

Data Protection with Hitachi Ops Center Protector on Cisco and Hitachi Adaptive Solutions for Converged Infrastructure

Implementation 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

Revision	Changes	Date
MK-SL-205-01	Included Cisco Intersight capabilities information	April 8, 2022
MK-SL-205-00	Initial release	August 6, 2020

Implementation Guide

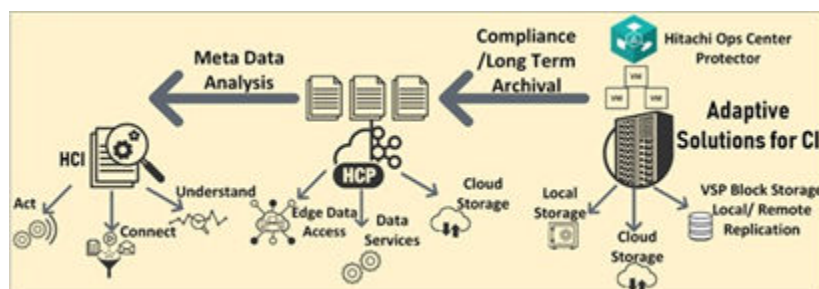
Back up virtual machines to Hitachi Virtual Storage Platform, Hitachi Content Platform, or Amazon Simple Storage Service bucket-based storage using Hitachi Ops Center Protector on a Cisco and Hitachi Adaptive Solutions for Converged Infrastructure solution.

Hitachi's proven leadership and joint innovations have accelerated enterprise IT initiatives for 80% of global Fortune 100 companies. Cisco and Hitachi Adaptive Solutions for Converged Infrastructure is a pre-validated, tested and rapidly deployable reference architecture. It's an agile data-driven foundation that supports a broad range of technologies and workloads and, when combined with continuous innovation, positions your organization to deliver better experiences and tap into new revenue streams on the same adaptable infrastructure solution provided by Hitachi and Cisco.

This solution and feature sets relate to Hitachi Ops Center Protector (formerly Hitachi Data Instance Director), Hitachi Content Platform, and Hitachi Content Intelligence on top of Adaptive Solutions for Converged Infrastructure. Also, these procedures may be utilized on Hitachi converged and hyperconverged infrastructure supporting VMware environments, such as the converged Hitachi Unified Compute Platform CI series, the hyperconverged Unified Compute Platform HC series, and the rack scale Unified Compute Platform RS series.

Ops Center Protector provides enterprise copy data management capabilities for VMware environments, using automated intelligence to provide a layered SLA-driven modern data protection schema. It provides a managed snapshot offload- based solution for granular virtual machine backup, cloning and recovery, and automated management of multi-datacenter replications. It supports VMware vSphere Storage APIs-based backup. It integrates with VMware vSphere tags to provide tag-controlled backup and copy data management capabilities.

The following figure provides a Hitachi Ops Center Protector, Hitachi Content Platform, and Hitachi Content Intelligence capability overview.



Use these procedures when using Hitachi Ops Center Protector local data protection services in conjunction with VMware on top of Hitachi Content Platform object storage and Hitachi Content Intelligence. Other Hitachi replication technologies will be mentioned in this guide, but not covered. These procedures were tested and validated on Cisco and Hitachi Adaptive Solutions for Converged Infrastructure with VMware to do the following:

- Bring the ability to protect your critical information
- Manage data growth
- Regulate compliance initiatives
- Build a private or hybrid cloud infrastructure

Utilizing Ops Center Protector with Hitachi Content Platform and Hitachi Content Intelligence facilitates search and discovery operations by analyzing, transforming, and indexing valuable information into actionable business information by improving data quality and increasing data awareness. This brings more insight into your business for increase operational efficiencies.

For more information about validated solutions using Cisco Unified Compute System (UCS) and Hitachi Virtual Storage Platform (VSP), see [Related documents \(on page 9\)](#).

This document is intended for the following:

- Storage administrators
- VMware administrators
- Sales engineers
- Field consultants
- Professional services staff
- Validated Hitachi and Cisco resale partners

Readers of this document should have a background in or understanding of the following:

- RAID systems and their functions
- VMware ESXi and VMware vCenter environments
- Converged infrastructures



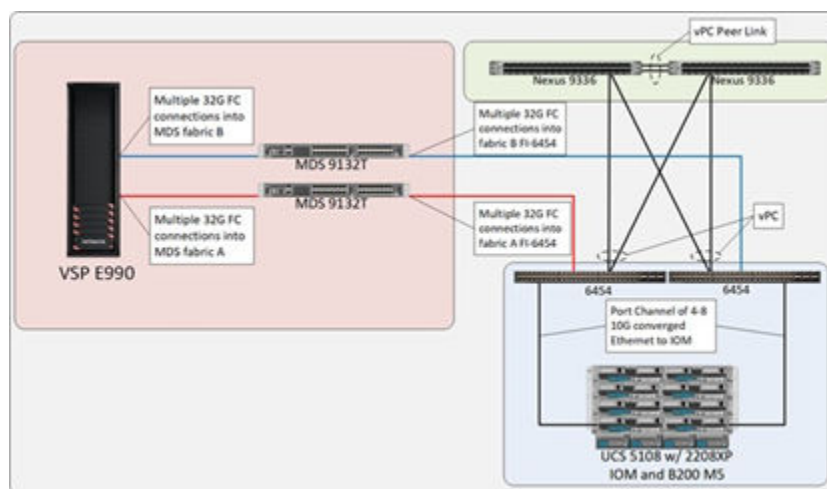
Note: Testing of these procedures was in a lab environment. Many things impact production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that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

Cisco Unified Compute System environment

Cisco and Hitachi Adaptive Solutions for Converged Infrastructure as a virtual server infrastructure is a best-practice datacenter architecture built on collaboration between Hitachi Vantara and Cisco to meet your enterprise needs using virtual server workloads.

This architecture uses Hitachi Virtual Storage Platform connecting to Cisco MDS multilayer switches that link to the Cisco UCS Fabric Interconnects and Cisco Unified Computing System (UCS) chassis. Northbound networking is enabled through the Cisco Nexus 9000 family of switches.

The following figure is the validated architecture for Cisco and Hitachi Adaptive Solutions for Converged Infrastructure.



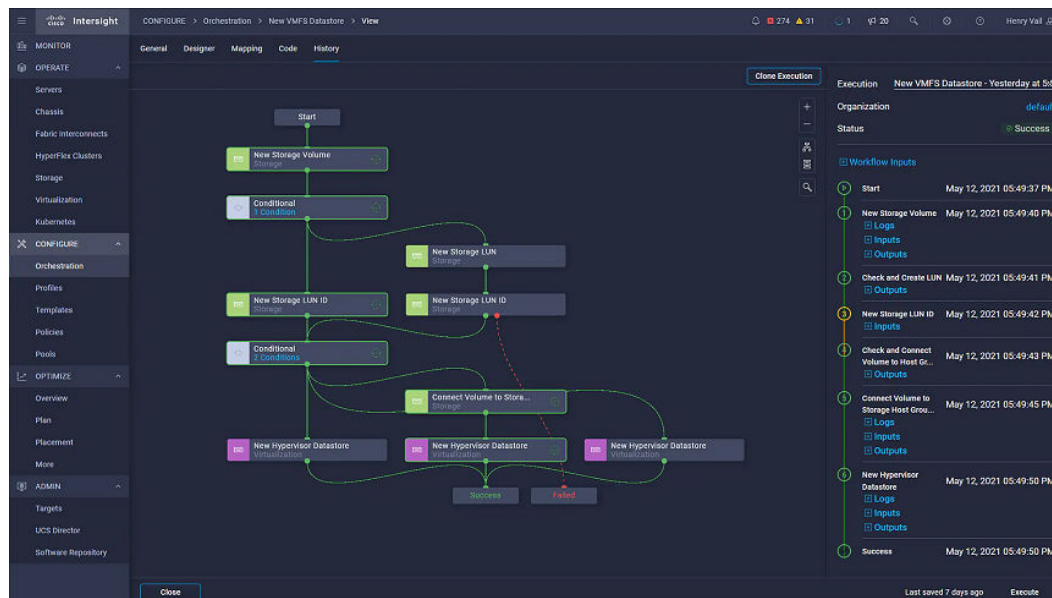
For more information about Cisco UCS hardware and software versions, refer to the [Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Design Guide](#).

Cisco Intersight Capabilities with Hitachi Virtual Storage Platform

Hitachi has enabled a magnitude of storage management capabilities that will now be able to be done using Cisco Intersight with the goal of saving administrators time and frustration.

Within the Cisco Intersight management platform, administrators can use the concept of tasks and workflows to easily manage their hybrid IT environments.

Tasks are essentially a library of functions that leverage API invoke calls that can be customized, or they can be provided by Cisco out of the box. These tasks can be compiled to create workflows to enable quick and easy automation of infrastructure without requiring code experts. This provides true single pane of glass orchestration through Cisco Intersight, reducing the need for datacenter administrators to host multiple screens to complete functions.



The following tables show the current capabilities of Hitachi Virtual Storage Platform (VSP) in orchestration with tasks and workflows provided by Intersight to end users.

Table 1 List of support tasks for Hitachi VSP

Tasks	Hitachi VSP
Compress Storage Pool	Y
Connect Initiators to Storage Host	Y
Connect Volume to Storage Host	Y
Copy Storage Volume	Y
Disconnect Initiators from Storage Host	Y
Disconnect Volume from Storage Host	Y
Edit Storage Pool	Y
Expand Storage Volume	Y
Expand Storage Pool	Y
Format Storage Volume	Y
New Storage Host	Y
New Storage Pool	Y
New Storage Volume	Y
Remove Storage Host	Y
Remove Storage Pool	Y
Remove Storage Volume	Y

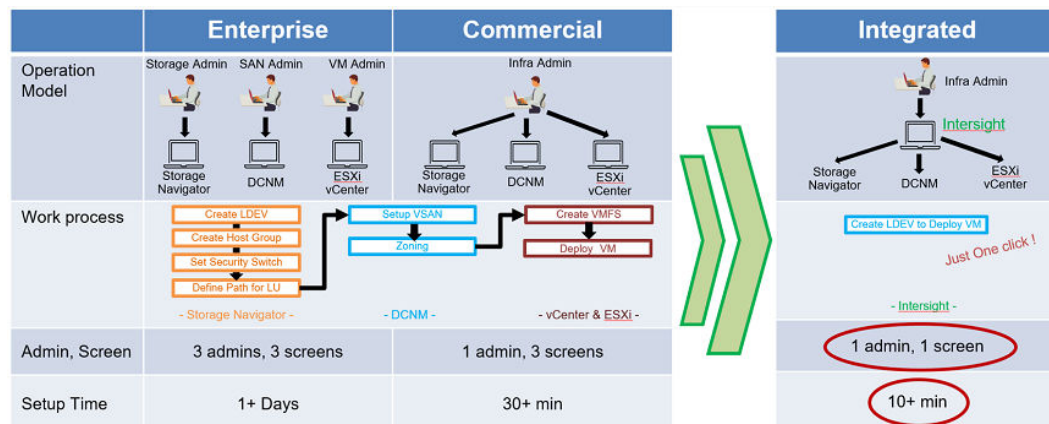
Table 2 List of supported workflows for Hitachi VSP

Storage Workflows	Hitachi VSP
New Storage Host	Y
New VMFS Datastore	Y
Remove Storage Host	Y
Update Storage Host	Y

With these capabilities administrators can complete a majority of day 0 to day N tasks to support their hybrid IT environment with Hitachi VSP storage systems.

A related reference architecture has been published as part of Cisco and Hitachi Adaptive solutions, at [Hitachi Virtual Storage Platform with Cisco Intersight Reference Architecture Guide](#). See [#unique_8](#) for more information.

This reference architecture explores the benefits of integrated management with Cisco Intersight compared to conventional methods using multiple management interfaces. When creating a virtual environment for enterprise workloads with Cisco Intersight with VSP integration, on average 50 hours of time is saved over the course of a year and 80% fewer screens are required to complete such operations, as shown in the following figure.



Hardware versions

The following table lists the hardware used to develop these procedures. Alterations can be made according to Hitachi's and Cisco's hardware compatibility lists.

Table 3 Hardware versions used for validation

Component	Version
Hitachi Virtual Storage Platform E990	93-02-01-60/00
Cisco MDS 9132T Fibre Channel switch	8.4(1a)
Cisco Nexus 9332-FX2 switch	NXOS7.0(3)I7(8)
Cisco Fabric Interconnect 6454	4.1(1b)
Cisco Unified Computing System B200 M5 Blade Servers	4.1(1b)
Cisco Unified Computing System 2208XP IOM	4.1(1b)

Software versions

The following table lists the software used to develop these procedures.

Table 4 Software versions used for validation

Component	Version
Hitachi Ops Center Protector	v7.0.0.80503-R7.0 or Newer
Hitachi Content Platform for VMware vSphere ESXi	v8.3.2.5 or Newer
Hitachi Content Intelligence	v1.6.1.27 or Newer
Command control interface (CCI)	01-56-03/03 or Newer
VMware ESXi 6.7 U3	6.7.0.14320388 or Newer
VMware vCenter Standalone (VCSA) 6.7 U3	6.7.0.41000 or Newer
VMware ESXi 6.7U3 nenic	1.0.31.0 or Newer
VMware ESXi 6.7U3 nfnic	4.0.0.48 or Newer

Release notes

Read the release notes before installing and using any of these products. They contain requirements or restrictions that are not fully described in this document, or updates, or corrections to this document. Release notes are available on the [Hitachi Vantara Knowledge portal](#).

Related documents

This document references the following documentation:

- [Hitachi Ops Center Protector Documentation](#)
- [Hitachi Content Platform Documentation](#)
- [Hitachi Virtual Storage Platform E990 Documentation](#)
- [Hitachi Content Intelligence Documentation](#)
- [Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Design Guide](#)
- [Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Deployment Guide](#)

Getting Help

Hitachi Vantara Support is the destination for technical support of products and solutions sold by Hitachi Vantara.

- To contact technical support, log on to Hitachi Vantara Support Connect for contact information, see [Customer Contact Us](#).
- To open a new support case, see [How to Create a New Case on the Support Website](#).

Why Hitachi?

Hitachi, Ltd. (TSE: 6501), headquartered in Tokyo, Japan, delivers innovations that answer society's challenges, combining its operational technology, information technology, and products/systems. The company's consolidated revenues for fiscal 2017 (ended March 31, 2018) totaled 9,368.6 billion yen (\$88.4 billion). The Hitachi Group is an innovation partner for the IoT era, and it has approximately 307,000 employees worldwide. Through collaborative creation with customers, Hitachi is deploying Social Innovation Business using digital technologies in a broad range of sectors, including Power/Energy, Industry/Distribution/Water, Urban Development, and Finance/Social Infrastructure/Healthcare. For more information on Hitachi, visit [Hitachi's website](#).

Founded in 1910, Hitachi has expanded social, economic, and environmental solutions for decades. Spanning eight different industries, Hitachi differentiates from your typical IT provider by blending information technology (IT) and operational technology (OT) to realize insights and advantages from a single partner. As an innovative global conglomerate Hitachi is poised to position solutions spanning beyond the datacenter floor and into industry solutions which meet and exceed the goals of our customers. Hitachi is an enterprise player in the technology solutions we offer, but our true value is shared by the business outcomes that only an IT and OT company can achieve. Our customers can leverage extensive IT products that drive business outcomes through resilient and agile IT infrastructure. So, no matter the problem, Hitachi is built to support all customers' industries in their quest to build a better tomorrow.

Hitachi Vantara, a wholly-owned subsidiary of Hitachi, Ltd., guides our customers from what's now to what's next by solving their digital challenges. Working alongside each customer, we apply our unmatched industrial and digital capabilities to their data and applications to benefit both business and society. More than 80% of the Fortune 100 trust Hitachi Vantara to help them develop new revenue streams, unlock competitive advantages, lower costs, enhance customer experiences, and deliver social and environmental value.

Hitachi Virtual Storage Platform E990

Hitachi provides an end to end portfolio of storage that satisfies small, large, and enterprise business demands. With the latest refresh, Hitachi has introduced the all-flash Hitachi Virtual Storage Platform E990, which bridges the gap between enterprise performance and at a modest small form factor price point. Paired with Hitachi Storage Virtualization Operating System RF (SVOS RF), which is optimized for NVMe flash systems, provides unmatched latency and IOPs.

Hitachi also uses a single operating system across its entire storage portfolio to enable you to manage and replicate your data seamlessly across various generations of storage systems. Hitachi enables you to get maximum ROI on legacy systems with the consolidation of Hitachi and third party systems virtualized behind Virtual Storage Platform E990 to provide the following:

- The benefit of Hitachi performance with deduplication and resilience.
- The ability to extend NVMe front-end performance to older, third-party storage.
- Simplified management with a single management interface.
- Modernized legacy systems with data reduction and extended service life.
- Improved agility with simplified migrations.
- Flexibility of being able to use Fibre Channel or iSCSI.

Virtual Storage Platform E990 is powered by Hitachi Ops Center, which provides you an artificial intelligence-driven management solution that provides the following:

- Simplified storage provisioning, analytics, and protection for artificial learning or machine learning.
- Containerized applications all through single pane of glass management.

Advantages of NVMe storage with Hitachi Virtual Storage Platform E990

While competitors have access to similar multicore processors and use the same NVMe low-latency command set and multiple queues per device for fast concurrent I/O, Hitachi Virtual Storage Platform E990 is differentiated by the efficiency of Hitachi Storage Virtualization Operating System. Containing only two controllers with 56 cores and 1 TB of cache memory, Virtual Storage Platform E990 has the lowest latency and double the maximum IOPS of the other five vendors when this document was published. Fewer cores and memory reduce power, cooling, and cost.

With these enhancements, Virtual Storage Platform E990 provides 99.9999% availability along with the following:

- Lower drive counts with NVMe drives which outperform SAS and SSD drives in IOPs, the number of gigabytes per second, and response time.
- NVMe storage provides higher read-intensive workloads.
- An 80% reduction in rebuild time for flash drives.
- Latency-sensitive host applications to stay up by avoiding write through mode.

Virtual Storage Platform E990 offers a future-proof solution that allows you to get more value from your investment.

Hitachi EverFlex consumption

Hitachi Vantara offers the EverFlex program to provide predictable pricing and flexible usage if you who do not wish to own and maintain a on-premises storage system. With the EverFlex program, expect the following:

- Aligned technology spending with usage and business outcomes, to reduce risk of over- or under-investing in the IT solutions that run your enterprise. Scale up or scale down as business demands change to keep IT costs in line with business needs.
- Eliminate dependence on accurate forecasts for capacity planning to achieve your efficiency targets and lower your effective asset acquisition costs.
- Reduce the cost of IT operations and improve service level delivery in terms of consistency and quality for the services your business relies on.
- Backed by Hitachi's 100% data availability guarantee.

With the EverFlex program, you now have another option to make the correct investment decisions to drive your business outcomes. For more information on the EverFlex program, visit [EverFlex from Hitachi Vantara](#).

Hitachi Ops Center Protector with VMware data protection capability matrix

The following table lists the capabilities of Hitachi Ops Center Protector used within the tested VMware virtual environment.

Table 5 Hitachi Ops Center Protector capabilities

Configur ation	Host Based Backup		Hitachi Block			Tier	Cloud
	Batch	Realtime /CDP	Snapsho t (Hitachi Thin Image)	Live Replicati on (SI, TC, HUR, GAD)	Batch Replicati on (HTI, SI)	Hitachi Content Platform	Amazon S3
VMware vCenter (Multiple ESXi)	X		X	X	X	X	X
VMware ESXi Server (Stand- alone)	X		X	X	X	X	X
Path (Client)	X	X				X	X

Solution components

These are the components used to implement these procedures.

Hitachi Ops Center Protector

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:

- [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.

This guide only covers Ops Center Protector data protection in conjunction VMware virtual environments. Protector enables you to take advantage of a range of data protection technologies while simplifying administration. It offers a unique policy-based copy data management workflow engine that allows you to combine the right tools to support complex, end-to-end data copy, recovery, retention, and repurposing requirements from a single solution.

Hitachi Ops Center Protector client software must be installed on each server participating in the data protection environment. Deploy the following node types when using Protector.

- Master Node is required per each instance of Ops Center Protector. This provides the user interface and central management of data protection capabilities.
- OS host nodes are automatically added to the inventory. These represent every server that has Ops Center Protector master or client software installed on it. A basic operating system host node can only be used as a file system data source or destination as a proxy repository using data flows.
- Repository nodes require an Ops Center Protector client with access to suitable disk storage capacity to act as a host. Repository nodes act as data destinations on data flows performing backup and continuous data protection operations.

Repositories are implemented entirely by software processes within the Ops Center Protector client. Depending on the backup data change rate and number of policies, servers that host repositories can be subject heavy network, processor, and disk I/O loading. It is important to specify the correct hardware and monitor performance.

- Hitachi Content Platform nodes do not require an Ops Center Protector client to be specified. They act as data destinations on data flows but must be accessed using a repository acting as a cache when performing tiering operations.
- VMware nodes do not require an Ops Center Protector client and are onboarded with IP. These nodes act as data sources on data flows, enabling users to specify VMware-specific policies and backup operations to repository or block storage destination nodes. During disaster recovery, VMware nodes are also destinations.
- Block device nodes require an Ops Center Protector client with block prerequisites installed to act as ISM block device nodes to be data sources and destinations on data flows. This enable you to specify block-specific policies with snapshot and replication operations to other block device nodes.
- Block host and logical block nodes can only be created by specifying the block device nodes upon which they are based. These nodes represent a subset of the resources available on the block device node upon which they are based.



Note: Block host and logical block nodes can only be used as data sources.

- Amazon Simple Storage System (S3) nodes provide a destination for batch backup of application, virtual machines, and file systems.
- Microsoft® Exchange, Microsoft SQL Server®, Oracle DB, and SAP HANA nodes require an Ops Center Protector client to be installed on the application server. Appropriate prerequisites must also be installed to enable Ops Center Protector to interact with the application.

These nodes act as data sources on data flows, enabling you to specify application-specific policies and backup operations to repository or block storage destination nodes.

The following is a brief overview of replication technologies offered using Hitachi Ops Center Protector. However, only local data protection services for VMware have been validated in this guide. For more information on Hitachi Ops Center Protector Installation and other local disaster recovery and remote replication technologies, see [Related documents \(on page 9\)](#).

VMware vSphere Storage APIs – Data Protection for LAN and SAN Backups

Hitachi Ops Center Protector provides data protection for VMware hypervisor environments by using VMware vSphere Storage APIs – Data Protection regardless of the datastore residing on Hitachi block storage.


Software protection using VMware vSphere Storage APIs – Data Protection provides a cost-sensitive approach with lower service levels, suitable for less-critical workloads and long-term retention. With VMware vSphere Storage APIs – Data Protection over LAN or SAN virtual machines on VMware ESXi hosts use the APIs to perform backups to a local Protector repository or directly to the storage system.



Note: Using VADP SAN backups users must use a physical repository with datastore access for SAN transport mode.

Ops Center Protector now offers a multi-streaming for VMware vSphere Storage APIs to allow users to specify how many virtual machines can be streamed concurrently to provide more effect use of bandwidth to complete batch backup operations in less time. With multiple data streams per client for data protection, the contents can be distributed to all backup streams transmitting the data in parallel to the storage media.

The following figure shows the data protection multi-stream capability.



Native Virtual Machine Disk Backups to cloud for Amazon S3 and Hitachi Content Platform

Hitachi Ops Center Protector offers native agentless virtual machine disk (VMDK) backups to data targets such as Amazon S3 and Hitachi Content Platform. Take advantage of VMware's native virtual hard disk format to allow seamless restores of virtual machines from multiple data targets to provide greater flexibility to administrators facing disaster recovery.

Hitachi customers can further ingest these native formats into Hitachi Content Intelligence, a framework solution that allows exploration of critical business data to unlock insights to make better business decisions.

Hitachi in-system heterogeneous replication bundle

Paired with all Hitachi Virtual Storage Platform, there is in-system replication capabilities to protect local datasets within the same storage system. With Hitachi Storage Virtualization Operating System paired with Hitachi Ops Center Protector, you can now take a visual, comprehensive, and automated approach to protecting virtual and physical datasets.

Hitachi ShadowImage® Clone


Hitachi ShadowImage® is part of the Hitachi In-System Replication bundle. It is also included in the foundation software package of many Hitachi Virtual Storage Platform systems.

Hitachi Ops Center Protector adds application-awareness and full automation to ShadowImage, alleviating replication management headaches. With ShadowImage, Protector creates virtual volumes of thin provisioned LDEVs created with Hitachi Dynamic Provisioning. The ShadowImage virtual volumes consume the same capacity as the P-VOL in a 1:1 ratio.

The following figure is a representation of a ShadowImage clone taking up 1:1 capacity.

Hitachi ShadowImage (replication software)

All data is saved from P-VOL to S-VOL



Consistent read/write access is available only in split states

ShadowImage cloning operations are independent of operating system, application, or device, allowing efficient and centralized replicated storage volume management. Application-consistent ShadowImage clones can be orchestrated using Protector.


Administrators can use Hitachi ShadowImage® technology to provide full copy block clones to recover locally from a disaster. ShadowImage clones provide best performance for local replication cloning as they do not operate on a common base. Backups are taken in a 1:1 capacity and can branch from and provide local full clones for remote replication branches.

Hitachi Thin Image snapshot

A Hitachi Thin Image snapshot rapidly creates up to one million point-in-time copies of mission-critical information within any Hitachi Virtual Storage Platform or virtualized storage pool without impacting host service or performance levels.

Because snapshots store only the changed data, the volume of storage capacity required for each snapshot copy is substantially smaller than the source volume. As a result, a Thin Image snapshot can provide significant savings over full volume cloning methods.

The following figure has Hitachi Thin Image-based snapshots showing data along with differential point in time virtual volumes.



Hitachi Thin Image snapshot copies are fully read/write compatible with other hosts and can be used for system backups, application testing, and data mining applications while your business continues to run at full capacity. With Thin Image, you can have multiple point-in-time copies of mission-critical data. Thin Image copies can primarily be used as test development copies and can branch from a Hitachi ShadowImage® clone along with remote replication technologies.

To learn more about Hitachi Thin Image features and benefits, see [Related documents \(on page 9\)](#).

Hitachi remote replication

If you require multisite capabilities, Hitachi Vantara offers replication technologies that can provide multisite data protection capabilities along with storage high availability for mission critical applications. With Hitachi Ops Center Protector, spend less time performing setup and deploy any remote replication technology.

Hitachi TrueCopy®

Hitachi TrueCopy® provides a remote mirror of any data with synchronous replication. The remote copy is always identical to the local copy and allows for fast restart and recovery of applications and datasets with no data loss. TrueCopy has no dependency on host operating systems or databases for file systems. Distance limit is a variable, but typically is less than 180 miles (290 km).

The following figure shows the round-trip of block data in conjunction with the Hitachi TrueCopy® primary volume (P-VOL) and secondary volume (S-VOL).




Hitachi Universal Replicator

Hitachi Universal Replicator provides added protection in case of a disaster with the ability to provide replication at an extended distance. With Universal Replicator, local and secondary journal caches are used to unload and transfer data to the secondary storage system. Records are written to the primary disk and metadata and copy data is placed in the journal cache which is then transferred to the secondary journal and is offloaded to the secondary volume at the secondary site.

In case of an outage, Universal Replicator uses a first in first out (FIFO) mechanism to process and commit replication data so that backups are consistent.

With Universal Replicator, the primary system has no overhead. This lets the system handle additional production processing.

The following figure depicts Hitachi Universal Replicator technology.

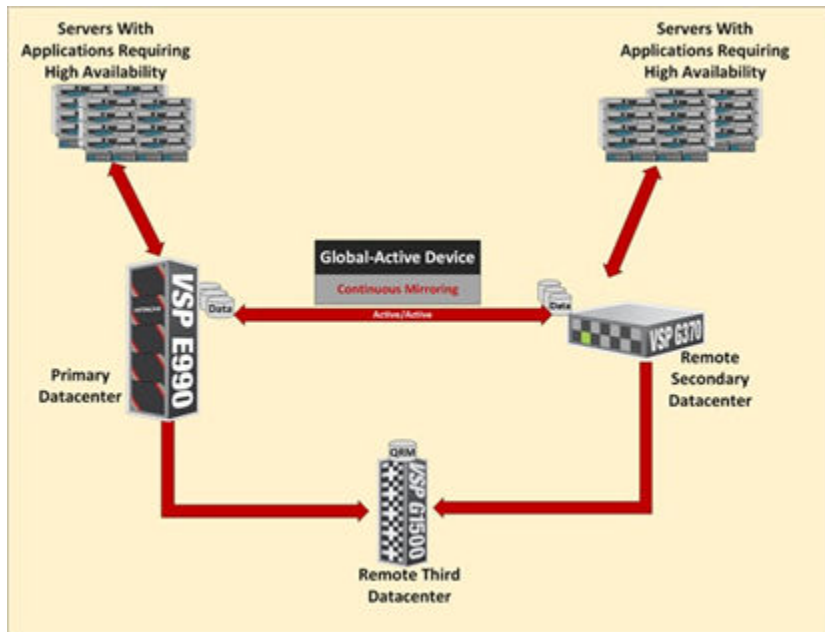


Global-active device

Global-active device enables you to create and maintain synchronous, remote copies of data volumes. Configure a virtual storage machine (VSM) in the primary and secondary storage systems using the actual information of the primary storage system, and the global-active device primary and secondary volumes are assigned the same virtual LDEV (logical device) number in the virtual storage machine. This enables the host to see the pair volumes as a single volume on a single storage system, with both volumes receiving the same data from the host.

A quorum disk, which can be in a third and external storage system or in an iSCSI-attached host server, monitors the global-active device pair volumes. The quorum disk acts as a heartbeat for the global-active device pair, with both storage systems accessing the quorum disk to check on each other. A communication failure between systems results in a series of checks with the quorum disk to identify the problem so the system can receive host updates.

The following figure shows a fully deployed high availability scenario with Cisco and Hitachi Adaptive Solutions for Converged Infrastructure.



To learn more about storage high availability with global-active device and Cisco and Hitachi Adaptive Solutions for Converged Infrastructure, see [Cisco and Hitachi Adaptive Solutions for Converged Infrastructure as a Stretched Data Center](#).

Hitachi Content Platform

Hitachi Content Platform is a massively scalable, multi-tiered, multi-tenant object store. This object store can be divided into thousands of virtual content platforms: Tenants and underlying namespaces have configurable attributes to deliver varying service levels for different users and applications. Network file system (NFS), common internet file system (CIFS), and representational state transfer (REST) protocols such as Amazon's S3 are supported, as well as Microsoft Active Directory® authentication.

Each piece of content is stored as an object, which is a container that includes the data and metadata used to define the structure and administration of that data. This provides IT with a deep understanding of the nature of the content and enables IT to assign policies and automate storage tiering with greater intelligence. Hitachi Content Platform can automatically apply data retention and disposition, deleting expired content and reclaiming storage. Content is accessed through HTTP/REST APIs like Amazon's S3, NFS, CIFS, SMTP, and more. Monitoring, reporting, and audit capabilities are built in and enable chargeback.

Hitachi Content Platform provides the ability to regulate and provide compliance with sensitive data that is required by governments and regulation entities.


Hitachi Content Platform for VMware vSphere

Hitachi Content Platform comes in a variety of deployment methods. Content Platform was validated on Cisco and Hitachi Adaptive Solutions for Converged Infrastructure on virtual server infrastructure. For more information about deployment types of Hitachi Content Platform, see [Related documents \(on page 9\)](#).

Hitachi Content Platform on a virtual machine storage infrastructure is highly available and fault tolerant. It is recommended for the physical servers that the VMware vSphere ESXi hosts run on are connected to shared SAN storage with RAID-6 protection. SAN storage for Hitachi Content Platform needs to have at least two paths to each logical unit number (LUN). Each LUN needs to have the same LUN number on each ESXi host. A datastore is created from each LUN or export, creating one VMware Virtual Machine File System (VMFS) volume per LUN or export.

A single datastore is not shared by Hitachi Content Platform on virtual machine nodes. However, Content Platform on virtual machine nodes can have multiple datastores. Each datastore is carved into one or multiple virtual disk files (VMDK), which are presented to the Hitachi Content Platform operating system as local disks. The Hitachi Content Platform operating system recognizes its storage as internal drives. The disks are controlled by the VMware Paravirtual SCSI controller (PVSCSI). VMware recommends PVSCSI for better overall performance.

The following figure describes the storage and network architecture of Hitachi Content Platform.



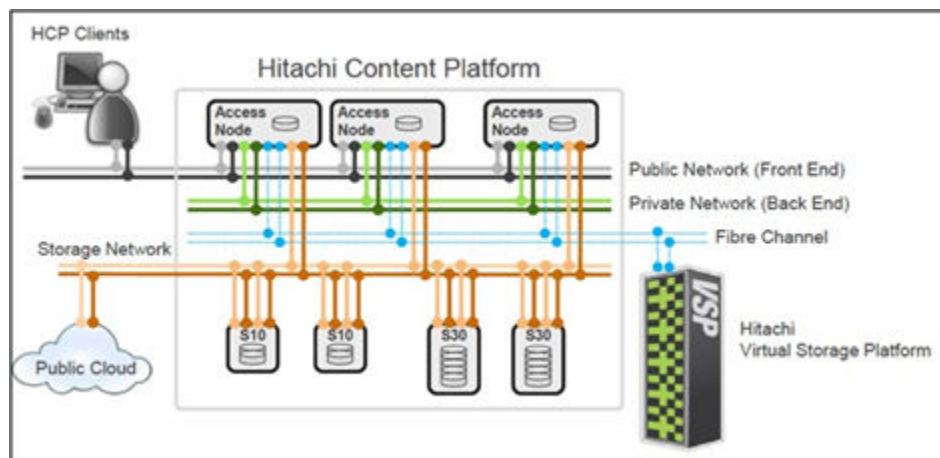
Hitachi Content Platform on a virtual machine network connectivity is provided to the Hitachi Content Platform guest operating system by VMware VMXNET3 or e1000 virtual network interface cards (vNICs), each VMware vNetwork Standard Switch (vSwitch), and each VMware vNetwork Distributed Switch (dvSwitch or vDS). The recommendation is to connect vNICs to a single vSwitch for back-end connectivity and a single vSwitch for front-end connectivity.

The Hitachi Content Platform front-end network is used for client and management access. For Content Platform front-end networks, the recommendation is to have the ESXi host create two vNICs on a second pair of physical NICs. Having two physical NICs dedicated to Content Platform ensures redundancy and consistent performance.

The Hitachi Content Platform private back-end network is used for internode communication and data transfer. The ESXi host has two virtual machine NICs which directly map to two physical NICs on the ESXi host server. The physical NICs dedicated to the back-end network must be connected to two physical switches on an isolated network. Physical NIC1 on all ESXi hosts must connect to the same physical switch, and Physical NIC-2 on all ESXi hosts must connect to the same second physical switch. The physical switches must be cabled for an inter-switch connection.

To guarantee data security and Hitachi Content Platform reliability, back-end switches must be configured with spanning tree disabled and multicast traffic enabled. The back-end switches must be at least 1 GbE and dedicated to Content Platform. Alternatively, Content Platform allows the utilization of S10 and S30 nodes as an appliance to expand object storage capacity independent of compute.

The following figure shows the frontend, backend, and Fibre Channel networks of Hitachi Content Platform.



To support Hitachi Content Platform on a virtual machine inter-node communication, enable multicast for the back-end network.



Note: For VMware environments using vDS, the multicast filtering mode on the switch must be set to IGMP/MLD snooping.

In most cases, enabling multicast on the switch is not enough to allow for multicast traffic. To allow multicast traffic between the Hitachi Content Platform on a virtual machine nodes switches, the following is required:

- Base additional configuration parameters from vendor specifications.
- If the Content Platform on a virtual machine back-end network is on a public network, have the Content Platform on a virtual machine system reside on its own VLAN.
- Configure the vDS switch to support multicasting for Content Platform backend communication. For information on how to deploy Hitachi Content Platform, see [Related documents \(on page 9\)](#).

Hitachi Content Intelligence

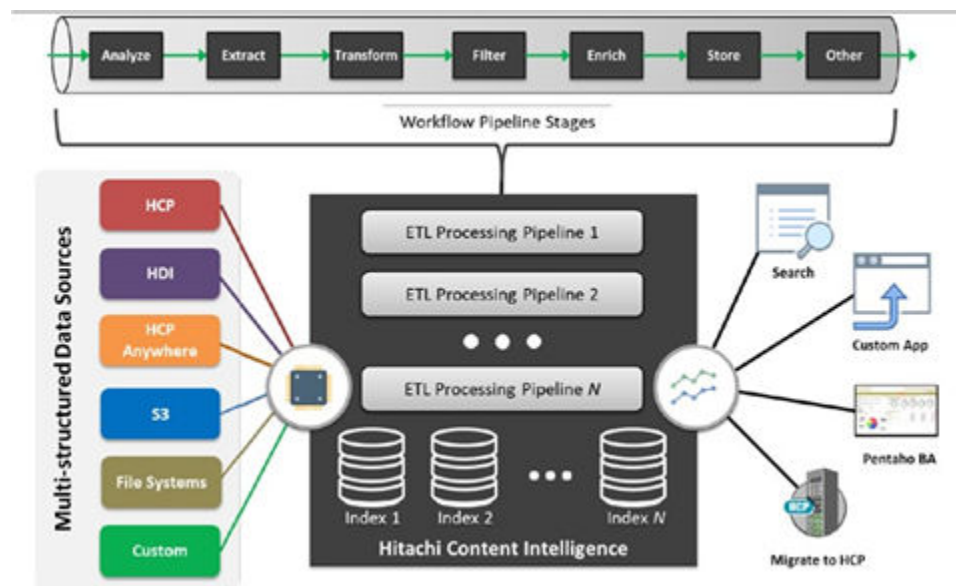
Hitachi Content Intelligence is a data collection, aggregation, processing, normalization, indexing, and analysis software suite. It implements processing as workflows, each consisting of the following:

- Data connections for the sources that are crawled
- Processing pipelines to extract, transform, and enrich the data
- Index collections to store the results

As shown in the following figure, a single Hitachi Content Intelligence deployment can handle multiple sources, processed through multiple pipelines, for multiple downstream users, including Hitachi Content Intelligence's included search capability.

Hitachi Content Intelligence search provides users with a powerful tool to perform analytics directly with the Hitachi Content Intelligence suite.

The following figure shows the Hitachi Content Intelligence data architecture.



The heart of Hitachi Content Intelligence capabilities is the ETL processing pipeline. Each pipeline can transform and enrich the selected files and objects. Pipelines are comprised of multiple stages executed serially and can be conditionally controlled. Hitachi Content Intelligence provides prebuilt stages while users can augment processing with their own custom stage plug-ins.

For information on how to deploy Hitachi Content Intelligence, see [Related documents \(on page 9\)](#).

VMware data protection best practice considerations

VMware vSphere, a leader in enterprise virtualization, is used by many enterprises to support their applications and workloads. It is highly recommended to have plans in place in case of a disaster, whether physical or virtual. To provide effective backups, administrators must consider the following:

- Are backups in native format?
- Can backups be stored in multiple locations?
- Are there multiple versions of critical data?
- Are backups hardware agnostic?
- Are backups complete copies?
- Is there the ability to do partial restores?
- Is there the ability to encrypt and replicate data?
- Are physical and virtual assets able for protection?
- Is data in compliance with government and local jurisdiction?

Hitachi Vantara's data protection solutions bring all these capabilities to your organization and enables administrators to not only follow best practices but also provide additional compliance capabilities.

Hitachi Unified Compute Platform CI, Unified Compute Platform HC, and Unified Compute Platform RS Series environments

Hitachi provides converged, hyperconverged, and rack-scale systems to handle a variety of enterprise-class workloads within one single architectural foundation. Hitachi Unified Compute Platform comes pre-tested and validated to ensure easy adoption and resilient platform awareness. It can solve complex business solutions on an agile and future-proof architecture.

To learn more about the supported compute, network, and storage configuration of Unified Compute Platform systems, refer to [Hitachi Vantara's interoperability matrix](#).

Solution Implementation

Here is how to implement data protection best practices in your Cisco and Hitachi Adaptive Solutions for Converged Infrastructure.

Using Hitachi Ops Center Protector

This is how to implement Hitachi Ops Center Protector.

Access Hitachi Ops Center Protector

After deploying Hitachi Ops Center Protector, open the user interface to set up policies and workflows. For information on how to install Ops Center Protector and its components, see [Related documents \(on page 9\)](#).

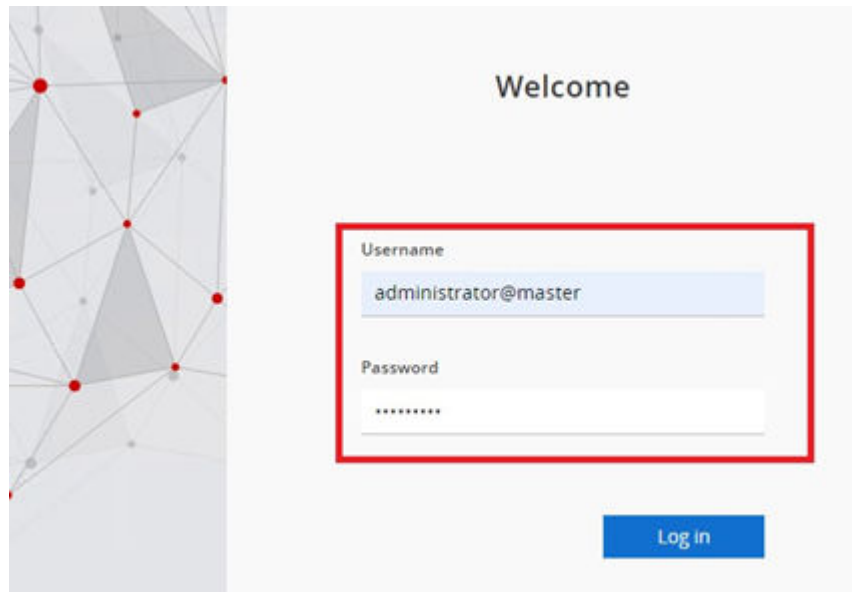


Note: Initial login to the master node uses the default credentials used when installing Ops Center Protector to the respective operating system.

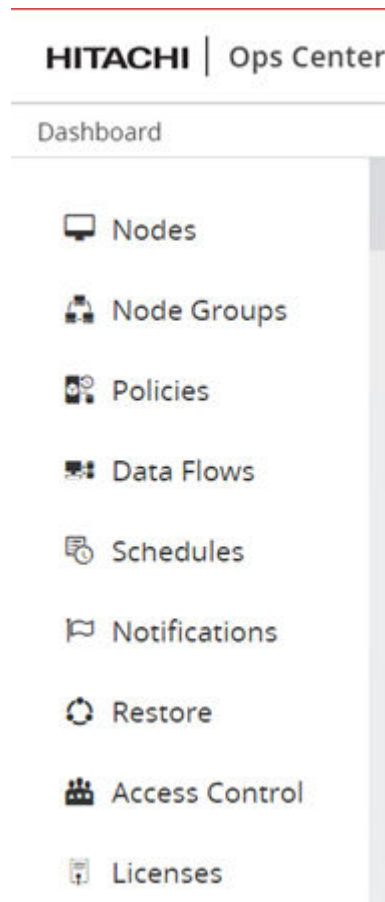
To access Ops Center Protector, do the following.

Procedure

1. Open a supported browser and open this URL: `https://Protector-Master-IP:443/`
2. Log on using the operating system credentials that were used during installation. For example, `administrator@master`.



When logged on, you will be presented with your dashboard, along with your navigation tree (the following figure). The navigation tree will be used throughout this document.



Authorize nodes

Prior to using Hitachi Ops Center Protector to create policies and data flows, authorize nodes, such as clients and proxy machines.

To authorize Ops Center Protector to utilize nodes, do the following.


Procedure

1. Select the Nodes container from your navigation tree.

An unauthorized node has a warning symbol on the top right corner of the node (the following figure).



2. Select the node or nodes that you want to authorize.
3. Scroll to the top of the page and, from the **Node Action** menu, select **Authorize** (the check mark). See the following figure.



Once authorized, you can use a node when creating policies and constructing data flows.

Create nodes

This is how to create a node. You must onboard your source and destination nodes to be able to protect your virtual or physical environment with Hitachi Ops Center Protector.

Client nodes

Hitachi Ops Center Protector client nodes, which can be either physical or virtual, to allow application-aware backups supporting Microsoft Exchange, Microsoft SQL Server, SAP HANA, and Oracle databases. Additionally, with the Ops Center Protector agent, you can perform file system-level protection on various operating systems.

For client nodes, install the Ops Center Protector agent. For directions on installing the Protector agent on your operating system, see [Related documents \(on page 9\)](#).



Note: Once onboarded, all client nodes must be authorized prior to creating policies and data flows.

Generation 1 and 2 repository nodes

Repository nodes are potential destination nodes which store data from one or more source nodes. On the machine where the repository is defined, mount and onboard a LUN to Hitachi Ops Center Protector using that directory.

Ops Center Protector introduces Generation 2 repositories that enable further features for administrators who utilize a virtual environment such as VMware. Protector still supports Generation 1 legacy repositories.

- Generation 2 supports batch backups directly from source nodes or by other Generation 2 repository nodes, Amazon S3 or Hitachi Content Platform nodes. Deduplication can be enabled repository-wide, supporting parallel backups from VMware nodes.
- Generation 1 supports batch backup directly from source nodes or by other Generation 1 repository nodes. Data can be tiered to Generation 1 Hitachi Content Platform nodes. Deduplication can be enabled on a per-store basis.

The repository node can be created on any machine that has the Ops Center Protector agent installed. This guide shows creating a virtual machine-named proxy created with Microsoft Windows Server 2016.

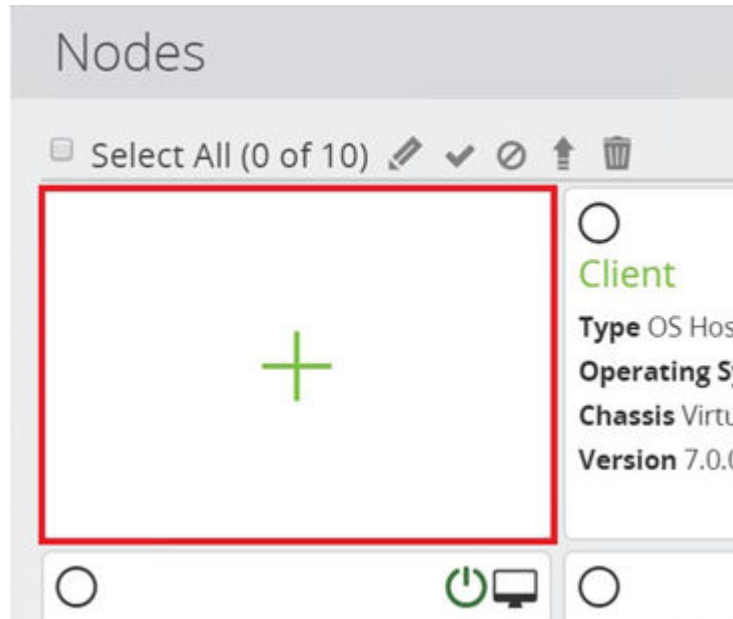


Note: Enabling encryption will require repositories to be mounted manually when the proxy is restarted. This will have overhead.

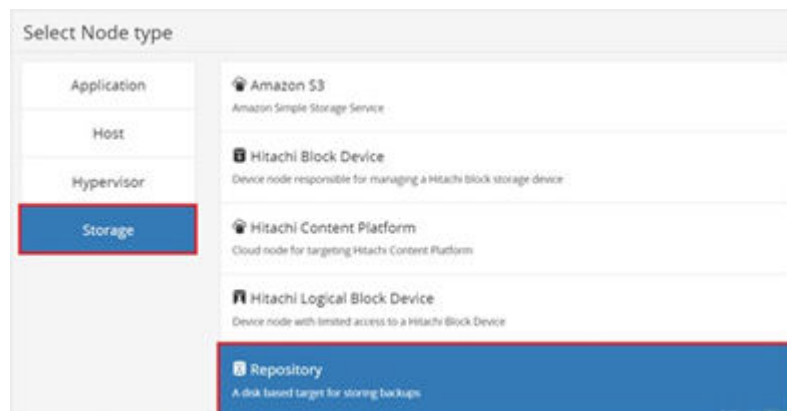
To create a generation 1 or 2 repository node, do the following.

Procedure

1. From the navigation tree, select **Nodes** and then the plus sign (+).




2. Create the node.
 - a. Under **Select Node Type**, select **Storage** and then **Repository**. Select **Next**.



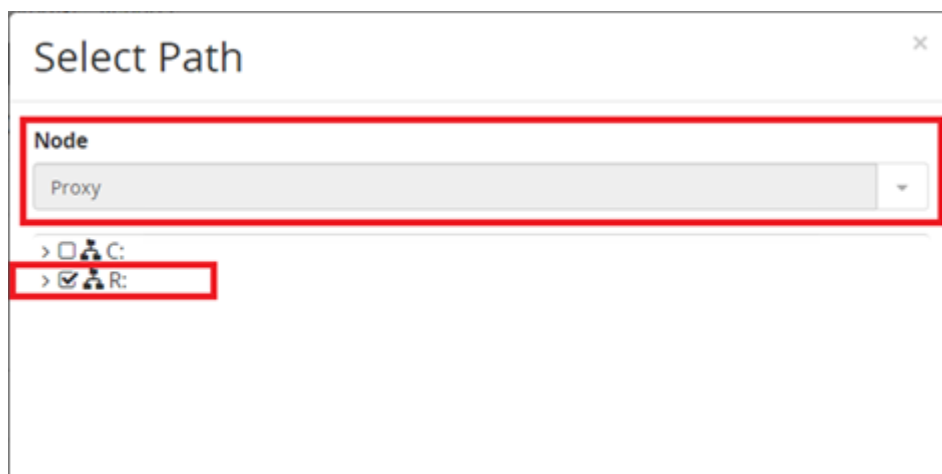
- b. Type the **Node Name** for the repository you are onboarding and select **Next**.



- c. Optionally, if your environment utilizes **Access Control Groups**, you can add them here. Select **Next**.
 - d. Select the machine which hosts the LUN you plan on defining as your repository. Select **Next**.



3. Create the repository.
 - a. Select **Create a new Repository**. Select **Next**.
 - b. Select your preferred Repository Generation, Generation 2 is recommended when using Protector. Select **Next**.
 - c. On **Specifying Repository directories**, select **Browse on Root Directory**.
 - d. Select the drive you have onboarded on to your guest operating system which will be used as the repository. Select **OK**. Data, Metadata, and Checksum will natively inherit the root directory specified. Select **Next**.



- e. (Optional) If you choose to enable repository encryption, select **Enable Repository Encryption** and provide a phase phrase. Select **Next**.
- f. (Optional) On **Configure Repository settings**, you can select the **Optimize for Cloud Replication** check box. Generation 1 Repositories do not have this option.



Note: Optimize for Cloud replication aligns block boundaries between the repository and the cloud to make cloud uploads more efficient.

- g. Define your **Capacity Warning Level** and **Capacity Critical Level**. By default, they are set to 85% and 95%, respectively.
- h. Select your transfer **Block Size**. This has an impact, depending on the size of data transfers. Select **Next**.

- i. View the summary of the repository and confirm correct drive selection. Select **Finish**. The repository is now onboarded and can be used as a destination for backups.

Hitachi Content Platform nodes

Hitachi Content Platform nodes allow another potential destination for file level backups via the operating system which have the Hitachi Ops Center Protector agent installed. Protector can now natively backup virtual machine VMDK files to Hitachi Content Platform object storage for long term archival.

A Hitachi Content Platform node in Ops Center Protector represents a single tenant, allowing the subset namespace to be created once the node has been defined.



Note: Prior to creating Hitachi Content Platform nodes, confirm that the management APIs are enabled at the system level for Content Platform as well as for the target tenant. See [Enable management APIs \(on page 55\)](#).

To create a Hitachi Content Platform node, do the following.

Procedure

1. From the navigation tree, select **Nodes** and select the plus sign (+).
2. Create a node.
 - a. Select **Storage** from **Select Node Type** and then select **Hitachi Content Platform**. Select **Next**.
 - b. Type the **Node Name**. Select **Next**.

- c. Optionally, if your environment utilizes access control groups, add them here. Select **Next**.

- d. Select either a Generation 1 or Generation 2 Repository. Generation 2 is recommended. Select **Next**.

- **Generation 2** supports batch backups directly from source nodes via other Generation 2 Hitachi Content Platform, Amazon S3, or repository nodes. This generation supports parallel backups from VMware nodes. Data is streamed through the proxy node to Content Platform.

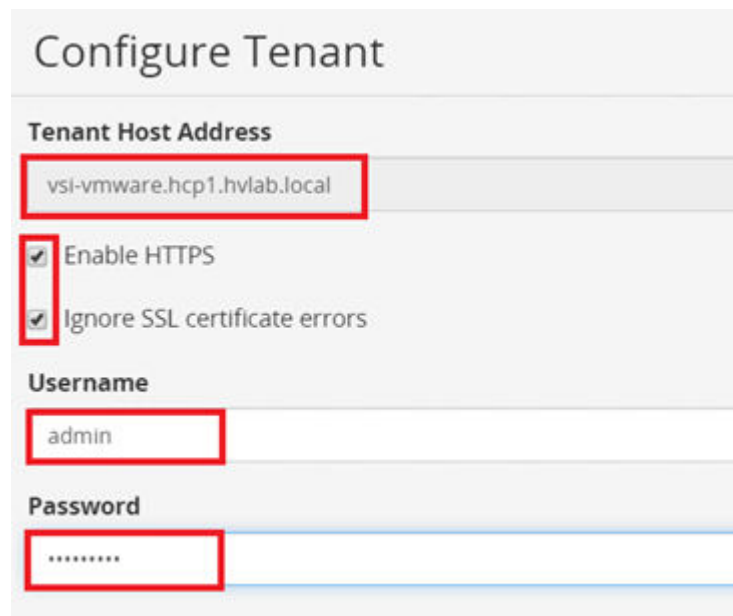
This requires Content Platform version 8.1.0.x or later. All the data for the node is stored in one Content Platform namespace, with statistics being available through Ops Center Protector.

- **Generation 1** supports tiering data from Generation 1 repository nodes. Data backed up to the repository is then uploaded to Hitachi Content Platform and deleted from within the repository. Restoring from a Generation 1 Content Platform node requires the original repository used to tier the data.

This requires Hitachi Content Platform version 7.3.2 or later. It is not compatible with Generation 2 repository nodes. Data is spread across multiple namespaces in Hitachi Content Platform. No statistics are available through Ops Center Protector.

3. Configure the new node.

- a. Select **Create a new HCP node**. Select **Next**.
- b. Select your Proxy node. Select **Next**.
- c. To select a metadata directory in relation to your proxy node, select **Browse**. Select **Next**.
- d. Define your tenant URL based on what was created on Hitachi Content Platform. Find your tenant URL by accessing the Hitachi Content Platform system. Enter the **Username** and **Password** for that tenant. Optionally users can enable HTTPS and to ignore SSL certificate errors. Select **Next**.



Configure Tenant

Tenant Host Address
vsi-vmware.hcp1.hvlab.local

☒ Enable HTTPS

☒ Ignore SSL certificate errors


Username
admin

Password

- e. Define the namespace. Select the Initial Quota in gigabytes, and then set Soft Quota Level. Select **Next**.

4. Review the setting summary and select Finish.

This onboards Hitachi Content Platform as a node. It can now be used as a destination. The namespace is created natively on Hitachi Content Platform (the following figure).



VMware nodes

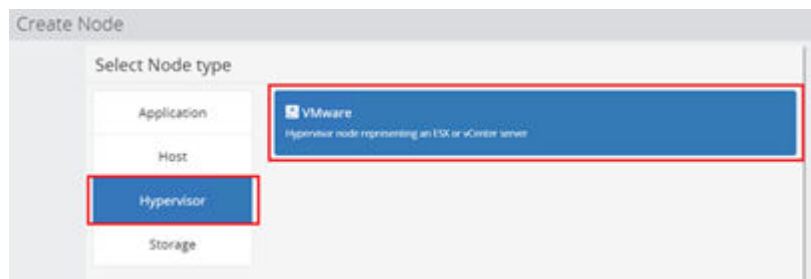
VMware nodes do not require an agent. Administrators can onboard single VMware ESXi instances, as well as VMware vSphere vCenter. Both onboarding methods require IP, user name, and password.

After onboarding VMware nodes, you can take advantage of change block tracking in VMware vSphere Storage APIs – Data Protection backups over a LAN to a local repository or over Fibre Channel SAN native to Hitachi Virtual Storage Platform.


To create a VMware node, do the following.

Procedure

1. From the navigation tree, select **Nodes** and select the plus sign (+).
2. Select **Hypervisor** from **Select Node type** and select **VMware** as your selection. Select **Next**.



3. Define the node name. Select **Next**.
4. Optionally, if your environment utilizes access control groups, add them here. Select **Next**.
5. From the **Select node to communicate with VMware server**, select the machine within your environment that has the Hitachi Ops Center Protector agent installed and hosts your repository. Select **Next**.



6. Specify the VMware server hostname or IP address. Select **Next**.
7. Enter VMware credentials with administrator or root privileges. Select **Next**.
8. Review the setting summary and select **Finish**.

You have now successfully deployed a VMware node which is considered a source data connection. VMware nodes can also be used for in place or out of place restore destinations.

Hitachi block devices

Hitachi block devices represent Hitachi Virtual Storage Platform systems which have been onboarded to Hitachi Ops Center Protector. Prior to onboarding, administrators must do the following:

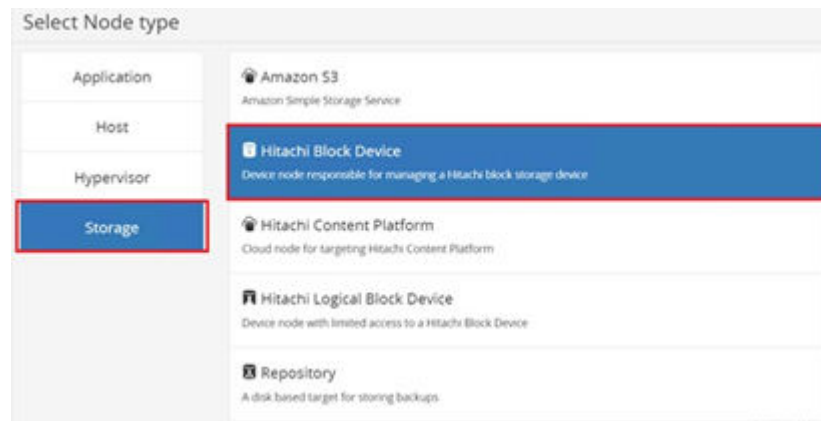
- Make sure a command device (CMD) from that respective system is RAW device mapped (in the case of virtual environments) and presented, in most cases, to the same guest operating system which holds the repository.
- Confirm installation of command control interface (CCI).

For more information on how to install Protector agents, and how to create CMDs, see [Related documents \(on page 9\)](#).

To create a Hitachi Block Device node, do the following.

Procedure

1. From the navigation tree, select **Nodes** and select the plus sign (+).
2. Select **Storage** from **Select Node type** and select **Hitachi Block Device**. Select **Next**.



3. Type a **Node Name** and select the check box confirming that you have read and understood the warnings around managing Ops Center Protector resources outside of the Protector user interface. Select **Next**.
4. Optionally, if your environment utilizes access control groups, you can add them here. Select **Next**.
5. Select the machine which hosts your CMD from your Hitachi Virtual Storage Platform. Select **Next**.



Note: Command control interface must be installed on the machine that hosts the CMD.

6. Choose a common metadata directory on the guest operating system which hosts your CMD. Protector places metadata files related to block snapshots and replications among all Virtual Storage Platform within this directory. Select **Next**.



Note: The metadata directory is defined once for this proxy and all storage nodes on this proxy will use this setting. It cannot be changed after initial configuration.

7. Select the **Select from detected storage devices** option and select the serial number of the Virtual Storage Platform you want to onboard from the **Select a Storage Serial Number** list.
 - If you do not see your Virtual Storage Platform in the list, confirm that you have the CMD presented to the guest operating system as well as confirming installation of command control interface. Select **Next**.



8. Configure the system.
 - a. Enter in your username and password to onboard the system. It is highly recommended you create an Ops Center Protector user on the respective Virtual Storage Platform system that has the roles of storage administrator and security administrator. Select **Next**.
 - b. Select your LDEV provisioning range as **All**. Select **Next**.
 - c. On **Configure CMD specification and priority**, keep default settings. Select **Next**.
 - d. Optionally you can specify LDEV ranges to associate with specific virtual storage machines (VSM). Select **Next**.
 - e. On **Specify ports used for provisioning**, you can predefine ports used for allocation. This can also be later defined when doing restore operations. Select **Next**.
 - f. View the summary and confirm settings. Select **Finish**.

You have now successfully deployed a Hitachi block device node which is considered a source and destination data connection.

Amazon S3 nodes

Use the Amazon Simple Storage Service (S3) storage node to transfer application data to the Amazon S3 cloud storage. The Amazon S3 storage node represents the storage in Amazon S3 cloud. When creating the node, you will define the necessary credentials to access the Amazon S3 service. These include Access Key ID, Secret Access Key, Bucket Name, and Region.

These are all important for your Amazon S3 installation:

- To reduce bandwidth and improve transfer speeds, leveraging Amazon S3 partial transfer with only the block changes are sent to Amazon S3.
- Proxy node will be the intermediary between the application and the Amazon S3 service.

- A key feature of Amazon S3 support is to keep data stored on Amazon S3 in a native format as much as practicable. Any files less than 4 GB in size has its data stored in a single object. This means that third party search, indexing, and analytics tools can process the data in Amazon S3 directly. By default, files bigger than 4 GB are split into 4 GB segments that can easily be stitched together, if required.
- If encountering connection issues with the Amazon S3 bucket, confirm that the tags in the following table are in place using Amazon S3 user interface (the following figure).

Table 6 AMAZON S3 TAGS

Key	Value
stream_max_parts	512
Product	Protector
Company	Hitachi Vantara
Resource_Owner	S3Admin

To create an Amazon S3 node, do the following.

1. From the navigation tree, select Nodes and select the plus sign (+).
2. Select Storage and select Amazon S3. Select Next.

3. Define a node name. Select Next.

4. Optionally, if your environment utilizes access control groups, you can add them here. Select Next.
5. Select Create a new S3 node. Select Next.
6. From the drop down select your proxy node. Select Next.
7. Select Browse and select a metadata directory for the S3 cache and select OK. Select Next.
8. On the Configure S3 menu, define your Access Key ID, Secret Access Key, Bucket Name, and Bucket Region. Select Next.
9. Review your settings and confirm selection. Select Finish. You have now successfully deployed an Amazon S3 node.

Create policies

A policy consists of the following:

- Classifications that specify what data is to be protected
- Operations that specify how that data is to be protected

With policies, administrators can tie in recovery point objectives and schedules to guarantee a time frame where Hitachi Ops Center Protector takes a backup along with the ability to run a prescript before or a postscript after the backup.



Note: Prior to setting up policies, you must have onboarded the nodes you plan to use as your source as well as destination.

For more information on policies, see [Related documents \(on page 9\)](#).

VMware vSphere Storage APIs - Data Protection over LAN and SAN

Hitachi Ops Center Protector can provide end-to-end protection for VMware vSphere using VMware vSphere Storage APIs – Data Protection (VADP), which can be applied regardless of whether the datastores reside on Hitachi storage. Ops Center Protector uses change block tracking to only backup changed blocks for further operational efficiencies. VMware vSphere Storage APIs – Data Protection backups over LAN has the destination of the local repository.

Backups utilizing VADP directly to SAN copies data directly to the storage system. Enabling change block tracking provides operational efficiencies during a backup.

To create a VMware vSphere Storage APIs – Data Protection over LAN or SAN policy, do the following.

Procedure

1. From the navigation tree, select **Policies** and select the plus sign (+).
2. From **Specify name and description**, type a **Name** and a **Description**. Select **Next**.

Specify name and description

Name
VADP_LAN

Description
VMware API for DP over LAN to REPO

3. Select the plus sign (+) to add a classification.

Add one or more Classifications

Select All (0) [icon]

No policy classifications. Click to add a classification.

+

4. From **Select Classification**, select **Hypervisor** and select **VMware**. Select **Next**.

Select Classification

Application

Physical

Hypervisor

VMware

5. From **Specify VMware classification attributes**, select your **VMware Node** and select **Add**.

Specify VMware classification attributes

VMware Node
vCenter_01

Include Items

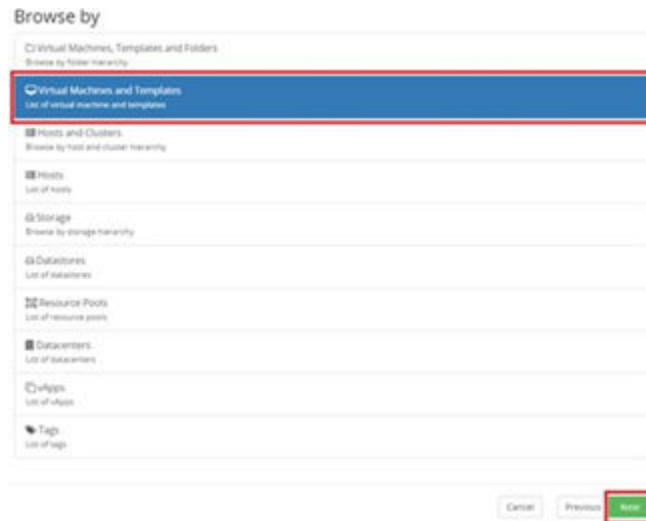
Name	Type
No items selected	

Exclude Items

Name	Type
No items selected	

Add

6. Add resources.
 - a. Using the selection inclusion wizard, select **Browse for resources**. Select **Next**.
 - b. Select **Virtual Machines and Templates**. Select **Next**.



- c. To select the virtual machines to include within the data protection policy, select the applicable check boxes. Select **Finish**.



- d. Select **Apply**.
e. Select **Next**.

7. Add operations.

- a. Select the plus sign (+) on **Add one or more Operations**.



- b. On **Select Operation**, select **Backup**. Select **Next**.

Select Operation

- Backup
- CDP
- Mount
- Replicate
- Snapshot
- Test

- c. On **Specify Backup operation attributes**, type a **Name** that is descriptive of your backup operation. Additionally, you can define your data protection schedule under **Run Options** by selecting the **Run an RPO** (recovery point object) option, the **Run on RPO and Schedule** option while selecting a schedule, or the **Run on completion of operation** (trigger schedule) option. Under **Schedule Options**, set the **Recovery Point Object** and **Retention** time for Protector. Select **Apply**.

Specify Backup operation attributes

Name: Backup_VADP_LAB

Run Options:

- Run on RPO
- Run on RPO and Schedule
- Run on completion of operation

Schedule Options:

Recovery Point Object: Days

Retention: Days

Source Options:

Quiesce configured applications before backup

Pre Script:

Post Script:

- d. Select **Apply**.

8. Select **Finish**.

The policy has now been defined. You can edit the policy from the policies container at your own discretion to change virtual machine selections, recovery point objects, schedules, and pre- or postscripts.

Native VMDK to Amazon S3

Hitachi Ops Center supports native backups to Amazon S3 via a batch operation.



Note: To further optimize data transfer speed in addition to delta block differencing, multiple objects are sent to Amazon S3 at the same time and multiple data blocks within those objects are sent separately to Amazon S3 in parallel.

To create an Amazon S3 policy, do the following.

Procedure

1. From the navigation tree, select **Policies** and select the plus sign (+).
2. On **Specify name and description**, type a **Name** and **Description** for the policy. Select **Next**.

Specify name and description

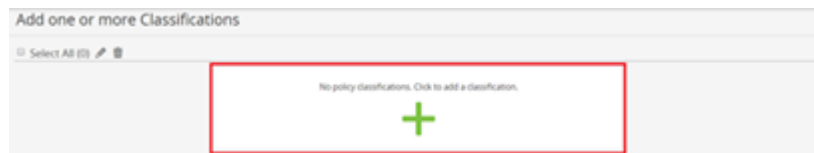
Name

VSI VMDK_S3

Description

Protects native VMDK files and pushes data off site to S3.

3. Select the plus sign (+) to create a classification.



4. From Select Classification, select Hypervisor and select VMware. Select **Next**.
5. On **Specify VMware classification attributes**, select your VMware node and select **Add**.
6. Select resources.
 - a. Using the selection inclusion wizard, select **Browse for resources**. Select **Next**.
 - b. Select **Virtual Machines and Templates**, Select **Next**.
 - c. To select the virtual machines, you want included within the data protection policy, select those check boxes. Select **Finish**.
 - d. Select **Apply**.
 - e. Select **Next**.
7. Select the plus sign (+) on **Add one or more Operations**.
8. Select **Backup Operation**. Select **Next**.
9. Type a **Name** that describes your backup operation. You can define your data protection schedule under **Schedule Options** and **Run Options** by defining a **Recovery Point Object** or trigger schedule. Also, under **Schedule Times**, define **Retention** times for Protector. Select **Apply**.
10. Select **Finish**.

The policy has now been defined. You can edit the policy from the policies container at your own discretion to change virtual machine selections, recovery point objects, schedules, and pre- or postscripts.

Native VMDK to Hitachi Content Platform with Generation 2 nodes

Hitachi Ops Center now supports native backups to the Hitachi Content Platform for long-term archival and object analysis via a batch operation.



Note: The following procedure is for Hitachi Content Platform with Generation 2 Nodes. For Generation 1 implementation, refer to [Generation 1 and 2 repository nodes \(on page 26\)](#).

Before you start, ensure that the proxy node defined within the wizard and the Hitachi Content Platform node are time synced with a common NTP server. If the time differs by over 5 minutes, node creation and future backups may fail.

To create a native VMware backup policy to the Hitachi Content Platform, do the following.

Procedure

1. From the navigation tree, select **Policies** and select the plus sign (+).
2. On **Specify name and description**, type a **Name** and **Description** for the policy. Select **Next**.

3. Select the plus sign (+) to create a classification.

4. From **Select Classification**, select **Hypervisor** and select **VMware**. Select **Next**.
5. On **Specify VMware classification attributes**, select your VMware node and select **Add**.
6. Select resources.
 - a. Using the selection inclusion wizard, select **Browse for resources**. Select **Next**.
 - b. Select **Virtual Machines and Templates**. Select **Next**.
 - c. To select the virtual machines you want included within the data protection policy, select those check boxes. Select **Finish**.
 - d. Select **Apply**.
 - e. Select **Next**.
7. Add a backup operation.
 - a. Select the plus sign (+) on **Add one or more Operations**.
 - b. Select **Backup Operation**. Select **Next**.
 - c. Type a **Name** that describes your backup operation. You can define your data protection schedule under **Schedule Options** and **Run Options** by defining a **Recovery Point Object** or trigger schedule. Also, under **Schedule Times**, define **Retention** times for Protector. Select **Apply**.

8. Select **Finish**.

The policy has now been defined. You can edit the policy from the policies container at your own discretion to change virtual machine selections, recovery point objects, schedules, and pre- or postscripts.

Create data flows

Data flows allow a whiteboard approach in defining how to protect your data. A combination of source and destination nodes are made available from which you will apply the policies created. After creating policies, you must active and distribute the rules among all nodes within the data flow.



Note: Prior to setting up data flows, you must have created your data collection policies to apply to the workflow.

For more information on data flows, see [Related documents \(on page 9\)](#).

VMware vSphere Storage APIs – Data Protection LAN or SAN flow

A VMware vSphere Storage APIs – Data Protection (VADP) over LAN or SAN flow allows over-the-network backups along with a direct to storage system.



Note: VADP to SAN is automatically decided by Hitachi Ops Center Protector, based on proxy datastore access which utilizes SAN mode transport.

To orchestrate VMware vSphere Storage APIs – Data Protection over LAN or SAN to a local repository for backup, do the following.

Procedure

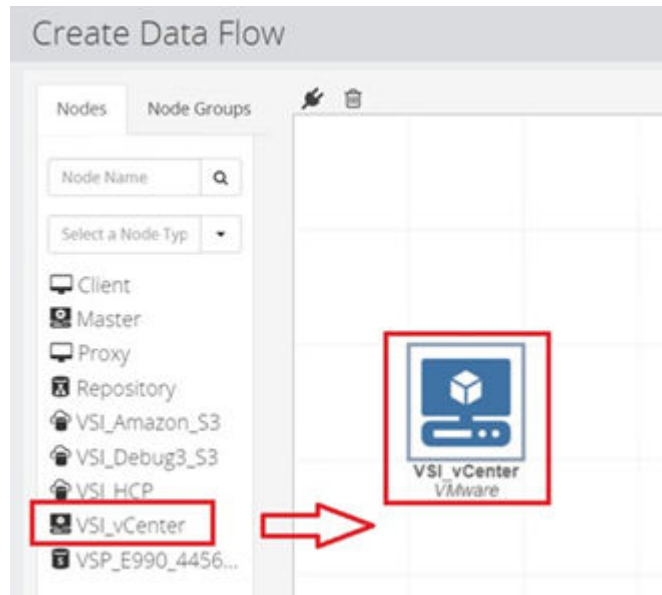
1. From the navigation tree, select **Data Flows** and select the plus sign (+).
2. Type a **Name** for the policy and type a description. Select **Next**.

The screenshot shows a configuration form with two main sections: 'Name' and 'Description'. The 'Name' section has a text input field containing 'VSI_VADP_LAN_Flow'. The 'Description' section has a larger text input field containing 'VADP backup of VM in vCenter pushed to REPO via LAN.' Both input fields are highlighted with red rectangular boxes.


3. Create the data flow.

You are presented with a blank workspace to draw out your data flow. Your list of source and destination nodes are listed on the left side.

- a. Drag your source **VMware** node to the workspace.



- b. Drag the **Repository** to your workspace through the **VMware** node. Make sure to drag the **Repository** node through the **VMware** node to create a visual connection.



- c. Select your **VMware** node to apply your VMware vSphere Storage APIs – Data Protection over LAN policy.



- d. Select your **Repository** node to apply the same VMware vSphere Storage APIs – Data Protection over LAN policy as your **VMware** node.




- e. After applying your policy, configure the operation properties of the **Repository** node associated with this backup.
- f. Select **Finish**.

4. Distribute the rules.

Once you have created your policy you must distribute the rules among nodes.

- a. Select your new data flow and select the **Play Activation** button (triangle).
- b. Once rules have compiled, select **Activate** to distribute among nodes within the data flow.



Once rules are distributed among nodes, your backup will kick off based on your schedules or RPO objectives.

Native VMDK to Hitachi Content Platform



Note: As of Hitachi Content Platform Generation 2 nodes, you can change the namespace size using the Hitachi Content Platform user interface.

To orchestrate native VMDK batch backups to Hitachi Content Platform Object storage, do the following.

Procedure

1. From the navigation tree, select **Data Flows** and select the plus sign (+).
2. Type a **Name** for the policy and type a **Description**. Select **Next**.

Specify name and description

Name

VSI_VMDK_to_HCP_Flow

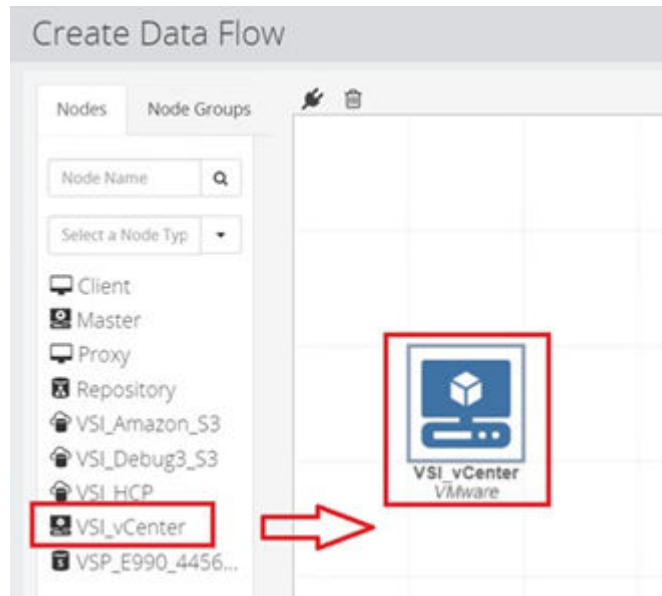
Description

Native batch backup of VM VMDKs to HCP NS to longterm retention.

3. Create the data flow.

You are presented with a blank workspace to draw your data flow. Your list of source and destination nodes are listed on the **Nodes** tab.

- a. Drag your source **VMware** node to the workspace.



- b. Select your **Hitachi Content Platform** node. Drag it through your **VMware** node to create a visual connection.



4. Apply your Hitachi Content Platform policy to your source VMware node and the destination Hitachi Content platform node.



5. Distribute the policy.

Once you have created your policy, you must distribute the rules among nodes.

- a. Select your new data flow and select the **Play Activation** button (triangle).
- b. Once the rules have compiled, select **Activate** to distribute the policy among nodes within the data flow. Once the rules are distributed among nodes, your backup will start based on your schedules or RPO objectives.

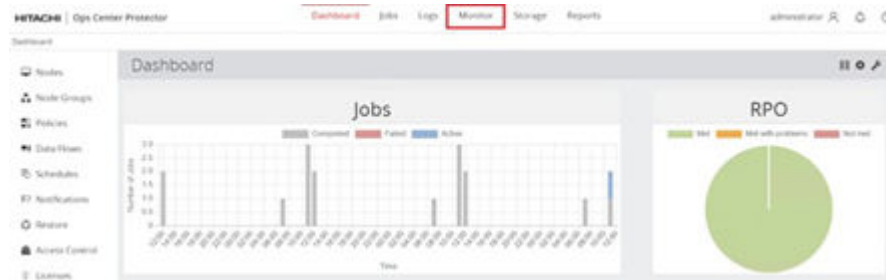
Manually trigger operations

Once activating data flows, you can start backups when you want rather than waiting for the trigger schedule or RPO objective.

To trigger a backup job manually, do the following.


Procedure

1. To log on to the Ops Center Protector user interface, type this in the Address bar of a browser: <https://Hitachi Ops Center Master IP:443>
2. From the Protector dashboard menu, select **Monitor**.



The **Monitor** page in Protector lists your active data flows.

3. From the **Monitor** page, select the name of the data flow you want to start.



4. Select your **source** node.
5. From the menu, select the **Thunderbolt** button.



6. From the wizard, select the backup policy. Select Run Now.



Once the job starts, you can follow its progress using the **Monitoring** menu. View job details by selecting the data flow.

View Jobs, Logs, and Storage Using Ops Center Protector

Once backup operations are in progress, view statistics related to the backup using Jobs, as well as viewing Hitachi Ops Center Protector backend operations using Logs. Once backups are completed, view native storage view on replica LDEVs within the storage container.

Jobs

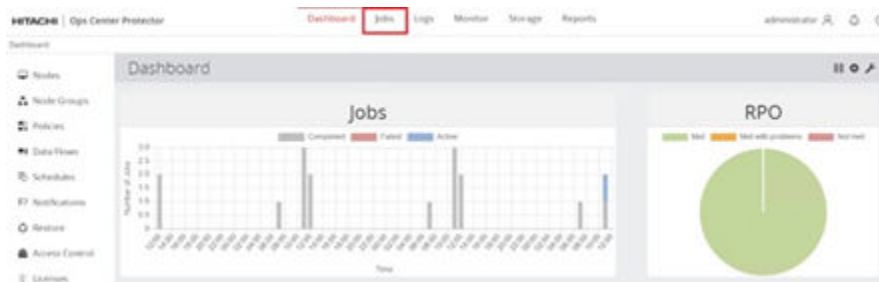
Jobs are tasks being processed by the system. They are either of the following:

- System-initiated tasks, such as a scheduled backup or report creation
 - User-initiated tasks, such as a restore or a repository analysis
- Jobs are split into three types:
- Backup
 - Restore
 - Other

Jobs are executed by various subsystems on various Hitachi Ops Center Protector nodes. They are tracked by the Master node so that it can provide progress and status information to users. A job typically goes through the following lifecycle:

1. Queued
2. In progress
3. Completed (succeeded or failed)

To view jobs related to a backup, select Jobs from the menu.

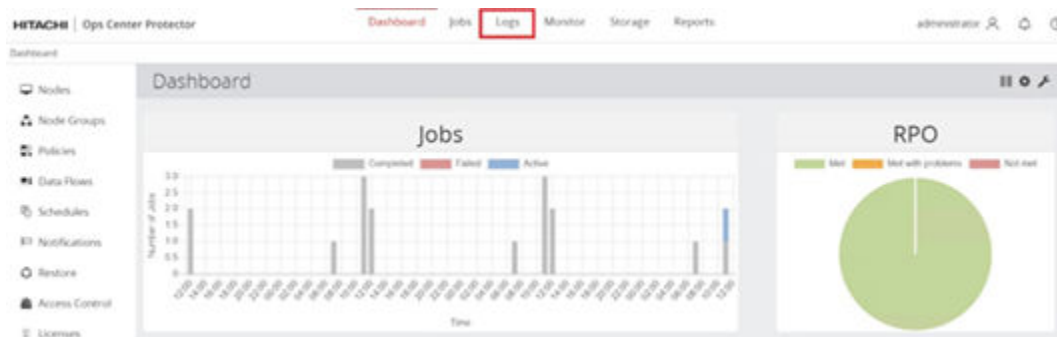


Logs

All log messages are stored on the Master node. They are made available using the user interface. Each log entry includes the following information in addition to a short textual description of the event:

- A unique ID assigned to every log entry, useful when discussing logs with colleagues and support engineers.
- An audit flag indicating a change to the system status. These cannot be modified or deleted, for compliance purposes. The identity of the user initiating the action is recorded.
- A date and time record when the log was generated on the originating node and received by the Master node.
- The level, identifying the importance of the log entry.
- An attachment enabling additional context information to be attached to the log entry.
- An ID identifying the position within the Protector code where the message was generated. The ID can be used to reference an engineering support database that provides additional description and identifies possible causes and solutions.

Log entries can be acknowledged by the user and marked with a comment so that there is a record that the event has been noticed and addressed. The entire log data base can be exported in various file formats so that it can be analyzed offline or presented in a report. To view logs related to a backup, select Logs from the menu.



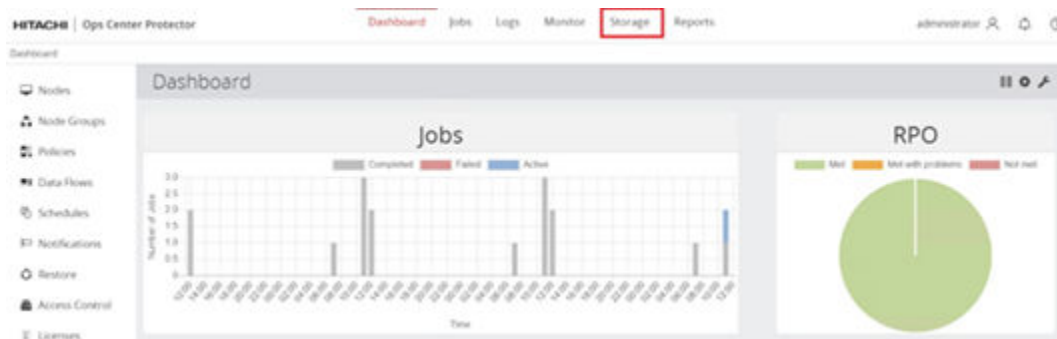
Storage

The storage inventory provides access to all the available data storage types managed by Ops Center Protector. Information such as logs, jobs, proxy status, and device configurations are made available. The following are the storage types:

- Repository — Administrators can view batch and continuous backups of applications, virtual machines, and file systems.
- Hitachi Block and Logical Block — Administrators can view snapshots and replications of applications, virtual machines, file systems, and LDEVs.
- Hitachi Content Platform — Administrators can view Jobs, Logs, Configuration, and Storage Details in relation to Hitachi Content Platform.
- Amazon S3 — Administrators can view Jobs, Logs, Configuration Storage Details, and Progress Details in relation to Amazon S3 buckets.

Additionally, restore operations can be initiated by selecting a storage type.

To view your storage container types related to a backup, select Storage from the menu.



Restoration using Ops Center Protector

The restore dashboard provides access to all the available backups created or managed by Protector. Backups can be viewed as one or by storage types, such as Hitachi Block or Repository.



Note: The process of restoring data can result in the destruction of some original data that exists on the restore target. Ensure that any critical data is copied to a safe location and/or is included in the data set being restored.

VMware vSphere Storage APIs – Data Protection Over LAN restoration

To access a VMware vSphere Storage APIs – Data Protection over LAN backup and initiate a restoration, do the following.

Procedure

1. From the navigation tree, select **Restore**.
2. Using the filter for each item, select your **Source** node, **Date Time Range** for the backup, and any applicable **Policy**. Select **Search**.

The screenshot shows the search interface for VMware vSphere Storage APIs – Data Protection Over LAN. The interface includes the following fields and controls:

- Capture Date:** Custom Range (06/23/2020 00:00:00 to 07/06/2020 23:59:59)
- Application or Origin Node:** VSI_vCenter
- Application Node Type:** VMware
- Virtual Machine:**
- Folder:**
- Datastore:**
- Storage Node:** Select a Node
- Storage Node Type:** Select Storage Node Type
- Data Flow:** Select a Data Flow
- Policy:** VSI_VADP_LAN
- Operation Name:** Operation Name
- Search:** Button

3. Select your **Backup**. From the menu, select **Restore**.

The screenshot shows the Restore interface for VMware vSphere Storage APIs – Data Protection Over LAN. The interface includes the following details:

- Restore:** Title
- Backup Item:** 07/06/2020 11:17:10
- Data Origin:** VSI_vCenter
- Application:** VMware
- Policy:** VSI_VADP_LAN (Backup_VADP_...)
- Storage:** Repository
- Expiry Date:** 07/07/2020 11:17:10

4. Set the restoration options.
 - a. Select the check boxes of the virtual machines to be restored, Select **Next**.



- For restore type, select **Clone**. Select **Next**.
- Define a clone prefix for your virtual machine and select the respective destination node. Restorations can be in- place or out of place. Select **Next**.
- Select your clone destination data center that supports your virtual cluster, Select **Next**.
- Select applicable compute resources. Select **Next**.



- Select an available datastore which supports your cluster. Select **Next**.

Name	Capacity	Free	Provisioned
datastore1	18.75 TB	18.75 TB	2.31 GB
datastore1 (1)	18.75 TB	18.75 TB	2.31 GB
datastore1 (2)	18.75 TB	18.75 TB	2.31 GB
datastore1 (3)	18.75 TB	18.75 TB	2.31 GB
DS_VF1	459.75 GB	452.74 GB	7.01 GB
E990-VMFS_1_Testing	499.75 GB	391.32 GB	125.46 GB
E990-VMFS_0	10.00 TB	9.59 TB	1.40 TB
HCP1_DS	1.20 TB	179.23 GB	1.02 TB

- Select the Powered OFF/ON state of the virtual machine after the restore operation. Select **Next**.
- Select **Finish**.


View your restoration details using **Jobs**.

VMware vSphere Storage APIs – Data Protection Over SAN restoration

To access a VMware vSphere Storage APIs – Data Protection over SAN backup and initiate a restoration, do the following.

Procedure

- From the navigation tree, select **Restore**.
- Using the filter for each item, select your **Source** node, **Date Time Range** for the backup, and any applicable **Policy**. Select **Search**.
- Select your **backup**. Select the **Restore** button in the menu.



4. Set the restoration options.
 - a. Select **User selection** as the restore scope. Select **Next**.
 - b. Select the check boxes of the virtual machines to be restored. Select **Next**.

Name	Type	Operating System	ID
<input checked="" type="checkbox"/> VSI_Windows_4	virtual machine	Microsoft Windows Server 2016 or later (64-bit)	vm-810
<input checked="" type="checkbox"/> VSI_Windows_1	virtual machine	Microsoft Windows Server 2016 or later (64-bit)	vm-807
<input checked="" type="checkbox"/> VSI_Windows_5	virtual machine	Microsoft Windows Server 2016 or later (64-bit)	vm-811
<input checked="" type="checkbox"/> VSI_LINUX_4	virtual machine	CentOS 7 (64-bit)	vm-805
<input checked="" type="checkbox"/> VSI_LINUX_1	virtual machine	CentOS 7 (64-bit)	vm-802
<input type="checkbox"/> VSI_Windows_2	virtual machine	Microsoft Windows Server 2016 or later (64-bit)	vm-808
<input type="checkbox"/> VSI_Windows_3	virtual machine	Microsoft Windows Server 2016 or later (64-bit)	vm-809
<input type="checkbox"/> VSI_LINUX_3	virtual machine	CentOS 7 (64-bit)	vm-804
<input type="checkbox"/> VSI_LINUX_5	virtual machine	CentOS 7 (64-bit)	vm-806
<input type="checkbox"/> VSI_LINUX_2	virtual machine	CentOS 7 (64-bit)	vm-803

- c. For **Restore Type**, select **Clone**. Select **Next**.
 - d. Type a clone prefix for the virtual machine and select the respective destination node. Restoration can be in-place or out of place. Select **Next**.
 - e. Select the clone destination datacenter that supports the virtual cluster. Select **Next**.
 - f. Select the applicable compute resources. Select **Next**.



- g. Select an available datastore which supports the cluster, Select **Next**.
 - h. Select the desired machine state after restoration. Select **Next**.


- i. Select your mounting operation.
 - Mounting the original and modifying overwrites any data that was captured by Protector.
 - Mounting a clone loses any changes made while using the cascade snapshots. To do this, select **Cascade Snapshot**.
- j. Select **Finish**.
You can now view the restoration details using **Jobs**.

Amazon S3 restore

To perform an Amazon S3 restore, do the following.

Procedure

1. From the navigation tree, select **Restore**.
2. Select the settings.
 - a. Using the filter for each item, select your **Source** node, the **Date Time Range** for the backup, and any applicable **Policy**. Select **Search**.
 - b. Select your **Backup**. Select the **Restore** button in the action's menu.



- c. Select the **virtual machines to restore**. Select **Next**.
 - d. Select **Clone**. Select **Next**.
Restores can be in place to the same node of the original source or out of place to a different node than the original source.
 - e. Select your **Cloned Virtual Machine Name Prefix**. Select a **Destination Node**. Select **Next**.
 - f. Select the restore destination. Select **Next**.
 - g. Select applicable **Compute Resource**. Select **Next**.
 - h. Choose a desired VMFS datastore for the restore destination. Select **Next**.
 - i. Select the power state of the restored virtual machine. Select **Next**.
3. Select **Finish**.


4. View your restore details using **Jobs**.

Hitachi Content Platform restore

To access Hitachi Content Platform backups and initiate a restoration, do the following.

Procedure

1. From the navigation tree, select **Restore**.
2. Select the settings.
 - a. Using the filter for each item, select your **Source** node, **Date Time Range** for the backup, and any applicable **Policy**. Select **Search**.
 - b. Select your **Backup**. In the Action menu, Select **Restore**.



- c. Select the resources that require restoration. Select **Next**.
 - d. Select **Clone**. Select **Next**.

Restores can be in place to the same node of the original source or out of place to a different node than the original source.
 - e. Select your **Cloned Virtual Machine Name Prefix**. Select a **Destination Node**. Select **Next**.
 - f. Select the restore destination. Select **Next**.
 - g. Select applicable **Compute Resource**. Select **Next**.
 - h. Choose a desired VMFS datastore for the restore destination. Select **Next**.
 - i. Select the power state of the restored virtual machine or virtual machines. Select **Next**.
 - j. Select **Finish**.
3. View your restore details using **Jobs**.

Use Hitachi Content Platform with Hitachi Ops Center Protector

This section describes how to implement Hitachi Ops Center Protector in conjunction with Hitachi Content Platform.

Access the Hitachi Content Platform user interface

After deploying Hitachi Content Platform, open the user interface to set up tenants, namespace, and the various services available to manage data. For information on how to install Hitachi Content Platform, see [Related documents \(on page 9\)](#).

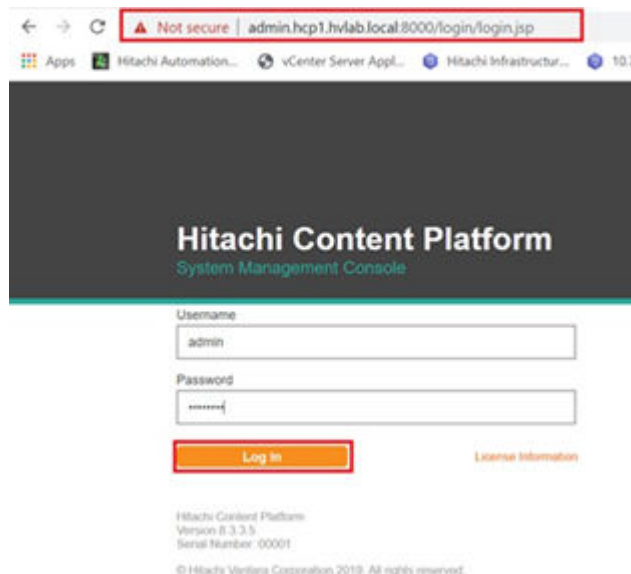
Note the following when accessing Content Platform:

- Initial log on to the Hitachi Content Platform user interface is with the built in Security account. This account allows the creation of other users and applies system permissions. Use the Security account to provide all roles, such as Monitor, Administrator, Security, Compliance, Service, and Search. For information on how to create users and apply permissions, see Related Documents.
- If your account is set up for local authentication, change your password in the System Management Console. When you change your password in this console, it also changes the password for any other Content Platform interfaces to which your user account gives you access.
- It is highly recommended you deploy Content Platform using FQDN rather than IP. Certain features, such as load balancing, require a secondary stub zone DNS server.

To configure Hitachi Content Platform with DNS, see [Related documents \(on page 9\)](#). To open the Hitachi Content Platform user interface, do the following.


Procedure

- Open a supported browser and type this URL in the Address bar: `https://admin.hcp-domain-name:8000`
- Enter your **Username** and **Password**, and select **Log In**.



Dashboard overview

When logged on to Hitachi Content Platform, you are presented with the system dashboard which gives insight into system events, status, and ingested objects.





















The top banner on the system gives a visual representation of overall health grouped by three categories.

- Hardware status correlates to the compute node health on the systems supporting the Hitachi Content Platform cluster.
- System Status gives insight to backend storage and its health supporting the Hitachi Content Platform system.
- Object Integrity provides administrators information whether any objects and its metadata have become corrupted. If any of these system components encounter issues, the buttons shift from green to red.

Run services on demand and disable system services

With Hitachi Content Platform, there are system wide, tenant, and namespace-based services.

System services typically run on a set schedule defined by the administrator. When administrators run services on demand, it is made available using the dashboard Services container. Press the Play button (triangle) next the service you want to run on demand. Disable a service system wide by selecting the Power button (off-on).

Services			
	Service	Status	Time
 	Capacity Balancing	Waiting	2/13/2020 4:21PM
 	Compression/Encryption	Completed	4/6/2020 10:02AM
 	Content Verification	Completed	4/5/2020 7:00PM
	Disposition	Disabled	2/13/2020 4:21PM
 	Duplicate Elimination	Completed	4/6/2020 12:04AM
	Garbage Collection	Completed	4/6/2020 12:00AM
 	Geo-distributed Erasure Coding	Waiting	2/13/2020 4:21PM
 	Protection	Completed	4/5/2020 12:00AM
 	Scavenging	Completed	4/5/2020 12:01AM
 	Storage Tiering	Completed	4/5/2020 12:00PM

Enable management APIs

Administrators need to enable management APIs (MAPI) at the system level for Hitachi Content Platform. MAPI is a REST interface that allows applications to access and modify some of the Hitachi Content Platform system settings.




Note: Enabling MAPI system wide is a prerequisite to enabling MAPI for tenants.


To enable MAPI system wide, do the following.

Procedure

1. Log on to Hitachi Content Platform.
2. Select **Security** and then select **MAPI**.



3. Under **Management API**, select the **Enable the management API** check box. Select **Update Settings**.

**Storage components for cloud adaptive tiering**

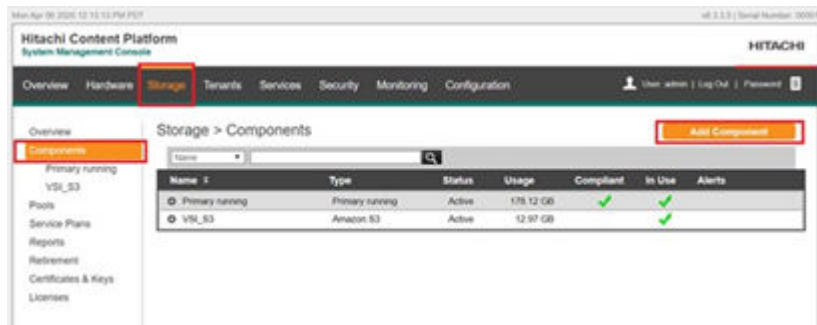
The Hitachi Content Platform provides administrators various data placement options once data has been ingested onto the system. To support hybrid cloud infrastructure, Content Platform supports the concept of cloud adaptive tiering, where administrators can push local on-premises data from Hitachi Virtual Storage Platform, s10 nodes, or s30 nodes to supported cloud infrastructure, such as Amazon Simple Storage Service, Google Cloud, Microsoft Azure®, NFS, s3 compatible silos, and ThinkOn.

Onboard Amazon Simple Storage Service (S3) account

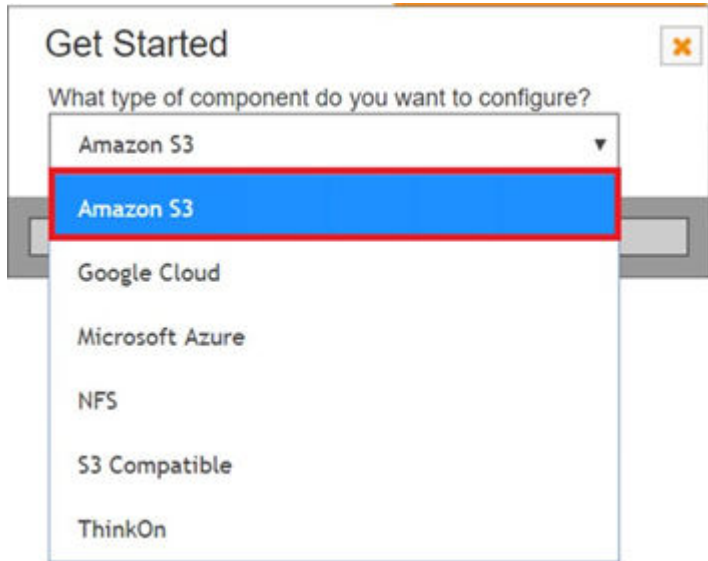
To onboard an Amazon Simple Storage Service (S3) account to enable off-premises data archival, do the following.

Procedure

1. Log on to Hitachi Content Platform.
2. Select **Storage**.
3. Select the **Components** tab and select **Add Component**.



4. From the list, select **Amazon S3**. Select **Go**.



5. Type a **Name** and **Description** for your account.
6. From **Connection Settings**, on **Advanced**, select the **Region** to which your buckets reside. The default is **US-East-1**.

7. Select **Next**.
8. Enter the **Account Label**, **Access Key**, and **Secret Key** for your account, Select **Next**.

9. Select your pre-created bucket or define a new bucket name. Select **Next**.
 - Administrators can select buckets natively created on Amazon Web Service.
 - Hitachi Content Platform allows administrators to create a bucket natively from the Hitachi Content Platform console.

10. Select **Finish**.
The Amazon Web Service account and bucket is onboarded to Hitachi Content Platform.

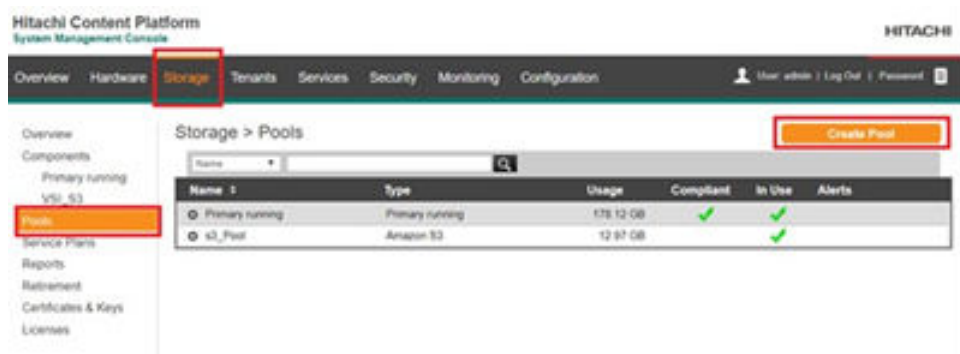
Create an Amazon S3 storage pool

After onboarding the Amazon S3 account, create an Amazon S3 pool which represents the storage tier of the Amazon Web Service account. This tier is used when creating service plans to define the transition of data


To create an Amazon Web Service storage pool, do the following.

Procedure


1. Log on to Hitachi Content Platform.
2. Select **Storage**.
3. Select the **Pools** tab. Select **Create Pool**.



4. From the pool type list, select **Amazon S3**. Select **Go**.



5. Type the **Name** and **Description** for your pool. Select **Next**.
6. Select your onboarded Amazon Web Service account. Select **Next**.



7. Select **Finish**.

Create service plans

Service plans define the transition of data among supported Hitachi Content Platform components. Once you define your cloud component and onboard its pool tier, you can create a service plan.

To create a service plan, do the following.

Procedure

1. Log on to the Hitachi Content Platform.
2. Select the **Service Plans** tab and select **Create Service Plan**.

Create Pool

Name	Component	Account	Bucket	Compliant
VSI_S3	VSI_S3	[Redacted]	vsi-s3-bucket	<input checked="" type="checkbox"/>

Data Storage Options

☐ Compress Data at Rest

Cancel < Previous **Next >** Finish

3. Type the **Name** and **Description** of your service plan. Select **Next**.
4. Select the primary ingestion tier, which refers to backend Hitachi Virtual Storage Platform. Typically, this is the **Default** tiering strategy. Select **Next**.
5. Select **Finish**.
6. Define the service plan.
 - a. From the list, select your new **Service Plan**. Select **Tiers**.
 - b. Select **Add Tier**.

Create Pool

Name	Component	Account	Bucket	Compliant
VSI_S3	VSI_S3	[Redacted]	vsi-s3-bucket	<input checked="" type="checkbox"/>

Data Storage Options

☐ Compress Data at Rest

Cancel < Previous **Next >** Finish

- c. Define your **Transition** time. Select **Next**.

Add Tier

Transition
Time since ingest to move objects to this tier:
7 Days

Threshold
Percent utilization of the ingest tier that must be reached before content can be moved to the next tier:
0 %

Replication before Tiering
Move only replicated objects to this tier:
☐ Always replicate objects before tiering

Cancel < Previous **Next >** Finish

- d. (Optional) Select whether you want to enable rehydration. Select **Next**.
- e. Select the new tier to which to transition, select the DPL on your new tier and the number of metadata copies on primary storage. Select the warning check box to confirm. Select **Next**.

Add Tier

Storage Pools

Name	Type	Data	Compliant
<input type="checkbox"/> Primary running	Primary running		<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> s3_Pool	Amazon S3	1 copy	

Metadata Copies on Primary Storage
1 copy

Warning: All data copies for this tier will be held in extended storage. This storage should be protected and secure and should have its health monitored closely.
☒ I understand

Cancel < Previous **Next >** Finish

7. Select **Finish**.

You can now associate service plans with namespaces to allow further granular selection of which datasets are pushed to the cloud.

Schedule system services

System services can be run on demand but are typically defined in a service schedule which runs 24 hours a day, 7 days a week.

To create a system-wide service schedule, do the following.

Procedure


1. Log on to Hitachi Content Platform.
2. Select **Services**. Select **Create New Schedule**.



3. Type a schedule name and select from **Blank Schedule**. Select **Save**.
4. Select a cell, define the services along with the duration, and select **Update Schedule**.



5. Once defining your schedule, select **Activate Schedule**.



Configure system services


Once you have defined a system schedule for services that run system wide, configure parameter settings around services to fit organizational needs.

Compression and encryption settings

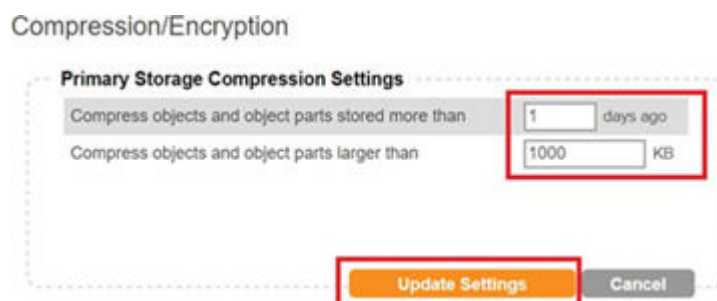
To modify compression and encryption settings, do the following.

Procedure

1. Log on to Hitachi Content Platform.
2. From **Services**, select **Compression/Encryption**.



3. Define the time duration and object size you want Content Platform on which to perform compression. Select **Update Settings**.



4. To exclude specific file types from compression, type a file type and select **Add**.




Content verification

Hitachi Content Platform supplies a hash-algorithm which the system continuously checks to make sure files have not become corrupted.

To define your content verification settings on how the system scans for objects, do the following.

Procedure

1. Log on to Hitachi Content Platform.
2. From **Services**, select **Content Verification**.
3. For Content Verification Mode, keep or select the default **Check all object and repair if needed** option.





Note: Changing this to the **Check only objects that can be repair and repair if needed** option or the **Do not check and repair objects** option results in less system overhead for Hitachi Content Platform. This lowers CPU cycles on the Hitachi Content Platform nodes, as the system knows what reparable objects it has ingested and will not spend cycles scanning. Select these options only if your organization needs it.

Duplicate elimination

Duplicate elimination is a service which strictly runs without any editing of its parameters. Hitachi Content Platform deduplicates like objects to save backend capacity on primary running storage. You can view statistics around duplicate elimination by selecting the Duplicate Elimination service.

Duplicate Elimination

Statistics	
Total objects and object parts merged:	95
Total bytes saved from duplicate elimination:	1.33 GB

Garbage collection

To define garbage collection settings for transaction records associated with the Hitachi Content Platform, do the following.

Procedure

1. Log on to Hitachi Content Platform.
2. From **Services**, select **Garbage Collection**.
3. As necessary, update the duration. By default, the deletion records are storage for 90 days.

Garbage Collection

☒ Keep deletion records in the transaction log for days

☐ Keep deletion records in the transaction log forever (Selecting this option may affect client performance over the long term.)

Update Settings Cancel

4. Select **Update Settings**.

Replication

Hitachi Content Platform allows for object replication to another Content Platform in case of a disaster at a site. If replicated, data is still available on the remote system. To onboard another Content Platform for replication and for more information on replication topologies, see [Related documents \(on page 9\)](#).

Shredding

The shredding service can be modified to make multiple passes over a backend disk to guarantee data is irrecoverable. To modify the shredding service, do the following.

Procedure

1. Log on to Hitachi Content Platform.
2. From **Services**, select **Shredding**.
3. From the **Shredding Rate** list, select the level needed.
4. Select **Update Settings**.




Tenants

Tenant creation is one of the main tasks of an administrator for Hitachi Content Platform. Creating a tenant allows enabling services which are passed down to the namespace, along with authentication types and capacity quotas.

Create a tenant From the system management console

To create a tenant on Hitachi Content Platform, do the following.

Procedure

1. Log on to Hitachi Content Platform.
 2. From the menu, select **Tenants**.
 3. Select **Create Tenant**.
- 
4. Type the **Tenant Name**.
 5. Set the **Hard Quota** for allocated capacity for the tenant.
 6. Set the **Soft Quota** as a warning threshold and set the **Namespace Quota**.
 7. For **Authentication Type**, select the **Local** check box, and define a **Username** and **Password** for the tenant administrator.
 8. In **Enable the features for this tenant**, pass down which services to the namespaces behind this tenant by selecting specific check boxes and which **Erasure Coding** option.
 - (Optional) If you plan on using cloud technologies as a potential tier, select the **All cloud-optimized namespaces** option.
 9. Select **Create Tenant**.

Create Tenant

Tenant Name:

Description (optional):

Initial Security Account:

Username:

Password:

Confirm Password:

Authentication Types:

☒ Local ☐ RADIUS ☐ Active Directory

Enable Features for this Tenant:

☒ Replication ☒ Retention Mode Selection ☒ Search ☒ Service Plan Selection ☒ Versioning

☐ Erasure Coding ☒ All Local optimized ☐ Namespaces ☐ Selected namespaces

Set tenant roles

To set a tenant's roles from the Tenant Management Console after creating it, do the following.

Procedure

1. Select the tenant from the available **Tenants** list. You are given a URL to open its **Tenant Management Console**.
2. Select the URL.

<https://hv.hcp1.hvlab.local:8000>

Features Enabled	Namespaces	Objects	Storage
Replication	Quota	Ingested	Quota
Retention Mode Selection	Used	Erasure coded	Used
Search	Available		Available
Service Plan Selection			
Versioning			

3. Log on to the **Tenant Management Console** using the local authentication credentials defined when creating the tenant.
4. If this is your first time opening this **Tenant Management Console**, type the **Existing Password** and a **New Password**. Confirm and select **Update Password**. Default settings in Content Platform require a password update the first time a console is opened.

Hitachi Content Platform
Tenant Management Console

Security Monitoring

You are required to change your password at this time.

Change Password for User: admin

Existing Password:

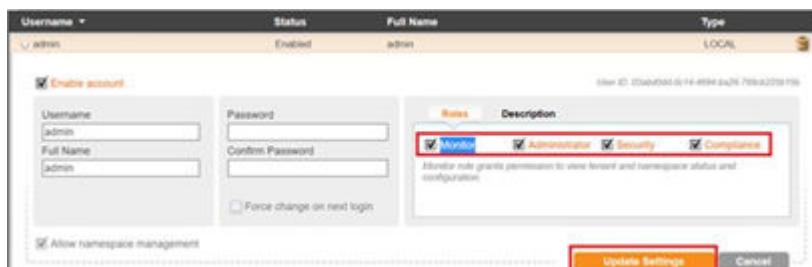
New Password:

Confirm New Password:

5. Assign permissions to the user who will administer this tenant. From **Security**, select **Users**.



6. Select your user, and select individual check boxes to set the **Monitor**, **Administrator**, **Security**, and **Compliance Roles**.
7. Select **Update Settings**.



The page reloads and provides the changed permissions to the user.



MK_SL_105

Namespaces are containers behind tenants which provide a data endpoint and allows administrators the granularity of assigning specific roles and feature sets to an individual namespace.

Create a namespace from the Tenant Management Console

To create a namespace from the Tenant Management Console, do the following.

Procedure

1. Log on to the **Tenant Management Console**
2. Select **Namespaces**.




3. Select **Create Namespace**.



4. Define settings for the namespace.
 - a. Type the **Namespace Name**.

- b. Select a **Hash Algorithm** for content verification.
- c. Set the **Hard Quota**, for the capacity you want to push to the namespace.
- d. Set the **Soft Quota**, for a warning level.
- e. Under **Configure the features for this namespace**, select the individual features for this namespace.
- f. (Optional) If you select **Versioning** as a feature for this namespace, set a duration of when to prune old versions.
- g. Select **Create Namespace**.



Create a namespace using Hitachi Ops Center Protector

To allow Ops Center Protector to natively create a NAMESPACE on Hitachi Content Platform, you must onboard the Hitachi Content Platform tenant as a node on the Protector user interface by designating its Tenant Management URL.

Enable namespace management APIs

You must enable the management API (MAPI) behind the namespaces used in conjunction with Protector. These MAPIs are used by Ops Center Protector to create Namespaces using REST.


To enable MAPIs at the namespace level, do the following.

Procedure

1. Log on to the **Tenant Management Console**.
2. From **Security**, select **MAPI**.



3. Select the **Enable management through APIs** check box.
4. Select **Update Settings**.



Configure namespaces

Namespaces allow granular configuration of how your data is handled and accessed through policies and services. Each namespace behind a common tenant can have different settings to satisfy client and organizational needs.

Policies

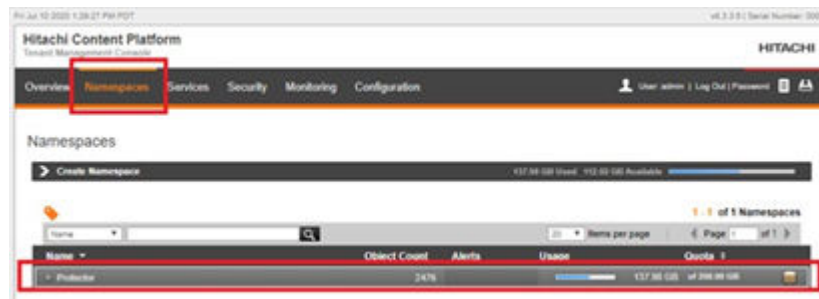
Policies are settings that influence how transactions, services, and internal processes work on objects. A setting can be a property of an object, such as retention, or a property of a namespace, such as versioning.

Indexing


Indexing is the Hitachi Content Platform policy that determines whether an object is included in the search index. To enable indexing behind a tenant, do the following.

Procedure

1. Log on to the **Tenant Management Console** for the tenant.
2. Select **Namespaces** and then select the specific namespace.



3. On the expanded Namespace, select the **Policies** tab.
4. Select the **Indexing** tab.
5. Select the **Index objects** check box. Select **Update Settings**.



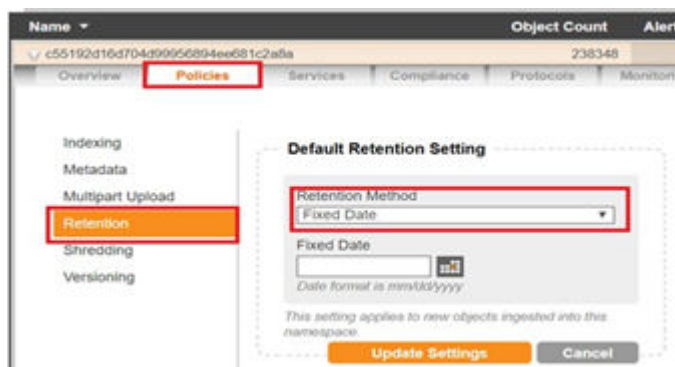
Fixed date retention

Objects have a retention property that determines how long the object must remain in the namespace before it can be deleted. This can range from allowing the object to be deleted at any time to preventing the object from ever being deleted. Retention to a fixed date can be set for object compliance.

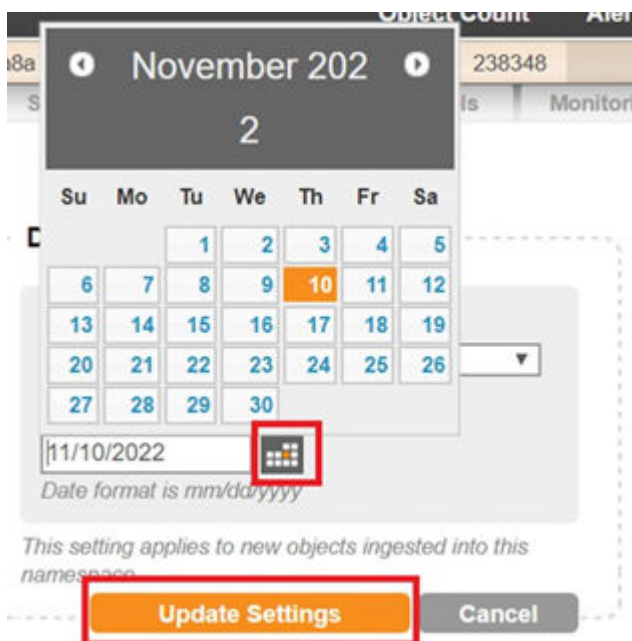
To specify a fixed date retention, do the following.

Procedure

1. Log on to the **Tenant Management Console** for the tenant.
2. Select **Namespaces** and select the specific namespace.
3. Select the **Policies** tab.
4. Select the **Retention** tab.
5. From **Retention Method**, select **Fixed Date**.



6. Select the **Calendar** button to choose the fixed date.
7. Select **Update Settings**.



Shredding

Shredding, also called secure deletion, is deleting an object and overwriting the places where its copies were stored so that none of its data or metadata, including custom metadata, can be reconstructed.

To enable shredding on a namespace, do the following.

Procedure

1. Log on to the **Tenant Management Console** for the tenant.
2. Select **Namespaces** and select the specific namespace.
3. Select the **Policies** tab.
4. Select the **Shredding** tab.
5. Select the **Shred on delete** check box.
6. Select **Update Settings**.



Versioning

Any given namespace can be configured to support versioning. When enabling versioning for a namespace, you can set a time for version pruning. Version pruning is the automatic deletion of previous versions of objects that are older than a specified amount of time.




Note: You cannot enable versioning for a namespace while the WebDAV, CIFS, NFS, or SMTP protocols are enabled for that namespace. Conversely, you cannot enable the WebDAV, CIFS, NFS, or SMTP protocol for a namespace while versioning is enabled for that namespace. While you can disable versioning at any time, this stops administrators from viewing versions of files, even the ones previously available with this setting on. Administrators can re-enable versioning at the namespace level to view previous file versions again.

To enable versioning on a namespace, do the following.

Procedure

1. Log on to the **Tenant Management Console** for the tenant.
2. Select **Namespaces** and select the specific namespace.
3. Select the **Policies** tab.
4. Select the **Versioning** tab.
5. Select the **Enable versioning** check box.
6. Define a time to prune older versions.
7. Select **Update Settings**.



Services

Services are responsible for optimizing the use of system resources and maintaining the integrity and availability of the data stored in the Hitachi Content Platform repository.

Disposition

Disposition is the automatic deletion of objects. Disposition can be enabled for objects with expired retention periods. This has the benefit of automatically freeing Hitachi Content Platform storage space for the creation of more objects. Without disposition, users need to explicitly delete qualified objects to free the occupied space.

To set disposition for a namespace, do the following.

Procedure

1. Log on to the **Tenant Management Console** for the tenant.
2. Select **Namespaces** and select the specific namespace.
3. Select the **Services** tab.
4. Select the **Disposition** tab.
5. Select the **Automatically delete objects with expired retention periods** check box.
6. Select **Update Settings**.



Replication

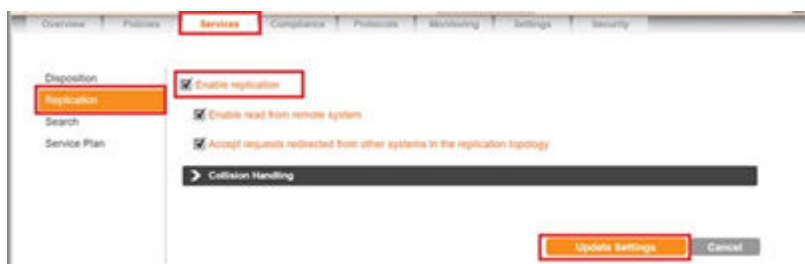
Replication is a process that supports configurations in which selected tenants and namespaces are maintained on two or more Hitachi Content Platform systems and the objects in those namespaces are managed across those systems. This cross-system management helps ensure that data is well-protected against the unavailability or catastrophic failure of a system.

In addition to replicating objects, Content Platform replicates tenant and namespace configuration, user and group accounts, retention classes, content classes, all compliance log messages, and most other tenant log messages.

To enable replication on a namespace if you have setup replication on the system management console, do the following.

Procedure

1. Log on to the **Tenant Management Console** for the tenant.
2. Select **Namespaces** and select the specific namespace.
3. Select the **Services** tab.
4. Select the **Replication** tab.
5. Select the **Enable Replication** check box.
6. Select **Update Settings**.



Service plan

Service plans allow administrators the ability to place their data on applicable tiers of storage, whether cloud or a local SAN infrastructure. To learn how to create a service plan using cloud adaptive tiering, see [Storage components for cloud adaptive tiering \(on page 56\)](#).

To apply a service plan to a namespace, do the following.

Procedure

1. Log on to the **Tenant Management Console** for the tenant.
2. Select **Namespaces** and then select the specific namespace.
3. Select the **Services** tab.
4. Select the **Service Plan** tab.
5. Select the **Service Plan** associated with applicable tier.
6. Select **Update Settings**.



Compliance

Hitachi Content Platform allows policy driven compliance settings to confirm entities are abiding by regulatory standards.

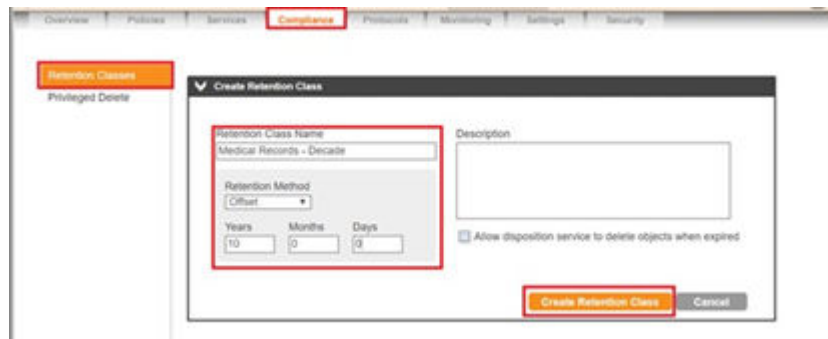
Retention classes

You can use retention classes to consistently manage data that must conform to a specific retention rule. For example, the law requires that medical records be kept for a specific number of years, use a retention class to enforce that requirement.

To create a retention class, do the following.

Procedure

1. Log on to the **Tenant Management Console** for the tenant.
2. Select **Namespaces** and select the specific namespace.
3. Select the **Compliance** tab.
4. Select the **Retention Classes** tab.
5. Select **Create Retention Class**.
6. Make these settings.
 - a. Type a **Retention Class Name**.
 - b. From the **Retention Method** list, select **Offset**.
 - c. Define duration in terms of **Years**, **Months**, and **Days**.
7. Select **Update Settings**.



Protocols and Data Access

Hitachi Content Platform namespace access protocols are categorized as either cloud protocols or non-cloud protocols. The cloud protocols are the REST, Amazon S3-compatible, and HSwift APIs. The non-cloud protocols are WebDAV, CIFS, NFS, and SMTP.

Protocol optimization improves namespace ingest performance. You can optimize namespaces for all protocols, which provides balanced performance across cloud protocols and non-cloud protocols.

Alternatively, you can optimize namespaces for cloud protocols only. Cloud-only optimization improves the ingestion rate of namespaces using cloud protocols. Cloud protocols, themselves, are further optimized for improved ingestion performance.

HTTP, HTTPS, and REST API protocols

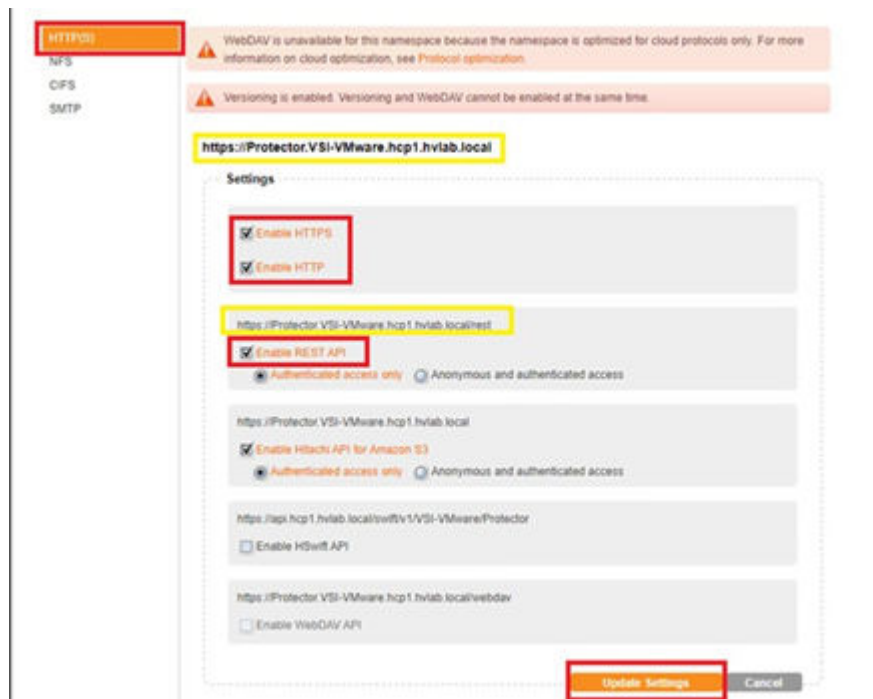
By using the HTTP, HTTPS, and the REST API protocols, you can store, view, retrieve, and delete objects. To enable HTTP, HTTPS and the REST API protocols, do the following.

Procedure

1. Log on to the **Tenant Management Console** for the tenant.
2. Select **Namespaces** and select the specific namespace.
3. Select the **Protocols** tab.
4. Select the **HTTP(S)** tab.
5. Under **Settings**, select these check boxes:
 - **Enable HTTPS**
 - **Enable HTTP**
 - **Enable REST API**

6. Select **Update Settings**.

Once enabled, you are presented with URLs to access using HTTP(S) as well as the Amazon S3 API and the REST API.



NFS

NFS is one of the industry-standard protocols Hitachi Content Platform supports for namespace access. To access a namespace through NFS either of the following ways:

- Write applications that use any standard NFS client library.
- Use the command line in an NFS client to access the namespace directly. To enable the NFS protocol, do the following.

Procedure

1. Log on to the **Tenant Management Console** for the tenant.
2. Select **Namespaces** and select the specific namespace.
3. Select the **Protocols** tab.

4. Select the **NFS** tab.
5. Select the **Enable NFS** check box.
6. Select **Update Settings**.

Once enabled, you are presented with URLs to access using NFS.



Note: NFS is unavailable if a namespace is optimized for cloud protocols only.

CIFS

You access a namespace through CIFS by mapping the namespace to a network drive or by adding the namespace as a network place on a CIFS client. You can have multiple directories mapped or added at the same time.

Once mapped or added, the namespace appears to be part of the local file system. You can perform any of the operations Hitachi Content Platform supports for CIFS.

To enable CIFS namespace access, do the following.

Procedure

1. Log on to the **Tenant Management Console** for the tenant.
2. Select **Namespaces** and select the specific **Namespace**.
3. Select the **Protocols** tab.
4. Select the **CIFS** tab.
5. Select the **Enable CIFS** check box.
6. Select **Update Settings**.

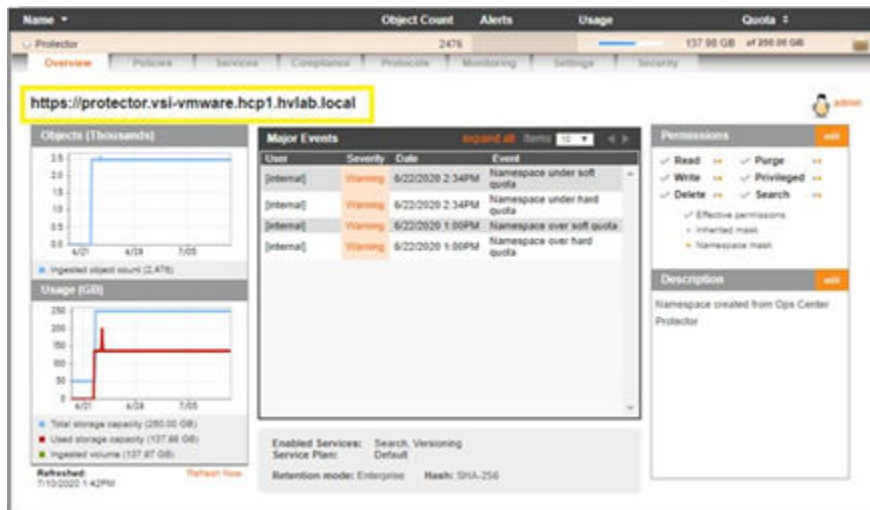
Once enabled, you are presented with URLs to access using CIFS.



Note: CIFS is unavailable if a namespace is optimized for cloud protocols only.

Natively access data using Tenant Management Console

When the HTTP or HTTPS protocols are enabled, by default the URL is available using the Overview tab on a namespace. Selecting this URL opens the namespace browser where you can view files, delete objects, view applied hash-algorithms, and see statistics.



Data migrator in Hitachi Content Platform

The data migrator in Hitachi Content Platform is a high-performance, multi-threaded, client-side utility for viewing, copying, and deleting data. Download the utility from the Tenant Management Console.

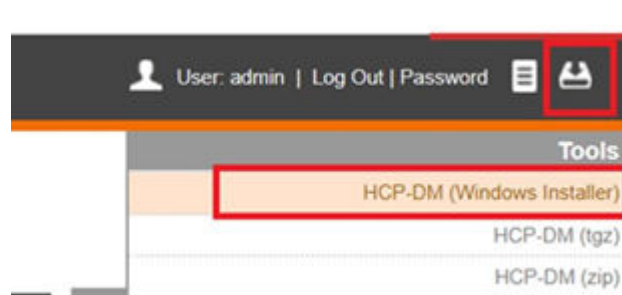
Download data migrator

To download the data migrator from the Tenant Management Console, do the following.

Procedure

1. Log on to any **Tenant Management Console**.
2. In the upper right corner, select the **Tools** button.
3. If you are using a Microsoft Windows or Windows Server computer, select **HCP-DM (Windows Installer)**.

For other operating systems, there are compressed files with the installation media.



Install and open data migrator

You need to download the data migrator utility installation file from the Tenant Management Console first.




Note: You must have Java installed prior to running Migrator.exe. For more information on prerequisites for the data migrator in Hitachi Content Platform, see [Related documents \(on page 9\)](#).

To install and open the data migrator, do the following.

Procedure

1. Locate and execute the HCPDM executable file. This installs the data migrator to the same directory as the executable file.




2. In the extracted directory, run Migrator.exe to open the data migrator.

Create a namespace profile

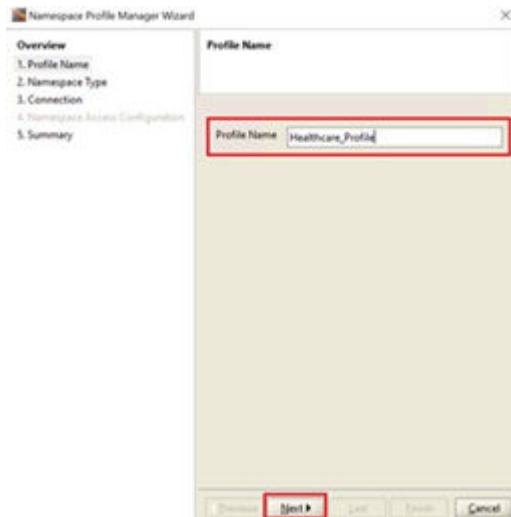
To begin moving data with the data migrator, you must onboard your namespace. To define a namespace as a migration target, do the following.

Procedure

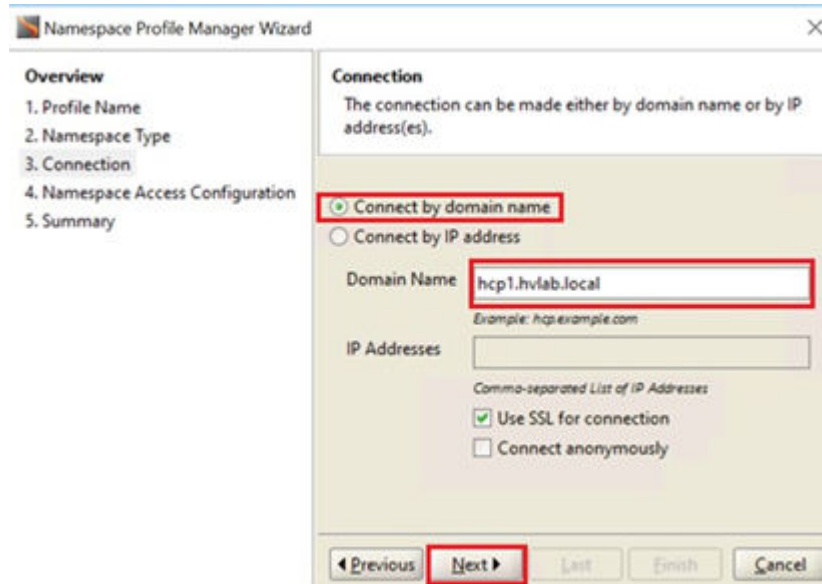
1. Run Migrator.exe.
2. From **File**, select **Namespace Profile Manager**.



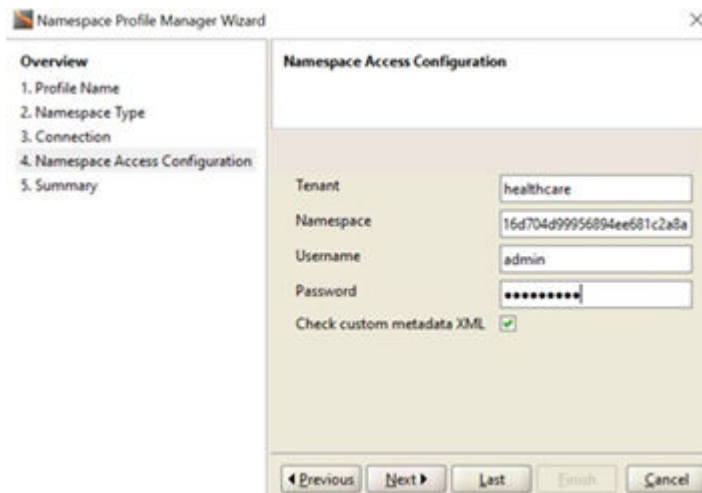
3. Select **Create**.
4. Type a **Profile Name**. Select **Next**.




5. Select **HCP namespace-6.0 or later**. Select **Next**.
6. Select the **Connect by domain name** option. Type your **Domain Name** for Hitachi Content Platform, as setup using DNS. Select **Next**.



7. Type the **Tenant, Namespace, Username, and Password**. Select **Next**.



8. To confirm connectivity to the namespace, select **Test**.
9. Select **Finish**.

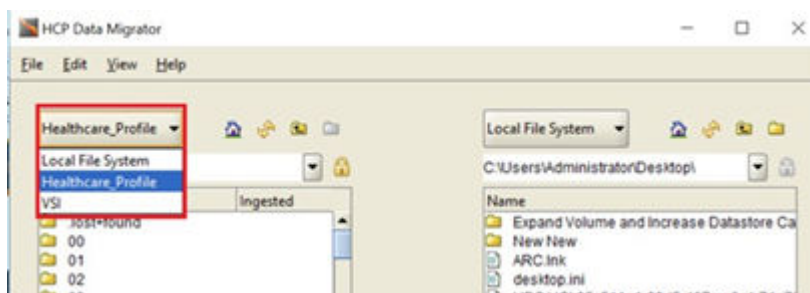


Move data with data migrator

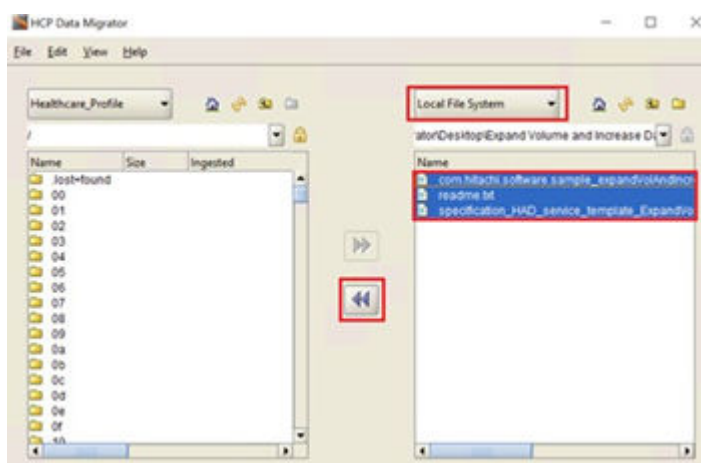
Once you have installed and configured the data migrator to communicate with a namespace, you can migrate data. To migrate data into a namespace with the data migrator, do the following.

Procedure

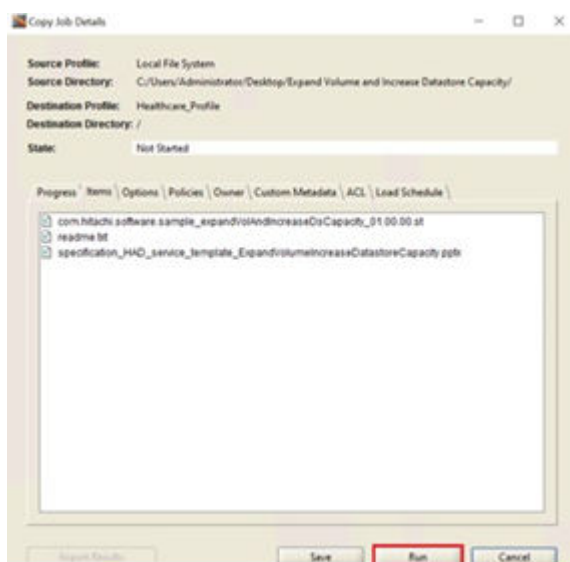
1. Run Migrator.exe as an administrator.
2. From the profile list on the left side, select the target namespace profile for migration.



3. For the profile list on the right side, keep the selection as **Local File System**.
4. Select the files you want to migrate to your namespace and then select the left arrows to move them to the namespace. This opens the **Copy Job Details** dialog box.



5. On **Copy Job Details**, confirm file selection and select **Run**. The job starts.



6. When the copy job is complete, close the **Copy Job Details** dialog box.

Using Hitachi Content Intelligence

The following section covers initial setup of Hitachi Content Intelligence Services.

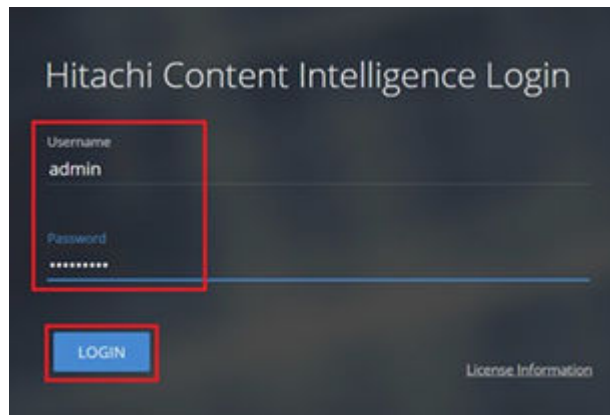
Access the Hitachi Content Intelligence WorkFlow Designer

Once Hitachi Content Intelligence is deployed, you can access the user interface to begin setting up data connections, processing pipelines and index collections. For information on how to install Hitachi Content Intelligence, see [Related documents \(on page 9\)](#).

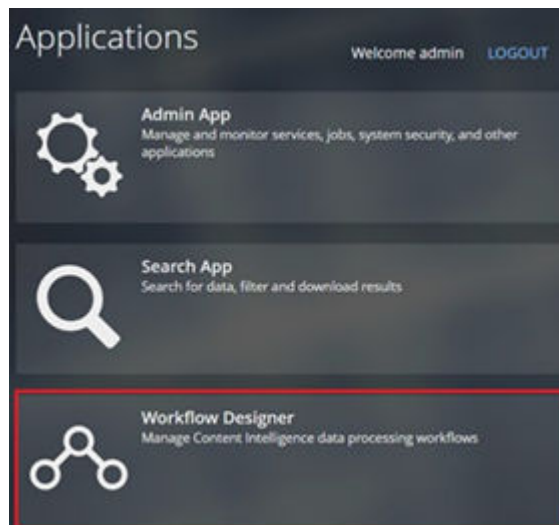
To open Hitachi Content Intelligence, do the following.

Procedure

1. In the Address bar of a browser, type the FQDN or IP address of the Hitachi Content Intelligence cluster:
https://Hitachi Content Intelligence FQDN or IP:8000
2. Enter your user name and password for the cluster. Select **Login**.



3. Select **Workflow Designer**.



Create data connections

A data source is a repository of files, such as a Hitachi Content Platform system. For your system to be able to access the files in a data source, create a data connection that contains all required access and authentication information for the data source.

Hitachi Content Intelligence system includes built-in data connections for connecting to several industry-standard data storage platforms. To connect to a type of system that it does not support, you can use the included software development kit (SDK) to write your own data connection plugins.

To onboard a Hitachi Content Platform system as your data connection, do the following.

Procedure

1. Log on to the **Workflow Designer**.
2. Select **Add Data Connection**.



3. To add a data connection, do the following:
 - a. From **Type**, select **HCP**.
 - b. Type a **Name** and **Description**.

 A screenshot of the 'Add Data Connection' configuration form. The form has a light blue header with the title 'Add Data Connection'. Below the header, there are four input fields:

- Type:** A dropdown menu with 'HCP' selected.
- Name:** A text input field containing 'HV-EngineeringNS'.
- Data source name:** A text input field containing 'Hitachi Vantara - Engineering Namespace'.
- Description:** A text input field containing 'Hitachi Vantara - Engineering Namespace'.

 At the bottom of the form, there is a label 'Data source description (optional)' which is currently empty.

4. For the connection information, do the following:
 - a. For **HCP System Name**, type the **domain name**.
 - b. Type the **HCP Tenant Name** and **HCP Namespace Name** exactly as it is on the Hitachi Content Platform system.
 - c. For **Use SSL**, select **Yes**.

Connection

HCP System Name
hcp1.hvlab.local
Name of the HCP system with domain (e.g. 'hcp1.hitachi.com')

HCP Tenant Name
hiv
Name of the HCP tenant

HCP Namespace Name
engineering
Name of the HCP namespace

HCP Root Directory
/
Directory path to crawl in the specified HCP namespace (e.g. '101/102')

Use SSL
☒

5. For authentication information, do the following:
 - a. Type the **User Name** and **Password** of the namespace as you would on the Hitachi Content Platform system to access the **Tenant Management Console**.
 - b. Select **Test**.

Authentication

User Name
admin
User name to use. Specify "all_users" and no password for anonymous access. To perform actions, this account must have the applicable HCP data access permissions.

Password

Password for the User

Test Create Cancel

6. If connection is successful, select **Create**.



Note: If the data connection is not successful, confirm MAPI is enabled on the namespace.

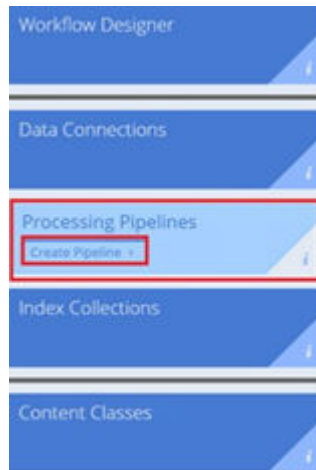
Create processing pipelines

Processing pipelines perform operations on the documents that the system extracts from your data sources. A processing pipeline is made up of one or more stages, each of which performs a specific type of operation.

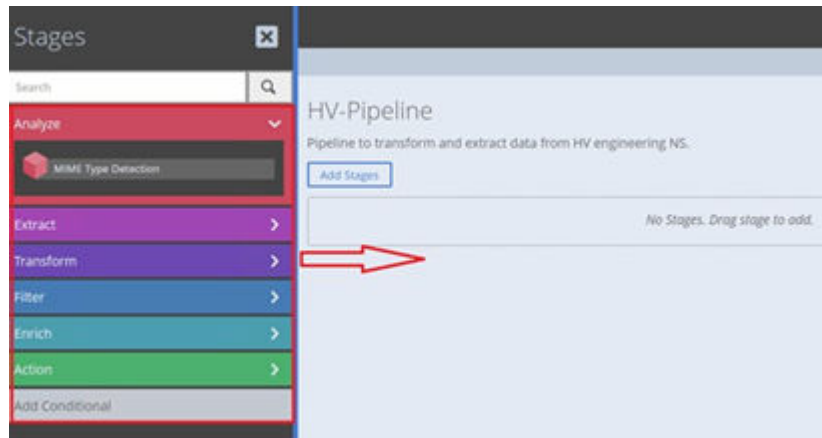
To create a processing pipeline, do the following.

Procedure

1. Log on to the **Workflow Designer**.
2. Select **Create Pipeline**.



3. Type the **Name** and **Description**. Select **Create**.
4. Select **Add Stages**. Your **Stages** are posted on the left side.
5. Drag stages that correlate to your data transformation goal to the right side.



6. Review settings for the respective stage and select **Update**.

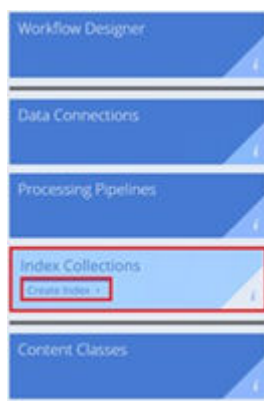
Create indexing collections

An index collection is a set of instructions for building a search index. Each index collection contains a schema, and query settings.

To create an index collection, do the following.

Procedure

1. Log on to the **Workflow Designer**.
2. Select **Create Index**.



3. Select **Hitachi Content Intelligence Index**.
4. Type and **Name** and **Description**. Keep the other default settings.
5. Select **Create**.

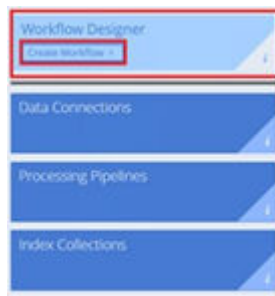
Create a workflow

Once you have created your data connections, processing pipelines, and index collections you can tie these components together using a workflow.

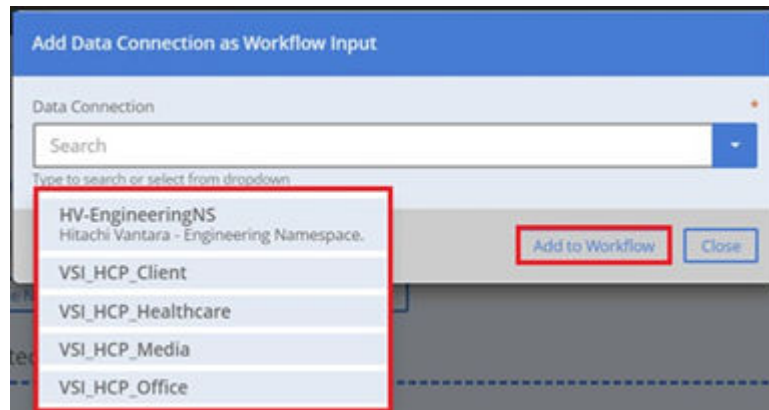
To create a workflow, do the following.

Procedure

1. Log on to the **Workflow Designer**.
2. Select **Create Workflow**.




3. Type a workflow **Name** and **Description**. Select **Next**.
4. Add a data connection.
 - a. Select **Select Data Connection**.
 - b. Select the **Data Connection** from the list.
 - c. Select **Add to workflow**.



d. Select **Close**. Select **Next**.

5. Add a pipeline.


- a. Select **Select Pipeline**.
- b. Select the **Pipeline** from the list and select **Workflow Agent**.
- c. Select **Add to workflow**.



d. Select **Close**. Select **Next**.

6. Select an index.

- a. Select **Select Index**.
- b. Select the **Index** from the list.
- c. From the **Selected Action** list, select **Index**.
- d. Select **Add to workflow**.



Add Index as Workflow Output

Index: HV_Index

Type to search or select from dropdown

Selected Action: Index

Indexes documents in Solr.

Add to Workflow Close

7. Select **Create**.

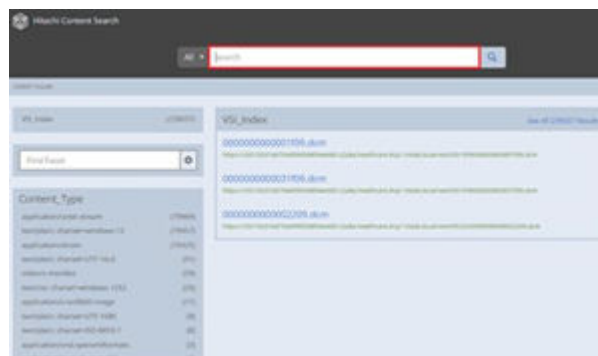
Access the Hitachi Content Intelligence Search console

Hitachi Content Intelligence includes a search app for querying your search indexes. Once you have created and executed your workflows you may query the index.

To access the Search console, do the following.

Procedure

1. In the Address bar of a browser, type the FQDN or IP address of the Hitachi Content Intelligence cluster:
`https:// Hitachi Content Intelligence FQDN or IP:8000`
2. Type your **Username** and **Password** for the cluster. Select **Login**.
3. Select **Search App**.
4. Use the **search bar** to query your indexes.



Solution references

For more information on Hitachi solutions and products please visit Hitachi Vantara.com along with the below solution references.

- Network
 - [Cisco Nexus 9000 Series Switches Data Sheets](#)
 - [Cisco MDS 9000 Series Multilayer Switches](#)
- Compute
 - [Cisco Unified Computing](#)
 - [Cisco UCS 6400 Series Fabric Interconnects Data Sheet](#)
 - [Cisco UCS 5100 Series Blade Server Chassis Data Sheet](#)
 - [Cisco UCS VIC 1440 Adapter Data Sheet](#)
 - [Cisco UCS Manager](#)
- Storage
 - [Hitachi Virtual Storage Platform F Series All-Flash Enterprise Cloud Solutions](#)
 - [Hitachi Virtual Storage Platform G Series Hybrid-Flash Midrange Cloud Solutions](#)
- Virtualization Layer
 - [VMware vCenter Server](#)
 - [VMware vSphere](#)
- Compatibility Matrixes
 - [Hitachi Interoperability Reports](#)
 - [VMware Compatibility Guide](#)
 - [Cisco UCS Hardware and Software Compatibility](#)

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