HITACHI
Inspire the Next

Hitachi Content Platform

HCP 7.3.1 Release Notes

HCP Software Version 7.3.1.35
HCP Operating System Version 7.2.1.149
April 5, 2017

Hitachi Data Systems

RN-RC001-47
5 April 2017

Contents

About this document

This document contains release notes for release 7.3.1 of **Hitachi Content Platform** (**HCP**). It describes new features, product documentation, and known issues, and provides other useful information about this release of the product.

Note: Throughout this document, the word *Unix* is used to represent all UNIX-like operating systems (such as UNIX® itself or Linux®), except where Linux is specifically required.

Release highlights for HCP 7.3.1

Release 7.3.1 of HCP fixes issues found in release 7.3 and earlier releases.

Resolving unavailable objects
With release 7.3.1 of HCP, the new fast object recovery service speeds up the clearing of unavailable status flags from objects which are no longer unavailable. An unavailable objects table has also been added to the Tenant Management Console.

Defining Active Directory domain controller filters
As of release 7.3.1 of HCP, system-level users with the security role can define their Active Directory domain controller filters through the Active Directory Configuration page in the System Management Console.

Moving database to optimal volumes
New with release 7.3.1 of HCP, service users can move an HCP node database to a larger available volume through the HCP Service menu to protect against volume capacity issues.

Improving non-replicating, temporary issue object reporting
As of release 7.3.1 of HCP, users can now report objects which have not been replicated due to temporary issues. A non-replicating objects table has also been added to the Tenant Management Console.

New field in the HCP MIB
The HCP 7.3.1 MIB contains a new field for the HCP node model. This field lists the node models in the system.

Release highlights for HCP 7.3

Release 7.3 of HCP fixes issues found in release 7.2.3 and earlier releases.

Swapping legacy hardware for G10 nodes
New with release 7.3 of HCP, you can swap legacy HCP nodes in SAIN systems with Quanta G10 nodes. The G10 nodes offer better performance and have more configuration options than the nodes previously used in HCP systems. To swap a node, first you use the HCP Service menu in the system console to run the node swap procedure. Then you physically replace the old node with the new node. Finally, you return to the system console to complete the node swap procedure.

Sending data access log messages to the syslog servers
As of release 7.3 of HCP, system-level users can choose to send log messages about HTTP data access events to specified syslog servers.

Changing system-level passwords with the management API
With release 7.3 of HCP, system-level users with the security role can now use the management API to change the password for any locally authenticated system-level user account, including their own accounts.

Expanding Keystone role support for HSwift
As of release 7.3, HCP now supports the _member_ and Member Keystone roles for HCP tenant data access.

New field in the HCP MIB
The HCP 7.3 MIB for use with SNMP contains a new field, syslogSendHttpAccessLogs. This field specifies whether HCP should send log messages about HTTP data access events to syslog servers.

Release highlights for HCP 7.2.3

Release 7.2.3 of HCP fixes issues found in release 7.2.2 and earlier releases.

Updated self-signed SSL certificate
Before release 7.2.3, HCP came with a self-signed SSL server certificate that used the SHA-1 algorithm. Because SHA-1 is being generally deprecated, all self-signed certificates in HCP release 7.2.3 and later use the SHA-256 algorithm.

Existing SHA-1 SSL certificates are not automatically changed to use SHA-256 when an HCP system is upgraded to release 7.2.3. If you were using a SHA-1 certificate prior to the upgrade, you need to generate a SHA-256 SSL server certificate for the HCP system after the upgrade is complete to ensure access to the HCP Management Consoles.

If the HCP system remains at a release version earlier than 7.2.3, the HCP Management Consoles and Namespace browser will not work with new versions of certain web browsers.

Improved PUT Object - COPY requests
PUT Object - COPY requests have been improved for HCP S Series ingest tiers, Zero Day Transition time S3 cloud storage pools and Azure storage pools. A PUT Object - COPY request sent to any of those storage pools now has to perform fewer read and write operations, resulting in faster completion.

Release highlights for HCP 7.2.2

Release 7.2.2 of HCP fixes issues found in release 7.2.1.

Release highlights for HCP 7.2.1

Release 7.2.1 of HCP fixes issues found in release 7.2 and earlier releases. Release 7.2.1 also introduces new VMware compatibility and HCP S Series Node tiering improvements.

In release 7.2.1, all newly installed HCP-VMs are configured to use VMXNET3 adapters. VMXNET3 supports both 1G and 10G network configurations. Existing HCP-VMs can also be configured to use VMXNET3 adapters.

With release 7.2.1, HCP-VM software is qualified to be hosted on ESXi 6.

With release 7.2.1, the maximum supported LUN size for HCP-VM software has increased from four to 16 terabytes.

With release 7.2.1, HCP improves Data at Rest and Data in Flight encryption performance for S Series Nodes and cloud storage components.

With release 7.2.1, HCP improves HNAS and HDI read performance from S Series Nodes when compression and encryption are disabled on S Series storage components.

Release highlights for HCP 7.2

The following sections describe the new features and enhancements included in release 7.2 of HCP.

Quanta D51B-2U servers

New with release 7.2, HCP introduces two new node models: HCP G10 Nodes with Local Storage and HCP G10 Nodes with Attached Storage. Both models use Quanta D51B-2U servers, which are versatile computing platforms that can assimilate to any existing HCP system configuration. The Quanta servers come with additional networking configuration options and optional SSD storage. The G10 Nodes with Local Storage can also be upgraded for additional storage.

Both node models are compatible with older generations of HCP systems. The G10 Nodes with Local Storage can be configured to any existing HCP RAIN system architecture, and G10 Nodes with Attached Storage can be configured to any existing HCP SAIN system architecture.

New network configuration options
HCP G10 Nodes support multiple variants of 1G and 10G network configurations. The network configurations can use either or both CAT-6 and SFP+ cabling. Release 7.2 of HCP boasts new 10G BASE-T front-end connectivity options and mixed 10G and 1G back-end options for maximum network versatility. 10G back-end networks are only supported with SFP+ cabling.

SSD
HCP G10 Nodes accommodate for an optional pair of 800GB SSDs. SSDs store metadata indexes on faster storage than traditional internal hard disk drives, which increases the overall performance of the HCP system.

Storage additions for HCP G10 Nodes with Local Storage
HCP G10 Nodes with Local Storage can be equipped with a second set of hard drives. The hard drives can be configured to act as a second RAID set which provides the node with extra storage capacity.

USB and DVD OS and software installs

New with release 7.2 of HCP, the HCP Appliance Operating System and HCP software installation can be performed with a single USB flash drive. This eliminates the need for an external USB CD drive and multiple installation CDs.

Rackless HCP G10

With release 7.2, new HCP systems RAIN and SAIN system can be ordered in a racked or rackless configuration. The systems are shipped with the latest HCP G10 Nodes, and come bundled with back-end switches, PDUs, and cables along with any other optional equipment specialized for the order, such as storage arrays and Fibre Channel switches.

HCP S Series storage integration enhancements

Release 7.2 offers enhancements to HCP S Series Node integration.

Direct write to HCP S Series Nodes
With release 7.2 of HCP, you can configure a service plan ingest tier to target HCP S Series Nodes instead of primary storage. All objects ingested through namespaces under the modified service plan are immediately sent from HCP to the specified S Series storage pools. Regardless of the ingest tier, HCP keeps all object metadata on primary running storage.

If one or more namespaces on a service plan are configured to allow CIFS, NFS, or WebDAV, the service plan ingest tier cannot be changed to HCP S Series Nodes. Likewise, NFS, CIFS, and WebDAV cannot be enabled on namespaces that have HCP S Series Nodes set as the ingest tier.

S30 Node integration
New with release 7.2, HCP introduces S30 Node integration. The HCP S Series Nodes, known as economy storage, serve as storage tiering platforms or alternative primary ingest tiers for HCP systems. The S30 node consists of two cooperating server modules that are standalone servers and

multiple high density disks in up to 16 enclosures. S Series Nodes use commodity hardware to ensure that the costs of growth and repair remain low.

S3 metadata
New with release 7.2 of HCP, HCP S Series and S3 compatible components use S3 metadata headers to store object metadata. The enhancement improves small object performance by reducing the number of operations required to store an object on an S3 or S Series tier.

AWS authentication enhancements

Release 7.2 offers enhancements to AWS authentication.

S3 signature V4
With release 7.2 of HCP, HCP introduces AWS Signature Version 4 signing process on inbound S3 gateways and S3 compatible storage components. With this enhancement HCP can tier to AWS regions created after January 30, 2014 including the eu-central-1 region. This protocol is more secure than the previous AWS Signature Version 2 signing process.

AWS security token service
With release 7.2 of HCP, you can now enable the AWS Security Token Service for your S3 compatible storage components. The service provides temporary, limited-privilege credentials which protect your cloud storage component credentials.

Replication verification service

With release 7.2 of HCP, you can enable the replication verification service to safeguard against data incongruity between systems in your replication topology. The service iterates every object in your HCP system and batch queries all other systems in the replication topology to confirm that they possess copies of your objects. If an object is missing from a replica system, the object is re-replicated. Any objects which cannot be replicated are listed in the Tenant Management Console **Namespace** page under the **Non Replicating Objects** tab. The replication verification service can be configured to run once or continuously.

Active Directory authentication enhancements

New with release 7.2, HCP fully supports Active Directory Single Sign On authentication with SPNEGO. If HCP is configured to use the feature, users logged into AD on their workstations can use Single Sign On with the HCP System Management Console, Tenant Management Console, Namespace Browser, and MQE Search UIs. Applications coded to use AD credentials with SPNEGO are now also supported.

A new REST API authentication header has been added to HCP cloud access protocols (HCP REST API, HS3, and HSwift) which allows users and applications to use AD credentials to authenticate with their storage tiers.

Active Directory Multi-Forest support

New with release 7.2, HCP supports multiple Active Directory forests. If the system is configured to use the feature, HCP system and tenant administrators can grant management and data access roles to groups in trusted forests.

Load Balancer REST API

New with release 7.2 of HCP, a /node_status resource has been added to the management API that can be used to check the health status of an HCP system or particular node. This resource is intended for use by load balancers to better manage pools of HCP systems.

Webhelp

New with release 7.2 of HCP, documentation has been converted from PDFs to Webhelp. The new format offers a sleeker interface and compiles the manuals into a single file. As a result, all of the books are cross-searchable.

The Webhelp can be downloaded or opened in a separate browser tab. Any section in the manuals can be linked to other users and viewed without logging into the System Management Console or Tenant Management Console.

Log download enhancements

New with release 7.2, you can specify the types of HCP system logs you want to download. By filtering out logs that are unnecessary to your use case, you reduce the size of the logs and the time it takes to download. You can perform the log download procedure through the HCP System Management Console or management API.

The system logs can be filtered by node type, specified node numbers, and log type. The possible log types are:

- **Access** — These logs capture all user requests related to the HTTP(s) gateway (REST, HS3, HSwift, MAPI), such as object **GET**, **PUT**, and **DELETE** requests.
- **System State** — These logs capture the current state of the HCP system OS and application states and provide useful information for diagnosing issues related to node restarts and unexpected service behaviors.
- **Service Procedure** — These logs capture information about service procedures performed on HCP such as the add node procedure, upgrade procedure, or add LUN procedure.
- **HCP Software** —These logs collect information about HCP software not listed in the other log types.

Miscellaneous enhancements

HCP release 7.2 offers these additional enhancements.

Removed support for TrueCopy Storage migration
HDS TrueCopy© is no longer supported as an array migration mechanism. Please refer to the *Administering HCP* manual for details on other supported methods of migrating arrays.

New and changed fields in the HCP MIB

The HCP 7.2 MIB for use with SNMP contains the new fields shown in the table below.

Field	Description
SSD volumes	
volumeSpinStateString	Reports if the volume is on a disk or SSD.
storageStatus	Reports if a SSD is degraded.
Replication verification service information	
nonReplicatingIrreparableObjects	Shows the total number of objects in the HCP system that currently cannot be replicated because they are irreparable.
nonReplicatingOpenObjects	Shows the total number of objects in the HCP system that currently cannot be replicated because they are open.
replicationLinkVerificationMode	Shows the current setting for replication verification service on the HCP system.
replicationLinkLastVerificationTime	Shows the last time a verification pass completed for the replication link.
Economy storage information	
economyStorageUsage	Shows the total amount of used space on all HCP S Series Nodes.
economyStorageCapacity	Show the total amount of storage capacity available on all HCP S Series Nodes.

Release highlights for HCP 7.1.2

Release 7.1.2 of HCP fixes issues found in release 7.1.1.

Release highlights for HCP 7.1.1

Release 7.1.1 of HCP fixes issues found in release 7.1.

Release highlights for 7.1

The following sections describe the new features and enhancements included in release 7.1 of HCP.

HCP S Series Node integration

New with release 7.1, HCP introduces HCP S Series Node integration. The HCP S Series Nodes, known as economy storage, serve as storage tiering platforms for HCP systems. Objects tiered to the S Series Nodes fall under the economy storage license. For more information about S Series Node licensing, see "Storage licensing" on the facing page.

HCP uses the S3 API, to write, retrieve, and otherwise manage objects on S Series Nodes. A single HCP system can seamlessly tier data across multiple S Series Nodes, thereby enabling scalability in both capacity and performance.

When an S Series Node is added to the HCP system, HCP creates an economy storage component in the System Management Console. The HCP S Series Nodes are a means of expanding the storage available to HCP. They do not function as backup storage.

Tiering to economy storage
The storage tiering service is responsible for moving object content between HCP storage and economy storage devices. It does this according to the rules defined in the service plans.

Information about the state of your economy storage components is displayed on the **Hardware ▸ Node** and **Storage ▸ Components** pages in the HCP System Management Console.

Economy storage pools
An economy storage pool is a grouping of one or more economy storage components. HCP creates economy storage pools when an HCP S Series Node is added to the system. When HCP creates an economy storage component, it adds the component to a storage pool.

In a service plan for which the rule is to move objects to economy storage, the rule specifies the target storage pool. The storage tiering service can move objects to any of the economy storage components in that pool.

Objects in economy storage
When moving an object to economy storage, the storage tiering service moves only the object data and a limited amount of system metadata. HCP keeps the rest of the metadata, including custom metadata, for the object in its own storage.

Depending on the DPL of the service plan, the object data in economy storage may be the only copy of the data. Alternatively, it may be one of multiple copies, with the other copies stored in HCP storage.

Encryption and compression
The configuration of the economy storage component specifies whether object data stored on that storage pool should be encrypted and/or compressed. If encryption is selected, the storage tiering service encrypts object data before writing it to the volume. When reading encrypted data from an economy storage component, HCP automatically decrypts the data.

If compression is selected, the storage tiering service compresses the object data before writing it to the volume. When reading compressed data from an economy storage component, HCP automatically decompresses the data.

Economy storage statistics
The economy storage statistics section is on the **Storage ▸ Overview** page in the HCP System Management Console. This page shows the total number of objects with data currently stored in economy and extended storage. It also shows the total number of bytes of object data currently stored in economy and external storage. These statistics are aggregates for all namespaces in a single HCP system.

Write and read statistics about any economy storage component can be found on the **Storage ▸ Components** page of the System Management Console. This information is presented in four graphs which show write and read processes in bytes and operations.

Storage licensing

With release 7.1, all HCP systems are required to have a storage license. Storage licenses are available for all system configurations: RAIN, SAIN, and VM. These licenses permit the use of a designated amount of data that can be stored on external devices.

The licenses do not prohibit you from using more storage space than they allot, but exceeding a license causes a warning message to appear on the System Management Console **Overview** page under System Status. Monitoring mechanisms such as syslog, system log messages, and SNMP send reminders every 24 hours to upload a new storage license.

Default licenses
Newly installed systems come with a storage license that has two terabytes of active storage and two terabytes of extended storage. HCP systems upgrading to version 7.1 of HCP or newer receive an unlimited storage license for both active and extended storage that lasts for one year. Contact your HDS account representative to collect your current license or procure a new license.

Uploading and monitoring licenses
You can upload a new license or view details about your current license from the **Storage ▸ Licenses** page in the System Management Console. The page displays information about your current license type, serial number, Quote number, capacity, expiration date, and status. From this page you can view your license history.

License types
The following are the three different types of storage licenses:

Active storage license — Covers the used primary storage capacity of an HCP system. This accounts for all objects stored on primary storage. If objects stored on primary storage are moved to an HCP S Series Node or other extended storage device, they are then covered by a different license.

Economy storage license — Covers the objects stored on HCP S Series Nodes. Although this license can be monitored through the HCP system, economy license management is done through the HCP S Series Management Console.

Extended storage license — Covers objects stored on devices outside of the HCP system. The devices include the following types of extended storage: Amazon S3, Google Cloud, Hitachi Cloud, Microsoft Azure, S3-compatible, and Verizon Cloud.

Advanced downstream DNS configuration

With release 7.1 of HCP, you can enable advanced downstream DNS configuration mode. Advanced downstream DNS configuration mode can be activated through the Management API, and it lets you directly access your

HCP DNS management files for any forward facing network, including [hcp_system]. With this mode enabled, you can modify the **named.config** entry and **Forward Zone** file for the network. In addition, you can enable TSIG by updating the zone entry.

Once enabled, an *advanced* downstream DNS configuration panel replaces the **Downstream DNS Configuration** panel of any individual forward facing network available on the **Network ▸ Network View** page. Instead of containing configuration options, the panel has a **Zone Entry** and **Forward Zone** field.

Using advanced downstream DNS configuration mode is only recommended if you have an extensive background in networking. If you need to remove your changes, it's suggested to revert your advanced downstream DNS configuration mode to its basic setting.

Increased total number of supported downstream DNS servers

Downstream DNS configuration basic mode now supports 32 downstream DNS servers. Advanced downstream DNS configuration mode does not limit your downstream DNS servers.

Isolating networks for tiering to storage components

With release 7.1 of HCP, you can use virtual networks for storage tiering. By assigning a storage component to a network, you can configure HCP to exclusively communicate with that component over the assigned network. Isolating a network increases tiering security by segregating the tiering process to individual networks instead of having all data tiered over one network.

Networks reserved for storage tiering can still be configured to perform replication. They can also serve as tenant data and management networks.

Protocol optimized namespaces

In release 7.1, HCP introduces the option to optimize namespaces for cloud access protocols. HCP has always optimized namespaces for balanced performance when accessed by any supported protocols, taking into account both cloud protocols (HCP REST API, HS3, and HSwift) and non-cloud protocols (CIFS, NFS, WebDAV and SMTP). Today, HCP retains the option

for balanced performance but also lets you optimize namespaces that only use cloud protocols to be further optimized for improved ingest performance.

Optimizing namespaces for cloud, in conjunction with upgrading HCP memory, can improve ingest performance of an HCP system with a high object and directory count.

OpenStack Swift API

Release 7.1 of HCP introduces the HCP OpenStack Swift API, also called HSwift. HSwift is a RESTful, HTTP-based API that's compatible with OpenStack Swift.

With HSwift, you work with containers and objects. Containers are equivalent to HCP namespaces. Containers created using HSwift can be managed through the HCP Tenant Management Console and management API, just like any other namespace.

Objects stored in containers can be accessed in the same ways that objects in other namespaces can be accessed — through any namespace access protocol, the Namespace Browser, the metadata query API, and the Search Console.

For you to create and manage containers using the HSwift API, the HCP management API must be enabled for the tenant.

What you can do with HSwift
Using the HSwift API, you can:

- Create a container (PUT Container)
- List existing containers (GET Account)
- Check the existence of a container and retrieve its metadata (HEAD Container)
- Set ACLs on a container (POST Container)
- List the objects in a container (GET Container)
- Delete an empty container (DELETE Container)
- Store an object and its metadata in a container (PUT Object)
- Replace object metadata (POST Object)

- Retrieve object metadata (HEAD Object)
- Copy an object and its metadata (COPY Object)
- Retrieve an object and its metadata (GET Object)
- Delete an object and its metadata (DELETE Object)

Keystone
HCP supports Keystone authentication. Keystone is an OpenStack identity service that provides token-based authentication. Keystone generates tokens that are used to authenticate users attempting to store and manage containers and objects in HCP.

Enabling HSwift for existing namespaces
You can enable the HSwift API for existing namespaces in the HTTP(S) panel of the namespace. To enable the HSwift API, you first need to open the HTTPS or HTTP port, the same way you do for the HTTP protocol (REST API).

Updates to the HCP Management API

The HCP management API has been enhanced to support other features that are new with HCP release 7.1.

New namespace property
The namespace data type has a new property, **optimizeFor**, which specifies whether the namespace is optimized for cloud protocols only or optimized for all protocols.

New system level Network resource
A new system level resource, Network, lets you toggle between advanced and basic downstream DNS configuration modes.

New networkSettings property
The networkSettings data type has a new property, **downstreamDNSMode**, which specifies whether the downstream DNS configuration is set to basic or advanced mode.

New Licenses resource for storage
A new resource for storage, Licenses, lets you retrieve information about a current storage license and upload new licenses.

New Licenses property
The Licenses data type has a new property, **Licenses**, which breaks up information about all the storage licenses listed.

License data type
The new License data type has the following property:

- **activeCapacity** — Specifies the active storage capacity in terabytes
- **expirationDate** — Specifies the storage license expiration date
- **extendedCapacity** — Specifies the extended storage capacity in terabytes
- **quoteNumber** — Specifies the storage license quote number
- **serialNumber** — Specifies the serial number of the HCP system the storage license is intended for
- **uploadDate** — Specifies the date which the license was uploaded

httpProtocol data type
The new httpProtocol data type has the following property:

hswiftEnabled — Specifies whether the HSwift API is enabled for the namespace.

hswiftRequiresAuthentication — Specifies whether user authentication is required or optional for access to the namespace through the HSwift API.

Miscellaneous enhancements

HCP release 7.1 offers these additional enhancements.

Verizon Cloud adaptive tiering
The list of current cloud services has been expanded to include Verizon Cloud.

New and changed fields in the HCP MIB
The HCP 7.1 MIB for use with SNMP contains the new fields shown in the table below.

Field	Description
Optimization	
protocolOptimization	Shows the current protocol optimization setting of the HCP system.
Storage license information	
storageActiveLicensedCapacity	The total active storage capacity licensed to your HCP system in gigabytes.
storageActiveUsedCapacity	The used active storage capacity licensed to your HCP system in gigabytes.
storageExtendedLicensedCapacity	The total extended storage capacity licensed to your HCP system in gigabytes.
storageExtendedUsedCapacity	The used extended storage capacity licensed to your HCP system in gigabytes.
storageLicenseExpirationDate	The storage license expiration date.
storageLicenseSerialNumber	The serial number on the storage license.
Advanced downstream DNS configuration	
advancedDNSNetwork	Shows whether advanced downstream DNS configuration mode is enabled.
Status of the OpenStack Identity Service	
osIdentityServiceEnabled	OpenStack Identity Service status.

Release highlights for HCP 7.0.1

Release 7.0.1 of HCP fixes several issues found in release 7.0.

Release highlights for HCP 7.0

The following sections describe the new features and enhancements included in release 7.0 of HCP.

Replication enhancements

Release 7.0 of HCP includes several replication-related enhancements.

Active/active links
New with release 7.0 of HCP, you can create replication links that replicate data in both directions between two systems. With such links, called **active/active links**, the tenants and namespaces being replicated are read-write on both systems. Active/active links are designed for use in a cloud storage environment, where applications need seamless access to namespaces regardless of where those applications are located.

After upgrading to release 7.0, you can leave existing replication links, called **active/passive links**, as they are, where the tenants and namespaces being replicated are read-write on only one system. Alternatively, you can change some or all of these links to active/active links.

Link creation wizard
In release 7.0 of HCP, link creation has been simplified by the introduction of a wizard. The wizard walks you through the process of creating the link itself. Selecting the content to be replicated is now separate from the link creation. For active/active links, you can use the HCP System Management Console for one of the systems involved in the link to select the content to be replicated from both systems.

Automatic collision handling
Because the tenants and namespaces involved in an active/active link are read-write on both systems, collisions can occur between changes made on the two different systems. The way HCP handles collisions depends on the type of collision. The general rule for namespace content collisions is that more recent changes have priority over conflicting less recent changes. The exception to this is retention changes. HCP always retains the longer retention period and keeps held objects on hold.

Object content collisions occur when a objects with the same name but different content are created on both systems. In this case, if versioning is not enabled, the more recently created object keeps its name and location, and the other object is either moved or renamed, depending on the setting for a new namespace configuration option.

When configuration changes made on one system collide with changes made on the other system, HCP automatically pauses replication of the applicable tenant.

Automatic failover and failback
HCP release 7.0 allows you to configure both active/active and active/passive links to automatically fail over to the other system after one system (the primary system for active/passive links) has been unavailable for a specified amount of time. You can also configure active/passive links to automatically switch to the complete recovery phase of data recovery when the up-to-date-as-of time for the link is less than a specified amount of time.

Automatic sharing of domains and certificates
With DNS failover, client requests to a failed system that identify that system by domain name are automatically redirected to another system in the replication topology. If the client request specifies HTTPS in the URL, the system to which the request is redirected must have an SSL server certificate for the domain specified in the request.

With release 7.0 of HCP, to ensure that the systems in a replication topology can service redirected HTTPS requests, you can configure each system to periodically send all the domains and SSL server certificates it has, including those shared with it by other systems, to each other system with which it participates as a sending system in a replication link.

HCP management API support for replication
You can now use the HCP management API to configure, monitor, and manage replication links. This new feature enables you to programmatically perform the same replication-related tasks as you can perform from the HCP System Management Console.

Time out-of-sync alert
A new HCP System Management Console alert warns you when the system times for the two systems involved in a replication differ by more than one minute. To avoid this situation, all HCP systems in a replication topology should use the same external time server.

Adaptive cloud tiering

As of release 7.0, HCP can tier objects not only to NFS storage volumes but also to the cloud and to any Amazon® S3™-compatible storage device. The currently supported cloud services are Amazon S3, Microsoft® Azure Blob Storage, Google Cloud Storage™, and Hitachi Cloud Service. Collectively, these types of storage, including NFS storage volumes, are called **extended storage**.

Adaptive cloud tiering enables storage capacity to be scaled as needed and reduces the costs of managing and retaining data. During their lifecycles, objects can be moved automatically among different storage locations with differing price/performance characteristics.

Storage components and storage pools
The physical storage that's managed by the nodes in the HCP system is called **primary storage**. By default, primary storage consists entirely of primary running storage. However, an HCP SAIN system can also be configured to use primary spindown storage for tiering purposes.

An HCP system can also be configured to use additional storage that's managed by devices outside of the HCP system. HCP uses this additional storage, called **extended storage**, for tiering purposes.

In the System Management Console, HCP uses **storage components** to represent primary running storage, primary spindown storage, and each physical device and cloud storage service that's used to access a specific type of extended storage. Each storage component is intended to represent all storage devices that share a common access point (whether that's the HCP system, an external device, or a cloud storage service) and that provide a specific set of data availability, price, and performance characteristics.

HCP uses **storage pools** to represent logical groups of storage components that can be used as storage tiers. Each storage pool consists of one or more storage components that are used to access the same type of storage. Each storage tier typically consists of only one storage pool, but a tier can be configured to use multiple storage pools. To store objects on a given tier, HCP uses all of the storage that's accessed using the storage components that are contained in the storage pools that are configured for the storage tier.

HCP uses predefined storage components and storage pools for primary running storage and primary spindown storage. However, to enable HCP to access any given type of extended storage device or cloud storage service, you need to define an extended storage component for that device or service. To define an **extended storage component**, you first need to specify the type of extended storage that's represented by the component (Amazon S3, Google Cloud, Hitachi Cloud Service, Microsoft Azure, S3-compatible, or NFS). You then need to specify the information that HCP needs to use to connect to the extended storage device or cloud storage service that's represented by the component. Finally, you need to configure

HCP to use one or more **access points** (mount points, buckets, containers, or namespaces) to access the physical storage that will be used to store object data.

To enable HCP to use the storage that's represented by a specific extended storage component, you need to define an **extended storage pool** that includes one or more extended storage component access points. You then need to configure a service plan to define a storage tier that contains the extended storage pool.

Enhanced service plans
Where the data for an object is stored during the object lifecycle is determined by the storage tiering strategy specified in the service plan that applies to the namespace that contains the object. A **storage tiering strategy** specifies storage tiers and transition criteria for moving objects from one tier to another.

A storage tier consists of one or more storage pools. The first tier in a storage tiering strategy is called the **ingest tier**. This tier always consists of only one pool, the primary running storage pool. You can configure any service plan to define one or more additional storage tiers. When defining a storage tier, you can specify whether objects become metadata-only on that tier. You can also specify rehydration behavior for objects on that tier.

Each service plan also specifies a **data protection strategy**. This strategy determines the number of copies of the object data that must be maintained on the storage that's represented by each storage pool that's defined for each storage tier. (Data protection strategies are a replacement for namespace DPL settings.)

System metadata, custom, metadata, and the metadata query engine index are always stored in primary running storage, thereby allowing objects to be easily indexed and quickly accessed in response to queries. For each storage tier, you can specify the number of copies of custom metadata to be maintained on primary running storage.

HCP release 7.0 includes wizards that step you through the process of creating service plans and defining storage tiering and data protection strategies for service plans.

Easier service plan assignment
If you are a system-level administrator, you can now assign a service plan to multiple tenants in a single operation. If you are a tenant-level administrator, you can now assign a service plan to multiple namespaces in

a single operation.

Retiring storage
If you no longer want to use storage in which objects have already been stored, you can retire that storage. You can retire extended storage pools, storage components, or individual storage volumes. Retiring storage causes HCP to remove all data from that storage. Once a pool, component, or volume has been retired, it can be deleted from HCP.

You can also retire primary storage by retiring nodes or storage arrays. Retiring primary storage is equivalent to performing a data migration with the migration service.

Abandoning storage
HCP lets you abandon storage components. Abandoning a storage component causes the component to be immediately deleted from HCP without first having data removed from it. This procedure is a last resort for deleting a storage component and should be used only after the data stored on the component is guaranteed to have copies in other locations.

Retiring and deleting service plans
You can now retire service plans. Retired service plans are enforced for the namespaces to which they apply but that they cannot be associated with any additional tenants or namespaces unless you choose to remove them from retirement.

You can delete a service plan only if it is not currently associated with any tenants or namespaces. If you try to delete a service plan that's in use, HCP prompts you for a replacement plan. The replacement plan is then associated with the tenants and namespaces that the now deleted plan was associated with.

Service plans after upgrades
As of HCP 7.0, both the namespace DPL setting and the system-level DPL setting have been replaced by the storage tier DPL settings specified in service plans. During an upgrade of an HCP system to release 7.0, preexisting service plans will be updated and new service plans will be automatically generated so that all previous combinations of storage tiering strategies and namespace DPL settings will be accounted for. The generated service plans will be marked retired.

Storage and service plan usage information
The new **Storage Overview** page in the HCP System Management Console reports on the usage of storage components, storage pools, and service plans. Other new pages that are specific to storage components and storage pools provide additional usage details.

Support for IPv6

New with release 7.0, HCP front-end networks can use IPv6 addresses for communication with other systems or devices. Each HCP front-end network has now an **IP mode** setting. The IP mode for a front-end network can be set to IPv4, IPv6, or Dual. If the IP mode is set to IPv4 or IPv6, the network can be configured to use only IPv4 addresses or only IPv6 addresses, respectively. If the IP mode is set to Dual, the network can be configured to use both IPv4 and IPv6 addresses. A network that supports IPv6 can be configured with both primary and secondary IPv6 addresses.

The IP mode options that are available to be selected for a front-end network depend on the HCP system configuration. Additionally, you can configure a user-defined network to use a specific type of IP address only if the [hcp_system] network is also configured to use that type of IP address.

Improved Active Directory reporting

HCP release 7.0 provides detailed information about issues that occur with the integration of HCP with Microsoft Active Directory®. New alerts warn of error conditions. New error messages and system log messages contain diagnostic information that you can use resolve problems. Additionally, the **Authentication** page in the HCP System Management Console now displays system log messages that relate to Active Directory while Active Directory is selected as the authentication method.

To enable better monitoring of the connection between HCP and Active Directory, the frequency with which HCP checks the health of that connection has been increased.

HCP search facility end of life

As of release 7.0, support for the HCP search facility has been removed from HCP. All options related to the HCP search facility have been removed from the HCP System and Tenant Management Consoles, and only the metadata query engine and the HDDS search facility are available to work with the HCP Search Console.

Important: Before HCP is upgraded to release 7.0, search nodes must be removed from the system. These nodes can then be repurposed to act as storage nodes.

7.0 miscellaneous enhancements

HCP release 7.0 offers these additional enhancements.

Single pruning option
With HCP release 7.0, you no longer need to provide a separate primary system and replica version pruning settings for namespaces. Instead, you specify a single version pruning that applies to the namespace on all HCP systems on which the namespace exists.

Namespace DPL moved to service plans
As of release 7.0 of HCP, DPL is no longer a namespace setting. Instead, the number of copies of object data that HCP must maintain for each object in a given namespace is determined by the service plan that applies to that namespace.

Additionally, HCP no longer has a system-level DPL setting. Instead, every service plan has a default ingest tier data protection setting (called **the ingest tier DPL**). This setting specifies the number of copies of object data that HCP must maintain on primary running storage. For an HCP RAIN system or an HCP-VM system, the default ingest tier DPL is two. For an HCP SAIN system, the default ingest tier DPL is one. You can configure any service plan to increase the ingest tier DPL, but you cannot specify an ingest tier DPL that's lower than the default value.

In the HCP management API, starting in HCP 7.0, the DPL property has been deprecated, so you can no longer use this property to set the DPL for any new or existing namespaces. However, to maintain backwards compatibility with applications running on HCP 6.0 and earlier releases, any attempt to use the DPL property to set the DPL of a namespace does not result in an error. Instead, the DPL property is simply ignored when it's specified on a PUT or POST request.

If an existing application currently uses the HCP management API to set the DPL for a given namespace or to set the default DPL for namespaces owned by a given tenant, then that application needs to be rewritten to use the service plan for each namespace to control the DPL for objects in that namespace.

Monitoring metadata query engine indexing with the HCP management API
In HCP release 7.0, you can use the HCP management API to monitor metadata query engine indexing. The new mqeIndexingTimestamp property of the namespace data type specifies the date and time before which objects are guaranteed to have been indexed by the metadata query engine.

Faceting on system metadata with the metadata query API
HCP release 7.0 has expanded the list of system metadata properties on which you can facet metadata query API query results. In the facets entry for an object-based query, you can now specify any system metadata property for which valid values numbers, dates, or Boolean values.

Enhanced chargeback report options in the System Management Console
New with release 7.0 of HCP, you can specify the reporting interval and time period for chargeback reports you generate from the HCP System Management Console.

Enhancements to downloading the HCP internal logs
As of release 7.0 of HCP, when downloading the HCP internal logs, you can choose to download the logs either for all nodes or only for selected nodes. Additionally, the downloaded logs are no longer encrypted.

Enhancement to HCP protection service processing
HCP uses the protection service to maintain the correct number of copies of the data and metadata for each object in the HCP repository to satisfy the data protection level (DPL) and metadata protection level (MPL) settings that are applied to each object.

When the number of copies of the data or metadata for an object goes below the number of copies that's required to satisfy the current DPL or MPL setting (respectively) that's applied to that object, the HCP protection service creates a new copy of the object data or metadata. When the number of copies of the data or metadata for an object goes above the number of copies that's required to satisfy the DPL or MPL setting (respectively) that's applied to that object, the HCP protection service deletes the extra copy of the object data or metadata.

In releases prior to HCP 7.0, when the protection service created a new copy of the data or metadata for an object, that object was marked as serviced. However, when the protection service deleted a copy of the data or metadata for an object, the service did not mark that object as serviced.

As of release 7.0, the protection service now marks an object as serviced each time the protection service creates or deletes a copy of the data or metadata for that object.

New and changed fields in the HCP MIB
The HCP 7.0 MIB for use with SNMP contains the new fields shown in the table below.

Field	Description
Node information	
nodeIPv6	The primary IPv6 address of the node
nodeIPv6Sec	The secondary IPv6 address of the node
Storage component information	
storageComponentName	The name of the storage component
storageComponentType	The storage component type
storageComponentStatus	The status of the storage component
Storage pool information	
storagePoolName	The name of the storage pool
storagePoolType	The storage pool type
storagePoolStatus	The status of the storage pool
Service plan information	
servicePlanCount	
Replication link information	
replicationLinkType	An indication of whether the replication link is active/passive or active/active

Related documents

The following documents contain additional information about Hitachi Content Platform:

- *Administering HCP* - This book explains how to use an HCP system to monitor and manage a digital object repository. It discusses the capabilities of the system, as well as its hardware and software components. The book presents both the concepts and instructions you

need to configure the system, including creating the tenants that administer access to the repository. It also covers the processes that maintain the integrity and security of the repository contents.

- *Managing a Tenant and Its Namespaces* - This book contains complete information for managing the HCP tenants and namespaces created in an HCP system. It provides instructions for creating namespaces, setting up user accounts, configuring the protocols that allow access to namespaces, managing search and indexing, and downloading installation files for HCP Data Migrator. It also explains how to work with retention classes and the privileged delete functionality.

- *Managing the Default Tenant and Namespace* - This book contains complete information for managing the default tenant and namespace in an HCP system. It provides instructions for changing tenant and namespace settings, configuring the protocols that allow access to the namespace, managing search and indexing, and downloading installation files for HCP Data Migrator. It also explains how to work with retention classes and the privileged delete functionality.

- *Replicating Tenants and Namespaces* - This book covers all aspects of tenant and namespace replication. Replication is the process of keeping selected tenants and namespaces in two or more HCP systems in sync with each other to ensure data availability and enable disaster recovery. The book describes how replication works, contains instructions for working with replication links, and explains how to manage and monitor the replication process.

- *HCP Management API Reference* - This book contains the information you need to use the HCP management API. This RESTful HTTP API enables you to create and manage tenants and namespaces programmatically. The book explains how to use the API to access an HCP system, specify resources, and update and retrieve resource properties.

- *Using a Namespace* - This book describes the properties of objects in HCP namespaces. It provides instructions for accessing namespaces by using the HTTP, WebDAV, CIFS, and NFS protocols for the purpose of storing, retrieving, and deleting objects, as well as changing object metadata such as retention and shred settings. It also explains how to manage namespace content and view namespace information in the Namespace Browser.

- *Using the HCP HS3 API* - This book contains the information you need to use the HCP HS3 API. This S3™-compatible, RESTful, HTTP-based API enables you to work with buckets and objects in HCP. The book introduces the HCP concepts you need to understand in order to use HS3 effectively and contains instructions and examples for each of the bucket and object operations you can perform with HS3.

- *Using the HCP OpenStack Swift API* - This book contains the information you need to use the HCP HSwift API. This OpenStack Swift, RESTful, HTTP-based API enables you to work with containers and objects in HCP. The book introduces the HCP concepts you need to understand in order to use HSwift effectively and contains instructions and examples for each of the container and object operations you can perform with HSwift.

- *Using the Default Namespace* - This book describes the file system HCP uses to present the contents of the default namespace. It provides instructions for accessing the namespace by using the HCP-supported protocols for the purpose of storing, retrieving, and deleting objects, as well as changing object metadata such as retention and shred settings.

- *HCP Metadata Query API Reference* - This book describes the HCP metadata query API. This RESTful HTTP API enables you to query namespaces for objects that satisfy criteria you specify. The book explains how to construct and perform queries and describes query results. It also contains several examples, which you can use as models for your own queries.

- *Searching Namespaces* - This book describes the HCP Search Console (also called the Metadata Query Engine Console). It explains how to use the Console to search namespaces for objects that satisfy criteria you specify. It also explains how to manage and manipulate queries and search results. The book contains many examples, which you can use as models for your own searches.

- *Using HCP Data Migrator* - This book contains the information you need to install and use HCP Data Migrator (HCP-DM), a utility that works with HCP. This utility enables you to copy or move data between local file systems, namespaces in HCP, and earlier HCAP archives. It also supports bulk delete operations and bulk operations to change object metadata. Additionally, it supports associating custom metadata and ACLs with individual objects. The book describes both the interactive window-based interface and the set of command-line tools included in HCP-DM.

- *Installing an HCP System* - This book provides the information you need to install the software for a new HCP system. It explains what you need to know to successfully configure the system and contains step-by-step instructions for the installation procedure.

- *Deploying an HCP-VM System* - This book contains all the information you need to install and configure an HCP-VM system. The book also includes requirements and guidelines for configuring the VMWare® environment in which the system is installed.

- *Third-Party Licenses and Copyrights* - This book contains copyright and license information for third-party software distributed with or embedded in HCP.

- *HCP-DM Third-Party Licenses and Copyrights* - This book contains copyright and license information for third-party software distributed with or embedded in HCP Data Migrator.

- *Installing an HCP SAIN System - Final On-site Setup* - This book contains instructions for deploying an assembled and configured single-rack HCP SAIN system at a customer site. It explains how to make the necessary physical connections and reconfigure the system for the customer computing environment. It also contains instructions for configuring Hi-Track® Monitor to monitor the nodes in an HCP system.

- *Installing an HCP RAIN System - Final On-site Setup* - This book contains instructions for deploying an assembled and configured HCP RAIN system at a customer site. It explains how to make the necessary physical connections and reconfigure the system for the customer computing environment. The book also provides instructions for assembling the components of an HCP RAIN system that was ordered without a rack and for configuring Hi-Track Monitor to monitor the nodes in an HCP system.

Upgrade notes

You can upgrade an HCP system to release 7.3.1 only from release 6.0 or later. In order to upgrade to release 7.3.1, you need at least 12GB of RAM, although it is recommended to have 32GB, prior to upgrading. You cannot downgrade HCP to an earlier release.

HCP upgrades can occur with the system either online or offline. During an online upgrade, the system remains available to users and applications. Offline upgrades are faster than online upgrades, but the system is unavailable while the upgrade is in progress. Work with your authorized service provider to determine which type of upgrade is better for you .

Note: During an online upgrade, data outages may occur as each node is upgraded. Whether data users are affected by an outage depends on the ingest tier DPL setting specified in the service plan that's assigned to the applicable namespace. No data is lost during a data outage, but users may experience some interruptions to data access.

Supported limits

HCP supports the maximum values listed in the table below.

Item	Limit
Hardware	
Maximum number of storage nodes	80
Maximum number of HCP S Series Nodes	80
Logical storage volumes	
SAIN systems	
Maximum number of data and shared volumes per storage node (SAIN)	63
Maximum logical volume size (SAIN)	15.999 TB
Maximum number of index volumes per node (SAIN)	27
Maximum number of spindown volumes per node (SAIN)	31
Maximum number of NFS volumes per node (SAIN)	31
RAIN systems	
Maximum number of data and shared volumes per storage node (RAIN)	4

(Continued)

Item	Limit
Maximum logical volume size (RAIN)	Depends on HDD capacity
Maximum number of NFS volumes per node (RAIN)	31
HCP-VM systems	
Maximum number of logical volumes per storage node (VM)	59 data LUNs
Maximum logical volume size VMDK for ESXi 5.5 or later	16 TB minus 1 GB
Data storage	
Maximum number of objects per storage node	800,000,000
Maximum number of objects per HCP system	64,000,000,000 (80 nodes times 800,000,000 objects per node)
Maximum number of directories per node if one or more namespaces are not optimized for cloud	1,500,000
Maximum number of directories per node if all namespaces are optimized for cloud	15,000,000
Maximum number of objects per directory	30,000,000
Maximum object size	By protocol: • HTTP: About 2 TB (2,194,719,883,008 bytes) • WebDAV: About 2 TB (2,194,719,883,008 bytes) • CIFS: 100 GB • NFS: 100 GB • HSwift: 5 GB • HS3: 5 GB

(Continued)

Item	Limit
Maximum number of tenants	1,000
Maximum number of namespaces	10,000
Maximum number of attachments per email for SMTP	50
Maximum aggregate email attachment size for SMTP	500 MB
User/Group accounts	
Maximum number of system-level user accounts per HCP system	10,000
Maximum number of system-level group accounts per HCP system	100
Maximum number of tenant-level user accounts per tenant	10,000
Maximum number of tenant-level group accounts per tenant	100
Maximum number of users in a username mapping file (default tenants only)	1,000
Custom metadata	
Maximum number of annotations per individual object	10
Maximum non-default annotation size with XML checking enabled	1 MB
Maximum default annotation size with XML checking enabled	1 GB
Maximum annotation size (both default and non-default) with XML checking disabled	1 GB
Maximum number of XML elements per annotation	10,000
Maximum level of nested XML elements in an annotation	100

(Continued)

Item	Limit
Maximum number of characters in the name of custom metadata annotation	32
Metadata query engine	
Maximum number of content classes per tenant	25
Maximum number of content properties per content class	100
Maximum number of concurrent metadata query API queries per node	5
Network	
Maximum number of user-defined networks (VNeM networks) per HCP system	200
Maximum downstream DNS servers	32
Maximum certificates and CSR per domain	10
Tiering	
Maximum number of storage components	100
Maximum number of storage pools	100
Maximum number of tiers in a service plan	5
Miscellaneous	
Maximum number of HTTP connections per node	255
Access Control Lists	
Maximum size of access control entries per ACL	1,000 MB

Supported clients and platforms

The following sections list clients and platforms that are qualified for use with HCP.

Windows clients

These Microsoft Windows 32-bit or 64-bit clients are qualified for use with the HTTP v1.1, WebDAV, and CIFS protocols and with the HCP HS3 API:

- Windows 7
- Windows 8
- Windows 2008 R2 (Standard and Enterprise Server editions)
- Windows 2012 (Standard and Datacenter editions)
- Windows 2012 R2 (Standard and Datacenter editions)

Note: Using the WebDAV protocol to mount a namespace as a Windows share can have unexpected results and is, therefore, not recommended.

Unix clients

These Unix clients are qualified for use with the HTTP v1.1, WebDAV, and NFS v3 protocols and with the HCP HS3 API:

HP-UX® 11i v1 (11.11) on Itanium®

HP-UX 11i v1 (11.11) on PA-RISC®

IBM AIX 7.1

Red Hat® Enterprise Linux ES 6.3

Sun Solaris® 10 SPARC®

Sun Solaris 11 SPARC

Note: HCP does not support NFS protocol v4.

Browsers

The table below lists the web browsers that are qualified for use with the HCP System Management, Tenant Management, and Search Consoles and the Namespace Browser. Other browsers or versions may work but have not been formally tested.

Browsers	Client Operating System
Internet Explorer® 11*	Windows
Firefox® 38.x.xesr, 40	Windows HP-UX IBM AIX Red Hat Enterprise Linux Sun Solaris
*The Consoles and Namespace Browser work in Internet Explorer only if ActiveX is enabled. Also, the Consoles work only if the security level is **not** set to high.	

Note: To correctly display the System and Tenant Management Consoles and the Namespace Browser, the browser window must be at least 1,024 pixels wide by 768 pixels high.

Note: Internet Explorer compatibility view mode may work but is not supported by HCP.

Client operating systems for HCP Data Migrator

These client operating systems are qualified for use with HCP Data Migrator:

- Microsoft 32-bit Windows:
  - Windows XP Professional
  - Windows 2003 R2 (Standard and Enterprise Server editions)
  - Windows 2008 R2 (Standard and Enterprise Server editions)

  - Windows 7
  - Windows 8
  - Windows 2012 (Standard and Datacenter editions)
- HP-UX 11i v1 (11.11) on Itanium
- HP-UX 11i v1 (11.11) on PA-RISC
- IBM AIX 7.1
- Red Hat Enterprise Linux ES 5 (32-bit)
- Red Hat Enterprise Linux ES 6.3 (64-bit)
- Sun Solaris 10 SPARC
- Sun Solaris 11 SPARC

Note: The Oracle Java Runtime Environment (JRE) version 7 update 6 or later must be installed on the client.

Platforms for HCP-VM

HCP-VM runs on these platforms:

- VMware ESXi 5.5
- VMware ESXi 6.0

Third-party integrations

The following third party applications have been tested and proven to work with HCP. HDS does not endorse any of the applications listed below, nor does HDS perform ongoing qualification with subsequent releases of the applications or HCP. Use these and other third party applications at your own risk.

HSwift tools

These tools are qualified for use with the HCP HSwift API:

- OpenStack Horizon UI

- OpenStack "swift" Python Client
- OpenStack Glance
- Cyberduck OpenStack Swift Browser
- Cloudberry Explorer for OpenStack Storage
- Duplicity Backup

HS3 tools

These tools are qualified for use with the HCP HS3 API:

- CloudBerry Explorer
- Cyberduck
- DragonDisk
- s3cmd
- s3Curl
- s3fs-c (works only with versioning enabled on the target bucket)

Mail servers

These mail servers are qualified for use with the SMTP protocol:

- Microsoft Exchange 2007
- Microsoft Exchange 2010 (64 bit)

NDMP backup applications

These NDMP backup applications are qualified for use with HCP:

- Hitachi Data Protection Suite 8.0 SP4 (CommVault® Simpana® 8.0)
- Symantec® NetBackup® 7 — To use NetBackup with an HCP system:
  - Configure NDMP to require user authentication (that is, select either the **Allow username/pwd authenticated operations** or **Allow digest-**

authenticated operations option in the **NDMP** protocol panel for the default namespace in the Tenant Management Console).

- Configure NetBackup to send the following directive with the list of backup paths:

```
set TYPE=openPGP
```

Windows Active Directory

HCP is compatible with Active Directory on servers running Windows Server 2003 R2, Windows Server 2008 R2, or Windows Server 2012 R2. In either case, all domain controllers in the forest HCP uses for user authentication must minimally be at the 2003 functional level.

Note: If the functional level of the AD domain controllers was raised to 2003 after Windows Server was upgraded from version 2000, the SSL certificate you export from AD for use with HCP must have been generated after the functional level was raised.

RADIUS protocols

HCP supports the following RADIUS protocols:

- CHAP
- EAPMD5
- MSCHAPv2
- PAP

Supported hardware

The following sections list hardware that is supported for use in HCP systems.

Note: The lists of supported hardware are subject to change without notice. For the most recent information on supported hardware, please contact your HCP sales representative.

Supported servers

These servers are supported for HCP RAIN systems:

Quanta D51B-2U
Hitachi® CR 220S
Hitachi CR 220

These servers are supported for HCP SAIN systems without internal storage:

Hitachi CR 210H
Hitachi CR 220
Hitachi CB 320

These servers are supported for HCP SAIN systems with internal storage:

Quanta D51B-2U
Hitachi CR 210H (with 1Gb Ethernet)
Hitachi CR 220S (with 10Gb Ethernet)

Server memory

It is required to have at least 32GB of RAM per node to use new software features introduced in release 7.0 or HCP or later. An HCP system can be upgraded to version 7.*x*.*x* of HCP with a minimum of 12GB of RAM per node and receive the patches and bug fixes that come with the upgrade, but the system cannot utilize the software features new to the release. Inadequate DRAM causes performance degradation and can negatively affect system stability. Please contact your HDS account team if you have less than 32GB RAM per node and would like to upgrade to a 7.*x*.*x* release.

Supported storage platforms

These storage platforms are supported for HCP SAIN systems:

Hitachi AMS 2100
Hitachi AMS 2300
Hitachi AMS 2500
Hitachi Unified Storage (HUS) 110
Hitachi Unified Storage (HUS) 130
Hitachi Unified Storage (HUS) 150
Hitachi Unified Storage (HUS) VM
Hitachi Unified Storage (HUS) T3
Hitachi VSP

Hitachi VSP-G 200
Hitachi VSP-G 400
Hitachi VSP-G 600
Hitachi VSP-G 1000
Hitachi VSP-G 1500

Supported back-end switches

These switches are supported for the back-end network in HCP systems:

Alaxala AX2430
Brocade® VDX® 6720 — SAIN systems only
Brocade® VDX® 6740
Brocade® ICX® 6430
Brocade® ICX® 6430-48
Cisco® Nexus® 5548UP
Cisco® Nexus® 5596UP
Dell PowerConnect™ 2824 (firmware version 1.0.0.42, A04)
HP ProCurve 4208VL (firmware version L.11.24)

Fibre Channel switches

These Fibre Channel switches are supported for HCP SAIN systems:

Brocade 5100
Brocade 6510
Cisco 9134
Cisco 9148
Cisco 9148S

Fibre Channel host bus adapters

These Fibre Channel host bus adapters (HBAs) are supported for HCP SAIN systems:

- Emulex® LPe 11002-M4
  (firmware version 2.82a4, boot BIOS 2.02a1)
- Emulex LPe 12002-M8
  (firmware version 1.10a5, boot BIOS 2.12a15)
- Emulex LPe 12002-M8 (GQ-CC-7822-Y)
  (firmware version 1.10a5, boot BIOS 2.02a2)

- Hitachi FIVE-EX 8Gbps
  (firmware version 10.00.05.04)

Issues resolved in this release

The table below lists issues that were resolved in this release of HCP. The issues are listed in order by reference number.

Ref. number	Description
HCP-28521	**False reports of missing service principal names (SPNs) for tenants and namespaces from chained replication links** When SSO is set up in a replication topology that contains chained replication links, certain HCP systems may report that the SPNs of tenants and namespaces on those links are missing, even though they are not. **Fix:** HCP no longer falsely displays the missing SPN warning.
HCP-29129	**Retired internal logs do not fail back after zero copy failover (ZCF)** When a node in a SAIN system fails, the logical volumes managed by that node fail over to the ZCF peer node for the failed node. When the failed node resumes operation, it's logical volumes fail back. If a log rotate event occurs during the failover, the peer node keeps the retired internal logs of the repaired node. **Fix:** The failed node now reclaims its retired internal logs from the peer node when its logical volumes fail back.
HCP-29241	**Acknowledging irreparable objects triggers non-replicating object alert** When HCP finds an irreparable object, the object is flagged as irreparable and non-replicating. If the object is permanently irreparable, on the **Tenant Management Console►Namespaces►Irreparable Objects** table you can acknowledge the object, which should remove the irreparable and non-replicating object alerts. If you acknowledge the irreparable object, HCP does not remove the non-replicating object alert. **Fix:** When you acknowledge an irreparable object, HCP now removes both the irreparable and non-replicating object alerts.
HCP-29279	**HCP system performance slowed due to repeated logged warning message** On an HCP G10 with Attached Storage system with multipathing enabled, HCP issues an unnecessary warning message to the internal logs every ninety seconds. Over time, the repeated message can bloat internal logs and affect system performance. **Fix:** HCP G10 with Attached Storage systems with multipathing enabled no longer issue the warning message to the internal logs.

Ref. number	Description
HCP-29285	**Offline upgrade failure in systems with mixed types of nodes** In HCP G10 systems that also use CR210H nodes, the offline upgrade procedure fails. **Fix:** The offline upgrade procedure no longer fails in systems with mixed node types.
HCP-29295	**Attempt to create custom annotation on object in default namespace causes node restart** Issuing an HTTP request to create a custom annotation on an object in the default namespace should result in a 400 bad request error, but instead the request causes the HCP node to restart. **Fix:** Issuing an HTTP request to create a custom annotation on an object in the default namespace now returns a 400 bad request error.
HCP-29304	**Node restart with invalid front-end message** When the HCP system tries logging exceptions that could not be retried, the node restarts with an invalid front-end message type 209. **Fix:** The node no longer restarts with an invalid front-end message type.
HCP-29502	**Slowed HCP system performance due to HTTP writes to a namespace that exceeded hard quota** If an HTTP write request is sent to a namespace with an exceeded hard quota, the request is declined and a warning message is issued to the internal logs. Repeated write requests continue to issue warning messages, which can bloat internal logs and slow system performance. **Fix:** Repeated HTTP write requests to an HCP namespace with an exceeded hard quota now generate only one internal log warning message per node every fifteen minutes.
HCP-29568	**HTTP GET request during node unavailability event returns incorrect status code** If an HCP node is unavailable, and you issue an HTTP GET request through the management API, the HCP system returns an HTTP 403 Forbidden error message instead of an HTTP 503 Service Unavailable error message. **Fix:** If an HTTP GET request is issued to an unavailable HCP node, the system now returns the correct response code.

Ref. number	Description
HCP-29626	**Migrating object with newer annotations on replica system causes node restart** On an HCP system with replication enabled, if you perform a migration procedure on a node in the primary system and the migration fails, HCP reattempts the migration using objects copied from the replica system. If an object on the replica system has newer annotations than its original object on the primary system, the copied objects causes the primary system HCP node to restart. **Fix:** If an object on a replica HCP system is copied for migration and has newer annotations than the original object, the primary system HCP node no longer restarts.
HCP-29685	**Nodes restart occasionally when writing objects to HCP S Series Nodes** When there is a service plan with more than one S Series storage pool, ingesting to HCP S Series Nodes can cause HCP nodes to restart. **Fix:** In service plans with more than one S Series storage pool, ingesting to HCP S Series Nodes no longer causes HCP nodes to restart.
HCP-29775	**HCP node restart due to protection service repairing too many deleted objects** On an HCP system with active/passive replication enabled, if an object is deleted from the primary HCP system, the protection service repairs the object using the replica system object copy and leaves the replica system object in an open state. If too many objects are left in an open state, the replica system HCP node restarts. **Fix:** HCP now closes the connection to the replica system object copy once the protection service has finished repairing the primary system object.
HCP-29785	**HCP system busy error message when downloading logs** When downloading internal logs from the **System Management Console ►Monitoring►Internal Logs** page, an error message is occasionally displayed which states that the system may be busy. **Fix:** The log download interval has been increased so logs have more time to begin downloading.
HCP-29888	**JVM restart due to HS3 PUT request on HCP system with direct write to HCP S Series Nodes** If an HCP system is configured to write directly to HCP S Series Nodes, issuing a PUT request using the HS3 protocol and AWS v4 signature causes JVM to restart. **Fix:** On an HCP system that is configured to write directly to HCP S Series Nodes, issuing a PUT request using the HS3 protocol and AWS v4 signature no longer causes JVM to restart.

Ref. number	Description
HCP-29900	**Upgrade from HCP version 7.1.0 or later to 7.3.0 reverts advanced downstream DNS to basic mode** On upgrade from HCP release version 7.1.0 or later to 7.3.0, if advanced downstream DNS mode is enabled, HCP reverts back to basic. **Fix:** When upgrading from HCP release version 7.1.0 or later to 7.3.1, HCP no longer reverts advanced downstream DNS to basic mode.
HCP-29910 and HCP-29050	**Node restart due to direct write and zero day tier to HCP S Series Node** An HCP system that's configured for direct write or zero day tier to HCP S Series Nodes occasionally causes an OutOfMemory exception and causes a node restart. **Fix:** HCP no longer encounters a OutOfMemory exception when using directwrite or zero day tiering to HCP S Series Nodes.
HCP-29914	**Node restart due to exception when rehydrating from cloudstorage endpoint** When rehydrating objects from a cloud storage endpoint through CIFS or NFS, certain exceptions cause nodes to restart. **Fix:** HCP now prevent the exceptions that cause nodes to restart.
HCP-29974	**Bad request error response when using .net S3 SDK** When using the .net AWS SDK to perform a create bucket operation, a bad request error response appears. **Fix:** The aws-chunked content encoding is now supported so the bad request error response no longer appears.
HCP-29997	**Apply patches to fix Struts 2 vulnerability** Struts2 was discovered to have a vulnerability. **Fix:** A patch has been applied to fix the following security vulnerability: • CVE-2017-5638 For more information about these vulnerabilities and the patches that have been applied to address these vulnerabilities, refer to the bash security advisory document found here: https://cwiki.apache.org/confluence/display/WW/S2-045
HCP-30028	**Failed node recovery procedure due to database location not found** On an HCP SAIN system, the node recovery procedure fails with the error message that database pgdata location could not be found. **Fix:** The node recovery procedure no longer fails with the error message that database pgdata location could not be found on HCP SAIN systems.

Issues resolved in release 7.3

The table below lists issues that have been resolved in release 7.3 of HCP. The issues are listed in order by reference number.

Ref. number	Description
HCP-28390	**Node unavailability due to attempted deletion of incomplete object metadata** When the garbage collection service tries to delete incomplete object metadata left by a failed CIFS or NFS request, the node on which the service is running becomes unavailable. **Fix:** The garbage collection service now deletes incomplete object metadata without causing the node to become unavailable.
HCP-29091	**Unavailability of added or recovered node due to reboot during online upgrade** During an online upgrade, if you reboot a node that was previously added to the system or recovered, the node becomes unavailable. **Fix:** You can now reboot an added or recovered HCP node during an online upgrade without the node becoming unavailable.
HCP-29210	**Node restart due to unavailability event during custom metadata tiering** Occasionally, if a node becomes unavailable while moving custom metadata to a storage tier, the node restarts. **Fix:** A node unavailability event during custom metadata tiering no longer causes the node to restart.
HCP-29212	**Node restart due to repeat of a canceled HTTP PUT object request with conditional headers** Canceling an HTTP PUT object request that includes conditional headers and then issuing the same PUT request causes the node processing the request to restart. **Fix:** Reissuing a canceled PUT object request with conditional headers no longer causes the node to restart.
HCP-29225	**Node restart due to enabling CIFS after disabling user-defined network** The data access network for a tenant can be a user-defined network. If you disable that network and then enable CIFS for a namespace owned by the tenant, all nodes in the HCP system repeatedly restart. **Fix:** Enabling CIFS in this situation no longer causes nodes to restart.

Ref. number	Description
HCP-29237	**Node restart due to replication of multiple changes to single object in short period of time** When replicating multiple changes to a single object in a short period of time, the node handling the replication restarts. **Fix:** Replication of multiple changes to a single object in a short period of time no longer causes the node to restart.
HCP-29258	**Truncated data returned in response to HTTP GET request for multiple versions** When a client uses the HTTP protocol to request multiple versions of an object, HCP returns truncated object data. **Fix:** GET requests for multiple versions no longer result in truncated object data.
HCP-29274	**Node restart due to too many objects in search results** An HCP node restarts when the total number of objects returned by a search exceeds the maximum supported result count. **Fix:** The maximum supported result count has been increased.
HCP-29330	**Node restart due to tiered object with unavailable metadata** If the content verification service runs on an object tiered to cloud storage and the object metadata is temporarily unavailable, the node running content verification restarts. **Fix:** Running content verification on tiered object with unavailable metadata no longer causes the node to restart.
HCP-29380	**Logs missing from download** When downloading the internal logs, if you set the end date to the current day, some logs are not included in the download. **Fix:** The HCP system now includes all requested logs in a log download.

Issues resolved in release 7.2.3

The table below lists issues that have been resolved in release 7.2.3 of HCP. The issues are listed in order by reference number.

Ref. number	Description
HCP-28371	**HCP software install fails on system with ten or more nodes** During an HCP software installation, if the HCP system has ten or more nodes, the HCP installation fails at the multicast precheck. **Fix:** The HCP software installation multicast precheck no longer fails due to the system having ten or more nodes.
HCP-28520	**Aborting compression service occasionally causes object to lose metadata and be ignored by scavenging service** If the compression service is aborted at the same time that an HCP node becomes unavailable, an object can be separated from its metadata. The affected object is ignored by the scavenging service. **Fix:** Object data that is separated from its metadata is now quarantined by the scavenging service.
HCP-28522	**HSwift and HS3 object metadata headers ignore lower case headers** HSwift and HS3 object metadata request headers are supposed to be case insensitive, but object metadata request headers that use the lower case are ignored. **Fix:** HSwift and HS3 object metadata request headers are now case insensitive.
HCP-28596	**SNMP MIB field reports incorrect server name** In the SNMP MIB, the nodeHardwareType field reports the HCP G10 server as Quanta D51B-2U when the value should be HCP G10. **Fix:** The nodeHardwareType field now correctly reports the HCP G10 server as HCP G10.
HCP-28808	**Tiering to NFS when primary storage is above 95% capacity causes HCP system restart** If primary storage is above 95% capacity and an object is tiered to NFS, the HCP system restarts. **Fix:** Once primary storage hits 95% capacity, it stops tiering to NFS without restarting the HCP system.
HCP-28813	**Replication collisions with repeated store and delete operations** In a replicated namespace with versioning disabled, if you store and then delete the same object multiple times a row within a short period of time, a replication collision can occur. **Fix:** Replication collisions no longer occur in this situation.

Ref. number	Description
HCP-28814	**Replication stalled on object with missing data** In a three site replication topology with two HCP systems that are set to metadata only mode, if the primary HCP system attempts to replicate an object that is missing data, the replication link stalls. **Fix:** The condition that causes the replication link to stall no longer occurs.
HCP-28823	**Scavenged object has custom metadata deleted if new custom metadata exists for object** If the scavenging service finds an object with custom metadata, the scavenging service might delete the custom metadata if another copy of that object with different custom metadata is found. Another object copy can exist if the HCP system is part of an active/active replication link. **Fix:** The scavenged object custom metadata is now placed in lost and found instead of being deleted.
HCP-28825	**Under specific conditions, custom metadata loss can occur during active/active replication** Under specific conditions, in an HCP system that is part of an active/active replication link, an object can be separated from its custom metadata. **Fix:** An object can no longer lose its custom metadata during active/active replication.
HCP-28850	**Network communication vulnerable to abuse due to TCP/IP issue** Due to a TCP/IP issue, network communication between two HCP systems may be compromised. **Fix:** This issue has been corrected within HCP.
HCP-28858	**Running protection service on object with missing data causes node restart** If the protection service runs on an object that is missing data, the HCP node restarts. **Fix:** If the protection service runs on an object that is missing data, HCP handles the error and no longer causes the node to restart.

Issues resolved in release 7.2.2

The table below lists issues that have been resolved in release 7.2.2 of HCP. The issues are listed in order by reference number.

Ref. number	Description
HCP-28261 and HCP-28274	**Object written directly to an HCP S Series Node has size set to zero** Rarely, the content verification service changes the size of an object written directly to an HCP S Series Node with compression or encryption enabled to zero. **Fix:** The content verification service no longer sets object size to zero unless the object size is actually zero. The content verification service also corrects any existing objects impacted by this issue.
HCP-28285	**Node restart due to multiple custom metadata PUT requests to the same object during HCP upgrade** Occasionally, during an upgrade from HCP 7.0.*x* or 7.1.*x* to 7.2.*x*, issuing multiple custom metadata PUT requests for the same object that has already been tiered to economy or external storage can cause an HCP node to restart. **Fix:** Multiple custom metadata PUT requests for the same tiered object during an upgrade no longer causes an HCP node to restart.
HCP-28310	**Node restart due to HEAD request error** Occasionally, when an object HEAD request fails, HCP nodes restart instead of handling the error. **Fix:** HCP now handles the HEAD request error without causing nodes to restart.

Issues resolved in release 7.2.1

The table below lists issues that have been resolved in release 7.2.1 of HCP. The issues are listed in order by reference number.

Ref. number	Description
HCP-26132	**Occasionally node operating system reboots without notification** Occasionally the operating system on a node running HCP release 7.0 or later reboots without logging that it has rebooted. **Fix:** The operating system on a node running HCP release 7.0 or later now logs each reboot.
HCP-27150	**Curl request does not accept a period (.) at the end of domain name** On an HCP system with replication enabled, sending a curl request with a period (.) at the end of the domain name and having the request redirected to a replica system returns a 403 error. **Fix:** A curl request can now be sent to an HCP system with a period (.) at the end of the domain name.

Ref. number	Description
HCP-27185	**Occasionally HCP system slows or becomes unresponsive** Some HCP system configuration changes send synchronous LDAP calls to Active Directory servers. If the AD servers are slow to respond, any further changes can be blocked for up to 30 minutes and can make the HCP system unresponsive. **Fix:** HCP system configuration changes no longer impair system performance.
HCP-27245	**Communication error makes all HCP system nodes unavailable** If a node in the HCP system experiences a communication error, all nodes in the system become unavailable instead of just the affected node. **Fix:** A communication error on one node no longer impacts the other nodes in an HCP system.
HCP-27283	**Tenants added to active/active replication links continually autopause under specific conditions** If a tenant contains a namespace that has ingested objects and has cloud optimization enabled, the tenant will continually autopause if it is added to an active/active replication link. **Fix:** Tenants added to active/active replication links no longer autopause.
HCP-27309	**Mounting external NFS volume through Windows NFS server causes tiering to fail** If file names contain special characters that Windows NFS servers do not support, an external NFS volume mounts successfully but fails to tier. **Fix:** HCP now checks if character mapping is enabled and tests if it can tier through NFS. If HCP cannot tier, the NFS volume mount fails.
HCP-27310	**Requesting a replicated object on an offline node causes node restart** If a replicated object that is stored on an offline node in a replica system is requested, the node that holds a copy of that object on the primary system restarts. **Fix:** Requesting a replicated object from an offline node no longer causes nodes to restart.
HCP-27313	**IP Allow/Deny filter causes HS3 requests to fail** The IP Allow/Deny filter for HS3 requests using path style URLs uses the incorrect ruleset, causing requests to fail. **Fix:** The HS3 IP Allow/Deny filter now uses the appropriate path style URL ruleset.
HCP-27314	**Requesting missing metadata causes connection error** If an object with multiple annotations is requested and one of the object annotations fails to open, an Exceeded Max Connections: 8192 error occurs. **Fix:** Objects with missing annotations no longer cause the Exceeded Max Connections: 8192 error when they are requested.

Ref. number	Description
HCP-27320	**Anonymous user cannot connect to HCP namespace through HCP Data Migrator** A valid anonymous request is passed through HCP-DM with an authentication header, but the cookie is denied. The system returns a 403 Forbidden error message. **Fix:** HCP now accepts valid anonymous requests with the authentication header, and a cookie is sent through HCP-DM.
HCP-27322	**Duplicate version IDs of same object cause node to restart** When versioning is enabled and two versions of the same object are written to an HCP system node at the same time, duplicate version IDs are generated which causes the node to restart. **Fix:** The HCP system now handles the duplicate version ID error without causing the node to restart.
HCP-27328	**Mishandling of Retries Exceeded exception causes JVM restart** The Retries Exceeded exception is not handled properly by the JVM, causing the JVM to restart. **Fix:** HCP now handles the Retries Exceeded exception without restarting the JVM.
HCP-27335	**Failback Zero copy failover (ZCF) may cause missing objects** When a node in a SAIN system fails, the logical volumes managed by that node fail over to the ZCF peer node for the failed node. When the failed node resumes operation, its logical volumes fail back. Occasionally, after the fail back, the peer node and repaired node both assume responsibility over the logical volumes. This may cause objects to lose their data. **Fix:** A repaired HCP node now assumes sole responsibility over its logical volumes.
HCP-27558 and HCP-28012	**Node hangs during failback after zero copy failover (ZCF)** When a node in a SAIN system fails, the logical volumes managed by that node fail over to the ZCF peer node for the node that failed. When the failed node resumes operation, its logical volumes fail back. Occasionally, the repaired node hangs when its logical volumes fail back. **Fix:** A repaired HCP node no longer hangs when its logical volumes fail back.
HCP-27337	**OutOfMemoryError: GC overhead limit exceeded error causes JVM to crash** On a HCP system with a lot of objects and metrics, an OutOfMemoryError: GC overhead limit error can occur and cause the JVM to crash. **Fix:** The chance of an OutOfMemoryError: GC overhead limit error and the subsequent JVM crash has been reduced on systems with a lot of objects and metrics.

Ref. number	Description
HCP-27340	**Duplicate Elimination service causes missing annotations** If the Duplicate Elimination service runs on duplicate objects with different annotations, the resulting object loses some of the annotations. **Fix:** The Duplicate Elimination service no longer loses annotations when removing duplicate objects from the HCP system.
HCP-27342	**Garbage Collection service does not remove certain non-replicating objects** Non-replicating objects that remain open for one week or longer are not removed from the HCP system by the Garbage Collection service. **Fix:** The Garbage Collection service now removes non-replicating objects that have been open for one week or longer from the HCP system.
HCP-27346	**Occasionally Active Directory user account is denied permissions and removed from AD group** An Active Directory user account can inherit permissions from the AD group that the user account belongs to. If an error denies the AD user account one of the permissions it is supposed to inherit from its groups, the user account is removed from all groups and does not inherit permissions. **Fix:** The AD user account is no longer removed from its AD groups if an error denies the AD user account one of its permissions. Instead, the AD user account inherits the permissions that did not error and remains a member of its AD groups.
HCP-27361	**After upgrade, replicating certain namespaces can cause nodes to restart** After an upgrade to release 7.2 of HCP, replicating a namespace that had versioning enabled and later disabled can cause nodes to restart. **Fix:** After upgrading, replicating a namespace that had versioning enabled and disabled does not cause nodes to restart.
HCP-27375	**Occasionally SPN status information does not update** If HCP checks the SPN status information of the wrong node, the SPN status information no longer updates. **Fix:** If HCP checks the SPN status information of the wrong node, HCP handles the error and continues to update the SPN status information.
HCP-27378 and HCP-27379 and HCP-27380	**Occasionally object reads and writes through CIFS cause HCP nodes to restart** A large amount of object reads and writes through CIFS can cause HCP nodes to restart. **Fix:** HCP nodes no longer restart from performing a large amount of CIFS object reads and writes.

Ref. number	Description
HCP-27383	**Metadata query engine stuck in Balancing state after indexing enabled** Occasionally the MQE indexer does not start when it is enabled. The indexer remains in the Balancing state. **Fix:** The MQE indexer now starts when it is enabled. The indexer transitions to the Running state.
HCP-27385	**Aborted CIFS or NFS large object ingest causes replication collision** If a large object ingest through CIFS or NFS into an HCP system with replication enabled is aborted and later retried, a replication collision occurs. **Fix:** CIFS or NFS large object ingest aborts no longer cause replication collisions.
HCP-27386	**Duplicate Elimination and Compression services skip objects on active/active replication link** Objects replicated over an active/active replication link are skipped by the Duplicate Elimination and Compression services. **Fix:** Objects replicated over an active/active replication link are no longer skipped by the Duplicate Elimination and Compression services.
HCP-27387	**Objects on the replica system in an active/active replication link fail to delete** If an object is ingested and deleted in a short period of time on an HCP system in an active/active replication link, the replicated version of the object fails to delete on the replica system. **Fix:** Objects on the replica system in an active/active replication link now delete when the counterpart object is deleted.
HCP-27408	**Occasionally HCP node hangs during restart** Occasionally an HCP node may hang during the restart process and require a manual reboot. **Fix:** HCP nodes no longer hang during the restart process.
HCP-27410	**Zero copy failover error causes node to hang** If a node goes offline and comes back online before its ZCF peer node claims ownership of the offline node volumes, the peer node hangs and needs to be rebooted. **Fix:** ZCF peer nodes no longer hang if they do not claim ownership of the offline node volumes.

Ref. number	Description
HCP-27413	**Storage trimming causes HCP system nodes to become unavailable** Occasionally, storage trimming causes a node on HCP systems with solid state disks or certain storage array models to become unavailable. The affected node remains unavailable until storage trimming finishes running. Storage trimming runs every hour. **Fix:** Running storage trimming on a node no longer causes the node to become unavailable.
HCP-27420	**Occasionally aborted PUT requests create unreplicatable open objects** Occasionally on a system with replication enabled, the Garbage Collection service does not remove left over data from aborted PUT requests. The left over data becomes an unreplicatable object. **Fix:** The Garbage Collection service now removes data left over from aborted PUT requests.
HCP-27421	**Occasionally an upgrade from 6.x to 7.x causes Allow IP list to be incorrect** After upgrading an HCP system from version 6.x to 7.x, IPv6 wildcards are not always correctly added to the IP Allow Lists. Incorrect wildcards can unintentionally restrict IP address access or give access permission to an unwanted IP address. **Fix:** Upgrading an HCP system no longer affects the IPv6 wildcards added to the IP Allow Lists.
HCP-27445	**snmpd crash causes failed SNMP requests** Occasionally snmpd crashes which causes SNMP requests to the HCP system to fail. **Fix:** The chance of an snmpd crash occuring has been reduced.
HCP-27450	**Sometimes replicated objects directly written to HCP S Series Nodes contain no data** When an object is replicated over a compressed replication link to a replica system where objects in the target namespace are supposed to be written directly to an S Series Node, the object is missing all of its data on the S Series Node. **Fix:** Objects replicated over a compressed replication link to a replica HCP system and written directly to an HCP S Series Node no longer lose data.
HCP-27466	**Replicated namespace inherits configuration settings from deleted namespace** When a namespace created through replication is deleted and a new namespace created through replication with the same name replaces it, the new namespace inherits its configuration settings from the deleted namespace instead of the primary HCP system namespace it was replicated from. **Fix:** A newly replicated namespace no longer inherits configuration settings from a deleted namespace with the same name.

Ref. number	Description
HCP-27499	**Disposition service fails to delete objects subject to replication collision** The Disposition service does not delete objects that were subject to a replication collision on the replica system in an active/passive replication link. **Fix:** The Disposition service now deletes objects that were subject to a replication collision on the replica system in an active/passive replication link.
HCP-27509	**Documentation link on Search Console returns 404 Not Found error** The documentation link on the Namespace Search Console returns a 404 Not Found error instead of linking to the *Searching Namespaces* manual. **Fix:** The documentation link on the Search Console now links to the *Searching Namespaces* manual.
HCP-27570	**Replication link failover status incorrect after HCP upgrade** If a failover occurs on an HCP system with release version 6.x of HCP and replication enabled then the replica system is upgraded to release version 7.x, there is no failover notification. **Fix:** If a failover occurs on the primary HCP system with replication enabled and the replica system is upgraded, the failover status is now preserved after the upgrade.
HCP-27626	**Storing too many internal logs causes node to restart** If an HCP node stores too many internal logs, the node restarts. **Fix:** HCP now deletes redundant messages to reduce the size of internal logs.
HCP-27676	**Operation-based MQE queries fail on certain tenants** If a tenant is not on the system domain, operation-based MQE queries that target the tenant fail. **Fix:** Setting a tenant to domains other than the system domain no longer causes operation-based MQE queries to fail.
HCP-27686	**Occasionally the HCP system becomes unavailable when an HCP-VM is restarted** Occasionally when a fast restart is performed on an HCP-VM, the restart causes the HCP system to enter an unavailable state. The system needs to be restarted. **Fix:** A fast restart of an HCP-VM no longer causes an HCP system to enter an unavailable state.
HCP-27727	**Replication link status incorrect when experiencing high error rate** If a replication link experiences a high error rate, the replication link status shows the broken link message instead of the high error rate message. **Fix:** The replication link status now accurately reflects the status of the replication link.

Ref. number	Description
HCP-27761	**Single sign on request fails on HCP system in an active/active replication link** A single sign on request fails if it is sent to an HCP system in an active/active replication link and redirected to the second system in the replication link. **Fix:** Single sign on requests no longer fail if they are sent to HCP systems in an active/active replication link.
HCP-27771	**Content-Type header returned in response to HEAD requests breaks HS3 third party applications** In release 7.2 of HCP, a Content-Type header response to a HEAD request that contains a charset parameter breaks certain third party applications. **Fix:** HS3 Content-Type header response charset parameters no longer break client applications.
HCP-27811	**Active Directory SSO and SPNEGO authentication fails for users who are members of many AD groups** When an AD user account is a member of a large number of AD groups, the user account fails SSO and SPNEGO authentication with a 413 Header Full error. **Fix:** The number of groups an AD user account can be a member of has been significantly increased.
HCP-27864	**Occasionally communication error causes node to hang** Occasionally if a node tries to communicate with the HCP system and another node is offline, the node that is sending data hangs and requires a reboot. **Fix:** A communication error no longer occurs if one node in the HCP system is offline.
HCP-27898	**Occasionally HCP node volumes fail to unmount during shutdown** HCP node volumes may fail to unmount during the shutdown procedure. The node needs to be restarted. **Fix:** An error message now appears if a volume failed to unmount. A node reboot is required.
HCP-27952	**Object replicated and tiered to HCP S Series Node contains incorrect metadata** Objects replicated through an active/passive replication link and then tiered through direct write from the replica system to an encrypted HCP S Series Node occasionally contain incorrect object size metadata. **Fix:** Objects no longer have incorrect metadata when tiered to an encrypted HCP S Series Nodes from a replica system.

Ref. number	Description
HCP-27962	**The Content Property extraction tool fails to extract content properties** The Content Property extraction tool fails to extract content properties for elements in XML that contain CDATA tags. **Fix:** The Content Property extraction tool has been fixed to extract character data from objects attached to custom metadata.
HCP-27968	**Certain AD user account requests fail** Occasionally AD user account roles are requested from an HCP node that has not loaded the list of roles. This returns a HTTP 503 error. **Fix:** HCP nodes now load the AD user account roles before handling the AD user account roles request.
HCP-28018	**Occasionally data cannot be accessed during a failover from the replica HCP system** Occasionally during a failover, data cannot be accessed from the replica HCP system in an active/active replication link. **Fix:** All data can now be accessed from a replica HCP system during a failover.
HCP-28020	**Occasionally MQE indexer hangs** MQE indexer hangs if multiple versions of an object have the same change time within seconds of each other. **Fix:** MQE indexer no longer hangs due to object change time.
HCP-28095	**Internal time gets set incorrectly** The HCP system assumes the internal time set is UTC instead of local, causing the time to be incorrect. **Fix:** HCP now assumes internal time is local.
HCP-28111	**Nodes restart if replication link stalls** If a replication link stalls on an object, there is a chance it may restart nodes. **Fix:** HCP handles a replication link stalling on an object without causing nodes to restart.

Issues resolved in release 7.2

The table below lists issues that have been resolved in release 7.2 of HCP. The issues are listed in order by reference number.

Ref. number	Description
HCP-24392 and HCP-26838	**Metadata Query Engine becomes unavailable after a node event** The Metadata Query Engine can become unavailable after a node event such as: Zero Copy Failover, node recovery, node addition, or node removal. **Fix:** If the Metadata Query Engine becomes unavailable, it automatically restarts.
HCP-25618	**When updating an object ACL in HS3, the object ACL is updated for all versions of the object** In S3, updating an object ACL only changes the ACL on the current version of the object. In HS3, the ACL update changes object permissions on all existing versions of the object. **Fix:** HS3 only changes the ACLs on the current version of the object when the object ACLs are updated.
HCP-25757	**Cannot rename Replication Network** The **Replication Network** name changes are not updated in the system configuration. **Fix:** The HCP **Replication Network** name is now updated correctly when it is changed through the System Management Console.
HCP-26008	**Offline upgrade takes longer than expected if objects are being compressed** The offline upgrade takes longer to complete if the object compression service is running when the upgrade starts. **Fix:** All services are automatically turned off when an upgrade begins.
HCP-26030	**Performing a HEAD command on an object that is exactly 4GB large returns two Content-length headers** If you perform a **HEAD** command on an object that is exactly 4GB in size, you receive two Content-Length headers in the response instead of one. The second header states that the object Content-Length is zero. **Fix:** HCP returns only one, accurate Content-Length header for objects.
HCP-26040	**Performing an HTTP GET command on an empty directory takes longer than expected** An HCP system takes a long time to respond to an HTTP **GET** command on an empty directory. **Fix:** HCP responds faster to an HTTP **GET** command on an empty directory.

Ref. number	Description
HCP-26053	**Storage retirement cannot retire components when objects tiered to NFS storage are on hold** If an object is tiering to an NFS storage component and set on hold, the garbage collection service cannot remove the object. Since the component cannot be emptied, it cannot be retired. **Fix:** The garbage collection service collects objects that are on hold, which lets storage retirement delete all objects from the retired component and remove it.
HCP-26060	**Data truncated on decompression with HS3 and HSwift** HCP truncates data from objects it decompresses that are gzipped and uploaded through HS3 or HSwift. **Fix:** HCP saves the entire, uncompressed object.
HCP-26138	**Cannot upload a storage license with no expiration in certain timezones** Valid storage licenses with no expiration cannot be uploaded into systems operating on timezones east of **Eastern Standard Time** (EST). **Fix:** Storage licenses with no expiration can be uploaded to an HCP system operating on any timezone.
HCP-26156	**Occasionally, domain service record sorting causes nodes to reboot** When HCP is under extremely high stress from sorting domain service records, the nodes occasionally reboot. **Fix:** HCP nodes do not reboot from sorting domain service records.
HCP-26262	**Broken replication link causes nodes to become unavailable** Occasionally, a broken replication link causes HCP nodes to become unavailable until the replication link is fixed or removed and the nodes are rebooted. **Fix:** A broken replication link no longer causes HCP system nodes to become unavailable.
HCP-26292	**Node restarts if an HTTP GET request for object custom metadata is aborted** If you perform an HTTP **GET** request for object custom metadata and the request is aborted while the node is still responding, the node restarts. **Fix:** An HTTP **GET** request for object custom metadata can be aborted without having the responding node restart.
HCP-26296 and HCP-26734	**Nodes may restart when HCP encounters an IO connection error** An HCP node may restart if it encounters an IO connection error with another node. **Fix:** HCP nodes can handle IO connection errors without restarting.

Ref. number	Description
HCP-26419	**Data from aborted object writes cause replication links to stall** HCP tries to hash data from an aborted object write, which makes the replication service attempt to replicate aborted write data. The replication link stalls because it cannot replicate the data. **Fix:** HCP does not hash aborted write data.
HCP-26444	**Failed object ingests through NFS or CIFS remain stuck in the open state** Objects with no data that failed to be ingested through NFS or CIFS remain in the open state instead of being removed by the garbage collection service. **Fix:** The garbage collection service removes objects that failed to be ingested through CIFS or NFS and remain in the open state.
HCP-26461	**Deleting objects with custom metadata does not remove them from all systems in an active/active replication link** An object with custom metadata that is deleted from a system in an active/active replication link does not get deleted on other systems in the replication link. **Fix:** An object with custom metadata that is deleted from an HCP system in an active/active replication link is deleted from all systems in the replication link.
HCP-26601	**Occasionally, objects are skipped during indexing** Occasionally, if an object is modified, some other objects are skipped when indexing resumes. **Fix:** Indexing does not skip objects when it resumes.
HCP-26757	**HCP SPN check flags existing SPNs as missing on replicated systems** When HCP checks SPNs on systems in a replication link that have different domain names, it marks existing SPNs on the replicated systems as missing. **Fix:** The SPN check no longer flags existing SPNs as missing.
HCP-26846	**Auto-exchange domains and certificates feature causes HCP quality issues** The auto-exchange domains and certificates feature does not perform its intended function and interferes with other HCP processes. **Fix:** The auto-exchange domains and certificates feature has been removed from HCP.
HCP-26890	**Setting the replication schedule to OFF on a replica system in an active/active replication link restarts nodes in the system** If the replication schedule is set to **OFF** on the replica system in an active/active replication link, nodes in the replica system restart. **Fix:** Setting the replication schedule to **OFF** on the replica system in an active/active replication link does not restart the system nodes.

Ref. number	Description
HCP-26971	**Reading an object that references a non existent retention class causes a node reboot** Occasionally, retention classes do not get replicated properly which causes objects to reference nonexistent retention classes. If an HCP node tries to read an object that references a nonexistent retention class, the HCP node reboots. **Fix:** An HCP node now handles the error when reading an object that references an nonexistent retention class without rebooting.
HCP-26982	**Adding nodes to an HCP system resets VLAN and VNeM configuration settings** Adding nodes to an HCP system resets saved VLAN and VNeM configuration settings to their default values. **Fix:** Nodes can be added to an HCP system without resetting the VLAN and VNeM configuration settings.
HCP-26991	**Mapping class content to namespaces over one hundred times on a system in a replication link causes a system reboot** If an HCP system has over one hundred class contents mapped to its namespaces and is part of a replication link, other systems in the replication link reboot. **Fix:** HCP systems in a replication link no longer reboot from having class contents mapped to namespaces.
HCP-27008	**Cannot add custom metadata to an object that is on primary and external storage** If an object is stored on primary storage and is tiered to an external storage component, you cannot add custom metadata to the object. **Fix:** You can add custom metadata to objects that exist on primary and external storage.
HCP-27103	**Data Protection Level (DPL) values are not saved on replica systems in an active/passive replication link when upgrading from release 6.1.2 of HCP to 7.*x* or 7.1.*x*** The namespace DPL values set for objects on a replica system in an active/passive replication link are reset to one when the systems are upgraded from release 6.1.2 of HCP to 7*x* or 7.1.*x*. **Fix:** The replica system in an active/passive replication link retains its namespace DPL settings when the systems are upgraded from release 6.1.2 of HCP to 7.*x* or 7.1.*x*.
HCP-27236	**Replication links with compression enabled cause node unavailability** Replication links with compression enabled occasionally cause node unavailability. **Fix:** Replication links with compression enabled no longer cause node unavailability.

Issues resolved in release 7.1.2

The table below lists issues that have been resolved in release 7.1.2 of HCP. The issues are listed in order by reference number.

Ref. number	Description
HCP-26120	**If a replication link fails over, tenants on a VLAN network cannot be accessed by the remote system** After a replication link fails over, accessing a tenant that's on a VLAN network through the remote system returns a 400, not found, error message. **Fix:** Remote systems can access tenants on a VLAN network after a replication failover.
HCP-26121 and HCP-26108	**Open state objects closed by the garbage collection service do not get replicated** Objects ingested through NFS and CIFS are lazy closed once they have been idle for 15 minutes. If there is a system unavailability event while an object is in an open state, the garbage collection service closes the idle, open state object. An object closed by garbage collection may not get replicated. Non-replicated objects might result in data loss if the primary system is retired. **Fix:** Replication now replicates all objects closed by garbage collection. The **Replication Verification Service** (**RVS**) has been added to the protection service. RVS identifies and replicates objects closed by garbage collection. RVS is configured to run once after an upgrade to 7.1.2, but can also be run at any time after upgrading or be configured to run continuously with the protection service. RVS can only identify non-replicated objects when all replication links are healthy and replicating. Please ensure all links are healthy and running protection service prior to upgrading to 7.1.2. If replication links go offline during the RVS run, please contact support to re-enable the service for the next protection service run.
HCP-26122	**Purging or pruning irreparable objects does not remove them from the irreparable object list** Purging or pruning an irreparable object does not remove it from the irreparable object list. **Fix:** Purging or pruning an object now removes it from the irreparable object list and clears any warnings associated with that object.

(Continued)

Ref. number	Description
HCP-26131	**Occasionally, purged or deleted objects can reappear in namespaces on an active/active replication link** Under certain circumstances, an object ingested on a system in an active/active replication link can be purged and marked for garbage collection and still remain on both systems. If this occurs, a purge record of the object exists, and it gets unmarked for garbage collection. **Fix:** Objects ingested on a system in an active/active replication link that are then purged and marked for garbage collection get removed from both HCP systems.
HCP-26145 and HCP-26160	**Changing tenant settings can inadvertently change the tenant service plan configuration** Changing tenant settings from the tenant's **Settings** tab on the **Tenant** page of the System Management Console may cause unintentional changes to that tenant's service plan configuration. These unintentional changes can set the tenant DPL to a different value, which can cause the system to be out of compliance. **Fix:** Making changes to a tenant's settings does not impact the tenant service plan configuration. Also, the system detects if a tenant service plan was inadvertently changed and restores it.
HCP-26163	**HCP is unable to join an Active Directory domain when LDAPS is enabled in the domain** HCP takes a long time, or fails, to join an AD domain with LDAPS enabled. **Fix:** LDAPS does not hinder HCP from joining an AD domain.
HCP-26164 and HCP-26165	**Entering certain special characters in the Shared Secret field on the RADIUS page under Security on the System Management Console causes subsequent system restarts to fail** The **Shared Secret** field on the **RADIUS** page under **Security** on the System Management Console cannot handle certain special characters. Entering one of these special characters and clicking on the **Add RADIUS Server** button causes the system restart to fail. **Fix:** Special characters entered in the **Shared Secret** field on the **RADIUS** page do not prevent the HCP system from restarting.
HCP-26216	**Adding nodes to a system reverts certain Downstream DNS configurations to their default values** Adding nodes to an HCP system resets the Downstream DNS Configuration by reverting the **Enable hidden master**, **Enable notify**, and **Downstream DNS servers** fields to their default values without notification. **Fix:** All Downstream DNS Configuration settings are preserved when nodes are added.

Issues resolved in release 7.1.1

The table below lists issues that have been resolved in release 7.1.1 of HCP. The issues are listed in order by reference number.

Ref. number	Description
HCP-22084	**On a namespace with the CIFS access protocol enabled, IP addresses listed on both the CIFS Allow list and Deny list are permitted access** On a namespace that has the CIFS access protocol enabled, IP addresses placed on both the CIFS Allow list and CIFS Deny list on the **Protocols ▸ CIFS** page in the namespace section of the Tenant Management Console are allowed to access the namespace. IP addresses placed on both lists should not be allowed to access the namespace. **Fix:** IP addresses that are on both the Allow and Deny Address Lists are denied access to the namespace.
HCP-24929	**HCP deletes Compute Rack Server SEL logs** HCP deletes Compute Rack Server SEL logs that can be used for troubleshooting HCP hardware issues. **Fix:** Unless they are on a Dell platform, SEL logs are saved.
HCP-25774	**Under specific conditions, an active/active replication link may stall** Replication over an active/active link may stall if all of the following conditions occur: • One of the systems on the active/active link has a namespace with versioning enabled and pruning set to 0 days • The default service plan on both systems is set to tier **Fix:** An active/active replication link no longer stalls under the specified conditions.
HCP-25798	**Occasionally, on the System Management Console, the replication Transfer Rate displays a negative transfer speed** The **Transfer Rate** on the **Status ▸ Overview** tab of the **Replication** page in the System Management Console occasionally displays a negative transfer speed. The replication speed cannot be lower than 0. **Fix: Transfer Rate** cannot show numbers less than 0.
HCP-25813 and HCP-23814	**HCP-VM doesn't log the media through which it is installed** The media through which an HCP-VM is install is not logged. This can lead to other problems later in the installation. **Fix:**The HCP-VM installation media type is logged.

(Continued)

Ref. number	Description
HCP-25854	**When first installed, an HCP VM's default Data Protection Level (DPL) and Metadata Protection Level (MPL) is two** The default DPL and MPL of a VM system should be one. When the HCP-VM is first installed, the DPL is two. **Fix:** When a VM is first installed, its default DPL and MPL is now one.
HCP-25856	**Purging objects in deleted directories can cause nodes to restart** Garbage collection can delete a directory before its deletion record has been removed from the transaction logs. If this occurs and an object inside the deleted directory is purged, the node may restart. **Fix:** Garbage collection does not delete directories until their deletion record is removed from the transaction log.
HCP-25869	**Occasionally, the "Stalled" status doesn't clear on working active/active and active/passive replication links** Occasionally, after an active/active or active/passive replication link stalls, the "Stalled" status does not clear even though replication has resumed. **Fix:** Active/active and active/passive links no longer display a "Stalled" status when replication resumes. If the issue recurs, clear the stall state with the clearStalledState admin command.
HCP-25872	**It's possible for the scavenging service to make extra copies, which become irreparable objects, of objects that have been in a replication collision** During a replication collision, the scavenging service occasionally takes objects that have collided and makes extra copies of them. The copied objects appear as irreparable objects. **Fix:** Scavenging no longer creates irreparable object copies of objects that have been in a replication collision.
HCP-25874 and HCP-25910	**When creating a Storage Report, setting the Start Date to a date before the start of data collection or setting a report period greater than one month makes the storage report fail without producing an error message** When setting the parameters for a Storage Report on the **Storage ▸ Report** page of the System Management Console, selecting a **Start Date** earlier than the date data collection started or setting a report period greater than one month causes an error that prevents the report from generating. No error message appears to notify you that the report has not been created. **Fix:** A blank report is generated if the **Start Date** is earlier than the date data collection started, and an error message is generated if the report period exceeds one month.

(Continued)

Ref. number	Description
HCP-25911	**Occasionally, volumes become unavailable when the storage manager exceeds its assigned virtual memory limit** If a system has a large number of volumes, the amount of virtual memory the storage manager needs might be greater than the permitted limit. If the storage manager requires more virtual memory than it's allotted, volumes become unavailable. **Fix:** The storage manager virtual memory limit has been removed.
HCP-25961	**A primary system in both an active/active and active/passive replication link sends duplicates of the Certificate Signed Record (CSR) to the outbound system, causing the outbound system to shutdown and not be able to restart** A primary system in both an active/active and active/passive replication link sends a replication certificate and multiple copies of the CSR to the outbound system. The outbound system restarts if it receives multiple copies of the CSR. **Fix:** CSRs are not sent to the outbound system during replication.
HCP-25965	**After enabling CIFS for a namespace, invalid IP addresses can be added to the CIFS Allow List which causes nodes to fail future reboots** On the **Protocols ▸ CIFS** page of the **Namespaces** section of the Tenant Management Console, enabling CIFS and then adding an invalid IP address to the Allow List causes nodes to fail their next reboot and remain inoperable. This does not issue an error message. **Fix:** Invalid IP addresses cannot be added to the CIFS Allow List.
HCP-25968	**Graphs on the Monitoring ▸ Resources page of the System Management Console sometimes do not provide statistics** Graphs on the **Monitoring ▸ Resources** page of the System Management Console sometimes show the "No statistics are available for the selected time period" error message instead of showing statistics. **Fix:** Graphs on the **Monitoring ▸ Resources** page of the System Management Console show their statistics.
HCP-25972	**SMTP port remains open on unavailable nodes causing any SMTP writes to hang** When a node is unavailable it is read-only. Occasionally, an SMTP port remains open on the namespace of an unavailable node. While the node is read only, email transactions written to the node time out. **Fix:** The SMTP port is closed when nodes are read-only.

(Continued)

Ref. number	Description
HCP-25991	**Fence check guard prints excess debug information** The fence check guard prints debug information on the volume manager logs every ninety seconds for each volume on an HCP system. On systems with a lot of volumes, the log file can grow too large. **Fix:** The debug information has been removed.
HCP-25995	**Replication between 7.0.1 and 7.1 systems fails** Replication between systems with 7.0.1 and 7.1 releases of HCP become blocked because 7.1 system messages are not downgraded when they are sent to the 7.0.1 system. **Fix:** Messages downgrade before being sent to systems with using older releases of HCP.
HCP-25999	**If a replication ping fails due to a connectivity issue, the replication connection doesn't close** If a replication ping fails due to a connectivity issue, the connection between the systems does not close. Over time, the open connections from failed pings can accrue which might cause a system restart. **Fix:** Replication connections close if the replication ping fails.
HCP-26016	**A new system domain cannot be named after a previously deleted system domain** Once a system domain has been created, named, and deleted, new system domains cannot use the former system domain's name even though it no longer exists. **Fix:** New system domains can use the names of previously deleted system domains.

Issues resolved in release 7.1

The table below lists issues that were resolved in release 7.1 of HCP. The issues are listed in order by reference number.

Ref. number	Description
HCP-20486	**Nodes restart if an object with ACLs on 100 or more versions of the object has its ACL updated** If an object with 100 or more versions of itself has its ACL updated, the ACL updates on all versions of the object which causes nodes to restart. **Fix:** Nodes do not restart if objects with one hundred or more versions has its ACL updated.

Ref. number	Description
HCP-21365	**Occasionally after chain upgrading HCP, the AD alert, "Active Directory communication disrupted on nodes" is raised although the error doesn't occur** Occasionally after performing chained upgrades from HCP 5.x to 6.x to 7.x, the AD alert, "Active Directory communication disrupted on nodes" remains in the system logs although there is no communication disruption. This occurs because Winbind sometimes doesn't start after the upgrade. **Fix:** If Winbind is unavailable, the Active Directory health checker tries to restart it. If the issue persists, health checker gives a more accurate error message, "Winbind unavailable".
HCP-21737	**Chargeback occasionally misses collection times and doesn't gather data for reports when Reporting Interval is set to Hour and system is under heavy load** If the **Reporting Interval** for a chargeback report is set to **Hour** and the HCP system is under a heavy load, some hourly reports show 0 for all field values or show incorrect field values. The data from hours is added to the data for the next hour, resulting in apparent spikes. This issue does not occur for chargeback reports with **Reporting Interval** set to **Day** or **Total**. **Fix:** Chargeback reports with their **Reporting Interval** set to **Hour** display the appropriate field values for all reporting intervals.
HCP-23801	**If an HCP system uses multiple array controllers for spindown storage, you cannot change the controller IP addresses** On HCP systems that are configured to use multiple storage array controllers for spin-down storage, attempts to use the System Management Console to change the controller IP addresses result in an error message. **Fix:** On an HCP system that uses multiple array controllers for spindown storage, the changes you make to controller IP addresses through the System Management Console take effect and do not produce an error message.

Ref. number	Description
HCP-24883	**After upgrading a system from HCP 4.1.1 to 5.x to 6.x to 7.0, attempting to start a data migration sometimes results in a 500 error** On an HCP SAIN system, after performing chained upgrades from HCP 4.1.1 to HCP 5.x to 6.x to 7.0, if you add storage to the SAN and then attempt to use the HCP System Management Console to migrate data from an existing storage array to the new storage array, the Console sometimes hangs and displays a processing message until it times out. If this happens, the data migration fails with a 500 error code (internal server error), and you have to restart the data migration procedure. **Fix:** An HCP SAIN system that has had chain upgrades from HCP 4.1.1 to 5.x to 6.x to 7.0 to 7.1 and has had storage added to its SAN can use the System Management Console to migrate data from existing storage arrays to new storage arrays without the data migration procedure timing out.
HCP-24885	**Add Node procedure occasionally changes saved downstream DNS configuration values** After an Add Node procedure, customized downstream DNS configuration values are occasionally reset to their default settings or deleted. **Fix:** DNS configuration values remain saved after the Add Node procedure.
HCP-24893 and HCP-24898	**DNS failover does not work for middle link in a replication chain** In a replication chain topology (A => B => C) with DNS failover enabled, if a system A fails over to system B, system B does not publish the domains used for the networks on system A. As a result, clients are not automatically redirected to system B. **Fix:** The 7.1 version of HCP doesn't experience longer shutdown times.
HCP-24993	**Occasionally, system memory is over allocated for compression** Running compression on an HCP node with more than six volumes causes the system to run out of memory. This causes the node to restart. **Fix:** The HCP system no longer devotes more memory than it has available to compression.
HCP-25142	**Shutdown time on 7.0 longer than on 6.1.2** On a 7.0 version of HCP the shutdown time of a system is longer than the shutdown time on 6.1.2 version of HCP. **Fix:** The 7.0 version of HCP no longer experiences longer shutdown times.

Ref. number	Description
HCP-25244	**Version 3.10 Linux kernel on Compute Rack servers may cause nodes to restart** The OS kernel in release 7.0 of HCP on Hitachi Compute Rack servers can occasionally cause system instability and node restarts. This is because the OS kernel has the NUMA (Non-Uniform Memory Access) balancing feature enabled by default and has reported issues associated with the feature that were resolved in later releases of the OS kernel. **Fix:** Version 7.1 of HCP has disabled NUMA balancing to resolved instability and node restart issues.
HCP-25259	**HCP cannot integrate with AD unless A-records exist in your DNS** HCP does not integrate with AD unless you have A-records in your DNS. This occurs even if the AD join would be otherwise successful. Without A-records, a preemptive forward zone lookup denies the connection. **Fix:** HCP can integrate with AD even no A-records exist but the connection is otherwise successful. If no A-records exist, a warning appears on your system log.
HCP-25267	**Apply patches to fix bash vulnerabilities** The bash that HCP currently uses was discovered to have some vulnerabilities. **Fix:** A bash patch has been applied to fix the following security vulnerability: • Shellshock (CVE-2014-7169) and (CVE-2014-6271) For information about these vulnerabilities and the patches that have been applied to address these vulnerabilities, refer to the bash security advisory document found here: https://fedoramagazine.org/flaw-discovered-in-the-bash-shell-update-your-fedora-systems/
HCP-25346	**Map balance may fail after increasing system region count on systems containing objects with ACLs** When increasing system region count, nodes have to rebalance. Occasionally the rebalance fails if the system is storing objects with ACLs. **Fix:** The node rebalance that occurs after the system region count is increased does not fail if objects with ACLs are stored on the system.

Ref. number	Description
HCP-25398	**Apply patches to fix known openssl vulnerabilities** The version of openssl that HCP currently uses for SSL communication was discovered to have some vulnerabilities. **Fix:** In the 2014 version of openssl that HCP uses for SSL communication, the following commonly known security vulnerabilities have been fixed. • POODLE: SSLv3 vulnerability (CVE-2014-2466) For information about these vulnerabilities and the openssl patches that have been applied to address these vulnerabilities, refer to the openssl security advisory document found here: https://access.redhat.com/articles/1232123.

Issues resolved in release 7.0.1

The table below lists issues that were resolved in release 7.0.1 of HCP. The issues are listed in order by reference number.

Ref. number	Description
HCP-22677 and HCP-24902	**Append functionality without the proper third party application can cause irreparable objects** Enabling the Append Support feature without the appropriate third party application causes errors such as irreparable objects. **Fix:**Changed the user interface so that there is no option to enable Append Support. Contact HCP customer support if you need to activate Append Support.
HCP-24378 and HCP-24963	**Certain settings in the Active Directory and HCP integration process cause objects to replicate indefinitely into namespace** Certain configuration setups for Active Directory to HCP integration cause Active Directory to not create an SPN for a namespace. This results in objects replicating indefinitely into the namespace. **Fix:** Fixed Active Directory configuration logic so that it creates SPNs for all namespaces.

Ref. number	Description
HCP-24967 and HCP-24941	**Incomplete upgrades to HCP 7.0 fail to restart on 500/500XL systems** If an upgrade (online or offline) to HCP 7.0 encounters a failure on an HCP 500 or HCP 500XL system, the upgrade will not restart properly and complete. This problem is limited to the restart process and only occurs if another problem prevents the upgrade from completing successfully. **Fix:** The HCP upgrade to 7.0 can now be restarted and complete properly.
HCP-24928 and HCP-24966	**Security Account reset feature requests immediate password reset** The Security Account reset feature logs users out, and forces them to reset passwords immediately after activation. This interrupts the user's workflow. **Fix:** HCP now allows a twenty minute delay before the log out occurs.
HCP-24911 and HCP-24955	**Replication page transfer rate graph doesn't display measurement units** The transfer rate graph's measurement units were displayed in the same font color as the page background color. **Fix:** Changed the font color of the measurement units.
HCP-24925 and HCP-24923	**Default namespace DPLs set to one change when HCP 6.0 or 6.1 upgrades to HCP 7.0** An HCP 6.0 or 6.1 system that has a default namespace with DPL set to one will not be preserved when upgraded to HCP 7.0. It is automatically changed to DPL two because it's the suggested DPL on RAIN systems. **Fix:** HCP now retains the default namespace DPL after upgrade.
HCP-24964 and HCP-24917	**Running NDMP third party applications when rebalancing nodes causes restart** If an NDMP third party back up is running while nodes rebalance, NDMP causes the nodes to restart. **Fix:**An error message now prevents the node rebalance from starting until the NDMP backup is completed.
HCP-24950 and HCP-24948	**Protection service running on HCP system with spindown disks may cause node to restart** The protection service may cause a node restart when writing to a spindown disk. **Fix:** Fixed the code so that the protection service can write to spindown disks.
HCP-24979 and HCP-24905	**Objects named after directories become moveable** Objects with names that are the same as directory names (active or deleted) can be moved over NFS. Objects should not be able to move. If they are moved it causes the node to go down. **Fix:** Corrected the code so objects can no longer be moved.

Ref. number	Description
HCP-24980 and HCP-24871	**Content Verification might try to recover objects in deleted namespace** Content Verification occasionally tries to recover objects in a purged namespace. This causes a node restart. **Fix:** Content Verification now ignores all objects in purged namespaces.
HCP-24991 and HCP-24945	**System may not address internal message overlap** Occasionally a node restarts instead of restarting the connection and resending faulty information. **Fix:** The system now behaves according to protocol. It closes the connection and retries the operation.

Issues resolved in release 7.0

The table below lists issues that were resolved in release 7.0 of HCP. The issues are listed in order by reference number.

Ref. number	Description
HCP-17907	**Cannot view list of service plans with management API** If a tenant is configured to allow it to assign service plans to its namespaces, you can use the HCP Management API to make these assignments. The API, however, does not provide a way to view the service plans that can be assigned. **Fix:** Updated the HCP management API to add a new /availableServicePlans resource to the tenant resource. If service plan selection has been enabled for a given tenant, then tenant administrators can use the /availableServicePlans resource for the applicable tenant resource to query the HCP system to determine which service plans are available for selection.
HCP-18758	**Offline upgrade time estimate inaccurate** Before performing an offline upgrade, HCP Setup displays an estimate of the amount of time the upgrade will take. This estimate may be off by a significant amount of time. **Fix:** This issue has been resolved. HCP Setup now displays a reasonably accurate estimate of the amount of time that an offline upgrade will take.

Ref. number	Description
HCP-20313	**Provide support for SMB2** HCP should support the use of SMBv2.0 for connecting to a CIFS share on an HCP system. **Fix:** CIFS clients are now able to negotiate SMB v2.0 while mounting a CIFS share from HCP. In addition, several clients such as Windows 2008, 2012, and Windows 7 use SMBv2.0 by default.
HCP-20820	**Default JVM memory settings are insufficient for large systems** The default JVM memory settings are not viable for very large systems. **Fix:** For some larger HCP systems, the default JVM memory settings were not sufficient. The maximum Java heap size has been increased to 1GB.
HCP-21741	**Wrong status code returned for GET of namespace you don't own** If you have the namespace management property and no roles, when you use the HCP management API to submit a GET request for a namespace you don't own, HCP incorrectly returns a 503 status code. **Fix:** If you submit a GET request for a namespace that you don't own, HCP now correctly returns a 403 (forbidden) status code.
HCP-22873	**Cannot use Namespace Browser to store objects in namespace with deleted default retention class** When you try to store an object through the Namespace Browser in a namespace for which the default retention setting is a retention class that has been deleted, the request is rejected with a 400 (Bad Request) error code, and the node servicing the request restarts itself. **Fix:** This issue has been resolved. You are now allowed to delete a retention class that's set as the default for an HCP namespace. In addition, if you try to store an object in an HCP namespace for which the default retention class has been deleted, HCP correctly returns an error message and does not store the object.
HCP-23030	**WebDAV PUT fails with a 500 error (java Null Pointer Exception)** When a client sends a WebDAV PUT request to HCP, if no default retention setting is specified in the PUT request, HCP returns a 500 error. **Fix:** This issue has been resolved. HCP no longer returns a Null Pointer Exception in the scenario described above.

Ref. number	Description
HCP-23263	**On an HCP system that uses CR210H servers as nodes, the kernel tick count is not properly synchronized with the HCP system clock** When a client sends a WebDAV PUT request to HCP, if no default retention setting is configured for the namespace that's specified in the PUT request, HCP returns a 500 error. **Fix:** On CR210H servers, an insufficiently designed calculation in the CPU accelerator could sometimes cause an arithmetic overflow that led to a kernel panic on an HCP system that used the Time Stamp Counter (TSC) clock source. Kernel panics occurred most frequently on HCP systems that were using Intel Xeon E5 processors that do not reset the TSC on soft power cycles. A kernel patch has been applied to modify the calculation so that the arithmetic overflow and subsequent kernel panic can no longer occur under these circumstances.
HCP-23464	**Storman restarts frequently when HCP is configured to use Strongbox as a storage tiering target** When HCP is configured to use Strongbox as a storage tiering target, storman frequently issues error messages and restarts. This can occur more than 25 times in a single day. **Fix:** This issue has been resolved. The errors and subsequent storman restarts can no longer occur.
HCP-23596	**External time server inaccessible** If an external time server specified in the HCP system configuration is identified by its fully qualified domain name, HCP may not be able to access that time server. If the HCP system is configured to use an external time server that's identified by its fully qualified domain name and the IP address of that time server is changed in the DNS, HCP will not be able to access that time server. If the system is restarted, it may or may not be able to access the time server. If all of the external time servers specified in the system configuration are inaccessible, HCP uses itself as a time server. **Fix:** After changing the IP address of an external time server in the DNS, to enable HCP to continue to access the external time server, you need to restart the HCP system.

Ref. number	Description
HCP-23624	**The storage tiering service should run constantly when it's scheduled to run for 24 hours** Currently, whenever any service is scheduled to run during a specific time period, the service runs only once during that time period. A scheduled service starts running at the beginning of the scheduled run time period, and if it finishes before the end of the scheduled run time period, the service sits idle until the next scheduled run time. The normal behavior for a scheduled service does not support the current use case for the storage tiering service. The behavior of the storage tiering service should be changed so that the service is always running and looking for "work" (objects to tier) during the entire period of time when it's scheduled to run. **Fix:** The behavior of the storage tiering service has changed. The service still runs once at the beginning of its scheduled run time period. However, if it finishes before the end of that time period, the service continues to look for objects that need to be moved between storage tiers, and restarts whenever it detects such objects.
HCP-23625	**When the storage tiering service runs during its scheduled run time period, it sometimes fails to service all of the objects stored in HCP** Sometimes, when the storage tiering service runs during its scheduled run time period, it stops before it finishes servicing all of the objects in the HCP repository. **Fix:** This issue has been resolved. The storage tiering service now attempts to service all of the objects in the HCP repository before it finishes its run.
HCP-23676	**When executing an snmpwalk command, HCP sometimes returns corrupted OIDs** Sometimes during an snmpwalk, HCP returns corrupted OIDs. The values of the OIDs are correct, but the OIDs themselves are incorrect. **Fix:** The snmp-net library occasionally added a random sequence of numbers after a valid OID. This issue has been resolved.
HCP-23728	**XML well-formed check can not handle UTF byte order mark and UTF16 characters** When ingesting an annotation file to an HCP namespace, if that file contains UTF characters in the range of U+10000-10FFFF (4 Bytes), an HTTP PUT request returns a 400 error code (Bad Request) because the annotation file is determined to contain "Invalid XML in custom metadata" by XML well-formed checking. However, the specifications for the XML1.0/1.1 formats allow the use of UTF-8 characters ranging between U+0000-10FFFF with a few exceptions. So, since these UTF-8 characters are valid characters in an XML file, using these characters in an annotation file should not cause the file to fail the check for well-formed XML. **Fix:** When checking an annotation file to verify that it contains only well-formed XML, HCP no longer generates an exception when a file contains valid UTF-8 characters.

Ref. number	Description
HCP-23764	**The disposition service fails to remove deleted objects from the HCP metadata query engine index** When disposition is enabled for a namespace, the disposition service deletes all expired objects as expected, but it retains the metadata query engine index checkpoints for the deleted objects. As a result, after an indexed object has been deleted, the System Management Console and the Tenant Management Console still include the deleted object in the applicable indexed object count graphs. **Fix:** Previously, the disposition service did not remove deleted objects from the HCP metadata query engine index. The disposition service now removes all deleted objects from the index. Note that for the fix to work a re-index is required. See *Administering HCP* for information about re-indexing.
HCP-23776	**HCP should generate system log messages when the final recovery of the primary system in an active-passive replication link is delayed due to outstanding active writes** When an active/passive replication link has been failed over from the primary system to the replica, if the primary system becomes available again, you can choose to fail back the link at any time. You can then initiate a final recovery procedure to copy all replicated tenants and namespaces from the replica back onto the primary system so that the primary system can once again be used to service client requests. To ensure that all replicated content is written to both systems participating in the link before and after it is failed over, a final recovery procedure cannot be completed while data is still being written to the replica. HCP should generate system log messages to inform the user when the final recovery of the primary system in an active-passive link fails to complete due to outstanding active write operations that are still being performed on the replica. **Fix:** If a final recovery procedure fails to complete because data is still being written to the replica, HCP generates this system log message: "Could not complete restore because there are still active writes."
HCP-23826	**ConcurrentModificationException error causes JVM to crash** On rare occasions, a ConcurrentModificationException error occurs and causes JVM to crash. **Fix:** This issue has been resolved. The error and the subsequent JVM crash can no longer occur.

Ref. number	Description
HCP-23864	**HCP Search Console incorrectly shows replicated objects with a status of "Not Replicated"** In the HCP Search Console, some replicated objects are shown with a status of "Not Replicated". **Fix:** The HCP Search Console now shows the correct "Replicated" status for each replicated object.
HCP-23866	**Cannot use HS3 to connect to an HCP tenant over any user-defined network that does not use the same domain as [hcp_system]** An HS3 connection to a tenant on a user-defined network fails if that network is not configured to use the same domain as the [hcp_system] network. **Fix:** Users can now use the HS3 API to connect to a namespace over any user-defined network, regardless of which domain is associated with the network.
HCP-23925	**HCP spindown storage spin-up and spin-down timeouts are too aggressive for large systems** Spindown volumes sometimes appear to become unavailable when the HCP nodes attempt to spin up or spin down the volumes because the timeout period for a spin-up or spin-down operation is too short. **Fix:** For spin-up and spin-down operations, the default timeout/retry time and the default volume availability wait time have both been increased. As a result, spindown volumes no longer appear unavailable due to spin-up/spin-down timeouts and retries.
HCP-23927	**An IllegalStateException sometimes occurs during a listing of the existing versions of an object** When a RESTful protocol is used to request a list of the existing versions for any given object, an IllegalStateException error sometimes occurs when the list of object versions that's returned to the client is particularly lengthy. **Fix:** This issue has been resolved. Returning a long list of object versions to a client no longer causes HCP to generate an exception.
HCP-24142	**HCP PUT-COPY operation fails on any object with custom metadata** After an upgrade from HCP 5.x to HCP 6.x, when a client attempts to perform a PUT-COPY operation on an object with custom metadata, if the object is contained in a namespace for which content verification is enabled, the PUT-COPY operation fails with a 503 error. **Fix:** This issue has been resolved. PUT-COPY operations no longer fail on objects with custom metadata.

Ref. number	Description
HCP-24380	**A node in an HCP SAIN system unexpectedly issued a java.lang.ArrayIndexOutOfBoundsException error and then rebooted** On rare occasions, a node in an HCP SAIN system may unexpectedly issue a java.lang.ArrayInIndexOutOfBoundsException error and then reboot. This happens because the SimpleDateFormat that's used by the HCP system calendar is not thread-safe. **Fix:** This issue has been resolved.
HCP-24395	**System Management Console incorrectly reports that Active Directory Single Sign-On is disabled for one or more tenants** The System Management Console sometimes issues a warning message that reads, "Active Directory SSO disabled for one or more tenants." When you look at the **Security Events** panel on the **System Events** page, you will see warnings that SPNs are missing for one or more tenants. However, the SPNs are not, in fact, missing. **Fix:** This issue has been resolved. The System Management Console no longer incorrectly issues the warning message, "Active Directory SSO disabled for one or more tenants."
HCP-24711	**Improve metadata query engine performance for operation-based queries** For operation-based queries, the metadata query engine performs better when the client does not specify any directories. Therefore, to maximize. metadata query engine performance, the query request processor should ignore any directory specified in a query request that is just a "/". **Fix:** When performing an operation-based query, the metadata query engine now ignores all empty directories that are specified in the query request.

Ref. number	Description
HCP-24784	**Apply 3rd-party patches to fix known openssl vulnerabilities** The version of openssl that HCP currently uses for SSL communication is known to be vulnerable to various types of attacks. We need to apply patches to openssl to eliminate known security vulnerabilities. **Fix:** In the 2014 version of openssl that HCP uses for SSL communication, the following commonly known security vulnerabilities have been fixed: • Heartbleed (CVE-2014-0160) • SSL/TLS MITM vulnerability (CVE-2014-0224) • DTLS recursion flaw (CVE-2014-0221) • SSL_MODE_RELEASE_BUFFERS NULL pointer dereference (CVE-2014-0198) • SSL_MODE_RELEASE_BUFFERS session injection or denial of service(CVE-2010-5298) • DTLS invalid fragment vulnerability (CVE-2014-0195) • Anonymous ECDH denial of service (CVE-2014-3470) For information about these vulnerabilities and the openssl patches that have been applied to address these vulnerabilities, refer to the openssl security advisory document found here: https://www.openssl.org/news/secadv_20140605.txt.

Issues with workarounds

The table below lists known issues in the current release of HCP for which workarounds exist. The issues are listed in order by reference number.

Ref. number	Description
HCP-5179	**Browser caching** When an object is added to a namespace, deleted, and then added again with the same name, it may appear to have the old content when viewed through a web browser. **Workaround:** To see the new content, clear the browser cache. Be sure to use the applicable browser option to do this rather than restarting the computer.

Ref. number	Description
HCP-7108	**Node restart with cross-mapped storage** In SAIN systems, if a cross-mapped node restarts while one of its physical paths to the storage array is broken, the node remains unavailable. **Workaround:** Fix the broken path and restart the node from the System Management Console.
HCP-13183	**SNMP version 2c traps sent for version 3 traps** HCP can be configured to use SNMP version 3. However, when configured this way, HCP sends version 2c traps instead of the expected version 3 traps. **Workaround:** To receive traps from HCP, have your SNMP application accept SNMP version 2c traps.
HCP-13574	**WebDAV does not correctly list objects with custom metadata** Namespaces can be configured to store WebDAV dead properties in `custom-metadata.xml` files. If regular custom metadata is stored for one or more objects in a directory before this configuration is set, subsequent WebDAV requests for listings of that directory fail with an XML parsing error. **Workaround:** Do not use `custom-metadata.xml` files to store WebDAV properties for an object if any objects in the same directory already have custom metadata.
HCP-16516	**Cannot log into HCP interface as local user in Internet Explorer** With Internet Explorer, if the Active Directory user account with which you're currently logged into Windows is not an account that's recognized by HCP and any of these applies, Internet Explorer displays a **Connect** window instead of the page with the link to the login page for the target interface: • You are trying to access the System Management Console, and support for Active Directory is enabled at the system level. • You are trying to access the Tenant Management Console for an HCP tenant, and Active Directory is enabled as an authentication type for the tenant. • You are trying to access the Namespace Browser for an HCP namespace, and Active Directory single sign-on is enabled for the namespace. If you enter credentials for an HCP user account in the **Connect** window, Internet Explorer returns an error message. **Workaround:** To access the target interface using an HCP user account, click on the **Cancel** button in the **Connect** window to display the page with the link to the login page for the target interface.

Ref. number	Description
HCP-18233	**Changed computer account not added to all applicable groups in Active Directory** When you enable HCP support for Active Directory, the HCP computer account you specify is automatically added to the groups in Active Directory that include the user account you specify. If you then remove the computer account from one or more of those groups and reconfigure Active Directory support with a new computer account, the new computer account is not automatically added to the groups from which the previous computer account was removed. **Workaround:** Do not remove the old computer account from the groups in Active Directory until after you have changed the computer account in HCP. If you have already removed the old computer account from one or more groups, resubmit the Active Directory configuration in HCP without changing the computer account. This puts that computer account back in the groups from which it was removed. When you subsequently change the computer account in HCP, the new computer account will be added to all the groups that include the user account.
HCP-19128	**Downloads with HTTPS fail in Internet Explorer 9** With Internet Explorer 9, attempts to download files (such as chargeback reports and SSL certificates) from URLs that use SSL security (that is, URLs that start with HTTPS) fail. **Workaround:** In Internet Explorer 9: **1.** On the **Tools** menu, select **Internet Options**. **2.** In the **Internet Options** window, click on the **Advanced** tab. **3.** On the **Advanced** page, under **Security**, deselect the **Do not save encrypted pages to disk** option. **4.** Click on the **OK** button.
HCP-20401	**Node restart due to large element content in annotation** While XML checking of custom metadata is enabled, if an annotation is added to an object where the content of an element in the annotation is very large, a node may restart itself. **Workaround:** Disable XML checking for the namespace that contains the object.
HCP-20827	**Delayed read from replica when external storage unavailable** In a replicated namespace, if the only copy of the data for an object is in external storage and that storage is unavailable, NFS and WebDAV requests for the object may time out for several tries before HCP retrieves the object from the replica. **Workaround:** Either bring the external storage back online, or retry the request in five minutes.

Ref. number	Description
HCP-23881	**Nodes may fail with the error message "Max Connections hit: Could not get a connection, pool is exhausted"** Even if the supported limit of 200 connections is not reached, if too many clients attempt to connect to the same namespace at the same time, one or more nodes in the HCP system may fail with the error message, "Max Connections hit: Could not get a connection, pool is exhausted". **Workaround:** Upgrade to release 7.1 or later of HCP and increase your system RAM on all nodes. At least 32GB of RAM needs to be added.
HCP-24155	**When performing an add-drives procedure on an HCP node, an existing node sometimes issues a "barrierWait" message and then hangs** When performing an add-drives procedure on an HCP node, if one node fails, its partner node may issue a "Waiting for others at barrierWait" message and then hang. **Workaround:** To get the existing node back into a working state, press Control-C to cancel the drive addition procedure. You can then restart the procedure.
HCP-24156	**Cannot use a domain name to connect to a namespace on an IPv6 or dual-mode HCP network** HCP-DM does not support the use of IPv6 addresses to connect to a namespace on an HCP system. HCP-DM can use IPv4 addresses to connect to a namespace on a dual-mode HCP network. However, if HCP-DM tries to use a domain name to connect to a namespace on a dual-mode network, when the DNS resolves the domain name, it will return both IPv6 and IPv4 addresses for the network. If HCP-DM then tries to use the IPv6 addresses to connect to the namespace, the connection will fail. **Workaround:** To ensure that HCP-DM can successfully connect to a namespace on a dual-mode HCP network, you need to configure HCP-DM to connect to that namespace using the IPv4 addresses for the network.
HCP-24222	**Enabling service plan selection for an existing tenant causes all namespaces owned by that tenant to inherit the Default service plan** When you enable service plan selection for an existing tenant, the Default service plan is automatically assigned to all namespaces owned by that tenant. If a different service plan was originally assigned to the tenant, the Default service plan replaces the original plan for all namespaces owned by the tenant. **Workaround:** After enabling service plan selection for an existing tenant, make sure the tenant administrator verifies that each namespace has the correct service plan assigned to it. If not, change the service plan assignment at the namespace level.

Ref. number	Description
HCP-24589	**When attempting to update annotations for objects that have been tiered to one or more types of cloud storage, HCP sometimes returns a 503 error** When a HCP attempts to update annotations on objects that have been tiered to cloud storage, the updates will fail with a 503 error if HCP is unable to connect to the applicable cloud storage service endpoints or if HCP is unable to access the applicable cloud storage buckets, containers, or namespaces. **Workaround:** Restore the connections between HCP and each applicable cloud storage service endpoint and make sure HCP can successfully access each applicable cloud storage bucket, container, and namespace. You should then be able to successfully update the annotations for any objects stored in each bucket, container, and namespace.
HCP-25388	**While HTTPS is enabled, HCP S Series Nodes fail to create storage components when added to the HCP system by virtual IP** An S Series node cannot use HTTPS when being added by virtual IP to HCP. While HTTP is enabled, HCP does not create storage components for S Series Nodes added by virtual IP address. **Workaround:** On the System Management Console, when adding an S Series Node through virtual IP, go to the **Connection** tab of the **Add Node** wizard, deselect the **Use HTTP for management** option and, under the **Advanced** panel, deselect the **Use HTTPS for data access** option before completing the add node procedure.
HCP-26037	**The Adding Logical Volumes service might fail if adding previously used, formatted LUNs** Occasionally during the add LUN service procedure, previously used, formatted LUNs might not be added to all nodes. If this occurs, the error message, "Failed to execute Partx" appears. **Workaround:** Restart the service procedure.

Ref. number	Description
HCP-26043	**Incorrectly shutting down and restarting a replication link when updating the signed certificate causes the replication link to fail** If you incorrectly shut down and restart a replication link while uploading an SSL certificate, the replication link refutes the certificate and fails. **Workaround**: Follow this certificate upload procedure: **1.** Upload a new certificate on the primary and replica systems. **2.** Remove the expired certificate from the primary and replica systems. **3.** Select the **Shut down all links** option from the **Replication settings** menu on the primary and replica systems. **4.** Select **Start up all links** on the primary and replica systems.
HCP-26066	**Log download fails under certain conditions** Log download initiated through the System Management Console could fail due to external issues such as networking. **Workaround:** Restart the log download.
HCP-26128	**HTTPS certificate errors appear during failover** Sending HTTPS requests to system A in a replication link that has failed over causes report certificate errors because the **Subject Common Name** in the certificate does not match the domain name in the request. **Workaround:** Add **Subject Alternative Name** entries to the certificates used by HCP for HTTPS.
HCP-26775	**Cannot download certificates from HCP through Internet Explorer 8** When you try to download a certificate from HCP using Internet Explorer (IE) 8, you may receive the "Unable to download." error message. This is caused by a known I.E 8 issue. **Workaround:** For more information on this issue, see the Microsoft Knowledge Base article, https://support.microsoft.com/en-us/kb/323308.

Ref. number	Description
HCP-27121	**Cannot download HCP system logs during an online upgrade to release 7.2 of HCP** During an online upgrade to release 7.2 of HCP, the HCP system logs cannot be downloaded from the System Management Console. **Workaround:** During the online upgrade, access the System Management Console by entering the IP address of a node that has already upgraded into your web browser. Perform the log download procedure through the targeted node.
HCP-27882	**After upgrade to 7.2, some third party applications receive HTTP 401 error to PUT requests** With release 7.2 of HCP, SPNEGO changes make certain third party applications incompatible with HCP. **Workaround:** Contact HCP support services to enable third party compatibility.
HCP-28092	**Occasionally HCP cannot join AD due to insufficient permissions** Occasionally HCP fails to connect to AD and returns the insufficient permissions to modify SPNs error even though all the permissions are set. **Workaround:** Contact HDS Support to apply a domain whitelist filter to a single Domain Controller during the join process or continue retrying the AD join process.
HCP-29088	**Modifications to active replication link may cause replication failure** Performing a change to an an unpaused tenant in a replication link may cause object replication failure. **Workaround:** Run the replication verification service or ensure that read from replica is enabled on the namespace.

HCP-29756	**Updates to AWS SDK cause HS3 PUT request failure**
	With release of AWS 1.9.3, if the HCP transport layer is SSL, Amazon sets the default setting to unsigned payload. HCP does not support unsigned payloads, which causes HS3 PUT requests to fail with an invalid signature error.
	
	Workaround: Prior to issuing an HS3 object PUT request, disable unsigned payloads in the Java SDK or wrap your PUT command with the following code:
	
	`s3.setS3ClientOptions(S3ClientOptions.builder().setPayloadSigningEnabled(true).build());`
	`s3.putObject(por);`
	`s3.setS3ClientOptions(S3ClientOptions.builder().build());`

Ref. number	Description
HCP-29915	**HCP installation fails due to imbalanced node volumes** If an HCP system does not have the same amount of archive volumes on each of its nodes, the HCP installation may fail. **Workaround:** Balance the volumes across all nodes in the HCP system before attempting the HCP install.
HCP-30102 and HCP-30107	**Unavailable objects due to new bucket account** When you select a new account for a bucket in a storage component, HCP does not automatically update the storage pool that uses that bucket. To force the storage pool update, you need to click on **Update Settings** in the **Settings** panel for the pool. Attempts to access objects in the bucket between the time you select the new account and the time you update the storage pool cause those objects to be marked unavailable. Clients using the HTTP, HS3, or HSwift protocol can access the affected objects by appending this query parameter to the request URL: allow_hash_mismatch=true Clients using other protocols cannot access these objects. Affected objects in S Series Node buckets become available again the next time the HCP content verification service processes those objects. **Workaround:** Before selecting the new account for the bucket, pause the storage component containing the bucket, and deselect all permissions on the **Security** ▸ **Permissions** page in the System Management Console. After updating the storage pool, resume the component and reselect the permissions you want.

Other known issues

The table below lists known issues in the current release of HCP for which no workarounds exist. The issues are listed in order by reference number.

Ref. number	Description
HCP-804	**HCP Data Migrator can set the value of the hold parameter to true, but not to false** HCP Data Migrator can be used to place an object on hold by updating the system metadata for the object to set the hold parameter to true. However, you cannot use the HCP Data Migrator to remove a hold from an object because the HCP Data Migrator cannot set the value of the hold parameter to false.

Ref. number	Description
HCP-5153	**False log messages with lowest-numbered node addition** When a new node is added to an HCP system, a message about it is written to the system log. If the number of the new node is lower than that of any existing nodes, the same message is written for each existing node, as if it were newly added.
HCP-7043	**Displaying UTF-16-encoded objects** Objects with content that uses UTF-16 character encoding may not be displayed as expected due to the limitations of some browser and operating system combinations. Regardless of the appearance on the screen, the object content HCP returns is guaranteed to be identical to the data before it was stored.
HCP-8385	**Exposed internal mechanism for dead properties for collections** HCP uses an internal mechanism for storing WebDAV dead properties for a collection. This mechanism entails the creation of a dummy object named `.webdav_properties`. This object is inappropriately: • Included in the count of objects in the namespace • Exposed through the HTTP, CIFS, and NFS protocols • Returned by searches for which it meets the search criteria If you are storing dead properties for collections, do not delete any `.webdav_properties` objects.
HCP-8570	**Missed log messages when no leader node** Normally, one node in an HCP system is responsible for writing messages to the system log. This node is called the **leader node**. Rarely, brief periods occur during which no leader node exists (for example, because the leader node has failed and a new leader node has not yet been established). During such periods, messages for which the leader node is responsible are not written to the log.
HCP-8665	**Shredding in SAIN systems** In SAIN systems, HCP may not effectively execute all three passes of the shredding algorithm when shredding objects. This is due to the fact that some storage arrays make extensive use of disk caching. Depending on the particular hardware configuration and the current load on the system, some of the writes from the shredding algorithm may not make it from the cache to disk.

Ref. number	Description
HCP-9212	**Log display skips messages** When you page through a display of log messages in the System or Tenant Management Console, some messages may be skipped. This happens because the Console retrieves the next or previous group of messages based on the message timestamps. Each time you request a next page of messages, the Console starts the new page with the message with the next later timestamp from the last message on the current page. If a page boundary falls between multiple messages with the same timestamp, retrieving messages starting with the next timestamp skips the messages that come after the page break. The equivalent process happens when you request a previous page of message. As additional messages are added to the log, the page boundaries change, with the result that previously skipped messages reappear.
HCP-9360	**Browser pages for large directories** You can view the contents of a namespace in a web browser through HTTP (default namespace only) or WebDAV. Some browsers, however, may not be able to successfully render pages for directories that contain a very large number of objects.
HCP-11317	**Using NFS to delete objects open for read** Using NFS, if you try to delete an object that is currently open for read on the same client, HCP returns this error: Read-only file system.
HCP-11667	**Appending to objects on unavailable nodes** If an object is open for append on a node that becomes unavailable, attempts to append to the object fail.
HCP-12089	**Cannot ingest very large email attachments** HCP fails to ingest email attachments substantially greater than 400 MB. In such cases, the client receives a 221 return code.
HCP-18352	**HCP unresponsive after Active Directory cache cleared while Active Directory is unavailable** If you clear the Active Directory cache while HCP cannot communicate with Active Directory, the HCP system becomes unresponsive for up to ten minutes.
HCP-18654	**No success or error message in response to action taken in Console** Occasionally, the System and Tenant Management Consoles do not display any success or error messages in response to an action that results in a fresh display of the page on which the action was taken.

Ref. number	Description
HCP-18700	**Standby spindown volumes possibly not reported in data migrations** In an HCP SAIN system with spindown storage, when you configure a data migration, the logical volumes (LUNs) reported for each node in the **Configuration** panel on the **Migration** page in the HCP System Management Console may not include standby volumes on spindown storage. Additionally, the standby spindown volumes may not be included in the downloaded configuration report. Although these volumes may not be reported, they are always properly included in the data migration.
HCP-19123	**Objects incorrectly reported as irreparable or unavailable after data migration** During a data migration, the migration service may incorrectly report one or more objects as irreparable or unavailable. After the data migration is complete, you can run the content verification service to clear these errors.
HCP-21365	**Alert about Active Directory connection with online HCP upgrade** In an HCP system that's configured to support Active Directory, during an online upgrade and for a short time after the upgrade is complete, the System Management Console may show an alert indicating a problem with support for Active Directory. This alert is most likely false and will go away on its own. If Active Directory authentication is working, the alert can be safely ignored.
HCP-21056	**Network interface event upon MTU change** When you change the MTU for a network, the network interface may go down and then come back up on nodes that are Dell 1950 servers.

Ref. number	Description
HCP-22241	**Username mappings applied to AD users of HCP namespaces** For the default namespace, Active Directory user authentication is implemented through the use of a username mapping file that associates AD usernames with UIDs and GIDs. If an AD user included in the username mapping file also has access to an HCP namespace, the objects that the user stores in the HCP namespace have the UID and GID specified in the username mapping file. As a result, a user using CIFS for authenticated access who is included in the username mapping file or a user using NFS has access to such an object only if one of these is true: • With CIFS, the UID for the user in the username mapping file matches the object UID. • With NFS, the user's UID matches the object UID. • With CIFS, the user is included in the AD group identified by the object GID. • With NFS, the user is included in the group identified by the object GID. • WIth CIFS, the user has been granted access to the object by the object ACL. • The object ACL grants all users access to the object. • The minimum data access permissions for the namespace grant all users access to all objects in the namespace. Users using HTTP, HS3, WebDAV, CIFS for authenticated access who are not included in the username mapping file, or CIFS for anonymous access have access to such an object regardless of the object UID and GID.
HCP-23070	**False alerts for network with same name as deleted network** If you create a network with at least one node IP address, then delete the network, and then create a new network with the same name as the deleted network and no node IP addresses, the **Overview**, **Hardware**, **Storage Node**, and **Networks** pages in the System Management Console display alerts indicating that a network error exists. Additionally, HCP writes this message to the system log: Network interface bond0.*xxxx* for network *network-name* is not functioning properly. When you subsequently assign IP addresses for the network to one or more nodes, the alerts disappear.
HCP-23115	**Incorrect release number in *Using HCP Data Migrator*** The preface in the version of *Using HCP Data Migrator* that comes with HCP release 6.1 incorrectly states that the book applies to release 6.0.1.

Ref. number	Description
HCP-24436	**Clearing the AD cache causes inconsistent directory permissions** Following a clearing of the AD cache on an HCP system that's accessing AD over CIFS, when users access a given CIFS share, they will find that their file permissions have changed to root/root, with the exception of the first user to access the share following the clearing of the AD cache. That user will see the original permissions on his/her file/folders, but all others will be root/root. All other users that connect to the CIFS share will only see root/root for existing files/folders.
HCP-24472	**When using AD for HCP authentication, if a user has an AD username that includes a % character, that user cannot access the HCP system** If you attempt to log into the HCP System Management Console or Tenant Management Console using an Active Directory username that includes a % (percent) character, the HCP user authentication fails.
HCP-24864	**HDvM cannot send updates to an HCP system with IPv6 only mode enabled** An HCP system with IPv6 only mode enabled can successfully connect to the HDvM server, but cannot receive HDvM updates.
HCP-24887	**When performing the TrueCopy storage array replication procedure, the san_ update command may fail** When using the TrueCopy procedure to replicate the HCP system OS and data LUNs to a different storage array, the san_update command may fail with an error that the file system on the source system differs from the file system on the second.
HCP-25595	**Pausing or failing an NFS write operation may cause HCP system processes to hang** Pausing or failing an NFS write operation increases the chances of HCP system processes hanging.
HCP-25602	**Occasionally while migration service runs, migration status shows incorrect values** Occasionally while the Migration service is running, the migration status values for the total number of bytes being migrated and the total number of objects being migrated are incorrect. This occurs regardless of how many bytes or objects are actually migrated. Once the migration completes, the migration status values become accurate.
HCP-25655	**Occasionally an offline upgrade procedure on HCP system with spindown fails** Occasionally an HCP system with spindown may fail an offline upgrade because the disk mounting procedure times out.
HCP-25697	**AD 100 Winbind error occasionally causes HCP nodes to restart** HCP system communication errors with AD may cause winbind to restart. If this happens more than 100 times, the HCP system restarts.

Ref. number	Description
HCP-25731	**Upgrade NTP to fix vulnerabilities** The NTP that HCP currently uses was discovered to have some vulnerabilities. For information about these vulnerabilities, refer to the NTP security advisory document found here: https://www.us-cert.gov/ncas/current-activity/2014/12/19/Vulnerabilities-Identified-Network-Time-Protocol-Daemon.
HCP-25761	**If your HCP system has data tiered to public cloud, the upgrade process to version 7.1 of HCP is extended** If your HCP system has data tiered to public cloud, metrics need to recompute when upgrading to version 7.1 of HCP. This extends the upgrade time.
HCP-25997	**Node recovery does not work for the HCP 500XL with new disks** Node recovery procedures fail with unformatted database drives.
HCP-26058	**Upgrading to HCP 7.2 or later prevents HCP from connecting to HCP Data Migrator** Release 7.2 and later of HCP use a different SSL cipher than previous releases. HCP-DM does not support these ciphers if HCP-DM is run with an outdated Java runtime.
HCP-26158	**Under specific conditions, creating an active/active replication link between two systems causes nodes to reboot** Under specific conditions, creating a replication link between two systems running version 7.1.1 of HCP that have ingested objects and have multiple namespaces causes nodes to reboot.
HCP-27176	**The Network page Advanced Settings tab appears blank when the HCP system is read only** When an HCP system is in a read only state, the **Advanced Settings** tab on the System Management Console **Configuration►Networks** page appears blank.
HCP-27417	**List directory request appears slow with versioning enabled and a lot of deletes** When HCP responds to a list directory request sent through any gateway, HCP needs to list all of its directories internally, including ones that have been deleted, before responding to the request. If versioning is enabled and a lot of directories have been deleted, the list of directories HCP produces may appear disproportionately small compared to the time it takes for HCP to process the request.
HCP-27456	**Storage retirement fails to complete on HCP S Series Node with irreparable objects** If storage retirement is run on an S Series Node that contains irreparable objects, the retirement process stops and restarts periodically but never completes.
HCP-27737	**HCP system raises full capacity alarm if a single volume is over 95% full** If a single volume in an HCP system becomes 95% full, the full file system warning is triggered for the system.
HCP-27757	**Active Directory node account not removed from HCP system when node retired** Running the retire node procedure on a node with Active Directory enabled removes the SPNs from AD, but the retire node procedure fails to remove the AD node account related to the retired node.

Ref. number	Description
HCP-27810	**Occasionally incorrect error message appears when switching tabs during replication schedule update** After creating an active/active replication link, clicking on the **Update Schedule** button on the System Management Console **Services►Replication►Schedule** page, and switching between the **local** and **remote** schedule tabs, an error may appear even though the replication link is working properly.
HCP-28129	**System Management Console Service Plan Tenant list scrolls infinitely** If the **Tenant list** on the System Management Console **Service Plan** page is populated with more tenants than can fit on the list and the list generates a scroll bar, the list will scroll infinitely, showing blank spaces after the inital collection of tenants.
HCP-29573	**Changing HCP-VM network adapter from e1000 to VMXnet3 causes VLAN performance issues** On an HCP-VM with VLANs enabled, converting from an e1000 to VMXnet3 network adapter causes VLAN performance issues.
HCP-30018	**Namespace browser cannot load directory due to ASCII characters in object name** The namespace browser cannot display the contents of a directory that contains an object with any of the following ASCII characters in its name: %00-%0F, %10-%1F, or %20.

Accessing product documentation

Product documentation is available on Hitachi Data Systems Support Connect: https://knowledge.hds.com/Documents. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

Hitachi Data Systems Support Portal is the destination for technical support of products and solutions sold by Hitachi Data Systems. To contact technical support, log on to Hitachi Data Systems Support Connect for contact information: http://portal.hds.com

Hitachi Data Systems Community is a global online community for HDS customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to http://community.hds.com, register, and complete your profile.

Note: If you purchased HCP from a third party, please contact your authorized service provider.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, California 95050-2627
U.S.A.
www.hds.com

Regional Contact Information

Americas
+1 408 970 1000
info@hds.com

Europe, Middle East, and Africa
+44 (0) 1753 618000
info.emea@hds.com

Asia Pacific
+852 3189 7900
hds.marketing.apac@hds.com

Hitachi Data Systems

RN-RC001-47