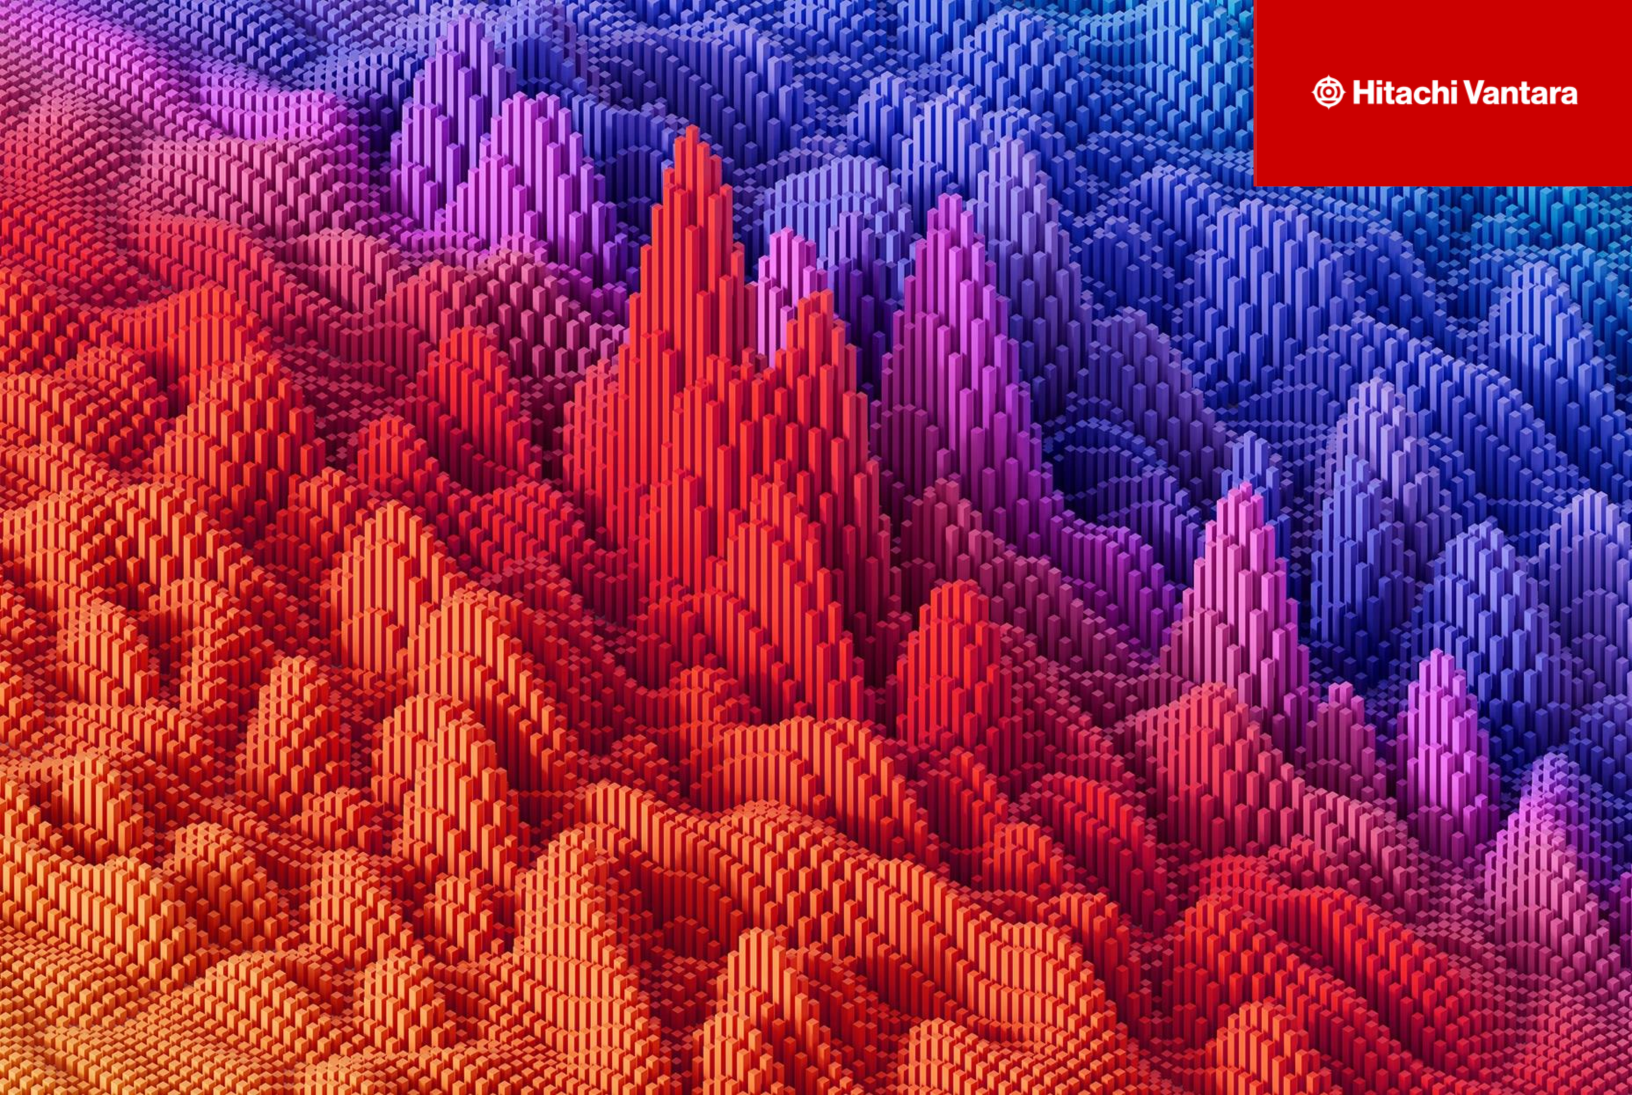

Hitachi Vantara Object Storage with HDPS

Solutions Reference Architecture & Best Practices

Solutions Engineering & Architecture

v1.1.0 – Feb 2025

Table of Contents

Solution Introduction

With Hitachi Data Protection Suite (HDPS), powered by Commvault, Hitachi Vantara and Commvault together deliver a unified, modern offering that facilitates the backup, recovery and management of enterprise and application data with industry-leading object storage offerings: Hitachi Content Platform (HCP) and Hitachi Content Platform for Cloud Scale (HCP CS). The joint solution offers the reliability required by the world's largest organizations, while featuring simplicity, cost-effectiveness and modern capabilities that are needed to remain agile and competitive. Hitachi Vantara's object storage offerings seamlessly extend the secured and guaranteed management of long-term data retention at petabyte scale.

This solution is currently supported with the following versions of Hitachi and Commvault offerings:

- Hitachi Content Platform 9.6
- Hitachi Content Platform for Cloud Scale 2.6
- Hitachi Content Protection Suite 11.32 LTS (powered by Commvault)

HDPS 11.32 LTS has been validated for compatibility with the Hitachi object storage offerings listed above.

No specific use cases (workload) have been validated at this time.

Rebranding and maintenance releases

Commvault delivers one Long-term support (LTS) platform release each year. Each LTS platform release is re-branded to HDPS by Commvault and provided to Hitachi Vantara. A re-branded HDPS release is functionally equivalent to Commvault of the same version. Commvault maintenance releases (MRs) are not re-branded. HDPS customers will update their system to the latest Commvault MR through CommServe without any impact to HDPS branding.

Important:

Commvault MR **11.32.38** includes a fix for automatic database sealing that is required for the solution to function properly. Without this fix, objects in HCP or HCP CS will not be pruned in a WORM scenario. Ensure that your HDPS system is on Commvault MR 11.32.38 or greater.

Partner certification program

(not applicable)

Prerequisites

This document assumes you have firm knowledge and understanding of:

- Features and functionality of the products and components in the solution
- How to install and configure the products and components in the solution
- Any support infrastructure used as part of the solution

Executive Summary

The validation process for solution compatibility covered **basic data lifecycle testing** (backup, restore, and delete functionality) of HDPS with HCP and HCP CS as backup targets.

The backup sources included:

- Microsoft SQL Server
- Oracle databases
- VMware ESXi VMs
- Linux servers
- Window servers

As this is currently a compatible solution, specific use cases, scale, and performance tests were not part of the validation.

Reference Model

The reference implementation used to validate **compatibility** of the solution included the following components:

- Backup sources:
  - Microsoft SQL Server 2016 v18.8
  - Oracle Server v21c
  - VMware VM (various OS)
  - Windows Server 2019
  - Linux Server CentOS 7
- Supporting infrastructure:
  - VMware ESXi Server 7.0
  - Microsoft Windows Server® Standard 2016 for Active Directory
  - 25GbE network with Brocade VDX 6740 switches for front-end and back-end networks

Reference implementation diagram

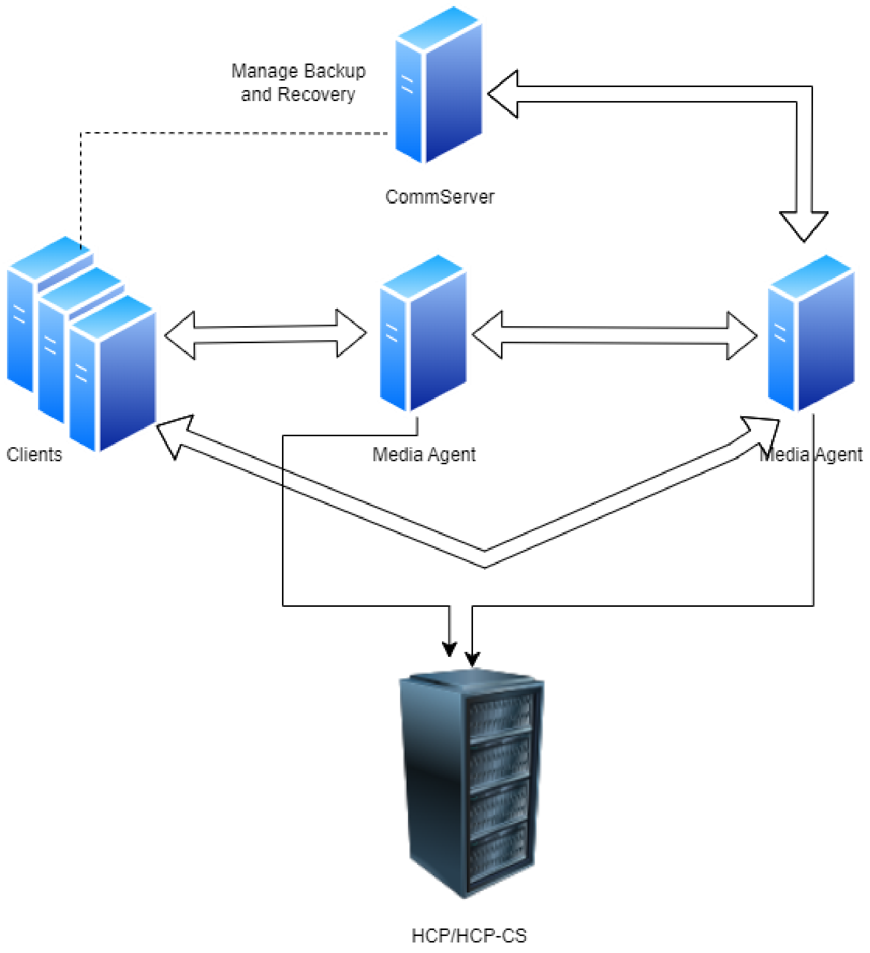

HDPS Configurations

This section summarizes the specific HDPS configurations that are important for solution interoperability.

HDPS configurations

Setting	Value	Notes
WORM Storage Lock	Storage Lock	WORM functionality within the storage backend. Requires S3 Object Lock enabled on HCP namespace or HCP CS bucket on creation.
Compliance Lock	HDPS level WORM	WORM functionality within HDPS that blocks premature deletions.
Max Object Size	With deduplication = 8 MB Without deduplication = 32 MB	These are default values and not end user configurable.
Deduplication Block Size	512 KB	
Chunk Size	4GB	
Media Agent		
DataMoverLookAheadLinkReaderSlot	1024	
SILookAheadAsyncIOBlockSizeKB	4096	
SILookAheadOvlIOThreadLimit	128	
SILookAheadOvlIoThreadsPerRdr	128	
DedupPrunerThreadPoolSizeCloud	60	Not required but suggested. Default is 10 – 20.

Overview of Best Practices for VSP One Object

This section summarizes the specific VSP One Object configurations that are important for solution interoperability.

HDPS MediaAgent configuration

Before VSP One Object can be configured as a Cloud Storage Library the following procedure must be completed on any Media Agents which will be used to access the VSP One Object bucket. For this procedure you will need to know the region that was specified during the VSP One Object installation.

Step 1. Find the Media Agent which will be accessing the VSP One Object system and open the properties window. Select the *Additional Settings* tab.

Step 2. Add the *sCloudDefaultRegion* setting with the following details:

- **Setting Name**: ‘sCloudDefaultRegion’
- **Category**: ‘MediaAgent’

- **Type**: ‘String’
- **Value: <**Provide the region that was specified during the VSP One Object Installation>
- **Enable:** checked

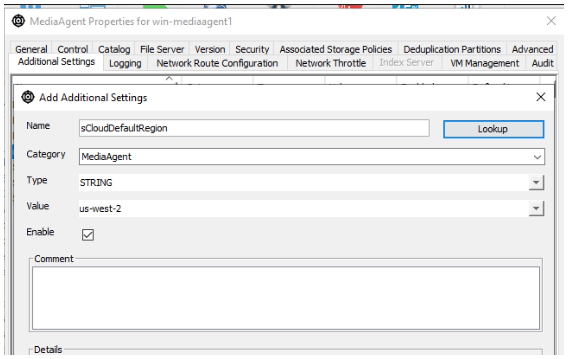

Name: sCloudDefaultRegion
Category: MediaAgent
Type: STRING
Value: us-west-2
Enable: ☑

Step 3. Repeat for all MediaAgents which will be accessing the VSP One Object system.

Step 4. (Optional) Add the *nCloudServerCertificateNameCheck* setting with one of the following values:

- 0: To skip checking the server’s CA Certificate.
- 2: To allow a self-signed certificate to verify the cloud server’s claimed identity.

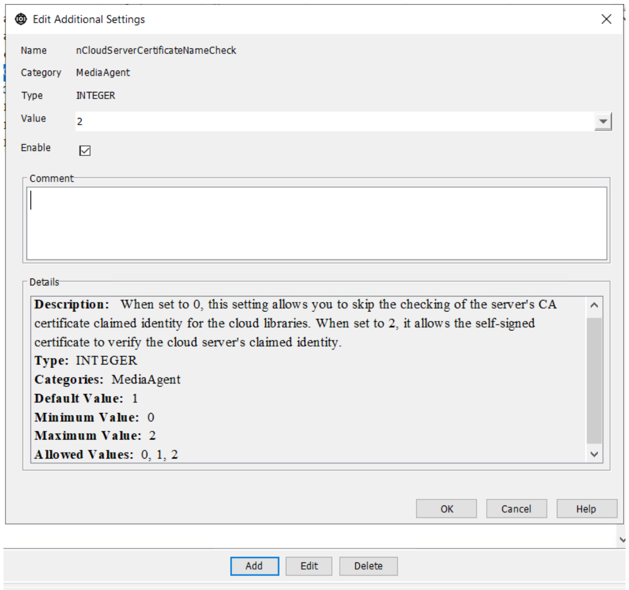

Name: nCloudServerCertificateNameCheck
Category: MediaAgent
Type: INTEGER
Value: 2
Enable: ☑

Description: When set to 0, this setting allows you to skip the checking of the server's CA certificate claimed identity for the cloud libraries. When set to 2, it allows the self-signed certificate to verify the cloud server's claimed identity.
Type: INTEGER
Categories: MediaAgent
Default Value: 1
Minimum Value: 0
Maximum Value: 2
Allowed Values: 0, 1, 2

Step 5. Use the Process Manager to stop and restart the MediaAgent services to apply the new settings.

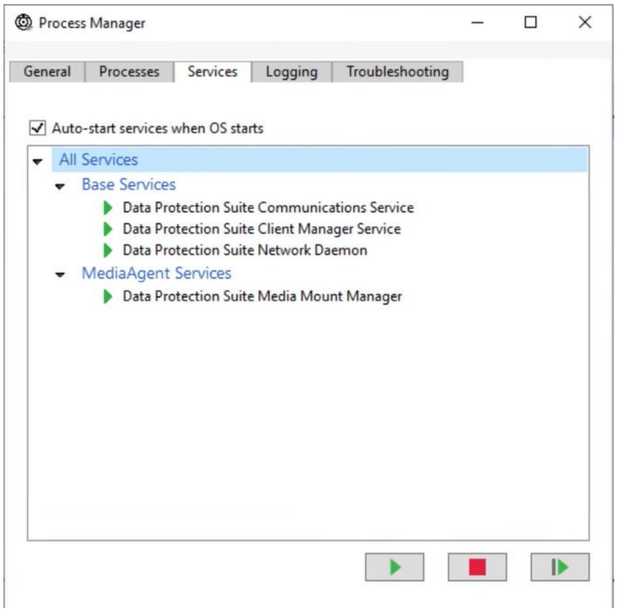

☑ Auto-start services when OS starts

- All Services
  - Base Services
    - Data Protection Suite Communications Service
    - Data Protection Suite Client Manager Service
    - Data Protection Suite Network Daemon
  - MediaAgent Services
    - Data Protection Suite Media Mount Manager

Adding VSP One Object as Cloud Storage

When adding VSP One Object as Cloud Storage you can use one of the following options:

- Hitachi Content Platform for Cloud Scale (S3)
- S3 Compatible Storage

When configuring the Cloud Storage type, enter the full VSP One Object URL in the Service

Host field using the following format: *https://s3.<regionHere>.<myDomain>.com*

Occasionally, you may encounter the following error: "The authorization mechanism you provided is not supported. Please use AWS4-HMAC SHA256. Resource." This occurs because HDPS attempts v4 authentication multiple times and then reverts to v2. To resolve this issue, restart the MediaAgent services.

VSP One Object configurations

Setting	Value	Notes
Object lock	enabled	If using HCP CS enabled WORM. Requires Storage Lock enabled in HDPS.

Overview of Best Practices for HCP

This section summarizes the specific HCP configurations that are important for solution interoperability.

HCP configurations

Setting	Value	Notes
Versioning	enabled	
S3 Overwrite	disabled	Only if Versioning is enabled. Does not overwrite, new objects are ingested even if enabled.
Pruning	unchecked	Only if Versioning is enabled. No new versions are created by HDPS.
S3 Object lock	enabled	If using HCP S3 object lock. Requires Storage Lock enabled in HDPS.
Namespaces	unbalanced	

Best Practices for Multi-Site Deployments

There are several ways to implement multi-site deployment of HDPS and HCP. The recommended implementation is to use the HDPS Auxiliary Copy feature. This allows backups to be transferred to secondary destinations as redundant copies which is useful for disaster recovery, data archiving, and ensuring data availability across different sites. Backup transfer can be triggered once a backup is complete or on a schedule. Deduplication Accelerated Streaming Hash (DASH) copy can also be enabled to optimize the replication process. DASH copy will allow for faster effective data transfer while maintaining data integrity.

HDPS Auxiliary Copy can be paired with HCP replication within each site. Doing so will provide protection against data unavailability from a single HCP system outage. A fully connected active/active replication topology can be implemented. This means that every HCP System will have its data replicated to all other HCP systems. In the event of an outage of the HCP system, requests to that system will be rerouted to another system to be fulfilled. This topology can have a maximum of 6 systems in it.

When configuring the HDPS storage pools, it is important to distribute the storage pools among the HCP systems. Meaning that each storage pool will be mapped to a single HCP system instead of being able to read/write directly to multiple HCP systems. Configuring the storage pools in this manner will prevent a situation where data is written to one HCP system and is requested to be read from a different system before it has been replicated throughout the topology. When this happens the request may fail and will have to be retried, which can be repeated depending on which system the request gets sent to next. Even if there is some replication lag across the replication topology requests can still be fulfilled as whole objects will be available on the HCP system each storage pool is configured to.

While not recommended, a multi-site deployment can be achieved using only HCP replication. In this configuration, each HDPS site would be configured to an HCP system and the HCP replication topology will facilitate data replication. Just like with Auxiliary Copy, it is important that every storage pool be configured to be a single HCP system. This will help avoid a scenario where data is requested from HCP, which has yet to be replicated to that system.

Using Erasure-Coded Replication to optimize HCP Storage Use

If HDPS is configured to an HCP replicated topology, storage use can be further optimized by leveraging HCP’s erasure-coded protection. With this feature, after an object is ingested into HCP, chunks for the object are calculated and then distributed among the rest of the HCP systems in the replication topology. A full copy of the object will remain on the original HCP only. This means that that HCP system can respond to requests without having to retrieve chunks to reconstruct the object. If this system were to become unavailable, other systems may fulfill requests by reading the chunks from the other systems to reconstruct the object. In this configuration the setting called erasure coding delay can be used to set the amount of time which the whole object will be available on the system which it was ingested. After that the whole object is replaced by a chunk as well saving additional storage. It is recommended to set this time to be equal to the retention time set in the HDPS storage policy as it is less likely to be read back by HDPS. In a 3 HCP system topology this configuration can store the same amount of data using have the storage compared using whole-object distribution.

Best Practices for Synthetic Full Backups

HDPS synthetic full backup is a backup type that has several benefits over full backups or full backup plus incremental. Synthetic full backups are constructed by combining the latest full backup (synthetic or initial full) with the incremental captured since, to create a new up to date full backup. Because this process only entails looking at backups already on HCP, it does not require communicating with the source device at all. This also means the backup window for the source device is limited to the time needed for an incremental, except the initial full backup.

Since synthetic full backups work by combining a previous backup with any incremental that have happened after the latest full backup, the creation of the synthetic full backup requires reading back all that data in addition to writing out the new full backup. This means that even if there are no recovery operations going on there may be more read traffic to HCP during this time than expected for taking a backup.

To mitigate the read intensiveness of performing synthetic full backups there are several things that can be done. Options include:

- Combined storage tiering can reduce the amount of read requests by ~20+%. See *Implementing Combined Storage Tier* for more details.
- For deployments leveraging multiple HCP systems with replication, assigning storage pools to use a single HCP system. See *Best Practices for Multi-Site Deployments* for more details.
- It is recommended that the frequency of synthetic full backups be equal to or greater than the retention setting in days. For example, when using a retention period of 30 days, synthetic full backups should be scheduled to occur every 30 days or more.

Implementing Combined Storage Tier

The combined storage tier feature for HDPS allows multiple storage devices to be used together. One tier, the warm tier, will contain indexing and metadata that is frequently accessed. The second cold tier contains the actual data and full copies of the indexes and metadata. For every read and write request from HDPS to HCP there are also metadata/index requests performed such as HEAD requests. The warm tier will be able to accommodate some of these requests to minimize any performance impact to the actual data requests to the cold tier HCP. When implemented with synthetic full backups you can expect up to 20% reduction in read requests coming to the HCP.

It is recommended to implement a NAS storage device, such as VSP One Block, as the warm tier while HCP is used at the cold tier. For assistance in configuring VSP One Block as a warm tier consult your Commvault sales representative.

Empty directories

Under some circumstances, container delete operations initiated by HDPS may leave objects and directories behind the HCP system. This is a known issue with HCP and is expected to be resolved in an upcoming HCP release.

Overview of Best Practices for HCP for Cloud Scale

This section summarizes the specific HCP CS configurations that are important for solution interoperability.

HCP CS configurations

Setting	Value	Notes
Object lock	enabled	If using HCP CS enabled WORM. Requires Storage Lock enabled in HDPS.

Appendix A – Troubleshooting

This section summarizes some of the common errors or issues with potential workarounds.

Error or issue	Workaround
Full restore on a windows data source will not restore in use system files. These will generate a warning.	This is **expected behavior**. A different data flow can be used to backup and restore these system files. See below links for guidance. • Full System Recovery: Windows File System Agent • Bare Metal Recovery Using 1-Touch for Windows
HDPS shows the following error: “Failed to run Data Aging job due to change in system time” and/or data is not removed from storage by HDPS	**Restart the HDPS services**. This is an issue with HDPS that arises after a de-sync between the Media Agent internal time and the system time. This can most commonly occur due to system outage or manual/automatic system time change. See Commvault community forum post Data Aging - Failed to run Data Aging job due to change in system time for additional information.

Appendix B – Resources

This section summarizes the resources available about the solution and/or used to validate the compatibility of this solution.

Hitachi Vantara resources

- Hitachi Vantara & Commvault alliance
- Documentation
  - Hitachi Content Platform (HCP)
    - See Installing an HCP System PDF for configuration and setup guidelines
  - Hitachi Content Platform for Cloud Scale (HCPCS)
    - See Installing HCP for Cloud Scale web document for configuration and setup guidelines
  - Data Protection Suite
    - See Data Protection Suite under Storage Software

Commvault resources

- Hitachi Data Protection Suite powered by Commvault
- Commvault on the Hitachi Vantara alliance
- Commvault HDPS documentation
  - See Cloud Storage configuration for instructions to configure HCP or HCPCS with HDPS
  - See Commvault Best Practices and CommCell Performance Tuning for guidance on performance parameters
- Commvault Maintenance Release 11.32.38 with fix for “Auto Sealing DDB's not acting as expected” (hotfix numbers 9900 and 9901) – access to release note requires a Commvault login

Appendix C – SRA Document History

Revision history for this document

Revision	Date	Summary
1.0.0	Feb 5, 2024	v1.0.0 of HDPS + HCP & CS v1.0 SRA (Compatibility only)
1.0.1	Mar 20, 2024	v1.0.1 of HDPS + HCP & CS v1.0 SRA (Compatibility only): ▪ updated whitepaper template; formatting changes only ▪ removed list of test cases as they are not pertinent to the best practices for the solution
1.1.0	Feb 7, 2025	V1.1.0 of Hitachi Vantara Object Storage with HDPS ● Added procedure for adding VSP One Object to HDPS ● Added best practices for configuring HCP replication with HDPS