



HPE STORAGE SOLUTIONS FOR SAP HANA NATIVE STORAGE EXTENSION

Using HPE Primera, HPE Nimble Storage, and HPE 3PAR for
SAP HANA Native Storage Extension

CONTENTS

Introduction.....	3
Data growth in SAP HANA.....	3
Data tiering decision parameters.....	4
Performance compared to cost.....	4
Ease of movement.....	5
NSE using HPE Storage.....	5
NSE compared with dynamic tiering and extension nodes.....	6
Configuring NSE.....	6
Sizing NSE.....	7
Example.....	7
NSE Advisor.....	8
Using NSE.....	8
Dataset.....	8
Defining warm data.....	9
Querying hot and warm data.....	9
Example.....	9
Using HPE Primera, HPE Nimble Storage, and HPE 3PAR for NSE.....	11
Summary.....	12



INTRODUCTION

Most enterprises today are data-driven. To continue to stay relevant in ever-evolving markets, enterprises are adopting new technologies including artificial intelligence, machine learning, internet of things, big data, quantum computing, and so on. Even though these are vastly different technologies, there is one commonality among them—they are data-hungry. This has led to a massive increase in the rate of data generation. According to [World Economic Forum](#), by 2025, 463 exabytes of data will be created each day, globally.

Businesses are considered successful not only when they collect and store data, but also when they possess the right infrastructure and tools to efficiently and quickly analyze this data and extract relevant information from it. As a result, data has become the new currency worldwide.

However, unlike monetary currencies used every day, data does not always hold the same value. As data grows, its usage and relevance changes. In general, it becomes less valuable as it ages. It has an expiration date, meaning specific data can lose its value at a certain point. In addition, the related technology, deployed hardware, and processes might be able to handle today's data, but they are not designed to meet the demands of tomorrow. For example, fuel efficiency has been a top priority for large-scale automotive manufacturers. They have numerous data collection mechanisms in place and have invested countless resources in improving their engine performance as much as possible. However, manufacturers have started investing resources in alternate options, such as electric cars. The focus then must shift to how electric cars can become as efficient as their gasoline-based counterparts, but data collected for gasoline-fueled cars would not be relevant in improving the efficiency of electric cars. This data has lost its value. New data collection mechanisms need to be employed. As business problems evolve, their solutions also change and so does the data that is required to arrive at positive outcomes.

SAP HANA® is a leading in-memory database system incorporating real-time analytics. Memory is expensive compared to disk storage, and therefore the value of data residing in memory should remain high. If this value decreases over time, there should be alternative storage options for this older or less-valued data, thereby making way for new and more important data to reside in memory, with the less-valuable data remaining accessible and usable.

This technical white paper describes one such offering from SAP® named *Native Storage Extension (NSE)* in which the storage used to store SAP HANA data is extended to make space for the less-valued data. This technology is native to SAP HANA and does not require a separate compute and storage tier to transfer data across.

This white paper talks in detail about SAP HANA NSE as a method to tier data. It assumes a fair knowledge of SAP applications, SAP HANA, and databases in general. It distinguishes between the three data tiering technologies that SAP offers. It also describes how HPE storage systems such as HPE Primera, HPE Nimble Storage, and HPE 3PAR can be used to leverage NSE technology.

DATA GROWTH IN SAP HANA

SAP classifies data in the SAP HANA database into three categories:

- **Hot:** Data that is of the highest importance and value and is critical to the business. The hardware that holds and processes this data should provide the highest performance possible.
- **Warm:** Data that is less frequently accessed or modified but still needs to be part of the database. The hardware that holds and serves this data can offer a lower level of performance than the hot store.
- **Cold:** Data that is infrequently accessed but still needs to be stored for archival or compliance purposes. The hardware needed to hold this data can be low-cost commodity storage hardware with low performance considerations.



DATA TIERING DECISION PARAMETERS

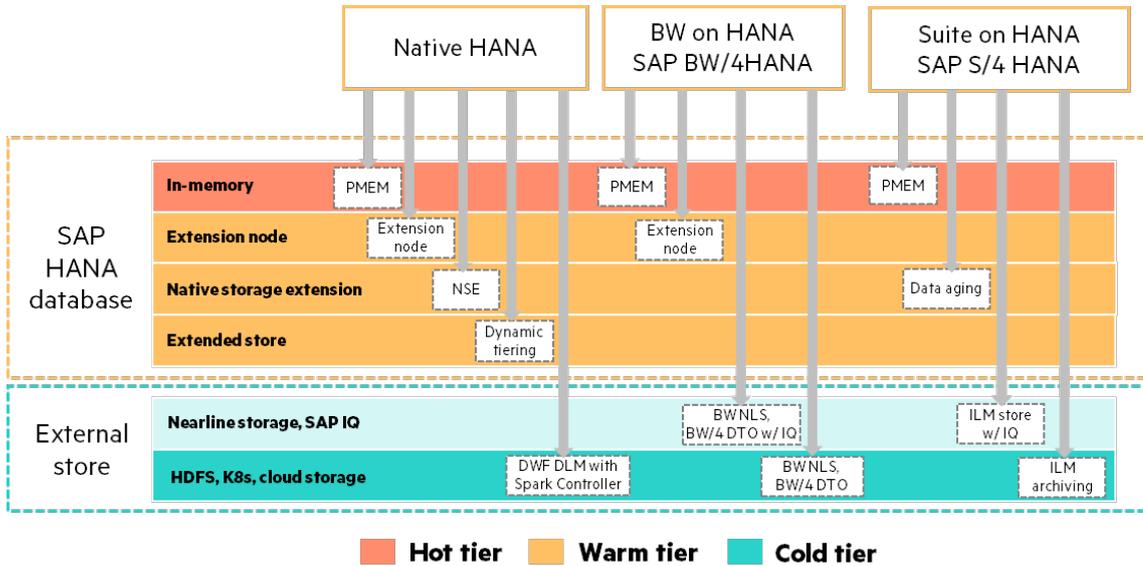


FIGURE 1. SAP HANA data tiering options for various workloads

After you classify data into tiers according to its value and importance, take the following parameters into consideration to decide which storage option to use to deploy the data:

- **Workload:** An SAP HANA database supports various SAP applications such as SAP Business Warehouse (BW), SAP Suite on HANA, and other native SAP HANA applications. The data contained in each of these workloads has varied characteristics, behavior, and usage.

As shown in Figure 1, SAP has broadly classified its data tiering approaches based on the workload being used.

- **Native HANA:** Hot data should reside in memory for native SAP HANA applications. One variation is to use persistent memory (PMEM), which has a lower total cost of investment and a larger footprint than traditional DRAM. This approach expands overall memory capacity and can maintain the state of data in PMEM after shutdown. The result can be improved database startup time because data does not need to be loaded in memory from the underlying storage.

Warm data management for native SAP HANA depends on the size of the warm data, the performance expectations of the warm tier, the cost investment planned for the warm tier, and the available hardware. Three methods exist for native SAP HANA:

- NSE
- Dynamic tiering
- Extension nodes

Cold data tiering for native SAP HANA workloads is possible using an external data source such as an Apache® Hadoop® cluster. The SAP technologies that help integrate SAP HANA with a Hadoop cluster are SAP Data Hub or SAP HANA Spark Controller.

- **SAP BW on HANA or BW/4HANA:** Warm data management for SAP BW on HANA workloads is typically handled using extension nodes. BW is usually a scale-out deployment and requires consistent performance even from the warm tier.
- **SAP Business Suite on HANA or S/4HANA:** For SAP Suite on HANA or S/4HANA workloads, warm data is managed using an SAP NetWeaver® technology known as *data aging*. It allows data movement within a database in order to use the available memory more efficiently.

Performance compared to cost

With any data tiering technology, the intent is to let the data reside at the most appropriate tier where it is least expensive and offers the best possible performance. As shown in the white paper [HPE Storage solutions for SAP HANA warm data tiering](#), SAP HANA dynamic tiering and extension node technologies physically move data from the hot tier to the warm tier. Although this movement is justified within these technologies, it is still an overhead. With NSE, SAP intends to avoid this movement by providing a technology that is native to HANA hardware and that extends the native storage to accommodate varying data temperatures. With NSE, you can increase the total amount of SAP HANA data that can be stored in the same hardware being used for SAP HANA.



Ease of movement

One of the characteristics of warm data is that it is part of the same SAP HANA system as the hot data, which makes it possible to move data from the hot tier to the warm tier and back again. Therefore, it is more efficient to use storage that provides faster reads to move data back to the hot tier in case a query or use case demands it. Ease of movement is a crucial factor in deciding which data tiering mechanism to use. With NSE, there is no movement of data between storage tiers. Instead, with the help of a user-defined criteria, SAP HANA determines which data should reside permanently in memory and which should be off-loaded to the disk. This operation is easier, faster, and more efficient than dynamic tiering and extension nodes.

NSE USING HPE STORAGE

NSE works on the same storage as the persistence layer without the need for a separate physical or virtual compute or storage tier. The data volume used for persistence hosts the hot as well as the warm data. The only difference is that hot data resides in memory and is periodically persisted on the disk, whereas warm data resides on the disk itself.

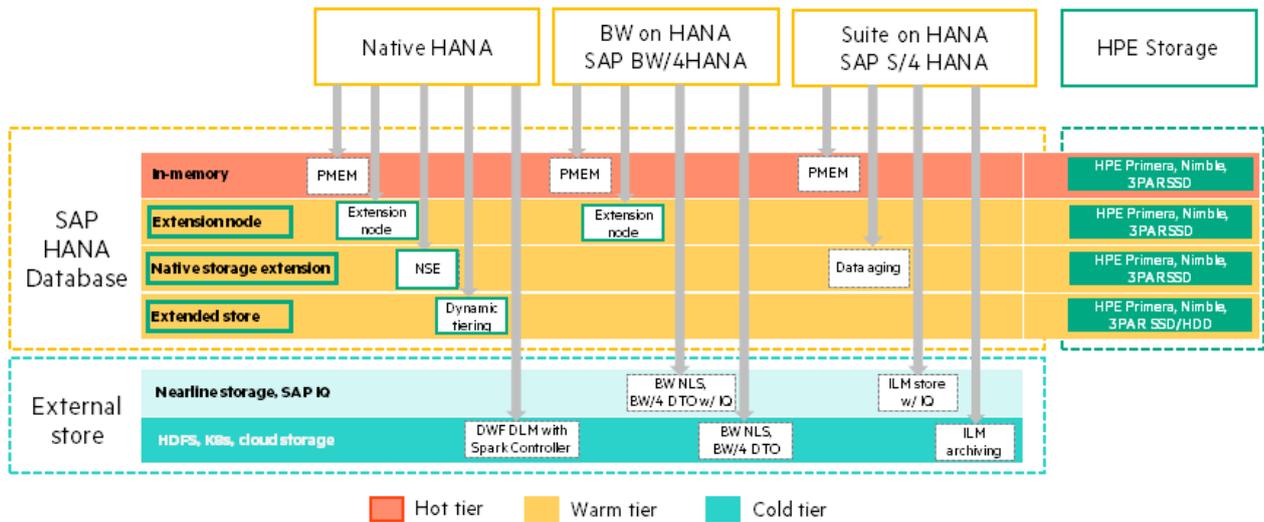


FIGURE 2. HPE Storage solutions for SAP HANA data tiering options

As shown in Figure 2, SSDs on HPE Primera, HPE Nimble Storage, and HPE 3PAR host the persistence layer as well as warm data using NSE. As a result, the storage layer offers the same read and write performance for hot and warm data. Whenever a query needs to access warm data, SAP HANA loads it in the form of pages from the disk to a special area in the memory called the *buffer cache*. The buffer cache is a place in the memory where warm data can temporarily reside, depending on how often it is accessed.



NSE COMPARED WITH DYNAMIC TIERING AND EXTENSION NODES

Table 1 provides a quick look at all the warm data tiering options that SAP HANA provides.

TABLE 1. SAP HANA warm data tiering options at a glance

Features	SAP HANA NSE	SAP HANA dynamic tiering	SAP HANA extension nodes
Detail	Built-in warm data store using extended storage	SAP IQ-based, disk-based extended store	An in-memory extension with a new worker node
Where is warm data stored?	On disk and in buffer cache	On disk	In memory and on storage
Hardware required	Existing SAP HANA hardware with additional storage	Commodity hardware	TDI certified hardware
Data movement between tiers	No physical movement of data between tiers; hot data resides in main memory and warm data is pushed to the disk	Using SQL or SAP Data Lifecycle Manager	Using SQL or SAP Data Lifecycle Manager and Data Distribution Optimizer
Licensing	No separate license needed	Separate license (usually part of SAP HANA Enterprise and Standard edition licenses)	No separate license needed
Technical difference from hot tier	Warm data is loaded in the buffer cache area of main memory on request	Warm data on esserver service	worker_dt worker group
Primary goal	To improve price: performance ratio	Cost savings from warm data	High performance from warm data
HPE 3PAR recommendation	SSD with RAID 6	HDD or SSD with RAID 6	SSD with RAID 6
Data protection	Planned	HPE StoreOnce Catalyst Plug-in for SAP HANA	HPE StoreOnce Catalyst Plug-in for SAP HANA and HPE Recovery Manager Central for data snapshots

CONFIGURING NSE

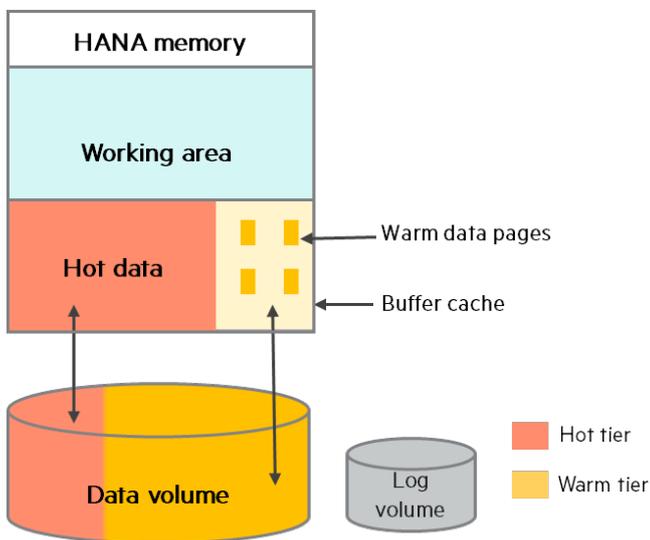


FIGURE 3. NSE configuration

SAP recommends that of the total amount of memory available in a system, 50% should be dedicated as working area in SAP HANA. This is used for data processing, storing temporary results, and other operating system tasks. The remaining 50% of main memory is used by SAP HANA to store data. This is the maximum size of the database that SAP HANA can hold.



The storage volumes (data, log, shared) are sized based on the total memory in the system as per the [SAP HANA Storage Requirements white paper](#). Without NSE, the entire 50% of the main memory is allocated to store hot data. However, with NSE, a special area called *buffer cache* is carved out from this main memory. This is the area where pages from warm data are loaded for query processing and are written back to the disk later.

Terminology used in NSE is defined as follows:

- **Buffer cache:** The buffer cache is an area of the main memory that is used to load warm data pages from the data volume for query processing. The buffer cache is automatically enabled in an SAP HANA system and by default, 10% of the total memory of the server is reserved for it. Note that this memory is reserved but not allocated unless some PAGE_LOADABLE data is accessed. Otherwise, 100% of the memory is available for SAP HANA for hot data.
- **Load behavior:** A table, partition, column, or an index within SAP HANA can be selected and enabled for paging—that is, designated as warm. These components must be labeled as PAGE_LOADABLE to be enabled for paging.
- **Page-loadable data:** A table, partition, column, or an index that is loaded in the memory page by page rather than in its entirety, thereby reducing its memory footprint. This data is loaded in the buffer cache part of the main memory of the SAP HANA system.
- **Column-loadable data:** The data that is frequently accessed or designated as hot is loaded entirely in memory. It can be a table, partition, column, or the part of a table that is not page-loadable. The query performance for a column-loadable component is usually better than page-loadable data because column-loadable data always resides in the main memory.
- **Tiering criteria:** As with other SAP HANA data tiering technologies, a distinct and definite criterion to identify and distinguish hot data from warm data—that is, column-loadable from page-loadable data—is required. The only difference between NSE and other SAP HANA tiering technologies is that no separate hardware tier must be installed and configured to hold warm data.

SIZING NSE

When using NSE, follow these guidelines to size the storage based on the amount of hot data compared with warm data in the system.

- NSE does not require any additional compute hardware.
- NSE does require additional storage. The current guidelines from SAP suggest that warm data size should be $\leq 4 \times$ hot data size.
- The amount of buffer cache is dependent on the size of the warm data according to this rule:
Buffer cache = warm data size / 8
- SAP HANA SPS 04 supports NSE only on a scale-up configuration.

Example

This example shows how sizing differs with and without NSE.

- Without NSE:
 - Total RAM = 2 TB
 - Work area = 1 TB (50% RAM)
 - HANA hot data = 1 TB
 - Data volume = $1.2 \times$ RAM = 2.4 TB
 - Total DB size = **1 TB**
- With NSE:
 - Total RAM = 2 TB
 - Work area = 1 TB (50% RAM)
 - Buffer cache = 200 GB
 - HANA hot data = 800 GB
 - HANA warm data = $200 \times 8 = 1.6$ TB
 - Data volume size = 2.4 TB + 1.6 T = 4 TB
 - Total DB size = 800 GB + 1.6 TB = **2.4 TB**



The total data stored within an SAP HANA system increases if it is divided into hot and warm categories using NSE. The data volume size increases accordingly.

NSE Advisor

SAP HANA NSE Advisor is a tool available with the SAP HANA Cockpit that helps you determine how to divide data into hot and warm categories. It provides suggestions about load units for tables, partitions, and columns based on their access frequency.

The goal of data tiering is to achieve the best possible cost-to-performance ratio for a given SAP HANA database. NSE Advisor creates recommendations in two steps:

1. Observes the existing queries from the application to the database and creates a data access pattern.
2. Applies heuristics on these data access statistics to generate recommendations.

For detailed information on NSE Advisor, refer to the [SAP HANA Administration Guide](#).

USING NSE

Use HDBSQL statements, SAP HANA Cockpit, or SAP HANA Studio to configure NSE.

Dataset

The industry-standard TPC-H database available on www.tpc.org was used to illustrate how load units are configured and NSE is implemented. TPC-H is a decision support benchmark. It has eight individual tables—ORDERS, LINEITEM, PARTSUPP, PART, SUPPLIER, CUSTOMER, NATION, and REGION. Figure 4 shows the database schema of the TPC-H dataset with relationships among separate tables. For more details, refer to the TPC Benchmark H Standard Specification Revision 2.18.0 available on www.tpc.org.

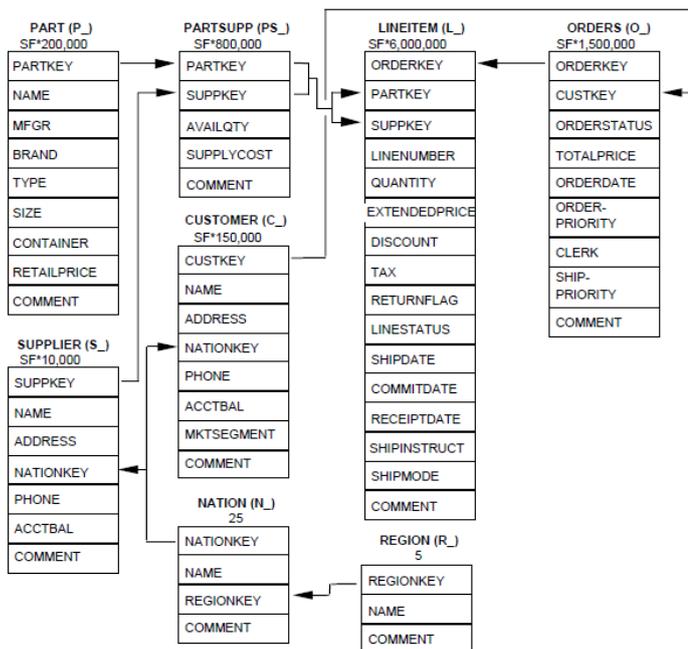


FIGURE 4. TPC-H v2.18.0 schema

For simplicity, the NATION, REGION, SUPPLIER, PART, PARTSUPP, and CUSTOMER tables were kept as column-loadable and the ORDERS and LINEITEM tables were made partially page-loadable.



Defining warm data

The TPC-H generated dataset contains data from Jan 01, 1992 to Dec 31, 1998. As it is with most enterprises, newer data is accessed more frequently than the relatively older data. In this example, all records from 1997 and 1998 were designated as hot and remaining data from 1992 to 1996 was designated as warm. All records in the ORDERS and LINEITEM tables can be classified with this criteria.

```
ALTER TABLE "TPCH"."ORDERS_CS"
PARTITION BY
RANGE ("TPCH"."ORDERS_CS"."O_ORDERDATE")
((PARTITION '1997-01-01' <= VALUES < '1999-01-01', PARTITION OTHERS));
```

```
ALTER TABLE "TPCH"."ORDERS_CS" ALTER PARTITION 2 PAGE LOADABLE;
```

The first ALTER TABLE SQL statement partitions the ORDERS table based on these criteria. SAP HANA automatically numbers the partitions as '1' for the partition with data from 1997 to 1998 and '2' for the partition containing the remaining data.

The next ALTER TABLE SQL statement then designates partition 2 as PAGE_LOADABLE. The first partition remains COLUMN_LOADABLE by default. This method designates data as hot and warm.

Similar SQL statements can be applied to the LINEITEM table:

```
ALTER TABLE "TPCH"."LINEITEM1_CS"
PARTITION BY
RANGE ("TPCH"."LINEITEM1_CS"."L_RECEIPTDATE")
((PARTITION '1997-01-01' <= VALUES < '1999-01-01', PARTITION OTHERS));
```

```
ALTER TABLE "TPCH"."LINEITEM1_CS" ALTER PARTITION 2 PAGE LOADABLE;
```

You can also perform this partitioning when creating the table with the CREATE TABLE SQL statement. You can view the table partitions and their respective load units from the M_TABLE_PARTITIONS system view in SAP HANA. Figure 5 shows how the tables are partitioned and how partitions have either COLUMN or PAGE as load unit values.

HOST	PORT	SCHEMA_NAME	TABLE_NAME	NODE_ID	PARENT_NODE_ID	PART_ID	LOAD_UNIT	STORAGE_TYPE	MEMORY_SIZE_IN_TOTAL	MEMORY_SIZE_IN_MAIN	MEMORY_SIZE_IN_PAGE_LOADABLE	MEMORY_SIZE_IN_DELTA	DISK_SIZE	
1	10.20.23.29	30,003	TPCH	LINEITEM1_CS	1	0	1	COLUMN	DEFAULT	102,391,608	102,295,856	0	81,208	27,264,225,280
2	10.20.23.29	30,003	TPCH	LINEITEM1_CS	2	0	2	PAGE	DEFAULT	73,972,879,692	73,972,783,764	67,619,016,704	81,208	72,862,535,680
3	10.20.23.29	30,003	TPCH	ORDERS_CS	1	0	1	COLUMN	DEFAULT	1,104	0	0	0	6,430,994,432
4	10.20.23.29	30,003	TPCH	ORDERS_CS	2	0	2	PAGE	DEFAULT	1,104	0	0	0	18,159,681,536

FIGURE 5. Table partitions and load units

In addition to a partition, you can set the load unit for a column or for an entire table. NSE supports only range and range-range heterogeneous partitioning schemes for setting the load unit for a partition. For more details, refer the [SAP HANA Administration Guide](#).

Querying hot and warm data

If you run a sample query on page-loadable and column-loadable data after partitioning and designating load units, the query performance is different because data is being read from disk in one case and memory in the other.

Example

Fetch all the records from the LINEITEM table for a particular year, between March 15 and April 1, where the item shipment mode was RAIL.

This query will be run for warm data (for example, year 1993) and then for hot data (for example, year 1998). First, use COUNT (*) to query the number of records for both these cases to make sure to extract relatively the same amount of data in both queries.



Query for warm data:

```
SELECT COUNT (*) FROM "TPCH"."LINEITEM1_CS" WHERE ("TPCH"."LINEITEM1_CS"."L_RECEIPTDATE"
>= '1993-03-15' AND "TPCH"."LINEITEM1_CS"."L_RECEIPTDATE" < '1993-04-01') AND
"TPCH"."LINEITEM1_CS"."L_SHIPMODE" = 'RAIL';
```

Result: 2167327

Query for hot data:

```
SELECT COUNT (*) FROM "TPCH"."LINEITEM1_CS" WHERE ("TPCH"."LINEITEM1_CS"."L_RECEIPTDATE"
>= '1998-03-15' AND "TPCH"."LINEITEM1_CS"."L_RECEIPTDATE" < '1998-04-01') AND
"TPCH"."LINEITEM1_CS"."L_SHIPMODE" = 'RAIL';
```

Result: 2165057

This example shows that the number of records are almost similar with a variation of less than 0.1 %. Now look at how much time SAP HANA takes to actually retrieve data in both these cases:

Query for page-loadable (warm) data:

```
SELECT * FROM "TPCH"."LINEITEM1_CS" WHERE ("TPCH"."LINEITEM1_CS"."L_RECEIPTDATE" >=
'1993-03-15' AND "TPCH"."LINEITEM1_CS"."L_RECEIPTDATE" < '1993-04-01') AND
"TPCH"."LINEITEM1_CS"."L_SHIPMODE" = 'RAIL';
```

Result: 2167327 rows selected (overall time 6303.031524 sec; server time 19.047972 sec)

Query for column-loadable (hot) data:

```
SELECT * FROM "TPCH"."LINEITEM1_CS" WHERE ("TPCH"."LINEITEM1_CS"."L_RECEIPTDATE" >= '1998-
03-15' AND "TPCH"."LINEITEM1_CS"."L_RECEIPTDATE" < '1998-04-01') AND
"TPCH"."LINEITEM1_CS"."L_SHIPMODE" = 'RAIL';
```

Result: 2165057 rows selected (overall time 5314.299225 sec; server time 6132.825 msec)

As observed, the page-loadable data took approximately 105 minutes to load but the column-loadable data was loaded in 88 minutes. This serves the purpose of NSE where priority is given to the hot data residing in main memory. Warm data, which is less frequently accessed, takes more time to load from the disk to the buffer cache.

The query performance of warm data depends on whether it is already available in the buffer cache, and if it is not, how quickly it can be read from the underlying storage. Because the storage is just an “extension” of the persistence volume, the IOPS, bandwidth, and latency observed for the warm data is the same as it would be for the load of hot data.

Query results in Table 2 show how NSE effectively helps prioritize hot data (1997 to 1998) with respect to warm data (1992 to 1996) in SAP HANA. The response time for Query 1 is high for the hot data mostly because this query gathers the entire record from the LINEITEM table for the given criteria. The response times are lower for Query 2, which seeks only the discount values from the ORDERS table for rail shipments and for Query 3, where the total number of records is far less.

These results are for demonstration only. Query performance in an SAP environment depends on the database schema and the efficiency of the underlying infrastructure.



TABLE 2. Query performance for hot compared with warm data

Number	Queries	Number of records retrieved	Hot data	Warm data
Query 1	All line item details between Mar 15, 1993 and Apr 01, 1993 with rail shipments	2167327	NA	105 min
	All line item details between March 15, 1998 and Apr 01, 1998 with rail shipments	2165057	88 min	NA
Query 2	Discount values for orders between Oct 15, 1995 and Nov 15, 1995 with rail shipments	3954637	NA	131 min
	Discount values for orders between Oct 15, 1997 and Nov 15, 1997 with rail shipments	3953610	17 min	NA
Query 3	All order details between July 02, 1994 to July 03, 1994 with urgent priority	44618	NA	99 min
	All order details between July 02, 1998 to July 03, 1998 with urgent priority	44586	14 min	NA

You also can view these results in a bar chart, as shown in Figure 6.

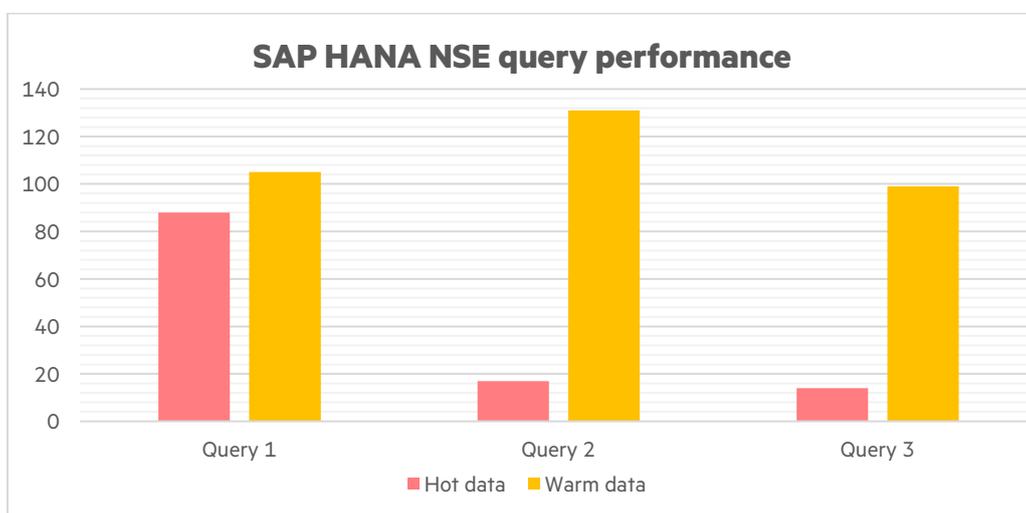


FIGURE 6: SAP HANA query performance for hot and warm data

Using HPE Primera, HPE Nimble Storage, and HPE 3PAR for NSE

You can use HPE Primera, HPE Nimble Storage, and HPE 3PAR systems to provision NSE. Because the storage tier for hot data stores the warm data as well, the only requirement to support NSE on a running SAP HANA system is to increase the data volume sizing based on these sizing guidelines. This means adding more SSDs on the back-end of the SAP HANA storage system. Tables 3, 4, and 5 list all the HPE Primera, HPE Nimble Storage, and HPE 3PAR models that are SAP HANA TDI certified.

TABLE 3. HPE Primera models for SAP HANA NSE

Feature	630	650 (two-node)	650 (four-node)	670 (two-node)	670 (four-node)
Max no. of SSDs	144	192	384	288	576
Max raw capacity	250 TiB	400 TiB	800 TiB	800 TiB	1.6 PiB



TABLE 4: HPE Nimble Storage models for SAP HANA

Feature	AF 40	AF 60	AF 80
Max no. of SSDs	48	48	48
Max raw capacity	167 TiB	502 TiB	1005 TiB

TABLE 5: HPE 3PAR models for SAP HANA NSE

Feature	8200	8400	8440	8450	9450	20450	20800 R2	20850 R2
Max no. of SSDs	120	240	480	480	576	576	1152	1152
Max raw capacity	838 TiB	1676 TiB	3351 TiB	3351 TiB	6000 TiB	4021 TiB	8043 TiB	8043 TiB

SUMMARY

This white paper describes how SAP HANA NSE technology works and how you can use HPE Storage to provision and employ it. The key goal with NSE is to use the existing SAP HANA hardware to improve the cost-to-performance ratio of the data that is stored within it.

This white paper also provides a sample dataset to illustrate how NSE is enabled, configured, and used. It considers an industry-standard dataset and provides several query examples that yield the expected results on hot as well as warm data.

Lastly, this paper describes various HPE Primera and HPE 3PAR models that are SAP HANA TDI certified and can be used to extend storage to store warm data with NSE. HPE Storage supports all possible warm data management options that SAP HANA offers, including NSE, dynamic tiering, and extension nodes.



Resources and additional links

- [SAP HANA Administration Guide](#)
- [HPE Primera All-Flash for SAP HANA Enterprise Storage](#)
- [HPE Primera 600 quick specs](#)
- [HPE 3PAR StoreServ Storage](#)
- [HPE 3PAR StoreServ Storage—SAP HANA TDI Certified](#)
- [HPE Nimble Storage](#)

LEARN MORE AT HPE AND SAP STRATEGIC ALLIANCE

h22168.www2.hpe.com/us/en/partners/sap/index.aspx

Check if the document is available
in the language of your choice.



Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



Share now



Get updates

© Copyright 2020 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

SAP HANA, SAP, and NetWeaver are trademarks or registered trademarks of SAP SE in Germany and in several other countries. Apache and Hadoop are registered trademarks of the Apache Software Foundation. All third-party marks are property of their respective owners.

a00095010ENW, January 2020