

2081591 - FAQ: SAP HANA Table Distribution

Version	16	Type	SAP Knowledge Base Article
Language	English	Master Language	English
Release Status	Released to Customer	Category	Problem
Component	HAN-DB (SAP HANA Database)	Released On	20.07.2018

Please find the original document at <https://launchpad.support.sap.com/#/notes/2081591>

Symptom

You have questions related to SAP HANA table distribution.

Environment

SAP HANA database

Cause

1. [Where do I find detailed information about table distribution?](#)
2. [How is data distributed in SAP HANA?](#)
3. [Why is table distribution necessary?](#)
4. [How are tables distributed across available index servers?](#)
5. [Are there specific privileges required to perform a table redistribution?](#)
6. [In which situation would you perform a table redistribution?](#)
7. [How is Landscape Redistribution done in SAP HANA?](#)
8. [How are tables distributed in SAP NetWeaver Business Warehouse on HANA system?](#)
9. [What are the parameters for table redistribution?](#)
10. [Where can I find information specific to table distribution on BW on HANA?](#)
11. [Can I stop the table distribution process? If so how? And does it roll-back all changes?](#)
12. [Should I plan a maintenance window for table distribution?](#)
13. [What can go wrong during the process and would a restore to saved config be an option or should I perform full database restore?](#)
14. [During a restore of a saved configuration, there were some failed tables, how can I determine the exact cause of failure?](#)
15. [Where can I find information regarding table location?](#)
16. [What information is available in the table distribution editor/ table runtime information?](#)
17. [Is there a need for a specific backup strategy when table redistribution is planned?](#)
18. [Is a complete database recovery necessary to revert table distribution?](#)
19. [How can I save a table distribution plan and send it to SAP for review before actually executing the plan?](#)
20. [What are the different algorithm IDs available and when they should be used?](#)
21. [Is there a way to estimate how long the process might take before starting the distribution?](#)
22. [Can I expect different execution times for the distribution on the pre-production system and production system?](#)
23. [How can I migrate from an <m> node environment to an <n> node environment?](#)
24. [Is there any tool designed to help in table distribution?](#)
25. [Where can I find additional information on Data Distribution Optimizer \(DDO\)?](#)
26. [Under which scenarios can I use Data Distribution Optimizer \(DDO\)?](#)
27. [Is there any dependency between DDO version and SAP HANA revision?](#)
28. [How is the content of DDO delivered?](#)
29. [How can the progress of table distribution activities be monitored?](#)

30. [How is the data distributed during optimize table distribution?](#)
31. [Which values are allowed for the LOCATION in TABLE PLACEMENT?](#)
32. [How can I gather information on table groups, number of records and other information relevant to table distribution via SQL?](#)
33. [Are there any special pre-requisites to be considered with regards to the disk space while planning SAP HANA table redistribution?](#)

Resolution

1. Where do I find detailed information about table distribution?

General information related to table distribution is contained in the [SAP HANA Administration Guide](#).

2. How is data distributed in SAP HANA?

SAP HANA supports two ways of distributing data

- Different tables assigned to different index servers (Database Partitioning / Table distribution)
- Same table is split across multiple index servers (Table Partitioning)

3. Why is table distribution necessary?

Tables are distributed among the available hosts to provide the necessary load balance.

In case you experience an OOM situation happening on one particular host within your landscape then it needs to be investigated if the load is properly distributed among all the available hosts.

Two kinds of scenarios are available for SAP HANA. Scale-up and Scale-out solution.

Scale-up means to use powerful single server to process the workload that fits within the server boundaries.

Scale-out is a different model which utilizes multiple processors as a single entity so a business can scale beyond the computer capacity of a single server. Table distribution is possible only in Scale-out scenario.

4. How are tables distributed across available index servers?

By default, new tables are distributed across available index servers using the round-robin approach; however it is also possible to explicitly specify that a table or partition is to be created on a specific index server.

5. Are there specific privileges required to perform a table redistribution?

You must have the system privilege RESOURCE ADMIN and at least the object privilege ALTER for all schemas involved to be able to run table redistribution.

To perform table distribution manually from one host to another, the following system privileges are required:

- System Privilege 'DATA ADMIN' or
- System Privilege 'CATALOG READ' and SQL privilege 'ALTER'

6. In which situation would you perform a table redistribution?

There are four different redistribution scenarios:

- Redistributing table before removing a host
- Redistributing table after adding a host
- Optimize current table distribution
- Optimize existing table partitioning

7. How is Landscape Redistribution done in SAP HANA?

Usually after adding or removing nodes, the Landscape Redistribution should be executed. Based on its configuration, it will suggest a new placement for tables and partitions in the landscape.

If the administrator confirms the plan, Landscape Redistribution will actually redistribute accordingly.

The redistribution operation also evaluates whether or not a partitioned table needs repartitioning based on its partitioning specification (that is, hash, round-robin, range and so on).

This is only relevant for column-store tables. System tables, temporary tables, and row-store tables are not considered.

8. How are tables distributed in SAP NetWeaver Business Warehouse on HANA system?

The BW master data is distributed across the slave servers using round-robin. This ensures a balanced utilization of the available CPU and memory resources.

Master Node - contains all row-store tables, ABAP system tables and general operational data like REPOSRC.

Worker Node - contains all BW master data, Cubes, DSO, PSA's.

9. What are the parameters for table redistribution?

Parameter	Description
global.ini --> [table_placement] --> same_num_partitions	If this parameter is set to true (default), all tables with the same name have the same number of level 1 partitions. The number of level 1 partitions is determined by the largest table within the group.
Indexserver.ini --> [table_redist] --> all_moves_physical	Default value true since SPS 09, for more info refer SAP Note 1880800 Administration guide
Indexserver.ini --> [table_redist] --> force_partnum_to_splitrule	Setting this parameter to true forces the execution of open table redistribution to change the number of level 1 partitions. For example, if a table has four level 1 partitions but should have three according to the table redistribution settings, the table redistribution process would not, by default, adjust this. Activating this option forces the adjustment.
global.ini --> [table_placement] --> method	The classification of row store tables is currently not taken into account. To ensure that these tables are always created on the master server and are moved there, you need to set the parameter method of [table_placement] of global.ini to "2".

Hint: All tables in BW that start with a '/' are always created on the slave servers. All other tables are only created on the master server if they are in one of the BW schema specified by parameter 'global.ini/table_placement/bw_schema'

10. Where can I find information specific to table distribution on BW on HANA?

Please refer to SAP Note [2143736](#).

11. Can I stop the table distribution process? If so how? And does it roll-back all changes?

There is a way to stop the table redistribution. Unfortunately there is no “stop” button. Open SAP HANA Studio --> Select Performance --> "Threads" Tab --> select the Thread with the name "LandscapeReorg Plan Exec" --> click right and "Cancel Operations"

The Table Redistribution will finish the actual move and split operation and end its processing. No Rollback is made.

12. Should I plan a maintenance window for table distribution?

As table distribution uses DDL operations locks are possible hence it is recommended to plan a maintenance window when performing this activity on a production system.

13. What can go wrong during the process and would a restore to saved config be an option or should I perform full database restore?

All performed operations (Table SPLIT/MERGE/MOVE) are transactional safe. Most problems come from missing privileges (alter privilege for schema) or locked tables.

It is recommended to do a full system backup before you start the table redistribution.

The “Save” and “Restore” option saves and restores the actual location for each table the executing user is able to “see”. This is not equivalent to a system restore and saves lot of effort and time.

14. During a restore of a saved configuration, there were some failed tables, how can I determine the exact cause of failure?

Every error will be logged to the trace files.

15. Where can I find information regarding table location?

You can find information about locations of tables in:

- Table Distribution viewer

Once table distribution editor is opened, search for a specific table to get the details of the table

- Table Runtime Information - Double click on a particular table to get the runtime information

16. What information is available in the table distribution editor/ table runtime information?

You can find the host in which the table is located and if it is partitioned, details about the partitions, size of the table, details regarding merge and if the table is loaded.

17. Is there a need for a specific backup strategy when table redistribution is planned?

Yes, it is highly recommended to perform a backup just before the table distribution is performed as well as immediately after the table distribution has been successfully completed.

18. Is a complete database recovery necessary to revert table distribution?

Not necessarily, you do not need to perform a database recovery to revert the table distribution if you would have saved the table distribution prior to the table redistribution.

- You can save the current table distribution before the redistribution as follows:
In the Administration editor, choose Landscape - Redistribution - Save the current table distribution by choosing Save.
- You may restore the table distribution from a previous point in time as follows:
In the Administration editor, choose Landscape - Redistribution - In the Executed Operations area, identify the operation that corresponds to the table distribution that you want to restore.

19. How can I save a table distribution plan and send it to SAP for review before actually executing the plan?

In order to save a plan, you need to perform these actions in SAP HANA SQL editor:

1. call `REORG_GENERATE(ALGORITHM_ID,");`
2. And within the same session execute, `select * from reorg_plan;`
3. Now you can export the results and could send it to SAP in case it needs to be reviewed.

20. What are the different algorithm IDs available and when they should be used?

There are six different algorithm IDs, each for a different purpose. Below is the list.

ALGORITHM ID	PURPOSE
1	Add host
2	Remove host
4	Save distribution
5	Restore table distribution
6	Optimize table distribution
7	optimize table partitioning

21. Is there a way to estimate how long the process might take before starting the distribution?

Unfortunately not. This is highly dependent on the infrastructure, the size and number of tables that have to be moved.

You can follow the progress by executing
`select IFNULL("STATUS", 'PENDING'), count(*) from REORG_STEPS where reorg_id=(SELECT MAX(REORG_ID) from REORG_OVERVIEW) group by "STATUS";`

22. Can I expect different execution times for the distribution on the pre-production system and production system?

Table distribution is always determined per system and thus might have different execution times on different systems because tables of different sizes might be involved. Additionally, the distribution optimization algorithm is non-predictive which means that you might even get a different distribution across the hosts even if the data is the same.

23. How can I migrate from an <m> node environment to an <n> node environment?

Yes, migration from m to n nodes is possible (where m is greater than n). The same process applies for migration from multi-node to single-node.

Refer SAP Note [2093572](#) for details.

24. Is there any tool designed to help in table distribution?

Yes, there is a more advanced tool that could be used for data distribution known as "Data Distribution optimizer"

The main purpose of the Data Distribution Optimizer (DDO) is to support the HANA administrator in managing the distribution of tables and optimize the allocation of HANA Memory in a HANA Scale-Out landscape.

More information in SAP Note [2092669](#)

25. Where can I find additional information on Data Distribution Optimizer (DDO)?

General information and documentation on how to install and usage of the tool can be found [here](#).

26. Under which scenarios can I use Data Distribution Optimizer (DDO)?

The SAP HANA Data Warehousing Foundation provides specific data management tools, to support large scale SAP HANA use cases like:

- Data Warehousing: SAP BW powered by SAP HANA, SAP HANA native, mixed scenarios
- Multiple applications on one SAP HANA database
- Data Temperature Management (Hot- to Warm- and Cold-Store including Bi-Directional Data Movement) based on Data Movement Rules

27. Is there any dependency between DDO version and SAP HANA revision?

Yes, below is the dependency:

- HANA DWF 1.0 SP00: DDO >= HANA Rev. 83
- HANA DWF 1.0 SP01: DDO >= HANA Rev. 96
- HANA DWF 1.0 SP02: DDO >= HANA Rev. 101
- HANA DWF 1.0 SP03: DDO >= HANA Rev. 110
- HANA DWF 1.0 SP04: DDO >= HANA Rev. 120

28. How is the content of DDO delivered?

The content of DDO is delivered within 2 delivery units.

- sap.com/HCO_HDM
- sap.com/HCO_HDM_DDO

You have to take care for the import order of the DUs. HCO_HDM.tgz has to be imported before HCO_HDM_DDO.tgz

29. How can the progress of table distribution activities be monitored?

Table distribution activities can be monitored in the following ways:

Tables	SQL statement (SAP Note 1969700)	Details
M_SERVICE_THREADS M_SERVICE_THREAD_SAMPLES HOST_SERVICE_THREAD_SAMPLES	SQL: "HANA_Threads_CurrentThreads" SQL: "HANA_Threads_ThreadSamples_FilterAndAggregation" SQL: "HANA_Threads_ThreadSamples_AggregationPerTimeSlice"	Table distribution related thread information can provide insight which tables are currently distributed and what kind of detailed activity is executed. See SAP Note 2114710 for more information.
M_JOB_PROGRESS	SQL: "HANA_Jobs_JobProgress"	Starting with SAP HANA SPS 10 the view M_JOB_PROGRESS contains information for the current distribution related activities (JOB_NAME = 'Table Redistribution Execute'. 'Table Redistribution Generate').

30. How is the data distributed during optimize table distribution?

During an optimized table distribution the data is moved physically (default) and not via delta merge.

You can increase the trace level to find the relevant command executed during reorg

```
global.ini -> [trace] -> landscapereorg = 'debug'
```

```
[32196]{-1}[-1/-1] 2017-01-17 08:54:42.318198 i LandscapeReorg LandscapeReorgPlanExecute.cpp(00062)
: --> ReorgPlanExecuter ID: 32196 started
[12013]{342860}[-1/-1] 2017-01-17 08:54:42.318882 d LandscapeReorg LandscapeReorgThread.cpp(00165)
: Start worker for reorg_step 3, reorg_group 4, SQL: ALTER TABLE "LHTEST2"."ATESTER" MOVE TO
LOCATION 'host2:30003' PHYSICAL
```

Please make sure that you have increased the number of trace files and size of the files as increasing the trace levels might also increase the amount of data written into the traces.

31. Which values are allowed for the LOCATION in TABLE_PLACEMENT?

The following values are allowed for LOCATION in TABLE_PLACEMENT

- master
- slave or slaves
- all
- location groups -> defined in global.ini [table_placement] with Parametername = Name of location group
- In SAP HANA 2.0, the name of the worker group is stored in LOCATION column in system view TABLE_PLACEMENT refer [Admin Guide](#) section: Configuration Using Worker Groups

- all locations like mentioned above can have suffix **#n** where $n = 1 \dots N$.
- all other entries are not allowed and will be ignored

32. How can I gather information on table groups, number of records and other information relevant to table distribution via SQL?

You can use SQL: *"HANA_Tables_ColumnStore_TableHostMapping"* which is now enhanced to provide better overview and analysis capabilities with regards to table distribution.

33. Are there any special pre-requisites to be considered with regards to the disk space while planning SAP HANA table redistribution?

Yes, during a table redistribution process, the system needs temporarily more disk space for data than calculated in the sizing phase. As a rule-of-thumb, the size of the disk space for data should be doubled during redistribution to avoid running in disk-full situations.

Refer, SAP HANA Storage whitepaper from SAP Note [1900823](#).

Keywords

Table distribution, redistribution, RSDU_TABLE_CONSISTENCY, DDO, Data Distribution Optimizer, DWF, privilege, RESOURCE ADMIN, DATA ADMIN, CATALOG READ, ALTER, SAP NetWeaver Business Warehouse, BW, table_placement, same_num_partitions, table_redist, all_moves_physical, force_partnum_to_splitrule, method, roll-back, stop, maintenance window, save a table distribution plan, REORG_GENERATE, Add host, Remove host, Save distribution, Restore table distribution, Optimize table distribution, optimize table partitioning

Products

SAP HANA, platform edition all versions

This document refers to

SAP Note/KBA	Title
2093572	SAP HANA Migration from Multi-Node to Single-Node
	SAP HANA Administration Guide
	SAP HANA SQL and System Views Reference Guide
	SAP HANA Table Distribution for BW on HANA Scale-out Systems