

Data Tiering Optimization with SAP BW/4HANA

SAP

Public

Agenda

Overview

Tiering options:

- Native Storage Extensions - Key Concepts and Deployment Options
- Extension Nodes - Key Concepts and Deployment Options






Modeling Best Practice

Positioning

Feature details

More information

SAP BW/4HANA Highlights & Design Principles



High Performance

10-100x faster query performance

- In-Memory Data Warehousing
- Algorithm Pushdown
- Advanced Analytics

Openness

Easier access to all information

- Big Data Warehousing
 - Simplified Data Integration
- Native SQL access

Simplicity

Reduced development efforts


- Simplified Data Structures
- Simplified Data Flows
 - Data Lifecycle Management

Modern Interface

New user interface for all users

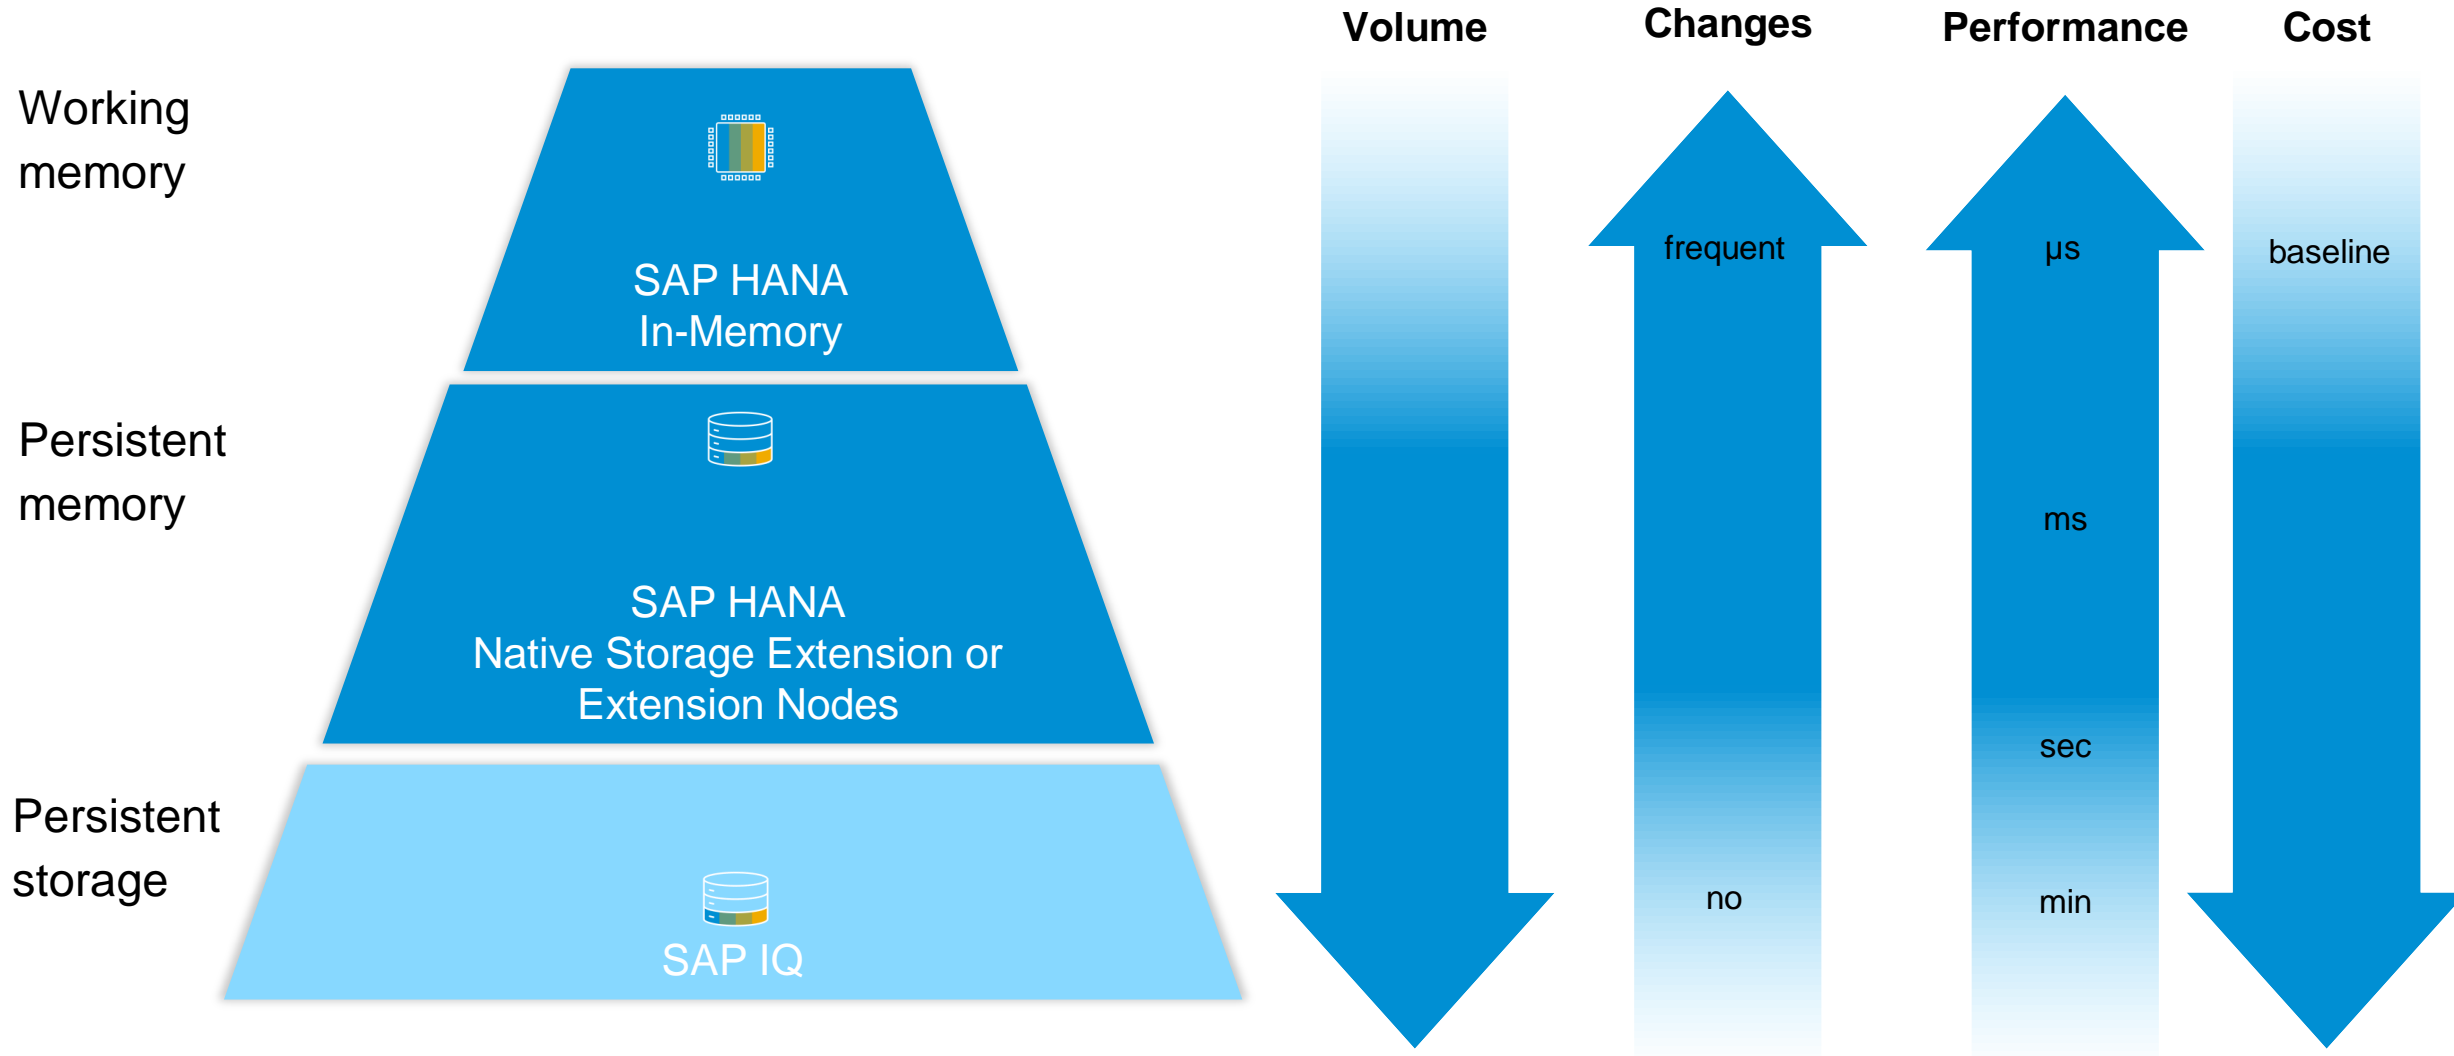
- New Business User UX
 - New Modeler UX
- New Administrator UX

SAP BW/4HANA Data Tiering Optimization Overview



SAP BW/4HANA Data Pyramid

Multiple storage options to meet your volume and performance needs



SAP BW/4HANA – Scalability with Data Tiering

Data Tiering is the assignment of data to various tiers/storage media based upon data type, operational usefulness, performance requirements, frequency of access and security requirements of the data.



Hot
Data

This tier is used to store mission-critical data for real-time processing and real-time analytics.

Data is retained in-memory of the SAP HANA database.



SAP HANA
In-Memory Store



Warm
Data

This tier is used to store data with reduced performance SLAs, which is less frequently accessed.

Data is stored on a lower cost storage tier, managed as a unified part of the SAP HANA database.



Extension Nodes or
Native Storage Extension





Cold
Data

This tier is used to store voluminous data for sporadic or very limited access. Data is stored on lower-cost storage tiers, like SAP IQ, managed separately from the SAP HANA database, but still accessible at any time.



Data Tiering with
External Storage

SAP BW/4HANA Data Tiering – technology per tier



Hot Store

Persistent Memory (PMEM) extends the in-memory storage capacity for hot data in HANA.

Warm Store

Native Storage Extension (NSE) is an intelligent, built-in disk extension for the SAP HANA in-memory database. It is the primary warm store option for HANA on-premise and HANA Cloud. Extension node as warm data tier will continue to be offered.

Native Storage Extension

Extension Node

Cold Store

SAP HANA cold data tiering provides persistence capabilities for HANA cold data in external data stores – SAP IQ.

NLS SAP IQ (SAP BW/4HANA only)

SAP BW/4HANA – Data Tiering Optimization Advantages

One concept for hot, warm and cold data

- Data Tiering based on Advanced DataStore Object Partitions
- Partition Temperature as local setting (no transport needed)
- Using **HANA Technology** such as SDA, Scale Out and disk storage in SAP IQ

Easy and central definition and implementation

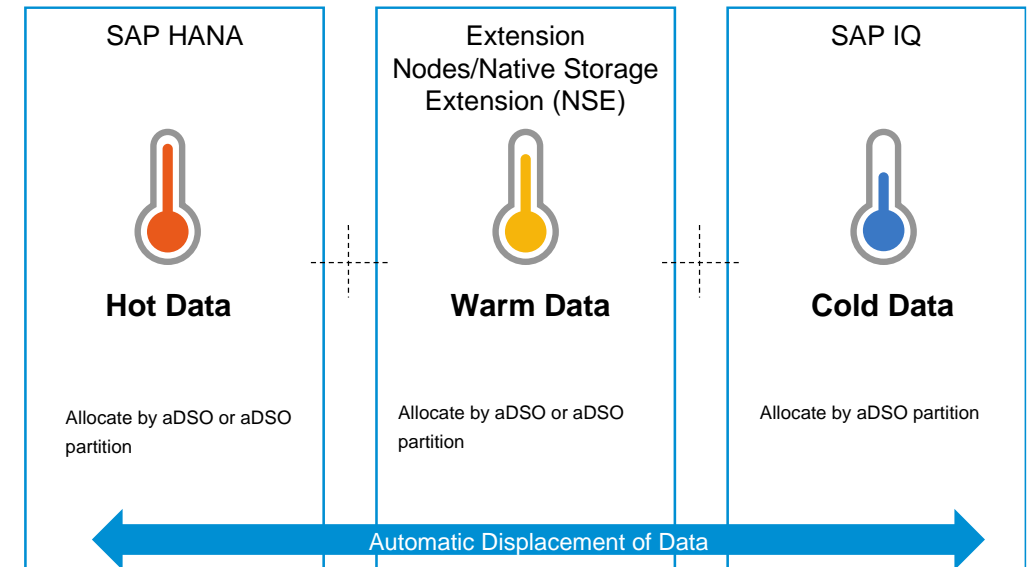
- Data Temperature defined in DataStore Object (advanced)
- No additional configuration of Data Archiving Processes

Displacement of data as simple and periodic housekeeping activity

- Single data tiering optimization job that periodically moves data to defined storages
- No complex process chain modeling for data archiving

Non-disruptive approach and protection of past investments

- Seamless conversion to or co-existence with existing SAP BW NLS IQ approach



SAP BW/4HANA Data Tiering Optimization (DTO)

Implementation Steps

Meta Data Maintenance in Advanced DataStore Object

Data Tiering Properties

Temperature Schema / Tiers

- ☒ Hot (SAP HANA Standard Nodes)
- ☒ Warm (SAP HANA Extension Nodes)
- ☒ Cold (External Cold Store)

Connection: SAPNLS.DTO Browse...

☒ Exceptional Updates to Cold Store

Temperature Maintenance

- ☐ On Object Level
- ☒ On Partition Level

Temperature Schema ①

Partitions

Field: OCALLDAY

☐ Static Partitioning

☒ Dynamic Partitioning

Partition Granularity: Calendar M

Settings: RHADS01

Field: OCALLMONTH

Lower Bound	Upper Bound
01.2016	12.2018
01.2014	12.2015
01.2010	12.2013

Buttons: Add, Edit, Split, Merge, Remove, Maintain Temperatures...

Partitioning Characteristic & Partition Ranges ②

Changes to Partitioning Schema = Remodeling

Data Tiering Optimization – Temperature Definition

Temperature Maintenance

DSO (adv.)	Partition	Partition Field	Planned Temperature	Current Temperature	DTO Status	S	Partition Operator	Low	High	Req. T...
RHADS01	1	OCALMONTH	Hot	Hot			()	01.2016	12.2018	
RHADS01	2	OCALMONTH	Warm	Hot			()	01.2014	12.2015	
RHADS01	3	OCALMONTH	Cold	Hot			()	01.2010	12.2013	
RHADS01	4	OCALMONTH	Hot	Hot						

Partition Temperature Maintenance ③

Plan

Temperature Maintenance - Execution

DataStore Object: Name RHADS01

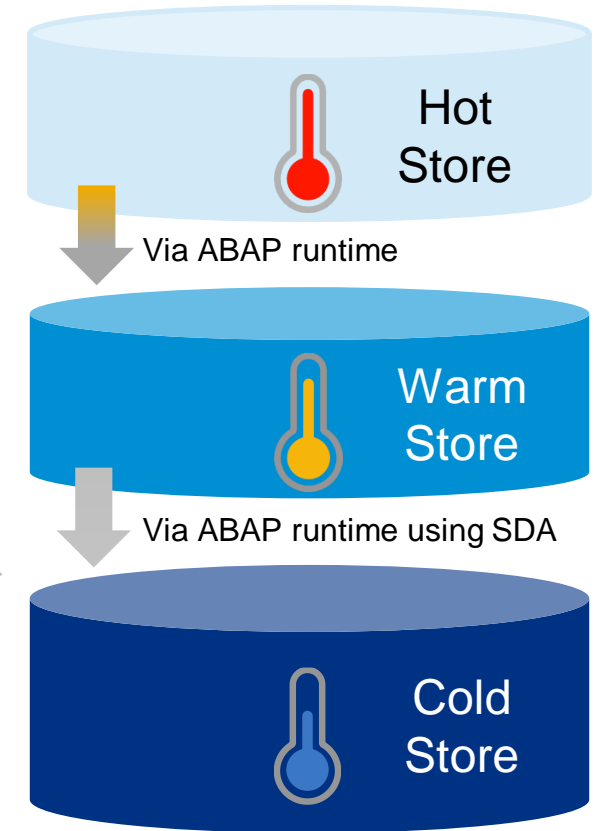
Field Name

☐ Run in Simulation Mode

Data Tiering Optimization Job ④

Changes to Temperature (only) = Data Tiering Optimization

Physical Storage



SAP BW/4HANA Data Tiering Optimization (DTO)

Native Storage Extension (NSE) Example

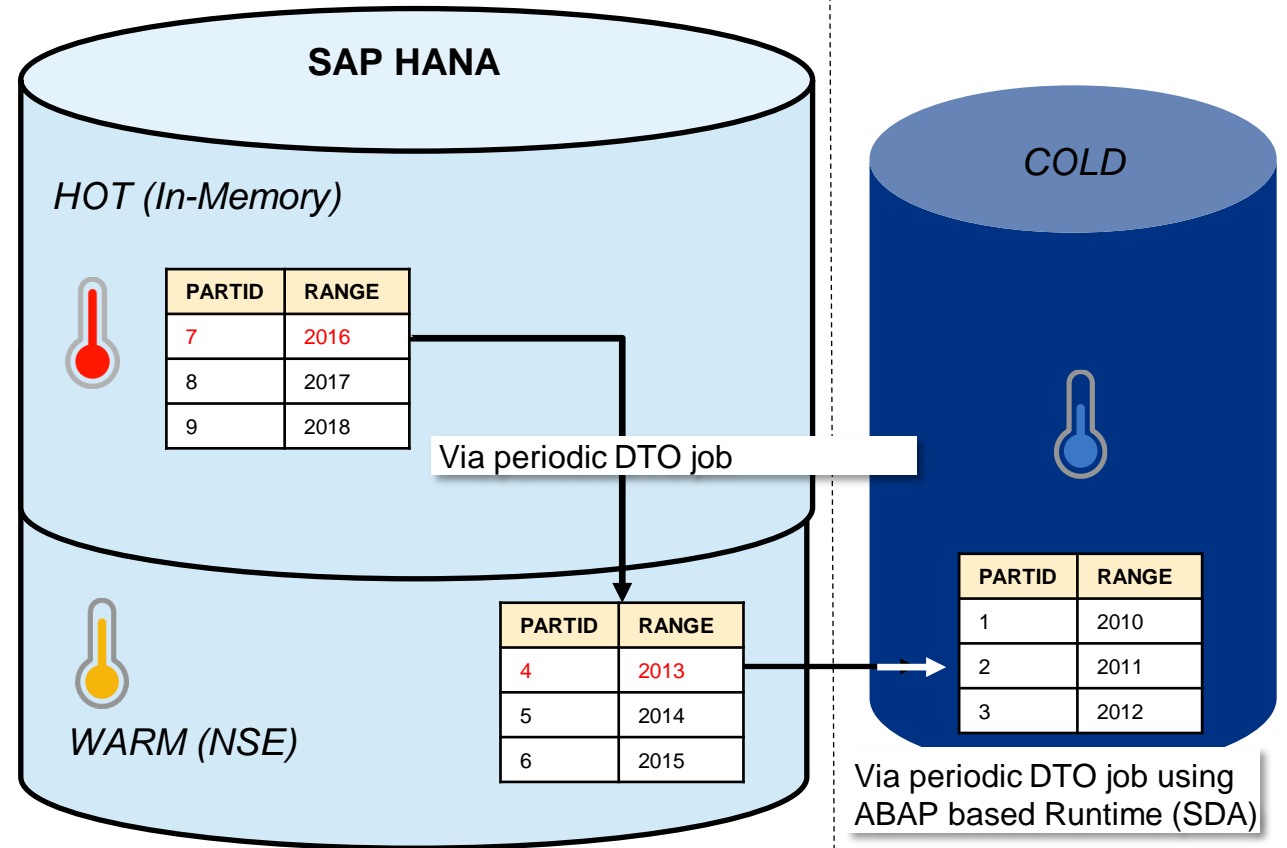
Partition Temperature in SAP BW/4HANA ("Plan")

ADSO Data Tiering Optimization (DTO)										
DSO (adv.)	Partiti	Partition Field	Planned Te	Current	Temperature	DTO Status	Status D	Low (incl.	High (excl	Number
ZDTO_GW	0001	OCALYEAR	Cold	Cold				2010	2011	11028
ZDTO_GW	0002	OCALYEAR	Cold	Cold				2011	2012	11245
ZDTO_GW	0003	OCALYEAR	Cold	Cold				2012	2013	11195
ZDTO_GW	0004	OCALYEAR	Cold	Warm				2013	2014	11131
ZDTO_GW	0005	OCALYEAR	Warm	Warm				2014	2015	11080
ZDTO_GW	0006	OCALYEAR	Warm	Warm				2015	2016	11151
ZDTO_GW	0007	OCALYEAR	Warm	Hot				2016	2017	10944
ZDTO_GW	0008	OCALYEAR	Hot	Hot				2017	2018	11025
ZDTO_GW	0009	OCALYEAR	Hot	Hot				2018	2019	11201

Updated entries
made by DTO
Admin...

2013 and 2016 to
be moved....

PARTID	RANGE	TEMPERATURE
1	2010	COLD
2	2011	COLD
3	2012	COLD
4	2013	COLD
5	2014	WARM
6	2015	WARM
7	2016	WARM
8	2017	HOT
9	2018	HOT



HANA Scale-Up*

Scale-out

External Storage
SAP IQ

SAP BW/4HANA Data Tiering Optimization (DTO)

Extension Node Example

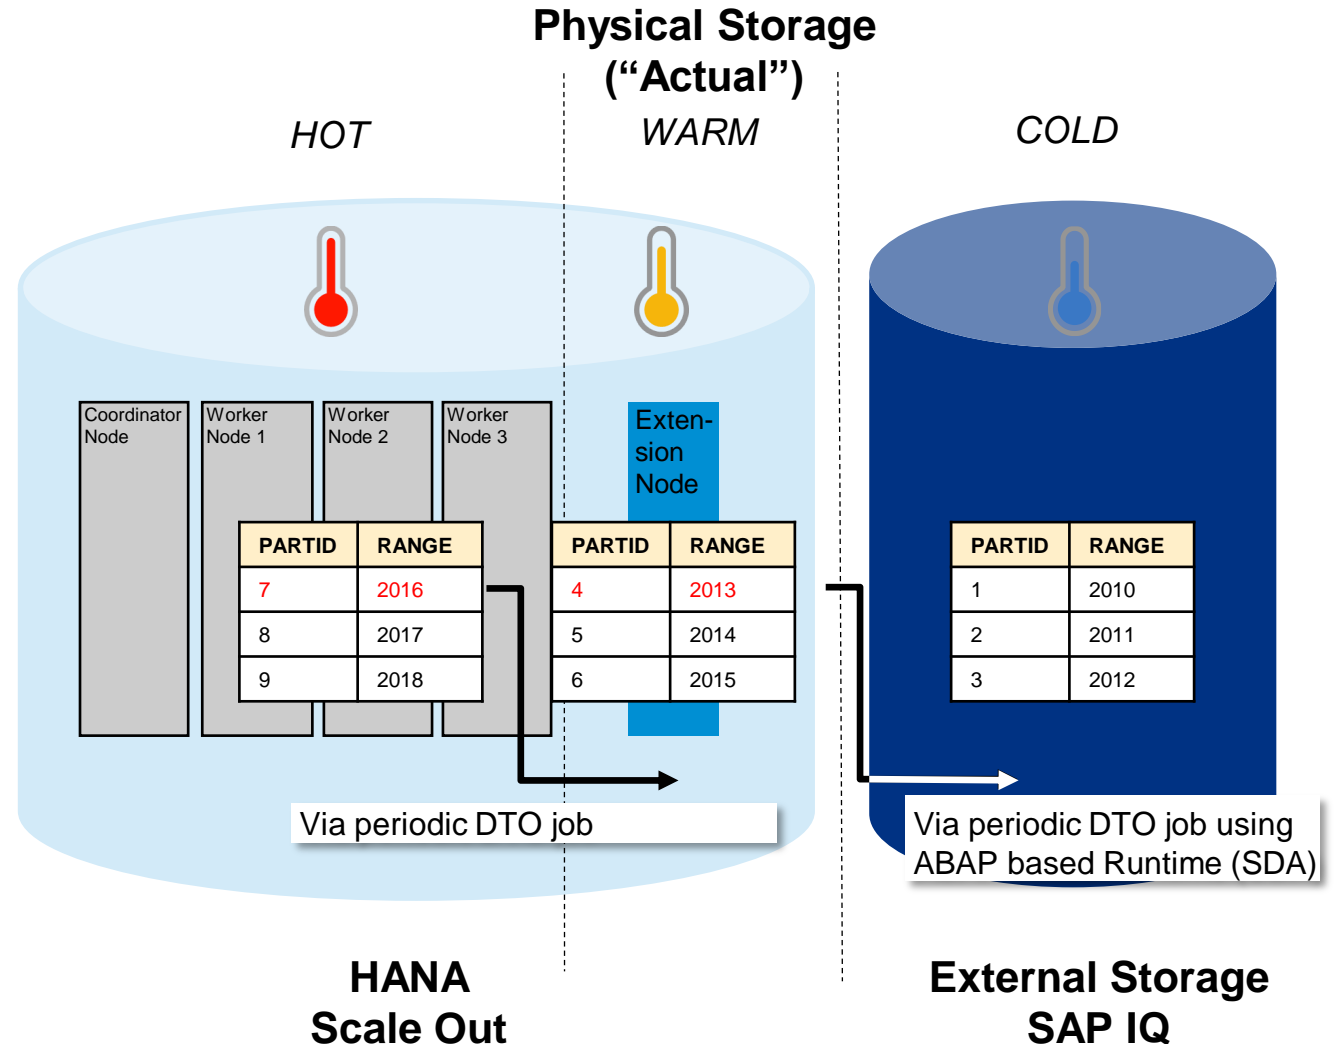
Partition Temperature in SAP BW/4HANA ("Plan")

DSO (adv.)	Partiti	Partition Field	Planned Te	Current	Temperature	DTO Status	Status D	Low (incl.	High (excl	Number	Last Changed On
ZDTO_GW	0001	OCALYEAR	Cold	Cold				2010	2011	11028	02.08.2017 11:52:46
ZDTO_GW	0002	OCALYEAR	Cold	Cold				2011	2012	11245	02.08.2017 11:52:46
ZDTO_GW	0003	OCALYEAR	Cold	Cold				2012	2013	11195	02.08.2017 11:52:46
ZDTO_GW	0004	OCALYEAR	Cold	Warm				2013	2014	11131	02.08.2017 11:54:09
ZDTO_GW	0005	OCALYEAR	Warm	Warm				2014	2015	11080	01.08.2017 15:53:48
ZDTO_GW	0006	OCALYEAR	Warm	Warm				2015	2016	11151	01.08.2017 15:53:48
ZDTO_GW	0007	OCALYEAR	Warm	Hot				2016	2017	10944	02.08.2017 11:54:09
ZDTO_GW	0008	OCALYEAR	Hot	Hot				2017	2018	11025	02.08.2017 11:41:52
ZDTO_GW	0009	OCALYEAR	Hot	Hot				2018	2019	11201	02.08.2017 11:41:52

Updated entries
made by DTO
Admin...


2013 and 2016 to
be moved....

PARTID	RANGE	TEMPERATURE
1	2010	COLD
2	2011	COLD
3	2012	COLD
4	2013	COLD
5	2014	WARM
6	2015	WARM
7	2016	WARM
8	2017	HOT
9	2018	HOT




SAP BW/4HANA Data Tiering Optimization

Native Storage Extensions – Key Concepts and Deployment options



SAP HANA native storage extension (NSE)



- Built-in data tiering feature of SAP HANA with a **simplified architecture** offering substantial scalability at a **low cost**
- **Full** SAP HANA functionality supporting **all data types**, data models, and applications
- **Scale-up supported since SAP BW/4HANA 2.0 SP06 and above**
- Minimize physical design compared to Extension Nodes
- Scale-out supported as of SAP HANA 2021 SP04 and SAP BW/4HANA 2021 SP04
- Should not be used purely as a cost-containment measure... should be SLA driven.
 - Docs [here](#) , FAQ [here](#)
 - Current limitations [here](#)

Cost Effective Solution for Data Capacity Expansion in SAP HANA



NSE allows the user to specify that certain data is “page loadable”

- **Hot “column loadable” data:**

- HANA is an in-memory database, and loads all data into memory for fast processing.
- Data is “column loadable” and resides completely in memory.

- **Warm “page loadable” data:**

- With NSE, less frequently accessed data may be specified as “page loadable”.
- “Page loadable” data is loaded into memory in granular units of pages as required for query processing.
- NSE will reduce memory footprint for “page loadable” data. Data is partly in memory, and partly on disk. Reduces memory usage by not requiring column/partition to be fully memory-resident
- Query performance on warm data may be somewhat reduced compared to hot data.



Only if
Warm
data is
being
accessed


Native Storage Extension (NSE) sizing

What is it: A built-in (native) data tiering feature of SAP HANA with a simplified architecture offering substantial scalability

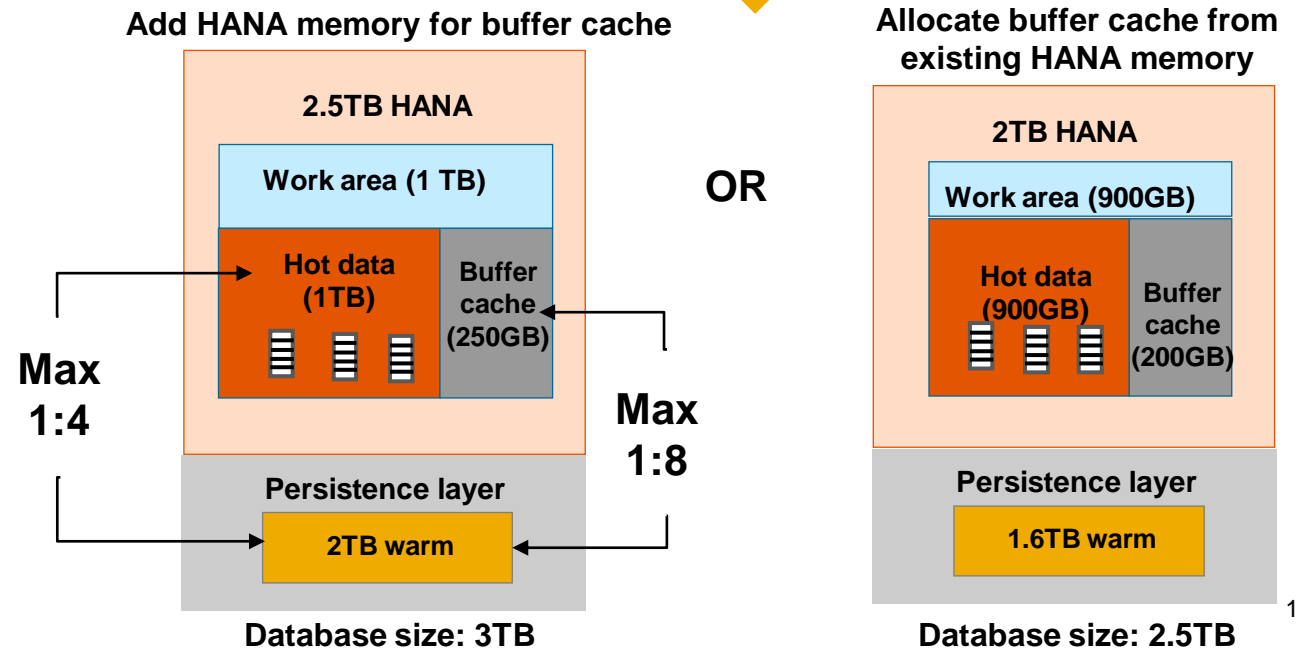
- Buffer cache memory:
 - ✓ Either add more HANA memory for buffer cache
 - or-
 - ✓ Use some of existing HANA memory for buffer cache (will reduce hot data volume)
 - Note: when you reduce the Hot Data capacity, the Work Area capacity can also be reduced to contribute to the Buffer Cache. See examples on following page.
- Two guidelines on sizing:
 - ✓ Total Warm data persistence should not exceed 4x “Hot Data” volume
 - ✓ Warm Data persistence should be no more than 8x Buffer Cache

NSE should be right-sized to the business scenario and not purely as a cost-control measure.

Example: begin with 2TB HANA system
Database size: 1TB



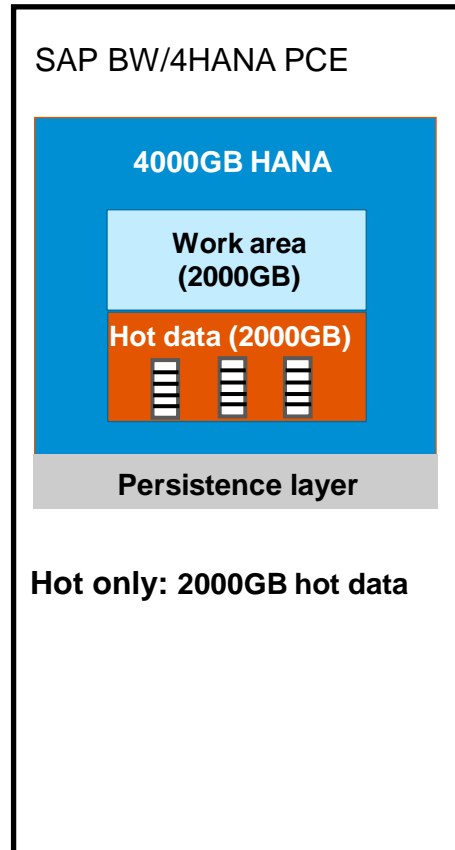
Expand HANA database capacity with warm data that is 2x the size of hot data



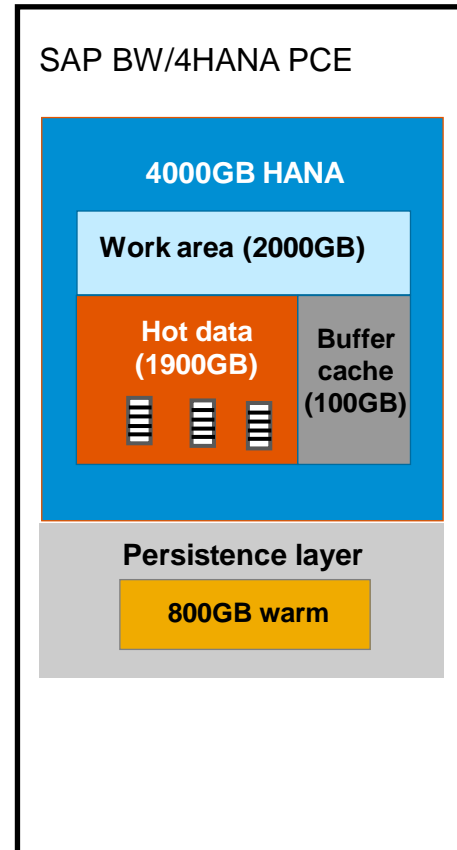
NSE sizing example

BW4 PCE Tier Size: 4000GB

Supported




Supported



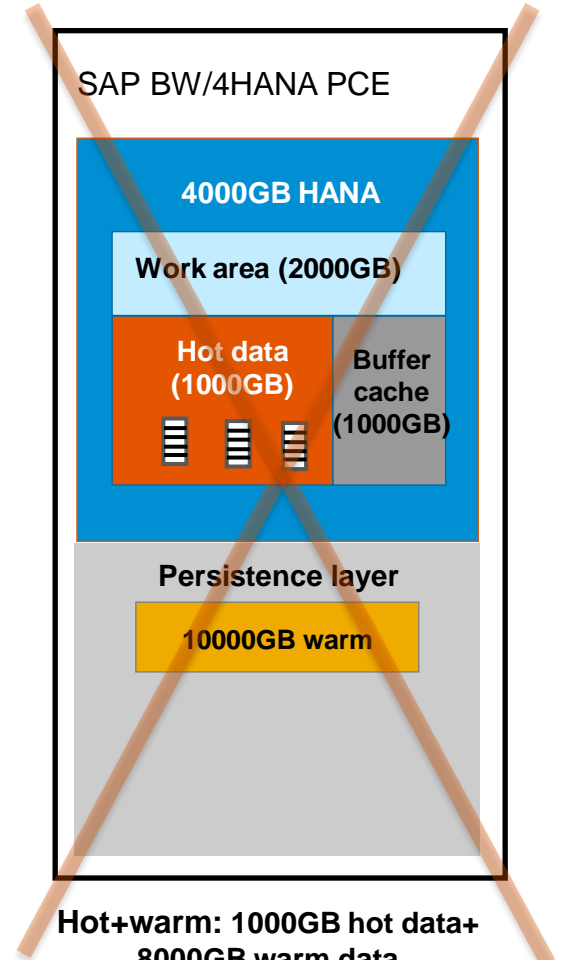
Hot+warm: 1900GB hot data+
800GB warm data

Supported



Hot+warm: 2000GB hot data+
4000GB warm data

NOT Supported




Hot+warm: 1000GB hot data+
8000GB warm data

**Reason: hot:warm ratio is >4x
and buffer:disk is >8x**


SAP BW/4HANA Data Tiering Optimization

Extension Nodes - Key Concepts and Deployment options



What is an Extension Node?

HANA extension node is a warm-data option of SAP HANA's data tiering capabilities to scale the total volume of HANA database systems with lower costs.

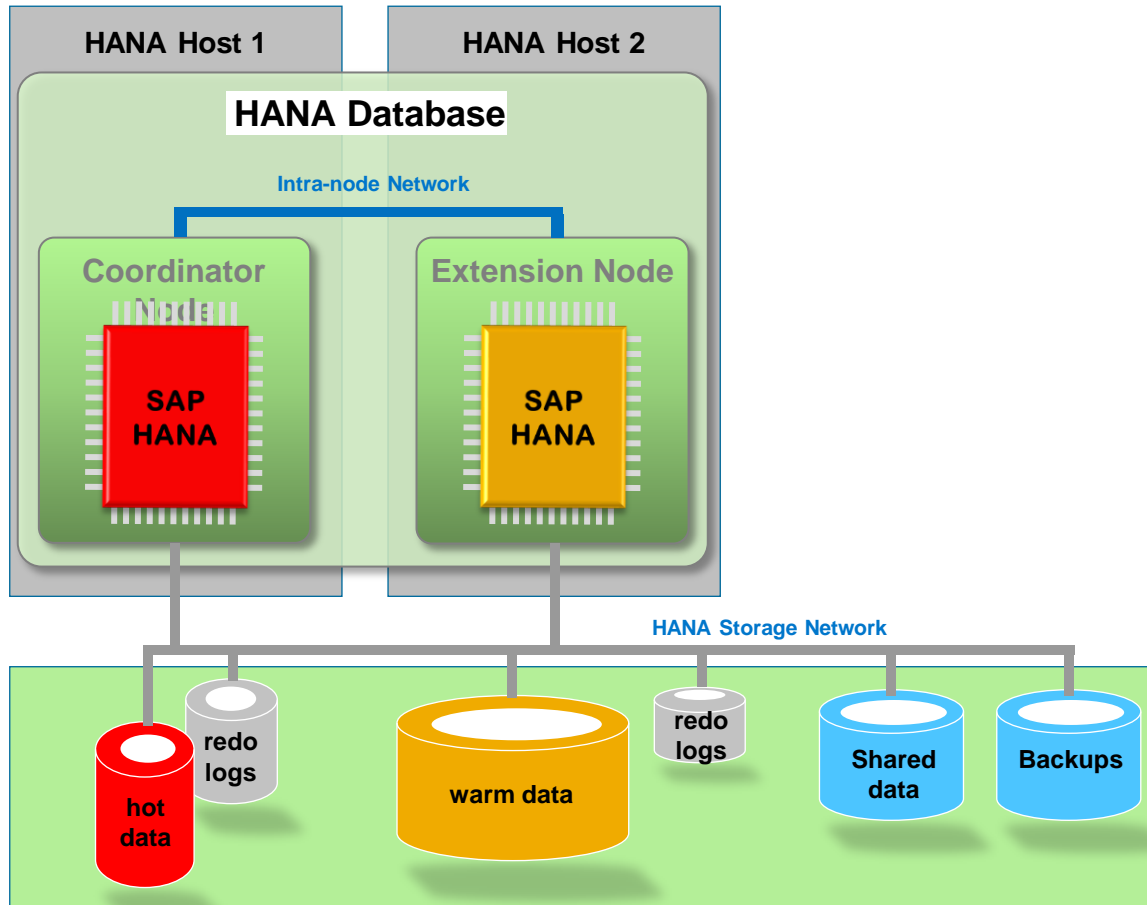


Characteristics:

- Extension Node is fundamentally based on the HANA scale-out feature
- One HANA worker node (slave node) in the scale-out landscape is reserved for warm-data storage and processing
- Support of all SAP HANA features for operations and data management
- Allows larger data footprint of 100% (default) of the node DRAM size
- Relaxed core/memory ratio with HANA TDI5
- Non-Volatile Memory (NVM) supported
- Availability:
 - HANA native applications with HANA 2 SPS03 onwards
 - SAP BW with HANA 1 SPS12 onwards

Extension Node - System Architecture

Overview



*) standby nodes not in this picture

■ Setup:

- Minimal: Coordinator-Node for hot data and Extension Node for warm data
- Optional: Additional worker nodes for hot data
- Optional: Multiple extension nodes

■ Standard scale-out topology:

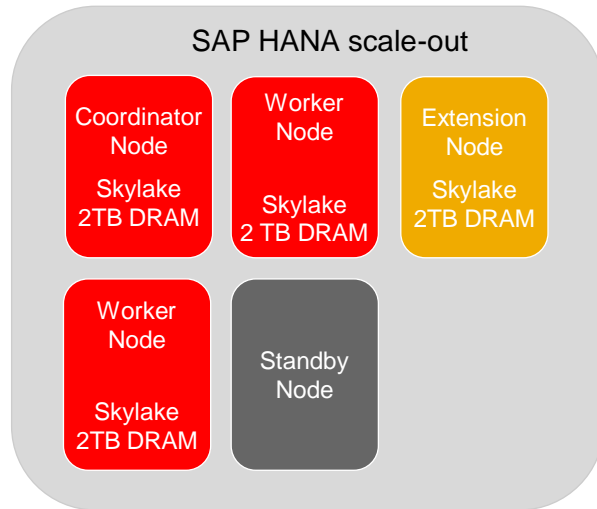
- Each node runs on its own local data set (no shared data)
- Shared file system for node binaries, trace-files, and config-files, dumps and backup
- Standby-node(s) can be shared between hot nodes and extension node, but requires symmetric HW-configuration of the scale-out landscape

SAP HANA Dynamic Tiering vs. Extension Node

SAP HANA Dynamic Tiering	Extension Node
<ul style="list-style-type: none"> ▪ SAP HANA Dynamic Tiering is an additional server, part of HANA landscape and easy to implement and manage 	<ul style="list-style-type: none"> ▪ Extension Node is a HANA scale-out configuration choice and easy to implement and manage
<ul style="list-style-type: none"> ▪ Low cost option to manage WARM data with slower performance than HANA in-memory nodes 	<ul style="list-style-type: none"> ▪ Higher cost option with almost in-memory performance
<ul style="list-style-type: none"> ▪ Less functional parity with HANA (advanced analytics functionality not supported) 	<ul style="list-style-type: none"> ▪ Full functional parity with HANA database
<ul style="list-style-type: none"> ▪ Stores up to 4x of HOT data, if HOT data volume is lower than 2.5TB; up to 8x, if HOT data is larger than 2.5TB 	<ul style="list-style-type: none"> ▪ Stores WARM data up to 100% of DRAM (optional: 200%); multiple extension nodes possible, but not default
<ul style="list-style-type: none"> ▪ Supported Applications: <ul style="list-style-type: none"> ○ SAP HANA native applications 	<ul style="list-style-type: none"> ▪ Supported Applications: <ul style="list-style-type: none"> ○ SAP BW on HANA and SAP BW/4HANA ○ SAP HANA native applications
<ul style="list-style-type: none"> ▪ May run on non-certified commodity HW 	<ul style="list-style-type: none"> ▪ Requires HANA TDI5 certified HW infrastructure

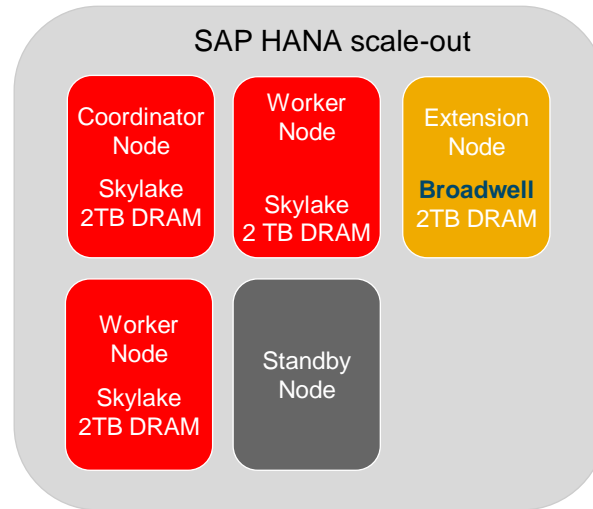
Extension Node - Basic Deployment Options

Symmetric



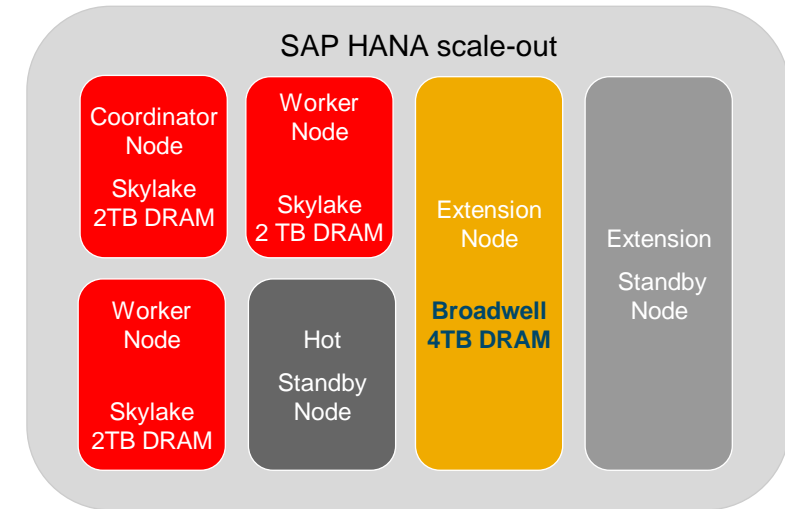
- Extension Node has identical HW-sizing like Hot-Nodes
- Minimal re-configuration required
- Standby node can be shared by all nodes
- Extension node allows larger data footprint (default: 100% of DRAM)
- Example data footprint:
{3x 1TB (Hot Nodes) + 1x 2TB EN} = 5TB
- BW (>=HANA 1 SPS12) and native HANA (>=HANA 2 SPS03) supported

Asymmetric CPU (>=SAP HANA 2 SPS03)



- Extension Node has same memory-sizing, but cheaper CPU (e.g. less cores)
- Minimal re-configuration required
- Standby node can be shared by all nodes, but explicit fallback recommended
- Extension node allows larger data footprint (default: 100% of DRAM)
- Example data footprint:
{3x 1TB (Hot Nodes) + 1x 2TB EN} = 5TB
- BW and native HANA supported (both with HANA 2 SPS03)

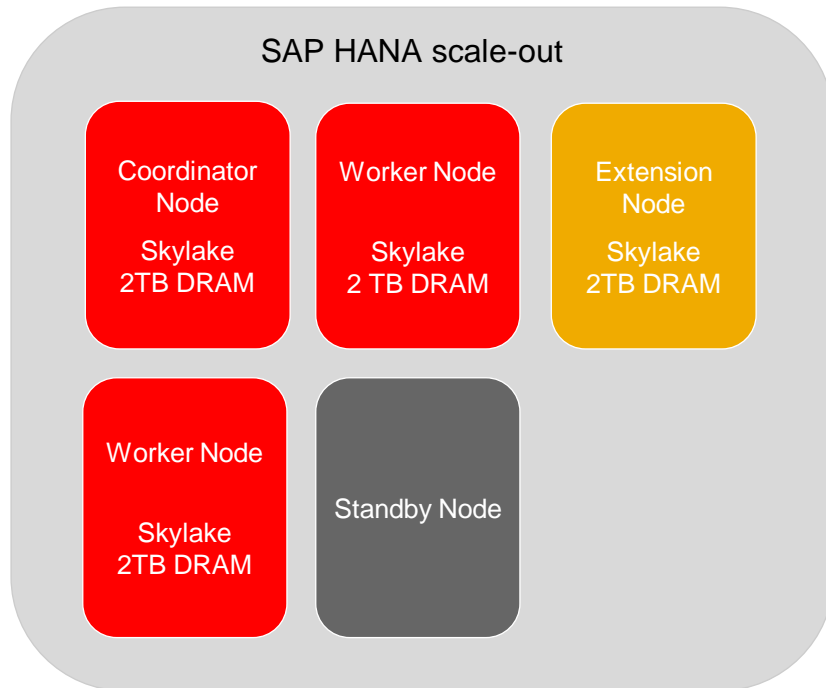
Asymmetric CPU + Memory-Size (>= SAP HANA 2 SPS03)



- Extension Node has higher memory-sizing and cheaper CPU (e.g. less cores)
- Advanced planning for HA and Backup/Recovery
- Extension Node requires its own standby node
- Extension node allows larger data footprint (default: 100% of DRAM)
- Example data footprint:
{3x 1TB (Hot Nodes) + 1x 4TB EN} = 7TB
- Only for BW with SAP HANA 2 SPS03 supported

Extension Node - Advanced Deployment Option with large data footprint

Advanced



Additional Remarks:

- Extension node with large data footprint (up to 200% of DRAM)
- Forced unloads to disk are now very common on the extension node
 - Disk-access degrades EN performance
 - Optimized data layout (partitioning) required
- Re-configuration on storage I/O level required (HW-partner)
- Standby node can be shared by all nodes, but explicit fallback recommended
- Large data volume on the extension node may now dominate the backup and recovery times of the HANA system
- More disk space required for data volume and backup volume (see sizing)
- Table partitioning on extension node required (see performance aspects)
- Example data footprint:
 $\{3 \times 1\text{TB (Hot Nodes)} + 1 \times 4\text{TB EN}\} = 7\text{TB}$
- Large data footprint for BW (\geq HANA 1 SPS12) and native HANA (HANA 2 SPS03) supported

SAP BW/4HANA Data Tiering Optimization Modeling Best Practice



SAP BW/4HANA Data Tiering Optimization (DTO)

Modeling Best Practice

Partitioning Characteristic

- Use preferably 0CAL* or 0FISC* as partitioning criteria for optimal query pruning if possible.
 - Partition / temperature pruning is always based on partitioning characteristics
 - Best Practice: Ensure that a filter on the partitioning characteristics is available in query design
- Please consider the “overwrite” capabilities of certain ADSO Types when adding a (time) characteristic to the key definition of your ADSO
 - Best Practice is to only add immutable characteristics to the key definition of “overwriting” ADSOs (e.g. Creation Date, Initial Loading Date, ...)

Partitions / Ranges

- Align frequency of Data Tiering Optimization job with partition ranges, also avoid very fine granular partitioning with many partitions
 - Fewer partitions needed on SAP HANA. As a consequence, Advanced DataStore Objects (ADSO) only allow for a single partitioning characteristic (from key)
 - The number of partitions is typically calculated as the product of the number of ranges (RANGE) and the available HANA Scala Out Nodes (HASH)
- Plan for sufficient (future) partitions with dynamic partition creation in order to avoid remodelling

Temperature Schema

- In **SAP BW/4HANA systems < SAP BW/4HANA 2021**: Do only use Hot and Warm with data movement on partition level when needed (in a SAP HANA Scale Out landscape)
 - Reason: With this combination, HASH partitioning and hence parallelization in a SAP HANA Scale Out system is not yet possible (but planned)
 - Example: **Do not simply use Hot, Warm and Cold as default temperature schema for all objects in your system**

Partitioning for very large objects

- For very large data volumes in single Advanced DataStore Objects (ADSO) please consider introducing Semantic Groups for mass creation of multiple ADSOs using a semantic split condition that is typically no time characteristic
 - As additional (DB) partitioning criteria for individual ADSOs, time characteristic may apply
 - See SAP Note [2374652](#) for further details


SAP BW/4HANA Data Tiering Optimization (DTO)

Details on Temperature Schema and Partitioning on SAP HANA

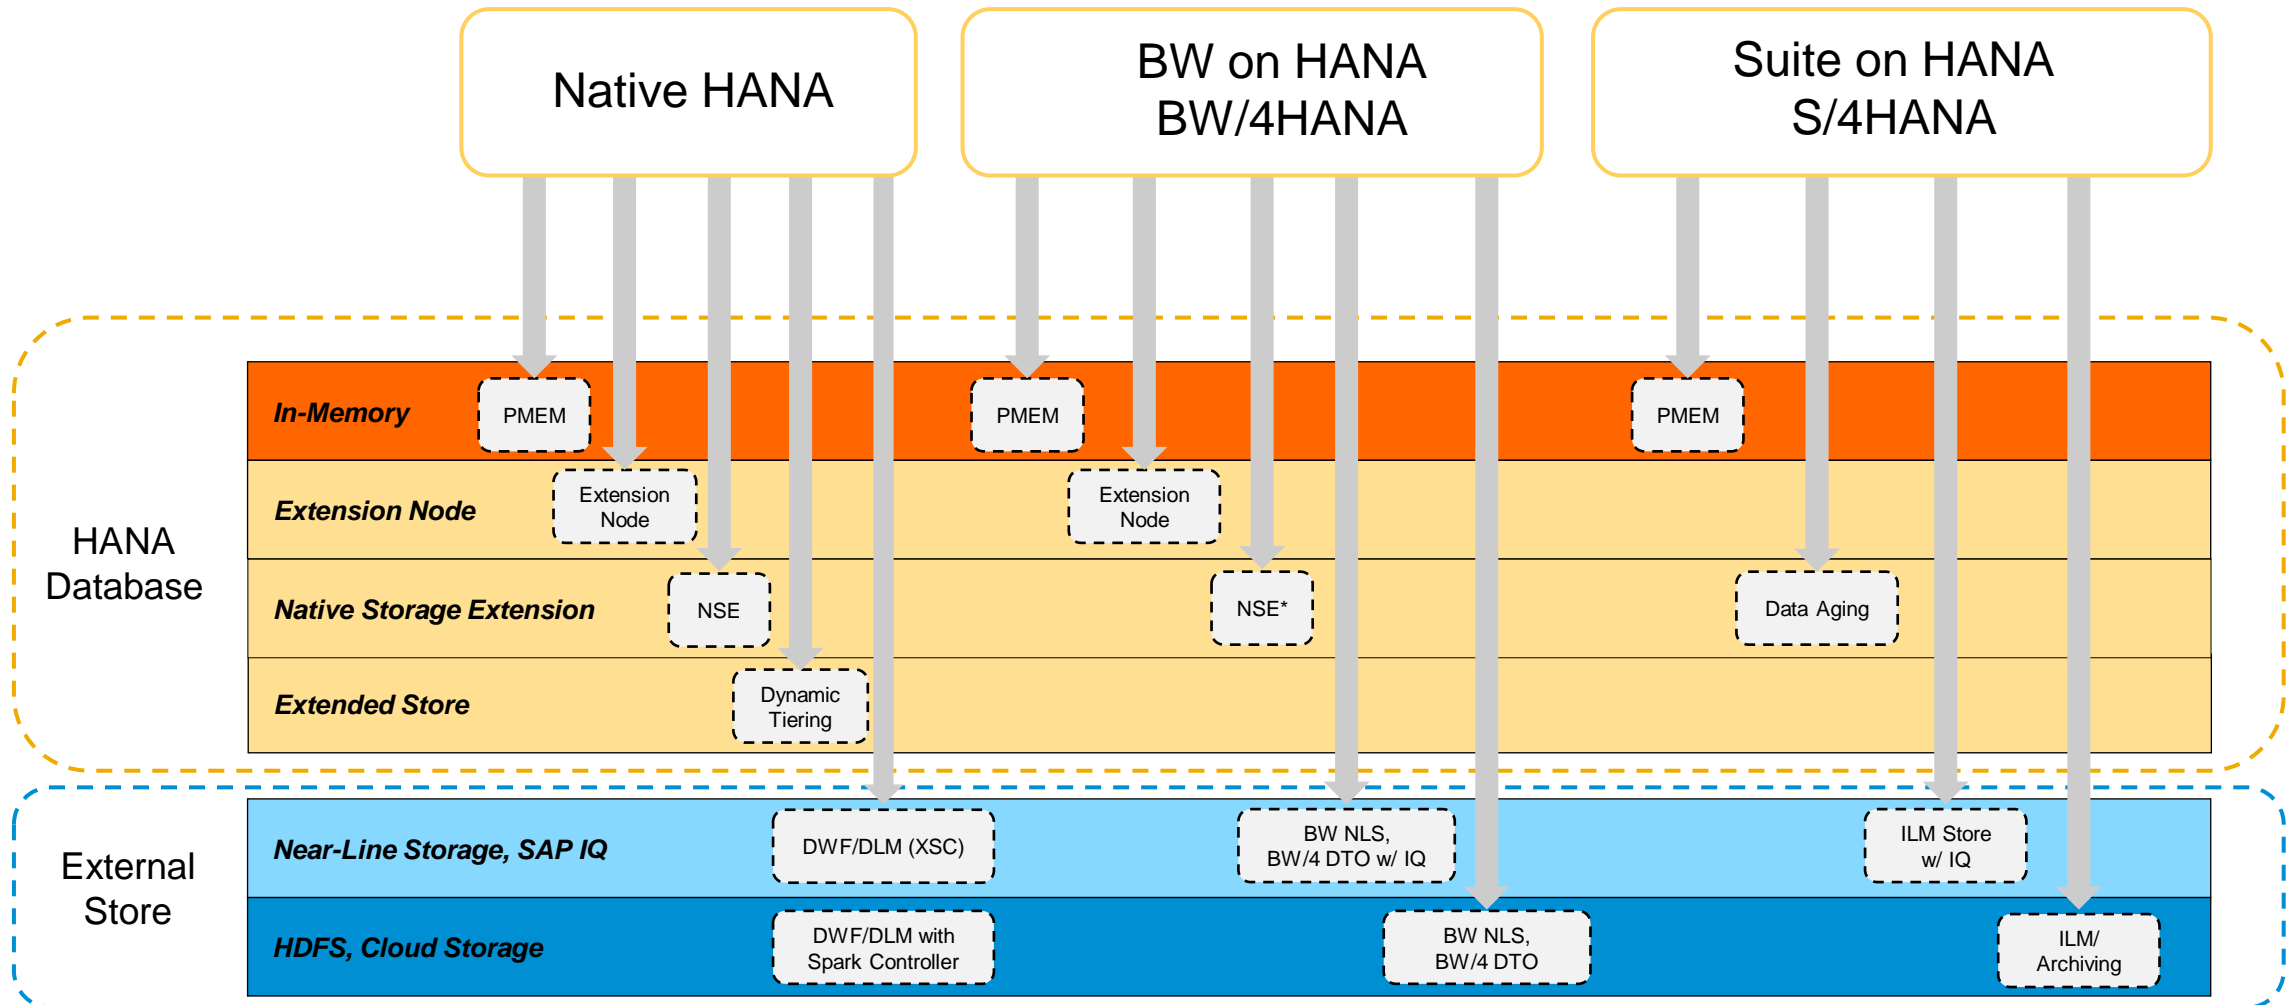
Temperature Schema	Temperature Maintenance	Partitioned (Model)	1 st level Partitioning	2 nd level Partitioning
Hot	N/A	No	Hash	-
Hot	N/A	Yes	Hash	Range
Warm	N/A	No	Hash	-
Warm	N/A	Yes	Hash	Range
Hot Warm	On Object Level	No	Hash	-
Hot Warm	On Object Level	Yes	Hash	Range
Hot Warm	On Partition Level	Yes	Range	Hash*
Hot Cold	On Partition Level	Yes	Hash	Range
Warm Cold	On Partition Level	Yes	Hash	Range
Hot Warm Cold	On Partition Level	Yes	Range	Hash*

* Available from SAP BW/4HANA 2021 onwards

SAP BW/4HANA Data Tiering Optimization Positioning





SAP BW/4HANA – Which Data Tier Should I Use ?



* Exclusively available for SAP BW/4HANA

Data Tiering in SAP BW/4HANA and SAP BW on HANA

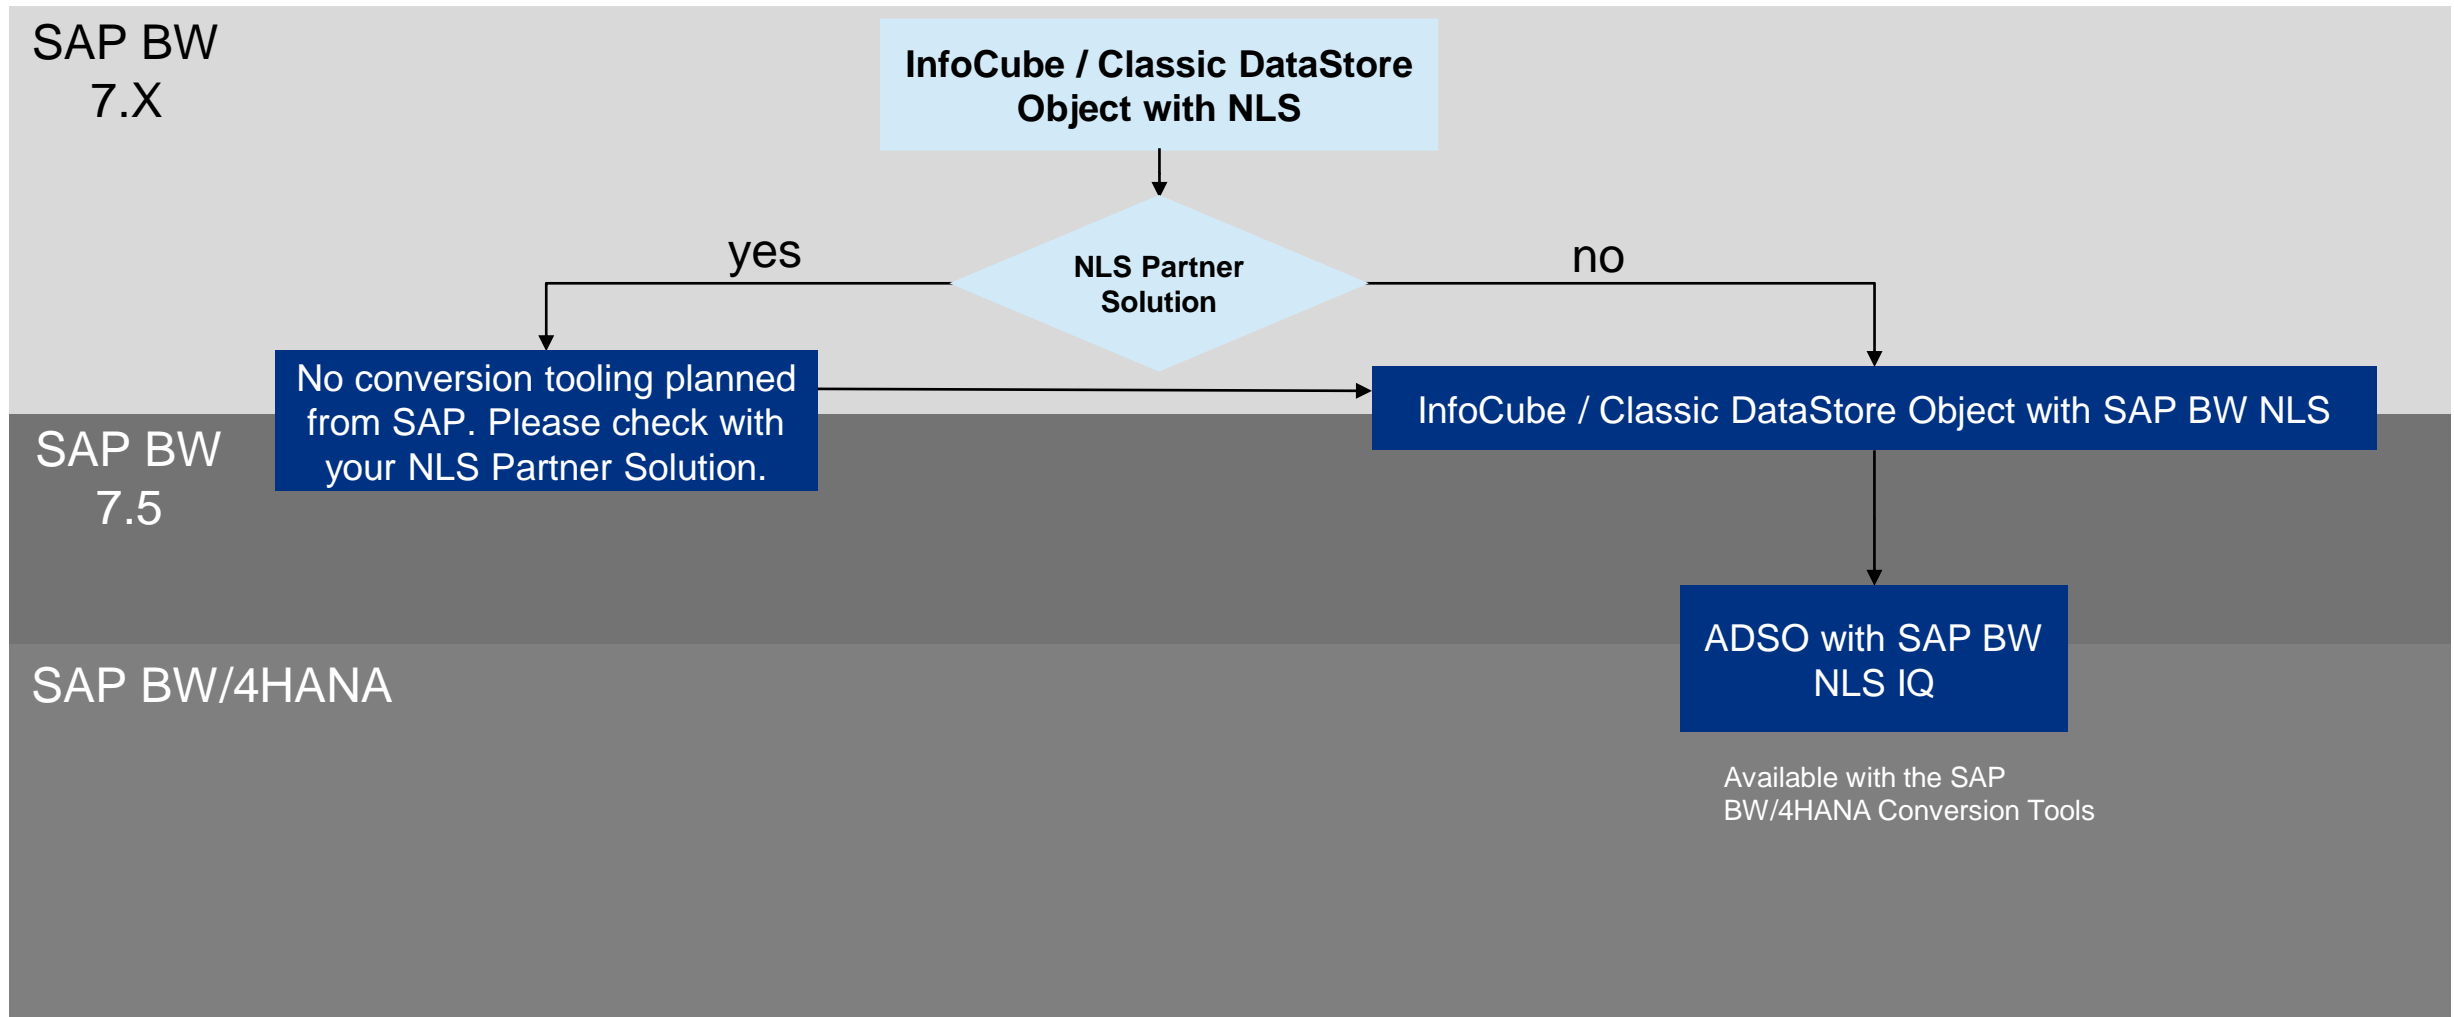
Overview on Options

Data Tier	Application	Technology	Data Storage Location	Supported Data Objects	Application Usage	Required (Minimum) Release
Warmstore 	SAP BW on HANA	Extension Nodes in HANA Scale Out (Maintenance via transaction RSHDBMON)	SAP HANA main-memory + on-disk (unloaded partitions)	PSA Tables, ADSO, Classic DSO (write optimized)	Data access: no restriction, all features available	BW 7.4 SP10 and higher HANA SPS12 and higher
	SAP BW/4HANA	DTO using Extension Nodes in HANA Scale Out or Native Storage Extension in HANA Scale Up or Scale Out *		ADSO	Less recent data and less frequent reporting access	SAP BW/4HANA 1.0 SP04 HANA SPS12 and higher
Coldstore 	SAP BW on HANA	a) SAP BW NLS + certified NLS Partner Solution	Depending on partner solution	All objects that are supported by BW NLS	Data access: read only	a) SAP BW 7.x
	SAP BW on HANA + SAP BW/4HANA	b) SAP BW NLS with SAP IQ	SAP IQ		Historic data and sporadic reporting access only	b) SAP BW 7.x / SAP IQ See SAP Note 1796393
	SAP BW/4HANA	Data Tiering Optimization (DTO) using External Storage in SAP IQ	SAP IQ	ADSO		SAP BW/4HANA 1.0 SP04


* Native Storage Extension in SAP HANA **Scale Out** available from SAP BW/4HANA 2021 **SP04** onwards

SAP BW/4HANA Data Tiering Optimization (DTO) – Cold Store

Conversion approaches for obsolete objects using Nearline Storage (NLS)



SAP BW/4HANA Data Tiering Optimization Feature Details



Data Tiering Optimization (DTO) - Overview

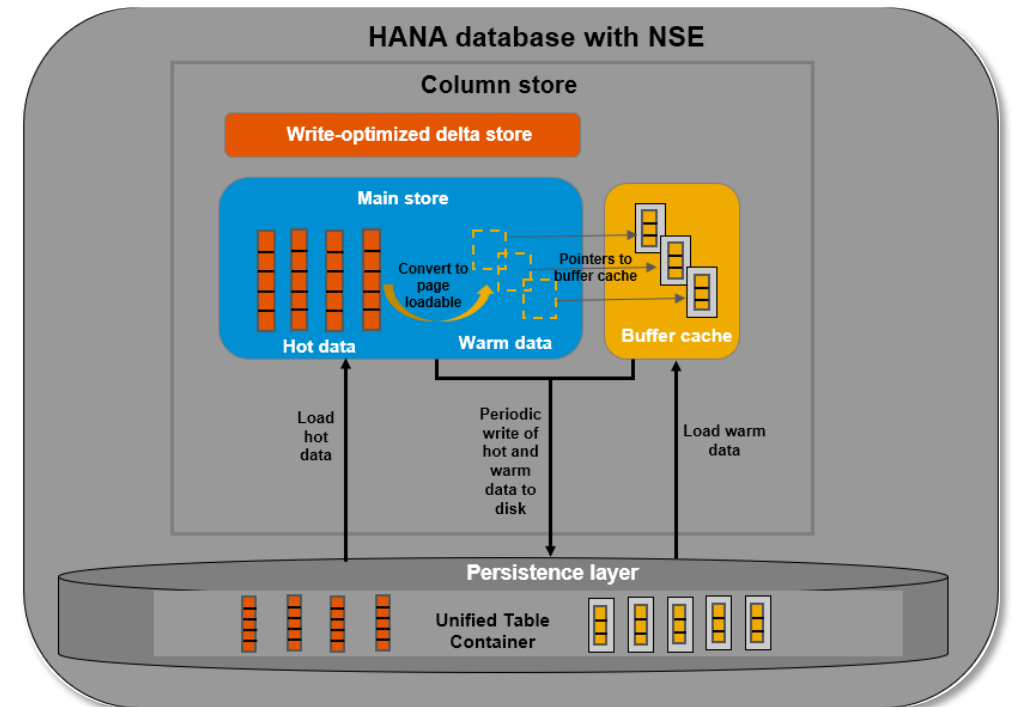
- **SAP HANA Native Storage Extension as warm storage location**
 - Option for SAP HANA Scale Up and Scale Out*
 - Increase HANA data capacity at low TCO
 - Deeply integrated warm data tier, with full HANA and BW functionality
- **Data Tiering Automation based on Temperature Rules**
 - Available for DataStore Objects with partitioning based on time characteristics or date fields
 - Definition of relative (lower, equal, greater) time condition per data tier (hot, warm, cold)
- **Dynamic Partitioning**
 - Automated creation of partitions in data activation process based on predefined partition granularity
 - Based on time characteristics or date fields
 - Static partition creation no longer needed

▼ Partitions

Field:

☐ Static Partitioning
☒ Dynamic Partitioning

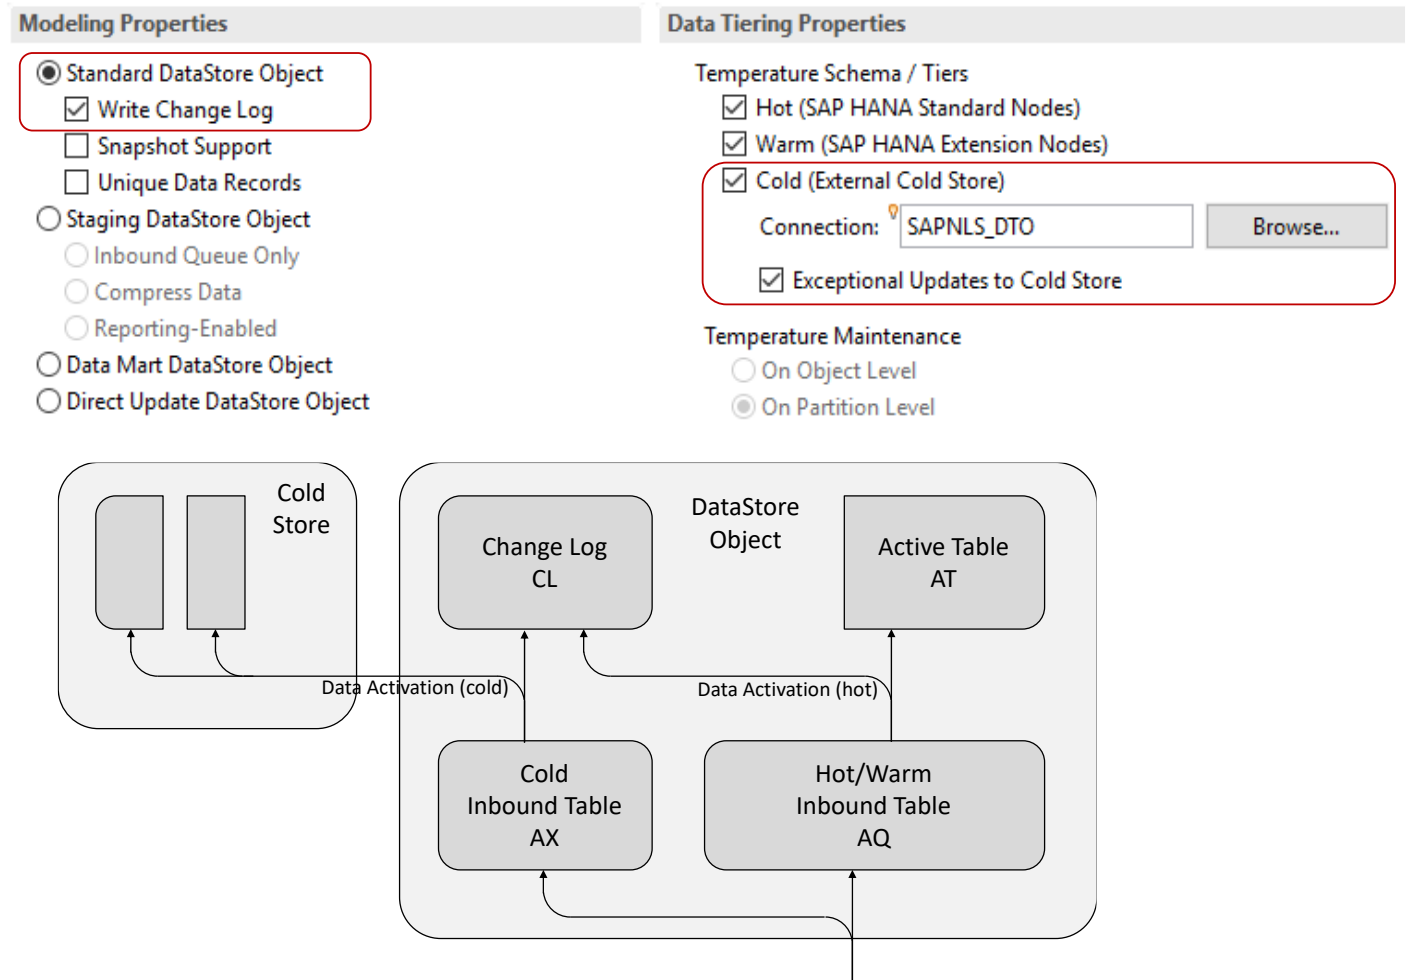
Partition Granularity: ▼



Data Tiering Optimization (DTO) - Exceptional Updates to cold store

As of SAP BW/4HANA 2.0 SP07

- General availability of exceptional updates to cold stores- see [SAP Note 2865304](#) for additional information
- Ability to process updates, inserts and deletes to cold store during cold activation process
- Supported for Standard DataStore Objects with Change Log
- Available on DataStore Object level independently by activating this feature in the DSO maintenance
- Data Store object activation process detects and handles exceptions on cold storage locations
- Available for SAP IQ as cold storage location



Data Tiering Optimization (DTO) - Administration

Deletion support for Cold Store

Feature

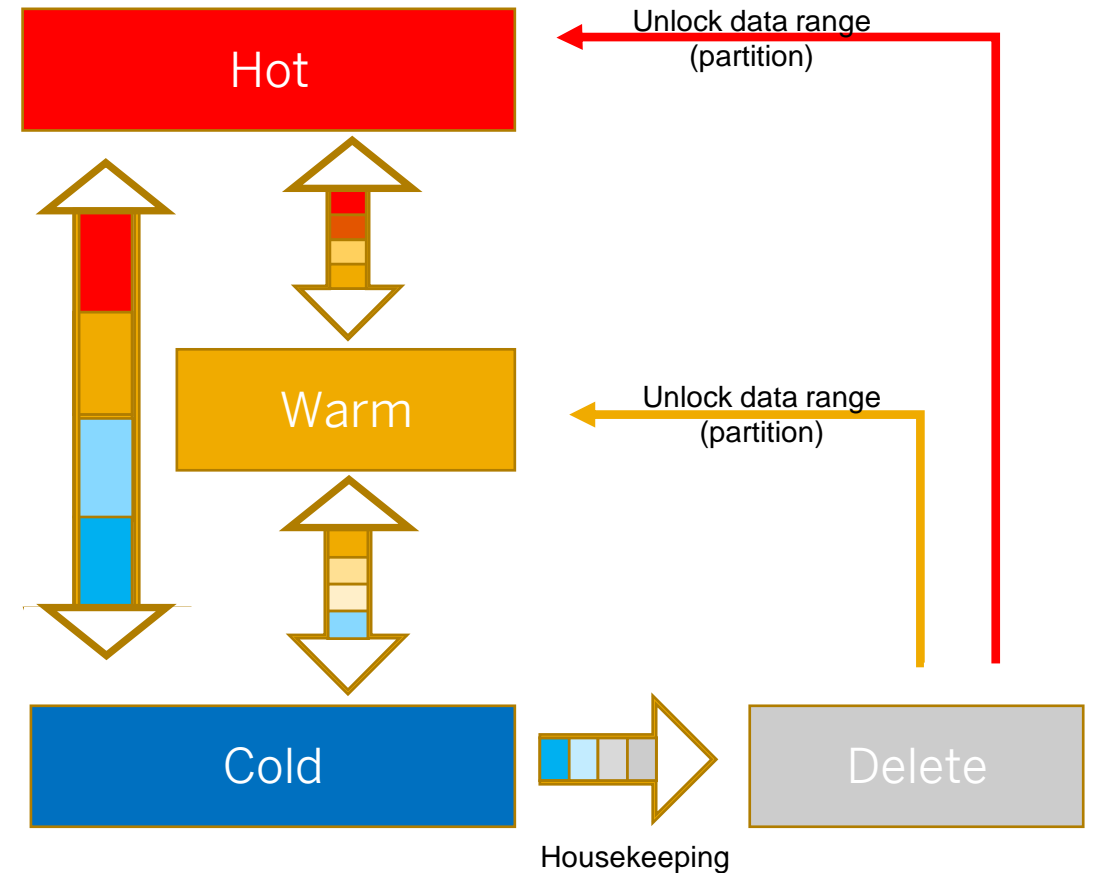
- New data temperature **Delete** for data deletion from the cold store (as housekeeping task for “end of life” data)

Details

- Separation of „logical“ and „physical“ deletion
 - Temperature “Delete” = partition in cold store is removed logically from the cold store = data is flagged for deletion and will no longer get considered during reporting
 - Additional physical deletion is possible on object level using the process type “Cleanup Archive” or as manual activity in the Temperature Maintenance

Considerations

- Only transitions from **Cold** to **Deleted** are supported, no direct transition from **Hot/Warm** to **Deleted**
 - Deletion from Hot / Warm Store is independent from this feature and can get achieved via request deletion (rollback of data loads) or selective deletion (business deletion of specific data ranges in Hot / Warm)
 - From SAP BW/4HANA 2021 SP04 onwards: When a partition is moved to cold storage, the corresponding SAP HANA partition is deleted (if temperature schema HWCP is used)
- Successful transition from **Cold** to **Deleted** can't get reverted
- Unlocking of Cold or Deleted partitions can get achieved by assigning the temperature **Hot / Warm** to the respective partition



Data Tiering Optimization (DTO)

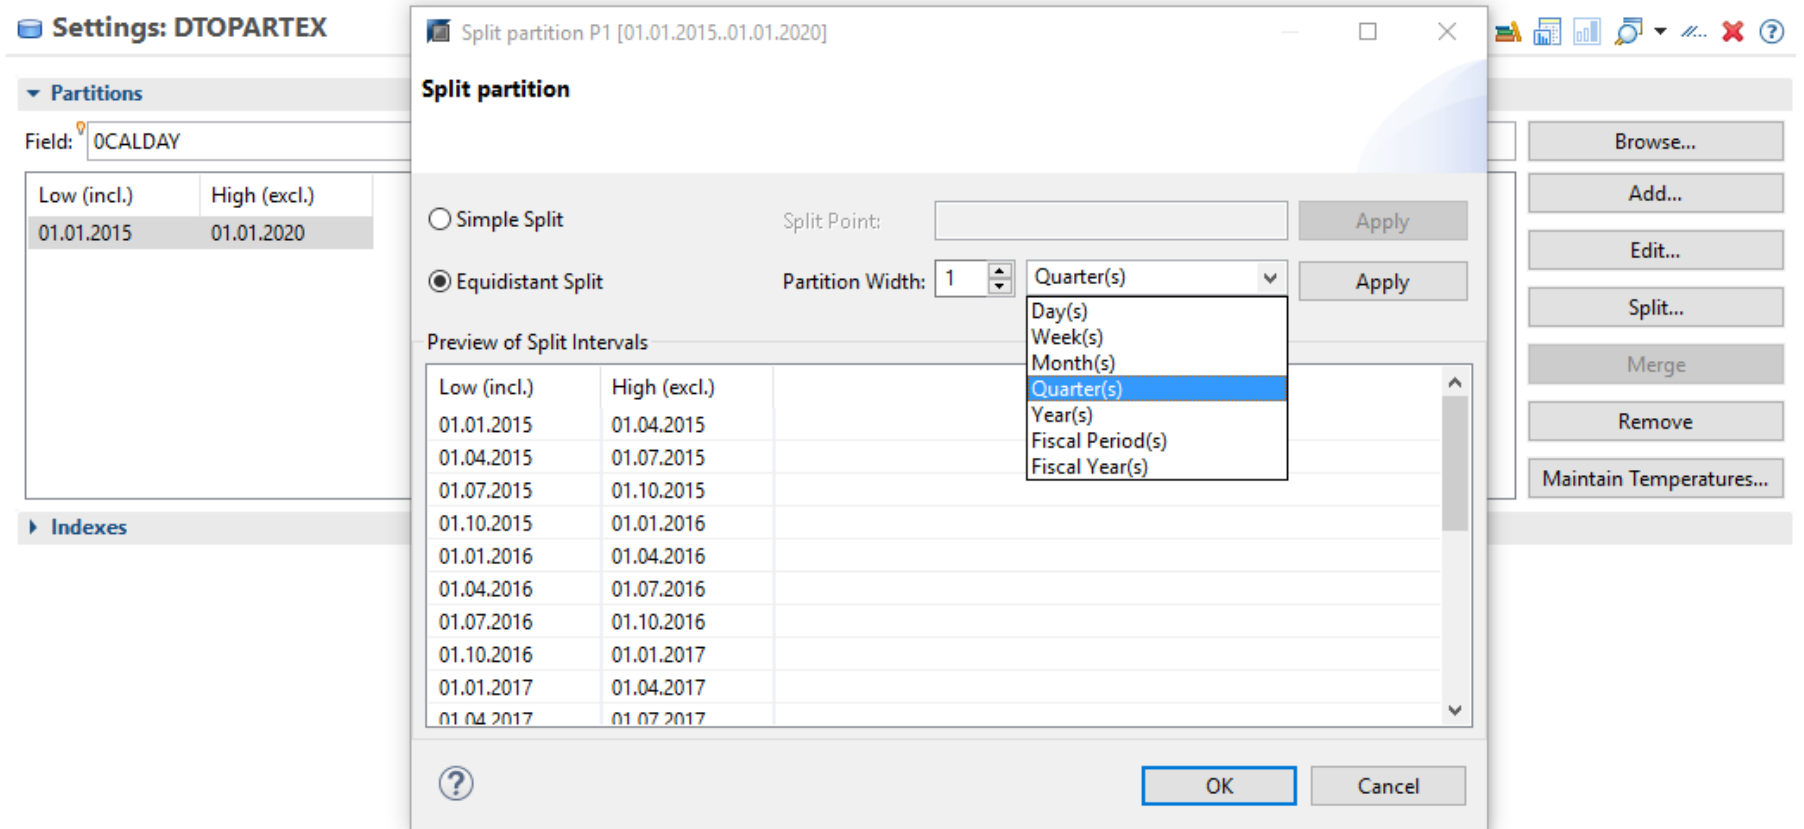
Enhanced Partition Split Capabilities

Equidistant Split for

- Date type
- Time Characteristics
 - 0DATE
 - 0CALDAY
 - 0CALWEEK
 - 0CALMONTH
 - 0CALQUARTER
 - 0CALYEAR
 - 0FISCPER
 - 0FISCYEAR

Equal Unit Split for

- Integer types
- Decimal types



Data Tiering Optimization (DTO)

Process Types

- Automated adjustment of DTO temperature change – New Process Type "Adjust Data Tiering" for Process Chains
- Housekeeping – deleting partitions from Cold Store (new temperature type "delete")

The image displays two screenshots of SAP process configuration interfaces.

Left Screenshot: Properties of Process

- General Tab:**
 - Process Type:** Data Target Administration - Adjust Data Tiering (highlighted with a red box)
 - Use Variant:** ☐
 - Description:** Adjust Data Tiering
- Properties Section:**
 - Data Store Objects:**

Name	Description
ZDTO_01	Sales Order Heade
ZDTOS_01	Sales Order Heade
 - Partitioning Parameter:**

Name	Name
No data	
 - Parallelization Factor:** (empty field)
 - Processing Option:** Synchronize

Right Screenshot: Administrative Data

- General Tab:**
 - Name:** RHTEST
 - Process Type:** Data Target Administration - Adjust Data Tiering
 - Description:** (empty field)
- Properties Section:**
 - Data Store Objects:**

Name	Description
No data	
 - Partitioning Parameter:**

Name	Name
No data	
 - Parallelization Factor:** 3
 - Processing Option:** Synchronize (dropdown menu is open, showing options: Synchronize, Repair, Repair and Synchronize)

Data Tiering Optimization (DTO) - Automation

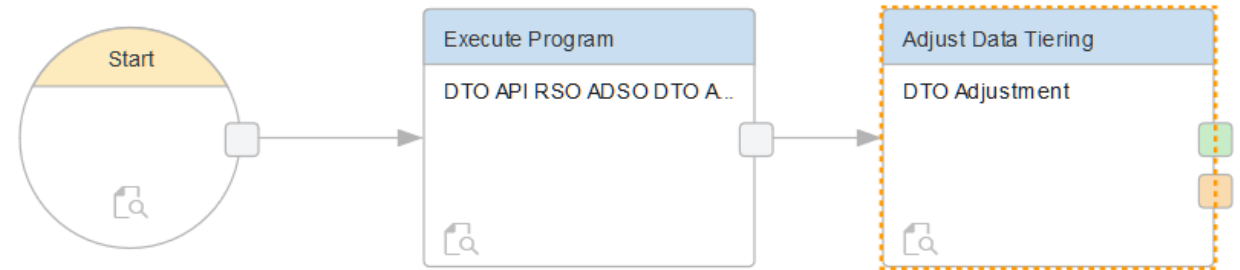
API for Temperature Maintenance and Adjustment

Feature

- API's to implement customer specific automated temperature adjustments for DTO

Typical use case

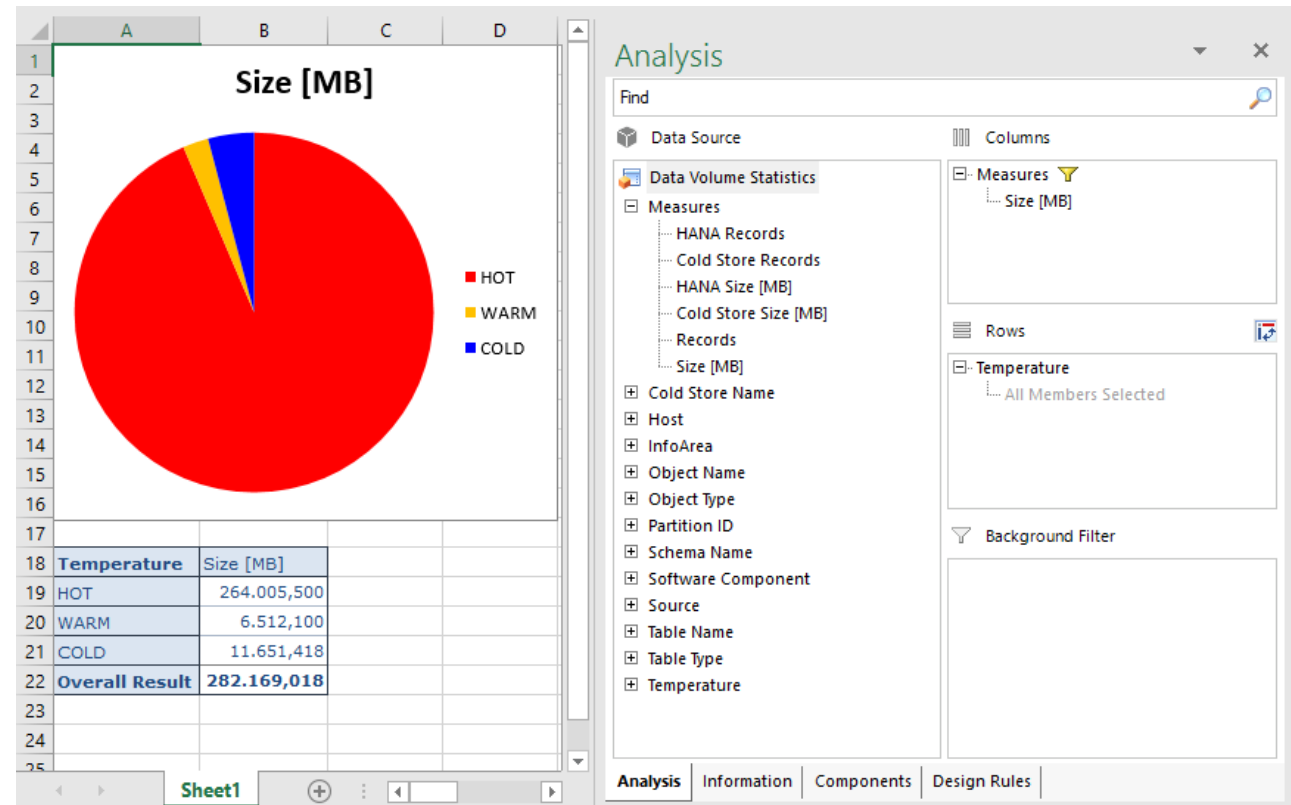
- Temperature adjustment by ABAP program using API RSO_ADSO.DTO_API.SET_TEMPS embedded in process chain with subsequent execution of the actual temperature change via process chain type DTO Adjustment
- For more extensive automation additional DTO API's are available
 - **RSO_ADSO.DTO_API.GET_INFO**
Reads the partition temperatures for a list of DataStore Objects
 - **RSO_ADSO.DTO_API.EXECUTE**
Executes the DTO adjustment run immediately



Data Tiering Optimization (DTO) - Performance and Monitoring

Enhanced Data Volume Statistics


- Data volume information to monitor the size of SAP BW/4HANA persistency objects including:
 - Table size in SAP HANA and the cold store
 - Raw size in SAP HANA and the cold store
 - Number of records in SAP HANA and in the cold store
 - Archiving rate (percentage share of archived data)
- Based on CDS technology for easier and simplified operations
 - Installation, loading and activation is NOT required
 - Real-time access
 - Analyze the entire volume of data in the system at the current point in time
- Cube View Rv_C_HanaNlsDbStat based on CDS View Rv_C_HanaNlsDbStatQuery*



*For more information please review the [SAP BW/4HANA Help page\(s\)](#)

SAP BW/4HANA Data Tiering Optimization

More information



SAP BW/4HANA Data Tiering Optimization (DTO)

Important Information - Setup and Installation

General Information

- [Blog SAP BW/4HANA Data Tiering Optimization](#)
- SAP Note [2296290](#) - New Sizing Report for BW/4HANA
- Reports include extension node sizing with “overload factor” for warm-data (planned: asymmetric extension node)

General Guidelines on the use of SAP HANA Native Storage Extensions

- SAP Note [2927591](#) - SAP HANA Native Storage Extension 2.0 SPS 05 Functional Restrictions
- SAP Note [2799997](#) - FAQ: SAP HANA Native Storage Extension (NSE)
- SAP Note [2374652](#) - Handling Very Large Data Volumes in SAP BW/4HANA

Warm Store with SAP HANA Extension Node

- SAP Note [2343647](#) How-To: Configuring SAP HANA for the BW Extension Node
- SAP Note [2453736](#) How-To: Configuring SAP HANA for SAP BW Extension Node in SAP HANA 2.0
- SAP Note [2644438](#) SAP HANA Extension Node – Master Release Note
- [SAP HANA Extension Node – Technical Overview presentation](#)

SAP BW/4HANA Data Tiering Optimization (DTO)

Important Information - Set Up and Installation

Cold Store with SAP IQ (fka Nearline Store)

- [SAP First Guidance - SAP BW: Implementation SAP-NLS/SDA with SAP IQ16.0](#)
 - Please note that the respective Remote Source (HANA SDA) to IQ (ODBC) needs to be “readwrite” enabled
- [Configuring SAP IQ as a Near-Line Storage Solution](#)
- SAP Note [1796393](#) - SAP BW near-line solution with SAP IQ
- [SAP IQ 16 Hardware Sizing Guide](#)
- [End-to-end Document — SAP IQ as a Near-line Storage Solution for SAP BW](#)

Cold Store with Hadoop File System (HDFS)

- SAP Note [2363218](#) - Hadoop NLS: Information, Recommendations and Limitations

SAP BW/4HANA Data Tiering Optimization (DTO)

Important Information – Implementation and Administration

General Information

- [SAP Help Portal: Data Tiering Optimization in SAP BW/4HANA](#)
- SAP Note [2517460](#) - SAP BW/4HANA Data Tiering Optimization: Information, Recommendations and Limitations

Warm Store with SAP HANA Extension Node

- SAP Note [2486706](#) FAQ - SAP BW/4HANA AND SAP BW-ON-HANA WITH EXTENSION NODES
- SAP Note [2643763](#) FAQ – SAP HANA Extension Node for SAP HANA native applications
- [SAP Help Portal: Using an Extension Node for Warm Data](#)
- [SAP Help Portal: Data Temperature: Extension Nodes](#)
- <https://blogs.sap.com/2016/04/14/important-update-data-lifecycle-management-for-bw-on-hana>
- <https://scn.sap.com/community/bw-hana/blog/2016/04/26/more-details--hana-extension-nodes-for-bw-on-hana>
- SAP Note [2334091](#) BW/4HANA: Table Placement and Landscape Redistribution
- SAP Note [2317200](#) Data lifecycle management for BW on SAP HANA and extension nodes

SAP BW/4HANA Data Tiering Optimization (DTO)

Important Information – Implementation and Administration

Cold Store with SAP IQ (fka Nearline Store)

- SAP Note [2165650](#) - FAQ: BW Near-Line Storage / Cold Store with HANA Smart Data Access
- SAP Note [2100962](#) FAQ: BW Near-Line Storage with HANA Smart Data Access: Query Performance
- SAP Note [1999431](#) SIQ: Setting up SSL for connections to IQ
- SAP Note [2133194](#) Can SAP IQ run in a cloud environment?

Cold Store with Hadoop File System (HDFS)

- SAP Note [2363218](#) - Hadoop NLS: Information, Recommendations and Limitations
- SAP Note [2165650](#) - FAQ: BW Near-Line Storage / Cold Store with HANA Smart Data Access

Thank you.

Contact information:

Product Management, SAP Data Warehousing



Abbreviations

ABAP	Advanced Business Application Programming	DWaaS	Data Warehouse as a Service	ODQ	Operational Delta Queue
ADSO	Advanced DataStore Object	ECC	Enterprise Core Component	ODS	Operational DataStore
AFL	Application Function Library	EDW	Enterprise Data Warehouse	OLAP	Online Analytic Processing
AGS	SAP Active Global Support	EIM	Enterprise Information Management	PSA	Persistent Staging Area
ASE	SAP Adaptive Server Enterprise Database	ELT	Extract, Load, Transform	RDBMS	Relational Database Management System
AWS	Amazon Web Services	ERP	Enterprise Resource Planning	SDI	SAP HANA smart data integration
BAPI	Business Application Programming Interface	ETL	Extract, Load, Transform	SLO	System Landscape Optimization
BEx	Business Explorer	FI,CO,SD,MM,HR	Financials, Controlling, Sales & Distribution, Material Management,, Human Resources	SLT	SAP Landscape Transformation
BI	Business Intelligence	HAP	SAP HANA Analytic Process	SOAP	Simple Object Access Protocol
BW	Business Warehouse	HEC	SAP HANA Enterprise Cloud	SP	Support Package
C4C	Cloud for Customers	HTML	Hypertext Markup Language	SPS	Support Package Stack
CRM	Customer Relationship Management	IQ	SAP IQ Database	SQL	Structured Query Language
DB	Database (Connect)	IoT	Internet of Things	SUM	Software Update Manager
DLM	Data Lifecycle Management	LSA	Layered Scalable Architecture	SWPM	Software Provisioning Manager
DMO	Database Migration Option	LSA++	Layered Scalable Architecture for SAP HANA	UD	Universal Data (Connect)
DSO	DataStore Object	ML	Machine Learning	UI	User Interface
DW	Data Warehouse	NLS	Near-line Storage	UI5	SAP UI Development Toolkit for HTML5
DWH	Data Warehouse	ODP	Operational Data Provisioning	UX	User Experience