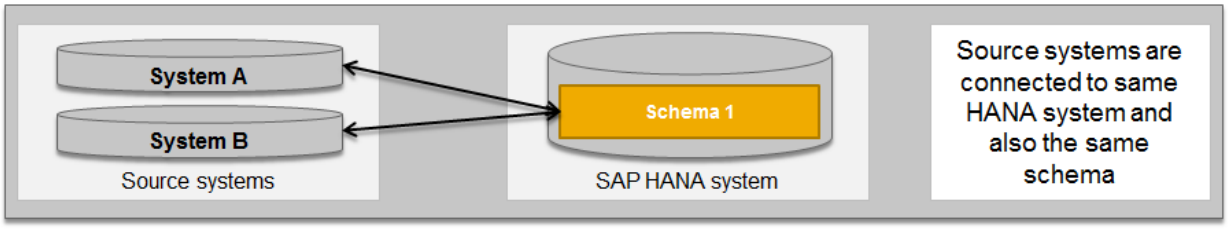

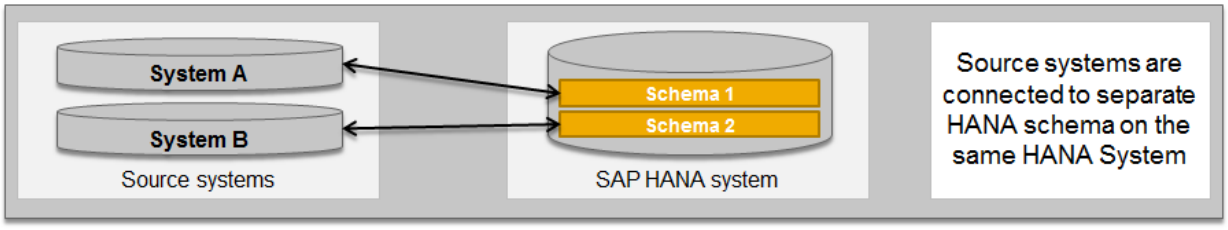

c) Multiple source systems are replicated to their own separate SAP HANA system.

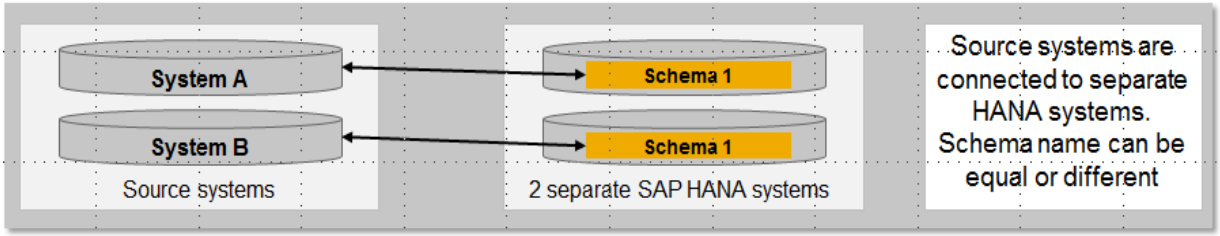

d) One source system is replicated to up to four separate schemas in more than one SAP HANA system.

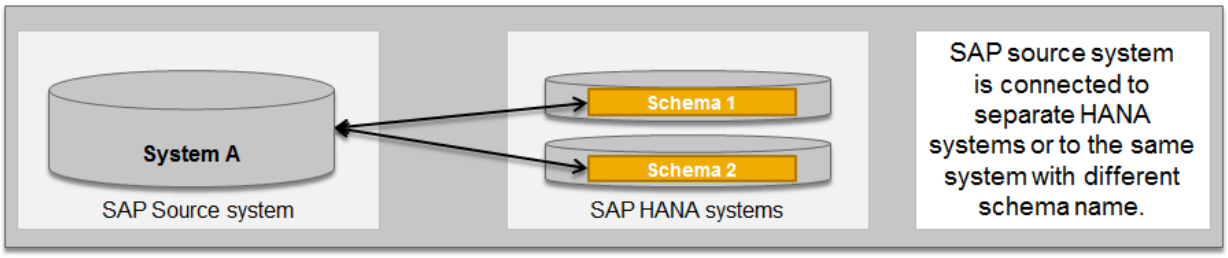

1.2 Factors that Influence the Performance

There are various factors which influence the performance of the replication. These include the following:

a) Version of SAP LT Replication Server and HANA Software

A major prerequisite for good SAP LT Replication Server performance is to have all SAP LT Replication Server related SAP Notes and recommendations outlined in the Installation Guide and in the SAP LT Replication Server central SAP Note applied to the system. For more information, see Chapter 5 "Miscellaneous".

b) Customer-Specific Configuration

The use of advanced replication settings, for example complex transformation rules for filtering data, advanced data mapping, or conversions requirements cause additional resource consumption and influence the response time. These scenarios are not covered in this sizing guide - we recommend that you engage the services of an SAP LT Replication Server consultant to support you when calculating these sizing requirements.

c) Network Factors

The performance of SAP LT Replication Server solution benefits from a network with low latency and a properly estimated bandwidth according to the Network Sizing Chapter in this sizing document.
In environments with frequent network failures, the replication will be repeated until the replication is successfully completed, therefore resulting in increased resource consumption and longer execution times.

d) Data reading type

SAP LT Landscape Transformation Server supports the following reading types:

Reading Type	Advantages	Disadvantages
Type 1 – Access Plan Calculation	• Fast data load if index exists • Parallel data load possible	• An additional index may be required. • Requires a key field which is sufficiently selective. • Calculation required before load
Type 3 (default) – DB_SETGET	• No separate index required • Parallel data load possible with DMIS_2010 SP07 or higher	• Additional consumption of database buffer
Type 4 & 5 – Index Cluster	• Very fast data load after data is extracted to table DMC_INDXCL • Minimal usage of database buffer	• Additional tablespace temporarily required in the source system

 For initial sizing, this sizing guide only covers reading type 3, which is the default option for SAP LT Replication Server. If you want to use another reading type, please contact SAP for support.

e) Appropriate Configuration of Source system, SAP LT Replication Server and HANA system

There is direct relationship between the number of replication jobs in the SAP LT Replication Server system and the number of background work processes for the ERP source system. The source system and SAP LT Replication Server should have enough background work processes to ensure parallel execution and full allocation of the hardware configuration which is sized for the expected replication.

 **An insufficient number of background work processes would lead to sequential execution of the replication, and therefore result in delays while replicating data changes to the SAP HANA system.**

For each SAP LT Replication Server 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 total, one mass transfer ID requires that the following background work processes are available in the SAP LT Replication Server system:

- 1 master controller job (/1LT/IUC_REP_CNTR_<MT_ID>)
- 1 job for defining migration objects (/1LT/IUC_DEF_COBJ_<MT_ID>)
- N jobs for calculating access plans (/1LT/IUC_CALC_<MT_ID>_<2digits>)
- N data transfer jobs (/1LT/IUC_LOAD_MT_<MT_ID>_<3digits>)

For each access plan calculation job in the SAP LT Replication Server system, there is a corresponding job in the source system with the naming convention /1LT/MWB_CALCS_<MT_ID>_<table name with maximum 10 characters>.

In the source system, the number of available dialog work processes which are reserved for the replication should be equal to the number of data transfer jobs running in the SAP LT Replication Server system.

 Note

Ensure that the source system and the SAP Landscape Transformation Replication Server has enough additional dialog and background work processes to handle their own tasks.

f) Data Volume

If the table size exceeds 2 billion records, you must split the table by using the available partitioning features. For more information, see the SAP HANA Administration Guide for SAP HANA Platform, section “Table Partitioning” on the SAP Help Portal at http://help.sap.com/hana.

2. SIZING FUNDAMENTALS AND TERMINOLOGY

SAP provides general sizing information on www.sap.com/sizing. For the purpose of this guide, we assume that you are familiar with sizing fundamentals. This section explains the most important sizing terms, as these terms are used extensively in this document.

Sizing

Sizing means determining the hardware requirements of an SAP application, such as network bandwidth, physical memory, CPU processing power, and I/O capacity. The size of the hardware and database is influenced by both business aspects and technological aspects. This means that the number of users using the various application components and the data load they put on the server must be taken into account.

Benchmarking

Sizing information can be determined using SAP Standard Application Benchmarks (www.sap.com/benchmark). Released for technology partners, benchmarks provide basic sizing recommendations to customers by placing a substantial load upon a system during the testing of new hardware, system software components, and relational database management systems (RDBMS). All performance data relevant to the system, user, and business applications are monitored during a benchmark run and can be used to compare platforms.

SAPS

The SAP Application Performance Standard (SAPS) is a hardware-independent unit that describes the performance of a system configuration in the SAP environment. It is derived from the Sales and Distribution (SD) Benchmark, where 100 SAPS is defined as the computing power to handle 2,000 fully business processed order line items per hour. (For more information about SAPS, see www.sap.com/benchmark → Measuring in SAPS).

Greenfield Sizing

Greenfield sizing refers to the sizing approach that provides statements about platform-independent requirements of the hardware resources necessary for representative, standard delivery SAP applications. The initial sizing guidelines assume optimal system parameter settings, standard business scenarios, and so on.

Expert Sizing

This term refers to a sizing exercise where customer-specific data is being analyzed and used to put more detail on the sizing result. The main objective is to determine the resource consumption of customized content and applications (not SAP standard delivery) by comprehensive measurements. More information can be found here.

Configuration and System Landscaping

Hardware resource and optimal system configuration greatly depend on the requirements of the customer-specific project. This includes the implementation of distribution, security, and high availability solutions by different approaches using various third-party tools. In the case of high availability through redundant resources, for example, the final resource requirements must be adjusted accordingly.

3. INITIAL SIZING FOR SAP LANDSCAPE TRANSFORMATION REPLICATION SERVER

If the source system is an SAP S/4HANA system (release 1909 or higher), the HANA-based Quick Sizer must be applied to determine the hardware requirements; therefore skip the sizing guidance in Chapter 3 and go to Chapter 4 “Sizing SAP Landscape Transformation Replication Server with Quick Sizer”, where you will find guidance on how to use the Quick Sizer tool.

Note that it is only possible to use the HANA-based Quick Sizer if data is replicated directly from one source table to one target table, it cannot be used for more complex scenarios. For example if you are replicating data to SAP Central Finance (cFIN) system, then the HANA-based Quick Sizer cannot be used.

3.1 Assumptions

This sizing approach is based on the use case where data from one source system is replicated to a single HANA schema without any complex data transformation.

When replicating data from multiple source systems to one or different HANA schemas, you should sum up the requirements of the single SLT configurations and calculate the overall expected initial load capacity or the expected maximum replication throughput rate as described in this chapter.

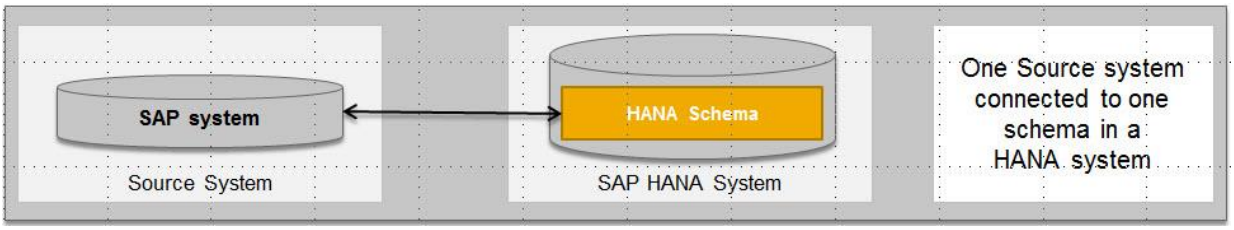

Categorization of Tables

As input for sizing of a replication scenario, you need to analyze tables which will be replicated, and classify them into categories. Determine the following information for all tables (or only for the most frequently modified (inserted, updated and deleted)):

- The weighted average number of table columns (one value C_W)

  The weighted average record length (one value, L_W)

Analyze the tables and determine the appropriate table category according to **Table 1**, i.e., either small (S), medium (M), large (L), or extra-large (XL).

Table 1: Table categories

Categories	Up to 150 columns	151 to 250 columns	More than 250 columns
< 1500 bytes per record	S	M	L
> 1500 bytes per record	M	L	XL

Example 1

Categorization of individual replication-relevant tables in the source system based their characteristics.

Table Name	# Columns	Length [bytes]	Category

COBK	41	510	S
BKPF	111	1350	S
ANLA	136	2130	M
MSEG	180	2170	L
BSEG	312	3250	XL

Example 2

Weighted categorization of replication-relevant tables in the source system based their characteristics and modification rate.

Weighted average number of columns (C_W) = (3.5*41 + 66*111 + 3.5*136 + 13.5*180 + 13.5*312)/100 = (143 + 7326 + 476 + 2430 + 4212)/100 = 14587/100 = ~146 columns

Weighted average bytes per record (L_W) = (3.5*510 + 66*1350 + 3.5*2130 + 13.5*2170 + 13.5*3250)/100 = (1785 + 89100 + 7455 + 29295 + 43875)/100 = 171510/100 = ~1715 bytes

Table Name	# Columns	Length [bytes]	Modification rate (per hour)	% of all modifications
COBK	41	510	50.000	3.5%
BKPF	111	1350	1.000.000	66%
ANLA	136	2130	50.000	3.5%
MSEG	180	2170	200.000	13.5%
BSEG	312	3250	200.000	13.5%
Total			1.500.000	100%
Weighted avg. number of columns (C_W)	146			
Weighted avg. bytes per record (L_W)		1715		
Weighted Category = M				

Throughput and number of data transfer jobs

For sizing, the target throughput (**T**) should be converted to metric “records per seconds”.

Example 3

If **NR** is the total number of replication-relevant records in all replication-relevant tables, and **D[h]** is the target duration of the initial load or scheduled replication in hours, then

```
T[records/second] = NR / (3600*D[h])
```

Note

- For “initial load”, NR is equal to the total number of records of all load-relevant tables.
- For “scheduled replication”, NR is equal to the number of modified (inserted, updated, deleted) records of all replication-relevant tables for the period since last replication completed.
- For “real-time replication”, NR is the total number of expected modified (inserted, updated, deleted) records per time interval (e.g., hour) for all replication-relevant tables.

The sizing result is linear to the number of data transfer jobs, which are required to achieve the expected throughput.

Based on the target throughput, i.e., records per second, the required number of data transfer jobs (**NJ**) should be calculated:

a) For regular tables:

```
NJ = T[records/second]/1000
```

b) For cluster tables:

```
NJ = T[records/second]/600
```

3.2 Template-Based Sizing Approach

In *Table 2* are provided typical use-cases with recommended configuration and precalculated sizing result. If your use-case deviates significantly from these templates, apply the sizing procedure as described in Chapter 3.3 “Sizing Guideline for SAP Landscape Transformation Replication Server”.

Table 2: Templates for typical use cases

	SMALL volume scenario	MIDRANGE scenario	UPPER MIDRANGE scenario
Number of source systems	1	3	10
Number of tables	50	200	500
Weighted category	S - M	M – L	M - XL
Expected throughput	1,000,000 records per hour	10,000,000 records per hour	50,000,000 records per hour
Number of Data Transfer Jobs in SLT	1	3	14
ABAP Work Processes on Source System, dedicated to replication	1 dialog process 2 background work processes	3 dialog processes 4 background work processes	14 dialog processes 8 background work processes
Source System(s) /net resources/	ABAP: • 90 SAPS • 60 MiB DB: 500 SAPS	ABAP: • 450 SAPS • 240 MiB DB: 2250 SAPS	ABAP: • 2520 SAPS • 1.4 GiB DB: 14000 SAPS
SLT Server (ABAP) /net resources/	• 370 SAPS • 60 MiB	• 1800 SAPS • 240 MiB	• 10500 SAPS • 1.4 GiB memory
SAP HANA System /net resources/	• 600 SAPS • 1250 MiB	• 2880 SAPS • 1900 MiB	• 16800 SAPS • 5.8 GiB memory

➡ Recommendation

If the net sizing result is lower than 1300 SAPS, 1 core for replication tasks will be sufficient.

3.3 Sizing Guideline for SAP Landscape Transformation Replication Server

The sizing result refers to the net resource requirement of the source system, the SAP LT Replication Server replication system and the target system. This net resource capacity should be available in addition to any other resource capacity, required for running other functions and tasks in that system in parallel. The minimal hardware requirements for the SLT Replication Server, if installed on separate ABAP system, are 2600 SAPS and 16 GB Memory. The sizing result should be considered on top.

CPU and Runtime Memory Sizing

Table 3: CPU Sizing

Table Category	Source ABAP Server [SAPS]	Source Database [SAPS]	SLT ABAP Server [SAPS]	HANA Database [SAPS]
S	60 * NJ	350 * NJ	250* NJ	400 * NJ
M	90 * NJ	500 * NJ	370* NJ	600 * NJ
L	150 * NJ	750 * NJ	600* NJ	960 * NJ
XL	180 * NJ	1000 * NJ	750 * NJ	1200 * NJ

Note

If the sizing result is lower than 1300 SAPS, then one core would be sufficient to fulfill the replication requirements.

Note

No scenario-specific CPU requirements to the SAP LT Replication Server own database

Table 4: Runtime Memory Sizing

Table Category	Source ABAP Server [MiB]	Source Database [MiB]	SLT ABAP Server [MiB]	HANA Database [MiB]
S	40 * NJ	Memory sizing should be performed according to the sizing procedures of the database vendors.	40* NJ	1000 + 200* NJ
M	60 * NJ		60* NJ	1000 + 250* NJ
L	80 * NJ		80* NJ	1000 + 300* NJ
XL	100 * NJ		100 * NJ	1000 + 350 * NJ

Note

No scenario-specific memory requirements to the SAP LT Replication Server own database.

How to use *Table 3* and *Table 4* for sizing is illustrated in *Example 4* and *Example 5*.

	Example 4	Example 5
Business usage	Weighted table category: L Overall change rate = 50,000,000 [records/hour]	Weighted Table Category: M Overall Change Rate = 10, 000,000 [records/hour]
Intermediate calculations	T = 13889 records/second NJ = T[records/second] /1000 = 13,9 → ~ 14 Data Transfer Jobs are required	T = 2778 records/second NJ = T[records/second] /1000 = 2,78 → ~ 3 Data Transfer Jobs are required
Sizing result	Source ABAP [SAPS] = 14 * 150 = 2100 Source DB [SAPS] = 14 * 750 = 10500 SLT ABAP [SAPS] = 14 * 600 = 8400 HANA DB {SAPS] = 14 * 960 = 13440 Source ABAP [MiB] = 14 * 80 = 1120 SLT ABAP [MiB] = 14 * 80 = 1120	Source ABAP [SAPS] =3 * 90 = 270 Source DB [SAPS] = 3 * 500 = 1050 SLT ABAP [SAPS] = 3 * 370 = 1110 HANA DB {SAPS] = 3 * 600 = 1800 Source ABAP [MiB] =3 * 60 = 180 SLT ABAP [MiB] = 3 * 60 = 180

	HANA DB [MiB] = 1000 + 14 * 300 = 5200	HANA DB [MiB] = 1000 + 3 * 250 = 1750

Disk and HANA Data Memory Sizing

Target HANA system

The data memory requirements for the target HANA are additional to the runtime memory requirements of the replication processing.

HANA is an in-memory database, which supports compression. For this reason, the table size in HANA database may differ from the size of same on the source system.

The disk sizing and HANA data memory sizing for the target HANA system should be performed using HANA-based Quick Sizer -> Technology Platforms -> Standalone HANA. As input there provide the size of the tables in the source system and, if the source system is not a HANA system, the expected compression factor.

Source system and SAP LT Replication Server system

On the source system and on the SAP LT Replication Server system there is no growth of business data; however, temporary disk space allocation for logging tables in the source system and in the SAP LT Replication Server system should be planned.

In normal situations, SAP LT Replication Server logging tables in the source system are usually empty (or close to empty). However, these logging tables may increase in size during the initial load if business processes change data in the source system. The logging tables can also increase in size if the replication is suspended for a long period of time, if there is a system outage, or if the SAP LT Replication Server system or the SAP HANA system is unavailable. For this reason, create a separate tablespace for the SAP LT Replication Server logging tables in the source system. We recommend including a safety buffer for unexpected incidents and using a tablespace with an initial size of 20 to 40 GB and an automatic extend feature is recommended.

Data replication from SAP source systems is managed by RFC connections. We recommend that you monitor the size of related RFC log files.

Network Sizing

The network requirements depend on the throughput and on the weighted avg. bytes per record (WL) and are calculated with the following formula:

$$\text{Network Bandwidth [Mbit/s]} = 8 * L_W * T[\text{records/second}] / 1000000$$

Example 6

The weighted average number of bytes per record (L_W) is 1715 bytes and 18.000.000 records should be replicated per hour.

1) Calculating the throughput in “records per second”:
   T [records/second] = 18.000.000/3600 s = 5000 records per second.
2) Applying the network sizing formula:
   Network Bandwidth [Mbit/s] = 8 * 1715 * 5000 / 1000000 = **~70 Mbit/s**

4. SIZING SAP LANDSCAPE TRANSFORMATION REPLICATION SERVER WITH QUICK SIZER

 Note that it is only possible to use the HANA-based Quick Sizer if data is replicated directly from one source table to one target table, it cannot be used for more complex scenarios. For example if you are replicating data to SAP Central Finance (cFIN) system, then the HANA-based Quick Sizer cannot be used. If the source system is an SAP S/4HANA system (release 1909 or higher), the HANA-based Quick Sizer must be applied.

The input fields which should be maintained for calculation of sizing result are as follows:

- No. of Tables (or partitions)
  Enter the number of tables or partitions that will be replicated.
- No. of records per table (or partition)
  Enter the number of records per table or partition that will be replicated.
- Avg. no. of fields per record
  Enter the average number of fields per record. If you like to replicate tables that have significantly different number of fields per record for each other, you must specify the different tables as separate input lines in the sizing table. To do so, use the standard feature of Quick Sizer tool to “insert line” in input table.
- S.t. and E.t.
  Enter the period described by start time and end time in which the replication should happen. Make sure to describe the expected behavior, i.e. for replications that will be overlapping provide overlapping Start Time and End Time period and for replications which will be non-overlapping, provide different Start Time and End Time period accordingly.

 Note

In Quick Sizer, the different tables and partitions that will be included in the replication job must be described in different input lines and the sizing result represents the total required resource to run the job with all described tables within the specified time interval **in hours**.

 Recommendation

If a table or a partition contains more than 999,999,999 records, it cannot be described on single input line in Quick Sizer; therefore, you need to provide the input for the same table over several input lines.

 Recommendation

There is limitation of total 11 * 10^9 fields (*“No. of records per table (or partition)” * “Avg. no. of fields per record”*) which can be replicated per hour per table or partition. If you exceed this data volume, you will get a notification from the QS tool. In such case, either split the table into smaller partitions or define longer processing time.

The sizing result is provided for each system separately for both CPU and Memory:

- Source System (the system, where the source tables are located)
  "Source SAPS (ABAP)", "Source Memory (ABAP, MiB)", “Source DB SAPS”, “Source DB Memory in MiB”
- SAP Landscape Transformation Replication Server
  "LTR SAPS (ABAP)", "LTR Memory in MiB". The LTR own database is not relevant for sizing with the replication functionality.
- Target Database (the HANA database where the target tables are located)
  "DB SAPS", "DB Memory in MiB".

Note

The data memory of HANA should be calculated using Quick Sizer -> “Standalone HANA”, providing the size of the tables in source system and estimated compression ratio as input for the calculation.

To configure sufficient SLT data transfer jobs, plan 1 job for each 36 MB of calculated memory result for SAP Landscape Transformation Replication Server and make sure the systems have available enough ABAP work processes to run the jobs.

5. MISCELLANEOUS

The following information is also available for SAP LT Replication Server:

- Installation Guide - Trigger-based Data Replication Using SAP LT Replication Server (https://help.sap.com/sapslt)

- SAP Note 1605140 (Central SAP Note for SAP LT Replication Server)

6. COMMENTS AND FEEDBACK

Both are very welcome!
Please use component CA-LT-SLT to address questions, comments and feedback to SAP.

www.sap.com/**contactsap**

THE BEST RUN SAP®