


SCM FOR HANA – HOW IT GENERATES A NEGATIVE TCO



Amal D Nair

Thejushree K

Richard Holmes

Product and Solutions Sales Specialist 4
Dell Technologies

The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across Dell's multiple technologies and products with both skill and outcome-based certifications.

Proven Professional exams cover concepts and principles which enable professionals working in or looking to begin a career in IT. With training and certifications aligned to the rapidly changing IT landscape, learners can take full advantage of the essential skills and knowledge required to drive better business performance and foster more productive teams.

Proven Professional certifications include skills and solutions such as:

- Data Protection
- Converged and Hyperconverged Infrastructure
- Cloud and Elastic Cloud
- Networking
- Security
- Servers
- Storage
- ...and so much more
-

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person Classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

Contents

- Executive Summary..... 4
- 1. Introduction 4
- 2. SAP HANA Workloads..... 4
 - 2.1 Introduction 4
 - 2.2 Features of SAP HANA..... 6
 - 2.3 Deployment Models..... 7
- 3. HANA Native Storage Extension (NSE)..... 8
 - 3.1 Cases that support the Implementation..... 10
 - 3.2 Difference with & without an NSE 10
- 4. Tiers in HANA NSE 11
 - 4.1 NSE Advisor – Identifying the Data 13
- 5. Value Proposition & Negative TCO 14
 - 5.1 Negative TCO..... 14
- 6. Conclusion..... 16
- 7. References 17

Executive Summary

With the modernization of solutions and all the organizations aligning themselves with Digital Transformation journeys, the need for analytics and business insights also increases for organizations. The organization deploys various kinds of applications on the system to have the upper edge over its competitors and be the best in the market. There are traditional applications and modern applications as well that need to be deployed on a common platform while the requirement of each of these varies. This article talks an overview of one such application workload called HANA and intends to take a bottom-up approach to understand the various subsets and how efficiently it can be deployed.

1. Introduction

Since the introduction of digital transformation & modernization, With the increasing amount of data being generated and stored, businesses and organizations need more robust and efficient storage solutions to manage and analyze their data. The market for storage solutions has seen a significant shift towards cloud-based storage, as well as the emergence of new technologies such as software-defined storage and hyperconverged infrastructure. Additionally, there is a growing demand for storage solutions that can support diverse application workloads, including big data analytics, artificial intelligence, and edge computing. In this whitepaper, we will address various technology topics and deep dive into subsets to understand SAP HANA, explore the Data tiering options available, and take a deep dive into implementation with and without NSA. But to have a clear understanding it is important to have a structured approach. This article is going to serve as the repository of single-stop information on how to generate a Negative TCO with the adoption and harnessing of the true potential of technology which is not properly documented and remains as tribal knowledge.

2. SAP HANA Workloads

2.1 Introduction

SAP HANA, short for High-performance Analytic Appliance, is a cutting-edge database system that utilizes in-memory technology to store and process data. Unlike traditional disk-based systems, SAP HANA keeps data in its memory, which greatly accelerates the speed at which data can be accessed and analyzed. This makes it possible to perform advanced analytics and real-time decision-making with large amounts of data, which would otherwise be too slow or complex to handle. This is achieved by storing the data in memory in a columnar format and using a specialized query engine that can work with the data in memory. This unique architecture allows for faster data processing, making it possible to perform complex analytics and reporting on large datasets in real time. This makes it suitable for various applications such as Business Intelligence, Predictive Analytics, and Real-time reporting.

An SAP HANA appliance comprises of integrated storage, computational, and networking components, it is certified and approved by SAP, manufactured by an SAP HANA hardware partner, and delivered to customers with all software components pre-installed, including the Operating System.

SAP HANA acts as a foundation for enterprise resource planning (ERP) software and other business tools and can be implemented either on a company's own servers, through cloud-based services, or through a combination of both in a hybrid cloud system. ERP is one of the most common adoptions that is seen in the market today and it provides software that organizations use to manage and automate various business processes, such as financial management, human resources, supply chain management, and customer relationship management. These systems help companies to integrate and manage data from

different departments and business functions in a centralized system, providing real-time visibility and control over their operations. Figure 1 displays the pictorial representation of what an ERP means and integrates



Figure 1: A Pictorial Representation of ERP

(Source: <https://www.sevenstarwebsolutions.com/erp-software/>)

This capability enables organizations to make better-informed decisions, improve efficiency, and reduce costs. It brings together data from various sections within a company, such as:

- (A) **Traditional business documents**, Including contracts and spreadsheets.
- (B) **UX/UI (User Experience/User Interface)** which includes a website, emails, forms, Interfaces, and other customer interaction modes.
- (C) Information from the **mobile devices** of customers that the organization is catering to and the workforce.
- (D) **IoT (Internet of Things)**, in which data from the many sensors that run in every aspect of a business, from warehouses to working machinery, to CCTV cameras and trucks to stores and offices. It is implemented in every subset one can think of to bring better integrations.

With these benefits it brings to the table it becomes an obvious question why is SAP HANA Important or why are organizations turning towards SAP HANA and gathering more and more knowledge based on planning and implementing the solution appropriately?

IMPORTANCE: If asked to mention the importance of SAP HANA, it can be best summarized by mentioning that many organizations have a large amount of data that remains unused in data warehouses, representing a huge untapped potential that could be utilized to improve business outcomes and deliver greater value to customers.

By using its high-speed, in-memory processing, and real-time data analytics powered by machine learning, SAP HANA can tap into a company's data and put it to work, optimizing processes, reducing errors, and providing benefits to clients, customers, and employees.

It can enhance customer satisfaction by providing real-time product availability that can be accessed from any smartphone. It can improve employee satisfaction by simplifying accounting processes, providing instant updates on benefits or vacation time, and enhancing collaboration platforms. It can also assist managers and executives in planning for the future by providing predictive analytics that can predict supply chain issues, manage cash flow, and integrate team workflows for improved efficiency.

2.2 Features of SAP HANA

To reiterate SAP HANA's in-memory, multi-model data management engine maximizes the capabilities of its hardware to minimize data movements, thus increasing the speed and flexibility of the system as it analyses real-time data. Depending on the needs of an organization, SAP HANA can be deployed on-premises, in the cloud, or as a hybrid system, which combines the privacy and control of the on-premises system with the cost-effectiveness, greater memory, and increased accessibility of the cloud.

Its ability to process large amounts of data efficiently makes it easily scalable for a growing business, without compromising security or stability. On the SAP HANA platform, developers can create their own tools and applications that integrate business logic, control logic, and the database layer with exceptional performance.

The following can be broadly classified as some of the features of SAP HANA.

- I. **Database Enhancement:** While traditional database systems are typically based on a single data model that constrains how data is organized, stored, and handled, SAP HANA's multi-model database allows for multiple data models to coexist on a single integrated backend, making it simple to establish connections among data points.
- II. **Data Management Enhancement:** SAP HANA's in-memory technology eliminates the need for time-consuming batch processing by keeping all data readily accessible with no delay, allowing all processes to occur in real time. This enables organizations to simplify their management systems and adapt them as necessary to improve efficiency.
- III. **Advanced Analytics Processing:** With real-time analytics, it is possible to access large amounts of data, including data from IoT sensors and mobile devices. SAP HANA's machine learning engine can extract and write data to the server in real time, immediately identifying issues and developing solutions across all aspects of an organization, including payroll, human resources, customer service, supplier management, and more.
- IV. **Application Development:** The SAP HANA database serves as the foundation for a software stack system that acts as a technology-agnostic server for any web-based applications, whether they are SAP or non-SAP, that a developer wishes to create. As the server can interact with all file types and exchange data with any software, businesses can quickly develop and implement custom queries and processes that are tailored to their specific needs.
- V. **Tools Enhancement:** SAP HANA's advanced analytical processing enables developers to create simple tools that business teams can use to generate customized reports without the assistance of IT or external consultants.

All of these can be considered as some of the features that make SAP HANA so compelling, but there are various deployment models too that one can choose while going with SAP HANA. The next subset of this whitepaper article covers the same.

2.3 Deployment Models

SAP HANA can be deployed in two different models, As an appliance model or as a Tailored Data-centre Integration (TDI) Model. The differences between both deployment models are as follows:

- **Appliance Model**
An SAP HANA appliance comprises of integrated storage, computational, and networking components. It is certified by SAP, manufactured by an SAP HANA hardware partner, and delivered to customers with all software components pre-installed, including the Operating System. Figure 2 displays the deployment architecture of the Appliance model.

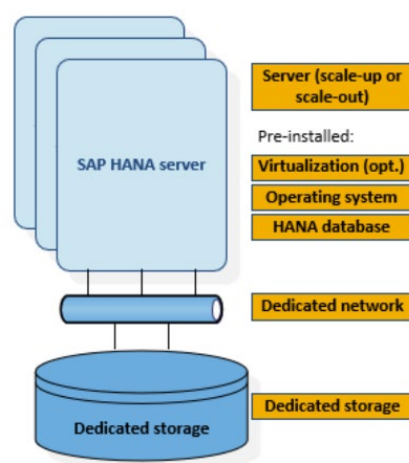


Figure 2: Appliance Model of SAP HANA
(Source: <https://infohub.delltechnologies.com/>)

However, the SAP HANA appliance model presents the following limitations for customers:

- (a) Restrictions on the selection of servers, networks, and storage
- (b) Inability to utilize current data centre infrastructure and procedures
- (c) Lack of control over the vital components in the appliance
- (d) Fixed sizes for SAP HANA server and storage capacities, resulting in higher costs from a lack of capacity and inability to quickly adapt to unexpected growth demands.

- **Tailored Data Centre Integration (TDI) Model**
The TDI deployment model gives customers the option to choose from a wide range of SAP HANA-certified servers that can be paired with SAP-certified network and storage components. Various workloads can share the storage and network components to minimize the total cost of ownership (TCO). Customers can easily integrate SAP HANA into their existing data center

operations such as disaster recovery, data protection, monitoring, and management, which reduces the cost, time-to-value, and risk of an SAP HANA adoption. Figure 3 displays the deployment architecture of the TDI model.




Figure 3: TDI Model of SAP HANA
(Source: <https://infohub.delltechnologies.com/>)

Now that the concepts of SAP HANA and its various subsets are laid out, let us look at how we can implement the SAP workloads more efficiently by using the various add-on features of HANA.

3. HANA Native Storage Extension (NSE)

To put the concept as simple as possible, HANA Native Storage Extension (NSE) is an add-on feature to the SAP HANA database that allows customers to manage and store large amounts of data on commodity hardware. The NSE feature is an alternative to traditional storage solutions, such as SAN or NAS, and enables customers to scale their storage capacity in a more cost-effective way.

With the HANA NSE, customers can store a large amount of data on low-cost hard disk drives (HDD) and access it at a much faster rate. The HANA NSE enables customers to take advantage of high-capacity storage, at a lower cost, and with high performance.

Although NSE is not a new idea, companies rarely use it. Based on various articles references and documentation we would strongly advise readers whose systems are already on HANA or planning to implement HANA soon to consider implementing NSE.

To cascade the messaging of SAP HANA in simple terms, it can be said that NSE gives the option to store less frequently accessed data in HANA to a disk-based column store, which reduces the total cost of ownership (TCO) for DRAM (Dynamic random-access memory).

Data can be stored at different levels in HANA depending on the usage frequency–

- **Hot Data** – It is used to save essential data for real-time processing and analytics. It is kept constantly in SAP HANA memory for fast performance and is stored in the highest performance (and highest TCO) storage.
- **Warm Data** – It is primarily used to store mostly read-only data that do not need to be accessed frequently. The data do not need to reside continuously in SAP HANA memory, but it is still managed as a unified part of the SAP HANA database, transactional consistency with hot data, participating in SAP HANA backup and system replication operations. It is stored in lower-cost storage options within SAP HANA.
- **Cold Data** – It is used to store read-only data that is accessed very rarely. Cold data is managed separately from the SAP HANA database, but it can still be accessed from SAP HANA through SAP HANA's data federation capabilities.

One of the major building blocks of HANA NSE is Buffer Cache, Buffer cache in HANA NSE is a memory area used to temporarily store data that is frequently accessed. The buffer cache is used to reduce the number of disk I/O operations that are required to access data and improve the performance of the system. When a user or application requests data, the buffer cache is checked first to see if the data is already in memory. If the data is not in the buffer cache, it is retrieved from the disk and placed in the buffer cache for future use.

The buffer cache is dynamically managed by HANA, and the data that is least frequently accessed is removed to make room for new data. Figure 4 shows a detailed look at how the HANA NSE Functions.

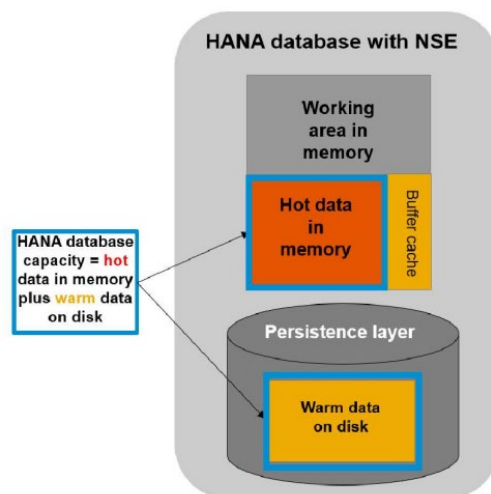


Figure 4: Building Blocks of HANA NSE

(Source <https://www.sap.com/documents/2019/09/4475a0dd-637d-0010-87a3-c30de2ffd8ff.html>)

3.1 Cases that support the Implementation

Some of the most typical examples or cases that support the implementation of HANA NSE are as follows:

- To aid in the reduced HANA memory requirements both in terms of performance and implementation cost for the nonproduction systems of the organization.
- In the case of an acquisition or merging of organization, support in favour of the acquisition but at the same time help keep the hardware cost in check.
- Reduce the over TCO of the organization by effectively placing the data in the correct tier.
- Keep the expense of adding more HANA Memory for the organization's organic growth under control.

There are a variety of cases or use cases that point the users towards the implementation of a HANA NSE, but these are the most common encounters that are seen in the most recent times regarding implementation.

3.2 Difference with & without an NSE

To accurately understand the value that HANA NSE brings to the table it is important to understand how a HANA database works without NSE and with NSE in a detailed side-by-side comparison. So, a HANA deployment without NSE has typically termed as simply plain HANA in-memory storage. We have discussed in detail what a HANA system with NSE brings to the table and how it benefits the customer. The term Hana In-memory storage (without NSE) is as follows.

HANA In-Memory Storage in direct terms means that all the data that is being processed is stored in the main memory (RAM) of the server, rather than on traditional disk-based storage. This allows for much faster access to the data, as the data does not need to be retrieved from disk and can be accessed directly from memory.

One of the key benefits of in-memory storage is that it allows for faster data processing, as the data can be accessed and manipulated much more quickly than if it were stored on disk. This makes it ideal for use cases such as real-time analytics, online transaction processing, and business intelligence, where fast access to data is critical.

The benefit of HANA in-memory storage is that it allows for much more advanced analytics, as the data can be processed and analyzed in real time. This allows for the creation of sophisticated, real-time analytics and reporting applications, as well as the ability to perform complex calculations and data modeling on-the-fly.

So, as we can see what HANA without NSE looks like, but the same in the case of HANA with NSE implementation we see that the layer extends itself onto the storage to use the capacity within the drives to be implemented or give a higher performance for the demand of the application as buffer cache.

Now even though there are various models of implementation, the use of SCM or Storage class memory drives can tremendously shift the performance requirement or saturation it puts up on the system. The deployment type can vary from using 100% SCM drives as a complete solution for the system (for mission-critical workloads), to SCM as a tier for the system to aid in the warm data stores or complete

SSD drives for the rest of the use cases of cold storage. Figure 5 best summarizes the differences between HANA with NSE and without NSE.

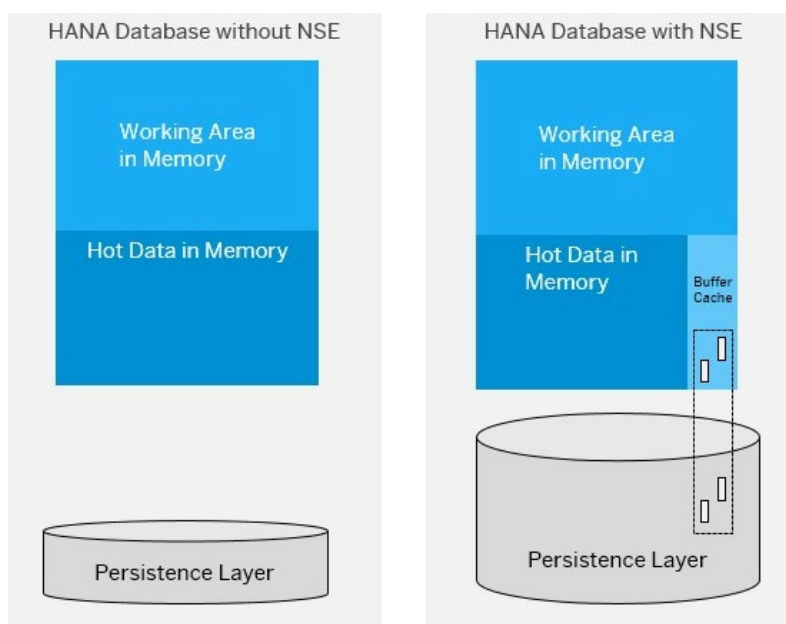


Figure 5: HANA In Memory Storage Vs HANA NSE

(Source <https://www.sap.com/documents/2019/09/4475a0dd-637d-0010-87a3-c30de2ffd8ff.html>)

4. Tiers in HANA NSE

- Hot Data

The data that is crucial for real-time analysis and decision-making is stored in the main memory of the SAP system, known as the HANA Native Storage Extension (NSE), for optimal performance. In the retail industry, this may include customer information from the last 2-3 years in a call center's CRM system, while older data is considered less important. According to SAP, the "Hot store" is data that is essential for real-time processing.

In this type of deployment, we see that the hot data is moved onto the drives that are the top tier in performance i.e., SCM or Storage Class Memory Drives. The next type of data for consideration is warm data.

- Warm Data

In as simple term as possible, warm data refers to data that is not frequently accessed. NSE can assist in reducing the total cost of ownership (TCO) by storing this type of data on a lower-cost storage tier, while still maintaining it as a unified part of the SAP HANA database. The usual deployment is that SCM drives (or any high-performing drives) as a tier, helping in placing the

hot data in high-performing drives that extend as an NSE and the warm data in lower-cost drives.

It is relevant to mention the terms OLTP (Online transaction processing) and OLAP (Online analytical processing) as they pertain to this context. OLTP captures, stores, and processes real-time transaction data, while OLAP utilizes complex queries to analyze historical data that has been aggregated from OLTP systems. With NSE, OLAP may not be as successful or may produce slow results. Table 1 Displays tiering overview in HANA.

HANA on-premise	DRAM		PMEM		Native Storage Extension – NSE		NLS IQ (BW only)	SAP HANA cold data tiering (Vora) Spark Controller / Hadoop
					Extension Node			
					Dynamic Tiering			
Hot Store			Warm Store			Cold Store		
Persistent Memory (PMEM) extends the in-memory storage capacity for hot data in HANA.			Native Storage Extension (NSE) is an intelligent, built-in disk extension for the SAP HANA in-memory database . It is an evolutionary successor of Paged Attributes. First version is planned for on-premise with HANA SPS04 in Q2 2019. NSE for HaaS will follow in Q4 2019. (Extension nodes and dynamic tiering will continue to be offered as on-premise options.)			SAP HANA cold data tiering (Vora) provides persistence capabilities for HANA cold data in external data stores, e.g. HDFS, GCP, S3, ADLS. Hadoop/Spark is a low cost option for on-premise HANA. Data Hub as a Service is the future cold tier for HaaS.		
HANA as a service (HaaS)	DRAM		PMEM (future)		Native Storage Extension – NSE		(Data Hub as a Service or Big Data Services - future)	

Table 1: Tiering Overview in HANA

(Source <https://www.sap.com/documents/2019/09/4475a0dd-637d-0010-87a3-c30de2ffd8ff.html>)

- Cold Data

In the case of Cold data, it can be termed or grouped as the data that is not at all frequently used or are archived for long term retention. Ad-hoc reporting is the major purpose for this data. Data is maintained separately from the SAP HANA database but always accessible on low-cost storage tiers, such as disc or Hadoop.

With SAP HANA NSE, you can identify specific warm data as "page loadable," even while hot data is "column loadable," living entirely in memory for quick processing and loaded from the disc into SAP HANA memory in columns. When needed for query processing, this page-loadable material is loaded into memory one page at a time. Page-loadable data does not have to live entirely in memory, in contrast to column-loadable data.

4.1 NSE Advisor – Identifying the Data

The obvious question that follows the NSE tiering is how we identify the temperature of the data and aid in tiering. HANA has a feature called NSE Advisor which can be employed for this process, The NSE Advisor is a feature in SAP HANA that provides recommendations for determining whether data is best suited for column-loading (hot data) or page-loading (warm data) on a table-, partition-, or column-level.

However, it is important to note that this feature is not yet available in SAP HANA Cloud. When enabled, the NSE Advisor tracks column scans and materializations. After it has been disabled, it categorizes accessed objects based on their scan density and makes load unit recommendations for them.

These categories include:

- Objects are categorized as Hot objects, with high scan density which are recommended for column-loading,
- Objects with average scan density, with no recommendation, and
- Cold objects with low scan density, are recommended for page loading.

The NSE Advisor can also be configured by setting thresholds for hot and cold objects in relation to all accessed objects during a run.

Once the data is identified they can be tiered and stored onto the extension of buffer cache for fast retrieval and improved performance of the system. Figure 6 depicts an accurate flowchart of how the data is tiered and cascaded onto the buffer cache.

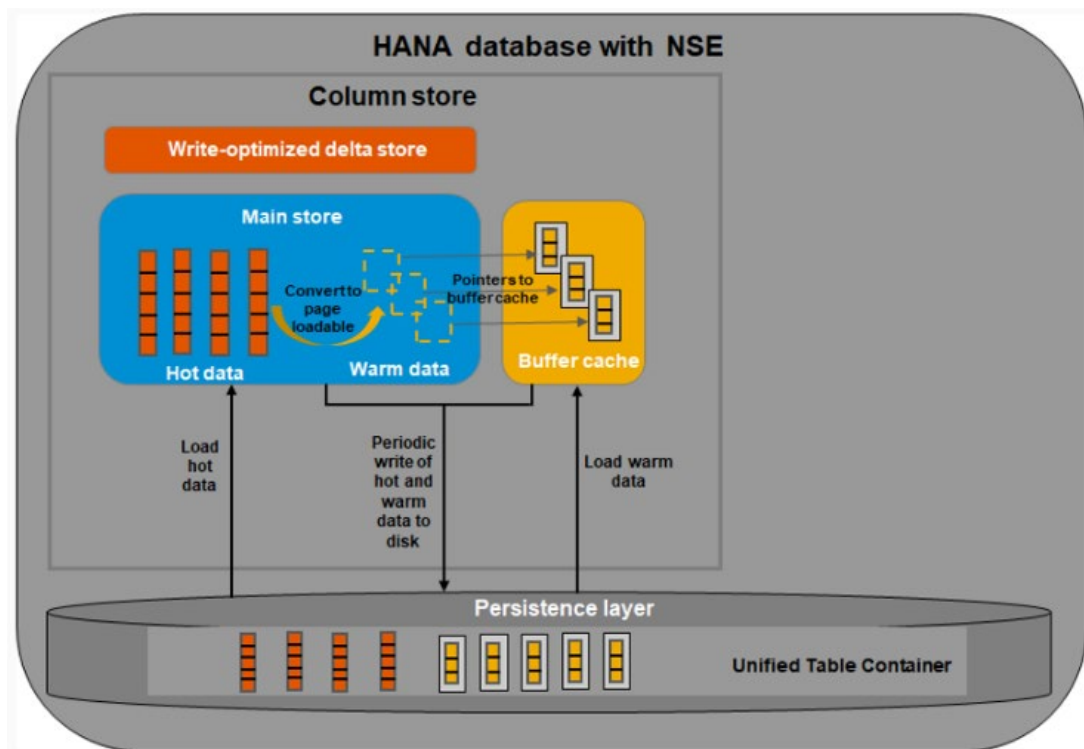


Figure 6: HANA NSE and Buffer Cache Relation

(Source <https://www.sap.com/documents/2019/09/4475a0dd-637d-0010-87a3-c30de2ffd8ff.html>)

5. Value Proposition & Negative TCO

The implementation of NSE it opens the door for organizations to leverage the best use of resources in the most efficient form by extending onto the layer and creating a buffer cache. The capabilities and benefits it brings to the table are immense. NSE also integrates and branches out by interoperability with other HANA functional layers such as column store, query execution engine, query optimizer etc.

Some of the values NSE brings among the many to the table are:

- (1) A remarkable increase in data capacity along with increased performance for high data volumes.
- (2) The distinctive ability that the NSE possesses to co-exist with the existing HANA in-memory column store, preserving SAP HANA memory performance.
- (3) The accessibility and control to monitor and manage buffer cache statistics via system views.
- (4) Can support SAP HANA Applications.
- (5) Availability of an advisor to collect statics about object access and provide load unit recommendations based on column store objects.
- (6) An intelligent buffer cache that manages memory pages in SAP HANA native storage extension column store tables

5.1 Negative TCO

Customers can work around optimizing SAP HANA with storage of their choice to achieve a negative TCO value for the complete system. As per the studies which are jointly done with Intel regarding using SCM for the “Warm” Store.

Dell Technologies PowerStore can aid the customers to achieve the same by providing a completely robust solution that can adopt this implementation model.

By adding an SCM drive which is the Storage Class Memory Drive (which is known to provide the sub-millisecond latency to HANA) the customers can reduce the host RAM requirement by up to 40%.

In a general situation where there are customers generally talking ATOS 8 or 16 socket servers with 1TB of RAM for each socket and a hefty price for SAP licensing per 1TB of RAM, it really adds up for just the compute that is contributed to the overall resource of the system.

As an example, Dell Technologies can position a solution with one PowerStore 5000T Appliance consisting of 21 Drives, each of 750GB SCM drives per HANA host with a list price like that of licensing 8 or 16 socket servers. Please note in this case it is not just the compute but the storage too that the user is getting. The solution is then positioned with additional PowerStore nodes with SSD for the Cold Store as well. With this approach or implementation model, the system portrays a 1/10 response time on the system for HANA workloads with SCM.

Providing an overall efficient solution for the customers that negates the TCO investment of the company on the system and typically PowerStore can be said as an “It pays for itself” solution while cascading the data to the users.

In a nutshell, we can redraw the tiering table in Table 2 with regards to the Powerstore as a solution to clearly cascade the message.

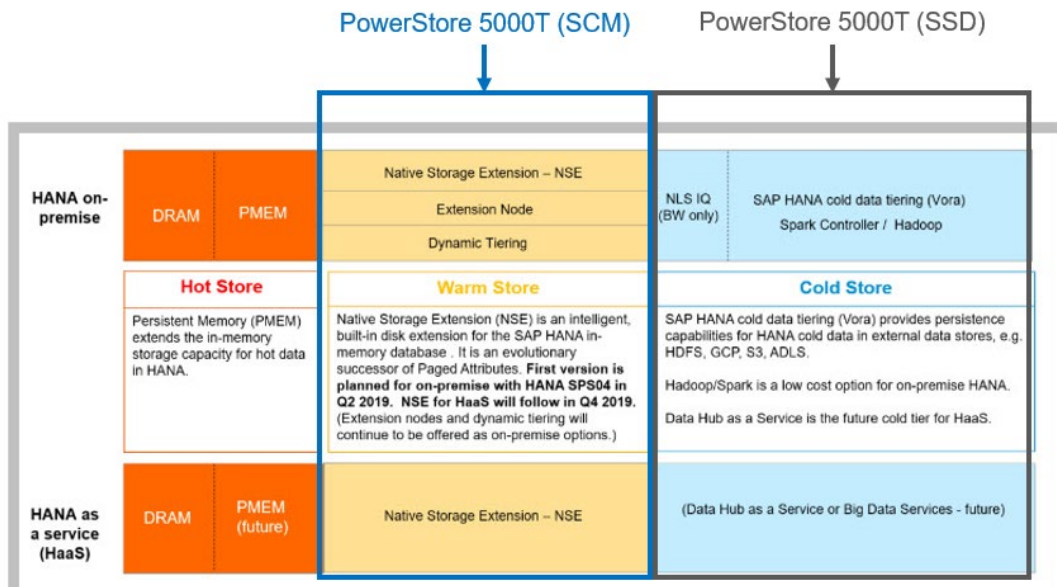


Table 2: PowerStore - Tiering Overview in HANA

Now in this case the PowerStore 5000T with SCM drives is deployed to cater to the performance requirements, HANA warm store, and reduce the expensive SAP licenses that are required which in turn reduces the overall investment cost of the solution. Furthermore, we also introduce a PowerStore 5000T with SSD drives which takes care of the HANA Cold Store and provides efficiency and consolidation for the overall solution.

6. Conclusion

As we have covered through this article, we do see that the scope of understanding of HANA NSE and its implementation to tier the data and leverage the technology are immense, but a fairly very low number of organizations are employing this capability and harnessing the true power of technology because of the lack of awareness and documentation that cascades this tribal knowledge onto the solution architects and product engineers. Cloud Analytics and the benefits it brings to the table are immense. It allows the diminution of TCO by simplifying IT landscapes that offer great performance, reducing the need for data duplication, and improving data governance and compliance management. Above all, NSE enables customers to access increasing volumes of data in a cost-effective and highly performant manner anytime in SAP HANA.

The article states that utilizing SAP HANA NSE can result in significant reductions in licensing expenses for the customer. It explains that mission-critical data that requires real-time processing and analysis is classified as "hot data" and is continually stored in SAP HANA memory for optimal performance, while data that is mostly read-only and infrequently accessed is considered "warm data" and can be temporarily stored on disk. Both "hot" and "warm" data are managed as a cohesive part of the SAP HANA database and are included in backup and replication processes, with the "warm data" being stored in more cost-effective locations within SAP HANA.

This paper serves as a one-stop destination for anyone intending to understand the capabilities of HANA NSE and is in the quest of technical documentation collating information from various articles, studies, and whitepapers. With the Dell Technologies PowerStore solutions, the users can harness the power of HANA and HANA NSE with the option to tier statistically and place the data in warm and cold-based storage resulting in the PowerStore solution portraying a negative TCO and ending up as an investment that pays for itself over time.

7. References

SAP HANA Administration Guide :

https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/4efaa94f8057425c8c7021da6fc2ddf5.html?version=2.0.04

SAP HANA Central: <http://www.hanaexam.com/2022/05/understanding-configuration-of-sap-hana.html>

Dell Technologies – HANA NSE: <https://infohub.delltechnologies.com/l/dell-validated-design-for-sap-hana-tdi-with-dell-powerstore-storage-systems-1/sap-hana-nse-17>

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect Dell Technologies' views, processes, or methodologies. Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice. THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license. © 2023 Dell Inc. or its subsidiaries. All Rights Reserved. Dell and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.