

Hitachi Ops Center Analyzer Best Practices for SAP HANA Infrastructure

Best Practices Guide

© 2024. Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara LLC (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, DB2, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, OS/390, PowerHA, PowerPC, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z15, z16, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, Microsoft Edge, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found in the product documentation, at <https://www.hitachivantara.com/en-us/company/legal.html>.

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

Changes	Date
Updated hyperlinks	April 2024
Initial release	May 2021

Best Practices Guide

A heterogeneous infrastructure for multiple SAP HANA systems is difficult to monitor and analyze the usage level for the whole infrastructure stack and its individual parts. Hitachi Ops Center provides a unified single solution that monitors the entire infrastructure stack of your SAP HANA landscapes and offers troubleshooting and problem analysis capabilities.

This best practices paper describes how to monitor the infrastructure of multiple SAP HANA systems consolidated on a Hitachi enterprise storage system under the umbrella of [SAP HANA Tailored Data Center integration](#) (TDI). This paper also includes guidelines for using Hitachi Ops Center Analyzer to monitor your SAP HANA TDI environment, setting thresholds and alerts to get the best out of your consolidated SAP HANA landscape, and ensuring the reliability and availability of your SAP HANA systems and infrastructure.

With an SAP HANA appliance deployment, the hardware vendor preconfigures all hardware components including storage and server. With Hitachi storage systems for an SAP HANA TDI deployment, you can customize each installation by using any supported servers and storage from a certified enterprise storage vendor. This means that you might be able to leverage your existing hardware to reduce your total cost of ownership (TCO) for a high-performance SAP HANA infrastructure.

Follow these directions and guidelines for setup and configuration when implementing an SAP HANA TDI infrastructure on Hitachi Virtual Storage Platform. For more information, see [SAP HANA Tailored Data Center Integration - Frequently Asked Questions](#).

SAP HANA Tailored Datacenter Integration

SAP increases flexibility and provides an alternative to SAP HANA Appliances with SAP HANA Tailored Data Center Integration, which consists of five phases. This solution uses various types of virtualization technology. Understanding the possibilities and requirements of SAP HANA TDI environments is crucial. One of these requirements is that all virtualized environments for SAP HANA are considered SAP HANA TDI by SAP. Every VM counts as one HANA node when counting the supported number of SAP HANA nodes connected to a single enterprise storage system. SAP provides the following documentation that explains the phases of SAP HANA TDI as well as the hardware and software requirements for the entire stack:

- [Overview Slide Deck](#)
- [Frequently Asked Questions](#)
- [SAP HANA Storage Requirements whitepaper](#)
- [SAP HANA Network Requirements whitepaper](#)

Taking this into account, Hitachi Vantara offers a solution for the SAP HANA TDI stack from hardware infrastructure to software deployment.

SAP HANA certified enterprise storage

SAP has introduced the certification of enterprise storage systems to enable consolidation of multiple HANA systems on a single enterprise storage system. All certified storage can be found in the [Certified and Supported SAP HANA Hardware Directory](#). Each entry lists the supported maximum number of production SAP HANA nodes. SAP HANA nodes can be either bare metal servers certified as SAP HANA Appliances, SAP HANA Intel Systems, or virtual machines running on a certified hypervisor.

The following table shows the supported maximum number of SAP HANA systems connected in parallel to a single Hitachi Virtual Storage Platform. While running storage key performance indicator (KPI) testing using the SAP HANA Hardware and Cloud Measurement Tool (HCMT), the storage was fully dedicated to the testing and not shared by other systems or applications.



Note: If you must have mixed applications running in parallel on the same storage system, SAP HANA and non-SAP HANA, you must test for the expected workload using the HCMT to meet SAP KPIs.

Storage	Drive Type	Maximum SAP HANA Systems	Minimum Parity Group	
			Data	Log
VSP G350 and VSP G370	SAS HDDs	3	3 RAID-6 (14D +2P)	1 RAID-6 (6D +2P)
VSP G350, VSP F350, VSP G370, and VSP F370	SAS SSDs	16	3 RAID-10 (2D +2D)	3 RAID-10 (2D +2D)
VSP G700	SAS HDDs	13	6 RAID-6 (14D +2P)	3 RAID-6 (6D +2P)
VSP G700 and VSP F700	SAS SSDs	34*	4 RAID-10 (2D +2D)	5 RAID-10 (2D +2D)
VSP G900 and VSP F900	SAS SSDs	40*	4 RAID-10 (2D +2D)	4 RAID10 (2D +2D)
VSP E590	NVMe SSDs	19	1 RAID-5 (6D +1P)	1 RAID-5 (6D +1P)
VSP E790	NVMe SSDs	34*	3 RAID-10 (2D +2D)	2 RAID-10 (2D +2D)
VSP E990	NVMe SSDs	50	7 RAID-10 (2D +2D)	7 RAID-10 (2D +2D)
VSP 5100, VSP 5100H	SAS SSDs	37	7 RAID-10 (2D +2D)	7 RAID-10 (2D +2D)

Storage	Drive Type	Maximum SAP HANA Systems	Minimum Parity Group	
			Data	Log
VSP 5500 or VSP 5500H (1 pair of nodes)	SAS SSDs	74	14 RAID-10 (2D +2D)	14 RAID-10 (2D +2D)
VSP 5500, VSP 5500H (2 pairs of nodes)	SAS SSDs	148	28 RAID-10 (2D +2D)	28 RAID-10 (2D +2D)
VSP 5500, VSP 5500H (3 pairs of nodes)	SAS SSDs	222	42 RAID-10 (2D +2D)	42 RAID-10 (2D +2D)
*Extrapolated scalability for maximum number of nodes using engineering validation data and information.				

When implementing an SAP HANA TDI infrastructure on Hitachi Virtual Storage Platform using NVMe SSDs, you do not have to use the exact storage design used in this best practices guide for enterprise storage KPI testing. However, the recommendation is to follow these guidelines. See [SAP HANA Tailored Data Center Integration - Frequently Asked Questions](#) for more details.



Note: Since the release of SAP HANA TDI in November 2013, several versions of HCMT have been published. To verify that the hardware configuration of your SAP HANA TDI infrastructure meets the SAP KPIs, you must use the same version of HCMT used during the certification of the hardware, compute servers, and storage system for your tests. [SAP Note 2493172](#) describes how to determine the right version of HCMT for your tests.

Storage sizing for TDI storage consolidation

If the SAP quicksizer and sizing reports are available, this information can influence proper storage sizing for HANA data and log volumes. If not, storage sizing is done using the minimum volume sizes for the operating system, HANA data, log, and shared volumes using the formulas listed in [SAP HANA Storage Requirements](#).

Hitachi recommends these minimum sizes, as calculated using the following formulas:

- Every HANA node requires approximately 100 GB capacity for the operating system.
- HANA data volume size requires at least the following: 1 × RAM of each HANA node

- HANA log volume size uses the following formulas:
 - Systems with equal or less than 512 GB memory: $1/2 \times \text{RAM}$
 - Systems with greater than 512 GB memory at least: 512 GB
- /hana/shared size uses formulas:
 - Single node (scale-up): $\text{MIN}(1 \times \text{RAM}; 1 \text{ TB})$
 - Multi-node (scale-out): $1 \times \text{RAM}$ of every 4 worker nodes

Servers for SAP HANA TDI environments

Only compute servers certified for use with SAP HANA and listed in the official SAP HANA Hardware Directory (HWD) are supported for use in an SAP HANA TDI environment. This follows the same bill of materials as the certified SAP HANA appliance server, but without the storage or local disks.

For servers certified as appliance servers with a certificate provided by SAP for a specific operating system, you can only use that operating system for SAP HANA in that implementation. This is different for the supported Intel systems listed in the HWD. For these servers, there is no separate certification for different operating system versions and service pack stacks. Instead, the server vendor's operating system certifications show which servers are supported for these TDI-only servers.

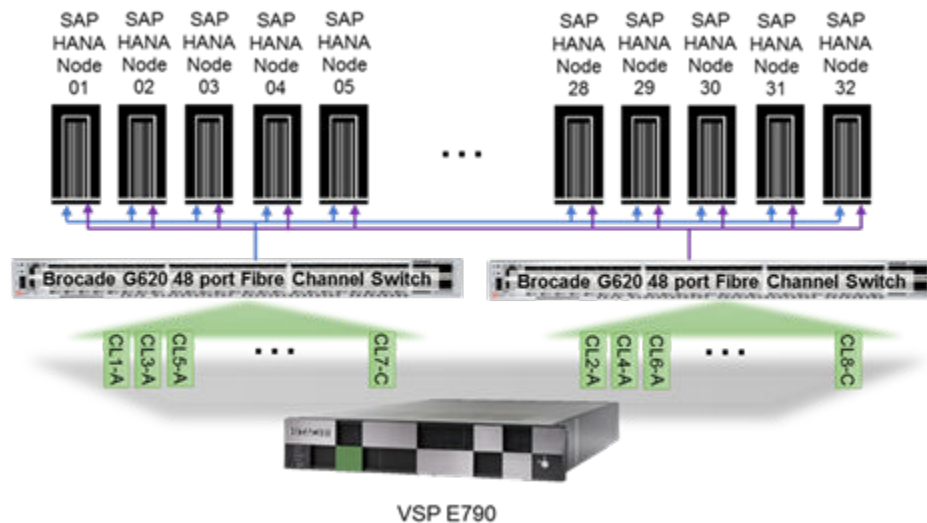
For a description of all certified servers and enterprise storage solutions, see [Certified and Supported SAP HANA Hardware](#) for. For more information on SAP HANA TDI, see [SAP HANA Tailored Data Center Integration - Frequently Asked Questions](#).

Solution components

This section describes an example configuration of an SAP HANA TDI consolidation landscape that is monitored using Hitachi Ops Center Analyzer. For detailed descriptions of the components in this configuration, see [Product descriptions \(on page 26\)](#).

Hardware components

An overview of the hardware used for the validation testing in the Hitachi Vantara lab environment is shown in the following figure. A VSP E790 was connected to 32 SAP HANA nodes using a pair of 48-port Brocade G620 Fibre Channel switches.



For details about the respective Hitachi Virtual Storage Platform models, see the following reference architecture guides:

- [SAP HANA TDI on VSP G Series and VSP F Series with Hitachi Storage Virtualization Operating System Reference Architecture](#)
- [SAP HANA Tailored Data Center Integration on Hitachi VSP E590 and VSP E790 with Hitachi Storage Virtualization Operating System RF](#)
- [SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform 5000 Series with Hitachi Storage Virtualization Operating System RF](#)

Software components

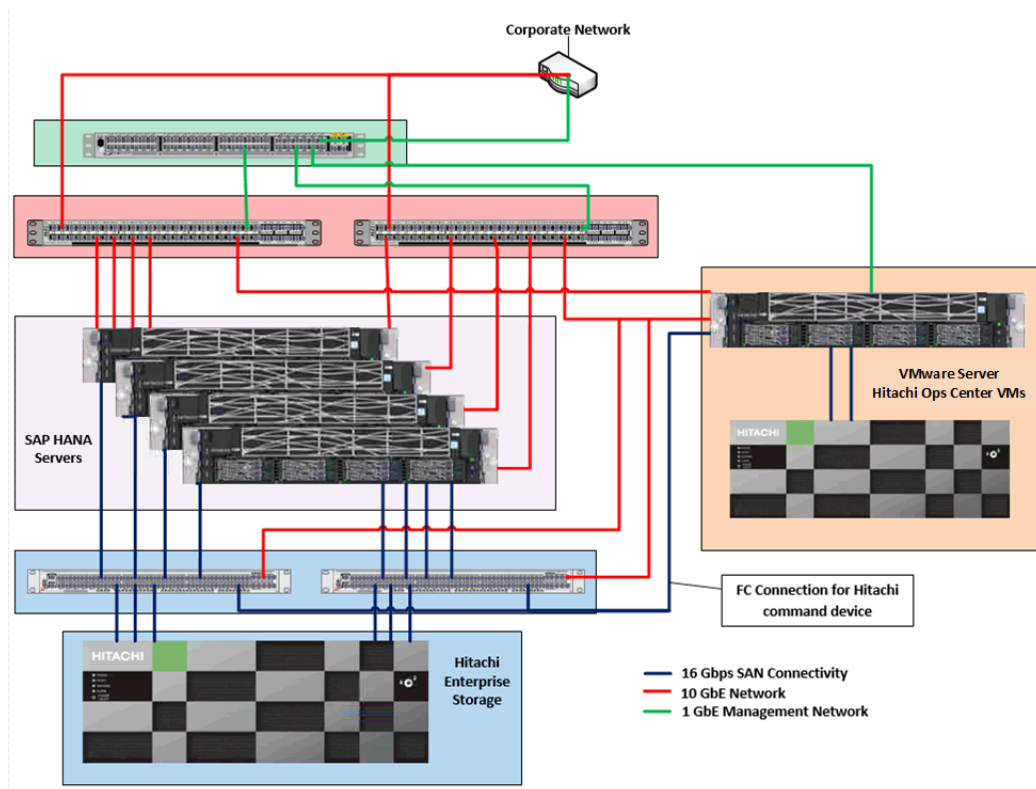
The following software was used in this environment. To ensure that you run certified and supported software versions in your solution, see the respective support matrices and certification listings.

- Hitachi Ops Center
- Hitachi Storage Virtualization Operating System RF
- SAP HANA
- Operating System Choices for the SAP HANA Servers:
 - SUSE Linux Enterprise Server (SLES) for SAP Applications
 - Red Hat Enterprise Linux (RHEL) for SAP HANA

Setup and implementation

This section summarizes the setup for the TDI storage consolidation landscape and provides insight into the configuration of Hitachi Ops Center.

The following figure shows an overview of the landscape including the VMware server and its storage, and all network and Fibre Channel connections.



Hardware setup

Follow the Hitachi Vantara best practices for your Hitachi enterprise storage system to set up the volumes for the SAP HANA nodes. The Hitachi Ops Center system requirements depend on the number of monitored resources and Hitachi Ops Center products that are installed. See *Hitachi Ops Center Requirements* in the [Product Documentation portal](#) for your version of Hitachi Ops Center.

Storage architecture

Configure your Hitachi enterprise storage system according to the best practices for each model described in the respective reference architecture guides. In general, each SAP HANA node needs the following storage layout:

- Operating system (OS) volume
- SAP HANA shared volume
- SAP HANA log volume
- SAP HANA data volume

For this, Hitachi best practices recommend using the following two dynamic provisioning pools created with Hitachi Dynamic Provisioning (HDP). This ensures maximum utilization and optimization at a lower cost than other solutions.

- OS_SH_DT_Pool for the following:
 - OS volume
 - SAP HANA shared volume
 - SAP HANA data volume
- LG_Pool for the following:
 - SAP HANA log volume

Provision the storage for each SAP HANA system using these pools. Determine the volume sizes according to the TDI sizing formulas and adapt them to additional requirements where applicable.

SAP HANA servers

Depending on the number of HANA servers used in a TDI landscape and the maximum supported node count of the storage model, the servers can either be direct attached storage (DAS) or connected using Fibre Channel storage area network (SAN) switches.

Ops Center server

For this environment, Hitachi Ops Center and its components are installed in a virtualized environment running VMware ESXi servers on a Hitachi certified server using the latest version of the Hitachi Ops Center OVA files.

For this setup, Advanced Server DS7040 servers running VMware ESXi were connected to a separate storage system, so the enterprise storage could be used exclusively for the SAP HANA nodes.

The ESXi server and the virtual machines running on it require the following physical network connection to all components of the storage consolidation setup that will be monitored:

- Ethernet connection to the SAP HANA server operating systems
- Ethernet connection to the SAP SAN switches
- Fibre Channel connection to the SAP HANA enterprise storage for the Hitachi command device

Optionally, additional management servers can be used as jump servers to access an infrastructure that is only accessible from these servers. A Hitachi Ops Center probe server must run on the ESXi server. In addition to the physical network connections, a network connection to the main Hitachi Ops Center server is required.

VMware ESXi installation and setup

The VMware ESXi installation and setup has no special needs when running an up-to-date version of ESXi that is supported for the hardware on which it is running. However, it should meet all the requirements for the Ops Center VMs it is running. See *Hitachi Ops Center Requirements* in the [Product Documentation portal](#) to verify that enough resources are available for the different Ops Center products running on the ESXi host.

To run all the VMs on the same ESXi host, a Fibre Channel storage connection to the Hitachi enterprise storage system is required for the command device. To enable using this command device, the virtual machine running the Ops Center probe server needs the command device LDEV mapped as a raw device before adding the storage probe.

SAP HANA node setup

Configure all SAP HANA nodes in this TDI landscape according to the following sections.

SAN operating system configuration

A scale-out configuration for SAP HANA requires SAN boot.

The operating system LUN is the primary boot device for each scale-out node using the 16 Gbps 2-port Fibre Channel PCIe cards. The operating system LUN holds partitions for the following:

- Operating system
- /usr/sap/ directory
- Linux swap space

Activate Device-mapper Multipath

This solution uses Device-mapper Multipath, a native component of the Linux operating system.

Using Device-mapper Multipath enables the configuration of multiple I/O paths between Hitachi Advanced Server DS7040 and Hitachi Virtual Storage Platform G700 or Virtual Storage Platform F700. Multipathing aggregates all physical I/O paths into a single logical path. The LUNs are always available unless all paths fail.

Use Device-mapper Multipath for the following I/O paths:

- Operating system volume LUN
- SAP HANA shared volume LUN
- SAP HANA data volume LUN
- SAP HANA log volume LUN

SAP HANA data volume configuration

The logical volume manager creates a single striped volume on which you create the XFS file system to store the SAP HANA data volume. The striped volume acts as the persistent layer for the SAP HANA server.

For each scale-out active SAP HANA node, create a single striped volume across the data LUNs. Format the volume with the XFS file system to store the SAP HANA data volume.

SAP HANA log volumes configuration

Logical Volume Manager creates a single striped volume on which you create the XFS file system to store the SAP HANA log volume.

For each scale-out active HANA node, create a single striped volume across the log LUNs. Format the volume with the XFS file system.

SAP HANA software installation

After configuring the file system for the SAP HANA data and log volumes, install the following SAP HANA software components on the SAP HANA nodes:

- SAP HANA database
- SAP HANA client
- SAP Host agent

Hitachi Ops Center implementation

This section summarizes the installation and setup of Hitachi Ops Center.

Install Ops Center

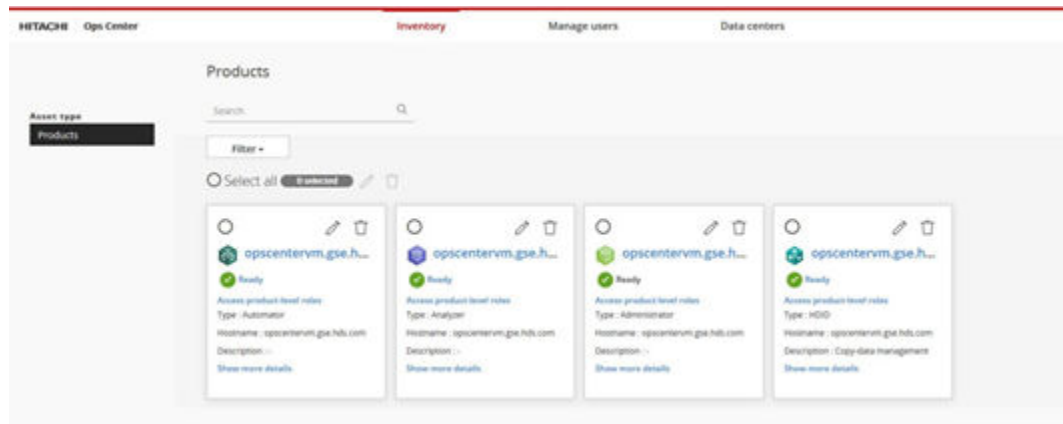
There are multiple options for installing Hitachi Ops Center:

- Installing Hitachi Ops Center products by using the OVA file in a VMware ESXi installation.
- Install Hitachi Ops Center products by using the installer on a bare metal management server.
- Installing Hitachi Ops Center on one VM using the installers on one virtual machine.

Detailed information can also be found in the Installation and Configuration Guide that is part of the Ops Center documentation.

Set up Ops Center

After going through the installation steps and the initial login, the installed Ops Center products are displayed in the inventory screen as shown in the following figure.



See *Getting started in your Ops Center Analyzer environment* on initial configuration of Ops Center Analyzer.

- To understand consumers, resources, and thresholds, see:
 - Overview of managing and monitoring infrastructure resources
 - Consumer settings
 - Resource monitoring settings
- To set up email and notification profiles, see:
 - Configuring the email server
 - Creating a condition profile
- To create consumers and resources, see:
 - Creating a consumer
 - Creating a resource assignment rule
 - Running a resource assignment rule
- To monitor resources, see:
 - Configuring resource monitoring
 - Monitoring using a static threshold
 - Monitoring using a dynamic threshold
 - Setting thresholds for user resources
 - Creating a user resource assignment rule
 - Running a user resource assignment rule

Add probes to the probe server

Using multiple, separate VMs for different Ops Center probe servers is possible, for example with an infrastructure that is only accessible from a management server. You can install a probe server on a management server that can access both - the infrastructure to monitor as well as the main Ops Center server that delivers the monitoring data.

A Fibre Channel connection is required from the probe server to the Hitachi enterprise storage for the Hitachi command device.

The recommended probes for an SAP HANA TDI landscape are:

- Adding Hitachi Enterprise Storage probe
- Adding Brocade FC Switch probe (if applicable)
- Adding Linux probe

Follow the prerequisites and instructions for adding the different types of probes. When configuring the RAID Agent for storage probes, make sure to configure Access Type 1 or 2, depending on your storage system model.

Then, connect to the Ops Center probe server with your Web Browser, click Add Probe, and complete the Add Probe Wizard.

SAP considers server OS probes to be third-party software/external tools on an SAP HANA server. SAP has specific rules for SAP HANA appliance servers as outlined in [SAP Note 1730928](#) and [SAP Note 1730929](#).

In a TDI environment, the rules are more relaxed, especially for Hitachi-only solutions where you get all components from Hitachi Vantara. But in all cases, SAP might request removal of the probe while investigating for a support ticket. This can be done using the probe server web interface without losing the data that was already collected. After the investigation is complete, the probe can be added and activated again.

Best practices

Hitachi Ops Center provides health checks, automated alerts, and an overview of performance metrics for the full infrastructure stack.

Matching SAP HANA alerts in Ops Center

Many non-application-focused SAP HANA alerts can easily be matched in Ops Center, especially storage and server CPU-related alerts. Corresponding infrastructure-related HANA alerts in Ops Center are also beneficial.

Ops Center provides one view for all components of the TDI landscape, from the operating system through switches, to the storage system. It provides a consolidated view of multiple servers/instances in a single dashboard across the entire infrastructure stack. With this, it is possible to go deeper into the stack to make sure you get the most out of your consolidation with regards to performance, reliability, and cost-optimization.

SAP HANA alert reference

The following table provides examples of SAP HANA alerts that can be matched, and therefore also provided in Hitachi Ops Center to provide a single point of reference for the infrastructure administrators who are not purely focused on application-centric alerts.

ID	Name	Description	Category	Unit	Thresholds / Severity	
2	Disk usage	Determines what percentage of each disk containing data, log, and trace files is used. This includes space used by non-SAP HANA files.	Disk	percent	Default thresholds: 90, 95, 98	Severity level(s): 2, 3, 4
5	Host CPU Usage	Determines the percentage of CPU idle time on the host and therefore whether or not CPU resources are running low.	CPU	percent	Default thresholds: 25, 15, 10	Severity level(s): 2, 3, 4
30	Check internal disk full event	Determines whether or not the disks to which data and log files are written are full. A disk-full event causes your database to stop and must be resolved.	Disk	state	Default thresholds: HANDLED, NEW	Severity level(s): 2, 4
51	Size of diagnosis files	Identifies large diagnosis files. Unusually large files can indicate a problem with the database.	Diagnosis Files	MB	Default threshold(s): 1024, 2048	Severity level(s): 2, 3
77	Database disk usage	Determines the total used disk space of the database. All data, logs, traces, and backups are considered.	Disk	GB	Default threshold(s): 300, 400, 500	Severity level(s): 2, 3, 4

For more details, see the [SAP HANA Alert Reference](#).

For example, the information provided by the SAP HANA disk usage or disk full alerts on the operating system side is the same as the information provided by the attributes *usedpercent* and *freepersent* of the Linux probe *IHostFileSystemDevice* resource for the following file systems:

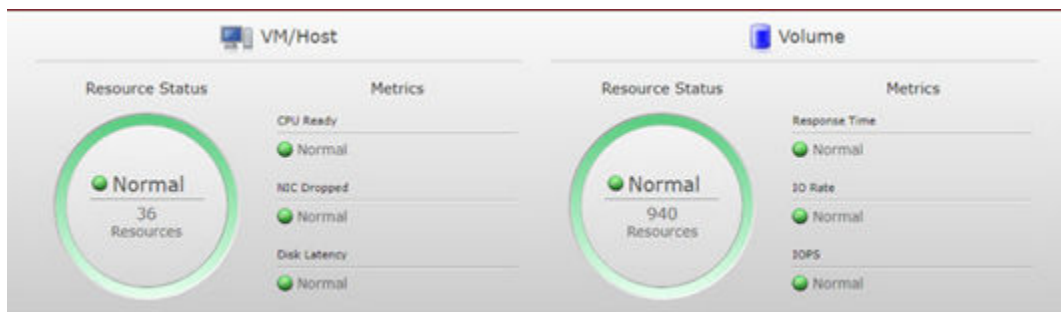
- HANA Data Volume
- HANA Log Volume
- HANA Shared Volume

Ops Center Analyzer performance metrics

There are a large number of pre-built performance reports that can be displayed in Ops Center Analyzer dashboards. The following figure shows the *System Resources* report.



The pre-built reports provide [Quick access to vital performance reports](#) and can be easily customized to your needs by adding [Additional Dashboard reports](#) as shown in the following figure.



This enables different users to see different reports and dashboards, depending on their needs and preferences.

Creating custom reports using report templates or even creating custom reports using the Analyzer detail view report builder can extend the existing monitoring capabilities and tailor them to your needs when displayed in a dashboard or sending dashboard reports to users by email.

The following table provides a list of recommended metrics to monitor your SAP HANA TDI storage consolidation environment.

Probe	Resource Type	Attribute ID	Description
Linux	IHostDisk	l_avgqu-sz	Average concurrency overall during an interval.
Linux	IHostDisk	l_await	Monitor redo logging I/O response time.
Linux	IHostDisk	l_svctm	Monitor redo logging I/O response time.
Hitachi Enterprise Storage	raidClpr	writePendingRate	Of the cache memory allocated for a CLPR, the percentage that is being used by write-pending data.
Hitachi Enterprise Storage	raidMPB	Utilization	Usage rate of the processor.
Hitachi Enterprise Storage	raidMpbClpr	writePendingRate	Of the cache memory allocated for each MP blade in the CLPR, the percentage that is being used by the write pending data of the LDEV that is owned by the MP Blade of the MP Blade ID field.
Hitachi Enterprise Storage	raidMP	utilization	Usage rate of the processor.
Hitachi Enterprise Storage	raidLdev	transferRate	Total data transfer rate.
Hitachi Enterprise Storage	raidDkaEsw	accessPathUsage	Usage rate of the access path between the disk adapter and the cache switch.
Hitachi Enterprise Storage	raidPG	utilization	Usage rate of the parity group. The total usage rate of the logical devices that associate with the parity groups.
Hitachi Enterprise Storage	raidPG	transferRate	Total data transfer rate.

Probe	Resource Type	Attribute ID	Description
Hitachi Enterprise Storage	raidPG	totalIOPS	Sum of read IOPS and write IOPS.
Hitachi Enterprise Storage	raidPool	transferRate	Total data transfer rate <ul style="list-style-type: none"> Read Xfer Rate: Total read data transfer rate for virtual volumes mapped to the Dynamic Provisioning pool. Write xfer Rate: Total write data transfer rate for virtual volumes mapped to the Dynamic Provisioning pool.
Hitachi Enterprise Storage	raidPool	readResponseTime	Average processing time per read request for virtual volumes mapped to the Dynamic Provisioning pool.
Hitachi Enterprise Storage	raidPool	writeResponseTime	Average processing time per write request for virtual volumes mapped to the Dynamic Provisioning pool.
Hitachi Enterprise Storage	raidPool	responseTime	<ul style="list-style-type: none"> Read Response Time Rate: Average processing time per read request for virtual volumes mapped to the Dynamic Provisioning pool. Write Response Time Rate: Average processing time per write request for virtual volumes mapped to the Dynamic Provisioning pool.
Hitachi Enterprise Storage	raidStorage	writePending	Of the cache memory, the cache memory size used by write pending data.
Brocade FC Switch (CLI)	fabBrocadeSwitchPort	fabPortRXUtil	The percentage of individual port (or trunk) bandwidth being used at the time of the most recent poll.

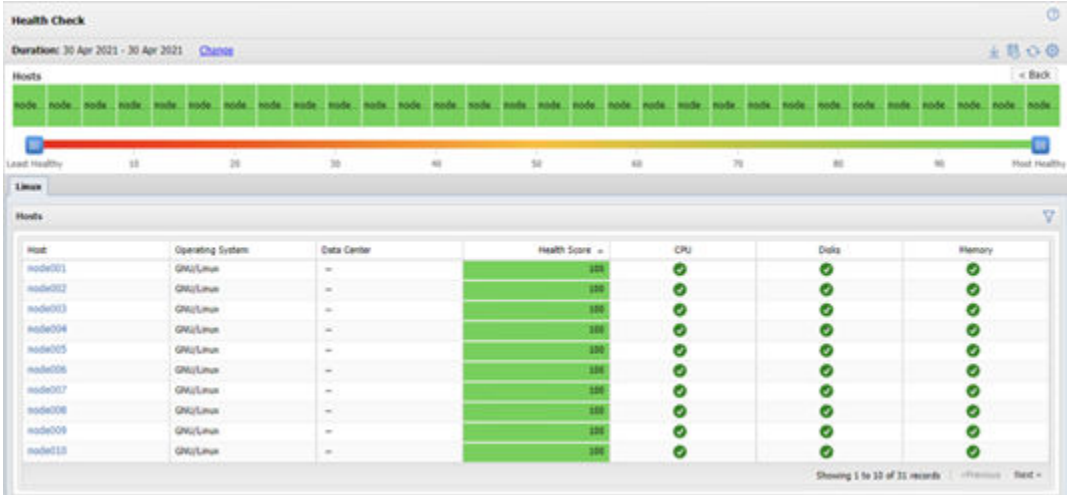
Probe	Resource Type	Attribute ID	Description
Brocade FC Switch (CLI)	fabBrocadeSwitchPort	fabPortTXUtil	The percentage of individual port (or trunk) bandwidth being used at the time of the most recent poll.

Ops Center Analyzer health checks

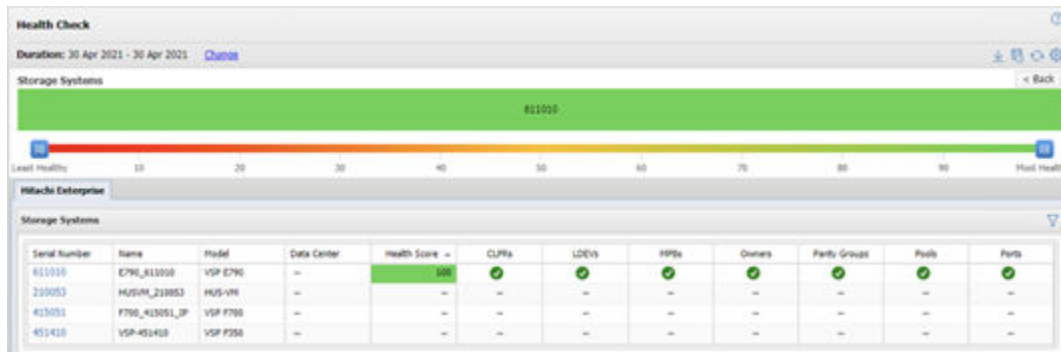
Using Analyzer detail view Health Check in Ops Center Analyzer detail view provides an overview of the health status of your entire infrastructure in a single view, as shown in the following figure.



By clicking the different categories, you can get a more detailed view of all the components.



Each host or storage system contains a link in the first column that leads to a detailed report containing standard performance and configuration data.



For example, by clicking the serial number of the Virtual Storage Platform E790 storage system, the following default report is opened.



This report provides a more detailed view of the storage performance figures and can be customized to extend the default figures.

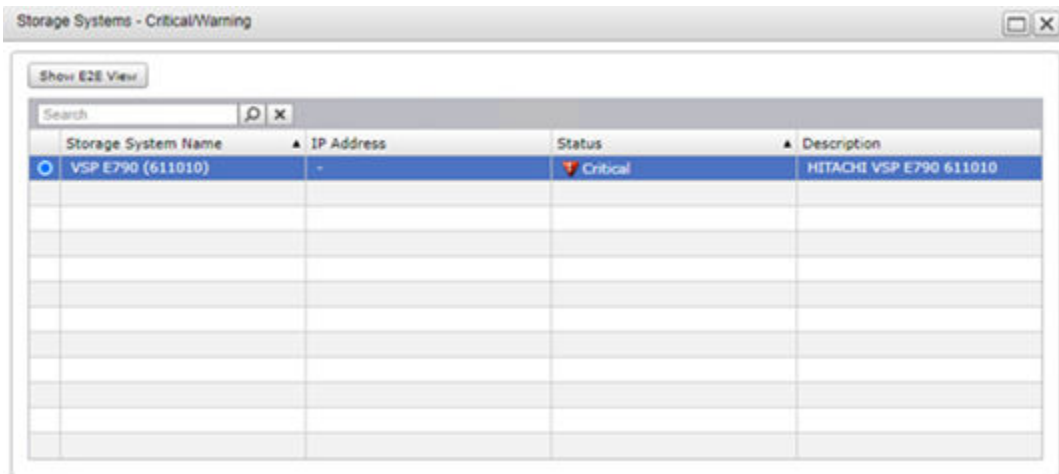
Health checks can be downloaded and sent via email on a regular schedule. This provides a ZIP file containing a summary and more detailed information for every category.

Investigating performance problems

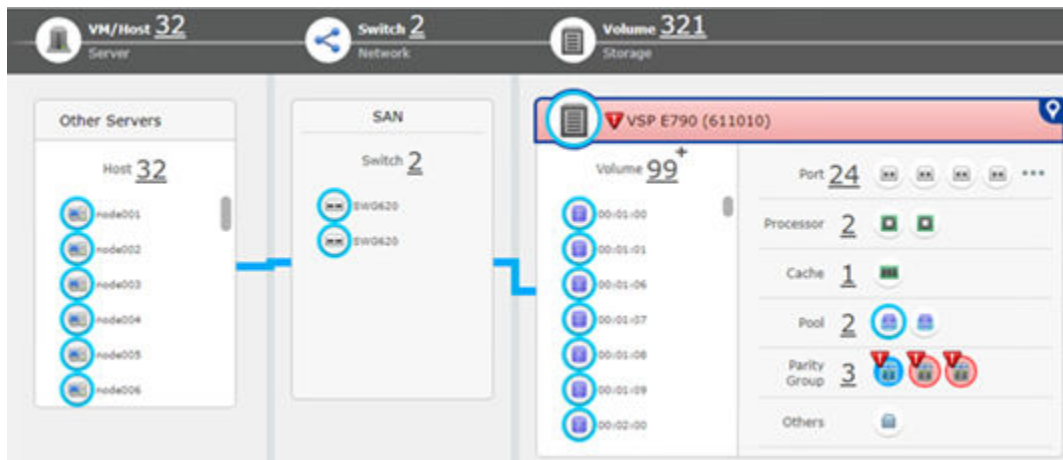
Dashboards also serve as a starting point for closer analysis in case of warnings or alerts. The following figure shows a dashboard that presents vital information about the Virtual Storage Platform E790 storage system including critical alerts.



To investigate the alerts, you can click any of the warning or alert signs in the dashboard to get a popup window from which you can access the end-to-end view (E2E view).



From there you can see all affected resources and go deeper into the analysis.



Engineering validation

For the Engineering validation testing of this SAP HANA storage consolidation project, 32 SAP HANA certified Hitachi servers were connected to a Hitachi Virtual Storage Platform E790 storage system using one pair of Brocade G620 SAN switches.

Testbed setup

The storage configuration tested in the Hitachi Vantara lab used NVMe SSDs for the persistent storage of the SAP HANA nodes according to Hitachi Vantara storage consolidation best practices. In this case, all 32 nodes had the same size memory, but variable sizes of main memory consolidated on the same storage system is also possible.

The following table lists the Hitachi Virtual Storage Platform E790 features.

Item		Specifications	
	Type	VSP E790	
System	SVOS Version		Hitachi Storage Virtualization Operating System v9.5
	Max. Drives	NVMe SSD	24
	Drive Options	NVMe SSD	1.9 TB, 3.8 TB, 7.6 TB, 15.3 TB
	RAID Level		RAID-10, RAID-5, RAID-6
	RAID group configuration	RAID-10	2D+2D
		RAID-5	3D+1P, 4D+1P, 6D+1P, 7D+1P
		RAID-6	6D+2P, 12D+2P, 14D+2P
	Max. number of LDEVs		49152
	Max. storage capacity		0.33 PB
Max. external configuration		192 PB	
Memory	Cache memory capacity		768 GB
Device I/F	Supported channel type	Fibre Channel	24
		Data transfer rate	8 Gbps, 16 Gbps, 32 Gbps
	Max. number of channel boards (CHB)		6

Item		Specifications
	Type	VSP E790
Non-stop maintenance	Control PCB	Supported
	Cache memory	Supported
	Cache flash memory	Supported
	Power supply, fan	Supported
	Microcode	Supported
	Flash drive	Supported
	Flash module drive	Supported

The hardware used in the Hitachi Vantara lab environment to test and validate this storage consolidation setup with a Hitachi Virtual Storage Platform E790 using NVMe SSDs is listed in the following table.

	Tested Number of SAP HANA Nodes	Storage Components	Fibre Channel Switch
VSP E790 (NVMe SSD)	32	<ul style="list-style-type: none"> ▪ CTL: 1 pair ▪ 32 Gbps 4-port CHB: 3 pairs ▪ MPU: 1 pair ▪ Cache: 768 GB ▪ NVMe SSDs: 24 × 1.9 TB 	48-port Brocade G620 Fibre Channel switch: 1 pair

Optional hardware that can be used with any of the Hitachi enterprise storage systems certified for use in an SAP HANA TDI landscape is listed in the following table.

Optional Item	Components
Management server	Hitachi Advanced Server DS120
1 GbE management network switch, using either one of these switches.	Brocade ICX 6430-48 port
	Cisco Nexus 3048
10 GbE network switch, using either one of these switches.	Brocade VDX 6740-48 port
	Cisco 93180YC-EX

Validation testing

Scalability and storage KPI testing was performed using [SAP HANA Hardware and Cloud Measurement Tool \(HCMT\)](#). For HCMT details, see [SAP Note 2493172 - Hardware and Cloud Measurement Tools](#) (SAP user credentials required).

The same tool was used to generate workloads for different numbers of nodes in parallel. To get a higher workload on the Virtual Storage Platform E790 storage system, the parameter *FileSize* was increased for all storage I/O tests in the HCMT JSON configuration file by a factor of 3, 4, and even 5.



Note: Testing of this configuration was performed in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that matches your production environment before your production implementation of this solution.

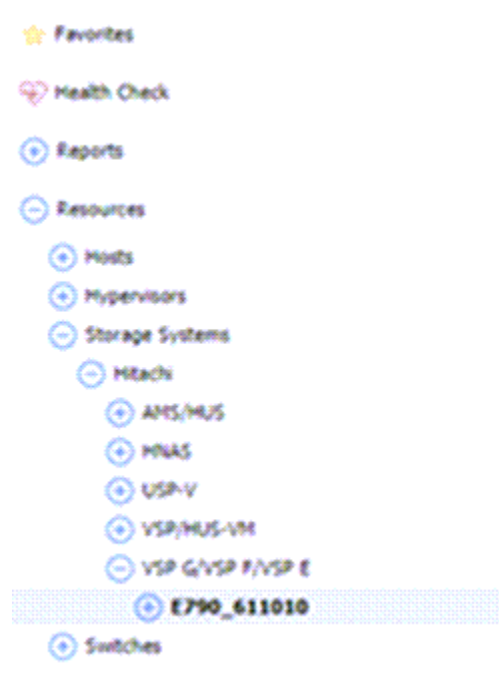
Hitachi Ops Center Analyzer provides a detailed view of the performance data for the Virtual Storage Platform E790 components.

The following section shows some of the performance figures, for an example run of the test set and how to access the performance reports.

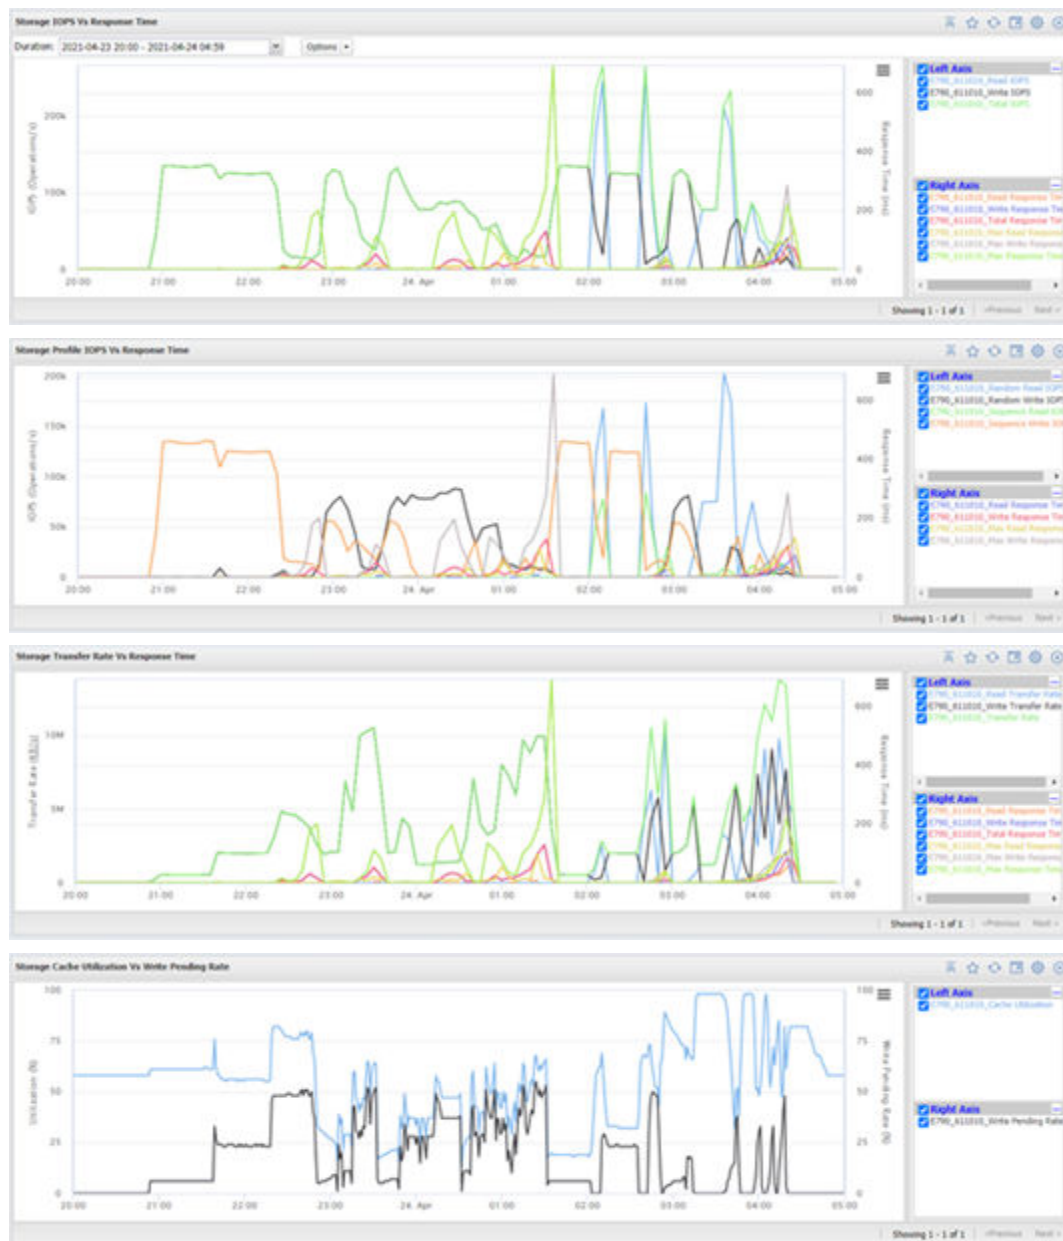
1. Open Hitachi Ops Center Analyzer detail view in your browser.
2. In the top bar, change the duration to match the start and end time of your tests.

This timeframe will be used for all the graphs and reports.

3. In the menu on the left, go to Resources > Storage Systems > Hitachi > VSP G/VSP F/VSP E and then select your VSP E790 storage system.



This opens the default performance report for this storage system, which shows a number of performance graphs.



The graphs can be filtered by the entries on the right side to create specific graphs. All these graphs are measured on a storage system level. For additional details, you can expand the menu for the different storage components to access graphs and figures on a component level.

- ⊖ Resources
 - + Hosts
 - + Hypervisors
 - ⊖ Storage Systems
 - ⊖ Hitachi
 - + AMS/HUS
 - + HNAS
 - + USP-V
 - + VSP/HUS-VM
 - ⊖ VSP G/VSP F/VSP E
 - ⊖ **E790_611010**
 - + MPBs
 - + Cache (CLPRs)
 - + Ports
 - + LDEVs
 - + Parity Groups
 - + Pools
 - + Owners
 - + JNL Groups
 - + Access Paths
 - + GAD V-Vols

For example, you can see the processor utilization is well distributed between the VSP E790 MPBs:



On the Parity Group Utilization graph, you can see that the HCMT tests stress the I/O for the data volumes and log volumes in separate tests:



The HDP pool for the data volume evenly distributes the load for the data volume among its two parity groups.

Especially with a large number of SAP HANA nodes connected to an enterprise storage system, it might help to adapt the existing default reports to maximize their benefit. As mentioned, custom reports can be created either based on existing templates or using the Query Builder to create new reports.

Product descriptions

The following information describes the hardware and software components used in these best practices.

Hitachi Virtual Storage Platform E series family

The Hitachi Virtual Storage Platform E series family provides agile and automated storage built upon the innovative technologies found in our high-end enterprise systems. The expansion of the VSP E series portfolio includes 2 new all-NVMe flash models that deliver super charged, ultra-low latency performance for the business-critical applications that small and midsized businesses rely on.

- Improve IT agility: “Faster-to-market” for IT projects with proven high-performance infrastructure. Brings “enterprise-class” features and benefits to customers of all sizes whose business is outpacing their existing infrastructure and supports modern business processes like DevOPs.
- Financial elasticity that aligns costs to business goals, growth, and use: Customers can “have it their way” with purchase, lease, or cloud-like consumption models.
- Improved workforce efficiency: a better digital experience which boosts customer satisfaction (with both internal LOBs and end-users) and increases business productivity and profitability.

Hitachi Virtual Storage Platform F Series family

Use [Hitachi Virtual Storage Platform F series family](#) storage for a flash-powered cloud platform for your mission critical applications. This storage meets demanding performance and uptime business needs. Extremely scalable, its 4.8 million random read IOPS allows you to consolidate more applications for more cost savings.

Hitachi Virtual Storage Platform F series family delivers superior all-flash performance for business-critical applications, with continuous data availability.

Hitachi Virtual Storage Platform G series family

The [Hitachi Virtual Storage Platform G series family](#) enables the seamless automation of the data center. It has a broad range of efficiency technologies that deliver maximum value while making ongoing costs more predictable. You can focus on strategic projects and consolidating more workloads while using a wide range of media choices.

The benefits start with Hitachi Storage Virtualization Operating System RF. This includes an all new enhanced software stack that offers up to three times greater performance than our previous midrange models, even as data scales to petabytes.

Hitachi Virtual Storage Platform G series offers support for containers to accelerate cloud-native application development. Provision storage in seconds, and provide persistent data availability, all the while being orchestrated by industry leading container platforms. Move these workloads into an enterprise production environment seamlessly, saving money while reducing support and management costs.

Hitachi Ops Center

Manage, optimize, orchestrate and protect your data with advanced IT analytics and automation using Hitachi Ops Center. Achieve new insights, accelerate resource delivery, eliminate risks, and speed innovation to modernize your data center operations.

Use the power of AI operations with the following:

- **Administrator:** Reduce storage management complexities for the Virtual Storage Platform. Intuitive graphical user interfaces (GUIs) and recommended configuration practices speed storage management operations so you spend more time on strategic efforts, not daily tasks.
- **Analyzer:** Improve IT operations with machine learning (ML) to drive resource service levels, utilization and automation at lower costs. Obtain operational visibility from virtual machines, servers, SAN switches to shared storage resources to optimize an application's full data path.
- **Automator:** Deliver resources up to 70% faster than manual processes. Free staff to focus on strategic initiatives.
- **Protector:** Meet tight service level requirements when protecting critical data and applications. Automatically support secondary business functions with data copies staff need to do their jobs. Make better use of backup data for activities, such as e-discovery and analysis. Simplify administration and replication management. Do it all with no disruption to production application availability and performance.

Hitachi Storage Virtualization Operating System RF

Hitachi Storage Virtualization Operating System RF powers the Hitachi Virtual Storage Platform (VSP) family. It integrates storage system software to provide system element management and advanced storage system functions. Used across multiple platforms, Storage Virtualization Operating System includes storage virtualization, thin provisioning, storage service level controls, dynamic provisioning, and performance instrumentation.

Flash performance is optimized with a patented flash-aware I/O stack, which accelerates data access. Adaptive inline data reduction increases storage efficiency while enabling a balance of data efficiency and application performance. Industry-leading storage virtualization allows SVOS RF to use third-party all-flash and hybrid arrays as storage capacity, consolidating resources for a higher ROI and providing a high-speed front end to slower, less-predictable arrays.

Operating system options for SAP HANA

SUSE Linux Enterprise Server for SAP Applications and Red Hat Enterprise Linux for SAP HANA are available operating systems when running SAP HANA.

- SUSE Linux Enterprise Server (SLES) for SAP Applications

Compete more effectively through improved uptime, better efficiency, and accelerated innovation using [SUSE Linux Enterprise Server for SAP Applications](#). This is a versatile server operating system for efficiently deploying highly available enterprise-class IT services in mixed IT environments with performance and reduced risk.

SUSE Linux Enterprise Server was the first Linux operating system to be certified for use with SAP HANA. It remains the operating system of choice for most SAP HANA customers.

- Red Hat Enterprise Linux (RHEL) for SAP HANA

Using the stability and flexibility of [Red Hat Enterprise Linux for SAP HANA](#), reallocate your resources towards meeting the next challenges instead of maintaining the status quo. Deliver meaningful business results by providing exceptional reliability and military-grade security. Use Enterprise Linux to tailor your infrastructure as markets shift and technologies evolve.

Changing the configuration settings is only supported along the guidelines of SAP and the operating system distributor and may otherwise cause significant performance problems. The following SAP Notes for SUSE Linux Enterprise Server and Red Hat Enterprise Linux are a good starting point for information on this topic:

- [1944799 - SAP HANA Guidelines for SLES Operating System Installation](#)
- [2009879 - SAP HANA Guidelines for Red Hat Enterprise Linux \(RHEL\) Operating System](#)

For more details, see "Updating and Patching the Operating System" by searching in the "View SAP HANA document" from [Technical Information and Best Practices](#).

SAP HANA

SAP HANA converges database and application platform capabilities in-memory to transform transactions, analytics, text analysis, predictive and spatial processing so businesses can operate in real-time. This combines database, data processing, and application platform capabilities in a single in-memory platform. Also, the platform provides libraries for predictive, planning, text processing, spatial, and business analytics — all on the same architecture. This architecture comes from leading hardware partners of SAP, including Hitachi Vantara. For more information, see <https://www.sap.com/products/hana.html>.

By eliminating the divide between transactions and analytics, SAP HANA allows you to answer any business question anywhere in real time.

As an SAP customer, you can [download more information](#), including the following:

- SAP HANA Master Guide

This is the central starting point for the technical implementation of SAP HANA. Use this guide for basic concepts and for planning.

- SAP HANA Server Installation and Update Guide

This guide provides an overview of how to install and update an SAP HANA system with the SAP HANA lifecycle management tools.

- SAP HANA Administration Guide

This guide explains how to configure, manage, maintain, and optimize your SAP HANA installation using SAP HANA administration tools.

[SAP HANA hardware directory](#) provides information about SAP HANA appliances certified by SAP hardware partners.

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact