HITACHI
Inspire the Next

SAP HANA Native Storage Extension Implementation with Hitachi Virtual Storage Platform

Best Practices Guide

By Hitachi Vantara

March 2021

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hitachivantara.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Revision History

Revision	Changes	Date
MK-SL-221-00	Initial release	March 10, 2021

Table of Contents

SAP HANA Native Storage Extension **1**

SAP HANA Data Classification **2**

SAP HANA Native Storage Extension (NSE) Support on VSP 5000 series, VSP E series, VSP F series, and VSP G series 2

Getting Started with SAP HANA NSE **3**

SAP HANA NSE Sizing **3**

Configuring SAP HANA NSE **6**

Recommendations **8**

SAP HANA Native Storage Extension Implementation with Hitachi Virtual Storage Platform

Best Practices Guide

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 the value of the data decreases over time, then there should be alternative storage options for this older data. This enables the new and more important data to reside in memory, and the older data can still be used and accessed.

This best practice document describes one such offering from SAP namely Native Storage Extension (NSE). With NSE, 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 separate compute hardware.

This paper describes SAP HANA NSE as a method to tier data. It assumes knowledge of SAP applications, SAP HANA, and databases in general and also Hitachi Vantara VSP storage systems such as VSP 5000 series, VSP E series, VSP F series, and VSP G series to leverage NSE technology.

SAP HANA Native Storage Extension

SAP HANA Native Storage Extension (NSE) is a general-purpose, built-in warm data store in SAP HANA that lets you manage less-frequently accessed data without fully loading it into memory. SAP HANA Native Storage Extension (NSE) uses techniques to only load the pages into memory that include data that is relevant to your search. Pages containing data that is not accessed by your query are not loaded from disk.

The capacity of a standard SAP HANA database is equal to the amount of hot data in memory. However, the capacity of a SAP HANA database with NSE is the amount of hot data in memory plus the amount of warm data on the disk. The following figure shows the difference between standard HANA database without NSE and standard HANA database with NSE.

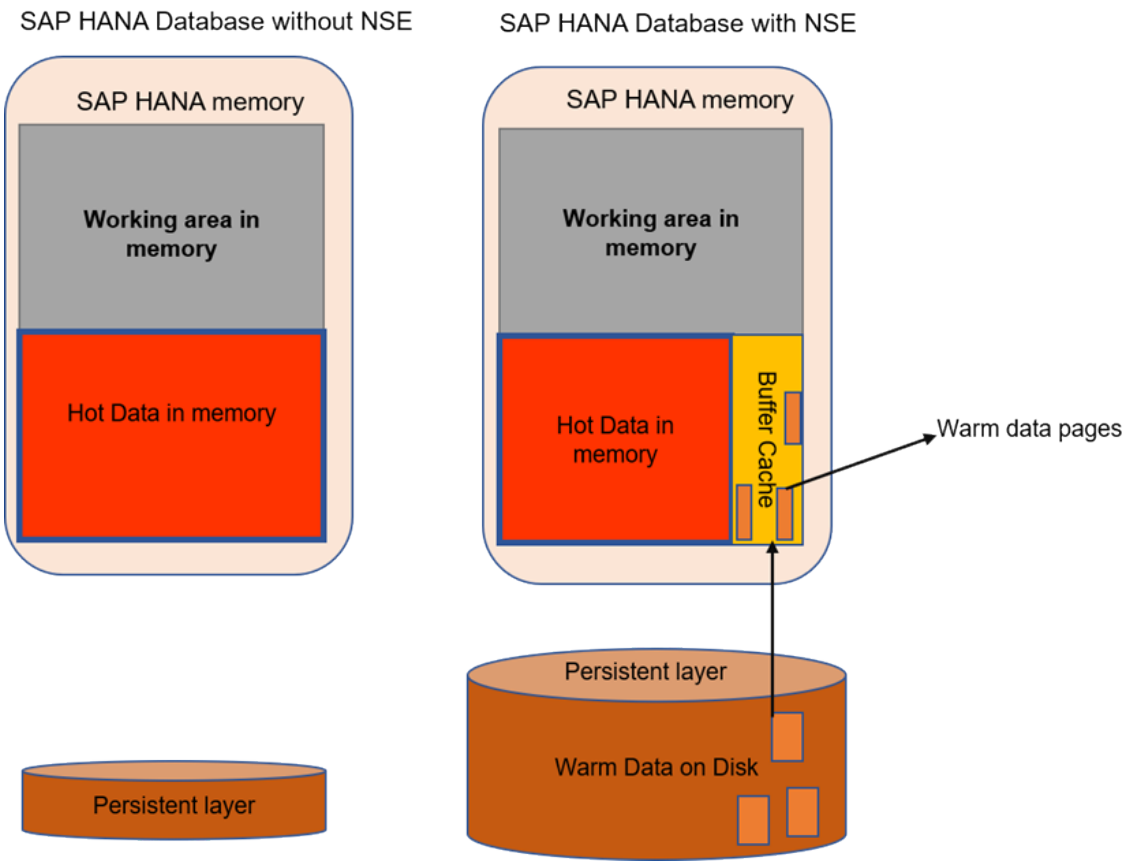

SAP HANA Data Classification

SAP HANA offers various software solutions to manage multi-temperature data (hot, warm, and cold), such as DRAMs for hot data, SAP HANA Native Storage Extension for warm data, and SAP HANA Cold Data Tiering for cold data. This document ifocuses on managing warm data leveraging SAP HANA NSE. SAP classifies data in the SAP HANA database into three categories:

- **Hot Data**: Hot data is of the highest importance and value, and it is critical to the business for real time processing and analytics. This data is loaded and retained continuously in SAP HANA memory for fast processing.
- **Warm Data**: This data is less frequently accessed and need not reside continuously in SAP HANA memory. However this data still needs to be a part of the SAP HANA database and is stored in lower cost stores within SAP HANA. NSE manages this warm data in the HANA database as an expanded disk capacity. Query performance on warm data might be reduced compared to hot data.
- **Cold Data**: Cold data is rarely accessed, but it still needs to be stored for archival or compliance purposes. This data is managed separately from the SAP HANA database. The hardware needed to hold this data can be low-cost commodity storage with low performance considerations.

Data can be converted between column loadable" and page loadable". This makes it possible to move the data back and forth from the hot tier to the warm tier. Therefore, it is more efficient to use storage that provides faster reads to move data back to the hot tier if a query or use case demands it. SAP HANA determines which data should reside permanently in memory and which should be offloaded to disk.

SAP HANA Native Storage Extension (NSE) Support on VSP 5000 series, VSP E series, VSP F series, and VSP G series

SAP HANA Native Storage Extension was introduced with SAP HANA 2.0 SPS 04.

- SAP HANA 2.0 SPS 04 features:
  - Scale-up only native SAP HANA applications supported
  - S/4HANA with data aging
- SAP HANA 2.0 SPS 05 features:
  - Scale-up and scale-out
  - Native SAP HANA, BW/4HANA applications supported
  - S/4HANA with data aging

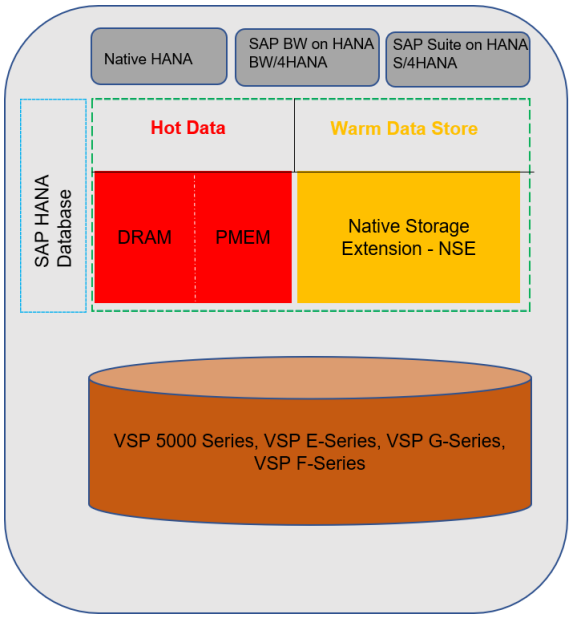

Note — S/4HANA can leverage SAP HANA NSE using the data aging method where NSE uses historical data, an attribute assigned by data aging at the time of table/partition/column creation, to load into buffer cache. See SAP Note 2816823 for more information. A SAP login is required.

Getting Started with SAP HANA NSE

The following terminology is used in NSE:

- **NSE Advisor**: The NSE Advisor is a tool that helps you determine hot data and warm data based on how frequently it is accessed at the query level for each column, (that is, columns, partitions, and tables) as candidates. You can use the NSE Advisor to provide adjustable thresholds for the following data:
  - Warm data, which is rarely accessed and stored on disk
  - Hot data which is frequently accessed and stored in-memory.

  See SAP HANA NSE Advisor for more information.

- **Load Units:** The different types of load units that can be configured are Column-loadable, Page-loadable, and Default**.**

  The system default load unit is Column-loadable.

  - **Column-loadable data:** Data that is frequently accessed is column-loadable, and it is loaded entirely into the memory. This data can be a table, a partition, a column, or the part of a table that is not page-loadable load unit. Column-loadable data is fully loaded into the memory from the disk. This is called as the hot data. The default load unit for the hot data is column-loadable.

    The query performance for a column-loadable component is usually better than the page-loadable data because column-loadable data always resides in the main memory. SAP HANA loads all the data into the memory for fast processing.

  - **Page-loadable data:** You can use NSE to specify a table, a partition, a column, or an index, which is less frequently accessed data as page loadable. This is called as the warm data.
    - In this load unit, the data is partly in-memory, and partly on the disk. Page loadable data is loaded from the disk into the memory in granular units of pages as required for query processing, thereby reducing its memory footprint. This data is loaded in the buffer cache part of the main memory of the SAP HANA system.
    - **Default:** Default indicates that there are no explicit load unit preferences configured for a database object. A database object with the tag DEFAULT LOADABLE inherits the load unit preference of its parent object.

- **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. Converting column-loadable data to page- loadable format moves the data into the buffer cache. This memory is reserved but not allocated unless the page-loadable data is accessed. When the buffer cache is full, it will eject pages intelligently based on user access patterns.

SAP HANA NSE Sizing

You should allocate enough buffer cache to avoid accessing the disk for queries. An application's performance varies based on the amount of data that is available in the buffer cache and how much data the application must read from the disk. Although the size of the buffer cache with respect to the warm data size varies across applications, the ratio of the buffer cache to the warm data on the disk should not be smaller than 1 is to 8. If your frequently accessed NSE data set is small, then you can experiment with a larger ratio of warm storage to buffer cache.

For example:

Buffer Cache by default is 10% of RAM and can be adjusted.

Buffer Cache = Warm Data/8 (Warm Data divided by 8)

Buffer Cache × 8 = Warm Data

The following figure considers 10% buffer cache.

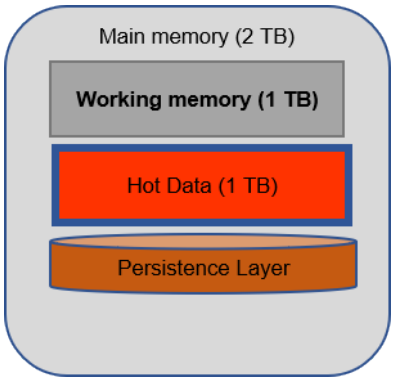

To accommodate buffer cache, do one of the following:

- Add more memory to the system for buffer cache.
- Reserve from existing HANA memory. This will reduce the hot data volume in memory..

This document concentrates on reserving buffer cache from existing memory, as it is more cost-effective and provides optimal performance. The following figure shows the buffer cache as part of memory.

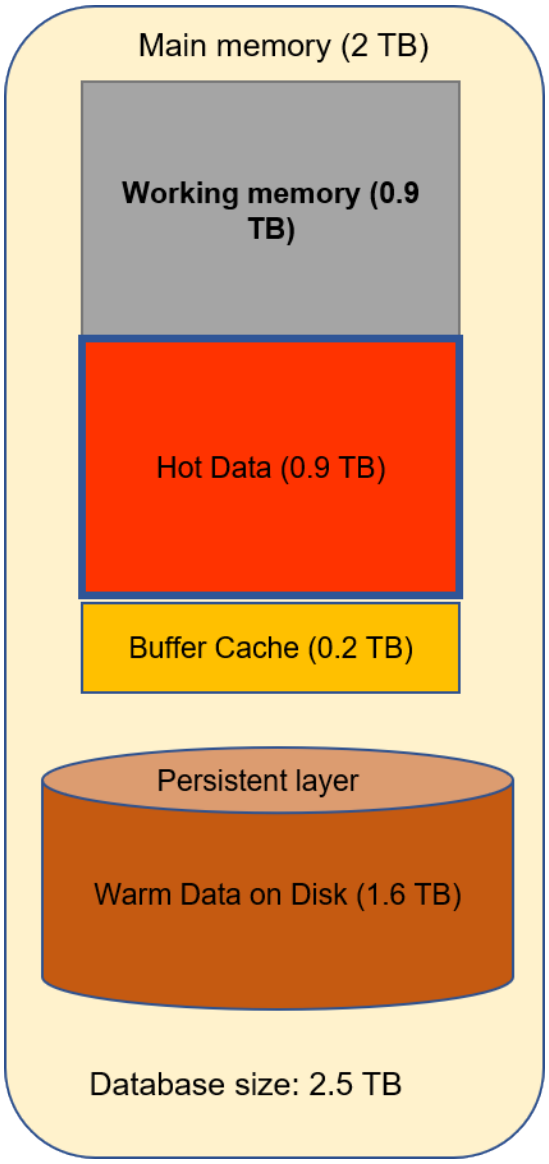

Configuring SAP HANA NSE

SAP HANA Cockpit, HANA Studio, or the HDBSQL command line utility, is used to configure SAP HANA NSE. The following figure shows the steps to configure SAP HANA NSE while designing a new schema and tables, or updating an existing application schema and tables.

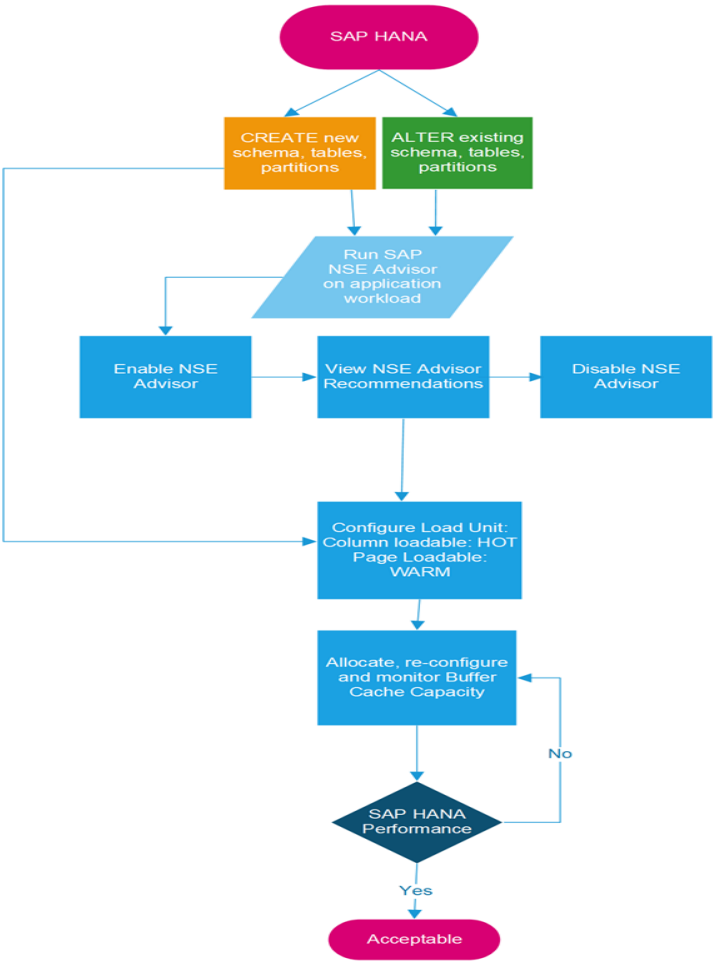

1. Run the NSE Advisor on an application workload.
   (1) Enable the NSE Advisor.
   (2) Get the NSE Advisor recommendations.
   (3) Disable the NSE Advisor.
2. Configure the Load Unit.
3. Allocate and Manage the Buffer Cache Capacity.
4. Monitor the Buffer Cache statistics and reconfigure if needed.

Run NSE Advisor

The SAP HANA NSE Advisor tool (NSE Advisor) determines the temperature of the data and uses rule-based algorithms to derive these recommendations. In particular, the NSE Advisor tool identifies hot and warm objects as candidates to be either page-loadable or column-loadable.

NSE Advisor creates recommendations in two steps:

1. It observes the existing queries from the application to the database and creates a data access pattern.
2. It applies heuristics to these data access statistics and generates recommendations.

Prerequisites

To use NSE Advisor, you must have the CATALOG READ or the DATA ADMIN system privilege to run the SQL statements used in this task.

The following queries enable NSE Advisor and help determine the temperature of the data based on how frequently it is accessed at the query level for each column.

Enable NSE Advisor

```
ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini','system') SET ('cs_access_statistics','collection_enabled') = 'true' WITH RECONFIGURE;
```

Get the recommendation results

```
SELECT * FROM SYS.M_CS_NSE_ADVISOR
```

Disable NSE Advisor

```
ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini','system') SET ('cs_access_statistics','collection_enabled') = 'false' WITH RECONFIGURE;
```

See the following links for more information:

Using the NSE Advisor with SQL Commands for Warm Data Recommendations

Using the SAP HANA Cockpit, use NSE Advisor for warm data recommendations

Note: After the workload has been running for a few hours, you can start getting recommendations from NSE Advisor indicating data temperature. See the SAP HANA NSE Advisor help page for more details.

Configure the Load Unit

Based on NSE Advisor recommendations, data can be specified as page-loadable at the table level, partition level, and column level by using a CREATE statement for new tables or ALTER statements for existing tables and partitions.

See the following links for more information:

Setting the Load Unit Using CREATE TABLE

Setting the Load Unit Using ALTER TABLE

Allocate and Manage Buffer Cache capacity

The buffer cache is automatically enabled in a SAP HANA system, and by default 10% of the total memory of the server is reserved for it. This memory is reserved, but not allocated unless the warm data is accessed. Otherwise, 100% of the memory is available for SAP HANA for the hot data.

Use the following parameters to dynamically expand or shrink the buffer cache size. Ensure that your buffer cache is large enough to manage your warm data set: max_size: Explicitly specifies the upper limit of the buffer cache, in MBs.

```
o  ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini', 'system') SET
   ('<cache_name>','max_size') = '<cache_size>'
```

- `max_size_rel`: Specifies the upper limit of the buffer cache as a percentage of the SAP HANA memory.

```
o  ALTER SYSTEM ALTER CONFIGURATION ('indexserver.ini', 'system') SET ('buffer_-
   cache_cs','max_size_rel') = '10' with reconfigure;
```

See SAP HANA NSE Buffer Cache for more details.

Monitor the Buffer Cache statistics

The following statements monitor the SAP HANA Buffer cache:

Select * FROM M_BUFFER_CACHE_STATISTICS

Select * FROM M_BUFFER_CACHE_POOL_STATISTICS

See Monitoring the SAP HANA NSE Buffer Cache for more details.

Recommendations

SAP HANA NSE supports any SAP HANA application. The NSE feature in SAP HANA does not require you to modify your applications. However, existing applications do not automatically use its capabilities, and thus database administrators are responsible for making database configuration changes to leverage the NSE feature, such as specifying a load unit in configuration parameters while creating or altering database tables.

SAP HANA 2.0 calculates 10% of memory for the buffer cache by default; you do not need to perform any additional steps to enable the buffer cache. This memory is only reserved and not allocated. SAP HANA accesses 100% of its memory, including the 10% reserved for the buffer cache if you are not using NSE.

SAP recommends the working memory to be the same size as the hot data in memory.

If there are page-loadable tables in your current version of SAP HANA and you move to another version, then only those tables that were designated as page-loadable in the earlier version use the buffer cache in the later version. This is up to the limit that was calculated in the original version of SAP HANA that you were running.

Depending on the usage pattern and buffer cache size, SAP HANA must read the warm data from the disk, which impacts the performance of the queries compared to the performance on the hot data. You can convert data from column-loadable to page-loadable and from page-loadable to column-loadable..

Objects that are small and accessed frequently should be kept in memory to improve performance.

Objects that are large and rarely accessed should be kept on the disk to reduce memory usage. These objects are brought in memory only when needed, and only for the pages that are accessed.

Follow the recommendations of Hitachi Vantara SAP HANA reference architecture guides for storage configuration with respect to RAID group, parity group, LUNs, and other settings for the corresponding solution.

Because the warm data on disk generates more IO traffic when it is accessed, consider provisioning two more LUN-paths between server and storage, if needed.

To move a portion of an unpartitioned table to SAP HANA NSE, you must first partition the table.

Using NSE for a given database object, that is, a table, a table column, or a table partition, depends on the workload relating to the object in the individual environment.

To use NSE with a given database object, you must assess and verify that the associated performance impact or other side effects are acceptable. Also, consult the SAP HANA database documentation regarding other implications of using NSE for database operations, monitoring, and other functions. A SAP login is required.

Systems leveraging SAP HANA NSE should be implemented first as a proof of concept (POC) in a test environment, and you should ideally work directly with your SAP account team for expert guidance. See the following SAP recommendations for more details.

For More Information

Hitachi Vantara Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Services website.

Demonstrations and other resources are available for many Hitachi products. To schedule a live demonstration, contact a sales representative or partner. To view on-line informational resources, see the Resources website.

Hitachi Academy is your education destination to acquire valuable knowledge and skills on Hitachi products and solutions. Our Hitachi Certified Professional program establishes your credibility and increases your value in the IT marketplace. For more information, see the Hitachi Vantana Training and Certification website.

For more information about Hitachi products and services, contact your sales representative, partner, or visit the Hitachi Vantara website.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
www.HitachiVantara.com | community.HitachiVantara.com

Regional Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact

© Hitachi Vantara LLC, 2021. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. VSP is a trademark or registered trademark of Hitachi Vantara LLC. Microsoft, Windows Server, and Microsoft Office are trademarks or registered trademarks of Microsoft Corporation. All other trademarks, service marks, and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Vantara LLC.

MK-SL-221-00, March 2021