

Hitachi and Veeam Backup and Recovery Solution for SAP HANA

Reference Architecture Guide

© 2022 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, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Edge, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, 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 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
Initial release	March 31, 2022

Reference Architecture Guide

This reference architecture guide describes the backup, recovery, and tiering of SAP systems and SAP HANA databases in a scale-up configuration using Veeam Backup & Replication, which is part of Veeam Availability Suite, and Hitachi Content Platform for cloud scale (HCP for cloud scale).

Continued exponential growth of data in the computing industry has led to an increased demand on storage capacity and optimization of business-critical processes, such as backup and recovery. Furthering the complexity of managing growth are the critical needs surrounding data security, risk mitigation, loss prevention, and compliance with industry standards.

To help users optimize their cloud data infrastructure, improve service levels for backups, and achieve immutability of critical recovery information, Hitachi Vantara and Veeam have partnered to deliver a *Veeam Object Ready – Immutable Backup* solution for enterprise applications such as SAP HANA.

The SAP HANA application is an in-memory database that keeps the bulk of its data in memory and uses persistent storage to provide a backup in case of a failure. However, if the persistent storage itself is damaged, for example due to drive failures or database corruption, then additional backups are needed to protect the data against these types of data loss.

The loss of business-critical system resources and services, such as the SAP HANA platform can translate directly into lost revenue. Therefore, it is critical to perform SAP HANA database backups to a secondary storage media, such as external storage or tape libraries, and to be able to restore the SAP HANA database from the backup in case of a failure.

Backup solution

Continuous usage of SAP enterprise systems produces massive amounts of data. These large databases increase storage costs and impact database performance, the backup window, and migration downtime.

Hitachi Content Platform for cloud scale (HCP for cloud scale) with Amazon Web Service (AWS) Simple Storage Service (S3) Object Lock is an on-premises cloud object-based storage (OBS) system that delivers a cost-effective, secure, and robust data platform to power many different enterprise use cases.

Veeam Backup & Replication, part of Veeam Availability Suite, delivers fast, flexible and reliable backup, recovery, and replication for virtual, physical, and cloud-based workloads. Veeam Cloud Tier, a feature of Veeam Backup & Replication Enterprise edition, provides a scale-out backup repository for massively scalable capacity and long-term data retention using OBS systems that are AWS S3 compliant, such as HCP for cloud scale.

Primary backups are tiered to HCP for cloud scale based on aging or capacity policies and are further made immutable by setting business-required retention policies. This helps you achieve the highest speed, lowest cost, and fully secured recovery data that is managed in a coordinated fashion between the Veeam Backup & Replication application and the HCP for cloud scale platform.

Because of the immutable nature of the backups leveraging Object Lock, external threats to corrupt “last resort to recover” data are denied and you can be assured that an “always recoverable data copy” will be preserved.

See [Making Veeam Backup & Replication and databases highly available](#) for more information.

Use this document to understand the example reference architecture for archival, backup, and recovery of SAP HANA in a scale-up configuration, and to help with deployment of the configuration.

This technical paper assumes familiarity with the following:

- Storage area network-based storage systems
- General storage concepts
- General network knowledge
- SAP HANA platform
- Common IT best practices
- SAP High-Performance Analytic Appliance (HANA)



Note: These procedures were developed in a lab environment. Many factors affect production environments beyond prediction or duplication in a lab environment. Follow recommended practice by conducting proof-of-concept testing for acceptable results before implementing this solution in your production environment. Test the implementation in a non-production, isolated test environment that otherwise matches your production environment.

Solution overview

This SAP HANA backup and recovery solution using Backint for SAP HANA, Veeam Backup & Replication, Veeam plug-in for SAP HANA, and Hitachi Content Platform for cloud scale uses the following components:

- Hitachi Advanced Server DS220 G2
- Hitachi Content Platform for cloud scale (HCP for cloud scale)
- Hitachi Virtual Storage Platform G350 (VSP G350)
- SAP High-Performance Analytic Appliance (HANA)
- Veeam Backup & Replication
- Veeam Plug-in for SAP HANA

Key solution elements

The key hardware and software components used in this reference architecture are described in the following tables. For detailed component information, see [Product descriptions \(on page 20\)](#).

Hardware elements

The following table describes the hardware configuration used in this solution.

Table 1 Hardware Elements

Hardware	Quantity	Configuration	Role	Implementation Type
Hitachi Advanced Server DS220 G2	1	CPU – 2 Intel 3rd Gen Intel® Xeon Scalable Processors: <ul style="list-style-type: none"> ▪ 8368 (38 cores, 270W, 2.4GHz) ▪ 8380 (40 cores, 270W, 2.3GHz) 2 heatsinks CPU 0/1 RAM per SAP HANA node: Refer to the supported configuration. 2 × USB 3.0 ports KVM connector (VGA, COM)	SAP HANA server	All
Hitachi Advanced Server DS220	1	CPU for each SAP HANA node: <ul style="list-style-type: none"> ▪ 2 Intel Xeon Platinum 8276 Base/Medium/Large Processor 28 cores, 2.2 GHz, 152 W or ▪ 2 Intel Xeon Platinum 8280 Base/Medium/Large Processor 28 cores, 2.70 GHz, 205 W 768 TB RAM	vSphere 6.7 U3	Hosting: Veeam Backup & Replication server

Hardware	Quantity	Configuration	Role	Implementation Type
Intel PCIe network cards	2	2 Intel E810-XXVDA2 2-port 25 GbE network cards	SAP HANA client network.	All
	2	Cisco SFP+ 3M twinaxial cables		
Hitachi Virtual Storage Platform G350 or F350	1	CTL: 1 pair 32 Gbps 4-port CHB: 1 pair MPU: 1 pair Cache: 128 GB	Storage for Veeam Backup & Replication server	Storage for Veeam Backup copies
Broadcom Emulex LPE 31002 PCIe HBA card	2	2-port LPE31002 16 Gbps	Required on the DS220 servers.	Fibre Channel Card
RAID controller card	1	Broadcom QS-3916 Type B MegaRAID (4G DDR) storage controller card	Required on the Hitachi Advanced Server DS220, when using internal storage.	Internal storage only
Internal storage drives on Hitachi Advanced Server DS220 G2	5	1.92 TB Intel S4510 SATA SSD	Supports less than or equal to 1024 GB RAM per HANA node.	Internal storage only
	9		For 2048 GB per HANA node.	
	11		Supports 4096 GB RAM per HANA node.	

Hardware	Quantity	Configuration	Role	Implementation Type
	1		Spare drive	
Cisco Nexus 92348 switch	1	48 × 1 GbE ports	Optional switch for management network.	All (optional)
Cisco Nexus 93180YC-FX	2	48 × 10/25 GbE ports	Optional switches for the client network or additional backup network.	All (optional)

Software elements

The following table describes the software used in this solution.

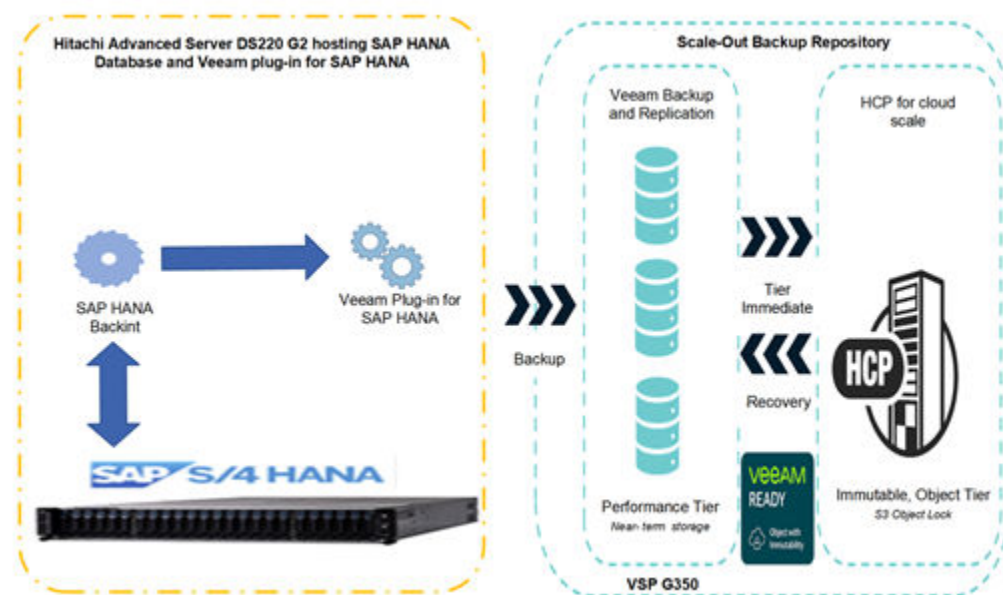
Table 2 Software Elements

Software	
Operating system	SUSE Linux Enterprise Server (SUSE) Linux Enterprise Server for SAP Applications 15SP2 Windows 2019 for Veeam Backup & Replication server
Database	SAP HANA 2.0 SPS05
Hypervisor	VMware vSphere 6.7 U3
Object storage	Hitachi Content Platform for cloud scale
Backup and replication	Veeam Backup & Replication 11
Veeam plug-in	Veeam Plugin-in for SAP HANA

Solution design

The detailed design for SAP HANA backup and recovery solution using Backint for SAP HANA, Veeam Backup & Replication, Veeam plug-in for SAP HANA, and Hitachi Content Platform for cloud scale includes the following, as shown in the following figure:

- SAP HANA installation
- Veeam Backup & Replication server setup
- Veeam plug-in for SAP HANA
- HCP for cloud scale setup and configuration
- Hitachi Content Platform configuration for cloud scale for Veeam as a capacity repository



Overview of the SAP HANA database backup components and configuration:

- SAP HANA 2.0 SPS05 was installed on the SLES15 SP2 operating system.
- TPC-H benchmarking data was loaded into the database before backup and recovery were performed.
- SAP HANA Studio, installed on the management server, was used to perform backup and recovery operations.
- Veeam Backup & Replication was installed and set up on a Windows 2019 virtual machine. See the [Veeam Backup & Replication Product Guides](#) for installation details.
- Veeam Plug-in for SAP HANA was installed on SAP HANA server.

- Veeam Plug-in acted as an agent between the SAP HANA server and Veeam backup repositories. The plug-in was integrated with databases through the SAP HANA Backint component. Backint for SAP HANA is an API that enables Veeam Plug-in to directly connect to the SAP HANA database and send the database backup files to Veeam repositories.
- SAP HANA backup parallelization can be performed for Backint, based on the environment as described in Multi-streaming Data Backups with Third-Party Backup Tools.

To install and configure the plug-in for SAP HANA, see the following resources:

- [Veeam Plug-in for SAP HANA manual](#)
- [Installing HCP for cloud scale](#)
- [Veeam performance repository configuration \(on page 9\)](#)
- [Veeam capacity repository configuration with HCP for cloud scale \(on page 12\)](#)
- [Veeam scale-out backup repository configuration \(on page 16\)](#)

A scale-out backup repository is a repository system with horizontal scaling support for multi-tier data storage. The scale-out backup repository consists of one or more backup repositories called performance tiers that can be expanded with object storage repositories for long-term and archive storage, called capacity tiers and archive tiers. All the storage devices and systems inside the scale-out backup repository are joined into a system, with their capacities summarized.

See [Tier Veeam Backups at Hyperscale with HCP for Cloud Scale](#) for more information.

Veeam performance repository configuration

A backup repository is a storage location where Veeam keeps backup files and metadata. The performance tier of a scale-out backup repository is the level used for fast access to this backup data. A performance extent is a backup repository (except an object storage repository) added to the scale-out backup repository. To configure a performance tier backup repository, you can use the following storage types:

- Direct-attached storage
- Network-attached storage
- Object storage
- Deduplicating storage appliance

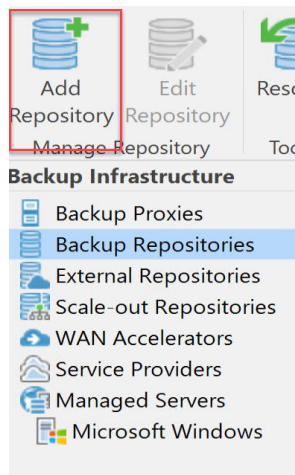
For this repository, Veeam supports any storage system that is currently certified for your hypervisor virtual environment. For our testing, LUs/volumes provisioned from VSP G350 were used as the backend storage for performance tier configuration.

In vCenter or on your ESXi host, edit your Veeam server VM to add a VMDK to support this new backup repository.

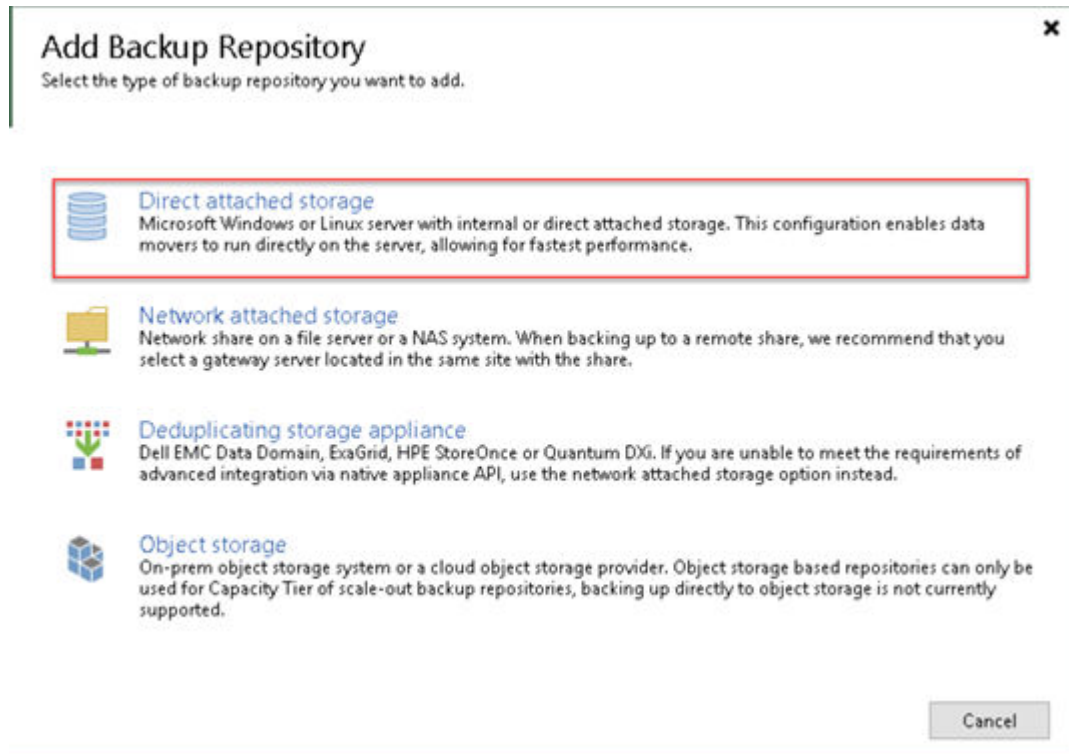
Procedure

1. Create an additional drive and add it to your Veeam server.
 - a. Within vCenter, right click on the Veeam Server and select **Edit Settings**.
 - b. Click **Add New Device** and select **Hard Disk** from the list.

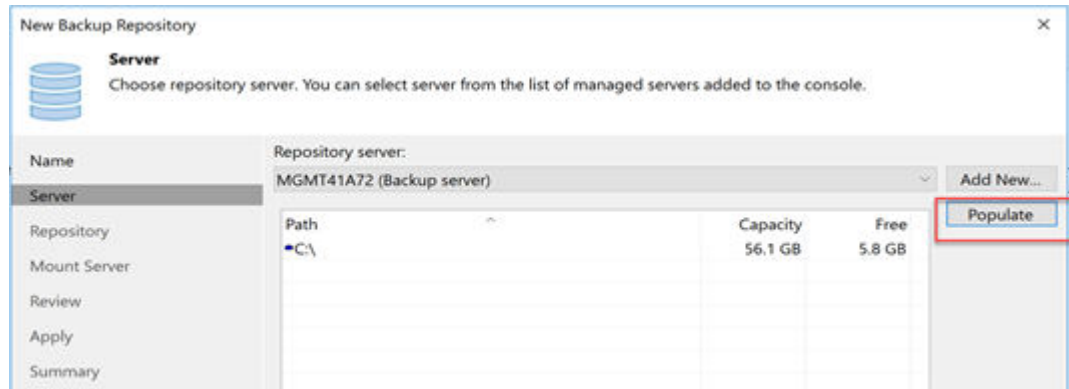
- c. Change the size of the new drive to accommodate the test, and then click **OK**. The actual storage should have ample space to back up the data.
2. Log in to the Veeam Server VM and set up the new storage so that it is available to the Windows environment.
3. Access **Disk Management** to configure the newly accessed storage.
4. Bring the drive online.
5. Create a new volume using the **New Simple Volume Wizard**.
6. Assign a drive letter and use the **ReFS** file system with a **64K** allocation unit size.
7. Within Veeam Backup & Replication, under the **Backup Infrastructure** view, click **Backup Repositories** in the Inventory pane and then click **Add Repository** in the ribbon.




8. In the **Add Backup Repository** screen, select the **Direct attached storage** option. On the following screen, select the Microsoft Windows repository type.



9. Create a name and description for the new backup repository. Click **Next** to continue.
10. On the **Server** screen, click **Populate** to display all available storage paths. Select the drive that you just created for this repository. Click **Next**.



11. In the Repository window, clear (uncheck) the **Limit maximum concurrent tasks** for check box to allow more concurrent operations. Click **Next**.



12. Finish the wizard using default settings and allow the repository updates to complete. When prompted to change the backup configuration to this repository, select **No** to keep the backup configuration on the default backup repository.

Veeam capacity repository configuration with HCP for cloud scale

The capacity tier is an additional tier of storage that can be attached to a scale-out backup repository. Applicable data from the scale-out backup repository performance extents can be transported to the capacity tier for long-term storage.

The capacity extent can be either a cloud-based object storage repository or an on-premises object storage repository, such as the following:

- S3-compatible object storage repository
- Amazon S3
- Microsoft Azure Blob storage
- Microsoft Azure Data Box
- IBM Cloud Object Storage
- Google Cloud Object Storage

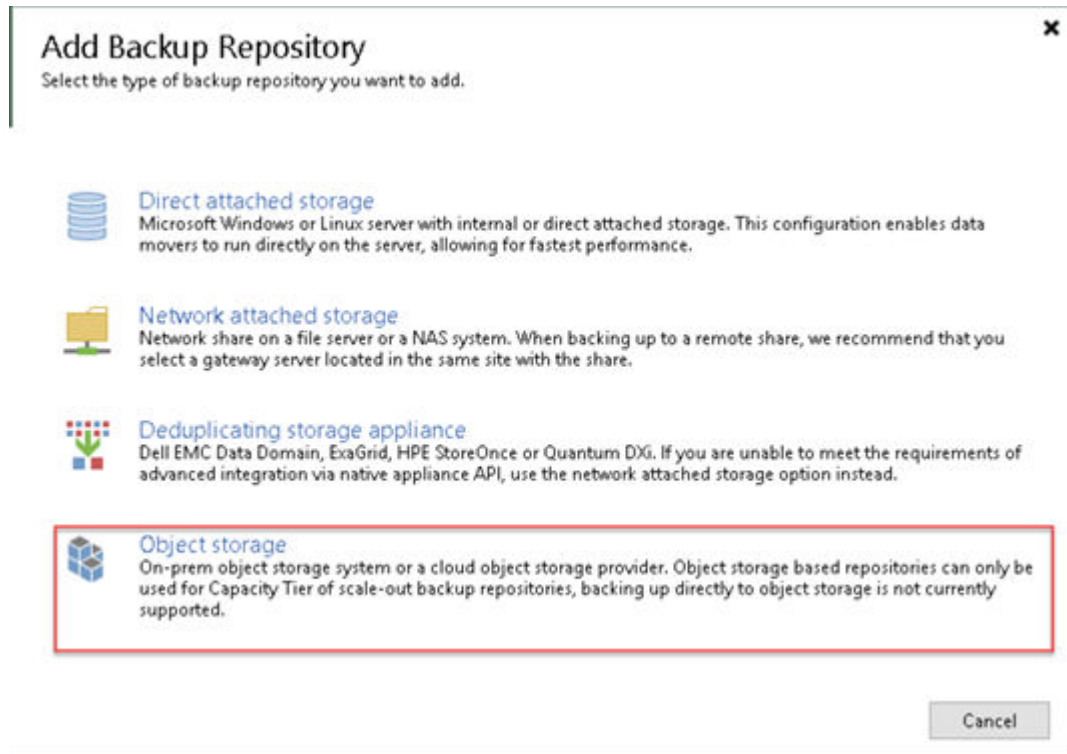
Hitachi Content Platform for cloud scale (HCP for cloud scale) with AWS S3 Object Lock is an on-premises cloud OBS system that delivers a cost-effective, secure, and robust data platform to power many different enterprise use cases.

Object storage repository configuration

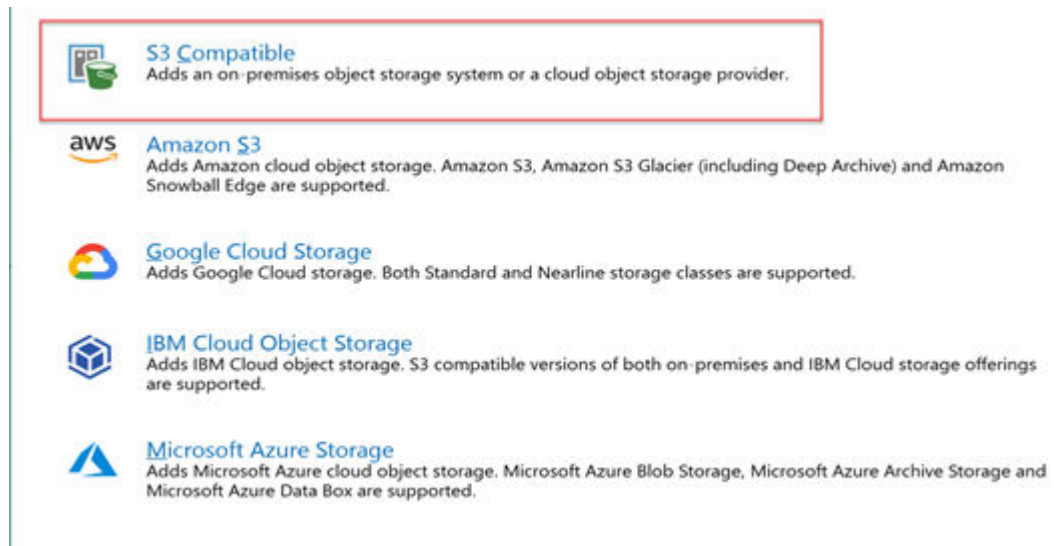
Use the following procedure to create an object storage repository for the capacity tier.

Procedure

1. Under the **Backup Infrastructure** view, right-click **Backup Repositories** and select **Add Backup Repository**.
2. In the **Add Backup Repository** window, select **Object Storage**.



3. Select **S3 compatible** object storage to start the Object Storage Repository wizard.



4. Enter a name and description of the new object storage repository. Click **Next**.
5. Enter your HTTPS endpoint address (for example, *https://<HCP-CS-Domain-Name>*) and region, and then select the credentials that have already been created.

The screenshot shows the 'New Object Storage Repository' wizard, specifically the 'Account' step. The window title is 'New Object Storage Repository' and it has a close button (X) in the top right corner. Below the title bar, there is a 'Account' icon and the text 'Specify account to use for connecting to S3 compatible storage system.' The main area is divided into a left sidebar and a right main panel. The sidebar has four items: 'Name', 'Account' (which is selected and highlighted), 'Bucket', and 'Summary'. The main panel contains the following fields and options:

- Service point:** A text box containing 'https://tryhcpforcloudscale.hitachivantara.com'.
- Region:** A text box containing 'us-east-1'.
- Credentials:** A dropdown menu showing 'DK3V9LaQH9gTR1BaLFFPnCOY4FY0ulaxNWswYUyCZKA23_FbD51N7Tb0vSUTJAqf ~' and an 'Add...' button. Below this is a link for 'Manage cloud accounts'.
- Use the following gateway server:** An unchecked checkbox. Below it is a dropdown menu showing 'MGMT41A72 (Backup server)'. A note below the dropdown reads: 'Select a gateway server to proxy access to the object storage system. If no gateway server is specified, all scale-out backup repository extents must have direct network access to the storage system.'

At the bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Leave the **Use the following gateway server** box unchecked, and then click **Next**. Accept any certificate warnings and click **Continue**.

6. In the **Bucket** window, identify the Bucket and click **Browse** to create a new folder for the capacity tier object storage. Check the box next to **Make recent backups immutable for:** and ensure that the number of days is set to **1 day**. Click **Next**.

New Object Storage Repository X

Bucket
Specify object storage system bucket to use.

Name	Bucket:	veeamtest	Browse...
Account			
Bucket	Folder:	SAP	Browse...
Summary			

Limit object storage consumption to: 10 TB
This is a soft limit to help control your object storage spend. If the specified limit is exceeded, already running backup offload tasks will be allowed to complete, but no new tasks will be started.

Make recent backups immutable for: 1 days
Protects recent backups from modification or deletion by ransomware, malicious insiders and hackers using native object storage capabilities. Object storage must support S3 Object Lock feature.

< Previous
Apply
Finish
Cancel

7. Review the **Summary** screen. Click **Finish** when done.

New Object Storage Repository X

Summary
You can copy the configuration information below for future reference.

Name	Summary:
Account	Object storage repository was successfully created.
Bucket	Name: Object storage repository 1
Summary	Description: Created by MGMT41A72\Administrator at 3/22/2022 12:03 PM.
	Type: S3-compatible
	Gateway server: not selected
	Service point: https://tryhcpforcloudscale.hitachivantara.com
	Region: us-east-1
	Bucket: veeamtest
	Concurrent tasks limit: unlimited
	Storage consumption limit: unlimited
	Recent backups will be immutable for 1 day

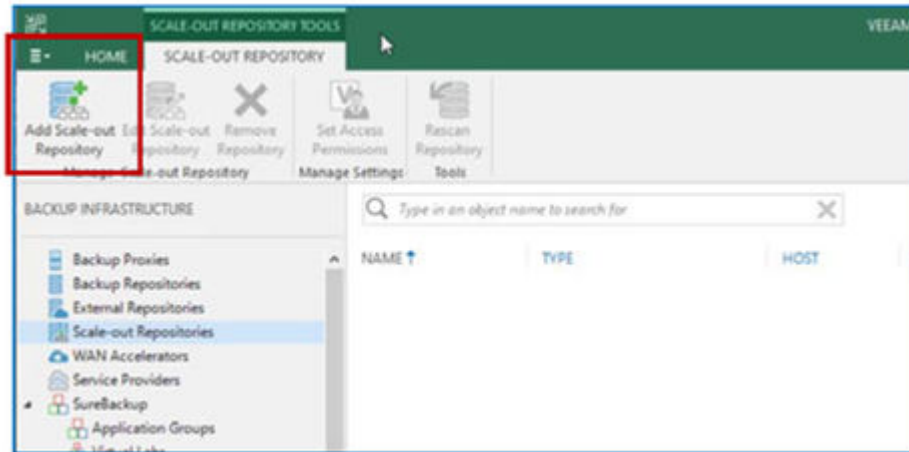
< Previous
Next >
Finish
Cancel

Veeam scale-out backup repository configuration

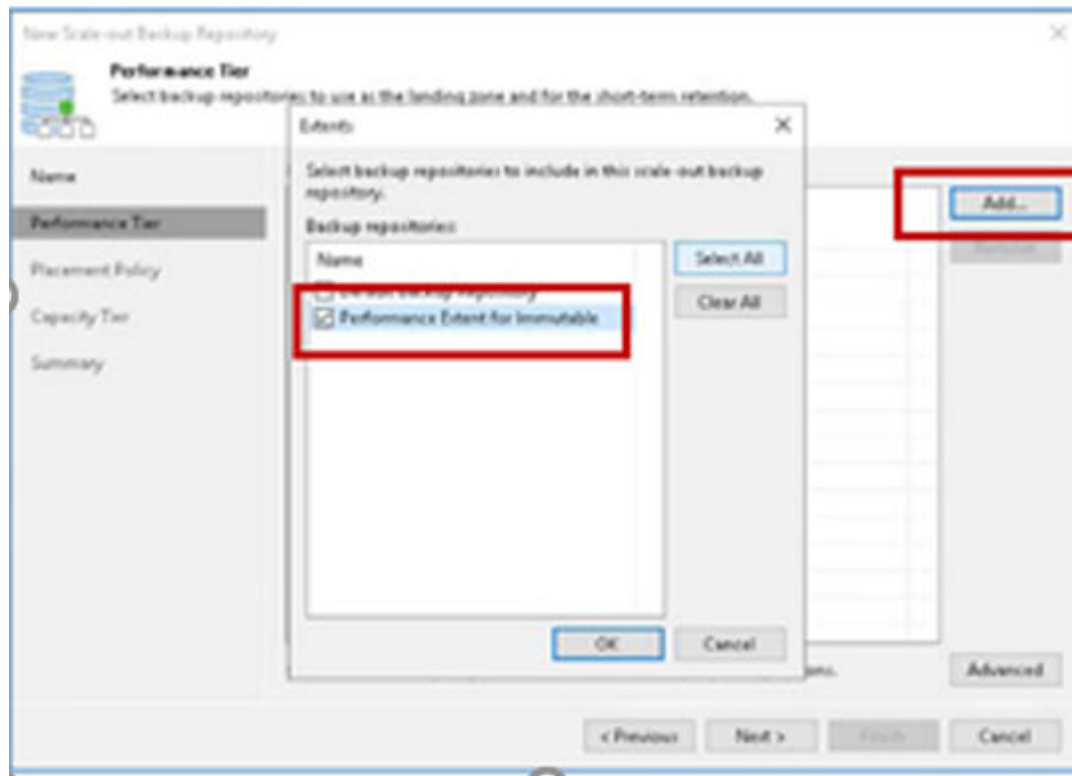
Use the following procedure to configure a scale-out backup repository.

Procedure

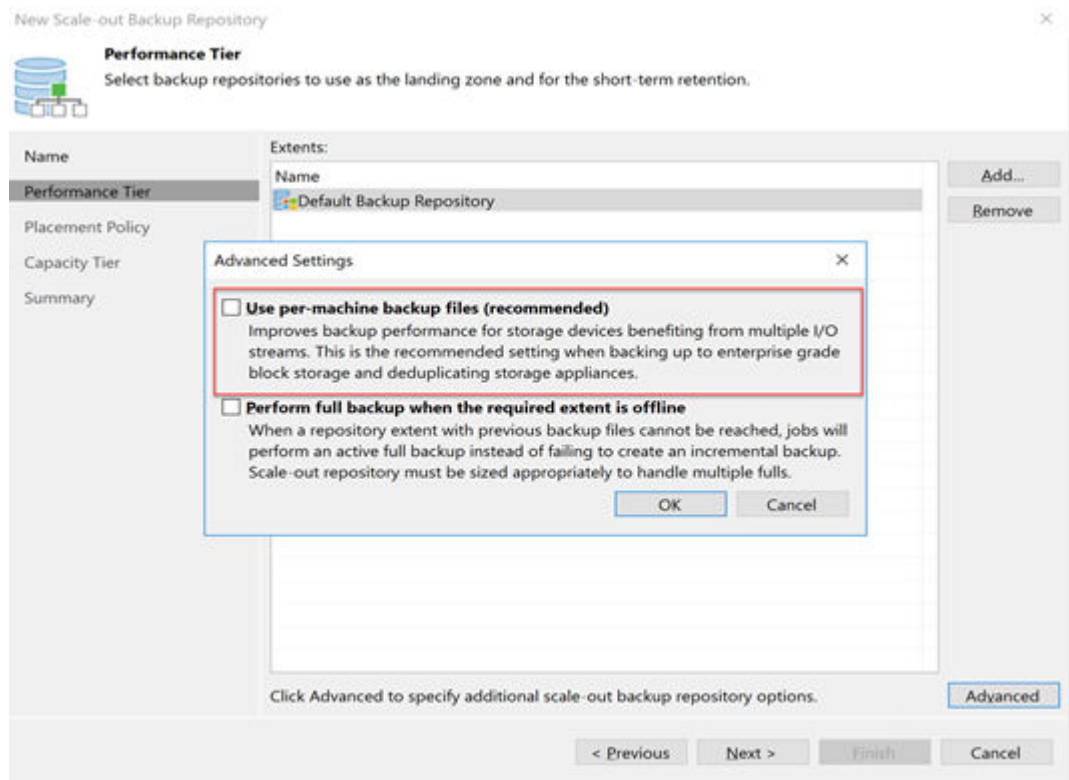
1. Under the **Backup Infrastructure** view, click **Scale-out Repositories** and then click **Add Scale-out Repository** on the ribbon.



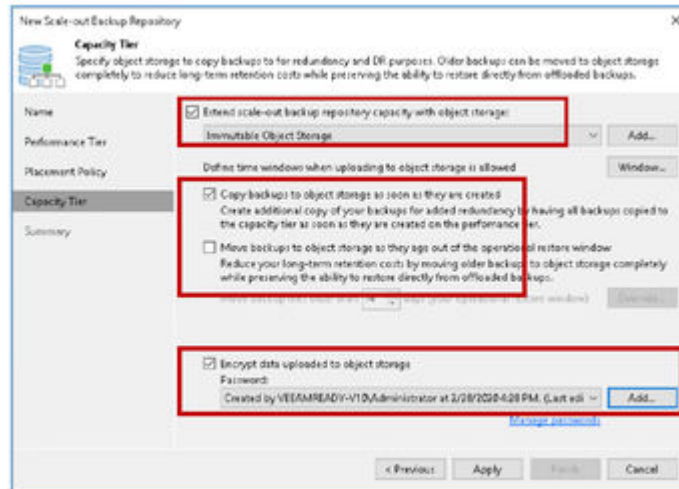
2. Type a name and description for this new scale-out repository. Click **Next**.
3. In the **Performance Tier** window, click **Add** to add a backup repository to the scale-out repository. Add the performance tier repository that was created for your immutability backups. Click **OK**.



- Click **Advanced** to view advanced settings for the performance tier. Uncheck the **Use per-VM backup files**, and then click **OK**. Click **Next**.



- In the **Placement Policy** window, keep the default **Data Locality** policy. This policy allows all backup files (full and incremental) to be located on the same extent. Click **Next**.
- In the **Capacity Tier** window, check the box next to **Extend scale-out backup repository capacity with object storage** and select the immutable object storage solution for repository from the list.
- Under the object storage selection, check the **Copy** box and uncheck the **Move** box. (For this test, we are going to copy backup data immediately to the object storage after it is created.)
- Check the box next to **Encrypt data uploaded to object storage**, and then click **Add** to add a password for your encryption. Click **Apply**.



9. In the **Summary** window, click Finish to create the scale-out repository for the test.

Engineering validation

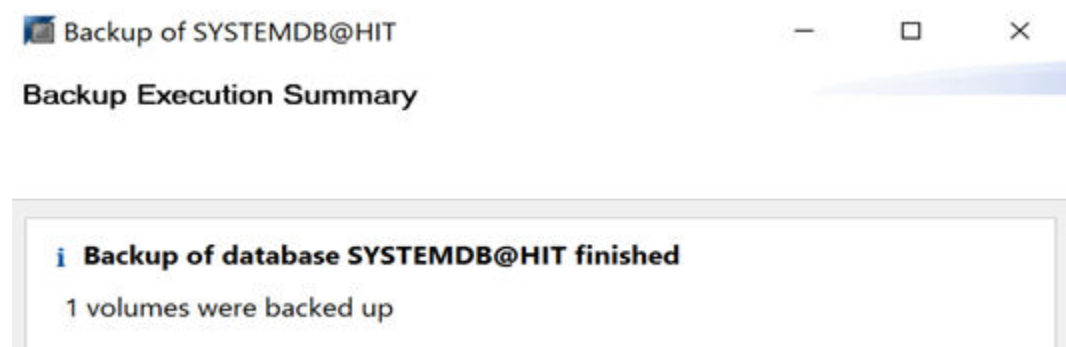
After configuration of all components was complete, backups and recovery of SAP HANA database to/from Veeam Backup & Replication were performed using SAP HANA Studio.

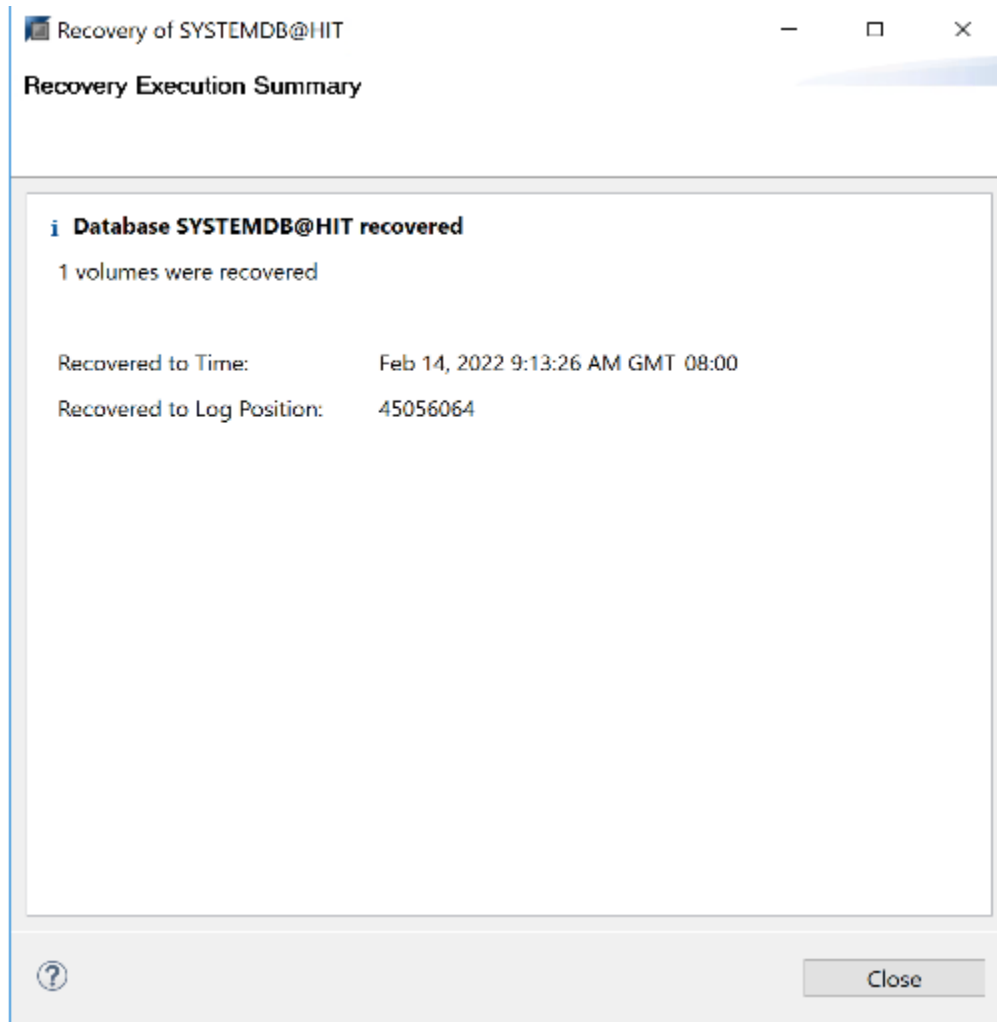
The following operations were successfully completed using SAP HANA Studio:

- Backup
 - Complete
 - Differential
 - Incremental backups
- Recovery
 - Recover the database to its most recent state.
 - Recover the database to a specific data backup performed from HANA Studio.
 - Tiering and recovery from HCP for cloud scale.

See [Database Backup using SAP HANA Studio](#).

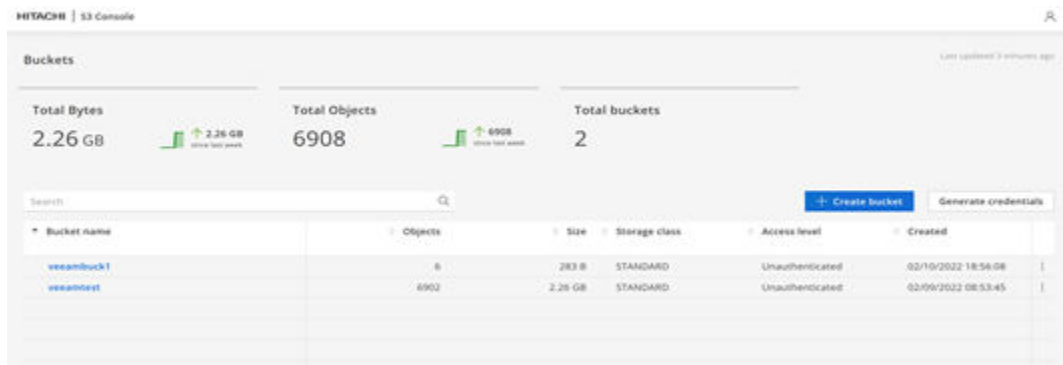
The following figures show a summary of backup information using SAP HANA Studio.





hanveeam SAP backint restore (Scale-out Backup Repository 1)	SAP Restore	Success	2/14/2022 9:15 AM	2/14/2022 9:16 AM
hanveeam SAP backint backup (Scale-out Backup Repository 1)	SOBR Taring	Success	2/14/2022 8:59 AM	2/14/2022 9:00 AM
hanveeam SAP backint backup (Scale-out Backup Repository 1)	SAP Backup	Success	2/14/2022 8:57 AM	2/14/2022 8:59 AM

The following figures show a summary of backup statistics on HCP cloud scale.





Product descriptions

The following information describes the hardware and software components used in this reference architecture.

Hitachi Advanced Server DS220 G2

With a combination of two Intel Xeon Scalable processors and high storage capacity in a 2U rack-space package, [Hitachi Advanced Server DS220 G2](#) delivers the storage and I/O to meet the needs of converged solutions and high-performance applications in the data center.

The Intel Xeon Scalable processor family is optimized to address the growing demands on today's IT infrastructure. The server provides 32 slots for high-speed DDR4 memory, allowing up to 4 TB memory capacity with RDIMM population (128 GB × 32) or 8TB (512 GB × 16) with Intel Optane Persistent Memory population.

DS220 G2 comes in three storage configurations to allow for end user flexibility. The first configuration supports 24 2.5-inch non-volatile memory express (NVMe) drives, the second supports 24 2.5-inch serial-attached SCSI (SAS), serial-ATA (SATA) and up to 8 NVMe drives, and the third supports 12 3.5-inch SAS or SATA and up to 8 NVMe drives. All the configurations support hot-pluggable, front-side-accessible drives as well as 2 optional 2.5-inch rear mounted drives. The DS220 G2 delivers high I/O performance and high capacity for demanding applications and solutions.

Hitachi Content Platform for cloud scale

Hitachi Content Platform for cloud scale (HCP for cloud scale) is a software-defined object storage solution that is based on a massively parallel microservice architecture, and is compatible with the Amazon S3 application programming interface (API). HCP for cloud scale is well suited to service applications requiring high bandwidth and compatibility with Amazon S3 APIs.

Veeam Backup & Replication

Veeam Backup & Replication is a comprehensive data protection and disaster recovery solution. With Veeam Backup & Replication, you can create image-level backups of virtual, physical, cloud machines and restore from them. Technology used in the product optimizes data transfer and resource consumption, which helps to minimize storage costs and the recovery time in case of a disaster.

Veeam Backup & Replication provides a centralized console for administering backup/restore/replication operations in all supported platforms (virtual, physical, cloud). Also, the console allows you to automate and schedule routine data protection operations and integrate with solutions for alerting and generating compliance reports.

For more information, see the [Veeam Backup & Replication User Guide](#).

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 a 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 a 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.

Veeam Plug-in for SAP HANA

Veeam Plug-in acts as an agent between an SAP HANA server and Veeam backup repositories. The plug-in is integrated with databases through the SAP HANA Backint component. Backint for SAP HANA is an API that enables Veeam Plug-in to directly connect to the SAP HANA database and send the database backup files to Veeam repositories.

Hitachi Virtual Storage Platform Family

Hitachi Virtual Storage Platform is the first 3-D scaling storage platform designed for all data types. Its storage architecture flexibly adapts for performance, capacity, and multi-vendor storage. Combined with the unique Hitachi Command Suite management platform, it transforms the data center.

- **Scale Up** — Meet increasing demands by dynamically adding processors, connectivity, and capacity in a single unit. Provide the highest performance for both open and mainframe environments.
- **Scale Out** — Meet multiple demands by dynamically combining multiple units into a single logical system with shared resources. Support increased demand in virtualized server environments. Ensure safe multi-tenancy and quality of service through partitioning of cache and ports.
- **Scale Deep** — Extend storage value by virtualizing new and existing external storage systems dynamically. Extend the advanced functions of Hitachi Virtual Storage Platform to multivendor storage. Offload less demanding data to external tiers to save costs and to optimize the availability of tier-one resources.

Using flash acceleration in Virtual Storage Platform is transparent to Hitachi Dynamic Provisioning, Hitachi Dynamic Tiering, in-system replication, and remote replication. Flash media efficiency increases with more workload consolidation when using flash acceleration.

Flash acceleration raises Virtual Storage Platform scalability. Express host I/O processing targets flash media or a large pool of disk drives. The performance boost increases random I/O throughput of the virtual storage director and lowers I/O response time.

Hitachi Vantara



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

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