

Migration Sizing for SAP BW on HANA

Version 2.1.6 – June 2015



SAP and SAP Customer Internal

DISCLAIMER

Sizing recommendations apply for certified hardware only. Please contact your hardware vendor for suitable hardware configuration.

Note that HANA is constantly being optimized. This might have impact on sizing recommendations, which will be reflected in this document. Therefore, check for the latest version of this document, the sizing software, and related notes.

The sizing guideline in this document refers to SAP BW on HANA only. Additional applications running on top of HANA are not covered in this document or in the report that is described. Please refer to application specific sizing guidelines for details in that case.

Actual memory requirements in HANA strongly depend on operational conditions which cannot be taken into account without analyzing the HANA system itself. The sizing results derived from this guideline are therefore to be understood as estimates based on best practices and experiences.

BW Sizing Approaches

„Green-field“ Sizing

- For customers starting „from scratch“
- No existing SAP BW system that could be analyzed
- Rough idea of data model available (largest InfoCubes / DSOs)
- Recommended Sizing Method: SAP QuickSizer for BW on HANA (<http://service.sap.com/quicksizer>)

Migration Sizing

- For customers who have an existing SAP BW running on a “traditional” database platform and plan migrating to HANA
- ABAP based BW on HANA Sizing report (SAP note 1736976): database platform independent, leverages BW semantic information, high level of accuracy

STRONGLY RECOMMENDED sizing approach for all SAP BW on HANA migration projects

Purpose and Pre-Requisites

Mission statement:

Provide an easy to use, source database independent tool to predict resource requirements for a BW system on a SAP HANA database.

The sizing report /SDF/HANA_BW_SIZING is a convenient method to estimate the memory requirements of a BW system after migration to SAP HANA. Major advantages of the ABAP report:

- Easy to deploy and use – no DB administrator required
- Versatile parameterization – control resource consumption and speed
- Independent of source database specific compression and data representation
- Considers user defined future growth and leverages BW semantic information

The report requires ST-PI 2008_1_7xx SP8 and SAP NetWeaver BW 7.0 SP 1 or higher. Updated versions can be obtained from note 1736976.

Important note for DB2 on iSeries (AS/400): Please implement note 1677958!

/SDF/HANA_BW_SIZING: How it works

Main characteristics of /SDF/HANA_BW_SIZING:

- Report obtains list of tables from ABAP dictionary (table DD02L)
- For each table, a **sample data set** is selected, and **data field contents** are analyzed.
NOTE: MaxDB does not support data sampling. Database statistics data are used for sizing in this case. Expect longer report execution times or less accurate results for this platform.
- Calculates table sizes based on sampled data in ABAP memory and total record count → **source database compression** is taken into account automatically
- Automatic uplift for non-unicode source systems
- Separation of tables in row store and column store
- Computation of overall SAP HANA memory requirements (including runtime memory) with **table type specific compression factors**.
- Parallel processing of tables to speed up processing time (degree of parallelism can be adapted to control utilization of DIA work processes)
- Computation of future resource requirements based on relative or absolute yearly growth (on demand)
- Subset sizing (either for a complete BW system containing specified subsystem, or for the set of specified objects as isolated add-on for an already existing BW)
- Data that is stored (or supposed to be stored) in Dynamic Tiering is taken into account

/SDF/HANA_BW_SIZING: Optimized Sizing Procedure

Table Type Specific Compression Factors

- Analysis of various PoC customer and internal test systems has revealed different average compression factors for InfoCubes, DSOs, PSAs, ChangeLogs, etc.
- Overall sizing result reflects data distribution across different table types much better now.

General Sizing Remarks

- The report determines for which tables an uplift needs to be considered to provide memory for temporary runtime objects, delta indexes, and merge results. For these tables, the table size is multiplied by a factor 2.
- Processing of row store also requires additional runtime memory. The runtime uplift is as large as the row store itself, however, it is capped by 100 GB.
- On each server node a fixed amount of memory has to be taken into account for
 - HANA statistics service, HANA name service (10 GB)
 - HANA system caches (40 GB)



/SDF/HANA_BW_SIZING: Recommendations and Remarks

Before you run the report /SDF/HANA_BW_SIZING, please consider the following:

- **Selection of parallel degree and availability of DIA work processes**
In order to run the report in parallel, make sure that for the specified degree of parallelism a corresponding number of dialog work process is available.
- **Result accuracy**
Due to the nature of data sampling, two consecutive executions of the report might deliver results which slightly differ from each other. However, the difference will not be significant enough to have an impact on sizing and can therefore be ignored.
We strongly recommend running the report on up-to-date database statistics!
- **Precision Settings**
The default precision ("Low") of the report results in a sampling rate which is large enough to calculate reliable table sizes. Tests have shown that "Medium" precision setting delivers results which are very close to low precision, but results in higher runtime. A setting of "High" is only recommended for very small systems (< 500 GB source DB size).

Report Selection Screen: Overview

Determine DB size relevant for BW on HANA Sizing

General

Store output in file ☒ File name

Number of parallel procs

Suppress tables < 1MB ☒

Unload inactive tables ☒

Target Release 7.40 ☒

Precision


High ☐

Medium ☐

Low ☒

Scenario / Subset Selection

Use system subset only ☐

List of top level InfoProv. to 

Subset for existing BW system ☐

New BW system using this subset ☒

Exclude listed objects ☐

Future Growth Simulation

Consider Growth ☐

Number of years


Relative growth (in %) ☒

Absolute growth (in GB) ☐

Growth value


Non-active Data

Consider non-active data ☒

WARM write-optimized DSOs to 

Dynamic Tiering

Use Dynamic Tiering ☐

Objects in Dynamic Tiering to 

Memory Configuration

244 GB ☐ 256 GB ☐ 512 GB ☒ 1024 GB ☒ 1536 GB ☒ 2048 GB ☒ 3072 GB ☒

Report Selection Screen: General Section

General	
Store output in file	<input checked="" type="checkbox"/> File name HANA_Sizing.txt
Number of parallel procs	20
Suppress tables < 1MB	<input checked="" type="checkbox"/>
Unload inactive tables	<input checked="" type="checkbox"/>
Target Release 7.40	<input checked="" type="checkbox"/>

Check flag „Store output file“ to save output to specified file in the DIR_HOME work directory of the application server (transaction AL11). File name MUST begin with „HANA_Sizing“.

Specify number of parallel processes to analyze tables. Make sure you have enough free work processes.

Tables with less than 1MB size can be suppressed in the output (but are still counted for sizing).

Check this box if your target release is 7.40 or higher. List of row store tables in earlier releases differs!

For execution on HANA only: If a table was completely unloaded, unload it again after analysis

Report Selection Screen: Precision, Scenario / Subset

The screenshot shows the 'Report Selection Screen' with two main sections: 'Precision' and 'Scenario / Subset Selection'. The 'Precision' section has three radio buttons: 'High', 'Medium', and 'Low', with 'Low' selected. The 'Scenario / Subset Selection' section has a checkbox 'Use system subset only' (unchecked), a text input field for 'List of top level InfoProv.' followed by 'to' and another text input field with a search icon, and three radio buttons: 'Subset for existing BW system' (unchecked), 'New BW system using this subset' (selected), and 'Exclude listed objects' (unchecked). Four yellow callout boxes provide additional information: 1. A callout from the 'Low' radio button explains that higher precision results in higher sampling rates and runtimes, while low precision delivers reliable results. 2. A callout from the 'Use system subset only' checkbox explains that checking this flag enables Scenario / Subset Selection functionality. 3. A callout from the 'List of top level InfoProv.' input field explains that users should enter technical names of top level InfoProviders, which are added automatically, and that wild cards are permitted. 4. A callout from the 'New BW system using this subset' radio button explains that the specified scenario can be used for a new BW system or added to an existing one, or excluded from sizing calculation.

Precision

High ☐

Medium ☐

Low ☒

Scenario / Subset Selection

Use system subset only ☐

List of top level InfoProv. to

Subset for existing BW system ☐

New BW system using this subset ☒

Exclude listed objects ☐

Select sampling precision. Higher precision results in higher sampling rates and higher runtimes. Usually, low precision delivers sufficiently reliable results.

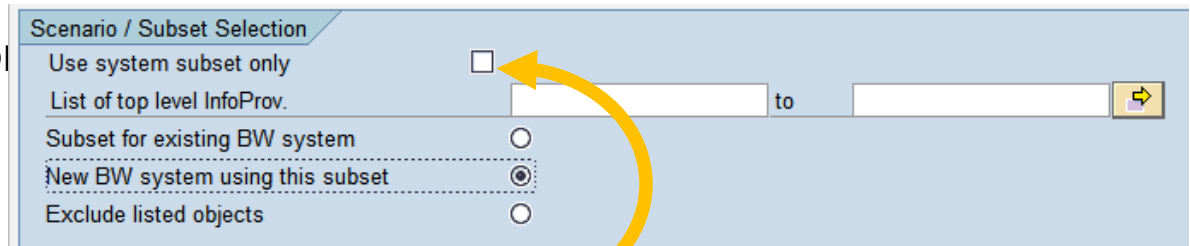
Check this flag to enable Scenario / Subset Selection functionality.

Enter list of **technical names** of top level InfoProviders which are contained in your scenario. Dependent objects will be added automatically. Wild cards are permitted.

Specified scenario can either be used for new BW system or can be added to an already existing one. Alternatively, you can exclude the specified list from sizing calculation.


Subset / Scenario Sizing

Subset / Scenario Selection



Scenario / Subset Selection

Use system subset only ☐

List of top level InfoProv. to 

Subset for existing BW system ☐

New BW system using this subset ☒

Exclude listed objects ☐

Switch on by checking flag “Use system subset only”.

You can determine the HANA size for parts of your source system by identifying all top level InfoProviders that need to be part of the scenario, or that you would like to exclude, respectively.

For inclusion scenarios, the report automatically considers all sizing relevant objects that depend on the provided list (e.g. all InfoObjects of an InfoCube, or all other InfoProviders in the downward data flow).

Object types that can be specified:

- InfoCubes
- DataStore Objects
- MultiProviders
- Semantically Partitioned Objects
- InfoObjects
- Hybrid Providers
- Open Hub Destinations
- InfoSets

Subset / Scenario Sizing (cont.)

Select type of subset:

Scenario / Subset Selection

☐ Use system subset only

List of top level InfoProv. to

☐ Subset for existing BW system

☒ New BW system using this subset

☐ Exclude listed objects

Subset for other BW System:

- Report includes objects specified in list AND all dependent objects, but no other BW related system tables.
- Select this option if you plan to move the scenario to an already existing BW system and you want to determine the additional amount of memory that the scenario will require in the target system

BW System using this subset:


- Report includes objects specified in list, all dependent objects and all BW related system tables
- Select this option if you plan for a new BW system that only contains the specified scenario.

Exclude listed objects:

- Report includes all tables except those belonging to the objects specified in list. Dependent objects will NOT be excluded.
- Select this option if you plan to move InfoCubes / DSO objects to Near Line Storage (NLS).

Subset / Scenario Sizing: Subset Types

Source BW System



- Business Data (Cubes, DSO, PSA, ...)
- BW System Tables
- Top Level Objects Specified in List
- Derived Dependent Objects (InfoObjects, etc.)

Subset for other BW System:



BW System using this subset:



Exclude listed objects:



Report Selection Screen: Future Growth, Non-active Data

The screenshot shows two sections of the SAP Report Selection Screen. The first section, 'Future Growth Simulation', contains a 'Consider Growth' checkbox (checked), a 'Number of years' input field (3), and radio buttons for 'Relative growth (in %)' (selected) and 'Absolute growth (in GB)'. Below these is a 'Growth value' input field (25,0). The second section, 'Non-active Data', contains a 'Consider non-active data' checkbox (checked) and a text input field for 'WARM write-optimized DSOs' with the value 'YTD*'. Callout boxes provide instructions for each field.

Future Growth Simulation

- ☒ Consider Growth: Check this flag to enable Future Growth calculation.
- Number of years: 3: Specify number of years to be considered for growth.
- Relative growth (in %) ☒ / Absolute growth (in GB) ☐: Decide if you want to specify yearly growth as an absolute value or relative to the current size.
- Growth value: 25,0: Enter value for relative or absolute growth.

Non-active Data

- ☒ Consider non-active data: Check this flag to Non-active Data calculation.
- WARM write-optimized DSOs: YTD*: Enter list of **technical names** of write-optimized DataStore Objects to be considered non-active. Wild cards are permitted.

Non-Active Data

hot	Data is read/written frequently In Memory, additional memory required for temporary objects (merge, intermediate results, etc.)	
warm	Infrequent access – no need to keep in memory all the time On disk, loaded to memory only on demand, good candidate for displacement if memory runs short	Relevant for non-active data concept
cold	Sporadic access – not stored in HANA DB Restricted to NLS capabilities	

Objects in SAP HANA can be marked as “non-active”

Tables / partitions belonging to these objects are ...

- displaced from RAM (per column) with higher priority in case of RAM shortage (but only then) or when a cleanup is triggered
- loaded into RAM only when accessed (read access, merges) as usual

SAP BW automatically marks all PSA tables and write-optimized DataStore Objects for early displacement, so no extra maintenance or tuning is necessary. All other BW objects are treated as usual.

To avoid aggressive downsizing, you are required to specify the technical names of all w/o DSO objects which are accessed only very infrequently. All w/o DSO objects which are not contained in the list will not be considered as non-active by the report. Wildcards may be used to specify name ranges.

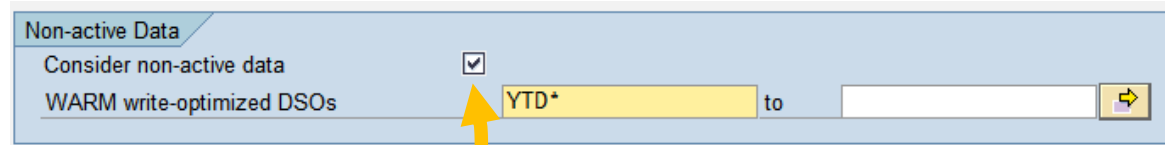
For more detail please refer to SAP note 1767880.

Requirements: SAP BW 7.30 SP08 and SAP HANA 1.0 SPS05

Non-Active Data: Recommendations and Remarks

Default settings:

„Non-active data“ switched OFF

A screenshot of the SAP 'Non-active Data' configuration interface. It features a blue header bar with the title 'Non-active Data'. Below the header, there are two main sections. The first section, 'Consider non-active data', contains a checkbox that is currently checked. A yellow arrow points from this checkbox to the text 'Switch on by checking flag „Consider non-active data“'. The second section, 'WARM write-optimized DSOs', contains a text input field with 'YTD*' entered, followed by a 'to' label and an empty text input field. A yellow arrow points from the text 'Enter list of write-optimized DataStore Objects' to the 'YTD*' input field.

Switch on by checking flag „Consider non-active data“ to use BW default settings (no additional maintenance in BW needed).

Enter list of write-optimized DataStore Objects that are considered to be “warm”. Wildcards are allowed.


- **NOTE: please enter technical object names, NOT the names of the database tables!**
- **Please refer to SAP note 1767880 for further information on the non-active data concept**

Avoid too aggressive downsizing: do not mark objects as non-active if they are accessed frequently!

Report Selection Screen: Dynamic Tiering, Memory Conf.

Dynamic Tiering

Use Dynamic Tiering ☐

Objects in Dynamic Tiering to 

Check this flag to enable Dynamic Tiering.

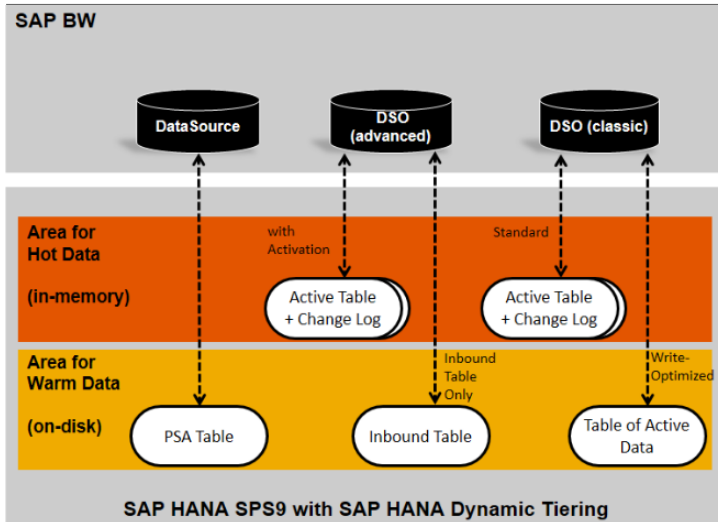
Enter list of **technical names** of objects which are supposed to be stored in Dynamic Tiering. Wild cards are permitted.

Memory Configuration

244 GB	<input type="checkbox"/>	256 GB	<input type="checkbox"/>	512 GB	<input checked="" type="checkbox"/>	1024 GB	<input checked="" type="checkbox"/>	1536 GB	<input checked="" type="checkbox"/>	2048 GB	<input checked="" type="checkbox"/>	3072 GB	<input checked="" type="checkbox"/>
--------	--------------------------	--------	--------------------------	--------	-------------------------------------	---------	-------------------------------------	---------	-------------------------------------	---------	-------------------------------------	---------	-------------------------------------

Select list of memory configuration options that the report is supposed to consider. Out of this list you will get a minimum configuration and a recommended configuration, depending on the sizing results.

Dynamic Tiering



Tables in Dynamic Tiering:

- Definition of table resides in HANA
- Data actually stored in Dynamic Tier (on-disk based layer)
- Size analysis of Dynamic Tiering tables provide input for Dynamic Tiering sizing (see SAP note 2086829)

The screenshot shows the 'Dynamic Tiering' selection screen. It has a tab labeled 'Dynamic Tiering'. Below the tab, there is a checkbox labeled 'Use Dynamic Tiering' which is checked. Below that is a text field labeled 'Objects in Dynamic Tiering' containing 'DT*' followed by a 'to' label and an empty text field. A yellow arrow button is on the right.

On the selection screen specify list of technical names of objects that you want to move to Dynamic Tier. The sizes of these tables are not taken into account for the overall HANA memory requirements.

Tables in on-disk layer will be displayed separately at the end of the sizing report output.

Requirements: SAP BW 7.40 SPS08, SAP HANA 1.0 SPS09

Please refer to SAP note 1983178 for more information

Allowed Memory Configurations

Memory Configuration													
244 GB	<input type="checkbox"/>	256 GB	<input type="checkbox"/>	512 GB	<input checked="" type="checkbox"/>	1024 GB	<input checked="" type="checkbox"/>	1536 GB	<input checked="" type="checkbox"/>	2048 GB	<input checked="" type="checkbox"/>	3072 GB	<input checked="" type="checkbox"/>

Select only those memory configurations that your hardware partner offers

Report calculates:

- Minimum configuration which can store all master node data in a single node (respecting future growth if selected)
- Recommended configuration which balances the number of server nodes and the amount of memory per node (this is not necessarily a landscape with the largest memory configuration available).
- Scale-out landscapes are recommended as long as they do not consist of more than a certain number of nodes (e.g. 8 nodes).
- Recommendations with an expected utilization of the memory on any node close to its capacity are avoided.

Sizing Report Results: Summary

The result screen of the Sizing Report contains very detailed information on the source system and its corresponding HANA sizing:

- Overview of the size of the source DB, based on sampled data. Here the sizes are based on an ABAP internal representation of the data.
- A summary of the resources that this system at minimum requires when running on a HANA database
- Extrapolated resource requirements in future, based on specified growth rate
- Overall sizes for row and column store tables, both for master and worker nodes.
- Detailed size information per table, including estimated ABAP size based on sample, derived HANA size, and record count. This detailed list can be used to directly determine database tables which should be targeted by housekeeping measures.

Sizing Report Results: How to Read

Section „SOURCE DB CONTENTS“:

- Summary on tables found in source DB
- Size figures reflect sizes of tables as if they were loaded into ABAP internal tables (“ABAP size”) – based on sample set of records read from each table.

Section „MINIMUM SIZING RECOMMENDATION “:

- Minimum amount of memory required to operate given system on a HANA database. This amount includes space for storing database tables as well as space for runtime objects and needs to be physically available on a target landscape (“HANA size”).

Section „SIZING DETAILS“:

- Memory share by master row store, master column store and worker column store, as well as data and runtime object shares. Important to identify potential memory shortages on a per-node basis (“HANA size”).

Section „TABLE DETAILS“:

- Detailed list of memory consumption figures per database table, including non-active and secondary index share (if applicable), but **without** runtime uplift. Important to identify impact of largest tables and tables subject to house keeping (“ABAP” and “HANA size”).

Sizing Report Results: Source DB Contents

Determine DB size relevant for BW on HANA Sizing

SOURCE DB CONTENTS

ABAP Size Row Store:	48.8 GB.	No. of tables:	2257
ABAP Size Column Store:	4747.8 GB.	No. of tables:	33237
Thereof:			
InfoCubes	3733.5 GB		3455
Std. DSO	702.4 GB		233
Change logs	0.4 GB		111
w/o DSO	1588.2 GB		56
PSA	16.8 GB		741
Master Data	166.4 GB		6456
Customer Tables	0.0 GB		0
Others	128.3 GB		22185
TOTAL:			35494
Tables excluded from sizing (aggregates, etc.)			0

NOTE: This section describes the contents of the **source database** as derived from data sampling. Figures refer to table **sizes in ABAP** and can vary from those reported by the source DB. Also note that indexes, temporary space, etc. of the source DB are **NOT** reflected here!

Shares of different table types that go to column stores (sizes in **ABAP**!)

Number of tables excluded from sizing (e.g. aggregates, dimension tables, etc.)

Total number and size of all tables in source DB (as per data sampling) with shares for row and column store.

Sizing Report Results: Minimum Sizing Recommendations

MINIMUM SIZING RECOMMENDATION - CURRENT

	minimum	recomm.
Phys. memory per node:	512 GB	1024 GB
Memory Requirement (Minimum Total):	2048 GB	1948 GB
Disk Space Requirement - data (Minimum Total):	1799 GB	1799 GB
Disk Space Requirement - logs (Minimum Total):	512 GB	512 GB
Number of Nodes incl. master (Minimum Total):	5	3

NOTE:

- Please carefully read documentation attached to SAP NOTE 1736976 for a detailed description of the sizing procedure and its results!
- Disk space requirement calculation no longer includes space for backups, dumps, etc. This space has to be provided on additional disk volumes. Guidance for sizing these additional requirements is described within the documentation attached to SAP note 1736976.

This section describes the minimum resource requirements for the system to run on a **HANA database**. The memory requirements include space for the tables, indexes (where applicable), and space for temporary database objects (e.g. intermediate query results).

Total memory, disk space, and number of nodes required for a system landscape with **minimum amount of memory per node**

Total memory, disk space, and number of nodes required for a system landscape with **recommended amount of memory per node**

Minimum vs. Recommended: HANA prefers a small set of nodes with less memory over a single node with more memory for cost effectiveness / flexibility reasons. Minimum memory amount is determined by requirements for master node. For Scale-out landscapes with more than 8 nodes the next larger server configuration is used (if available).

Sizing Report Results: Sizing Details

SIZING DETAILS				
=====				
(For 1024 GB nodes)	data [GB]	total [GB] incl. tmp.	total [GB] (non-act.)	util.
MASTER:				
Row Store (incl. ind.)	100	150	150	39 %
Master Column Store	93	186	186	
Caches / Services	50	50	50	
TOTAL (MASTER)	243	386	386	
SLAVES:				
Slave Column Store	732	1463	1462	79 %
Caches / Services	100	100	100	
TOTAL (SLAVES)	832	1563	1562	
<hr/>				
TOTAL (All servers)	1075	1949	1948	

This section describes the shares of memory consumption for row and column store tables on master and worker nodes of the **HANA database**. The figures include space for services and caches of HANA on each server node instance. The utilization figures should help to determine how much memory is utilized on each node.

Memory requirements for master node. If more than 400GB is needed for master, you should analyze table details to determine housekeeping measures!

Memory requirements for all worker nodes. Right column shows memory consumption with non-active data considered, column in the middle shows requirements without considering non-active data.

Sizing Report Results: Future Growth Extrapolation

SIZING RECOMMENDATION - FUTURE GROWTH			
=====			
Time considered for growth:	3 years		
Assumed yearly growth rate:	30.0 %		
	minimum	recomm.	
Phys. memory per node:	512 GB	1024 GB	
Memory Requirement (after 1 year):	2555 GB	2455 GB	
Disk Space Requirement - data (after 1 year):	2257 GB	2257 GB	
Disk Space Requirement - logs (after 1 year):	512 GB	512 GB	
Number of Nodes incl. master (after 1 year):	6	4	
Memory Requirement (after 2 years):	3062 GB	2912 GB	
Disk Space Requirement - data (after 2 years):	2715 GB	2715 GB	
Disk Space Requirement - logs (after 2 years):	512 GB	512 GB	
Number of Nodes incl. master (after 2 years):	7	4	
Memory Requirement (after 3 years):	3569 GB	3369 GB	
Disk Space Requirement - data (after 3 years):	3173 GB	3173 GB	
Disk Space Requirement - logs (after 3 years):	512 GB	512 GB	
Number of Nodes incl. master (after 3 years):	8	4	

This section shows an estimation on the future growth of the HANA system, based on the specified growth rate. Anticipated growth rate of source system should be specified. Growth is then split among system tables (10%) and tables with business data (90%).

Total memory, disk space, and number of nodes required for a system landscape with **minimum / recommended amount of memory per node**, year by year, based on specified growth rate.

Sizing Report Results: System Information

SYSTEM INFORMATION

```
=====
Report version :      2.1.6
Execution date and time: 13.07.2015 12:41:09
Runtime:             00:09:43
Parallel degree:     20
Precision settings:   L
Consider non-active data: YES
Use system subset:    NO
RDBMS:               ORACLE
Operating System:     Linux
DB hostname:          ld8508
SID:                  PTP
Release:              SAP_BW 730 SP0010
Use 7.4 row store set: YES
```

This section shows information on the system on which the sizing report was executed, runtime information and the most important parameters of the report.

Number of row store tables was significantly reduced with release 7.40. Report considers new list by default, unless flag on selection screen is removed (see page 9).

Sizing Report Results: Table Details

NON-ACTIVE DATA =====

All PSA and Change Log tables are considered 'warm' by default.

TABLE DETAILS =====

MASTER NODE

Type	Table Name	estimated ABAP Size [MB]	estimated Ind. Size [MB]	estimated HANA Size [MB]	rec count [million]
Row	RSBATCHDATA	4573.7	4573.7	9147.5	1.636
Row	REPOLoad	3899.3	3899.3	7798.6	0.130
Row	RSODSO_RUNTIME	2280.2	2280.2	4560.5	0.950
Row	RSBMREQ_DTP	2213.4	2213.4	4426.7	42.022
Row	DDPRS	1287.1	1287.1	2574.2	14.956

Row	RSSTATMANPART	0.5	0.5	1.0	0.003
Row	CWBCIDATAOBJ	0.5	0.5	1.0	0.000
more	...				

Type	Table Name	estimated ABAP Size [MB]	estimated HANA Size [MB]	rec count [million]
Col	ZFHPSDBLC02	25048.6	17891.8	23.281
Col	ZDHPSDBLC033	10930.2	7807.3	0.700
Col	ZDHPSDBLC043	6891.7	4922.7	0.441
Col	ZFHPSDORC01	6690.0	4778.6	36.154
Col	ZEHPSDBLC03	5841.2	4172.3	43.418
Col	ZPSPROT_DETAIL	5442.0	3887.1	34.199
Col	ZDHPSDBLC023	5115.0	3653.5	0.327
Col	ZFHPSDDLC02	4600.1	3285.8	43.931
Col	WBTESTCROSS	1.0	0.7	0.020
Col	SSYNTAXSTRUCTURE	1.0	0.7	0.001

SLAVE NODES

Type	Table Name	estimated ABAP Size [MB]	estimated HANA SIZE [MB]	nonact. [MB]	na. [million]	rec count [million]
Col	/BIC/FZCRM_SAL2	1545907.5	220843.9	220843.9		3148.711
Col	/B49/ABENCHD0340	708484.8	141697.0	141697.0		1110.000
Col	/BIC/FHPSDORC02	739521.3	105645.9	105645.9		183.755
Col	/BIC/FHPSDDLC02	631602.5	90228.9	90228.9		463.135
Col	/BIC/FHPSDBLC03	615562.1	87937.4	87937.4		466.418

List of write-opt. DSOs considered as "warm" (if specified in selection screen)

Details of **row store tables on master node**: estimated ABAP size (based on sampling), estimated size of HANA indexes, estimated HANA size of and indexes (based on table type specific compression), and record count.

Details of **column store tables on master node**: estimated ABAP size (based on sampling), estimated HANA size of table (based on table type specific compression), and record count.

Details of **column store tables on worker nodes**: estimated ABAP size (based on sampling), estimated HANA size of table (based on table type specific compression), estimated HANA size after considering non-active data, and record count.

Tables with non-active data are marked as '(w)'.

Sizing Report Results: Table Attributes

Co1	/BIC/ATBREM0000	(W)	472.1	59.0	11.8	(M)	0.992
Co1	/BIC/B0000370000	(C)	469.3	58.7	58.7	(w)	1.021
Co1	/BIC/B0000149000	(P)	439.2	54.9	54.9	(w)	0.999
Co1	/BIC/B0000282000	(P)	438.3	54.8	54.8	(w)	0.997
Co1	/BIC/B0000159000	(P)	401.7	50.2	50.2	(w)	1.950
Co1	/BIC/B0000159000	(P)	401.7	50.2	50.2	(w)	1.950

Additional classification of tables:

- (W): write-optimized DataStore Object
- (C): Change Log of standard DSO
- (P): PSA table

“Non-active data” Flag:

- (M): write-optimized DataStore Object, marked as non-active on selection screen
- (w): PSA table or Change Log, automatically considered non-active.

Sizing Report Results: Dynamic Tiering

DYNAMIC TIERING: List of Tables

Type	Table Name	ABAP Size [MB]	HANA Size [MB]	IQ Size [MB]	rec count [million]
ES	/BIC/AYTDWD00100 (W)	5075.1	634.4	0	11.493
ES	/BIC/AYTDWD00200 (W)	439.4	54.9	0	0.995
ES	/BIC/B00000150000 (P)	24.0	3.0	0	0.121
more ...					
TOTALS		5538.5	692.3	0.0	12.609

This section shows information on those tables which are supposed to go to a Dynamic Tier, as specified on the selection screen.

Details of **tables in Dynamic Tier**: estimated ABAP size (based on sampling), estimated HANA size of table (based on table type specific compression) which is excluded in HANA memory sizing, size of tables already stored in Dynamic Tier, and record count.

See page 17 for more information on Dynamic Tiering!

*) Only relevant if sizing report was executed on a HANA based system with Dynamic Tier attached.

Special Sizing Situations

POSDM

- Migration to HANA might require a migration of the TLOGS to the TLOGF table. Substantial increase of the data volume is possible. Please refer to the “SAP POS DM Sizing Guidelines” on Service Marketplace: https://websmp204.sap-ag.de/~form/handler?_APP=00200682500000002672&_EVENT=DISPLAY&_SCENARIO=01100035870000000122&_HIER_KEY=501100035870000020093&_HIER_KEY=601100035870000253253&_HIER_KEY=601100035870000253358&_HIER_KEY=701100035871000577571&_

BPC

- Business Planning and Consolidation mostly affects the sizing of the NetWeaver application server. All BPC related tables are taken into account by the sizing report. For more details on BPC sizing please refer to the BPC Sizing Guide on Service Marketplace: https://websmp208.sap-ag.de/~sapidb/011000358700000250732010E/BPC_10_0_NW_HANAV3_2.pdf

BW on HANA Sizing: Scale Out

If a single HANA node cannot accommodate data due to limited memory, data has to be distributed across multiple nodes (scale-out).

- **Symmetric solution:** 1 master node, n worker nodes, all on identical hardware
- **Master node** will handle system load and transactional load: ABAP system tables and general operational data of the BW are stored on the master node. Note that this includes both **column store** and **row store** data. DDL statements are executed on this node, global locks are acquired here.
- **worker nodes** will handle OLAP queries as well as loading/staging/activation/merging. BW data (master data + cubes/DSOs/PSAs – all tables that have been generated by BW) is distributed across the **column stores** of all workers. This ensures a balanced utilization of the available CPU and memory resources. **Note that in general no column store data (except system tables) may be stored on the master node! For exceptions see SAP note 1908075.**
- **Easy extensibility:** When a new worker is added, BW data can quickly be reorganized to fit the new system. More information here: http://help.sap.com/hana/hana_db_part_en.pdf
- Optional: **stand-by node(s)** can take over in case of node failure
- For more detailed information please refer to document “HANA_BW_Sale_Out” attached to note 1736976.

BW on HANA Sizing: Scale Out

Sizing Report delivers necessary information at a glance:

MINIMUM SIZING RECOMMENDATION - CURRENT				
Phys. memory per node:		minimum 512 GB	recomm. 1024 GB	
Memory Requirement (Minimum Total):		2048 GB	1948 GB	
Disk Space Requirement - data (Minimum Total):		1799 GB	1799 GB	
Disk Space Requirement - logs (Minimum Total):		512 GB	512 GB	
Number of Nodes incl. master (Minimum Total):		5	3	
NOTE: - Please carefully read documentation attached to SAP NOTE 1736976 for a detailed description of the sizing procedure and its results! - Disk space requirement calculation no longer includes space for backups dumps, etc. This space has to be provided on additional disk volumes. The guidance for sizing these additional requirements is described within the documentation attached to SAP Note 1736976.				
SIZING DETAILS				
(For 1024 GB nodes)	data [GB]	total [GB] incl. tmp.	total [GB] (non-act.)	util.
MASTER:				
Row Store (incl. ind.)	100	150	150	
Master Column Store	93	186	186	
Caches / Services	50	50	50	
TOTAL (MASTER)	243	386	386	39 %
SLAVES:				
Slave Column Store	732	1463	1462	
Caches / Services	100	100	100	
TOTAL (SLAVES)	832	1563	1562	79 %
TOTAL (All Servers)	1075	1949	1948	

Total memory including data, runtime, services on each node

Number of nodes (workers + 1 master for scale-out configurations)

Detailed split of memory requirements for data, runtime objects and services with / without impact of non-active data. If scale-out landscape is required, master and worker memory requirements are shown separately. Average utilization per node in percent alerts if nodes are already filled close to their capacity.

BW on HANA Sizing: How to Size HANA Systems

HANA cannot allocate the complete physical memory of each node. Some space must be reserved for the operating system and other services. As a rule, 10% of the first 64 GB and another 3% of the remaining memory will be reserved exclusively for OS purposes.

So the remaining net capacity for HANA data and temporary space is (after reduction by 50GB for HANA services and caches - all values in GB):

Available memory	Reserved for OS	Available for HANA	Available for data and temp.
256	12	244	194
512	20	492	442
1024	35	989	939
1536	52	1484	1434
2048	66	1982	1932
3072	97	2975	2925

In case of scale-out configurations, row store and non-generated column store tables must fit in one node (the master node). This requirement usually determines the minimum amount of memory per node in a scale-out landscape.

If total master storage requirements are larger than 400 GB, **consider house keeping measures** (see SAP notes 706478, 1729988, and 1829728 for details)

BW on HANA Sizing: Disk Sizing

HANA requires three different disk areas (usually different disk volumes):

- HANA data files
- HANA log files
- HANA executables, dumps, exports, backup, etc.

The size of the data files will be approximately as large as the overall memory requirement.

For small systems (< 512 GB) half the memory requirement should be reserved. All other systems require a total of 512 GB for log files.

For executables, dumps, exports, etc. we recommend reserving at least 2x-3x the total memory requirement.

Please refer to the [HANA Storage Whitepaper](http://www.saphana.com/docs/DOC-4071) (<http://www.saphana.com/docs/DOC-4071>) for further details.

BW on HANA Sizing: How to Size a Single Node System

MINIMUM SIZING RECOMMENDATION - CURRENT			
=====			
Phys. memory per node:	512 GB	1024 GB	
Memory Requirement (Minimum Total):	943 GB	843 GB	
Disk Space Requirement - data (Minimum Total):	943 GB	843 GB	
Disk Space Requirement - logs (Minimum Total):	943 GB	843 GB	
Number of Nodes incl. master (Minimum Total):	3	1	
SIZING DETAILS			
=====			
(For 512 GB node)	data [GB]	total [GB] incl. dyn.	total [GB] (non-act.)
<u>MASTER:</u>			
Row Store	45	89	89
Master Column Store	5	10	10
Caches / Services	50	50	50
TOTAL (MASTER)	100	149	149
<u>SLAVES:</u>			
Slave Column Store	357	700	693
Caches / Services	100	100	100
TOTAL (SLAVES)	457	800	793
TOTAL (All Servers)	557	949	943

Example: Node with 1024 GB

Available memory (see p.33): 989 GB

Memory for services / caches: 50 GB

Free memory for data: 939 GB

Master Row Store Size:

89 GB

Master Column Store Size:

10 GB

Worker Column Store Size:

693 GB

Total Data Size:

792 GB

Memory for services / caches:

50 GB

Total Memory Requirement (data + services):

842 GB

< 939 GB 

BW on HANA Sizing: How to Size a Scale Out System

MINIMUM SIZING RECOMMENDATION - CURRENT			
=====			
Phys. memory per node:	512 GB	1024 GB	
Memory Requirement (Minimum Total):	943 GB	843 GB	
Disk Space Requirement - data (Minimum Total):	943 GB	843 GB	
Disk Space Requirement - logs (Minimum Total):	943 GB	843 GB	
Number of Nodes incl. master (Minimum Total):	3	1	

SIZING DETAILS			
=====			
(For 512 GB node)	data [GB]	total [GB] incl. dyn.	total [GB] (non-act.)
MASTER:			
Row Store	45	89	89
Master Column Store	5	10	10
Caches / Services	50	50	50
TOTAL (MASTER)	100	149	149
SLAVES:			
Slave Column Store	357	700	693
Caches / Services	100	100	100
TOTAL (SLAVES)	457	800	793
TOTAL (All Servers)	557	949	943

Example: Nodes with 512 GB

Available memory (see p.33): 492 GB

Memory for services / caches: 50 GB

Free memory for data: 442 GB

Master Node:

Row Store + Column Store: 99 GB < 442 GB 

Memory for services / caches: 50 GB

Total Master Memory: 149 GB

Worker Nodes:

Column Store: 693 GB

Workers: $693 / 442 = 1,6$ → 2 workers

Memory for caches / services: 100 GB

Total Worker Memory: 793 GB

Grand Total:

149 GB + 793 GB = 942 GB → Scale-out: 3 nodes with 512 GB each

Thank You!

Frequently Asked Questions (1)

- **Do I have to apply any further calculation on the report results to get the final memory requirements?**

No, the sizing report does all the math for you. You can find the overall memory and disk space requirements for different hardware configurations in the section “**SIZING RECOMMENDATION – CURRENT**” of the report. The information in section “**SIZING DETAILS**” helps understanding how the overall memory requirements are mapped to master and worker servers.

- **Why is the size of a table calculated by the report different to the size reported by the database (e.g. in transaction DB02)?**

The sizing report reads a sample set of records from the DB, analyzes the sample size in its ABAP representation, and extrapolates the total size of the table if it was entirely loaded into ABAP memory, using the total row count information of the DB. By looking at the ABAP representation, the impact of DB specific features like intrinsic compression or special data formats on sizing is automatically eliminated, and the report can estimate the size of the table in ABAP by applying a suitable compression factor. In the output of the report you can find the ABAP size of tables along with the estimated HANA size. Due to the different representation of data in ABAP and the underlying DB system, the DB table sizes may significantly differ from the ABAP sizes.

- **In section “SIZING RECOMMENDATION – CURRENT” the report prints “n/a”, but section “SIZING DETAILS” shows proper results**

The values in section “**SIZING RECOMMENDATION – CURRENT**” are calculated for specific hardware configurations (currently for server nodes with 512GB and 1024GB RAM). While the overall result is valid, a mapping to the existing server configurations was not possible. One reason may be that the master data cannot be stored in a single server node due to its size.

- **I am only getting the first 10 pages of output when saving from spool**

This is a system default setting which you have to change before saving the file. Press Ctrl-Shift-F10 in the spool output controller and change “To page” from 10 to 99.

Frequently Asked Questions (2)

- **Why do the results of the latest version of the sizing report differ from those of a previous version?**
The report is constantly being updated to reflect recent changes in HANA and to improve the sizing algorithms. In some cases this can lead to significant changes in the sizing results. We recommend relying on the results of the very latest version of the sizing report.