IBM® Storage

IBM FlashSystem with SAP HANA Native Storage Extension

Version 1.0

IBM Storage Team & WW SAP Technical Enablement Team

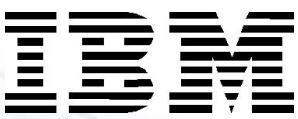

Contents

About this document

This paper is intended as an IBM solution reference for using IBM FlashSystem with SAP HANA Native Storage Extension (NSE) in an SAP HANA tailored data center integration (SAP HANA TDI) environment. SAP HANA TDI allows the SAP customer to attach external storage to the SAP HANA server.

This document is written for customers and business partners with knowledge about SAP HANA and IBM System Storage.

Scope

This document was developed using the following software tools:

- SAP BW/4HANA
- IBM Storage Insights

This technical report does not:

- Replace any official manuals and documents issued by IBM
- Explain installation and configuration of SAP HANA or SAP BW/4HANA and NSE

Prerequisites

This document refers to SAP documentation that is available for download from the SAP websites:

- SAP HANA NSE: https://www.sap.com/documents/2019/09/4475a0dd-637d-0010-87a3-c30de2ffd8ff.html

It is assumed that you are familiar with and have basic knowledge of the following products:

- IBM FlashSystems 9200
- SAP BW/4HANA
- SAP HANA Native Storage Extension (NSE)

SAP HANA Native Storage Extension (NSE)

With SAP HANA 2.0 SPS 04 the warm data tier NSE (Native Storage Extension) was introduced and made available to customers. IBM worked together with SAP to be ready to roll out NSE in SAP HANA 2.0 SPS 04.

A significant part of the SAP HANA database server cost is related to the memory (DRAM), because the more data the customer needs to have in memory, the much higher the infrastructure costs are. Depending on the license model, the in-memory footprint of HANA may also affect the SAP software license cost.

Here an NSE use case example:
A company is running an SAP HANA application, S4/HANA The initial sizing of the server was for a database of 4TB RAM. After a year of running the application, it was determined, that the database growth was larger than the RAM size could accommodate. Here NSE, in combination with a fast storage system, is the right technology to mitigate the effects of such database growth.

The SAP HANA NSE feature implements a separate buffer cache inside the allocated main memory from SAP HANA. The default size of the buffer cache is 10% of the HANA memory. It's activated by default, but will only be used, when the first data objects are NSE enabled. Data objects could be complete tables, parts of tables or only specific columns of tables. NSE enabled data objects are removed from the Column Store Buffer and reside on disk. Only small parts of NSE enabled tables are stored in the NSE buffer cache when required.

By using the optimal $/Performance IBM FlashSystem FS5100 or the high performance FS7200 or FS9200 with FlashCore Modules (FCM), Storage Class Memory (SCM) technology in combination with NSE three objectives are met:

- Significant savings of RAM capacity can be achieved without impacting performance of your SAP application.
- Better protection in critical situations where warm data is temporarily accessed in a hot manner.
- SAP HANA start up times are improved by factors compared to traditional SSDs.

The overall architecture of NSE is shown in Figure 1. NSE has a separate buffer cache inside the allocated main memory from SAP HANA. NSE enabled data objects are removed from the Column Store Buffer and only small parts of NSE enabled tables are stored in the NSE buffer cache. The rest is placed on the persistence (disk).

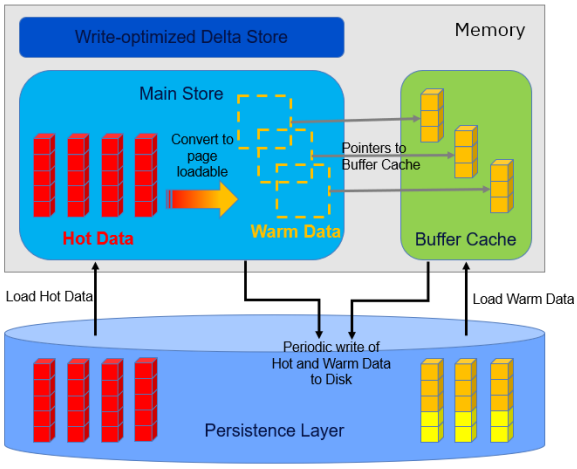

Figure 1. NSE architecture

IBM FlashSystem 9200

IBM FlashSystem helps to optimize SAP HANA environments, efficiently managing high performance data storage with SAP HANA Native Storage Extension.

The IBM FlashSystem family combines the performance of flash and end-to-end Non-Volatile Memory Express (NVMe) with the reliability and innovation of IBM FlashCore technology, the ultra-low latency of Storage Class Memory (SCM), the rich features of IBM Spectrum Virtualize and AI predictive storage management and proactive support by Storage Insights. Built in a powerful 2U enterprise-class, blazing fast storage all-flash array, as shown in Figure 2.

NVMe protocol inside FlashSystem 9200
NVM Express (NVMe) is an optimized, high-performance scalable host controller interface designed to address the needs of systems that utilize PCI Express-based solid-state storage. The NVMe protocol is an interface specification for communicating with storage devices. It is functionally analogous to other protocols, such as SAS. However, the NVMe interface was designed for extremely fast storage media, such as flash-based solid-state drives (SSDs) and low-latency non-volatile storage technologies.

NVMe storage devices are typically directly attached to a host system over a PCI Express (PCIe) bus. That is, the NVMe controller is contained in the storage device itself, alleviating the need for an additional I/O controller between the CPU and the storage device. The architecture results in lower latency, throughput scalability, and simpler system designs. NVMe protocol supports multiple I/O queues, versus legacy SAS and SATA protocols that use only a single queue.

These all-flash systems include IBM Spectrum Virtualize software and introduce remarkable new features in comparison to the predecessor models:

- End-to-end **NVMe** support: NVMe is a logical device interface standard from 2011 for accessing non-volatile storage media that is attached via a PCI Express bus.
- **Lower latencies** through RDMA: Direct memory access from the memory of one node into that of another without involving either one's operating system.
- **Data reduction pools** (DRP) represent a significant enhancement to the storage pool concept. Now with the introduction of data reduction technology, compression, and deduplication, it has become more of a requirement to have an uncomplicated way to stay “thin”.

- **FlashCore Modules** (FCMs) or industry standard NVMe drives can be used for the IBM FlashSystems. If the FCM option is chosen, then the user can take advantage of the built-in hardware compression, which will automatically try to compress the stored data when written to the drives.
- Thin-provisioned **IBM FlashCopy** uses disk space only when updates are made to the source or target data, and not for the entire capacity of a volume copy.
- **HyperSwap** capability enables each volume to be presented by two IBM FlashSystems. This high-availability configuration tolerates combinations of node and site failures, using host multipathing driver, based on the one that is available for the regular IBM FlashSystem.
- The IBM FlashSystem 9200 supports the new low latency, high speed Storage Class Memory (SCM). SCM is a non-volatile memory device that performs faster (~10µs) than traditional NAND SSDs(100µs), but slower than DRAM (100ns).
- ***IBM Storage Insights is an additional part of the monitoring capability of the IBM FlashSystem 9200 system and supplements the views available in the GUI.***

Figure 2: IBM FlashSystem 9200 control enclosure

For more information about the IBM FlashSystem family see the following resources:

- IBM FlashSystem 9200 and 9100 Best Practices and Performance Guidelines: http://www.redbooks.ibm.com/abstracts/sg248448.html?Open

- SAP HANA certified and supported IBM storage systems https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/enterprise-storage.html#categories=certified%23International%20Business%20Machines%20Corporation

IBM Storage Insights

IBM Storage Insights is offered free of charge to customers who own IBM block storage systems. It is a secured IBM Cloud storage service that monitors IBM block storage and provides advanced functionality for alerting.

It provides single-pane views of IBM block storage systems, such as the Operations dashboard and the Notifications dashboard. With the information that is provided, such as the diagnostic event information, key capacity and performance information, and the streamlined support experience, you can quickly assess the health of your storage environment and get help with resolving issues. And, on the Advisor page, IBM Storage Insights provides recommendations on the remedial steps that can be taken to manage risks and resolve issues that might impact your storage services.

All IBM FlashSystem performance graphs in this paper where created with IBM Storage Insights.

For more information about IBM Storage Insights see the following resources:

- https://www.ibm.com/products/analytics-driven-data-management

IBM SAP BW/4HANA Setup

For the testing covered by this document the following setup was used, as shown in Figure 3 . The hardware setup consists of:

- IBM FlashSystem 9200 with:
  - 1563 GB System Memory per I/O Group
  - 12 x 9.6 TB (22 TB) 2.5” NVMe FCM (FlashCore Modules)
  - 1 x 16Gb FC 4 Port Adapter Pair

- IBM Power System E980
- 16Gb SAN infrastructure

Two LPARs were used for the SAP BW/4HANA and the SAP HANA database. The LPARs were attached to IBM FlashSystem 9200 volumes from a DRAID 6 pool. This pool was made of 12 FCM drives.

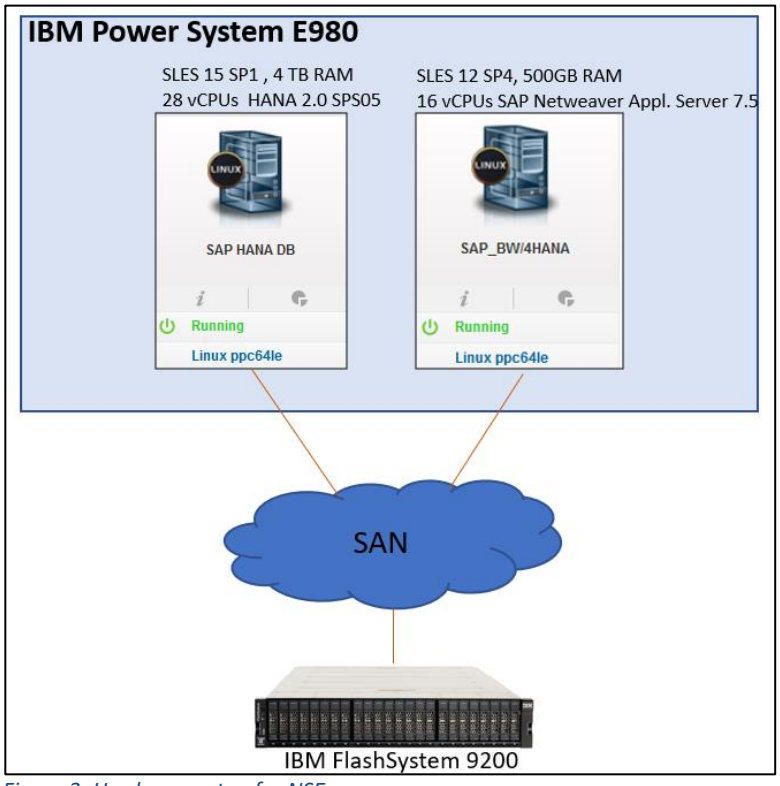

Figure 3: Hardware setup for NSE

SAP HANA database server storage configuration

In the test environment, the filesystems for storing SAP HANA data volumes and SAP HANA transaction logs are located on LVM logical volumes. For log and for data dedicated LVM volume groups exists. The LVM physical volumes of these volume groups are provided by IBM FlashSystem 9200.
We only measured the IO behavior of the disks from the “data” volume group – SAP HANA log IO is out of scope as it is not impacted by NSE.

Here is a quick overview of the structure from the /hana/data filesystem down to the corresponding IBM FlashSystem VDisks:

1) Filesystem on LVM logical volume /dev/hn_dt_vg/hn_dt_lv mounted as /hana/data

2) This logical volume consumes 100% space of LVM volume group “hn_dt_vg”
3) The following multipath devices are used as LVM physical volumes, as shown in Figure 4:

```
lsh40042:~ # pvs
  PV                        VG       Fmt  Attr PSize   PFree
  /dev/mapper/mpatha-part2  rootvg   lvm2 a--  255.99g 198.99g
  /dev/mapper/mpathb        datavg   lvm2 a--  256.00g      0
  /dev/mapper/mpathc        hn_dt_vg lvm2 a--  750.00g      0
  /dev/mapper/mpathd        hn_lg_vg lvm2 a--   64.00g      0
  /dev/mapper/mpathe        hn_lg_vg lvm2 a--   64.00g      0
  /dev/mapper/mpathf        hn_lg_vg lvm2 a--   64.00g      0
  /dev/mapper/mpathg        hn_lg_vg lvm2 a--   64.00g      0
  /dev/mapper/mpathh        hn_lg_vg lvm2 a--   64.00g      0
  /dev/mapper/mpathi        hn_lg_vg lvm2 a--   64.00g      0
  /dev/mapper/mpathj        hn_lg_vg lvm2 a--   64.00g      0
  /dev/mapper/mpathl        hn_dt_vg lvm2 a--  750.00g      0
  /dev/mapper/mpathm        hn_dt_vg lvm2 a--  750.00g      0
  /dev/mapper/mpathn        hn_dt_vg lvm2 a--  750.00g      0
  /dev/mapper/mpatho        hn_dt_vg lvm2 a--  750.00g      0
  /dev/mapper/mpathp        hn_dt_vg lvm2 a--  750.00g      0
  /dev/mapper/mpathq        hn_dt_vg lvm2 a--  750.00g      0
  /dev/mapper/mpathr        hn_dt_vg lvm2 a--  750.00g      0
  /dev/mapper/mpaths        hn_lg_vg lvm2 a--   64.00g      0
```

Figure 4: LVM Physical Volumes

4) The “multipath -ll” command lists the UUID of these devices, as shown in Figure 5:

```
lsh40042:~ # multipath -ll | grep mpath
mpathr (3600507681081019a1000000000000019) dm-16 IBM,2145
mpathe (3600507681081019a100000000000001c) dm-3 IBM,2145
mpathq (3600507681081019a1000000000000018) dm-15 IBM,2145
mpathd (3600507681081019a100000000000001b) dm-2 IBM,2145
mpathp (3600507681081019a1000000000000017) dm-14 IBM,2145
mpathc (3600507681081019a1000000000000012) dm-1 IBM,2145
mpatho (3600507681081019a1000000000000016) dm-13 IBM,2145
mpathb (3600507681081019a100000000000000f) dm-0 IBM,2145
mpathn (3600507681081019a1000000000000015) dm-12 IBM,2145
mpatha (3600507680c8086498000000000000084e) dm-18 IBM,2145
mpathm (3600507681081019a1000000000000014) dm-11 IBM,2145
mpathl (3600507681081019a1000000000000013) dm-10 IBM,2145
mpathk (3600507681081019a1000000000000022) dm-9 IBM,2145
mpathj (3600507681081019a1000000000000021) dm-8 IBM,2145
mpathi (3600507681081019a1000000000000020) dm-7 IBM,2145
mpathh (3600507681081019a100000000000001f) dm-6 IBM,2145
mpathg (3600507681081019a100000000000001e) dm-5 IBM,2145
mpaths (3600507681081019a100000000000001a) dm-17 IBM,2145
mpathf (3600507681081019a100000000000001d) dm-4 IBM,2145
```

Figure 5: Multipath devices

5) The corresponding data VDisks are listed in the IBM FlashSystem Web GUI, as shown in Figure 6.

Name	UID	Capacity
V-lsh40042-ID00	600507681081019A100000000000000F	256.00
V-lsh40042-ID01	600507681081019A1000000000000012	750.00
V-lsh40042-ID02	600507681081019A1000000000000013	750.00
V-lsh40042-ID03	600507681081019A1000000000000014	750.00
V-lsh40042-ID04	600507681081019A1000000000000015	750.00
V-lsh40042-ID05	600507681081019A1000000000000016	750.00
V-lsh40042-ID06	600507681081019A1000000000000017	750.00
V-lsh40042-ID07	600507681081019A1000000000000018	750.00
V-lsh40042-ID08	600507681081019A1000000000000019	750.00

Figure 6: List of IBM FlashSystem mapped VDisks

NSE and IBM FlashSystem configuration

NSE is enabled by default for all SAP HANA versions >= SAP HANA 2.0 SPS 04. Per default all database tables are loaded into the column store. Data objects could be NSE enabled/disabled based on the complete table, only a partition of a table or some columns of a table.

Workload

In this test setup we basically run two different kinds of database workload:

- Multiple test runs based on complex queries
- Multiple massive test runs based on simple queries

NSE testing configuration

We ran the described workload against three different NSE configurations. For all performance test we measured the following points:

- Number of random queries across all database tables in a given time frame
- Storage utilization captured by IBM Storage Insights.
- Comparison for different NSE configurations:
  1) No NSE enabled. All HANA database tables are located in the main memory, which is the default for SAP HANA 2.0 SPS 05.
  2) A large data object (the greatest one) is NSE enabled.
  3) All data objects recommended by the application owner are NSE enabled

The expectation is to measure significant increasing storage IO impact from NSE configuration 1 to NSE configuration 3 without impacting the overall performance of the BW workload.

Results:

We'll start with a comparison of the overall number of database queries during the test runs. In all three test configurations, we measured nearby identical values. This means that using the NSE feature in recommended configurations has no impact on the database performance but allows customers to run SAP HANA using a significant lower memory footprint. In our configuration 3, for example, we used 25% less RAM compared to configuration 1.

Table 1 is showing a comparison of the NSE performance impact on the IBM FlashSystem 9200, depending on the number of tables being NSE enabled.

Table 1: Queries performance impact on IBM FlashSystem 9200

	NSE enabled tables			
	no	minimal	recommended	all tables NSE enabled
complex queries	no measurable impact	no measurable impact	no measurable impact	~5% node utilization on FlashSystem 9200
simple queries	no measurable impact	no measurable impact	no measurable impact	~3% node utilization on FlashSystem 9200

In Figure 7, we illustrate that enabling NSE doesn't impact the performance of the SAP HANA database in each of the tested configurations.

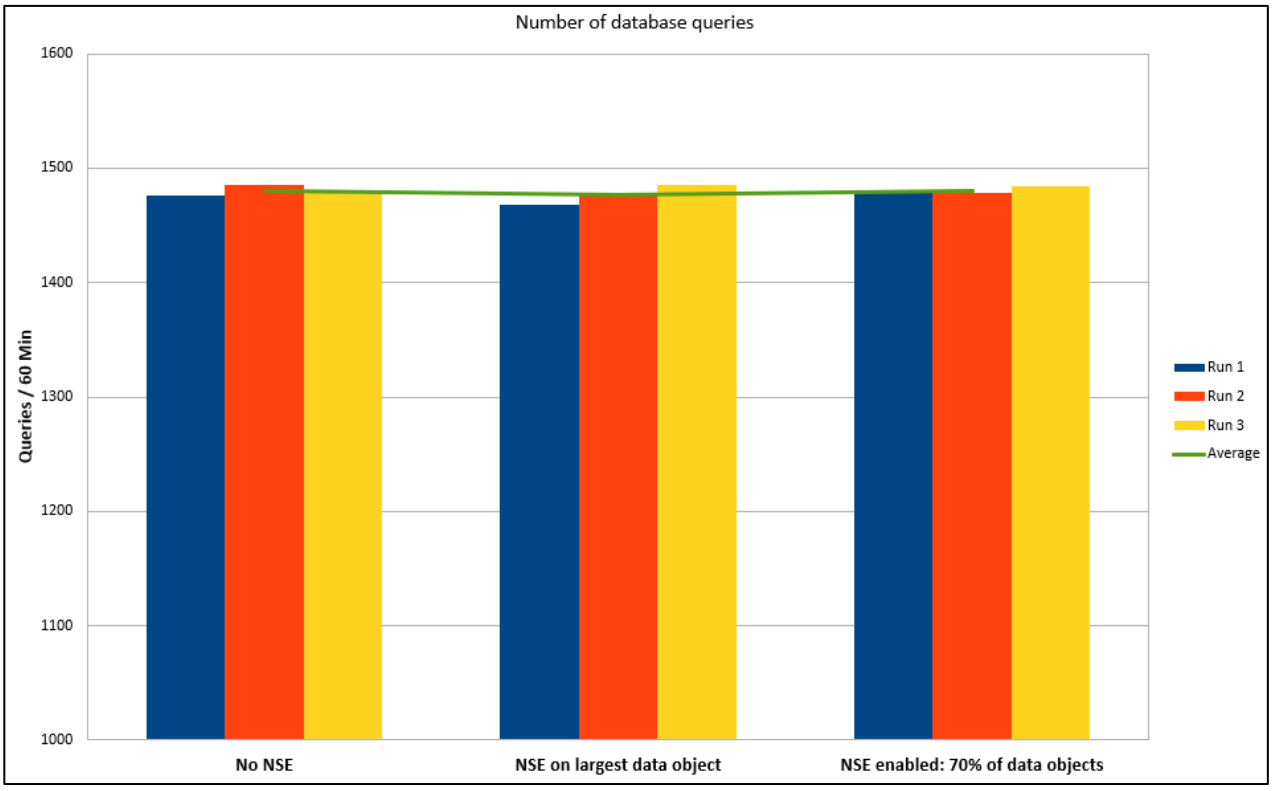

Figure 7: Comparison of different NSE enabled tables

The more tables are NSE enabled, the more IO operations need to be performed by the database host. To avoid misconfigurations, please refer to all relevant SAP notes regarding host and database configuration.

We will demonstrate now, how the use of IBM FlashSystem helps the customer not to get additional performance impact in the data path.
IBM FlashSystem is the recommended storage solution for SAP HANA workloads including NSE.
The IBM® System Storage™ Architecture and Configuration Guide for SAP® HANA™ Tailored Datacenter Integration document provides all information required to set up IBM FlashSystem for SAP HANA in the best way.

Most important factor, however, is the overall performance of the storage subsystem. IBM FlashSystem is a perfect choice for getting best performance in combination with SAP HANA NSE.

For more detailed information about the tested NSE configurations, we will show technical diagrams of some key performance values in the following section.

From Standard to recommended use of NSE

In this test area the following configurations are executed:

a. Run without use of the NSE feature
b. Only the biggest recommended table was NSE enabled. This table takes nearly 25% size of the complete database
c. All recommended table are NSE enabled. This tables takes nearly 70% size of the hole database.

All 3 test scenarios have nearly the same impact on the system performance and the test results. This is due to the fact, that the selection of tables for NSE is optimal and couldn't be better. The read IO diagram shows only minimal read operation to the disk subsystems. Although the CPU utilization of the database host is absolute ideal. The following graphs are representing the host perspective of the test with all recommended tables NSE enabled. Figure 8 and Figure 9 were created with nmon.

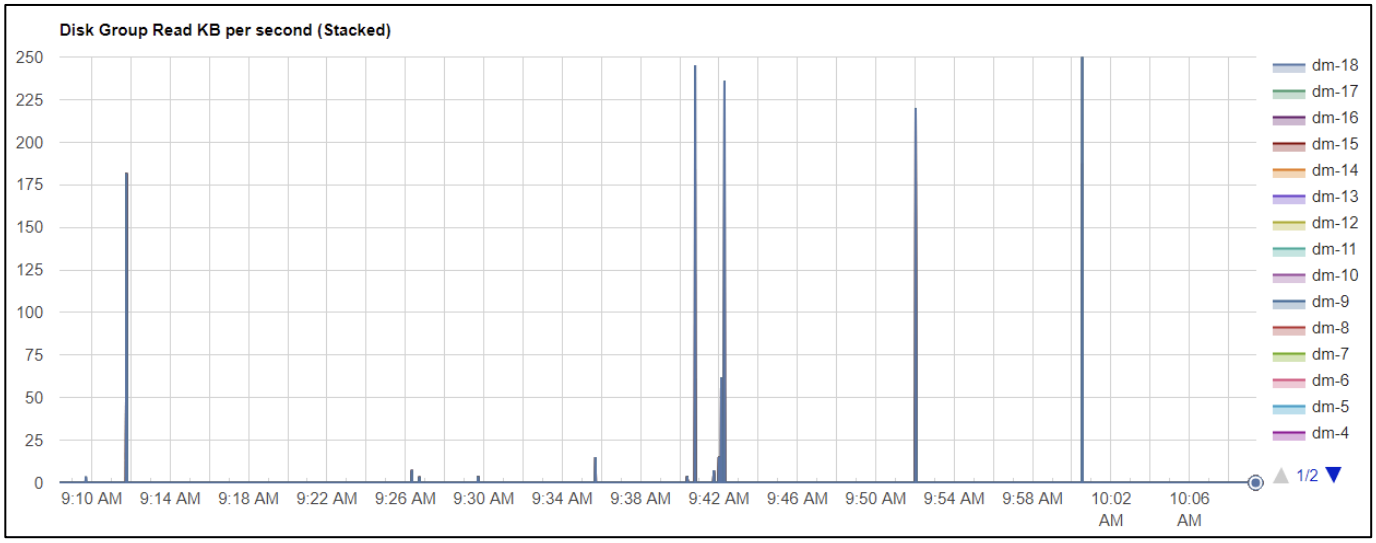

Figure 8: Read IO during a run with all recommended tables NSE enabled

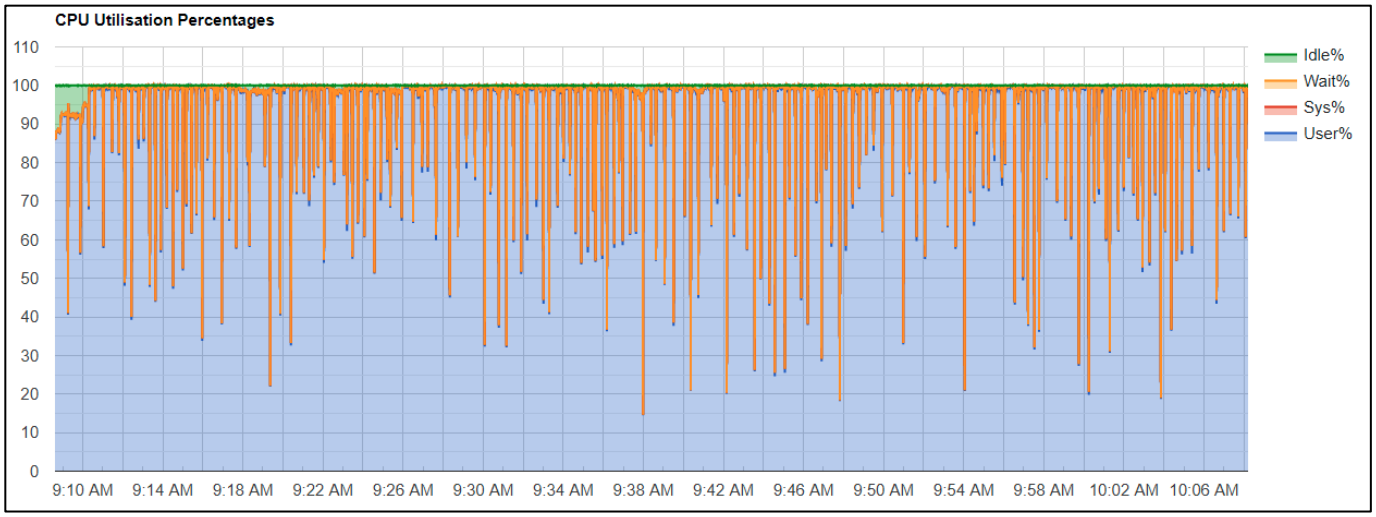

Figure 9: CPU Utilization during a test with all recommended tables NSE enabled

For the third test all tables, recommended by the SAP HANA NSE advisor tool, were NSE enabled. This means that all tables that contain warm data are moved out of the memory and are placed on IBM FlashSystem storage.

The SAP HANA NSE feature allows massive memory capacity savings when it is used in this way:
With nearly around 70% of data now being NSE enabled, the overall RAM memory savings are 25%!
Consider, that 1 TB of NVMe storage capacity is about 10 times less expensive than 1 TB of RAM in a current Power server.
As already shown in Figure 8 the read data rate from the host point of view is almost zero. For comparison we show in Figure 10 the read data rate of the data volumes from the storage system point of view. As expected, there is no visible data rate to show.

Figure 10: Read Data Rate of Data Volumes (NSE advisor recommended tables)

Worst case scenario

In addition to the tested configurations we decided to run another scenario, where all HANA tables are NSE enabled. This is of course a neither recommended nor supported NSE configuration. The goal of this scenario

is to simulate a situation when former warm data becomes hot at a sudden, for example the reinvestigation of older business data.

In this scenario the query rate decreases significantly – which is expected as all the database objects needs now get loaded into the NSE buffer cache and SAP HANA does not longer benefit from persistent data in memory. To load the data as fast as possible the storage system must not become a bottleneck. The indicators here are the read data rate of the volumes. The Read IO increases significant, shown the Figure 11. The CPU Utilization will decrease slightly, which is an indicator for the massive CPU bound workload in the recommended NSE setups.

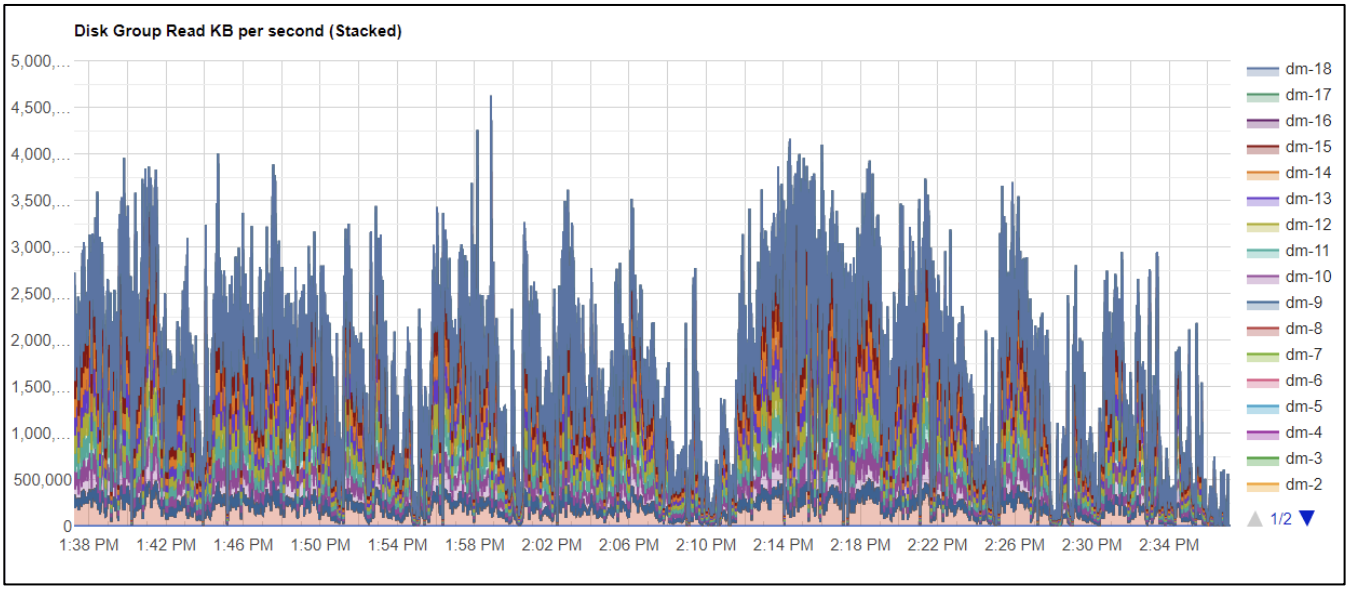

Figure 11: Read IO diagram from the host part at the worst case scenario

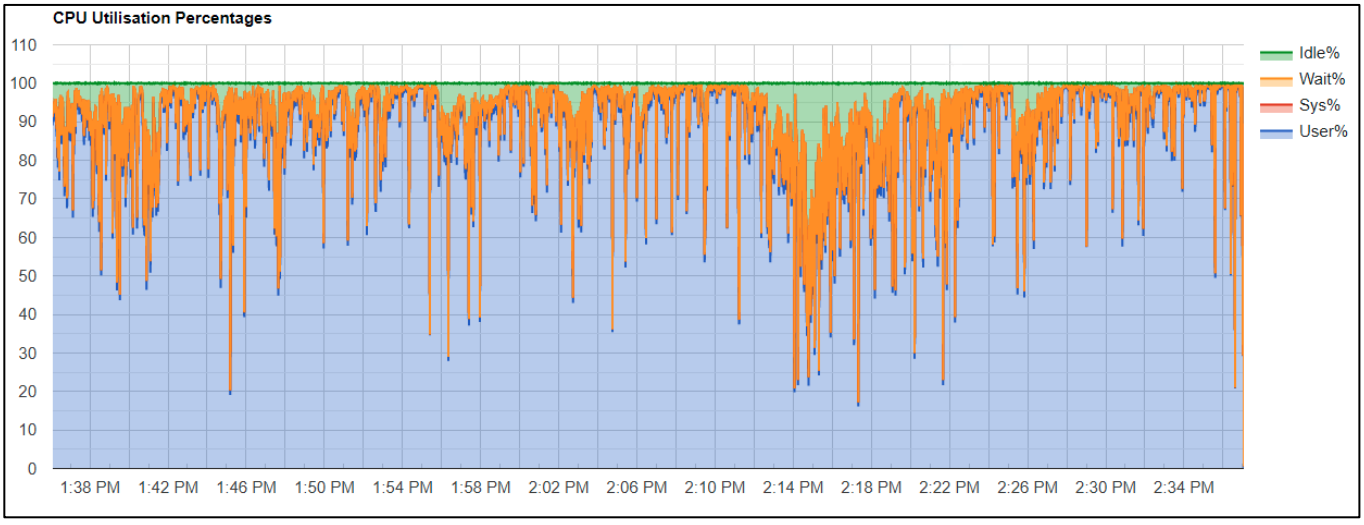

Figure 12: Host CPU Utilization in the worst case scenario

Even in this scenario we still see an excellent performance of the IBM FlashSystem 9200 with random queries run, as shown for the “Read Data Rate” in Figure 13 , the “Read Response time” in Figure 14 and finally the “Node CPU Utilization Rate” in Figure 15. “Read Data Rate” and “Read Response Time” has been measured for the HANA data volumes only.

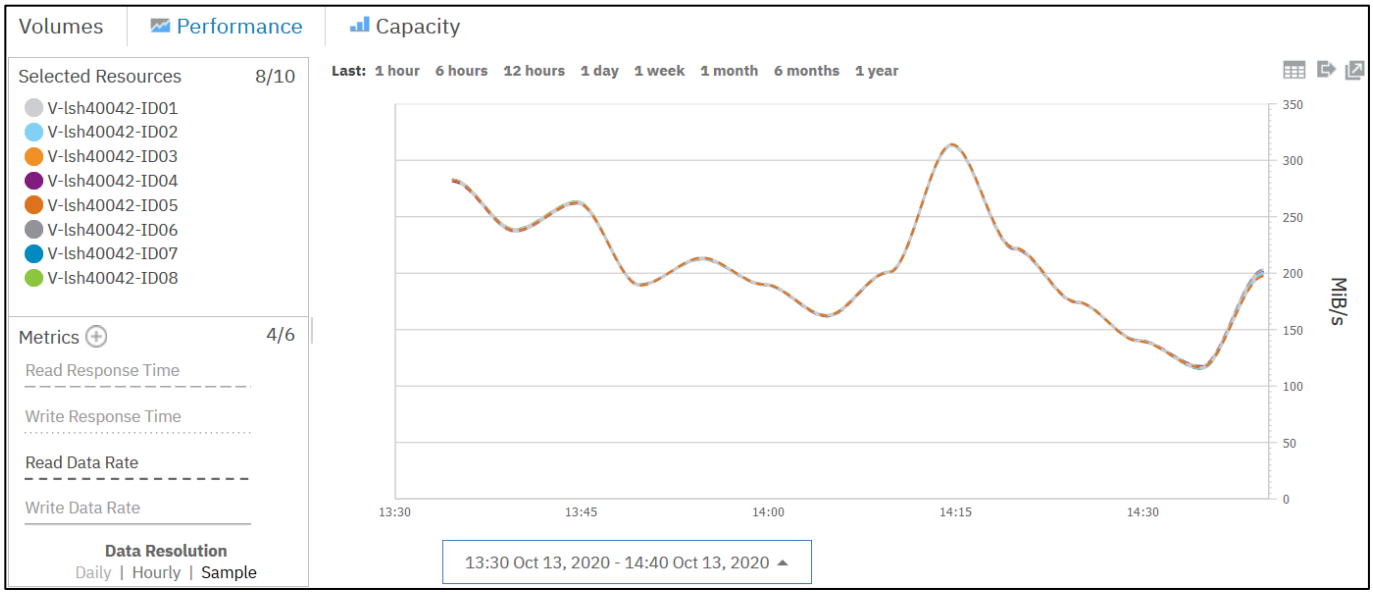

Figure 13: Read Data Rate when all tables are NSE enabled

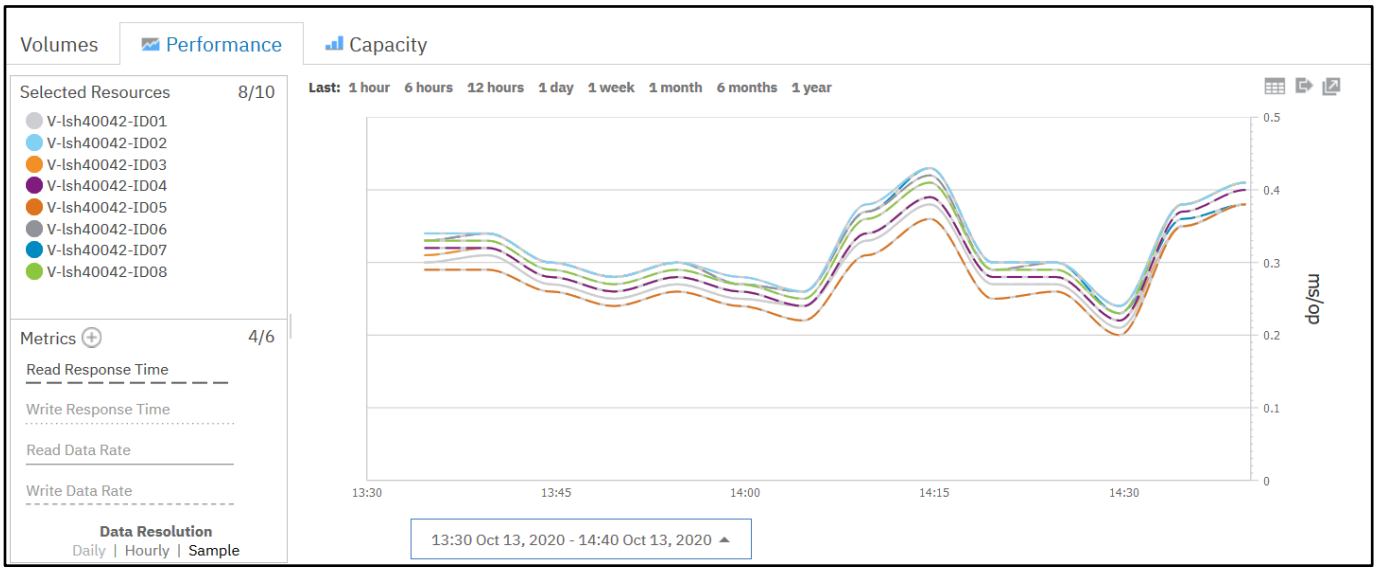

Figure 14: Read response time when all tables are NSE enabled

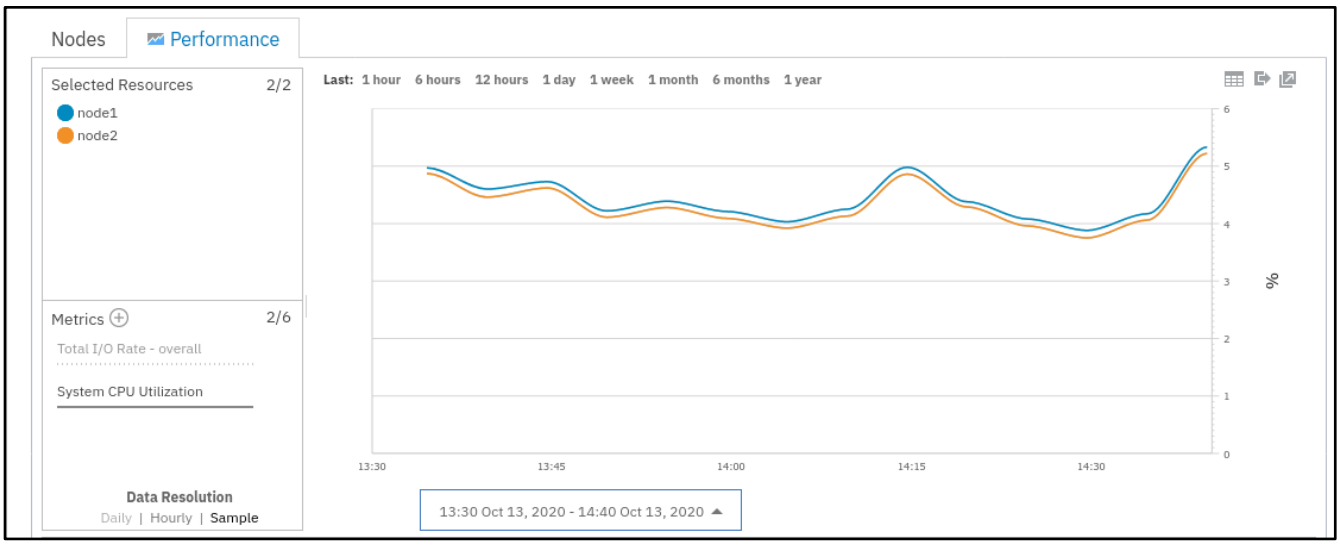

Figure 15: Node CPU Utilization

The “Read Data Rate” is an indicator for the required bandwidth – we can state that this isn’t the maximum bandwidth of the system.
The most interesting indicator however is the Read response time – an overloaded system is no longer able to respond to a read request in less than 1 ms.
In fact, the peak load was around 40.000 IO/S with an average IO Size of 65Kb. This is 12% of the IBM FlashSystem 9200 capacity as configured for our tests. The system could run the same workload nine times.
The tested workload could also be handled by an IBM FlashSystem 5100 – it would be 25% busy.

Summary

With IBM FlashSystem in combination with SAP HANA NSE, it can be shown that there is no performance impact when enabling the NSE feature on the database tables, recommended by the SAP HANA advisor tool.
The IBM FlashSystem could cover the additional NSE load without any impact. The measured values prove, that IBM FlashSystem 9200 could run a lot of additional workload in combination with NSE workload.
In the additionally tested “worst case” scenario, we could prove that the IBM FlashSystem 9200 wasn’t stressed in bandwidth or response times.
By using the recommended SAP HANA NSE settings 25% of the former required DRAM capacity can be saved. Furthermore, in combination with IBM FlashSystem with FCM technology the database keeps up and running even in unexpected NSE situations.

Notices

reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®
DB2®
Easy Tier®
FlashCopy®
IBM®
IBM FlashCore®
IBM FlashSystem®
IBM Spectrum™
IBM Spectrum Protect™
IBM Spectrum Storage™
IBM Spectrum Virtualize™
POWER®
Power Systems™
POWER9™
PowerLinux™
Redbooks (logo) ®
Storwize®
System Storage®
Tivoli®

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

October 2020

Please recycle