



From Near-line Storage to Cold Store

Data Lifecycle Strategy for BW/4HANA

Dr. Peter Zimmerer, SAP
Roland Kramer, SAP
June 22nd, 2018



Disclaimer

Die Informationen in dieser Präsentation sind vertraulich und urheberrechtlich geschützt und dürfen nicht ohne Genehmigung von SAP offengelegt werden. Diese Präsentation unterliegt weder Ihrem Lizenzvertrag noch einer anderen Service- oder Subskriptionsvereinbarung mit SAP. SAP ist in keiner Weise verpflichtet, in dieser Präsentation oder einem dazugehörigen Dokument dargestellte Geschäftsabläufe zu verfolgen oder hierin wiedergegebene Funktionen zu entwickeln oder zu veröffentlichen.

Diese Präsentation oder jedes dazugehörige Dokument über die Strategie von SAP und mögliche zukünftige Entwicklungen, Ausrichtungen und Funktionen von Produkten und/oder Plattformen kann von SAP jederzeit aus beliebigen Gründen ohne vorherige Ankündigung geändert werden. Die Informationen in dieser Präsentation stellen keinerlei Zusage, Versprechen oder rechtliche Verpflichtung zur Auslieferung von Materialien, Code oder Funktionen dar. Diese Präsentation wird ohne jegliche Gewähr, weder ausdrücklich noch stillschweigend, bereitgestellt. Dies gilt insbesondere, hinsichtlich der Gewährleistung der Marktgängigkeit und der Eignung für einen bestimmten Zweck sowie für die Gewährleistung der Nichtverletzung geltenden Rechts. Diese Präsentation dient zu Informationszwecken und darf nicht in einen Vertrag eingebunden werden. SAP übernimmt keine Verantwortung für Fehler oder Unvollständigkeiten in dieser Präsentation, es sei denn, solche Schäden wurden von SAP vorsätzlich oder grob fahrlässig verursacht. Sämtliche vorausschauenden Aussagen unterliegen verschiedenen Risiken und Unsicherheiten, durch die die tatsächlichen Ergebnisse von den Erwartungen abweichen können.

Die vorausschauenden Aussagen geben die Sicht zu dem Zeitpunkt wieder, zu dem sie getätigt wurden. Dem Leser wird empfohlen, diesen Aussagen kein übertriebenes Vertrauen zu schenken und sich bei Kaufentscheidungen nicht auf sie zu stützen.

Agenda

Looking back

- Milestones

Data Tiering Optimization (DTO) with BW/4HANA 1.0 Feature Pack 8

- The concept
- Supported Cold Stores
- What's the difference to classical Near-line Storage (NLS)?

Looking today

- SAP NLS with IQ and the conversion to BW/4HANA

Looking forward

- Roadmap

Agenda

Looking back

- Milestones

Data Tiering Optimization (DTO) with BW/4HANA 1.0 Feature Pack 8

- The concept
- Supported Cold Stores
- What's the difference to classical Near-line Storage (NLS)?

Looking today

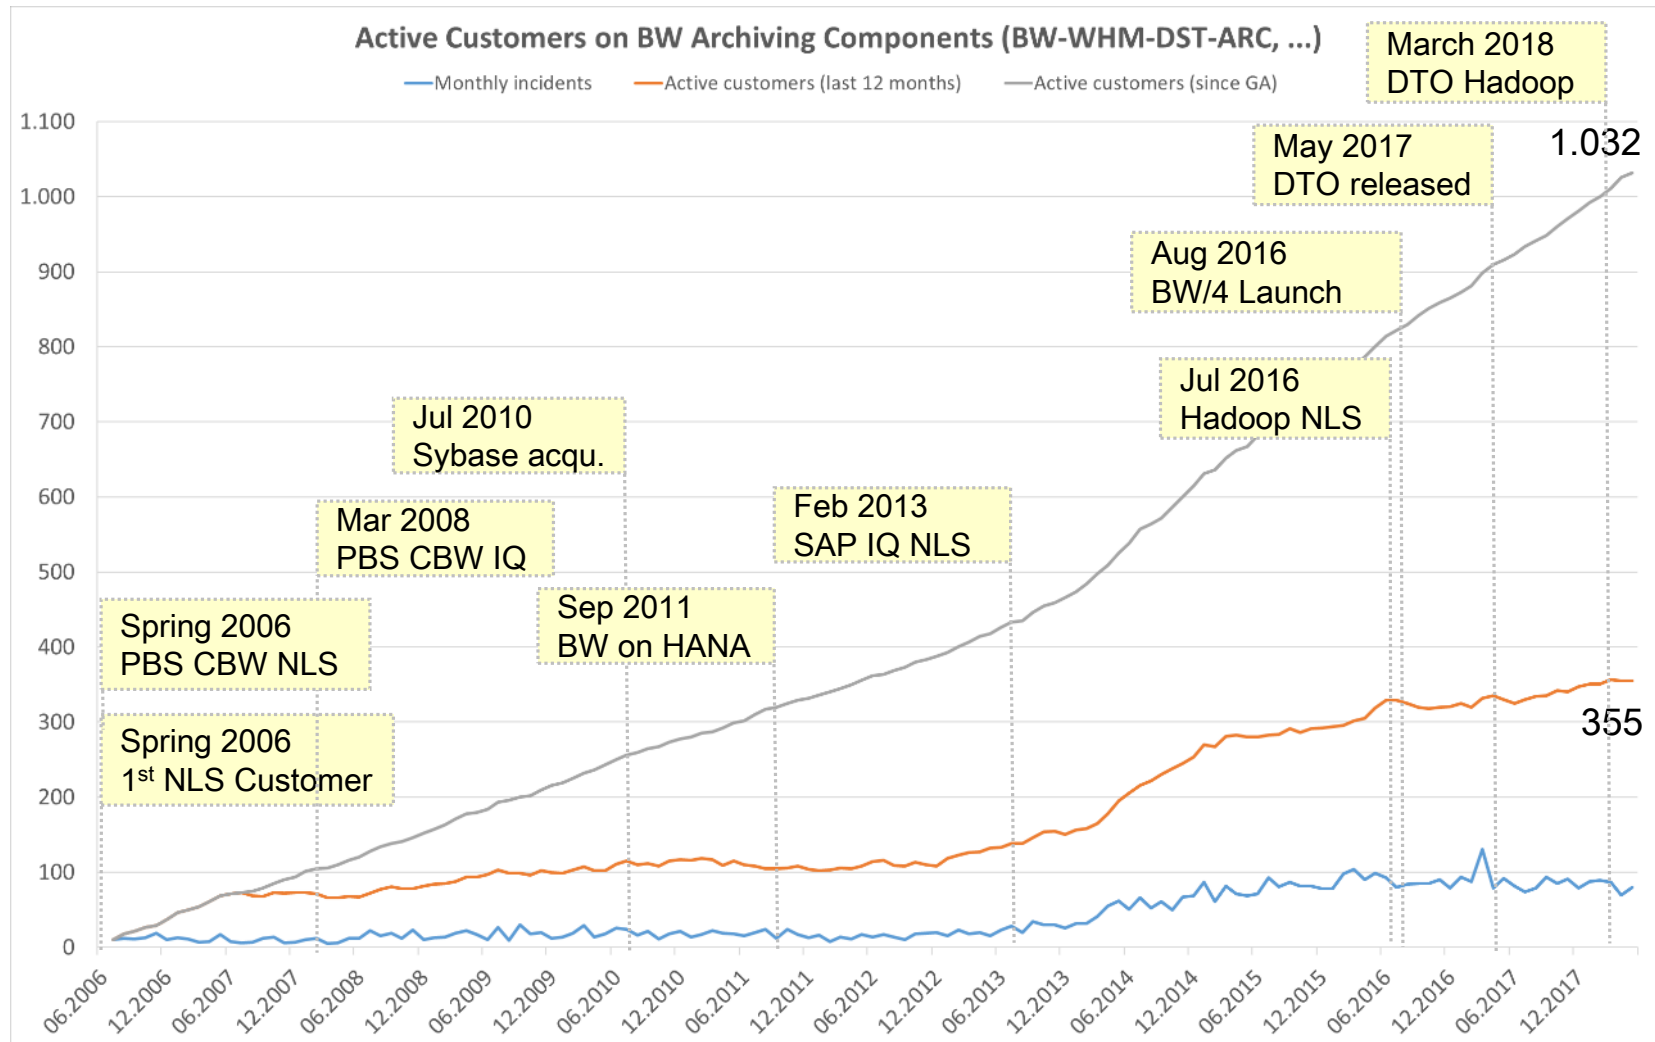
- SAP NLS with IQ and the conversion to BW/4

Looking forward

- Roadmap

12 Years of BW Nearline Storage / 2 Years BW/4HANA / 1 Year DTO

Celebrating 1.000 customers



Agenda

Looking back

- Milestones

Data Tiering Optimization with BW/4HANA 1.0 Feature Pack 8

- The concept
- Supported Cold Stores
- What's the difference to classical Near-line Storage?

Looking today

- SAP NLS with IQ and the conversion to BW/4

Looking forward

- Roadmap


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 HANA Data Tiering

The Data Growth Challenge

Strong coupling between data and hardware

- More data → more RAM → more CPUs
- Impacts on hardware configuration

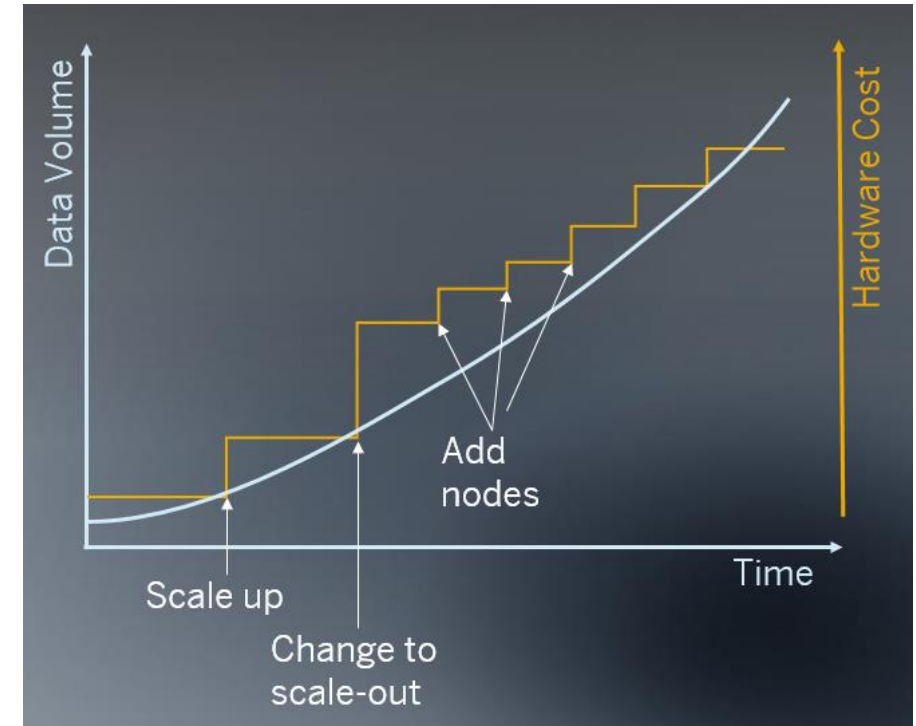
Growing SAP HANA hardware

- Scale-up ends at 8 (analytical) or 8-20 (transactional) TB
- Larger systems require scale-out hardware

Challenges for Very Large Systems

- Does all my data justify hardware and license cost?
- Is it possible to de-couple hardware growth from data growth?

SAP HANA as In-Memory Database



SAP BW/4HANA Data Tiering Optimization (DTO)

Motivation and Overview

One concept for hot, warm and cold data

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

Easy and central definition and implementation

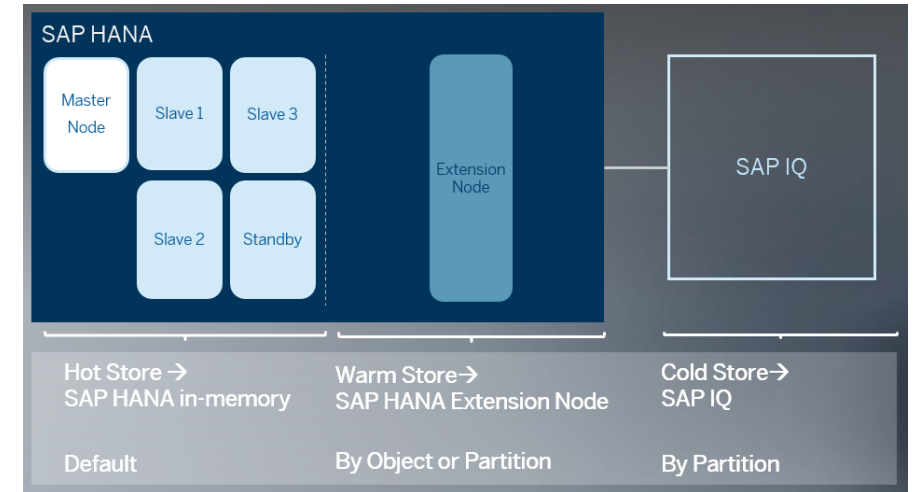
- Data Temperature defined in Advanced DataStore Object only
- No additional configuration of Data Archiving Processes

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 or co-existence with existing SAP BW NLS IQ / Hadoop approach as of sharing some central technical concepts for cold data storage

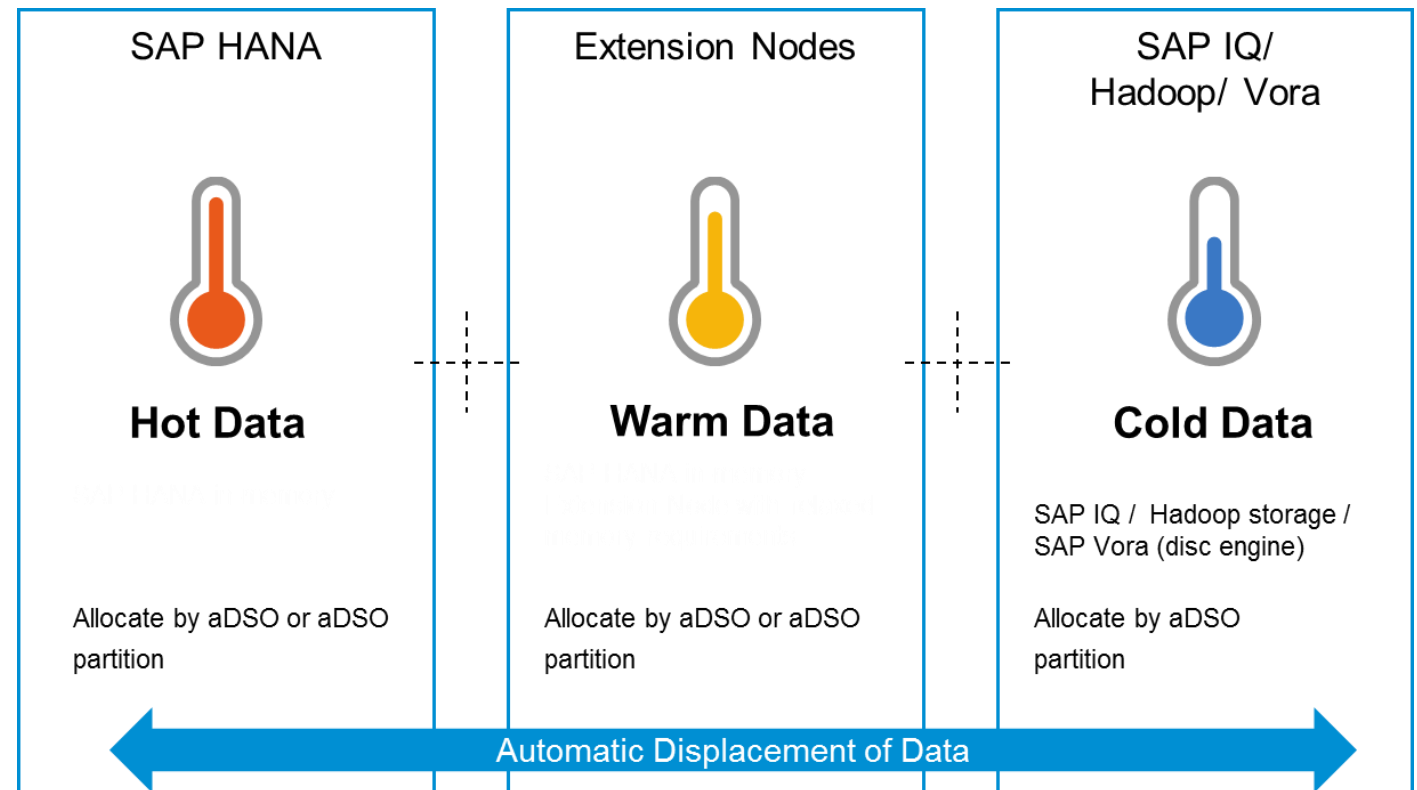


Simplicity: Simplifying Data Lifecycle Management



New Storage Options with FP08

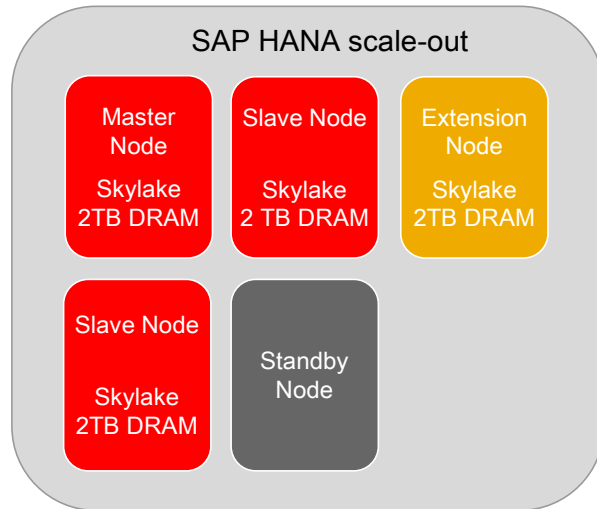
- Support for cold storage in the Hadoop File System (HDFS)
- Support for cold storage in SAP HANA Vora (disk engine)*
- Support for multiple extension nodes



*Validation with upcoming Vora release still pending (see note [2608405](#) for updates)

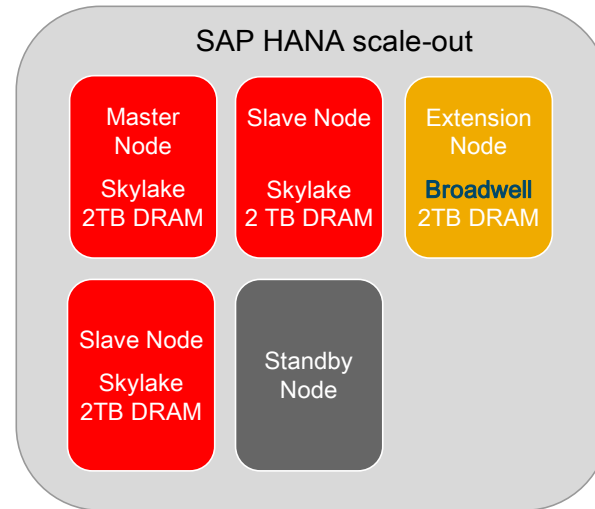
Basic Deployment Options

Symmetric



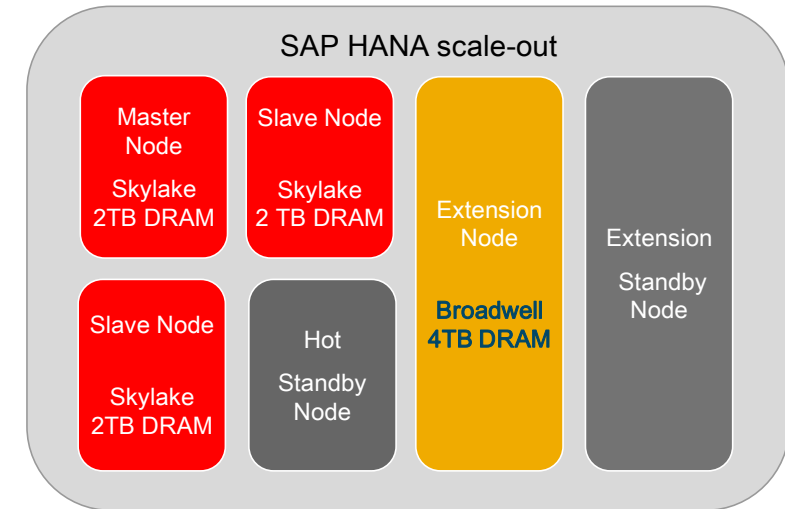
- Extension Node has identical HW-sizing like Hot-Nodes
- Minimal re-configuration required
- Standby node can be shared by all nodes
- Extension node allows larger data footprint (default: 100% of DRAM)
- Example data footprint: {3x 1TB (Hot Nodes) + 1x 2TB EN } = 5TB
- BW (>=HANA 1 SPS12) and native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

SAP BW/4HANA Data Tiering Optimization (DTO)

Implementation Steps

Meta Data Maintenance in Advanced DataStore Object

General: RHADSO1

DataStore Object (advanced)

Technical Name: RHADSO1
Description: RHADSO1

External SAP HANA View

Modeling Properties

Activation:

- Activate Data
- Write Change Log
 - Keep Inbound Data, Extract from Inbound Table
 - Unique Data Records
 - Snapshot Support

Special Types:

- Direct Update
- All Characteristics are Key, Reporting on Union of Inbound and Active Table
- Planning Mode
- Inventory

The selected properties match template 'Data warehouse layer - delta calculation'.

Data Tiering Optimization:

Temperature Schema

- Standard Tier (Hot)
- Extension Tier (Warm)
- External Tier (Cold)

Temperature Maintenance (Hot and Warm)

- On Object Level
- On Partition Level

Settings: RHADSO1

Partitions

Field: OCALMONTH

Lower Bound	Upper Bound
01.2016	12.2016
01.2014	12.2015
01.2010	12.2013

Temperature Schema ①

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...
RHADSO1	1	OCALMONTH	Hot	Hot			()	01.2016	12.2018	
RHADSO1	2	OCALMONTH	Warm	Hot			()	01.2014	12.2015	
RHADSO1	3	OCALMONTH	Cold	Hot			()	01.2010	12.2013	
RHADSO1	4	OCALMONTH	Hot	Hot						

Partition Temperature Maintenance ③

Plan

Data Tiering Adjustment

Execute Execute in Background Show Logs

DataStore Object: Name: RHADSO1

Partition Parameter:

Run in Simulation Mode

Parallelization factor:

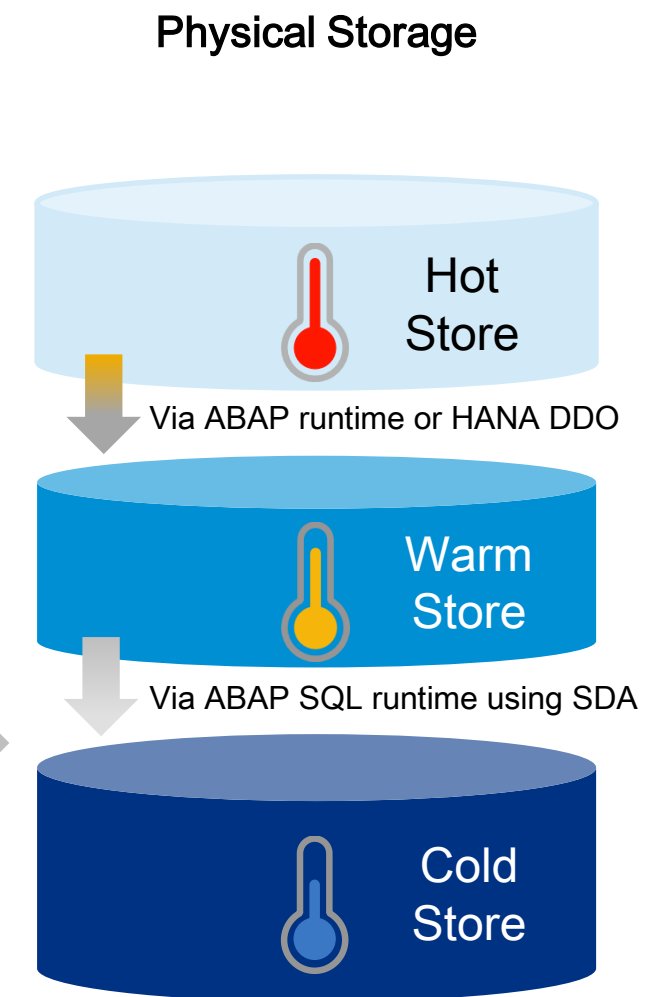
Adjust

Repair

Repair and Adjust

Data Tiering Optimization Job ④

Changes to Temperature (only) = Data Tiering Optimization



SAP BW/4HANA Data Tiering Optimization (DTO)

Overview – Database Perspective

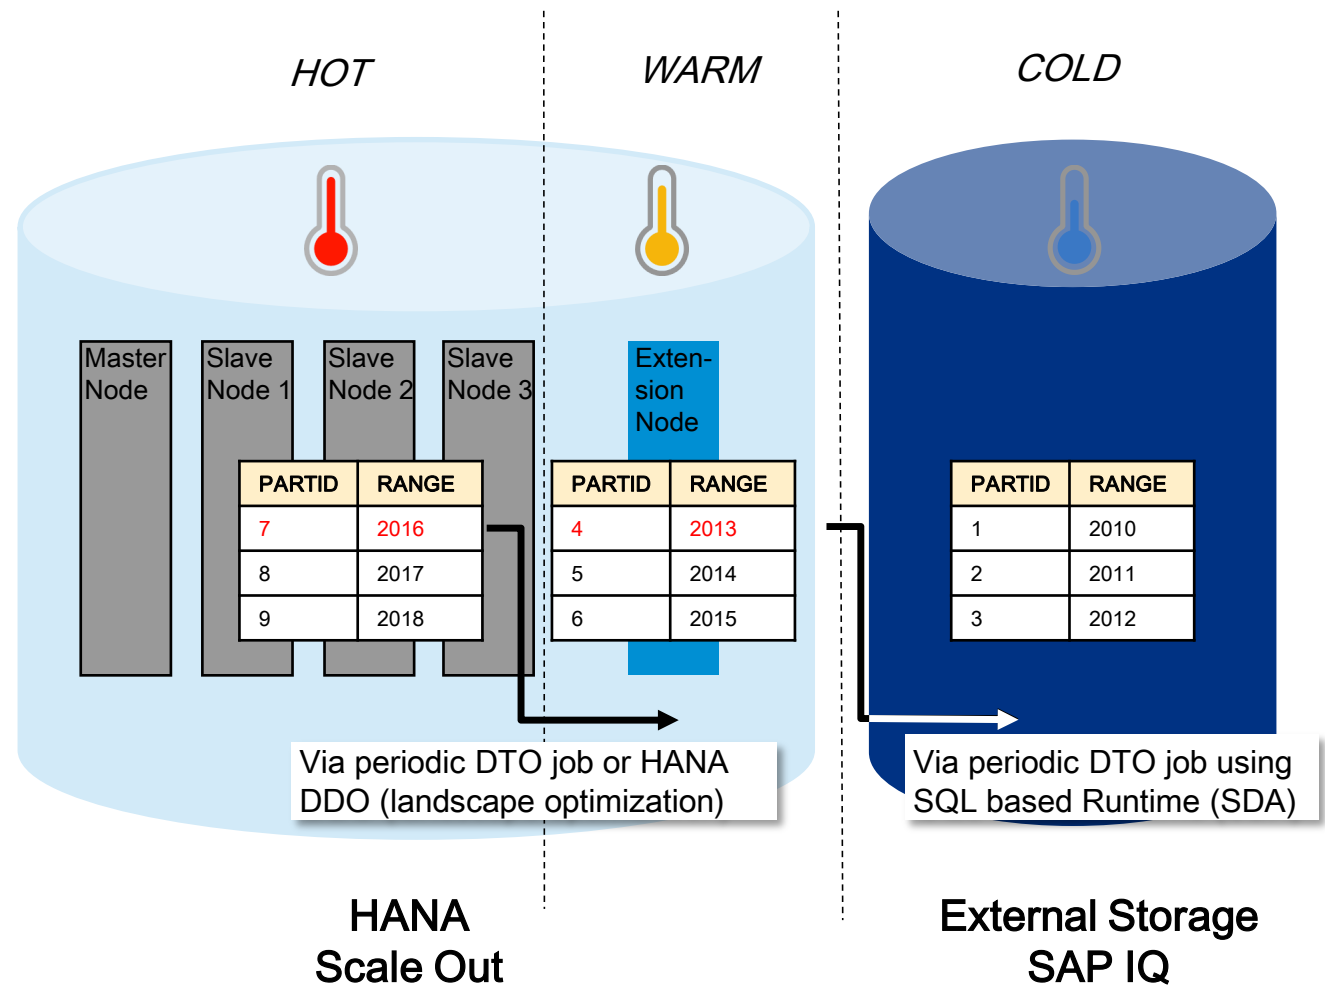
Partition Temperature in SAP BW/4HANA (“Plan”)

ADSO Data Tiering Optimization (DTO)

DSO (adv.)	Partiti	Partition Field	Planned Te	Current	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


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

Physical Storage (“Actual”)



IQ + Hadoop/Spark - Data Movement Performance (preliminary results)

Load ADSO Data to Archive



	10Mio	20Mio	30Mio	40Mio	50Mio	60Mio	70Mio	80Mio	90Mio	100Mio
[CGC14] : Archived to SybaseIQ (NLS)	382,979	615,385	596,686	605,043	608,109	690,096	709,860	684,086	775,120	731,708
[CGC14] : Archived to Hadoop (textfile_ABAP)	191,631	303,897	367,467	401,250	440,249	450,189	455,866	475,430	469,772	488,705
[CGC18] : Archived to SybaseIQ (NLS)	272,727	529,412	755,245	804,470	833,334	710,527	842,810	802,229	778,847	839,162
[CGC18] : Archived to SybaseIQ (SDA)	45,743	47,493	46,875	47,493	50,322	50,931	50,139	50,535	50,467	52,083
[CGC18] : Archived to Hadoop (textfile_ABAP)	189,474	259,928	180,602	302,521	309,279	319,527	326,003	331,798	331,967	341,556
[CGC18] : Archived to Hadoop (textfile_HANA)	292,683	411,429	473,685	521,740	555,556	580,646	584,687	607,595	617,143	627,178
[CGC18] : Archived to Hadoop (orc_ABAP)	218,182	289,157	323,354	345,324	360,722	369,863	375,559	380,449	386,174	389,611
[CGC18] : Archived to Hadoop (orc_HANA)	74,534	138,196	216,000	266,174	289,389	318,115	352,448	357,320	351,410	396,913
[CGC18] : Archived to Hadoop (parquet_ABAP)	83,527	133,333	172,800	195,122	209,546	213,650	241,149	243,038	237,189	256,960
[CGC18] : Archived to Hadoop (parquet_HANA)	75,157	145,455	201,117	247,849	288,462	326,284	340,541	344,498	393,682	393,873

Size of archived Data

Hardware Info for Cold Store

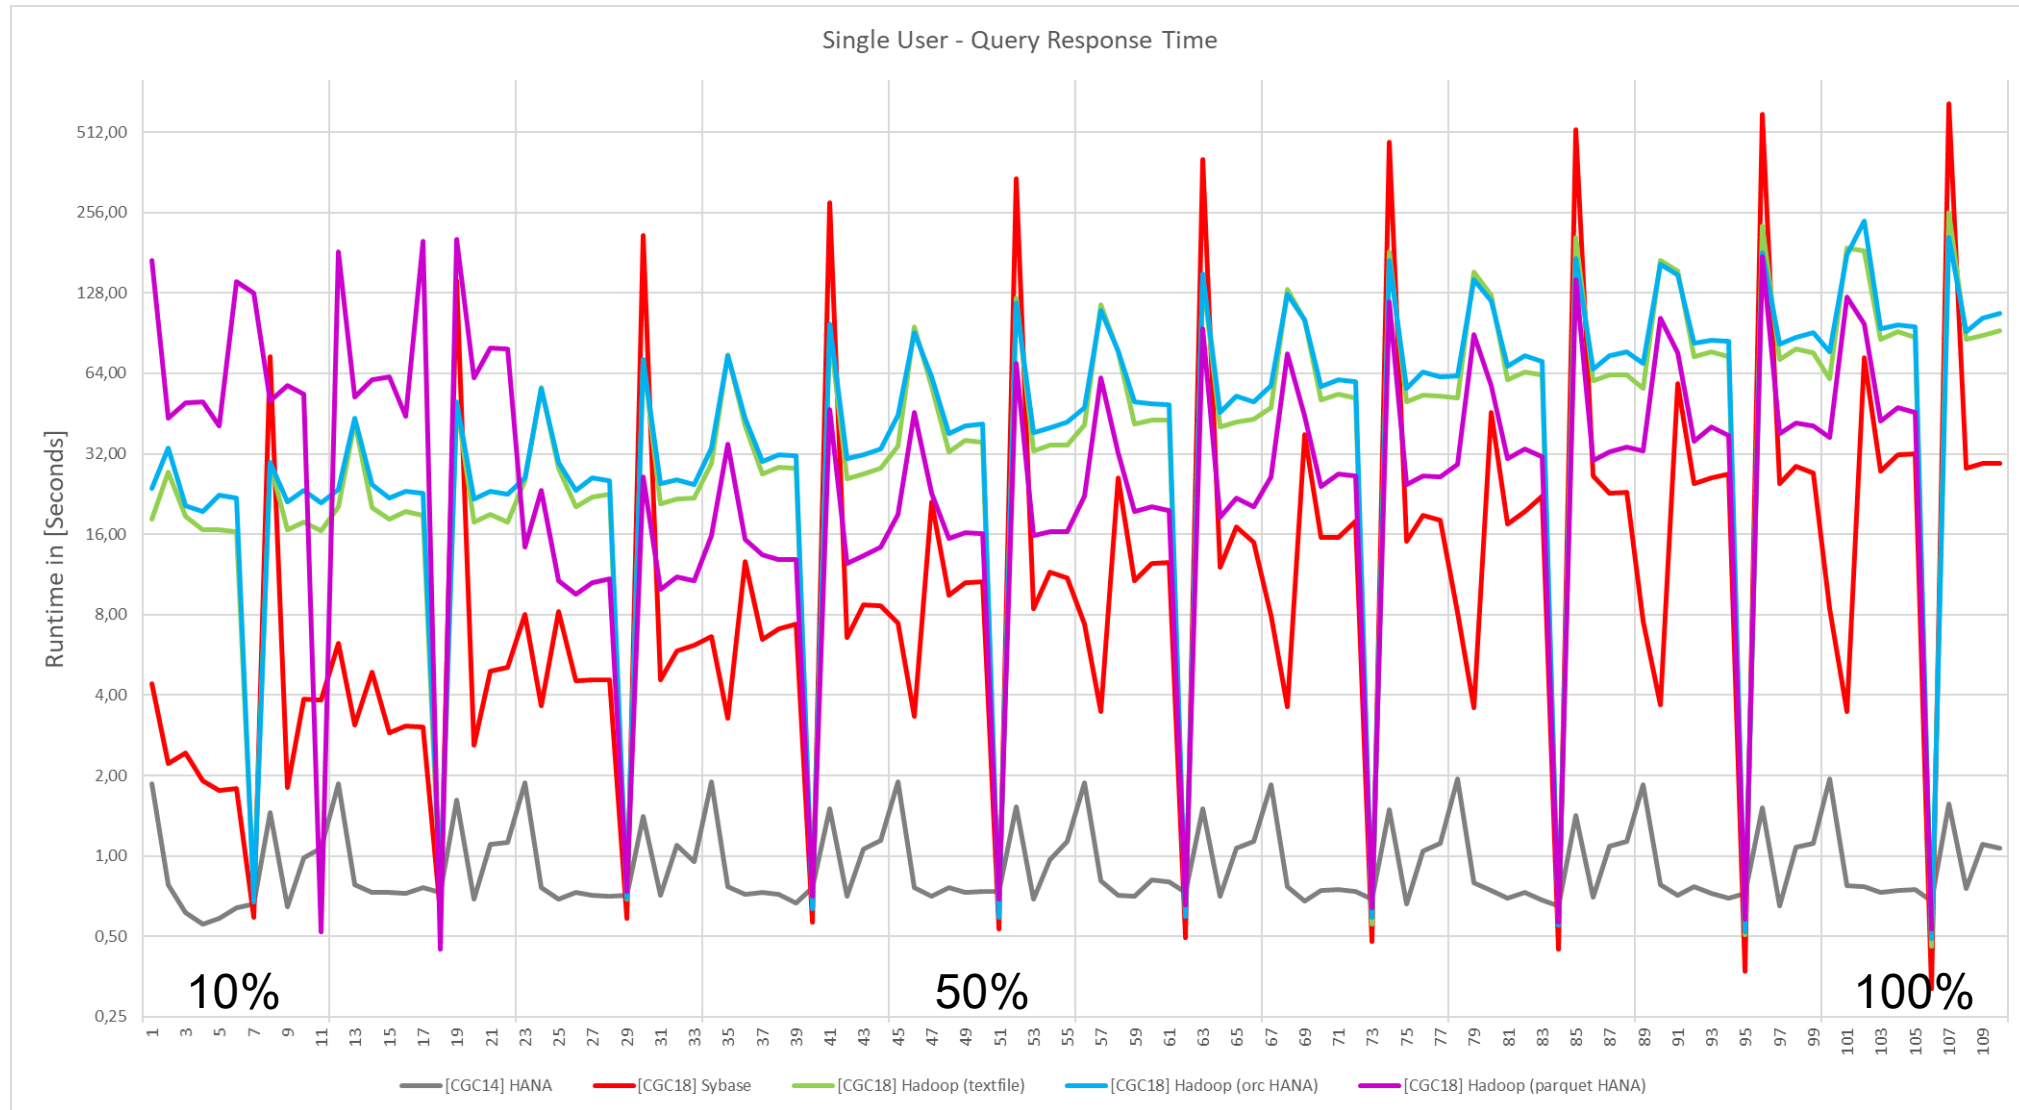
Architecture Intel Broadwell

CPU Cores 4*22 = 88 + HT

Total Memory 2 TByte

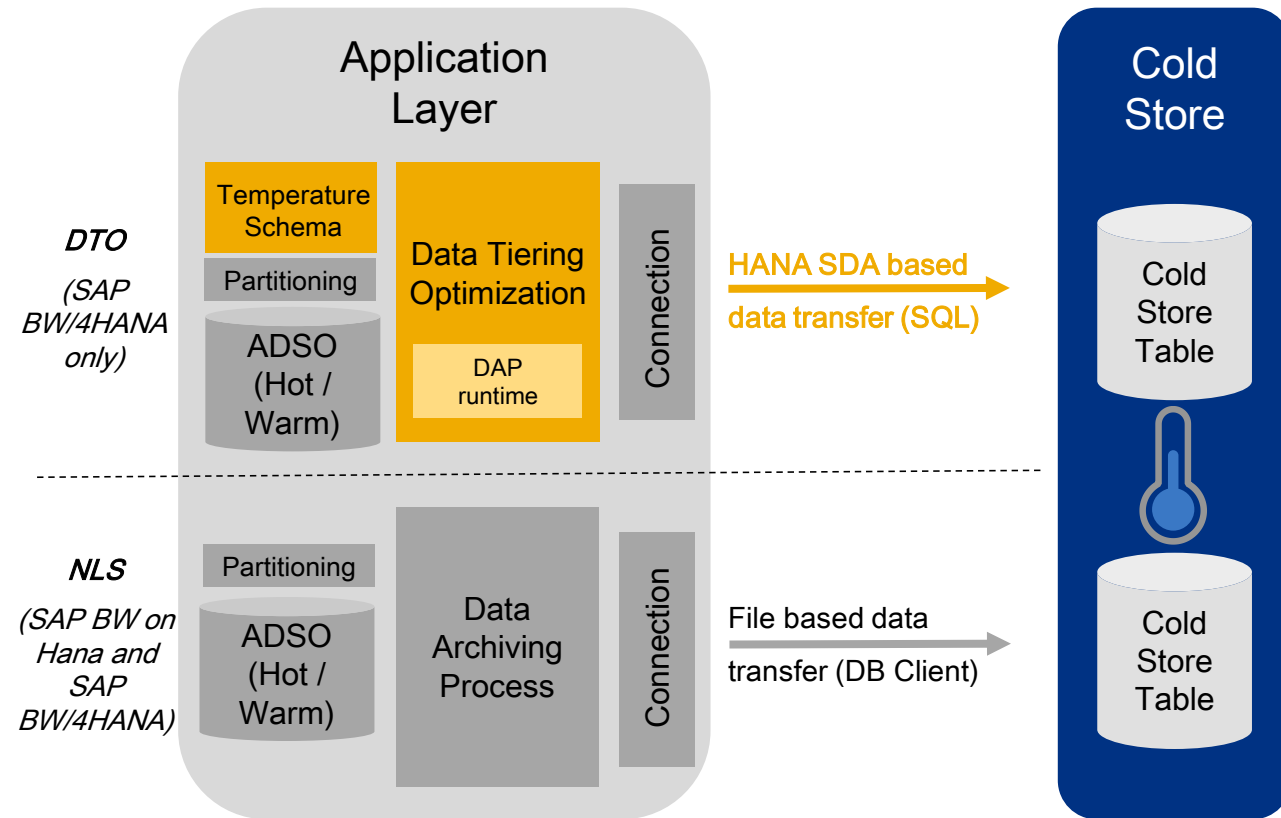
Component	Version
CGC14	2.00.012.01.1501753869 (fa/hana2sp01) HANA_2.0_SPS01_Rev12_RC1_500019_final_gcc6
CGC18	2.00.030.00.1522209842 (fa/hana2sp03) HANA_2.0_SPS03_Rev30_RC2_500005_gcc6
Hadoop/Spark	1.6.3.2.6.4.0-91-hadoop2.7.3.2.6.4.0-91
Spark Controller	2.0 SP02 PL0

IQ + Hadoop/Spark – Query Performance (prelimin. results)



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

Comparison to SAP BW Nearline Storage (NLS)



Simplification in Architecture is **highlighted**

Non highlighted objects (e.g. Cold Store Connection (fka NLS Connection) or Definition of Cold Store Table (fka NLS Table) remain **unchanged** (available product information still applies).

Common properties DTO to NLS

- Same Cold Store table layout (e.g. IQ: "ON"/"OR" table + "OU" view)
- Reuse of existing Near-line connection possible (SAP IQ NLS + Hadoop)

Implementation differences DTO to NLS

- Partitioning definition in ADSO becomes mandatory when using DTO with Cold Store
- DTO Cold Store partitions are limited to HANA range partitioning
 - Only single column/dimension
 - Only key column
 - Only single-value partitions and right-open intervals
- No maintenance and scheduling of data archiving processes
- New temperature schema definition required in ADSO

Functionality planned in DTO (but possible with NLS in SAP BW/4HANA)

- Definition of archive conditions for characteristics not in key of object
- Relative time selections
- Multiple cold store connections

For further details see SAP Note [2517460](https://www.sap.com/support/2517460)

SAP BW/4HANA Data Tiering Optimization (DTO)

Roadmap

Today

SAP BW/4HANA 1.0 FP08

- Support for cold storage in the Hadoop File System (HDFS)
- Support for cold storage in SAP HANA Vora (disk engine)*
- Implementation, Administration and Performance Improvements, e.g.
 - Mass partition creation support
 - Integration with Process Chains
- API for Temperature Maintenance that enables Data Tiering Automation Implementation
- Support for multiple extension nodes

Conversion Support

- Transfer of Classic DataStore Objects / InfoCubes to Advanced DataStore Object incl. Nearline Store

Mid Term – Estimated for Q1 2019**

SAP BW/4HANA 2.0

- Data Tiering Optimization Cockpit as Web User Interface
- Support for (exceptional) update operations to data in the cold store
- Coverage of additional Data Store Objects types in cold store (without Activation, Inventory)
- Enhanced Data Tiering Automation
 - Relative or rule based conditions
 - Data Access Statistics
- Make Cold Store connection configurable on object level
- Simplification of temperature schema

Future Direction

- Enhanced Data Tiering Automation
 - Statistics based data movements
 - Dynamic Partitioning Schema
- Conversion of Data Archiving Processes (NLS) to SAP BW/4HANA Data Tiering Optimization
- Multi-dimensional partitioning scheme
- Deeper SAP Data Hub Integration

*Validation with upcoming Vora release still pending

** This is the current state of planning and may be changed by SAP at any time

Agenda

Looking back

- Milestones

Data Tiering Optimization (DTO) with BW/4HANA 1.0 Feature Pack 8

- The concept
- Supported Cold Stores
- What's the difference to classical Near-line Storage (NLS)?

Looking today

- SAP NLS with IQ and the conversion to BW/4

Looking forward

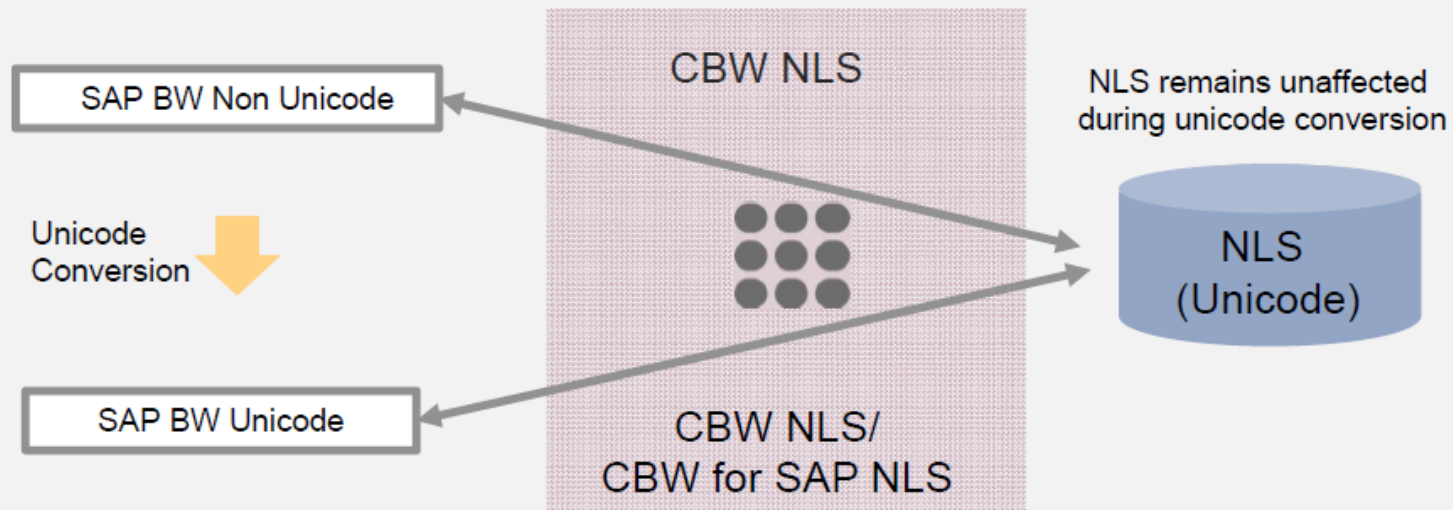
- Roadmap

CBW NLS for SAP IQ - using the “classical PBS NLS Adapter”

Digression: Optimization of Unicode Conversion

Unicode is a prerequisite for the migration to BW on HANA.

Most of the data is removed from non unicode BW to NLS using CBW NLS



Data is accessible via both CBW NLS and CBW for SAP NLS after the unicode conversion

CBW NLS for SAP Nearline Storage with IQ - switch to the SAP Interface

Data Transfer Procedure


Source: Former existing NLS connection used so far for archiving

Target: New NLS connection which should be used for archiving after the migration

Problem: In all DAPs source connection -> no access to the target connection until the completion of the migration

Solution: Virtual connection which assigns all accesses to NLS during the migration

Nearline connection			
Source	CBW_IQ	Target	H75NLS
log. destination	PBSRIQ_E_SAPNLSDB01@G0NZ0	log. destination	



NLS Access	Connection		User group
... to read	H75NLS		2
... to write	CBW_IQ		

CBW NLS for SAP Nearline Storage with IQ - switch to the SAP Interface

BW Archiving: Connection to a Nearline Storage

Name of Class	/PBS/CL_NLSA_CONNECTION
Connection Mode	Productive Mode
DB Connection	SAPNLS
Destination	PBS_RIQ
Conn. Parameter	LOAD_STRIPE_SIZE=4;LOAD_STRIPE_WIDTH=4;BACKUP_REQUEST_EVENT=BACKUP_REQUEST

BW Archiving: Connection to a Nearline Storage

Name of Class	CL_RSDA_SYB_CONNECTION
Connection Mode	Productive Mode
DB Connection	SAPNLS
Destination	
Conn. Parameter	LOAD_STRIPE_SIZE=4;LOAD_STRIPE_WIDTH=4;BACKUP_REQUEST_EVENT=BACKUP_REQUEST



Diagnosis

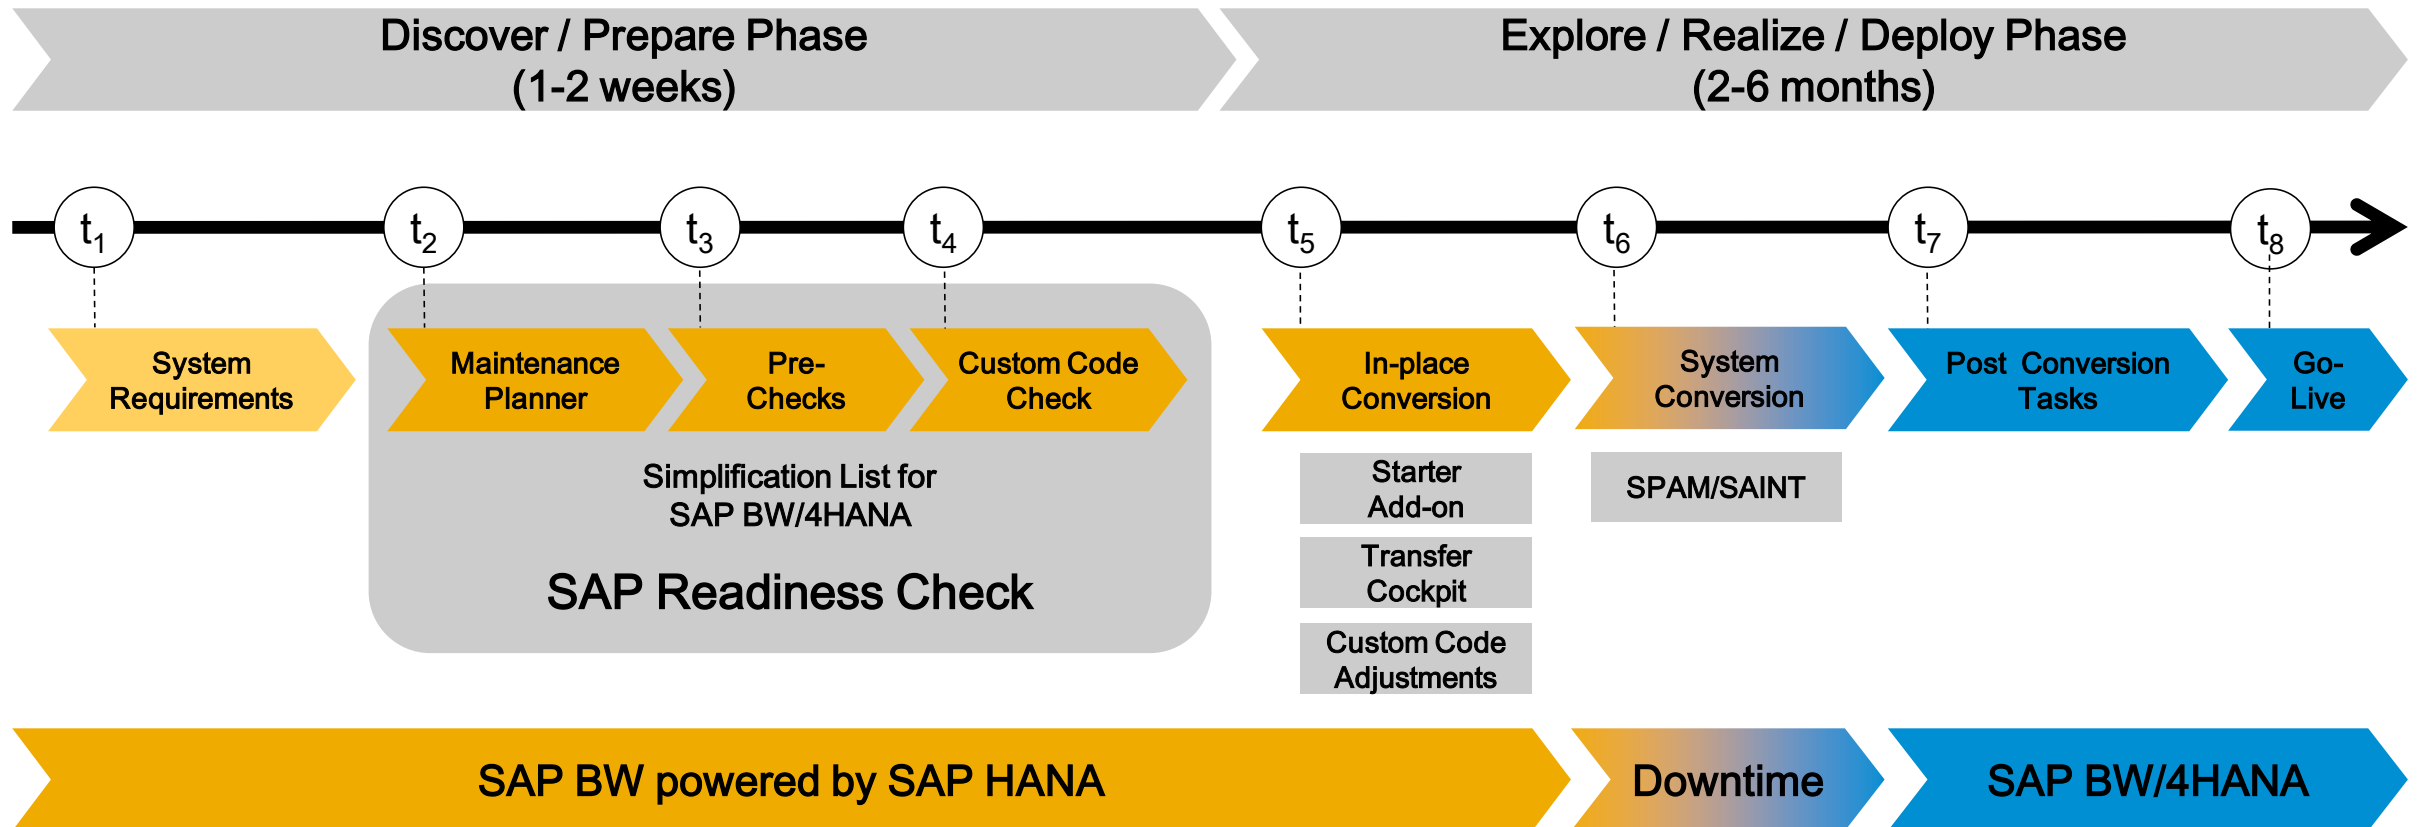

Component	Status	Release	Level	Patch	Short Description of Component
SAP_BW_NEARLINE_IQ	○○○	740	17	SAPKW74017	Near-line implementation for Sybase IQ

----- SP_IQSTATUS -----

Ty...	Message Text	LTxt
■	DB connection SAPNLS opened (work process 3 of server It5006_HBW_03)	
■	LOAD striping: Width = 4, size = 4 for 64 of 64 logical processors	
■	----- SP_IQSTATUS -----	?
■	SAP IQ (TM) Copyright (c) 1992-2016 by SAP AG or an SAP affili	
■	Version: 16.0.110.2335/10361/P/sp11.05/Enterprise Linux64 -	

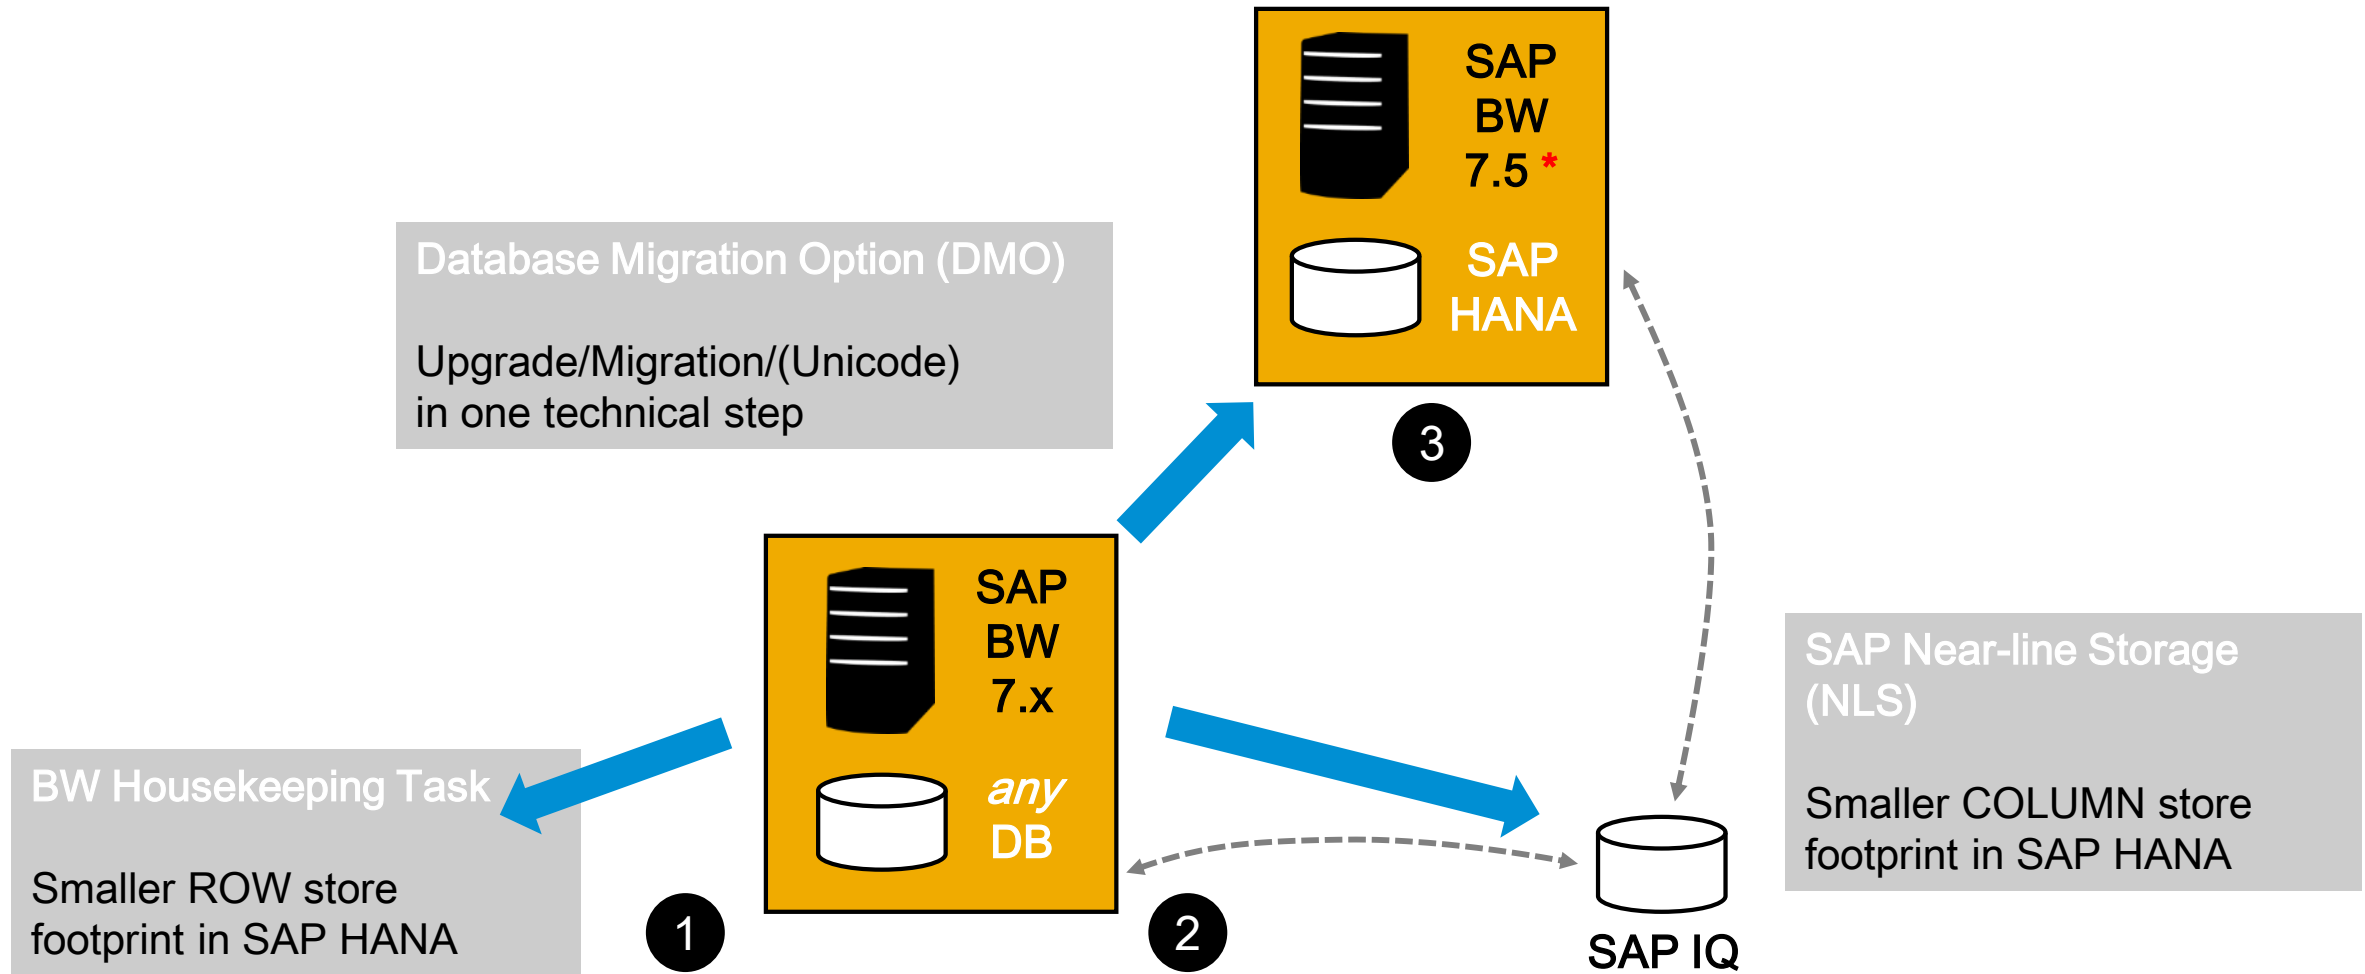

In-place Conversion

Basic Sequence



In-place Conversion

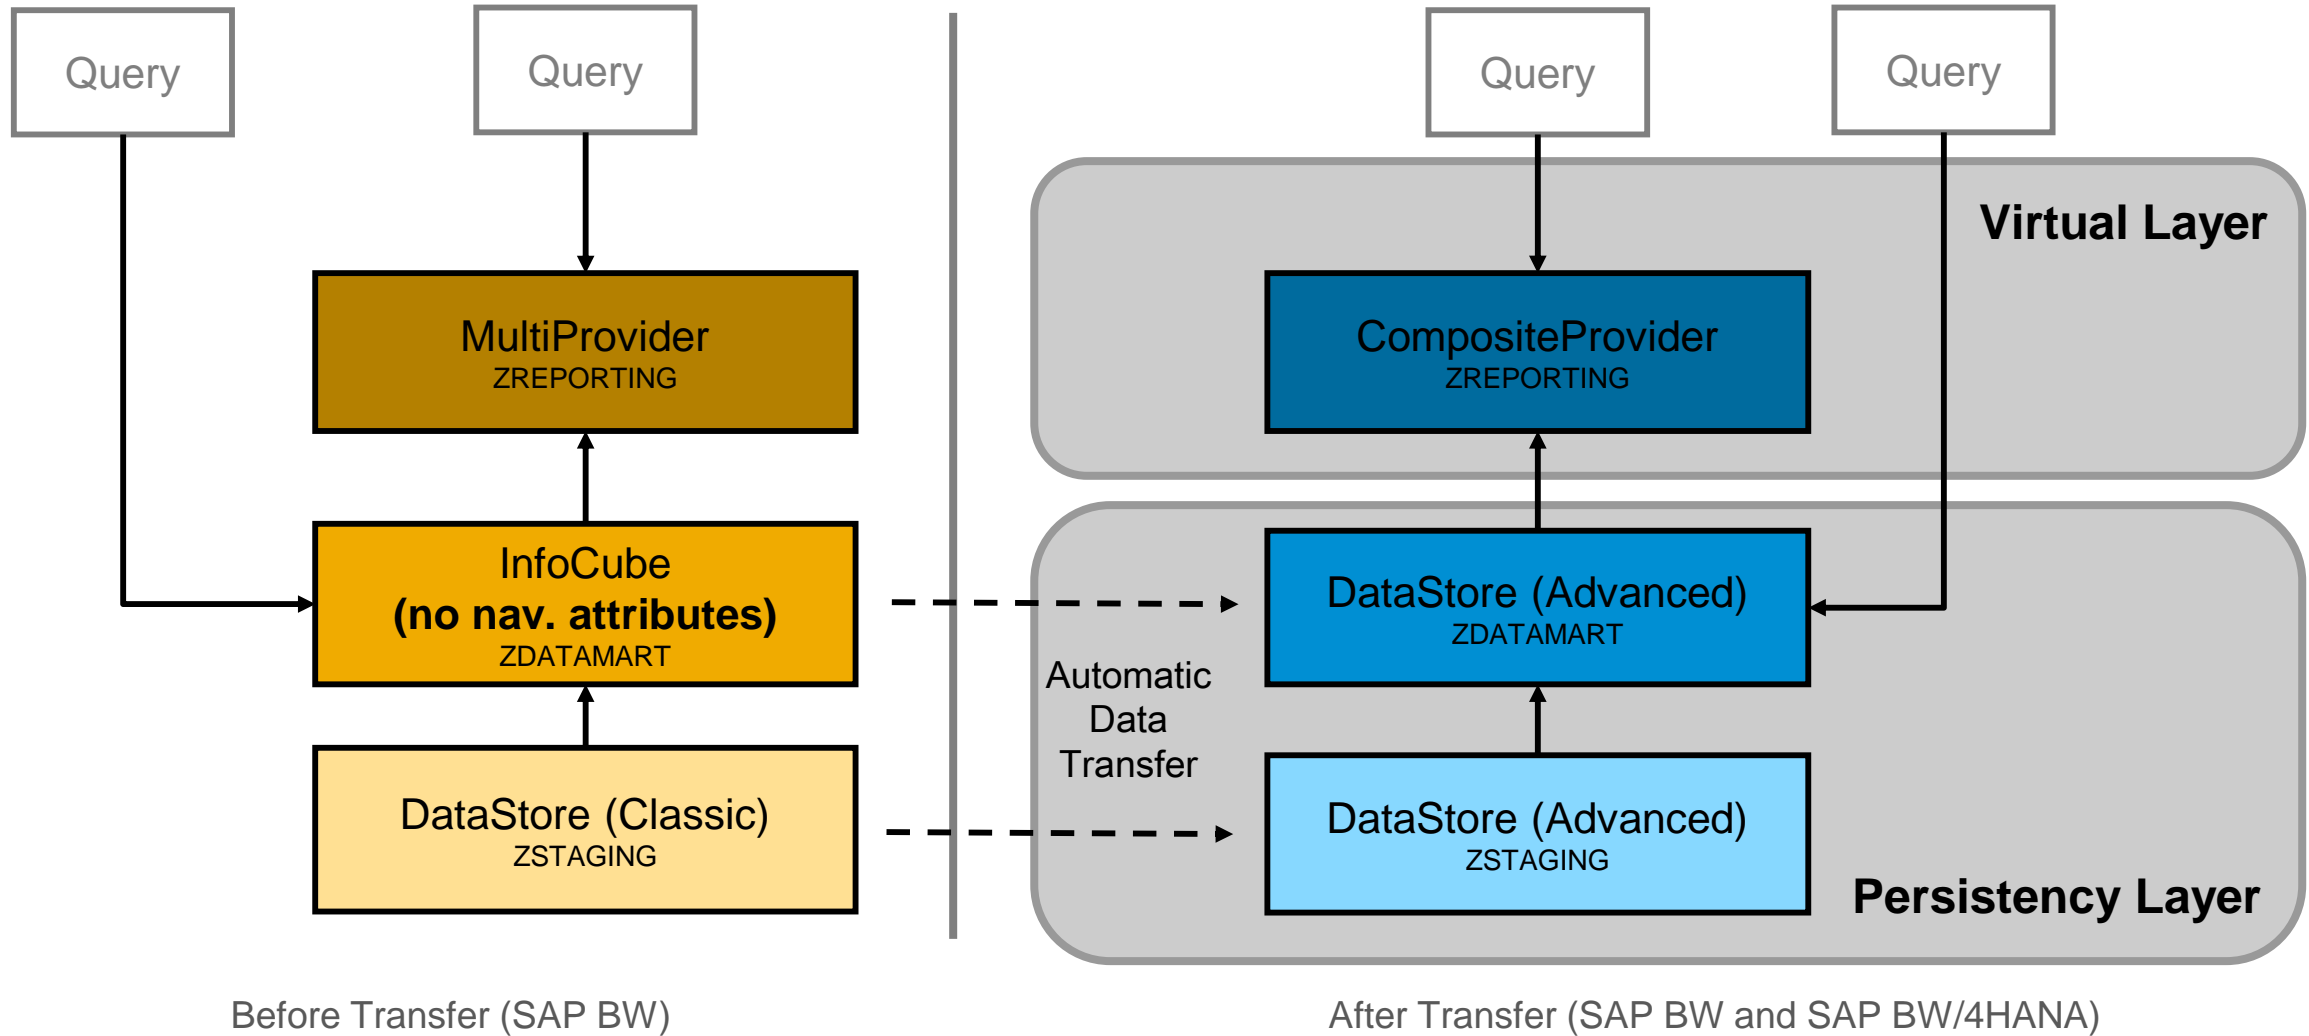
Migration to SAP BW powered by SAP HANA



* Original system is already Unicode with 7.0x

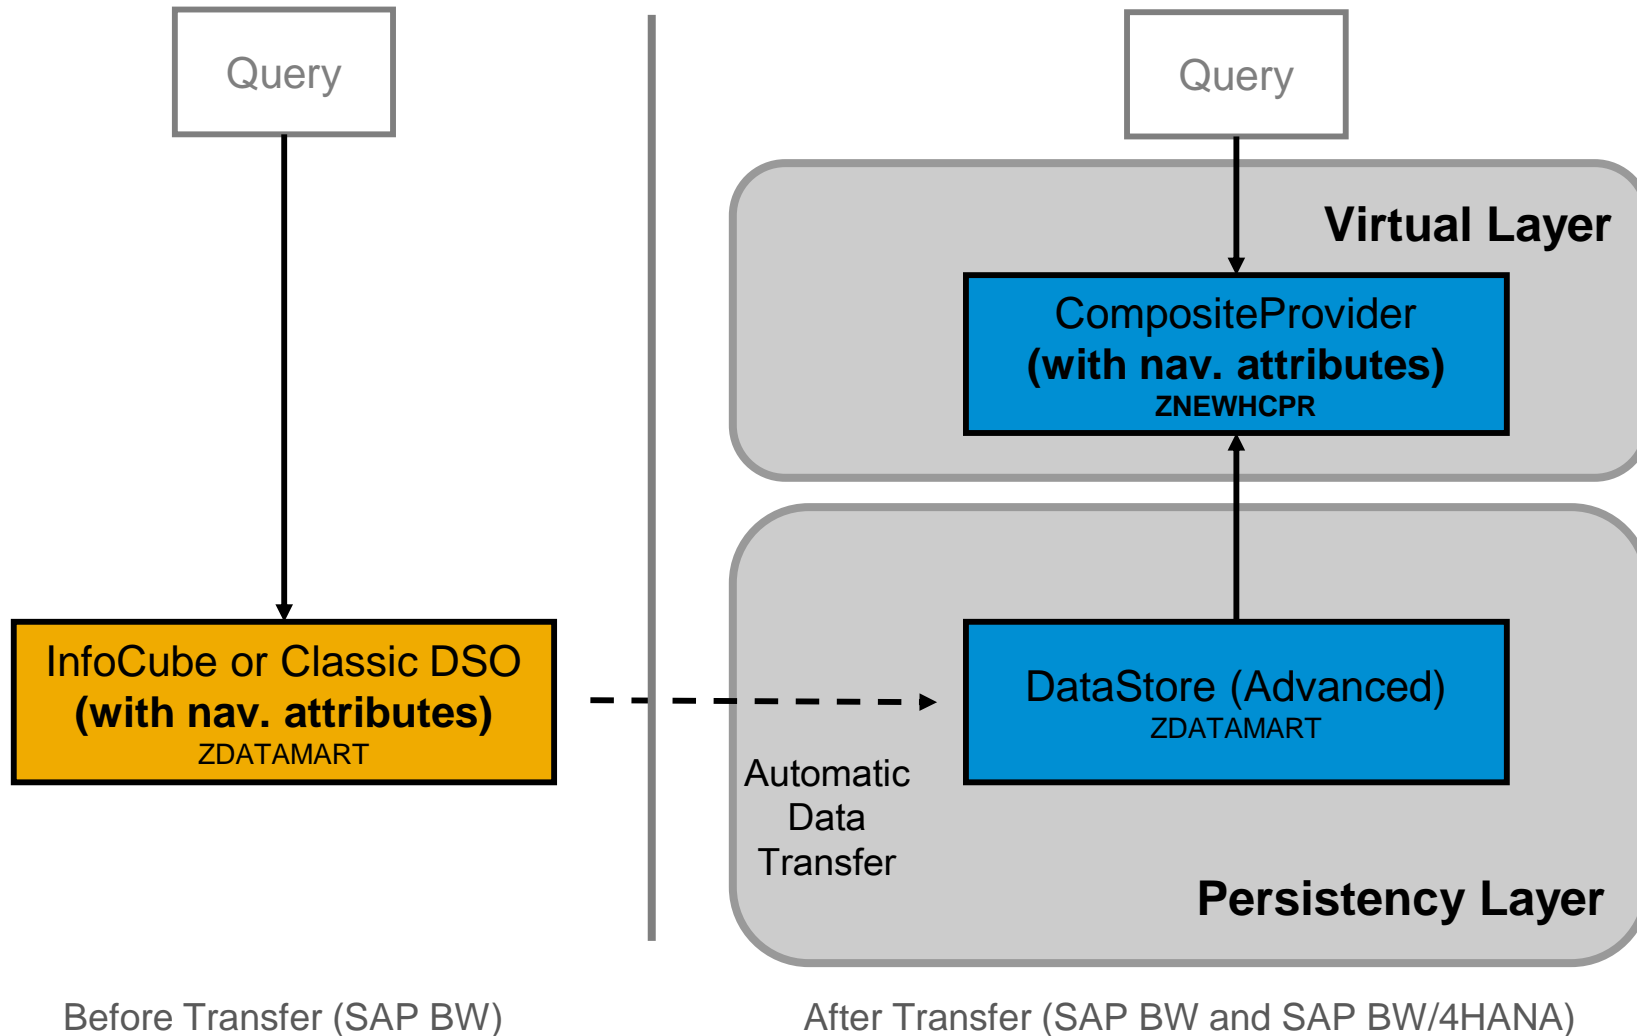

In-place Conversion

Meta Data and Data Conversion



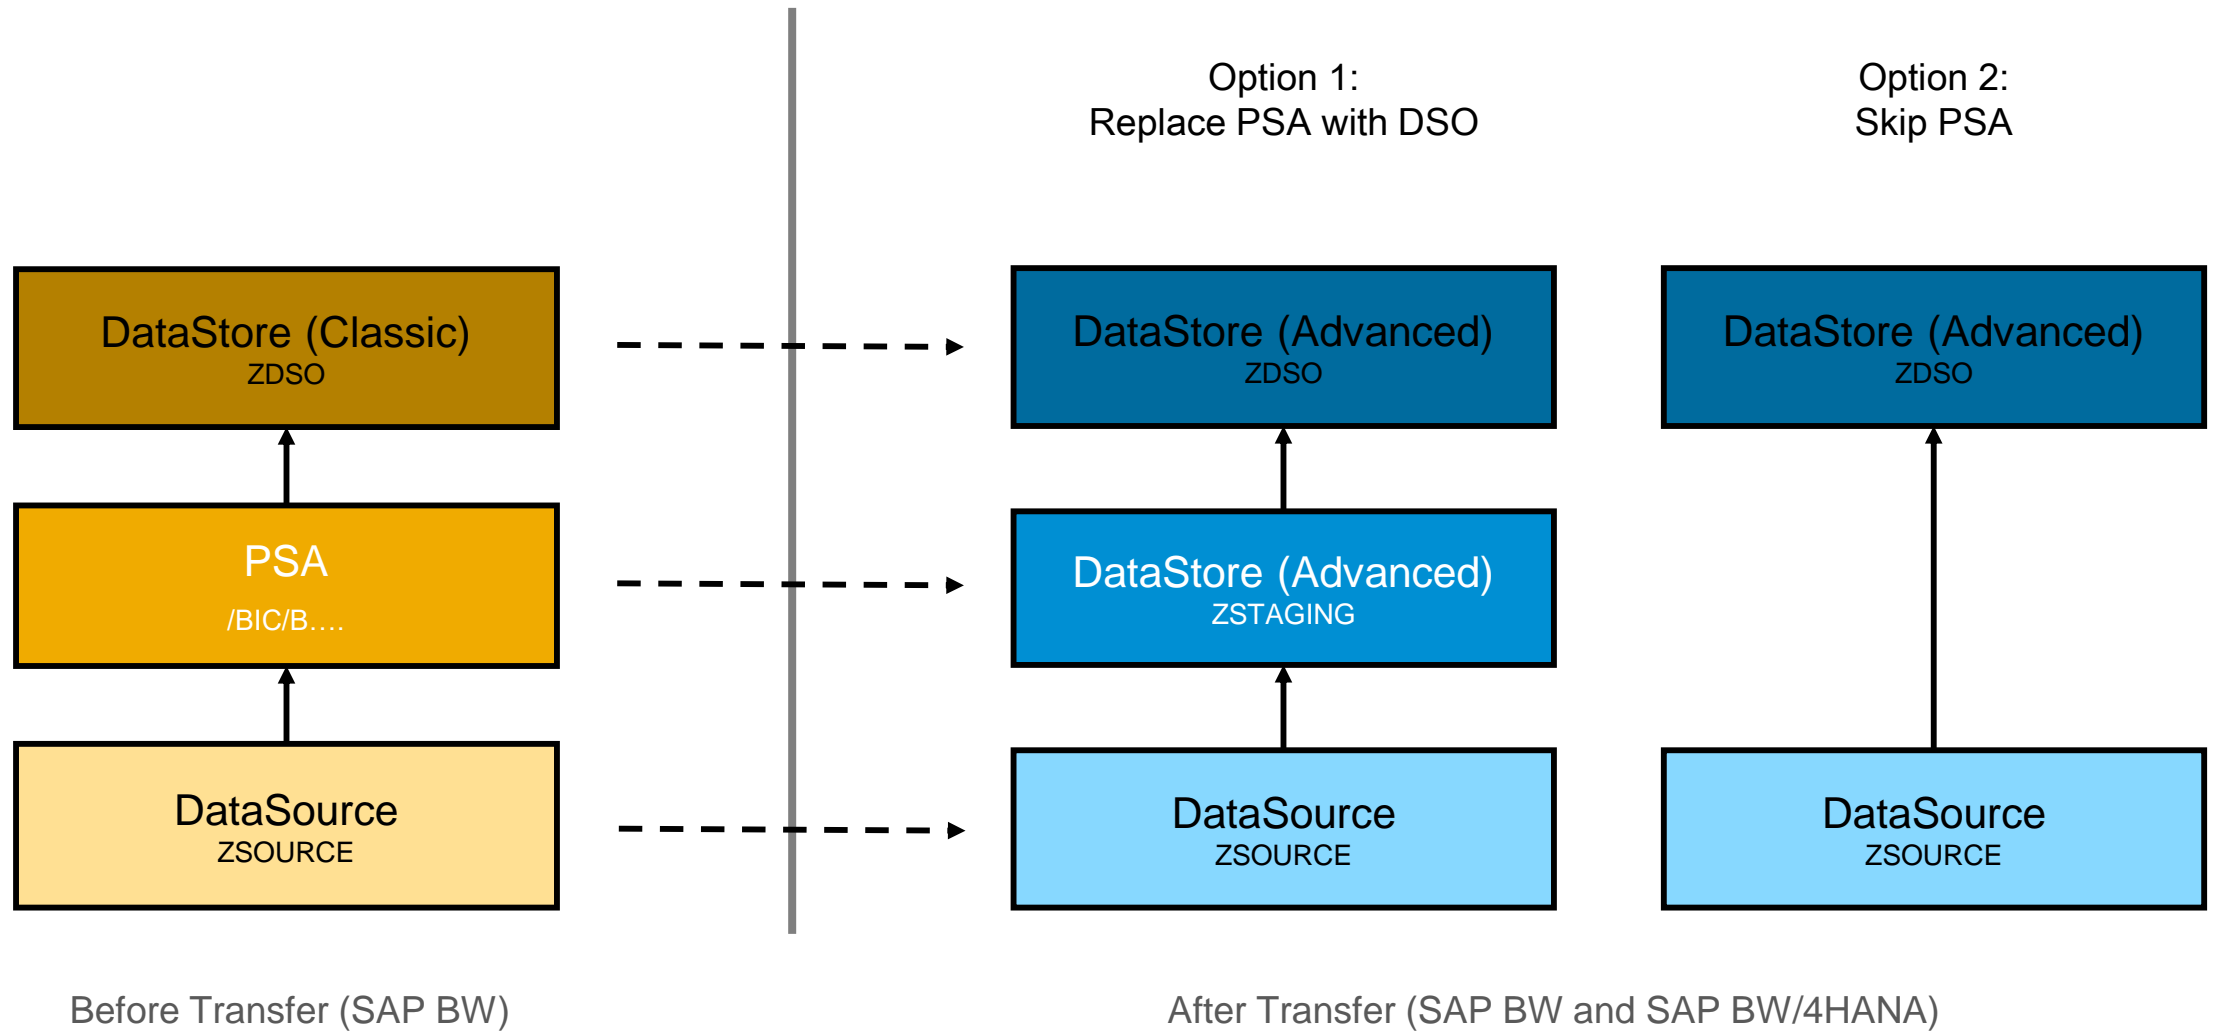

In-place Conversion

Meta Data and Data Conversion



- DataStore Objects (advanced) do not contain any settings regarding navigational attributes
- If navigational attributes are configured for an InfoCube or classic DataStore Object and used in at least one query, then the conversion process will introduce a new Composite-Provider which contains the corresponding attributes.
- Queries will then reference the CompositeProvider
- Reasons
 - Clear separation of virtual and persistency layers
 - Option to create different CompositeProviders with different nav. attributes

In-place Conversion PSA Conversion



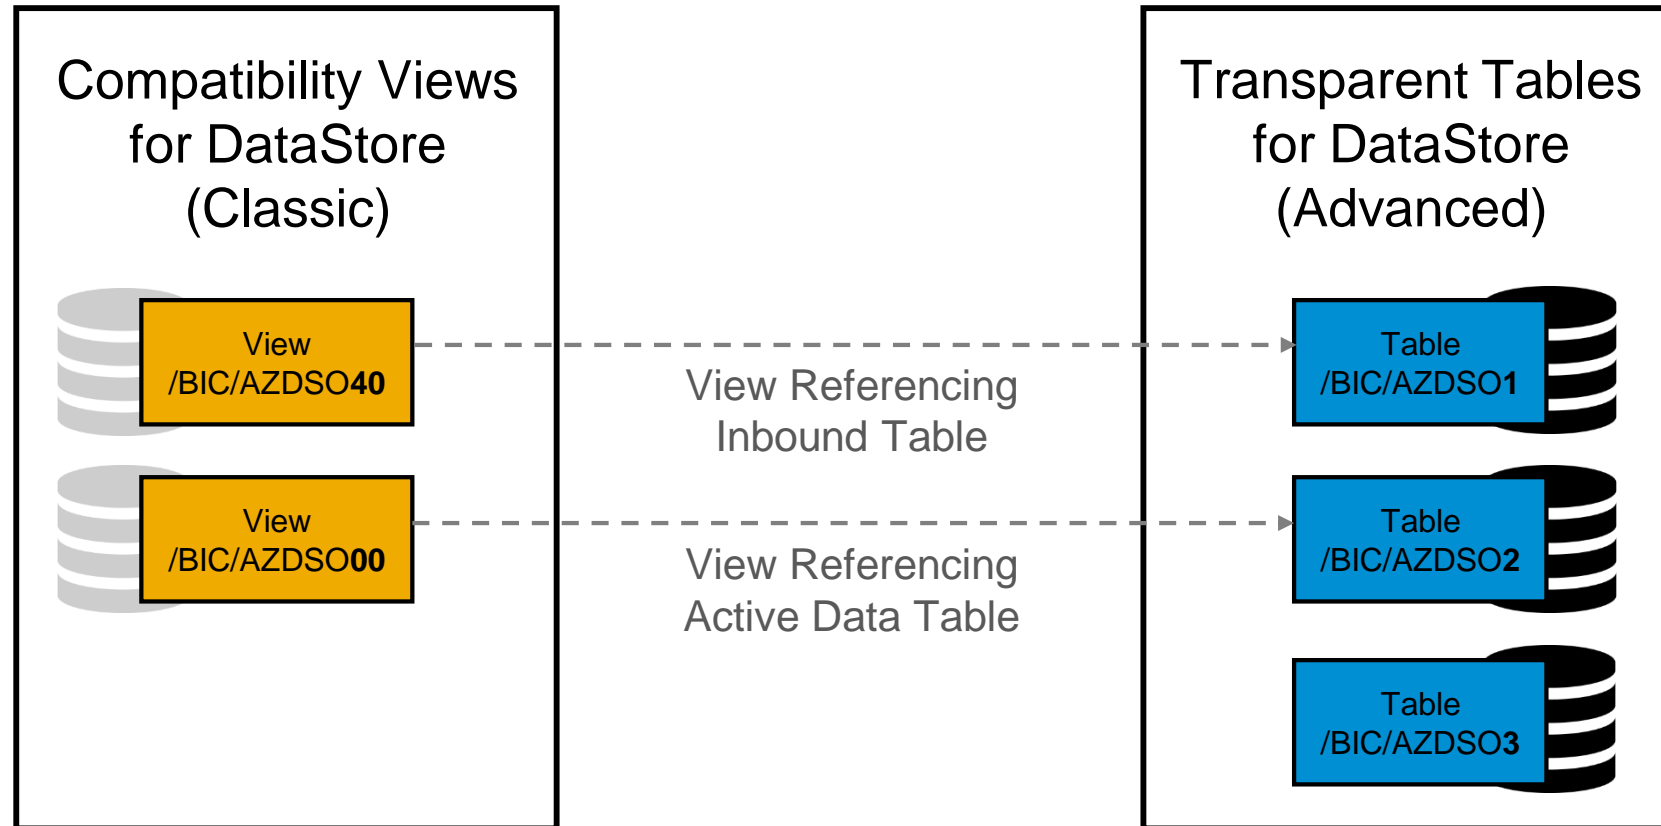
In-place Conversion

Example with Compatibility Views





Conversion of classic to advanced DataStore Objects generates Compatibility Views for the activation queue and active data table of the classic DSO (SAP Note [2539205](#) required).

→ Less effort to adjust custom coding



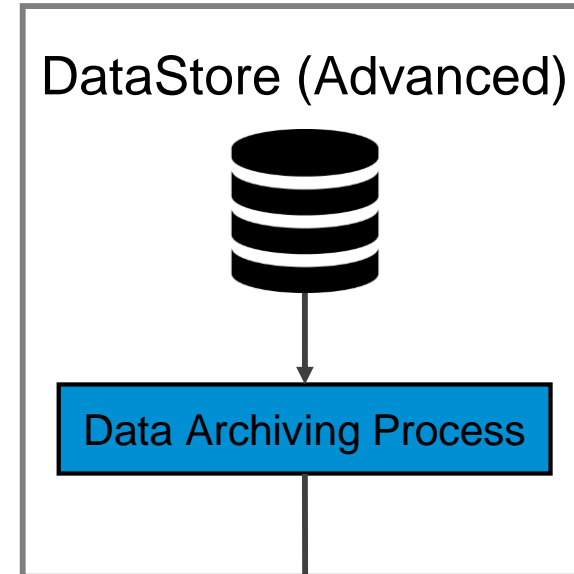
In-Place Conversion

Near-line Storage (SAP IQ)



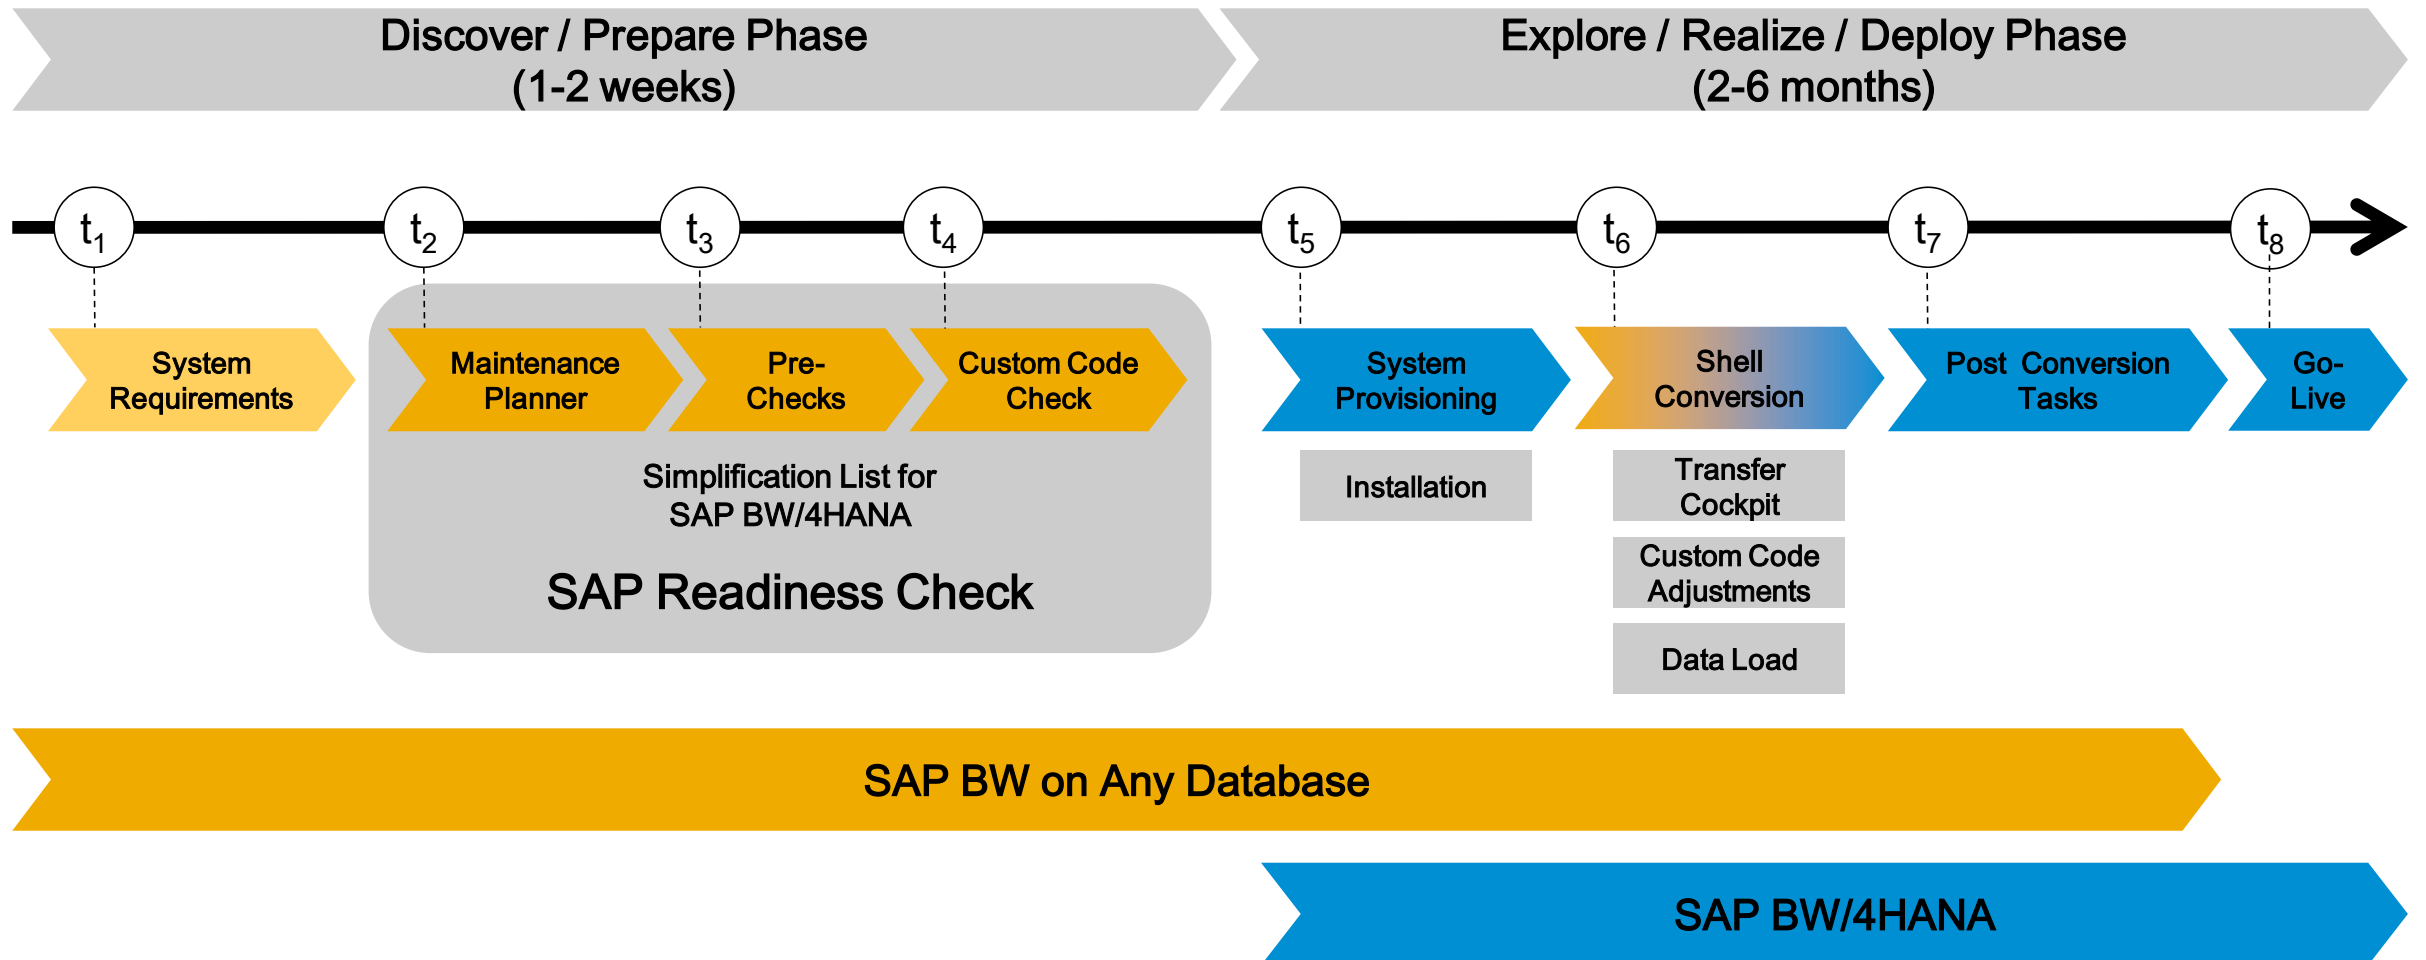

Near-line Storage
SAP IQ

- 1 Switch NLS to read-only
- 2 Convert NLS Request Management
- 3 Delete Data Archiving Process
- 4 Transfer Classic to Advanced DSO
- 5 Create new Data Archiving Process
- 6 Convert NLS Archive
(Request ID → Request TSN)



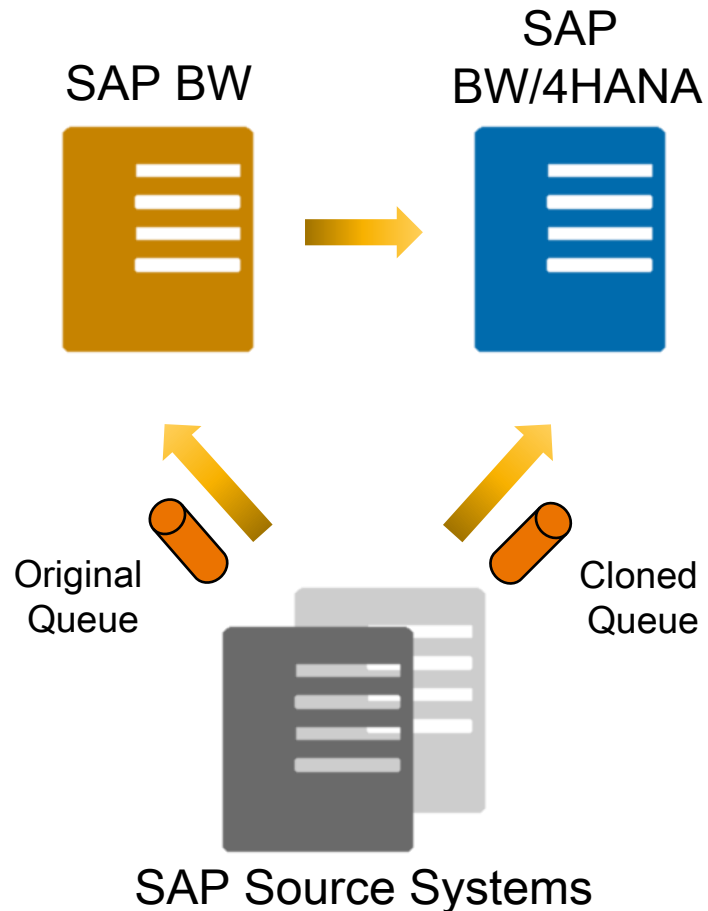

Near-line Storage
SAP IQ

Shell Conversion Basic Sequence



Remote/Shell Conversion

Automated Delta Queue Cloning and Synchronizing



Scope


- Enable SAP BW/4HANA system to use same source systems (remote connection)
- Prepare to easily bring new SAP BW/4HANA system in consistent state regarding data loading
- Minimal downtime of original SAP BW system
- Parallel operation (temporarily) of both systems and comparison of performance possible

Key Features of Remote Conversion

- Automatic delta queue cloning
- Checks for consistent delta states
- Automatic synchronization of both delta queues
- No downtime in productive source systems

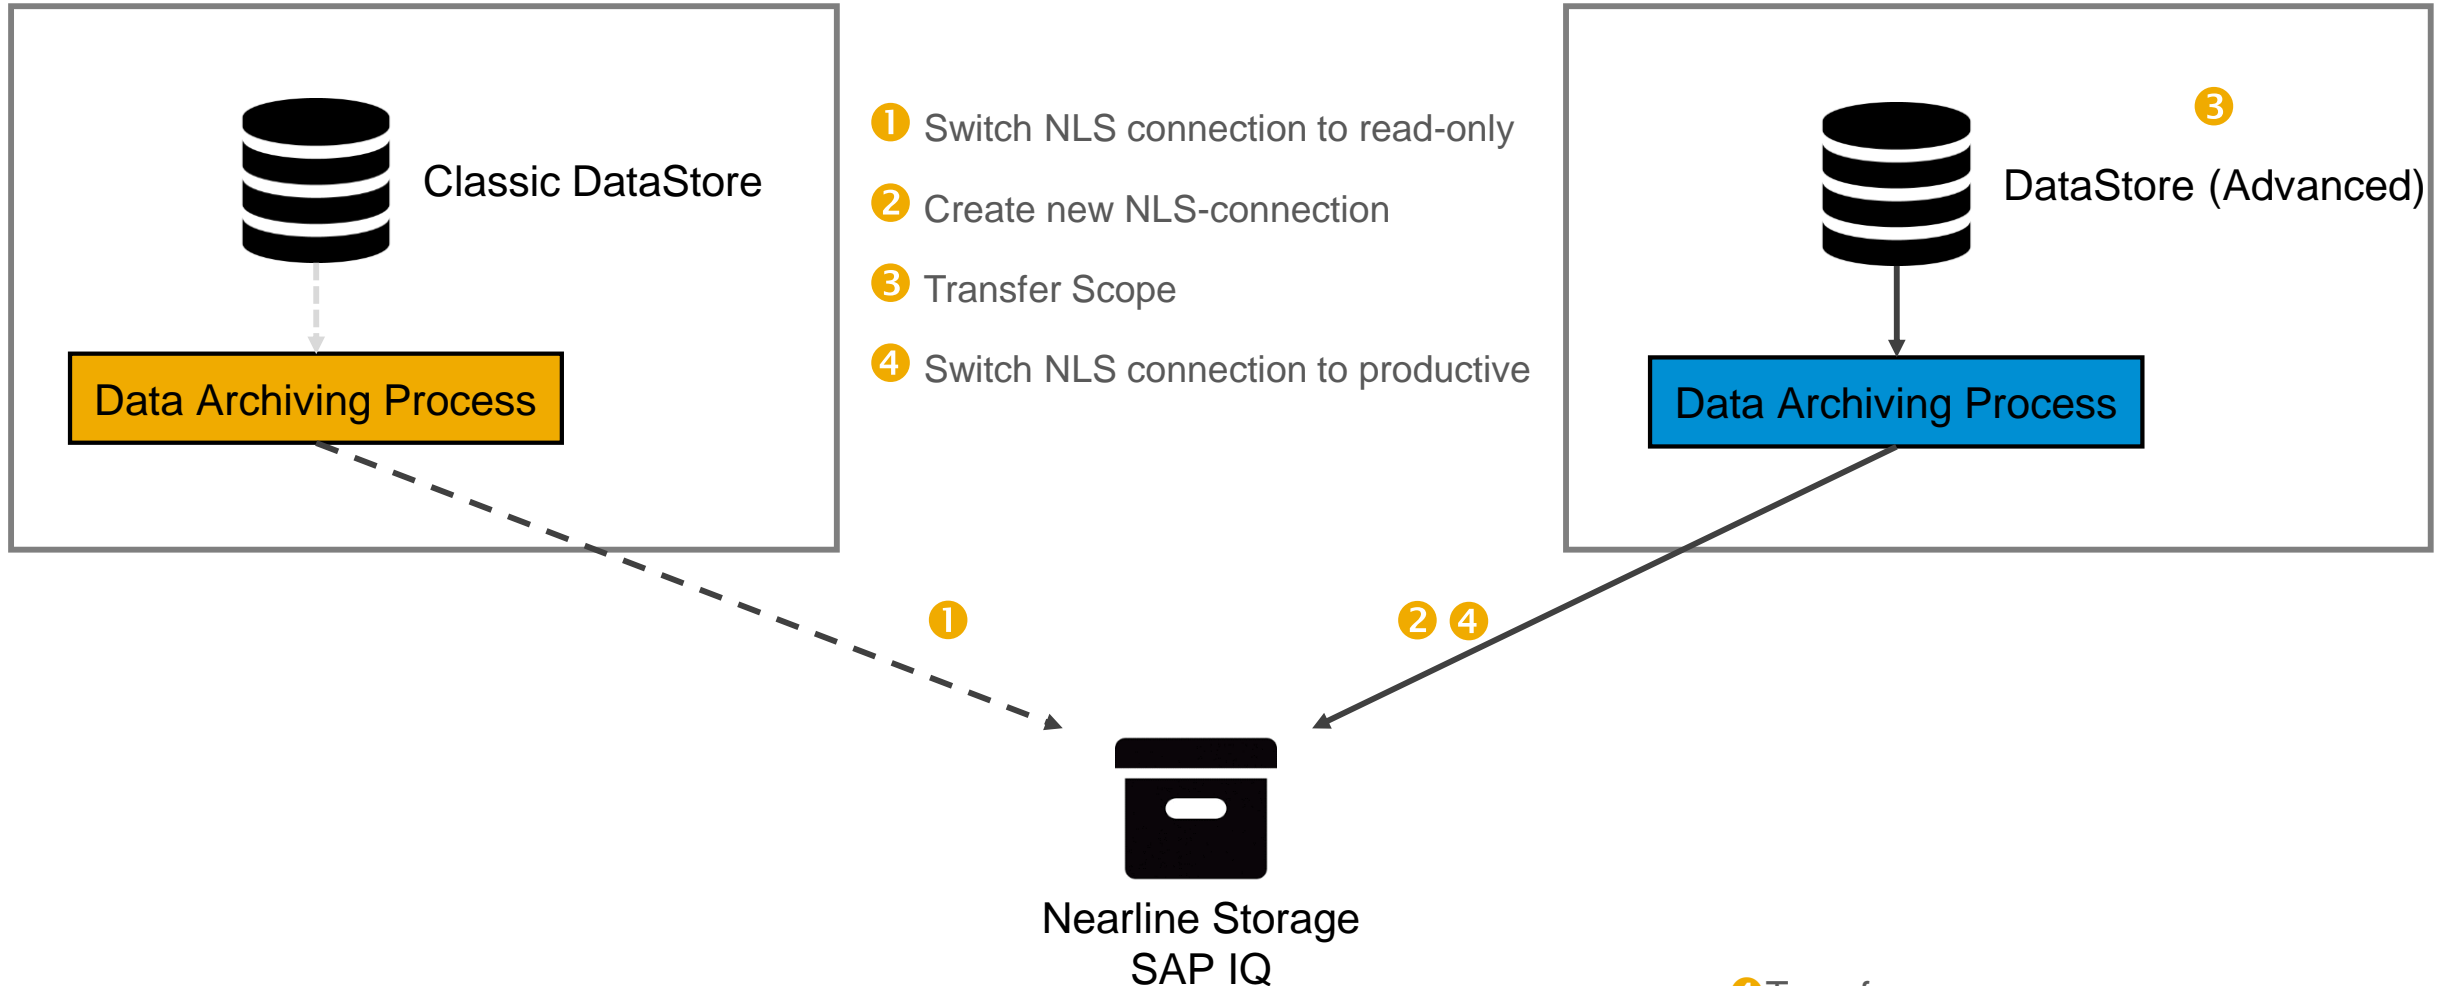
Remote/Shell Conversion

Near-line Storage (SAP IQ) – Reuse NLS Archive




Sending System (SAP BW)

Receiving System (SAP BW/4HANA)



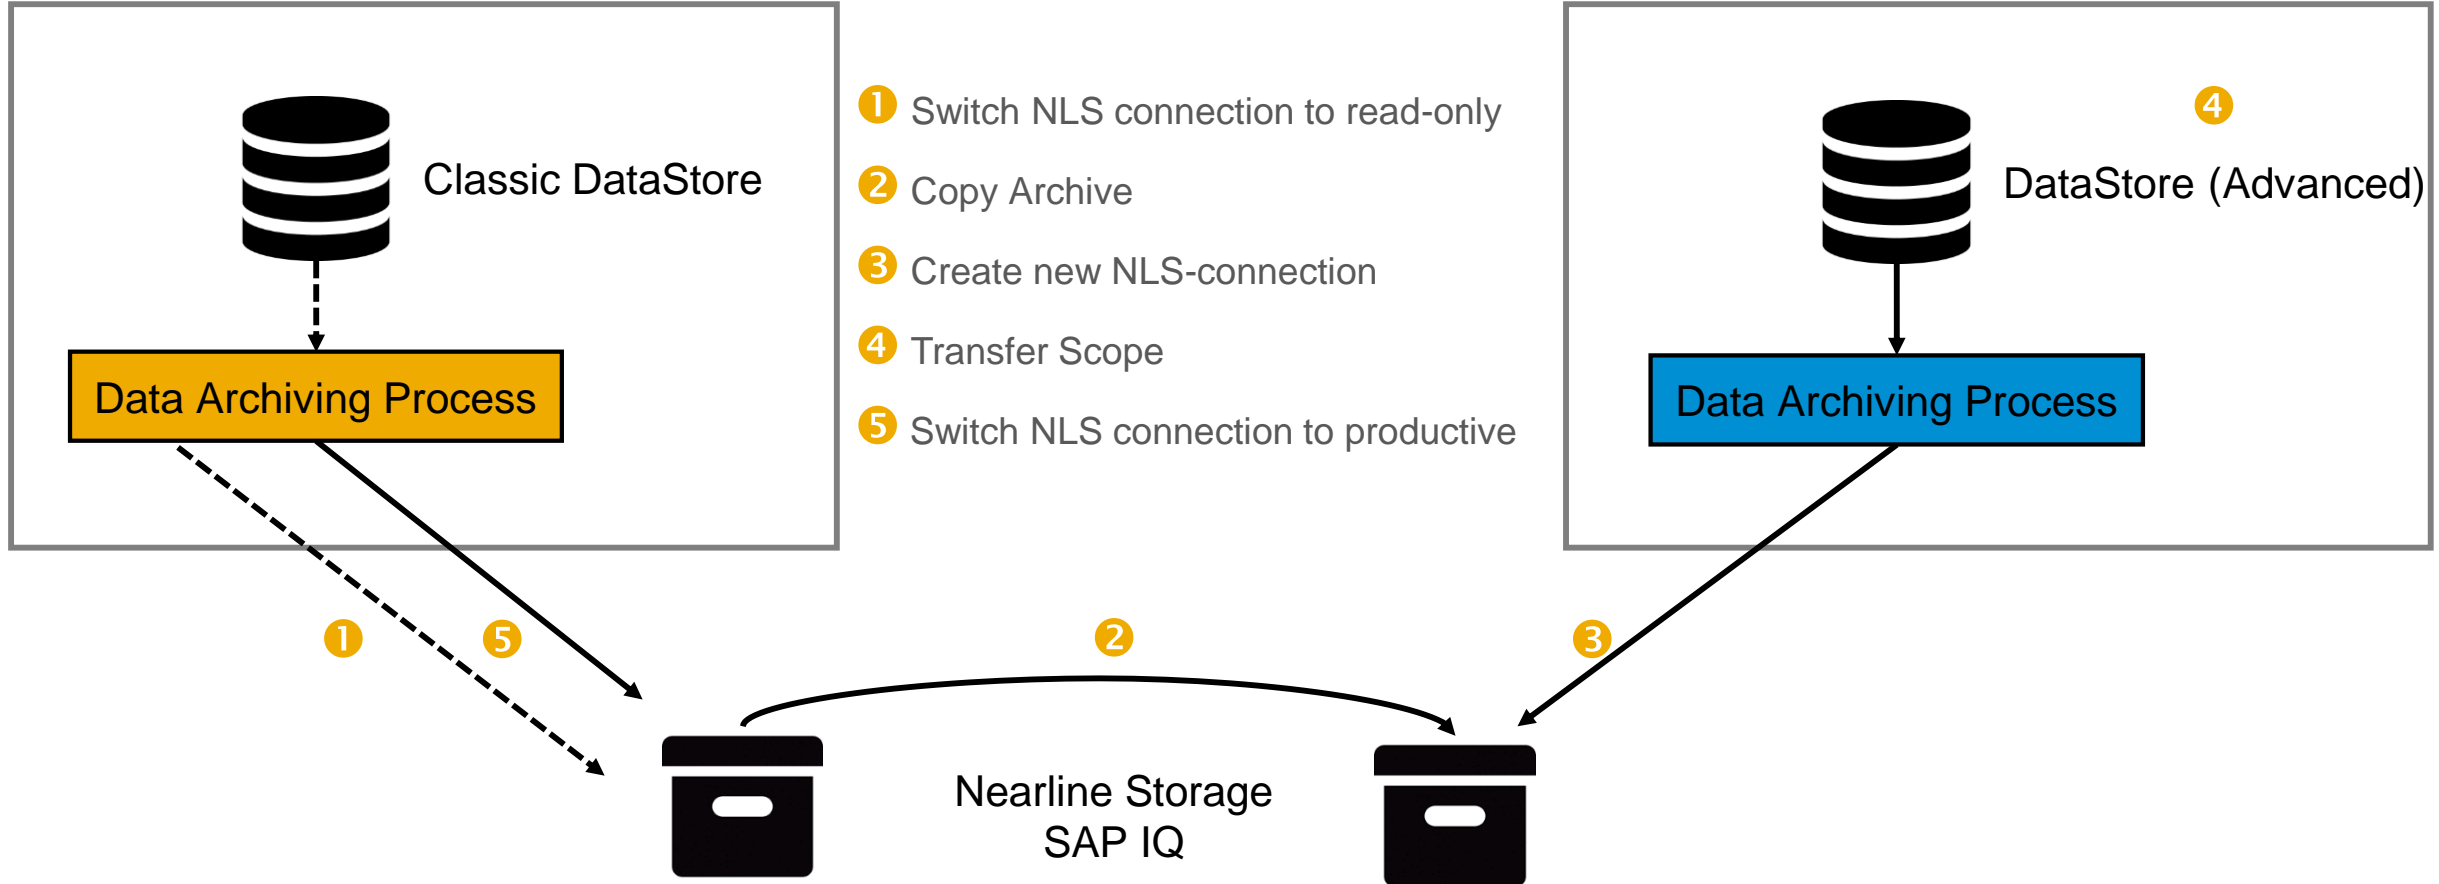
Remote/Shell Conversion

Near-line Storage (SAP IQ) – Copy NLS Archive (Alternative)



Sending System (SAP BW)

Receiving System (SAP BW/4HANA)



Agenda

Looking back

- Milestones

Data Tiering Optimization with BW/4HANA 1.0 Feature Pack 8

- The concept
- Supported Cold Stores
- What's the difference to classical Near-line Storage?

Looking today

- SAP NLS with IQ and the conversion to BW/4

Looking forward

- Roadmap

SAP BW/4HANA Data Tiering Optimization (DTO)

Roadmap

Today

SAP BW/4HANA 1.0 FP08

- Support for cold storage in the Hadoop File System (HDFS)
- Support for cold storage in SAP HANA Vora (disk engine)*
- Implementation, Administration and Performance Improvements, e.g.
 - Mass partition creation support
 - Integration with Process Chains
- API for Temperature Maintenance that enables Data Tiering Automation Implementation
- Support for multiple extension nodes

Conversion Support

- Transfer of Classic DataStore Objects / InfoCubes to Advanced DataStore Object incl. Nearline Store

Mid Term – Estimated for Q1 2019**

SAP BW/4HANA 2.0

- Enhanced Data Tiering Automation
 - Relative or rule based conditions
 - Data Access Statistics
- Data Tiering Optimization Cockpit as Web User Interface
- Support for (exceptional) update operations to data in the cold store
- Coverage of additional Data Store Objects types in cold store (without Activation, Inventory)
- Make Cold Store connection configurable on object level
- Simplification of temperature schema

Future Direction

- Enhanced Data Tiering Automation
 - Statistics based data movements
 - Dynamic Partitioning Schema
- Conversion of Data Archiving Processes (NLS) to SAP BW/4HANA Data Tiering Optimization
- Multi-dimensional partitioning scheme
- Deeper SAP Data Hub Integration

*Validation with upcoming Vora release still pending

** This is the current state of planning and may be changed by SAP at any time

Thank you.

Contact information:

Dr. Peter Zimmerer
Chief Development Architect
SAP SE

Roland Kramer
PM EDW (BW/HANA/IQ)
SAP SE
@RolandKramer

