



SAP BW/4HANA AND SAP BW-ON-HANA WITH SAP HANA EXTENSION NODES

ABSTRACT

This document contains the Frequently Asked Questions regarding the SAP BW/4HANA and SAP BW-on-HANA with SAP HANA Extension Nodes

Dr. Klaus Nagel,
Gordon Witzel

Table of Content

Abstract.....	3
General Remarks.....	3
Value Proposition.....	3
What is the SAP HANA Extension Node Value Proposition?	3
Questions	4
Release state.....	4
1. What are the recommended releases?	4
2. What is the release state of SAP HANA Extension Nodes?.....	4
3. Are SAP HANA Extension Nodes released also for non-BW scenarios/applications?	4
Positioning	4
4. How is SAP Dynamic Tiering with Extended Storage positioned in relation to the SAP HANA Extension Nodes?.....	4
5. How does dynamic tiering compare to extension nodes for non-BW scenarios/applications?.....	5
6. How are SAP HANA Extension Nodes positioned in relation to SAP BW-NLS?.....	5
Sizing	5
7. How do I size BW for SAP HANA Extension Nodes?	5
8. What SAP BW data can be moved to SAP HANA Extension Nodes?	5
9. How much data can I move to the SAP SAP HANA Extension Nodes?	5
Limitations etc.....	6
10. What are the drawbacks of moving data to the SAP HANA Extension Nodes?.....	6
11. Are there functional limitations for data on SAP HANA Extension Nodes?.....	6
12. Are the BW-generated SAP HANA Views for external consumption supported with SAP HANA Extension Nodes?.....	6
13. Are there restrictions about SAP HANA operations with SAP HANA Extension Nodes?	6
14. What are supported SAP HANA scale-out settings for SAP HANA Extension Nodes?	6
15. Is it possible to have more than 1 SAP HANA Extension Node possible?	7
16. Is it possible to rollback the changes after implementing SAP HANA Extension Nodes?.....	7
17. What kind of objects and which BW version(s) are supported?	8
18. Are SAP HANA Extension Nodes for SAP BW supported in HEC?	8
Performance	8
19. Will there be any performance impact on the current SLA for data loading when using SAP HANA Extension Nodes?	8
Others	9

20.	Where do I find details about how to set up SAP HANA with SAP HANA Extension Nodes?	9
21.	How should the Development and QA environment look like when Extension Node is configured in the production environment?	9
22.	Which HW partners offer setups with SAP HANA Extension Nodes?.....	9
23.	How does the license model for data in SAP HANA Extension Nodes look like?	9
Additional information/Links		9
Blogs:.....		9
Documentation:		9
Notes:.....		10

Abstract

The SAP HANA Extension Nodes offer a simple way to use SAP HANA memory more efficiently. Data with the right access pattern and a partitioning specification can be moved to the SAP HANA Extension Nodes where more data can be stored than using the usual 50:50 ratio of SAP HANA for data footprint versus working memory.

With BW's tight control over its data it is the primary application to make use of the SAP HANA Extension Nodes. BW data can be classified into HOT-WARM-COLD categories, depending on its different access patterns and SLAs. The SAP HANA Extension Nodes can be used to hold the WARM data of a SAP BW system.

General Remarks

This FAQ currently covers the options 1&2 of the "detail"-blog below. Option 3 (heterogeneous scale-out) is not yet released for customers.

<https://blogs.sap.com/2016/04/14/important-update-data-lifecycle-management-for-bw-on-hana/>

<https://blogs.sap.com/2016/04/26/more-details-hana-extension-nodes-for-bw-on-hana/>

Value Proposition

What is the SAP HANA Extension Node Value Proposition?

SAP wants customers to have architectural choices to solve their specific needs in data warehouses that include tradeoffs between speed, availability and price.

SAP HANA Extension Nodes offer another layer in the SAP Data Tiering Optimization (Multi-Temperature Data Management) strategy that provides several benefits to the SAP BW on HANA and SAP BW/4HANA customer:

- In-Memory tiering that provides higher performance than disk based SAP HANA Dynamic Tiering
- Permits customers to extend the SAP HANA framework to store more data in a SAP HANA node than typical for SAP HANA compute nodes today
- SAP HANA Extension Nodes are easier to set up and administer than disk based SAP HANA Dynamic Tiering
- SAP HANA Extension Nodes offers, right out of the box, all features of SAP HANA with respect to operations, updates, and data management.

Questions

Release state

1. What are the recommended releases?

SAP BW7.4 SP12 is the minimum SAP BW release, but we strongly recommend to use SAP BW7.50 SP05 (or higher) or SAP BW/4HANA.

SAP HANA1.0 SP12 is the minimum SAP HANA release, but we strongly recommend to use SAP HANA2.0 SP01 (or higher).

2. What is the release state of SAP HANA Extension Nodes?

Based on the above releases the SAP HANA Extension Nodes are general available (GA) for all customers.

3. Are SAP HANA Extension Nodes released also for non-BW scenarios/applications?

The BW-case is FULL supported.

For HANA native scenarios, extension nodes are under controlled availability. If you want to use extension nodes in your native scenario, please create a customer ticket on component HAN-DB (SAP Note [2415279](#)).

The Data Lifecycle Manager (DLM) tool included with the SAP HANA Data Foundation (DWF) Support Package 05, Patch 2 has been enabled the usage of SAP HANA Extension Nodes to support Data Tiering for the HANA SQL Data Warehouse Use-Case. DLM is supporting to move table partitions between the HANA Primary Slave nodes and SAP HANA Extension Nodes for HANA RANGE partitioned column store tables, where the HANA datatype of the partition column need to be DATE, TIMESTAMP or SECONDDATE. DLM is offering date-function support for automatic partition movement when executing the according DLM-profile.

The Modeling and Partitioning of the according table can be either done via an integrated tool set using SAP PowerDesigner (PD) / SAP Enterprise Architecture Designer (EAD) (available later in 2017) or via the WebIDE integrated graphical Table-Editor, where the access patterns are based on the HANA virtual data model (VDM) or the SQL-Statements executed by the HANA certified Frontend tool.

Positioning

4. How is SAP Dynamic Tiering with Extended Storage positioned in relation to the SAP HANA Extension Nodes?

The option to use Extended Storage for warm data in BW is still supported in SAP BW-on-HANA, but not in SAP BW/4HANA. We however strongly recommend to use the SAP HANA Extension Nodes also in SAP BW-on-HANA as it is the simpler approach with far less restrictions and in, most cases, superior performance.

Mixing Extended Storage and SAP HANA Extension Nodes in a SAP BW-on-HANA system landscape is NOT supported.

5. How does dynamic tiering compare to extension nodes for non-BW scenarios/applications?

Dynamic tiering is the integration of a disk-backed columnar store with the HANA database. Dynamic tiering is based on technology acquired from Sybase (SAP IQ), and therefore has some functional gaps when compared with a pure HANA system – HANA with an extension node in this case. However, because dynamic tiering has its roots as a disk-backed columnar database, it is well suited for the management of large volumes of warm data at an economical price point. Dynamic tiering runs on commodity hardware, and you may scale memory, cores, and storage independently of each other. Dynamic tiering is currently supported for native HANA data marts without restriction.

6. How are SAP HANA Extension Nodes positioned in relation to SAP BW-NLS?

For SAP BW, we clearly see the need for a 3-tier data approach: HOT-WARM-COLD. The COLD store should be an external store with a basically infinity storage capacity which is NOT part of all HANA operations (like backup&recovery, system copies, ...) and SLAs. The COLD store comes with the better price per data, but of course with a significantly higher negative performance impact. The WARM tier is part of all the HANA operations, does not have functional limitations, but a moderate negative impact is acceptable.

Of course, not every BW system needs to implement all 3 tiers, but depending on the scenario 1 or 2 tiers is absolutely sufficient. If e.g. only a relatively small (few TBs) portion of the data qualifies for the COLD store, this data can also be placed on the SAP HANA Extension Nodes and a COLD store can be omitted completely.

Sizing

7. How do I size BW for SAP HANA Extension Nodes?

The standard SAP BW sizing tool (Version 2.3.0 or higher) can be used (SAP Note [2296290](#) for SAP BW-on-HANA and SAP Note [2363248](#) for SAP BW/4HANA).

8. What SAP BW data can be moved to SAP HANA Extension Nodes?

This depends on the SAP BW release/version. An overview can be found in the “details” blog below.

9. How much data can I move to the SAP HANA Extension Nodes?

There is first the limitation on the BW application where only certain InfoProviders fulfill the technical criteria to be qualified and modeled as “warm”. This is release/version dependent: SAP BW/4HANA certainly covers the biggest scope with selected partitions of several types of (advanced) DSOs. Since “(advanced) DSOs with Activation” are supported for SAP HANA Extension Nodes only with BW7.5 we recommend (see above) to upgrade at least to SAP BW7.5 and convert existing scenarios with InfoCubes and (classic) DSOs to (advanced) DSO to make full usage of the “warm” data qualification.

In addition, there is the limitation of how much I can “overload” the SAP HANA Extension Node – this depends on whether option 1 or 2 is used.

Given these boundary conditions a “normal” ratio of hot:warm data is between 80:20 and 50:50 – anything beyond this is highly questionable if really the data can be qualified as “warm”.

Limitations etc.

10. What are the drawbacks of moving data to the SAP HANA Extension Nodes?

Of course, there are some performance implications for all access types to data sitting on SAP HANA Extension Nodes – that’s why they are classified as “warm” compared to “hot”. In the best-case the access (read and write) only touches the latest partition which is not yet removed from the memory of the SAP HANA Extension Node, in this case processing the data is as fast as for data on the “hot” nodes. In the worst case the complete table must be loaded to memory as the filter criteria do not match with the partitioning specification of the table. In this case, there is a delay until the data is loaded to memory, which can result in delay times up to minutes depending on the size of the table and the I/O of the HANA system.

11. Are there functional limitations for data on SAP HANA Extension Nodes?

There are no functional limitations – **ALL** BW processes are supported on data in SAP HANA Extension Nodes.

12. Are the BW-generated SAP HANA Views for external consumption supported with SAP HANA Extension Nodes?

Yes (but!) – HANA Views can be generated also for BW InfoProviders/Queries if the data is located on the SAP HANA Extension Nodes. Here special care must be taken to ensure all accesses hit the partition criteria of the table which are distributed across “hot” nodes and an SAP HANA Extension Node. Also, the InfoProvider pruning in a CompositeProvider must be ensured by the access to the HANA Views to avoid unnecessary access to the SAP HANA Extension Node.

Clearly the HANA Views are designed for reporting purposes only and classifying reporting data as “warm” should only be done under close surveillance with an understanding of the impact.

Since the BW Analytic Engine has more metadata knowledge about the underlying InfoProviders compared to the HANA Views, a SAP BW access can be “smarter” in filtering out in advance unnecessary access to data on the SAP HANA Extension Nodes.

13. Are there restrictions about SAP HANA operations with SAP HANA Extension Nodes?

All major HANA operation functionality like Backup&Recovery, system replication, active/active, ... are supported.

14. What are supported SAP HANA scale-out settings for SAP HANA Extension Nodes?

The standard scale-out configuration where the SAP HANA Extension Nodes fit best is a scale-out system with 4 or more active nodes. Here simply an existing node can be configured as SAP

HANA Extension Node (Option 1 or 2), the data classified in BW and is then re-distributed. By this you free up additional memory for your “hot” data.

In specific cases, also a smaller (minimal) setup may make sense (and is supported), e.g. 1 “hot” node + 1 SAP HANA Extension Node. If both nodes have the same amount of RAM, this would however mean that the system has a 1:2 ratio (Option 1) of hot:warm data or even a 1:4 ratio (Option 2). As mentioned above such a data distribution may not be realistic.

15. Is it possible to have more than 1 SAP HANA Extension Node possible?

Looking at the most common use cases and data distributions in BW we assume a single SAP HANA Extension Node to be sufficient for almost all scenarios. Should your appropriate sizing show the clear need of more than one SAP HANA Extension Node please create a customer message on component HAN-DB-ENG-BW to discuss whether it is better to increase the data on the single SAP HANA Extension Node (increase the “overload” factor) or to configure more than one SAP HANA Extension Node.

16. Is it possible to rollback the changes after implementing SAP HANA Extension Nodes?

Yes, changing data from “hot” to “warm” as well as configuring and de-configuring SAP HANA Extension Nodes is reversible.

17. What kind of objects and which BW version(s) are supported?

Required SAP HANA prerequisite: SAP HANA SP12 or SAP HANA 2.0.

The following BW objects can be classified as “warm” in BW:

BW Object	Available release
InfoCubes	not available
Classic DSOs	not available
DataSources/PSA tables	BW7.4 SP10
Write-optimized DSOs	BW7.4 SP10
Advanced DSOs w/o Activation	BW7.4 SP10 & BW/4HANA
Advanced DSOs w/ Activation	BW7.5 SP01 & BW/4HANA
Advanced DSOs with reporting access	BW7.5 SP01 & BW/4HANA
RANGE- Partitions of Advanced DSOs	BW/4HANA

18. Are SAP HANA Extension Nodes for SAP BW supported in HEC?

SAP HANA Extension Nodes for SAP BW are not yet part of a standard offering of SAP HEC. Please contact your HEC representative with your demand.

Performance

19. Will there be any performance impact on the current SLA for data loading when using SAP HANA Extension Nodes?

There will be an impact on performance, see the respective remarks in the blogs dedicated to SAP HANA Extension Nodes (see below). The SAP HANA Extension Node concept relies on the following principles: a) only a (small) portion of the data is accessed at the same time, and b) the time needed to upload the data into memory is acceptable for the process. As far as PSA tables are concerned, load operations affect only the latest partitions, as well as extract operations in case of delta extraction. But if used in high-frequency load scenarios or with frequent full load extractions, they should not be classified as “warm”. Write-optimized DSOs have a similar partitioning and access pattern as PSA tables, but they should only be classified as “warm” if they are of type Corporate Memory, i.e. no reporting access, no heavy look-up activities, etc. In contrast to a normal HANA node, where everything is typically kept in memory, in a SAP HANA Extension Node a table/partition needs to be loaded into memory first. As memory to CPU ratio is higher in an SAP HANA Extension Node, there might be also an impact on operations, like HANA execution of DTPs.

Others

20. Where do I find details about how to set up SAP HANA with SAP HANA Extension Nodes?

The technical setup is described in SAP Note [2343647](#) (and SAP Note [2453736](#)).

21. How should the Development and QA environment look like when Extension Node is configured in the production environment?

From a functional point it is not necessary to configure an Extension node in Dev and QA landscapes. Configurations with and without Extension Nodes are supported along the transport landscape. For production-like test environments we strongly recommend setting up an Extension Node also in the QA system. Supposed the QA landscape is regularly updated from production (system copy) an identical system set up is essential for meaningful performance tests. Placing an indexserver for warm data and another indexserver for hot data onto a database host is not supported in productive systems and not recommendable for production-like test environments. Such a setup would lead to performance issues as an Extension Node can carry up to factor 2 or 4 of data compared to a standard node (option 1 or 2) and therefore needs to be accommodated physically.

22. Which HW partners offer setups with SAP HANA Extension Nodes?

Since SAP HANA Extension Node setups with Option 1 and Option 2 use standard HANA HW configurations ALL hardware partners have matching offerings (the configuration of an SAP HANA Extension Node is then done following the guidance in SAP Note [2343647](#) / SAP Note [2415279](#)). Some partners may have offerings that go beyond this with special configurations – please check this with your hardware contacts.

23. How does the license model for data in SAP HANA Extension Nodes look like?

For all license questions, please contact your SAP account partner.

Additional information/Links

Blogs:

<https://blogs.sap.com/2016/04/14/important-update-data-lifecycle-management-for-bw-on-hana/>

<https://blogs.sap.com/2016/04/26/more-details-hana-extension-nodes-for-bw-on-hana/>

Documentation:

SAP BW7.5 documentation:

https://help.sap.com/saphelp_nw75/helpdata/en/f6/99e81daef24dfa8414b4e104fd76b7/frameset.htm

SAP BW/4HANA documentation (FP04):

<https://help.sap.com/viewer/107a6e8a38b74ede94c833ca3b7b6f51/1.0.4/en-US/f2a4eb578452482fbbcb9078a8e51551.html>

SAP BW/4HANA data lifecycle management with DTO (FP04):

<https://help.sap.com/viewer/107a6e8a38b74ede94c833ca3b7b6f51/1.0.4/en-US/9d76dae79ab047099ee81b50208d5945.html>

Notes:

[2317200 - Data lifecycle management for BW on SAP HANA and SAP HANA Extension Nodes](#)

[2296290 - New Sizing Report for BW on HANA](#)

[2343647 - How-To: Configuring SAP HANA for the BW Extension Node](#)

[2415279 - How-To: Configuring SAP HANA for the SAP HANA Extension Node](#)