

Hitachi Content Platform for Cloud Scale

v2.0.0

HCP for Cloud Scale Administration Guide

This document contains instructions for configuring, managing, monitoring, and troubleshooting the Hitachi Content Platform for cloud scale (HCP for cloud scale) software, storage components, and users. It describes the Object Storage Management and System Management applications.

© 2020 Hitachi Vantara LLC. All rights reserved.

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

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

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

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

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

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

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

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

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

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

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

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

EXPORT CONTROLS - Licensee will comply fully with all applicable export laws and regulations of the United States and other countries, and Licensee shall not export, or allow the export or re-export of, the Software, API or Materials in violation of any such laws or regulations. By downloading or using the Software, API, or Materials, Licensee agrees to the foregoing and represents and warrants that Licensee is not located in, under the control of, or a national or resident of any embargoed or restricted country.

Contents

Preface.....	13
About this document.....	13
Intended audience.....	13
Product version.....	13
Release notes.....	13
Related documents.....	14
Document conventions.....	14
Accessing product documentation.....	16
Getting help.....	16
Comments.....	16
Chapter 1: Getting started.....	17
Introducing HCP for cloud scale.....	17
Storage components, buckets, and objects.....	17
Data access.....	18
Data access control.....	19
Data security.....	19
Bucket synchronization.....	21
Object locking.....	21
S3 Select.....	22
S3 event notification.....	23
Supported limits.....	23
High availability.....	23
Scalability of instances, service instances, and storage components....	24
Site availability.....	24
Service availability.....	25
Metadata availability.....	25
Object data availability.....	25
Network availability.....	26
Failure recovery.....	26
Instance failure recovery.....	26
Service failure recovery.....	26
Storage component failure recovery.....	27
Support for the Amazon S3 API.....	27

HCP for cloud scale APIs.....	34
Object Storage Management API.....	34
System Management API.....	35
Security and authentication.....	35
User accounts.....	35
API access.....	36
Network isolation and port mapping.....	37
Logging in.....	37
HCP for cloud scale applications.....	38
Switching between applications.....	39
Providing unseal keys to KMS service.....	39
Serial number.....	40
Entering your serial number.....	40
Object Storage Management application instructions.....	40
Related REST API methods.....	41
Displaying your serial number.....	41
Object Storage Management application instructions.....	41
Related REST API methods.....	41
License.....	41
Uploading your license.....	41
Object Storage Management application instructions.....	41
Related REST API methods.....	42
Enabling encryption.....	42
Object Storage Management application instructions.....	42
Related REST API methods.....	43
Defining subdomain for S3 Console application.....	43
Chapter 2: Storage components.....	45
Adding a storage component.....	45
Object Storage Management application instructions.....	45
Related REST API methods.....	47
Modifying a storage component.....	47
Object Storage Management application instructions.....	47
Related REST API methods.....	48
Activating a storage component.....	48
Object Storage Management application instructions.....	48
Related REST API methods.....	48
Deactivating a storage component.....	48
Object Storage Management application instructions.....	49
Related REST API methods.....	49
Marking a storage component as read-only.....	49
Object Storage Management application instructions.....	49

Related REST API methods.....	49
Marking a storage component as read-write.....	50
Object Storage Management application instructions.....	50
Related REST API methods.....	50
Viewing storage components.....	50
Object Storage Management application instructions.....	51
Related REST API methods.....	51
Displaying storage component analytics.....	51
Displaying counts of storage components.....	51
Object Storage Management application instructions.....	51
Related REST API methods.....	52
Metrics.....	52
Displaying the active object count.....	52
Object Storage Management application instructions.....	52
Displaying metrics.....	52
Object Storage Management application instructions.....	53
Available metrics.....	53
Examples of graphed metrics.....	59
Tracing requests and responses.....	61
Displaying traces.....	61
Traceable operations.....	61
Chapter 3: Services.....	68
Service categories.....	68
HCP for cloud scale services.....	68
Viewing services.....	85
Viewing all services.....	85
Viewing individual service status.....	86
Related CLI commands.....	87
Related REST API methods.....	87
Listing service ports.....	87
Managing services.....	87
Moving and scaling services.....	88
Relocating services.....	88
Related CLI commands.....	90
Related REST API methods.....	90
Scaling Metadata Gateway instances.....	91
Configuring service settings.....	91
Configuring service settings.....	92
Related CLI commands.....	92
Related REST API methods.....	92
Repairing services.....	92

Configuring TLS cipher suite.....	93
Avoiding Message Queue shutdown.....	93
Chapter 4: Monitoring.....	95
Monitoring instances.....	95
Viewing all instances.....	95
Viewing the services running on an instance.....	96
Related CLI commands.....	97
Related REST API methods.....	97
Monitoring services.....	97
Viewing all services.....	98
Viewing individual service status.....	98
Related CLI commands.....	99
Related REST API methods.....	99
Monitoring processes.....	99
Monitoring service operations.....	99
Related CLI commands.....	100
Related REST API methods.....	100
Monitoring system processes.....	100
Related CLI commands.....	100
Related REST API methods.....	101
System events.....	101
Related CLI commands.....	101
Related REST API methods.....	101
HCP for cloud scale events.....	102
Alerts.....	105
Viewing alerts.....	109
Object Storage Management application instructions.....	110
Related CLI commands.....	110
Related REST API methods.....	110
Related REST API methods (HCP for cloud scale).....	110
Email notification rules.....	110
Creating email notification rules.....	111
Related CLI commands.....	112
Related REST API methods.....	112
Creating syslog notification rules.....	112
Creating syslog notification rules.....	113
Related CLI commands.....	114
Related REST API methods.....	114
Logs and diagnostic information.....	114

Chapter 5: Security..... 119

Granting access to users.....	119
Setting the session timeout limit.....	119
Related CLI commands.....	119
Related REST API methods.....	120
Setting the refresh token timeout limit.....	120
Related CLI commands.....	120
Related REST API methods.....	120
Setting the CORS Allowed Origins.....	120
Related CLI commands.....	121
Related REST API methods.....	121
Identity providers.....	121
Adding identity providers.....	121
Related CLI commands.....	121
Related REST API methods.....	122
Identity provider configuration settings.....	122
User information caching.....	124
Related CLI commands.....	125
Related REST API methods.....	125
Viewing identity providers.....	125
Related CLI commands.....	125
Related REST API methods.....	125
Deleting identity providers.....	126
Related CLI commands.....	126
Related REST API methods.....	126
Groups.....	126
Adding groups.....	126
Related CLI commands.....	127
Related REST API methods.....	127
Viewing groups.....	127
Related CLI commands.....	127
Related REST API methods.....	127
Assigning roles to groups.....	128
Related CLI commands.....	128
Related REST API methods.....	128
Deleting groups.....	128
Related CLI commands.....	128
Related REST API methods.....	128
Roles.....	129
Creating roles.....	129
Related CLI commands.....	130

Related REST API methods.....	130
Viewing roles.....	130
Related CLI commands.....	130
Related REST API methods.....	130
Editing roles.....	130
Related CLI commands.....	131
Related REST API methods.....	131
Deleting roles.....	131
Related CLI commands.....	132
Related REST API methods.....	132
Permissions.....	132
Changing the admin account password.....	136
Related CLI commands.....	136
Related REST API methods.....	136
Certificates - HCPCS.....	136
Viewing installed certificates.....	137
Related CLI commands.....	137
Related REST API methods.....	137
Adding data source certificates.....	137
Uploading data source certificates manually.....	137
Related CLI commands.....	138
Related REST API methods.....	138
Changing the system certificate.....	138
System certificate considerations.....	138
Installing a certificate you created.....	139
Installing a new self-signed certificate.....	140
Creating a CSR and installing the returned certificate.....	141

Chapter 6: System management..... 145

Setting the system hostname.....	145
Changing system hostname.....	145
Related CLI commands.....	145
Related REST API methods.....	146
System scaling.....	146
Networking.....	146
Load balancing.....	148
Script to enable S3 traffic redirection.....	149
Script to disable S3 traffic redirection.....	150
Script to list S3 traffic redirection.....	151
Handling network changes.....	151
Safely changing an instance IP address.....	152
After a network change.....	152

Volumes.....	152
Viewing volumes.....	154
Viewing job volumes.....	154
Viewing service volumes.....	154
Instances.....	155
About master and worker instances.....	155
Single-instance systems versus multi-instance systems.....	155
Requirements for running system instances.....	156
Hardware requirements.....	156
Operating system and Docker minimum requirements.....	157
Operating system and Docker qualified versions.....	157
Docker considerations.....	158
SELinux considerations.....	159
Supported browsers.....	159
Time source.....	159
Adding new instances.....	159
Before adding a new instance.....	159
Install Docker on each server or virtual machine	160
Configure Docker on each server or virtual machine.....	160
(Optional) Configure Docker volume drivers.....	161
Configure maximum map count setting.....	161
(Optional) Enable or disable SELinux on each server or virtual machine.....	161
Configure the firewall rules on each server or virtual machine.....	161
Install and configure NTP.....	162
Run Docker on each server or virtual machine.....	162
Unpack the installation package.....	162
Set up networking.....	163
Run the setup script on each server or virtual machine.....	163
Start the application on each server or virtual machine.....	165
Configure services and jobs on the new instances.....	166
Viewing instances.....	166
Viewing all instances.....	166
Viewing the services running on an instance.....	167
Related CLI commands.....	168
Related REST API methods.....	168
Removing instances.....	169
(Optional) Shut down the instance you want to remove.....	169
Remove the shut down instance from the system.....	170
Replacing a failed instance.....	170
Plugins.....	171
Viewing installed plugins.....	171

Related CLI commands.....	171
Related REST API methods.....	171
Upgrading plugin bundles.....	172
Related CLI commands.....	172
Related REST API methods.....	172
Setting the active plugin bundle version.....	172
Related CLI commands.....	173
Related REST API methods.....	173
Deleting plugin bundles.....	173
Related CLI commands.....	173
Related REST API methods.....	173
Packages.....	173
Exporting packages.....	174
Related CLI commands.....	174
Related REST API methods.....	174
Importing packages.....	174
Related CLI commands.....	175
Related REST API methods.....	175
Setting a login welcome message.....	176
Related CLI commands.....	176
Related REST API methods.....	176
Updating the system.....	176
Applying a system update.....	178
Related CLI commands.....	178
Related REST API methods.....	178
Viewing update history.....	179
Related CLI commands.....	179
Related REST API methods.....	179
Removing the system.....	179

Chapter 7: Best practices..... 181

Best practices for distributing services.....	181
Best practices for maintaining system availability.....	181

Chapter 8: Reference..... 182

Troubleshooting.....	182
CLI reference.....	183
Accessing the CLI tools on a system instance.....	183
Installing CLI tools on your computer.....	183
Syntax.....	184
Viewing available commands.....	185
Viewing request models.....	185

Editing configuration preferences.....	186
System error responses.....	187
REST API Reference.....	188
Input and output formats.....	188
Access and authentication.....	188
Viewing and using REST API methods.....	190
Error responses.....	191

Preface

About this document

This document contains instructions for configuring, managing, monitoring, and troubleshooting the Hitachi Content Platform for cloud scale (HCP for cloud scale) software, storage components, and users. It describes the Object Storage Management and System Management applications.

Intended audience

This document is intended for people who are configuring, managing, or monitoring HCP for cloud scale systems. It assumes you have some experience using computer software through a graphical user interface (GUI) or application programming interface (API).

Product version

This document applies to v2.0.0 of Hitachi Content Platform for cloud scale.

Release notes

Read the release notes before installing and using this product. They may contain requirements or restrictions that are not fully described in this document or updates or corrections to this document. Release notes are available on Hitachi Vantara Support Connect: <https://knowledge.hitachivantara.com/Documents>.

Related documents

This is the set of documents containing information about HCP for cloud scale. You should have these documents available before using the product.

- *Hitachi Content Platform for Cloud Scale Release Notes* (RN-HCPCS004-07): This document is for customers and describes new features, product documentation, and resolved and known issues, and provides other useful information about this release of the product.
- *Installing Hitachi Content Platform for Cloud Scale* (MK-HCPCS002-07): This document gives you the information you need to install or update the HCP for cloud scale software.
- *Hitachi Content Platform for Cloud Scale Administration Guide* (MK-HCPCS008-03): This document explains how to use the HCP for cloud scale applications to configure and operate a common object storage interface for clients to interact with; configure HCP for cloud scale for your users; enable and disable system features; and monitor the system and its connections.
- *Hitachi Content Platform for Cloud Scale S3 Console Guide* (MK-HCPCS009-00): This document is for end users and explains how to use the HCP for cloud scale S3 Console application to use S3 credentials and to simplify the process of creating, monitoring, and maintaining S3 buckets.
- *Hitachi Content Platform for Cloud Scale Management API Reference* (MK-HCPCS007-05): This document is for customers and describes the management application programming interface (API) methods available for customer use.





Document conventions

This document uses the following typographic conventions:

Convention	Description
Bold	<ul style="list-style-type: none"> ▪ Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click OK. ▪ Indicates emphasized words in list items.
<i>Italic</i>	<ul style="list-style-type: none"> ▪ Indicates a document title or emphasized words in text. ▪ Indicates a variable, which is a placeholder for actual text provided by the user or for output by the system. Example: <code>pairedisplay -g group</code> <p>(For exceptions to this convention for variables, see the entry for angle brackets.)</p>

Convention	Description
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairedisplay -g oradb</code>
< > angle brackets	Indicates variables in the following scenarios: <ul style="list-style-type: none"> Variables are not clearly separated from the surrounding text or from other variables. Example: <pre>Status-<report-name><file-version>.csv</pre> Variables in headings.
[] square brackets	Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a b } indicates that you must choose either a or b.
vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b.

This document uses the following icons to draw attention to information:

Icon	Label	Description
	Note	Calls attention to important or additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).
	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

Accessing product documentation

Product user documentation is available on Hitachi Vantara Support Connect: <https://knowledge.hitachivantara.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 Vantara Support Connect](https://support.hitachivantara.com/en_us/contact-us.html) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to Hitachi Vantara Support Connect for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

[Hitachi Vantara Community](https://community.hitachivantara.com) is a global online community for Hitachi Vantara 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 community.hitachivantara.com, register, and complete your profile.

Comments

Please send us your comments on this document to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Thank you!

Chapter 1: Getting started

Hitachi Content Platform for cloud scale (HCP for cloud scale) is a software-defined object storage solution that is based on a massively parallel microservice architecture and is compatible with the Amazon Simple Storage Service (S3) application programming interface (API).

Introducing HCP for cloud scale

HCP for cloud scale is especially well suited to service applications requiring high bandwidth and compatibility with the Amazon S3 API.

HCP for cloud scale can federate S3 compatible storage from virtually any private or public source and present the combined capacity in a single, centrally managed, global name space.

You can install HCP for cloud scale on any server, in the cloud or on premise, that supports the minimum requirements.

HCP for cloud scale supports S3 event notification through a graphical user interface or through GET and PUT Bucket Notification configuration.

HCP for cloud scale lets you manage and scale storage components. You can add storage components, monitor their states, and take them online or offline for purposes of maintenance or repair. HCP for cloud scale provides functions to send notification of alerts, track and monitor throughput and performance, and trace actions through the system.

Storage components, buckets, and objects

Storage components

A *storage component* is an Amazon S3 compatible storage system, running independently, that HCP for cloud scale manages as a back end to store object data. To an S3 client using HCP for cloud scale, the existence, type, and state of storage components are transparent.

HCP for cloud scale supports the following storage systems:

- Amazon S3
- Hitachi Content Platform (HCP)
- HCP S Series Node
- Any Amazon S3 compatible storage service

Buckets

A *bucket* is a logical collection of secure data objects that is created and managed by a client application. An HCP for cloud scale bucket is modeled on a storage service bucket. HCP for cloud scale uses buckets to manage storage components. You can think of an HCP for cloud scale system as a logical collection of secure buckets.

Buckets have associated metadata such as ownership and lifecycle status. HCP for cloud scale buckets are owned by an HCP for cloud scale user and access is controlled on a per-bucket basis by Amazon access control lists (ACL) supporting the S3 API. Buckets are contained in a specific region; HCP for cloud scale supports one region.



Note:

1. HCP for cloud scale buckets are not stored in storage components, so HCP for cloud scale clients can create buckets even before adding storage components.
2. Storage component buckets are created by storage component administrators and are not visible to HCP for cloud scale clients.
3. To empty and reuse a bucket, don't just delete the bucket and create a new one with the same name. After a bucket is deleted, the name becomes available for anyone to use and another account might take it first. Instead, empty and keep the bucket.

Objects

An *object* consists of data and associated metadata. The metadata is a set of name-value pairs that describe the object. Every object is contained in a bucket. An object is handled as a single unit by all HCP for cloud scale transactions, services, and internal processes.

For information about Amazon S3, see [Introduction to Amazon S3](#).

S3 Console application

HCP for cloud scale includes an S3 Console application that provides convenient functions for bucket users as an alternative to using S3 API methods:

- Obtaining S3 credentials
- Managing bucket synchronization, policies, and rules
- Creating S3 event notifications to synchronize buckets
- Managing objects in buckets

For more information, see the *S3 Console Guide*.

Data access

HCP for cloud scale supports the Amazon S3 API, which lets client applications store and retrieve unlimited amounts of data from configured storage services.

Data access control

HCP for cloud scale uses ownership and access control lists (ACLs) as data access control mechanisms for the S3 API.

Ownership is implemented as follows:

- An HCP for cloud scale bucket is owned by the user who creates the bucket and the owner cannot be changed.
- A user has full control of the buckets that user owns.
- A user has full control of the objects that user creates.
- A user can list only the buckets that user owns.

ACLs allow the assignment of privileges (read, write, or full control) for access to buckets and objects to other user accounts besides the owner's.

Data security

HCP for cloud scale supports encryption of data sent between systems (that is, data "in flight") and, as a licensed feature, data stored persistently within the system (that is, data "at rest").

Certificate management

HCP for cloud scale uses Secure Sockets Layer (SSL) to provide security for both incoming and outgoing communications. To enable SSL security, two types of certificates are needed:

- System certificate: the certificate that HCP for cloud scale uses for its GUIs and APIs (for incoming communications)
- Client certificates: the certificates of IdPs, storage components, and SMTP servers (for outgoing communications)

When the HCP for cloud scale system is installed, it generates and automatically installs a self-signed SSL server system certificate. This certificate is not automatically trusted by web browsers. You can choose to trust this self-signed certificate or replace it by using one of these options:

1. Upload a PKCS12 certificate chain and password and apply it as the active system certificate.
2. Download a certificate signing request (CSR) and then use it to obtain, upload, and apply a certificate signed by a certificate authority (CA).
3. Generate a new self-signed certificate and apply it as the active system certificate.

For client certificates, you need to upload the certificate of each client HCP for cloud scale needs to access using SSL.

You can manage certificates, and view details of installed certificates, using the System Management application.

Data-in-flight encryption

HCP for cloud scale supports data-in-flight encryption (HTTPS) for all external communications. Data-in-flight encryption is always enabled for these data paths:

- S3 API (HTTP is also enabled on a different port)
- Management API
- System Management application graphical user interface (GUI)
- Object Storage Management application GUI

You can enable or disable data-in-flight encryption for the data paths between HCP for cloud scale and:

- An identity provider (IdP) server
- Each application using TLS or SSL
- Each managed storage component
- Each SMTP server using SSL or STARTTLS

Communication between HCP for cloud scale instances does not use data-in-flight encryption. Depending on your security needs, you might need to set up an isolated internal network for HCP for cloud scale at your site.

Data-at-rest encryption

HCP for cloud scale stores these kinds of data persistently:

- HCP for cloud scale services data
- HCP for cloud scale metadata and user-defined metadata
- User data (object data)

The first two kinds of data are handled by the hardware on which HCP for cloud scale instances are installed. If needed, you can install HCP for cloud scale on servers with encrypted disks.

Object data is handled by storage components. HCP for cloud scale supports system-wide encryption, using AWS SDK client-side encryption and strong encryption ciphers, as a licensed feature. Encryption and decryption are transparent to users. Each storage component has a separate master key. Storage components that use hardware acceleration for encryption and decryption are supported.

To manage encryption master keys, HCP for cloud scale supports the HashiCorp Vault key management system (KMS) through a KMS client that is automatically deployed as a service when encryption is enabled. After you set up a Vault server, you can enable encryption support on HCP for cloud scale as a global setting and then manage the encryption client service as needed. For information about Vault and how to set up a server, see <https://www.hashicorp.com/products/vault>.



Note: Once enabled, encryption can't be disabled.

As an alternative, you can use individual storage components that support data-at-rest encryption. Storage components can self-manage their keys, or HCP for cloud scale can facilitate keys you supply following the Amazon S3 API specification.

Bucket synchronization

Bucket synchronization to a bucket (*bucket sync-to*) allows automatic, asynchronous copying of objects in buckets in an HCP for cloud scale system to external storage systems. Bucket synchronization from a bucket (*bucket sync-from*) allows automatic, asynchronous copying of objects in buckets in external storage systems to an HCP for cloud scale bucket.

An external storage system can be another HCP for cloud scale system, AWS, or any S3 compatible system.

Bucket sync-to offers the following advantages:

- Data protection: Data is well protected against the unavailability or catastrophic failure of a system. Buckets can be synchronized to multiple remote systems of different types. This arrangement can provide geographically distributed data protection (called *geo-protection*).
- Data availability: AWS services can access synchronized data directly from AWS.

Bucket sync-from offers the following advantages:

- Data consolidation: Transformed data can be stored on an HCP for cloud scale system. An HCP for cloud scale system can synchronize data from multiple remote systems of different types.
- External update: Data can be updated directly in an external system and stored on an HCP for cloud scale system.

Access to bucket synchronization is controlled on a per-user basis by *role-based access control (RBAC)*. Use the System Management application to define users, groups, and roles.

Access to an external resource might need an SSL certificate. You can upload an SSL certificate using the System Management application, the same as for uploading SSL certificates for storage components and IdPs.

For information on bucket synchronizations, see the *S3 Console Guide*.

Object locking

HCP for cloud scale supports object locking, which prevents specified objects from being deleted. A bucket owner can lock or unlock objects or lock them for a specified time period. This feature implements legal hold and retention period requirements.

Object locking is enabled at the bucket level, either when or after a bucket is created. Once enabled, object locking can't be disabled.

Object locking offers the following advantages:

- Locked objects can't be deleted. This implements write once, read many (WORM) behavior, which protects objects from accidental or malicious changes.
- A bucket owner can lock objects until a specified date and time. This implements retention periods, which complies with record retention policy. The retention period can be up to 100 years in the future.



Note: Once set, a retention period can be extended, but not shortened or turned off.

- A bucket owner can lock an object indefinitely, and then turn the lock off. This complies with legal hold requirements. If a legal hold is placed on an object it can't be modified, versioned, moved or deleted, even if it has an expired retention period (that is, a legal hold overrides a retention period). A legal hold never expires, but must instead be removed. An object can have multiple legal holds placed on it.

HCP for cloud scale implements compliance mode as described by the Amazon S3 specification. It does not support governance mode.



Note: Using S3 PUT Object Lock methods in HCP for cloud scale v1.4 and earlier is not supported. Using the methods might return an HTTP status code of 200 but will not produce the expected behavior. Only use S3 object lock methods after updating to v1.5 or later.

For information on how to lock and unlock objects, see the *S3 Console Guide*.

S3 Select

HCP for cloud scale supports the S3 Select feature.

HCP for cloud scale supports the S3 Select Object Content method, which allows retrieval of a portion of a structured object by an S3 client such as Apache Spark, Apache Hive, and Presto. The portion of the object returned is selected based on a structured query language (SQL) query sent in the request. The query is performed by S3 storage components that support pushdown. Selecting only the data needed within an object can significantly improve costs, time, and performance.

A request can select serialized object data in these formats:

- Apache Parquet

A request can return data in these formats:

- Comma-separated values (CSV)

The client application must have `s3:GetObject` permission. S3 Select supports reading of encrypted data. The SQL expression can be up to 256 KB. Returned data can be up to 1 MB.

Here is a simple example of a SQL query against a Parquet object. The query returns data for salaries greater than 100,000:

```
select salary from s3object s where s.salary > 100000
```

S3 event notification

HCP for cloud scale supports the S3 PUT Bucket notification configuration and GET Bucket notification configuration methods.

HCP for cloud scale can send notifications of specified events in a bucket to a message server for applications to consume. This is a more efficient way to signal changes than periodically scanning objects in a bucket.

HCP for cloud scale supports event notification to signal specified events in buckets. Notifications can be sent to SQS Standard services. A retry mechanism assures highly reliable notifications.

Supported limits

HCP for cloud scale limits the number of instances (nodes) in a system to 160.

HCP for cloud scale does not limit the number of the following entities.

Entity	Minimum	Maximum	Notes
Buckets	None	Unlimited	
Users (external)	None	Unlimited	The local user has access to all functions including MAPI calls and S3 API calls. However, it is best to configure HCP for cloud scale with an identity provider (IdP) with users to enforce role-based access control.
Groups (external)		Unlimited	
Roles		Unlimited	
Objects	None	Unlimited	The default size limit for a single PUT or POST object call is 5 GB.
Storage components	1	Unlimited	

High availability

HCP for cloud scale supports high availability for multi-instance sites.

High availability needs at least four service instances: three master instances, which run essential services, and at least one worker instance. The best practice is to run the three master instances on separate physical hardware (or, if running on virtual machines, on at least three separate physical hosts) and to run HCP for cloud scale services on more than one instance.

Scalability of instances, service instances, and storage components

You can increase or decrease the capacity, performance, and availability of HCP for cloud scale by adding or removing the following:

- Instances: physical computer nodes or virtual machines
- Service instances: copies of services running on additional instances
- Storage components: S3 compatible systems used to store object data

In a multi-instance site, you might add instances to improve system performance or if you are running out of storage space on one or more instances. You might remove instances if you are retiring hardware, if an instance is down and cannot be recovered, or if you decide to run fewer instances.

When you add an instance, you can also scale floating services (such as the Metadata Gateway) to the new instance. When you scale a floating service, HCP for cloud scale automatically rebalances itself.

In a multi-instance site, you can manually change where a service instance runs:

- You can configure it to run on additional instances. For example, you can increase the number of S3 Gateway service instances to improve throughput of S3 API transactions.
- You can configure it to run on fewer instances. For example, you can free computational resources on an instance to run other services.
- You can configure it to run on different instances. For example, you can move the service instances off a hardware instance to retire the hardware.
- For a floating service, instead of specifying a specific instance on which it runs, you can specify a pool of eligible instances, any of which can run the service.

Some services have a fixed number of instances and therefore cannot be scaled:

- Metadata Coordination

You might add storage components to a site under these circumstances:

- The existing storage components are running out of available capacity
- The existing storage components do not provide the performance you need
- The existing storage components do not provide the functionality you need

Site availability

An HCP for cloud scale site has three master instances and thus can tolerate the failure of one master instance without interruption of service. HCP for cloud scale services can continue to function if two or even all three master instances fail. However, you cannot move or scale service instances until master instances are restored.

Service availability

HCP for cloud scale services provide high availability as follows:

- The Metadata Gateway service always has at least three service instances. When the system starts, the nodes "elect a leader" using the raft consensus algorithm. The other service instances follow the leader. The leader processes all GET and PUT requests. If the followers cannot identify the leader, they elect a new leader. The Metadata Gateway service tolerates service instance failure, and functions without loss of data, as long as at least two service instances are healthy.
- The Metadata Coordination service always has one service instance. If that instance fails, HCP for cloud scale automatically starts another instance. Until startup is complete, the Metadata Gateway service cannot scale.
- The Metadata Cache service always has one service instance. If that instance fails, HCP for cloud scale automatically starts another instance. Until startup is complete, overall performance decreases.
- To protect messaging consistency, the Message Queue service always has three service instances. To prevent being split into disconnected parts, the service shuts down if half of the service instances fail. In practice, messaging stops if two of the three instances fail. Do not let the service run with only two instances, because in that scenario if one of the remaining instances fails, the service shuts down. However, when one of the failed instances restarts, messaging services recover and resume.
- To maintain access to the encryption key vault, the Key Management Server service uses an active-standby model. One service instance is the active instance and any other service instances are kept as standbys. If the active vault node becomes sealed or unavailable, one of the standbys takes over as active. You can scale up to the number of instances in the HCP for cloud scale system or your acceptable performance limits.

The rest of the HCP for cloud scale services remain available if HCP for cloud scale instances or service instances fail, as long as at least one service instance remains healthy. Even if a service that has only one service instance fails, HCP for cloud scale automatically starts a new service instance.

Metadata availability

Metadata is available as long as these services are available:

- S3 Gateway
- Metadata Gateway

Object data availability

Object data is available as long as these items are available:

- The S3 Gateway service (at least one instance)
- The storage component containing the requested object data
- At least two functioning Metadata Gateway service instances (of the required three)

For high availability of object data or data protection, you should use a storage component with high availability, such as HCP, HCP S Series Node, or AWS S3.

Network availability

You can install each HCP for cloud scale instance with both an internal and an external network interface. To avoid single points of networking failure, you can:

- Configure two external network interfaces in each HCP for cloud scale instance
- Use two switches and connect each network interface to one of them
- Bind the two network interfaces into one virtual network interface in an active-passive configuration
- Install HCP for cloud scale using the virtual network interface

Failure recovery

HCP for cloud scale actively monitors the health and performance of the system and its resources, gives real-time visual health representations, issues alert messages when needed, and automatically takes action to recover from the failure of:

- Instances (nodes)
- Product services (software processes)
- System services (software processes)
- Storage components

Instance failure recovery

If an instance (a compute node) fails, HCP for cloud scale automatically adds new service instances to other available instances (compute nodes) to maintain the minimum number of service instances. Data on the failed instance is not lost and remains consistent. However, while the instance is down, data redundancy might degrade.

HCP for cloud scale adds new service instances automatically only for floating services. Depending on the remaining number of instances and service instances running, you might need to add new service instances or deploy a new instance.

Service failure recovery

HCP for cloud scale monitors service instances and automatically restarts them if they are not healthy.

For floating services, you can configure a pool of eligible HCP for cloud scale instances and the number of service instances that should be running at any time. You can also set the minimum and maximum number of instances running each service. If a service instance failure causes the number of service instances to go below the minimum, HCP for cloud scale starts another service instance on one of the HCP for cloud scale instances in the pool that doesn't already have that service instance running.

Persistent services run on the specific instances that you specify. If a persistent service fails, HCP for cloud scale restarts the service instance in the same HCP for cloud scale instance. HCP for cloud scale does not automatically bring up a new service instance on a different HCP for cloud scale instance.

Storage component failure recovery

HCP for cloud scale performs regular health checks to detect storage component failures.

If HCP for cloud scale detects a storage component failure, it sets the storage component state to INACCESSIBLE, so that HCP for cloud scale will not try to write new objects to the storage component, and sends an alert. While a storage component is unavailable, the data in it is not accessible.

HCP for cloud scale continues to check a failed storage component and, when it detects that the storage component is healthy again, automatically sets its state to ACTIVE. HCP for cloud scale sends an alert when this event happens as well. Once the storage component is repaired and brought back online, the data it contains is again accessible and HCP for cloud scale can write new objects to it.

Support for the Amazon S3 API

HCP for cloud scale is compatible with the Amazon Simple Storage Service (Amazon S3) REST API, which allows clients to store objects in containers called buckets. A bucket is a collection of objects that has its own settings, such as ownership and lifecycle. Using HCP for cloud scale, users can perform common reads and writes on objects and buckets and manage ACL settings through the client access data service.

For information about using Amazon S3, see the [Amazon S3 API documentation](#).

For information about obtaining S3 user credentials, see the *S3 Console Guide*.

The following tables list the supported Amazon S3 API features and describe any implementation differences between the Amazon and HCP for cloud scale S3 APIs.

Authentication and addressing

Feature	Implementation differences
Authentication with AWS Signature Version 4	Fully implemented.
Addressing virtual host (such as <code>http://bucket.server/object</code>)	Fully implemented.
Addressing Path style (such as <code>http://server/bucket/object</code>)	Fully implemented.
Signed/unsigned payload	Fully implemented.
Chunked request	Fully implemented.

Feature	Implementation differences
Presigned URL	Fully implemented.

Service

Feature	Implementation differences
LIST buckets (GET Service)	Fully implemented.

Buckets

Feature	Implementation differences
GET Bucket (list objects) V1	Fully implemented.
GET Bucket (list objects) V2	Fully implemented.
PUT Bucket	To support legacy S3 buckets, HCP for cloud scale supports bucket names of less than three characters. When anonymous requests to create or delete a bucket use a bucket name that isn't valid, Amazon S3 verifies access first and returns 403. HCP for cloud scale returns 400 if the bucket name validation fails.
DELETE Bucket	
HEAD Bucket	
PUT Bucket ACL	In Amazon S3 each grantee is specified as a type-value pair, where the type is one of the following: <ul style="list-style-type: none"> <code>emailAddress</code> if the value specified is the email address of an AWS account <code>id</code> if the value specified is the canonical user ID of an AWS account <code>uri</code> if granting permission to a predefined group HCP for cloud scale does not support <code>emailAddress</code> . HCP for cloud scale fully supports <code>id</code> . HCP for cloud scale supports <code>uri</code> for the predefined groups <code>Authenticated Users</code> and <code>All Users</code> . HCP for cloud scale does not support the <code>aws-exec-read</code> canned ACL. HCP for cloud scale does not mirror or mirror back ACLs or policies.
GET Bucket ACL	
List Multipart Uploads	Fully implemented.

Feature	Implementation differences
GET Bucket Lifecycle (except transition action)	HCP for cloud scale supports the latest API for bucket lifecycle management. Old and deprecated V1.0 methods are not supported. HCP for cloud scale does not support Object Transition actions. Including these actions causes a Malformed XML exception.
PUT Bucket Lifecycle (except transition action)	
DELETE Bucket Lifecycle (except transition action)	
PUT Bucket Notification Configuration	A configuration can have up to 100 rules. Amazon S3 considers that two rules overlap if both apply to the same object and share at least one event type. HCP for cloud scale supports notification from the same object to multiple targets. However, rules are blocked if they send a message for the same event to the same target. All notification message fields are returned except <code>Region</code> and <code>Glacier Storage</code> . The field <code>awsRegion</code> is returned but left empty.
PUT Bucket Replication	Amazon S3 allows only one-to-one replication. HCP for cloud scale supports one-to-many mirroring and many-to-one mirroring back. The bucket Amazon Resource Name (ARN) is replaced by configuration settings. For mirroring back, HCP for cloud scale supports one queue server, <code>AMAZON_SQS</code> . Sending encrypted data to a remote bucket is not supported.
GET Bucket Versioning	Version listing requests do not strictly comply with documented behavior for <code>NextKeyMarker/NextVersionIdMarker</code> . Amazon S3 documentation currently states that these values specify "the first key not returned that satisfies the search criteria." However, HCP for cloud scale specifies the last key returned in the current response. S3 V1 object listings do not call out as specific a requirement and V2 object listings use a continuation token that is opaque to the caller. Internally, HCP for cloud scale shares the same listing logic across all listing types.
GET Bucket Object Versions	Fully implemented.
GET Bucket Location	You must be the bucket owner. The only supported location is <code>us-west-2</code> .
GET Bucket Notification Configuration	Fully implemented.

Object

Feature	Implementation differences
GET Object	<p>If a lifecycle policy is configured for a bucket, HCP for cloud scale displays the expiration date of an object (in the x-amz-expiration header) fetched using the subresource <code>?versionId</code>. Amazon displays this only when performing unversioned GET requests.</p> <p>Legal hold is fully implemented.</p> <p>Object retention is fully implemented.</p> <p>Object names can't contain NUL or backslash (\) characters. GET operations on objects so named fail with a 400 error.</p>
HEAD Object	<p>If a lifecycle policy is configured for a bucket, HCP for cloud scale displays the expiration date of an object (in the x-amz-expiration header) fetched using the subresource <code>?versionId</code>. Amazon only displays this when performing unversioned HEAD requests.</p>
PUT Object	<p>Amazon S3 limits the maximum file size for a single PUT or POST object call to 5 GB. In HCP for cloud scale, this value is configurable and the default is 5 GB.</p> <p>Amazon S3 is extremely liberal in what is accepted for the Content-Type of an object. HCP for cloud scale adds additional content-type validation.</p> <p>Bucket synchronization is supported.</p> <p>Legal hold is fully implemented. The bucket owner can set a legal hold without restriction; that is, AWS object lock permissions are not supported.</p> <p>Object retention is implemented, but not governance mode; that is, a retain-until date can be set but not changed. The bucket owner can set object retention without restriction; that is, AWS object lock permissions are not supported.</p> <p>Object locking can be applied to a bucket even after it's created. To enable object locking, in the S3 API <code>PUT Bucket ObjectLockConfiguration</code>, include the URI request parameter <code>x-amz-bucket-object-lock-token</code> (with any string).</p> <p>Object names can't contain NUL or backslash (\) characters. PUT operations on objects so named fail with a 400 error.</p>
PUT Object (Copy)	<p>Conditional headers are not supported. Server-side encryption is not supported. Multiple AWS regions are not supported; as a result, cross-region limitations are not supported.</p>
PUT Object (Part Copy)	<p>Conditional headers are not supported. Server-side encryption is not supported.</p>

Feature	Implementation differences
Object and version encoding	<p>Amazon S3 object and version listing documentation mentions the ability to pass an encoding parameter so the object name in the response XML to the client can be escaped to avoid names containing XML characters that aren't valid. Encoding is documented only as it applies to object names, not Owner/DisplayNames. Also, escaping for Bucket Listing requests isn't mentioned.</p> <p>The Owner/DisplayName is a concern because user display names might contain characters that can cause XML parsing issues. Amazon might be able to restrict display names, though it does not currently return a display name for all regions. HCP for cloud scale uses IdPs and thus doesn't control restriction.</p> <p>Bucket name restrictions should prevent problematic bucket names from being created. For security, HCP for cloud scale passes the user display name through a URI encoder before returning the name in an XML response.</p>
Object tagging	<p>Amazon S3 wraps eTags in double quotes. For XML listings (v1 object, v2 object, version) double quotes are escaped. For example:</p> <pre><ETag>&quot;32c81604d07395b1aa39a7e206c3af06\$&quot;</ETag></pre> <p>HCP for cloud scale doesn't need to do this because only attributes, not double quotes, need to be escaped within content.</p>
	<p>Expiration date URL encoding (x-amz-expiration header)</p> <p>HCP for cloud scale URL-encodes the <code>RuleID</code> portion of the x-amz-expiration header using the same encoding strategy that Amazon suggests for V4 authentication. This strategy can result in encoded strings that do not exactly match how Amazon encodes RuleIDs. However, decoding them should always return the original strings.</p>
	<p>HCP for cloud scale mirrors and mirrors back object tagging and tag updates.</p>
GET Object ACL	Bucket synchronization is not supported.

Feature	Implementation differences
PUT Object ACL	<p>In Amazon S3, each grantee is specified as a type-value pair, where the type is one of the following:</p> <ul style="list-style-type: none"> ▪ <code>emailAddress</code> if the value specified is the email address of an AWS account ▪ <code>id</code> if the value specified is the canonical user ID of an AWS account ▪ <code>uri</code> if granting permission to a predefined group <p>HCP for cloud scale does not support <code>emailAddress</code>. HCP for cloud scale fully supports <code>id</code>. HCP for cloud scale supports <code>uri</code> for the predefined groups <code>Authenticated Users</code> and <code>All Users</code>.</p> <p>HCP for cloud scale does not support the <code>aws-exec-read</code> canned ACL.</p>
DELETE Object	<p>Bucket synchronization of deletion of an object or a specific version of an object is not supported.</p> <p>To improve performance, if the current version of an object is a delete marker, HCP for cloud scale doesn't create another delete marker.</p>
DELETE Multiple Objects	<p>Fully implemented.</p> <p>Bucket synchronization is not supported.</p>
POST Object	<p>Fully implemented.</p> <p>Amazon S3 limits the maximum file size for a single PUT or POST object call to 5 GB. In HCP for cloud scale, this value is configurable and the default is 5 GB.</p> <p>Bucket synchronization is supported.</p>
POST Select Object Content	<p>Scan range is supported.</p> <p>HCP for cloud scale supports the use of <code>*</code> by itself with no alias reference. For example, this SQL query is supported:</p> <pre>select *, first_name from s3object s where s.salary > 100000 limit 10</pre> <p>HCP for cloud scale supports a wider range of date-time formats than AWS. The full list is available at https://docs.oracle.com/javase/8/docs/api/java/time/format/DateTimeFormatter.html.</p> <p>HCP for cloud scale supports nested aggregate functions. For example, this expression is supported: <code>count (sum (s.salary))</code></p>

Feature	Implementation differences
	<p>HCP for cloud scale SQL queries on columns are case sensitive, while AWS SQL queries are case insensitive. For example, given an object <code>s</code> with the columns <code>ID</code>, <code>iD</code>, and <code>id</code>, an SQL query to <code>select s.id</code> will return column <code>id</code> in HCP for cloud scale but column <code>ID</code> in AWS.</p> <p>Only input serialization of Parquet is supported. Requests for CSV or JSON objects are not supported and return an error.</p> <p>Parquet compression is managed automatically, so the <code>CompressionType</code> argument is not needed, and if specified returns an error.</p> <p>Only CSV output is supported. Specifying another output format returns an error.</p>
Initiate/Complete/Abort Multipart Upload	<p>Fully implemented.</p> <p>Bucket synchronization is supported.</p>
Upload Part	Fully implemented.
List Multipart Uploads	Fully implemented.

Unsupported S3 API methods

HCP for cloud scale does not support the following Amazon S3 API methods:

Authentication

- Authentication v2 (deprecated by AWS)

Bucket

- GET/PUT/DELETE Bucket website
- GET/PUT/DELETE Bucket policy
- GET/PUT/DELETE Bucket tagging
- GET/PUT/DELETE Bucket CORS (cross-origin resource sharing)
- PUT Bucket versioning (with HCP for cloud scale versioning is always on)
- GET/PUT Bucket logging
- GET Bucket notification
- GET/PUT Bucket requestPayment
- GET/PUT/DELETE Bucket Inventory
- List Bucket Inventory Configurations
- GET/DELETE Bucket metrics
- List Bucket Metrics Configurations
- GET/PUT/DELETE Bucket analytics

- List Bucket Analytics Configurations
- PUT/GET Bucket accelerate
- Server-side encryption with customer-provided encryption keys (SSE-C)
- Server-side encryption with storage-managed encryption keys (SSE-S3)

Object

- Options object
- GET/POST Object torrent

HCP for cloud scale APIs

The Hitachi Content Platform for cloud scale (HCP for cloud scale) system includes RESTful application programming interfaces (APIs) that you can use for writing applications that exercise its functions and manage the system.

Anything you can do in the Object Storage Management, S3 User Credentials, or System Management application GUIs you can also do using APIs.

Object Storage Management API

The Object Storage Management application includes a RESTful API to administrative functions such as managing storage components, configuring Amazon S3 settings, and obtaining or revoking S3 user credentials. For more information on the Object Storage Management API, see the *MAPI Reference*.

System Management API

The System Management application includes a RESTful API to system management functions such as system monitoring, service monitoring, user registration, and configuration. For more information on the System Management API, see the online help in the System Management application.

Amazon S3 API

Unless otherwise noted, HCP for cloud scale is fully compatible with the Amazon S3 API.

Object Storage Management API

The Object Storage Management application includes a RESTful API interface for the following functions:

- Managing storage components and Amazon S3 settings
- Managing administrative resources such as serial numbers and system events
- Managing user resources such as S3 user credentials and OAuth tokens

The Object Storage Management API is served by the MAPI Gateway service from any HCP for cloud scale node.

You can execute all functions supported in the Object Storage Management application using the API.



Note: The system configuration, management, and monitoring functions included in the System Management application can be performed using the System Management API.

All URLs for the API have the following base, or root, uniform resource identifier (URI):

`https://hcpcs_ip_address:9099/mapi/v1`

System Management API

The System Management application provides a RESTful API for managing the following:

- Alerts
- Business objects
- Certificates
- Events
- Instances
- Jobs
- Licenses
- Notifications
- Packages
- Plugins
- Security
- Services
- Setup
- Tasks
- Updates

You can execute all functions supported in the System Management application using the API.

Security and authentication

HCP for cloud scale controls access to system functions through user accounts, roles, permissions, and OAuth tokens, where user accounts are stored in an external identity provider. HCP for cloud scale controls access to data by S3 API requests through S3 credentials, ownership, and access control lists. HCP for cloud scale supports in-flight encryption (HTTPS) for all external communications.

User accounts

The initial user account, which has all permissions, is created when you install HCP for cloud scale. The initial user account can perform all HCP for cloud scale functions. After the initial user account is created, you can change its password any time, but you cannot disable the account and you cannot change its permissions.

The initial user is the only local account allowed and is intended only to let you configure an identity provider (IdP). HCP for cloud scale can communicate with IdPs using HTTP or HTTPS. HCP for cloud scale supports multiple IdPs:

- Active Directory
- OpenLDAP
- 389 Directory Server
- LDAP compatible

HCP for cloud scale supports external users defined in the IdP. External users with the appropriate permissions can perform some or all of these functions:

- Log in to the Object Storage Management application and use all functions
- Log in to the System Management application and use all functions
- Get an OAuth token to use all API calls for the Object Storage Management and System Management applications
- Log in to the S3 User Credentials application and get S3 credentials to use the S3 API

HCP for cloud scale discovers the groups in each IdP and allows assigning roles to groups.

HCP for cloud scale uses OAuth2 as a service provider to authenticate single sign-on (SSO) access. SSO lets you use one set of login credentials for all HCP for cloud scale applications, so you can switch between applications without logging in again.

API access

Object Storage Management application API methods need a valid OAuth access token for a user account with suitable permissions, or else the requests are rejected. With one exception, System Management application API methods also require a valid OAuth access token for a user account with suitable permissions, or else the requests are rejected. (The exception is the API method to generate an OAuth token, which requires only a username and password in the body of the request.)

Before using either the Object Storage Management or System Management APIs, you need to obtain an OAuth token. You can generate an OAuth token by sending a request to the OAuth server with your account credentials. Then you can supply the OAuth token in the Authorization header in each request. OAuth tokens are valid for five hours.



Note: An administrator can revoke all OAuth tokens for any other HCP for cloud scale user. You would do this, for example, if an employee leaves the company, you delete the user account, and you do not want to wait for the account tokens to expire.

S3 API requests generally require valid S3 credentials for users with the right privileges, that is, access control lists (ACLs). (Exceptions are methods configured to allow anonymous access and pre-signed requests.) HCP for cloud scale supports AWS Signature version 4 authentication to include S3 credentials in S3 requests.

Users with a valid account and suitable permissions can generate S3 credentials. You can generate an unlimited number of S3 credentials, but only the last credentials generated are valid. These credentials are associated only with your account. S3 credentials do not have an expiration date, so they are valid until revoked.

Users with a valid account and suitable permissions can revoke all S3 credentials of any user. That is, you can revoke your own S3 credentials or the S3 credentials of any other user. Revocation removes all S3 credentials associated with the account.



Note: Deleting a user account from the IdP does not revoke S3 credentials, and if a user's S3 credentials are revoked the user can still generate new credentials. The best practice is to delete the user account from the IdP and then revoke the S3 credentials.

Network isolation and port mapping

When you install HCP for cloud scale, you can set up network isolation by configuring one external network and one internal network.

HCP for cloud scale software creates a cluster using commodity x86 servers that are networked using Ethernet. The software uses two networks on the operating system hosting the HCP for cloud scale software. These networks can also use link aggregation defined by the OS administrator.

While two networks provide optimal traffic isolation, you can deploy the software using a single network. The OS administrator must make and implement networking decisions before you install HCP for cloud scale.

HCP for cloud scale services use a range of network ports. You can configure services to use different ports instead of the default ports. Installation is the only opportunity to change the default ports used by services.



Note: The following services must be deployed with their default port values:

- Message Queue
- Tracing Agent
- Tracing Collector
- Tracing Query

For information about installing HCP for cloud scale, see *Installing Hitachi Content Platform for Cloud Scale*.

Logging in

User accounts reside in an external identity provider (IdP). To log in you need this information:

- The IP address of the HCP for cloud scale instance that you're using
- Your user name as assigned by your system administrator
- Your password as assigned by your system administrator
- The security realm where your user account is defined

Procedure

1. Open a web browser and go to `https://instance_ip_address:8000`
instance_ip_address is the IP address of the HCP for cloud scale instance you're using
2. Type your username and password.
3. In the **Security Realm** field, select the location where your user account is defined.
To log in using the local administrator account, without using an external IdP, select **Local**. If no IdP is configured yet, **Local** is the only available option.
4. Click **LOGIN**.

Result

The Applications page opens.

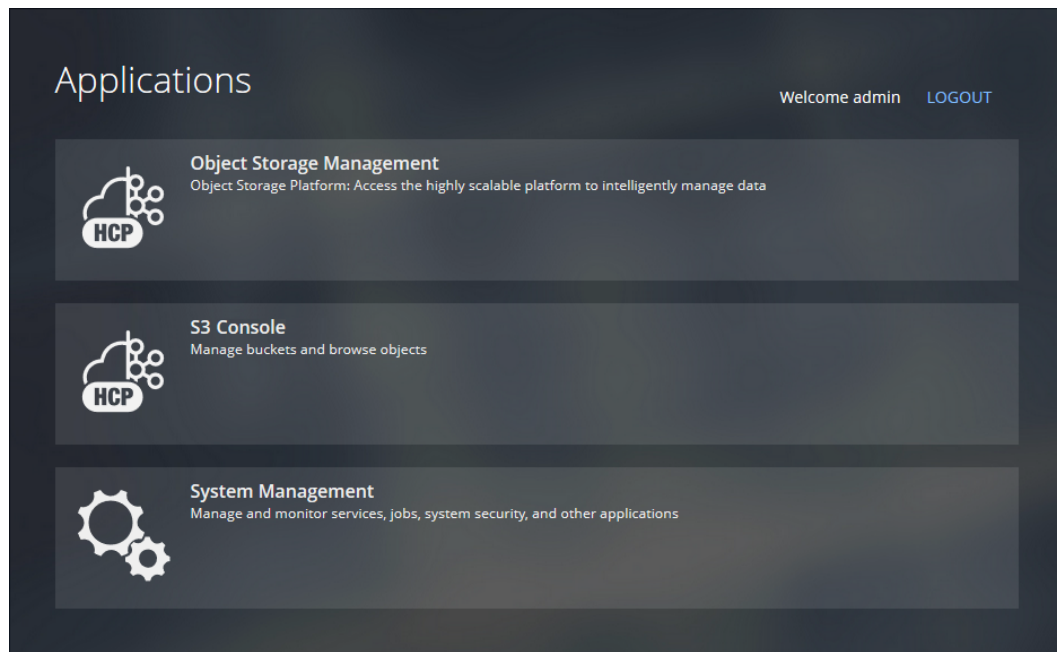


Note: When a new user is created and added to a group, that user might not have immediate access to HCP for cloud scale. Instead, login fails with the message "Not authorized. Please contact your system administrator." Verify the credentials. If the condition persists, the system administrator can use the API method `security/clearCache` to allow immediate login.

HCP for cloud scale applications

After you log in, the HCP for cloud scale **Applications** page shows you the applications you are authorized to use, such as:

- Object Storage Management: Manage and monitor storage components, data objects, alerts, and regions
- S3 Console: Generate S3 access and secret keys; conveniently create and manage buckets, bucket synchronization, and bucket policies; manage S3 event notification; and browse objects in buckets
- System Management (sometimes referred to in the application as the Admin App): Manage and monitor cluster instances, software services, system security, user accounts, and other cluster configuration parameters




From the **Applications** page, or from within each application, you can switch back and forth between applications as needed.

Switching between applications

HCP for cloud scale uses OAuth2 as a service provider to authenticate single sign-on (SSO) access. You only need one set of login credentials for all HCP for cloud scale applications, so you can switch between applications without logging in again.

Depending on the permissions assigned to your account role, you can have access to one or more HCP for cloud scale application. To switch between applications:

Procedure

1. Click the Open menu () in the right corner of the top navigation bar, and select the application you want to use.



Note: The **System Management** application is also identified in the user interface as **Admin-App**.

The application opens.

Providing unseal keys to KMS service

When encryption is enabled for a HCP for cloud scale system, the Key Management System service provides encryption keys for storage components. If the service restarts, the key repository vault closes and data objects can't be decrypted. If a vault instance becomes sealed, you must provide a quorum of unseal keys (three of the five provided when encryption was first enabled) to reopen the vault and resume encryption and decryption.



Caution: Don't try to initialize the vault manually outside of HCP for cloud scale. Doing so results in data loss.

To unseal the vault:

Procedure

1. Select **Global Settings**.
The **Global Settings** page opens.
2. Click **Unseal Vault**.
The **Unseal Vault Instances** window opens.
3. Enter the first unseal key (a master key portion) by pasting or typing it into the field **Unseal Key 1**.
The key is validated. You can't leave the field blank.
4. Enter a second unseal key into the field **Unseal Key 2**.
The key is validated. You can't leave the field blank. Each key must be different.
5. Enter a third unseal key into the field **Unseal Key 3**.
The key is validated. You can't leave the field blank. Each key must be different.
6. Click **Unseal**.

Result

The vault is unsealed.

Serial number

You can use the Object Storage Management application or API to enter and display your HCP for cloud scale serial number.

A serial number is required to activate the HCP for cloud scale software. You must enter the serial number before you can use the system or its licensed features.

Entering your serial number

The Object Storage Management application displays the product serial number. An administrative account with appropriate permissions can enter or edit this number.

Object Storage Management application instructions

To enter your product serial number:

Procedure

1. Select **Global Settings** and then click the Edit icon next to the **Serial Number** field.
The **Add Serial Number** window opens.
2. Type your serial number and then click **Add**.

Related REST API methods

```
POST /serial_number/set
```

For information about specific API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

Displaying your serial number

You can use the Object Storage Management application or API to displays the product serial number.

Object Storage Management application instructions

The product serial number is displayed in the upper right corner of the **Dashboard** page.

Related REST API methods

```
POST /serial_number/get
```

For information about specific API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

License

You can use the Object Storage Management application or an API method to enter and validate your HCP for cloud scale license.

A license is required before you can activate certain HCP for cloud scale features. You must enter your serial number before you can upload your license.

Uploading your license

The Object Storage Management application displays your product licenses. An administrative account with appropriate permissions can upload a license file.

Object Storage Management application instructions

To upload a product license file:

Procedure

1. Select **Global Settings** and then click **Upload** in the **License** field.
The **Upload License** area opens, displaying the status of your licenses.
2. Do one of the following:
 - Drag a license file into the **Upload License** area.
 - Click **Choose File**, select a license file, and then click **Open**.

3. After you've selected the license file click **Submit**.

Result

The license file is decrypted and validated. If the license is valid you see the message, "License is valid." If the license is invalid or expired you see an error message.

Related REST API methods

```
POST /license/list
```

For information about specific API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

Enabling encryption

Encryption is an available licensed feature. You must obtain and upload a license to enable encryption.

An administrative account with appropriate permissions can enable encryption. All objects added to the system after encryption is enabled are encrypted on all storage components.



Note: Encryption is a global setting. Once enabled, you can't turn off encryption or unencrypt either storage components or the objects stored on them.

Before you enable encryption, set up a Vault server. Enabling encryption generates encryption keys, an initial root token, and a set of five unseal keys. When establishing a connection to the Vault server, HCP for Cloud Scale provides the initial root token for authentication and root access.

Vault doesn't store the generated master key. Instead, each time the Vault server starts it uses the unseal keys to regenerate the master key, which is then used to return storage component encryption keys. If the Vault server goes down, it seals the vault, and to regenerate the master key you must provide a quorum of at least three valid unseal keys.



Caution: If you don't provide a quorum of unseal keys to reconstruct the master key, Vault remains sealed, so the master key is unavailable and encrypted storage components can't be decrypted. To ensure encryption security, the best practice is to encrypt and store unseal keys separately.

You can enable encryption using the Object Storage Management application or a management API method.

Object Storage Management application instructions

To enable encryption:

Procedure

1. Select **Global Settings**.
The **Global Settings** page opens. The page displays information about the key management server connection.
2. In the **Encryption** section of the page, select **Encryption**.
You are prompted that turning on encryption is permanent and can't be disabled.
3. Click **Proceed**.
If the key management server is not yet available you receive an error message. Try again when the server is available.
The **Vault Unsealing Information** window opens, displaying five unseal keys and your initial root token.
4. Click **Copy** for each unseal key and save the keys elsewhere.
This window is the only time that all of this data is ever known by Vault and also the only time that the unseal keys should ever appear together. To minimize the possibility of multiple keys becoming unavailable, the best practice is securely distribute, encrypt, and store the unseal keys in as many different locations.
5. Click **Copy** for the initial root token and save it elsewhere.
6. Click **Close**.
You are warned that you won't have another opportunity to record the unseal keys and the initial root token.
7. Click **Okay**.
The Key Management Server service is initialized, a connection to the KMS server is established, and storage component encryption keys are generated and applied.

Result

You have enabled encryption.

Related REST API methods

```
POST /s3_encryption/set
```

For information about specific API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

Defining subdomain for S3 Console application

The S3 Console application uses a subdomain of the HCP for cloud scale system.

The S3 Console application uses a subdomain within the HCP for cloud scale system, such as `s3.hcpcs.Company.com`. For user convenience, you can modify the `hosts` file on systems used to call the S3 Console application.

Procedure

1. On a system that calls the S3 Console application, open the `hosts` file in an editor.
On a Windows system, the hosts file is normally located at `C:\Windows\System32\Drivers\etc\hosts`. On a Linux system, the hosts file is normally located at `/etc/hosts`.
2. Associate the IP address of the HCP for cloud scale system with the S3 subdomain.
`10.24.19.54 s3.hcps.Company.com`
3. Save the file.
4. Repeat Steps 1-3 for every system used to call the S3 Console application.

Chapter 2: Storage components

Hitachi Content Platform for cloud scale (HCP for cloud scale) lets you manage and monitor storage components.

Your starting point is the **Storage Component** page of the Object Storage Management application.

Adding a storage component

You can use the Object Storage Management application or an API method to add a storage component to the HCP for cloud scale system.



Tip: To improve performance and availability, and to avoid transfer fees, add storage components that are local to the HCP for cloud scale site.

The storage component must contain an HCP for cloud scale bucket before you can add the storage component to the HCP for cloud scale system.

To add a storage component, it must be available and you need the following information about it:

- Storage component type
- Endpoint information (host name or IP address)
- If used, the proxy host and port and the proxy user name and password
- API port
- S3 credentials (the access key and secret key to use for access to the storage component bucket)

Object Storage Management application instructions

The Add Storage Component wizard helps you add a storage component.


The storage component must contain an HCP for cloud scale bucket.

To add a storage component:

Procedure

1. From the **Storage Component** page, click **Add Storage Component**.
The **Add Storage Component** wizard opens. The first page describes the process and the information needed.
2. Click **Start**.
The **Connection** page opens.

3. Specify the following:
 - a. **Storage Component Name** (optional): Type the display name you choose for the storage component, up to 1,024 alphanumeric characters.
If you leave this field blank, the storage component is listed without a name.
 - b. **Storage Type**: Select **AMAZON_S3**, **HCP_S3**, **HCPS_S3**, or **GENERIC_S3**.
 - c. **Region** (optional): Type a region name of up to 1,024 characters.
HCP for cloud scale doesn't validate this field except for its length.
 - d. **Host**: Type either the IP address or the host name of the storage component.
Type as many as 255 URI unreserved characters A-Z, a-z, 0-9, hyphen (-), period (.), underscore (_), and tilde (~) only.
4. Click **Next**.
The **Connection Advanced** page opens.
5. Specify the following:
 - a. Select the protocol used, either **HTTPS** (the default) or **HTTP**.
 - b. If **Use Default** is selected, the applicable default port number is filled in. If you cancel the selection **Use Default**, type the **Port** number.
 - c. If you select **Proxy**, type values in the fields **Proxy Host** (following the same naming rules as **Host**) and **Proxy Port**, and if the proxy needs authentication, type the **Proxy User Name** and **Proxy Password**.
6. Click **Next**.
The **Activation** page opens.
7. Specify the following:
 - a. **Bucket Name**: Type the name of the bucket on the storage component. The name can be from 3 to 63 characters long and must contain only lowercase characters (a-z), numbers (0-9), periods (.), or hyphens (-).

 **Note:** The bucket must already exist on the storage component.

 - b. (Optional) To use path-style URLs to access buckets, select **Use Path Style Always**.
 - c. **Authenticate**: Select the AWS Signature version, either **V2** or **V4**.
 - d. Type your **Access Key**.
 - e. Type your **Secret Key**.
8. Click **Next**.
The **Review** page opens.
9. Review the configuration of the storage component.
 - If the information is not correct, click **Back** to return to the wizard page and correct the information.
 - If the information is correct, click **Create**.

Result

The storage component is added, the Storage Component page is updated, and the storage component is added to the list.



Note: After you add the storage component, if its state is UNVERIFIED, edit the definition parameters.

If verification of the storage component fails, this error message appears:

Error activating the storage component. Please check and update your configuration before trying again.

The message gives additional details that you can use to troubleshoot the problem.

Related REST API methods

POST /storage_component/create



Note: After you define the storage component, if its state is UNVERIFIED, check the parameters you used when defining it.

For information about specific API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

Modifying a storage component

You can use the Object Storage Management application or an API method to modify a storage component.

Object Storage Management application instructions

You can modify the configuration of a storage component.

Procedure

1. From the **Storage Component** page, click the **Edit Component** icon by the storage component you want to modify.
The **Edit Storage Component** wizard opens.
2. Edit connection information as needed. When you're finished click **Next**.
The **Connection Advanced** page opens.
3. Edit advanced connection information as needed. When you're finished click **Next**.
The **Activation** page opens.
4. Edit activation information as needed. When you're finished click **Next**.
The **Review** page opens.
5. Review the edited configuration of the storage component.
 - If the information is not correct, click **Back** to return to the wizard page with the information to correct.
 - If the information is correct, click **Create**.

Related REST API methods

```
POST /storage_component/update
```

For information about specific API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

Activating a storage component

You can use the Object Storage Management application or an API method to activate a storage component.

A storage component is displayed as UNVERIFIED if HCP for cloud scale cannot reach the storage component with the supplied parameters.

Object Storage Management application instructions

You can activate a storage container that is in the state INACTIVE.

To activate a storage component:

Procedure

1. Select **Storage**.
The **Storage Component** page opens.
2. For the storage component you want to activate, click **Activate Now**.
The storage component state changes to **ACTIVE**.

Related REST API methods

```
POST /storage_component/update_state
```

For information about specific API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

Deactivating a storage component

You can use the Object Storage Management application or an API method to deactivate a storage component.

You might deactivate a storage component for maintenance purposes.

After you mark a storage component as INACTIVE, read, write, and healthcheck requests are rejected.

Object Storage Management application instructions

You can deactivate a storage container that is in the state **ACTIVE**.

To deactivate a storage component:

Procedure

1. Select **Storage**.
The **Storage Component** page opens.
2. For the storage component you want to deactivate, click **Yes, Inactivate**.
The storage component state changes to **INACTIVE**.

Related REST API methods

```
POST /storage_component/update_state
```

For information about specific API methods, in the Object Storage Management application, click the profile icon and select REST API.

Marking a storage component as read-only

You can use the Object Storage Management application or API calls to mark a storage component as read-only.

You might mark a storage component as read-only if it is nearly full.

Once you mark a storage component as read-only, write requests are directed to different storage components.

Object Storage Management application instructions

To mark a storage component as read-only:

Procedure

1. From the **Storage Component** page, select the **Read-Only** box for the storage component you want to mark.
A window opens, displaying details about the storage component endpoint and showing you the buckets affected by the change, and prompts you to confirm your action.
2. Click **Mark Read-only**.
The storage component is marked as read-only.

Related REST API methods

```
PATCH /storage_component/update
POST /storage_component/update_state
```

For information about specific API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

Marking a storage component as read-write

You can use the Object Storage Management application or API calls to mark a storage component as read-write.

This makes the storage component available for writing new objects.

Object Storage Management application instructions

To mark a read-only storage component as read-write:

Procedure

1. From the **Storage Component** page, clear the **Read-Only** box for the storage component you want to mark.
A window opens, displaying details about the storage component endpoint and showing you the buckets affected by the change, and prompts you to confirm your action.
2. Click **Open for Writes**.
The storage component is marked as read-write.

Related REST API methods

```
PATCH /storage_component/update
POST /storage_component/update_state
```

For information about specific API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

Viewing storage components

You can use the Object Storage Management application or an API method to list the storage components defined in the system.

For each storage component, the list gives its name, type, region, and state.

The storage component types are:

- AMAZON_S3: An Amazon Web Services S3 compatible node
- HCP_S3: A Hitachi Content Platform node
- HCPS_S3: An HCP S Series node
- GENERIC_S3: An S3 compatible node

The possible storage component states are:

- **ACTIVE:** Available to serve requests
- **INACTIVE:** Not available to serve requests (access is administratively paused)
- **INACCESSIBLE:** Available to serve requests, but HCP for cloud scale is having access issues (for example, network, authentication, or certificate issues)
- **UNVERIFIED:** Not available to serve requests (unreachable by specified parameters, or awaiting administrative activation)

Object Storage Management application instructions

The storage components defined in the HCP for cloud scale system are listed in the Storage Components section of the Storage Components page.

Related REST API methods

```
POST /storage_component/list
```

For information about specific API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

Displaying storage component analytics

The **Storage Component** page includes an Analytics section that displays counts of active, inactive, and unverified storage components.

The states displayed are:

- **ACTIVE:** Available to serve requests
- **INACTIVE:** Not available to serve requests (access is administratively paused)
- **UNVERIFIED:** Not available to serve requests (unreachable by specified parameters, or awaiting administrative activation)

Displaying counts of storage components

You can use the Object Storage Management application or an API method to display counts of storage components in the system.

Object Storage Management application instructions

To display storage counts, select Storage.

The infographic displays the count of active, inactive, and unverified storage components.

Related REST API methods

```
POST /storage_component/list
```

For information about specific API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

Metrics

HCP for cloud scale uses a third-party, open-source software tool, running over HTTPS as a service, to provide storage component metrics through a browser.

The Metrics service collects metrics for these HCP for cloud scale services:

- S3 Gateway
- MAPI Gateway
- Policy Engine
- Metadata Coordination
- Metadata Gateway

By default the Metrics service collects all storage component metrics and you cannot disable collection. By default, the Metrics service collects data every ten seconds (the Scrape Interval) and retains data for 15 days (the Database Retention); you can configure these values in the service by using the System Management application.



Note: Metrics related to HCP for cloud scale instances and services are collected and provided by the System Management application. Collection of these metrics cannot be disabled.

Displaying the active object count

The Object Storage Management application displays a count of active objects stored in the system.

Object Storage Management application instructions

To display the Active Object Count report, select Dashboard.

The report displays a line graph showing the total number of active objects in the system over the past week.

Displaying metrics

You can use the metrics service to display or graph metrics, or use the service API to obtain metrics.

Object Storage Management application instructions

You can display and graph metrics using the metrics GUI.

To display metrics, select Dashboard and then select the Metrics window. The metrics tool opens in a separate browser window.

The metrics tool is a third-party, open-source package. For information about using the metrics tool, see the documentation provided with the tool.

Available metrics

Metrics collect information about the operation of a service. Metrics are collected while the service is active. If a service restarts, its metrics are restarted.

Metrics from all services

The following metrics are available from all services.



Note: If a metric is measured over an interval (for example, `http_s3_servlet_requests_latency_seconds`), but doesn't have at least two data points, the value is reported as NaN.

Metric	Description
<code>http_healthcheck_requests_total</code>	The count of the total number of requests made to the health verification API.
<code>http_monitoring_requests_total</code>	The count of the total number of requests made to the monitoring API.
<code>scrape_duration_seconds</code>	The duration in seconds of the scrape (collection interval).
<code>scrape_samples_post_metric_relabeling</code>	The count of samples remaining after metric relabeling was applied.
<code>scrape_samples_scraped</code>	The count of samples the target exposed.
<code>up</code>	1 if the instance is healthy (reachable) or 0 if collection of metrics from the instance failed.

Message Queue

The Message Queue service provides a large number of general metrics. Information on these metrics is available at <https://github.com/rabbitmq/rabbitmq-prometheus/blob/master/metrics.md>.

Metadata Coordination

The following metrics are available from the Metadata Coordination service.

Metric	Description
mcs_copies_per_partition	The count of number of copies of each metadata partition of each key space (to verify protection).
mcs_disk_usage_per_instance	The total disk usage of each metadata instance.
mcs_disk_usage_per_partition	The disk usage of each metadata partition of each key space.
mcs_failed_moves_per_keyspace	The count of number of unsuccessful requests for metadata partition moves per keyspace.
mcs_failed_splits_per_keyspace	The count of number of unsuccessful requests for metadata partition splits per keyspace.
mcs_moves_per_keyspace	The count of number of successful requests for metadata partition moves per keyspace.
mcs_partitions_per_instance	The count of total number of metadata partitions per metadata instance (to verify balance).
mcs_splits_per_keyspace	The count of number of successful requests for metadata partition splits per keyspace.

Metadata Gateway

The following metrics are available from the Metadata Gateway service.



Note:

1. Client count metrics are an approximation and might not correspond to the actual count.
2. Depending on when garbage collection tasks run, the ratio of client objects size to stored objects size might show a discrepancy.

Metric	Description
async_action_count	The count of actions performed.

Metric	Description
async_action_latency_seconds_bucket	A histogram for the duration, in seconds, of actions on buckets. For actions comprising multiple steps, this is the total of all steps.
async_action_latency_seconds_count	The count of action latency measurements taken.
async_action_latency_seconds_sum	The sum of action latency in seconds.
async_concurrency	A gauge for the number of concurrent actions.
async_duq_latency_seconds_bucket	A histogram for the duration, in seconds, of operations on the durable update queue.
async_duq_latency_seconds_count	The count of durable update queue latency measurements.
async_getwork_database_count	The number of driver work checks accessing the database.
async_getwork_optimized_count	The number of driver work checks avoiding the database.
async_duq_latency_seconds_sum	The sum of actions on durable update queue in seconds.
metadata_clientobject_active_count	The count of client objects in metadata that are in the ACTIVE state.
metadata_clientobject_active_encrypted_count	The count of encrypted client objects in metadata that are in the ACTIVE state.
metadata_clientobject_active_unencrypted_count	The count of unencrypted client objects in metadata that are in the ACTIVE state.
metadata_clientobject_and_part_active_space	the space occupied by client objects and parts in metadata that are in the ACTIVE state.
metadata_clientobject_part_active_count	The count of client object parts in metadata that are in the ACTIVE state.
metadata_storedObject_active_space	The space occupied by stored objects on the back-end storage components.
update_queue_inprogress	The count of update queue entries in progress.
update_queue_size	The size of the update queue.

Policy Engine

The following metrics are available from the Policy Engine service.

Metric	Description
scheduled_policy_items_accepted	The count of items processed by each policy. This is the count of items examined which had operations outstanding.
scheduled_policy_items_examined	The count of items examined by each policy.

S3 Gateway

The following metrics are available from the S3 Gateway service.

Metric	Description
http_s3_servlet_errors_total	The count of total number of errors returned by the s3 servlet, grouped by error.
http_s3_servlet_get_object_response_bytes_total	The count of total bytes in the body of S3 GET object responses.
http_s3_servlet_ingest_object_bytes_per_bucket	The count of total objects ingested for the specified bucket.
http_s3_servlet_operations_total	The count of total number of S3 operations made to the s3 servlet for each method, grouped by operation.
http_s3_servlet_post_object_bytes_total	The count of total bytes of objects posted to S3.
http_s3_servlet_put_copied_bytes_total	The count of total bytes of objects PUT copied to S3.
http_s3_servlet_put_object_bytes_total	The count of total bytes of objects put to S3.
http_s3_servlet_put_object_part_bytes_total	The count of total bytes of PUT part operations to S3.
http_s3_servlet_requests_histogram_latency_seconds	The latency in seconds as measured by a histogram timer, grouped by operation.
http_s3_servlet_requests_histogram_latency_seconds_bucket	The latency in seconds as measured by a histogram timer, grouped by bucket.

Metric	Description
http_s3_servlet_requests_histogram_latency_seconds_count	The count of s3 servlet request observations; used with sum to determine average.
http_s3_servlet_requests_histogram_latency_seconds_sum	Sum of s3 servlet request latency in seconds; used with count to determine average.
http_s3_servlet_requests_latency_seconds	The latency in seconds as measured by a summary timer, grouped by operation.
http_s3_servlet_requests_latency_seconds:hour_average	The latency in seconds over the last hour as measured by a summary timer.
http_s3_servlet_requests_latency_seconds_count	
http_s3_servlet_requests_latency_seconds_sum	The sum of request latency in seconds.
http_s3_servlet_requests_per_bucket	The count of total put, get, or delete requests made to the specified bucket.
http_s3_servlet_requests_total	The count of total number of requests made to the s3 servlet, grouped by method.
http_s3_servlet_unimplemented_api_request_total	The count of total number of requests made for unimplemented S3 methods.
http_s3_servlet_unimplemented_bucket_api_request_total	The count of total number of requests made for unimplemented S3 methods per bucket, grouped by API.
http_s3_servlet_unimplemented_object_api_request_total	The count of total number of requests made for unimplemented S3 methods per object, grouped by API.
http_s3_servlet_unimplemented_service_api_request_total	The count of total number of requests made for unimplemented S3 methods per service, grouped by API.
http_s3_servlet_unknown_api_requests_total	The count of total number of requests made for unknown S3 methods, grouped by API.
s3select_total_bytes_scanned	The count of bytes scanned in the object
s3select_total_bytes_processed	The count of bytes processed by the request

Metric	Description
s3select_total_bytes_returned	The count of bytes returned from the request
s3select_input_type	Count of requests by file type
s3select_output_type	Count of responses by file type
sync_from_bytes	The count of bytes synchronized from external storage (sync-from) by this instance.
sync_from_bytes_copied	The count of bytes synchronized by full copy from external storage (sync-from) by this instance.
sync_from_bytes_putcopied	The count of bytes synchronized by put-copy from external storage (sync-from) by this instance.
sync_from_object_count_failed	The count of objects that failed to synchronize from external storage (sync-from) by this instance.
sync_from_object_count_succeeded	The count of objects synchronized from external storage (sync-from) by this instance.
sync_from_object_size_total	Total size of object data synchronized from external storage (sync-from) by this instance.
sync_from_objects	Total size of object data synchronized from external storage (sync-from) by this instance.
sync_to_bytes	The count of bytes synchronized to external storage (sync-to) by this instance.
sync_to_bytes_copied	The count of bytes synchronized by full copy to external storage (sync-to) by this instance.
sync_to_bytes_putcopied	The count of bytes synchronized by put-copy to external storage (sync-to) by this instance.
sync_to_object_count_failed	The count of objects that failed to synchronize to external storage (sync-to) by this instance.

Metric	Description
sync_to_object_count_succeeded	The count of objects synchronized to external storage (sync-to) by this instance.
sync_to_objects	The count of objects synchronized to external storage (sync-to) by this instance.
sync_to_object_size_total	The total size of object data synchronized to external storage (sync-to) by this instance.

Examples of graphed metrics

By using metrics in formulas, you can generate useful information about the behavior and performance of the HCP for cloud scale system.

Growth of active-object count

The following expression graphs the count of active objects (`metadata_clientobject_active_count`) over time. (This is similar to the graph displayed on the Dashboard.) You can use this formula to determine the growth in the number of active objects.

```
sum(metadata_clientobject_active_count)
```

Sum of update queues

The following expression graphs the size of all update queues. You can use this formula to determine whether the system is keeping up with internal events that are processed asynchronously in response to S3 activity. If this graph increases over time, you might want to increase capacity.

```
sum(update_queue_size)
```

Changes in S3 put requests over time

The following expression graphs the count of S3 put requests, summed across all nodes, at one-minute intervals. If you remove the specifier `{operation="S3PutObjectOperation"}` the expression graphs all S3 requests.

```
sum(rate(http_s3_servlet_operations_total{operation="S3PutObjectOperation"}[1m]))
```

Request time service levels

The following expression divides the latency of requests (`async_duq_latency_seconds_bucket`) in seconds by the number of requests (`async_duq_latency_seconds_count`), for the bucket `getWork` and requests less than or equal to 10 ms, and graphs it over time. You can use this formula to determine the percentage of requests completed in a given amount of time.

```
sum(rate(async_duq_latency_seconds_bucket{op="getWork",le="0.01"}[1m])) /
sum(rate(async_duq_latency_seconds_count{op="getWork"}[1m]))
```

Here is a sample graph of data from a lightly loaded system:



Request time quantile estimates

The following expression estimates the quantile for the latency of requests (`async_duq_latency_seconds_bucket`) in seconds for the bucket `getWork`. You can use this formula to estimate the percentage of requests completed in a given amount of time.

```
histogram_quantile(.9,
sum(rate(async_duq_latency_seconds_bucket{op="getWork"}[1m])) by (le))
```

Here is a sample graph of data from a lightly loaded system:



Tracing requests and responses

HCP for cloud scale uses an open-source software tool, running over HTTPS as a service, for service tracing through a browser.

The Tracing service supports end-to-end, distributed tracing of S3 requests and responses by HCP for cloud scale services. Tracing helps you monitor performance and troubleshoot possible issues.

Tracing involves three service instances:

- Tracing Query: serves traces
- Tracing Agent: receives spans from tracers
- Tracing Collector: receives spans from Tracing Agent service using Tchannel

Displaying traces

You can display traces using the tracing service GUI.

To begin tracing, select Dashboard and then click the Tracing window. The tracing tool opens in a separate browser window.

When tracing, you can specify:

- Service to trace
- Operation to trace (all or specific) for each service
- Tags
- Lookback period (by default, over the last hour)
- Minimum duration
- Number of results to display (by default, 20)

The service displays all found traces with a chart giving the time duration for each trace. You can select a trace to display how the trace is served by difference services in cascade and the time spent on each service.

For information about the tracing tool, see the documentation provided with the tool.

Traceable operations

The following operations are traceable.

Component	Operation
async-policy-engine	Action Pipeline Action: BucketIdToNameMapAction
	Action Pipeline Action: BucketLookupForAsyncPolicyAction

Component	Operation
	Action Pipeline Action: BucketOwnerIdToNameMapAction
	Action Pipeline Action: BucketUpdateSecondaryAction
	Action Pipeline Action: ClientObjectDispatchRemoveBackReferencesAction
	Action Pipeline Action: ClientObjectLookupAction
	Action Pipeline Action: ClientObjectModifyInProgressListAction
	Action Pipeline Action: ClientObjectModifyListAction
	Action Pipeline Action: ClientObjectUpdateSecondaryAction
	Action Pipeline Action: DequeueAction
	Action Pipeline Action: MetadataAction
	BUCKET
	CLIENT_OBJECT
	STORED_OBJECT_BACK_REFERENCE
balance-engine	BalanceCluster
	BalanceEngineOperation
	controlApi.ControlApiService
	RefreshCluster
client-access-service	Action Pipeline Action: BucketAuthorizationAction
	Action Pipeline Action: BucketCountLimitAction
	Action Pipeline Action: BucketCreateAction
	Action Pipeline Action: BucketRegionValidationAction
	Action Pipeline Action: BucketUpdateAclAction
	Action Pipeline Action: ClientObjectInitiateMultipartAction

Component	Operation
	Action Pipeline Action: ClientObjectListInProgressMultipartAction
	Action Pipeline Action: ClientObjectListVersionsAction
	Action Pipeline Action: ClientObjectSizeLimitAction
	Action Pipeline Action: ClientObjectTableLookupAction
	Action Pipeline Action: ClientObjectUpdateAclAction
	Action Pipeline Action: CompleteMultipartUploadAction
	Action Pipeline Action: DataContentAction
	Action Pipeline Action: DataDeletionAction
	Action Pipeline Action: NotAnonymousAuthorizationAction
	Action Pipeline Action: ObjectAuthorizationAction
	Action Pipeline Action: ObjectDataPlacementAction
	Action Pipeline Action: ObjectGetCurrentExpirationAction
	Action Pipeline Action: ObjectGetMultipartAbortDateAction
	Action Pipeline Action: ObjectGetUndeterminedExpirationAction
	Action Pipeline Action: ObjectLookupAction
	Action Pipeline Action: PartDataPlacementAction
	Action Pipeline Action: PutAclAction
	Action Pipeline Action: RequestBucketLookupAction
	Action Pipeline Action: RequestVersionIdValidationAction
	Action Pipeline Action: UploadIdValidationAction
	Action Pipeline Action: UserLookupBucketsAction

Component	Operation
	Action Pipeline Action: VersionIdNotEmptyValidationAction
expiration-rules-engine	EvaluateOperation
foundry-auth-client	FoundryAuthorizeOperation
	FoundryValidateOperation
jaeger-query	/api/dependencies
	/api/services
	/api/services/{service}/operations
	/api/traces
mapi-service	GET
	POST
metadata-client	BucketService/Create
	BucketService/List
	BucketService/ListBucketOwnerListing
	BucketService/LookupBucketNameById
	BucketService/LookupByName
	BucketService/UpdateACL
	ClientObjectService/CloseNew
	ClientObjectService/ClosePart
	ClientObjectService/DeleteSpecific
	ClientObjectService/List
	ClientObjectService/LookupLatest
	ClientObjectService/LookupSpecific
	ClientObjectService/OpenNew
	ClientObjectService/OpenPart
	ClientObjectService/setACLOnLatest
	ClientObjectService/Delete
	ConfigService/List

Component	Operation
	ConfigService/LookupById
	ConfigService/Set
	StoredObjectService/Close
	StoredObjectService/Delete
	StoredObjectService/List
	StoredObjectService/Lookup
	StoredObjectService/MarkForCleanup
	StoredObjectService/Open
	UpdateQueueService/SecondaryEnqueue
	UserService/LookupById
	UserService/LookupOrCreate
	UserService/UpdateAddAuthToken
metadata-coordination-service	Status.Service/GetStatus
metadata-gateway-service	Status.Service/GetStatus
	BucketService/Create
	BucketService/List
	BucketService/ListBucketOwnerListing
	BucketService/LookupBucketNameById
	BucketService/LookupByName
	BucketService/UpdateACL
	ClientObjectService/CloseNew
	ClientObjectService/ClosePart
	ClientObjectService/DeleteSpecific
	ClientObjectService/List
	ClientObjectService/LookupLatest
	ClientObjectService/LookupSpecific
	ClientObjectService/OpenNew
	ClientObjectService/OpenPart

Component	Operation
	ClientObjectService/setACLOnLatest
	ConfigService/Delete
	ConfigService/List
	ConfigService/LookupById
	ConfigService/Set
	StoredObjectService/Close
	StoredObjectService/Delete
	StoredObjectService/List
	StoredObjectService/Lookup
	StoredObjectService/MarkForCleanup
	StoredObjectService/Open
	UpdateQueueService/SecondaryEnqueue
	UserService/LookupById
	UserService/LookupOrCreate
	UserService/UpdateAddAuthToken
metadata-policy-client	PolicyService/ExecutePolicy
metadata-policy-service	ServiceStatus/GetStatus
	PolicyService/ExecutePolicy
	ScheduledDeleteBackendObjectsJob
	ScheduledDeleteFailedWritesJob
	ScheduledExpirationJob
	ScheduledIncompleteMultipartExpirationJob
	ScheduledStorageComponentHealthCheckJob
storage-component-client	InMemoryStorageComponentVerifyOperation
	InMemoryStorageDeleteOperation
	InMemoryStorageReadOperation
	InMemoryStorageWriteOperation
storage-component-manager	StorageComponentManager Operation: Create

Component	Operation
	StorageComponentManager Operation: List
	StorageComponentManager Operation: Lookup
	StorageComponentManager Operation: Update
tomcat-servlet	S3 Operation

Chapter 3: Services

Services perform functions essential to the health and function of the Hitachi Content Platform for cloud scale (HCP for cloud scale) system. The System Management application supports management of services.

For example, the S3 Gateway service serves S3 API methods and communicates with storage components, while the Watchdog service ensures that other services remain running.

Services provide cluster management and coordination, metadata coordination and caching, and external gateways.

Internally, services run in Docker containers on the instances of the system. The container orchestration framework supports cloud or on-premise deployment.

HCP for cloud scale supports an adaptive service deployment model that can change based on workload.

Service categories

Services are grouped into categories depending on what actions they perform.

Services are grouped into these categories:

- *Product services* enable HCP for cloud scale functions. For example, the S3 Gateway service serves S3 API methods and communicates with storage components. You can scale, move, and reconfigure product services.
- *System services* maintain the health and availability of the HCP for cloud scale system. For example, the Watchdog service ensures that other services remain running. You cannot scale, move, or reconfigure system services.

HCP for cloud scale services

This table describes the services that HCP for cloud scale runs. Each service runs within its own Docker container. For each service, the table lists:

- Configuration settings: The settings you can configure for the service.
- RAM needed per instance: The amount of RAM that, by default, the service needs on each instance on which it's deployed. For all services except for System services, this value is also the default Docker value of Container Memory for the service.

- Number of instances: Shows both:
 - The minimum number of instances on which a service must run to function properly.
 - The best number of instances on which a service should run. If your system includes more than the minimum number of instances, you should take advantage of the instances by running services on them.
- Whether the service is persistent (that is, it must run on a specific instance) or supports floating (that is, it can run on any instance).
- Whether the service is scalable or not.



Note: For HCP for cloud scale services, you cannot set the size of Container Memory larger than the value of the setting Max Heap Size. For other services, you should not set the size of Container Memory larger than the value of the setting Max Heap Size.

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
Product services: These services perform HCP for cloud scale functions. You can move and reconfigure these services.			
Cassandra Decentralized database that can be scaled across large numbers of hardware nodes.	Container Options: Default <ul style="list-style-type: none">▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 2.4 GB.▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1.	RAM needed per instance:	2.4 GB
	Service Options <ul style="list-style-type: none">▪ Max heap: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 1200m.	Number of instances:	Minimum: 3 Best: All

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
	<ul style="list-style-type: none"> ▪ Heap new size: Heap size for the Java virtual machine. Valid values are integers representing a number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 512 MB. 	Service unit cost:	10
		Persistent or floating?	Persistent
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes
		Service unit cost:	10
Chronos Job scheduling.	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 712 MB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	RAM needed per instance:	712 MB
		Service unit cost:	1
	Service Options <ul style="list-style-type: none"> ▪ Heap settings: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 356 MB. 	Number of instances:	Minimum: 1 Best: 1
		Persistent or floating?	Floating
		Supports volume configuration?	Yes

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
		Single or multiple types?	Single
		Scalable?	No
		Service unit cost:	1
Elasticsearch Data indexing and search platform.	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 2 GB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	RAM needed per instance:	2 GB
		Service unit cost:	25
	Service Options <ul style="list-style-type: none"> ▪ Heap settings: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 1024 MB. ▪ Days to keep logs: The number of days to keep service logs, including access and metrics indexes. The default is 30 days. ▪ Index Protection Level: The number of additional replicas (copies) to keep of each index file (shard). Replicas are kept on separate instances. You can set this value for every shard. The default is 1 replica (which means that two copies are kept). The maximum is the number of instances less one. 	Number of instances:	Minimum: 3 Best: All
		Persistent or floating?	Persistent
		Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	Yes
		Service unit cost:	25
Kafka	Container Options: Default	RAM needed per instance:	2 GB

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
Stream processing platform for handling real-time data streams.	<ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 2 GB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	Service unit cost:	5
		Number of instances:	Minimum: 3 Best: All
	Service Options <ul style="list-style-type: none"> ▪ Heap settings: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 1 GB. 	Persistent or floating?	Persistent
		Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	Yes
		Service unit cost:	5
Key Management Server Manages storage component encryption keys	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 2 GB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.5. 	RAM needed per instance:	2 GB
		Service unit cost:	10
		Number of instances:	Minimum: 1 Best: 2 or more
		Persistent or floating?	Floating
		Supports volume configuration?	Yes

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
		Single or multiple types?	Single
		Scalable?	Yes
		Service unit cost:	10
Logstash Collection engine for event data. Can perform transformations on the data it collects and then send that data to a number of outputs.	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 700 MB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	RAM needed per instance:	700 MB
	Service Options <ul style="list-style-type: none"> ▪ Max heap: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 350 MB. 	Number of instances:	Minimum: 1 Best: 1
		Persistent or floating?	Floating
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	No
		Service unit cost:	10
MAPI Gateway Serves MAPI methods.	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 768 MB. 	RAM needed per instance:	768 MB

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
	<ul style="list-style-type: none"> ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	Service unit cost:	5
		Number of instances:	Minimum: 1 Best: All
		Persistent or floating?	Floating
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes
		Service unit cost:	5
Message Queue Provides system queueing services.	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 2048 MB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.5. 	RAM needed per instance:	2048 MB
		Service unit cost:	10
	Service Options None.	Number of instances:	Minimum: 3 Best: 3
		Persistent or floating?	Floating

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	No
		Service unit cost:	10
Metadata Cache Cache for HCP for cloud scale metadata	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 1024 MB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	RAM needed per instance:	1024MB
	Service Options <ul style="list-style-type: none"> ▪ Max Heap Size: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 384 MB. 	Number of instances:	Minimum: 1 Best: 1
		Persistent or floating?	Persistent
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	No
		Service unit cost:	10

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
Metadata Coordination Coordinates Metadata Gateway service instances and coordinates scaling and balancing.	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 768 MB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	RAM needed per instance:	768 MB
	Service Options <ul style="list-style-type: none"> ▪ Max Heap Size: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 384 MB. 	Number of instances:	Minimum: 1 Best: 1
		Persistent or floating?	Floating
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	No
		Service unit cost:	5
Metadata Gateway Stores and protects metadata and serves it to other services.	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 4096 MB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	RAM needed per instance:	768 MB
		Service unit cost:	50

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
	Service Options <ul style="list-style-type: none"> ▪ Max Heap Size: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 2048 MB. 	Number of instances:	Minimum: 3 Best: All
		Persistent or floating?	Persistent
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes
		Service unit cost:	50
Metrics	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 768 MB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	RAM needed per instance:	768 MB
		Service unit cost:	10
	Service Options <ul style="list-style-type: none"> ▪ Prometheus Scrape Interval: The time interval between runs of the metrics collection task. Type an integer number of seconds. You can optionally specify the suffix s (seconds). The default is 10 seconds. ▪ Prometheus Database Path: Storage location for prometheus local time-series db. Type a path. The default is tsdb/. 	Number of instances:	Minimum: 1 Best: 1
		Persistent or floating?	Persistent
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
	<ul style="list-style-type: none"> ▪ Prometheus Database Retention: The number of days to retain files. Type an integer number of days. You can optionally specify the suffix d (days). The default is 15 days. 	Service unit cost:	10
Policy Engine Executes system policy operations and asynchronous metadata updates.	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 2048 MB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	RAM needed per instance:	768 MB
		Service unit cost:	25
	Service Options <ul style="list-style-type: none"> ▪ Max Heap Size: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 1024 MB. 	Number of instances:	Minimum: 1 Best: All
		Persistent or floating?	Floating
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes
S3 Gateway Serves S3 API methods and communicates with storage components.	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 768 MB. 	RAM needed per instance:	768 MB

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
	<ul style="list-style-type: none"> ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	Service unit cost:	25
	Service Options <ul style="list-style-type: none"> ▪ Max Heap Size: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 384 MB. 	Number of instances:	Minimum: 1 Best: All
	HTTP Options: <ul style="list-style-type: none"> ▪ Enable HTTP: Select to enable HTTP connections. ▪ Max Http Request Headers: The maximum number of HTTP request headers to allow. Type an integer. The default is 100 request headers. 	Persistent or floating?	Floating
		Supports volume configuration?	No
		Single or multiple types?	Single
	HTTPS Options: <ul style="list-style-type: none"> ▪ SSL Ciphers: A comma-separated list of ciphers used to encode SSL traffic. Changing the list causes the service to redeploy. 	Scalable?	Yes
		Service unit cost:	25
Tracing Agent Supports end-to-end distributed tracing for S3 API calls and MAPI calls.	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 768 MB. 	RAM needed per instance:	768 MB

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
	<ul style="list-style-type: none"> ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	Service unit cost:	1
	Service Options <ul style="list-style-type: none"> ▪ Collector TChannel Hostname: Type a host name. The default is localhost. ▪ Collector TChannel Port: Type a port number. The default is 14267. 	Number of instances:	Minimum: All
		Persistent or floating?	Floating
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes
		Service unit cost:	1
Tracing Collector Supports end-to-end distributed tracing for S3 API calls and MAPI calls.	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 768 MB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	RAM needed per instance:	768 MB
	Service Options <ul style="list-style-type: none"> ▪ ElasticSearch Hostname: Type a host name. The default is localhost. ▪ ElasticSearch Port: Type a port number. The default is 9200. 	Number of instances:	Minimum: 1 Best: All
		Persistent or floating?	Floating

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
	<ul style="list-style-type: none"> ▪ Sampling Rate: The sampling rate for all clients implementing remote sampling. Type a number between 0 and 1 inclusive. The default is 1. ▪ Max open Index age: How long to keep tracing indexes open in the database, in days. Type a value from 1 to 365 days inclusive. The default is 30 days. ▪ Max Index age: How long to keep tracing indexes in the database, in days. Type a value from 1 to 365 days inclusive. The default is 60 days. 	Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes
		Service unit cost:	10
Tracing Query Supports end-to-end distributed tracing for S3 API calls and MAPI calls.	Container Options: Default <ul style="list-style-type: none"> ▪ Container Memory: The hard memory limit for the service's Docker container, in MB. The default is 768 MB. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container receives more CPU resources than other processes (including other service Docker containers) running on the instance. Type a decimal number. The default is 0.1. 	RAM needed per instance:	768 MB
	Service Options <ul style="list-style-type: none"> ▪ ElasticSearch Hostname: Type a host name. The default is localhost. ▪ ElasticSearch Port: Type a port number. The default is 9200. 	Number of instances:	Minimum: 1 Best: All
		Persistent or floating?	Floating
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes
		Service unit cost:	5

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
System services: These services manage system resources and ensure that the HCP for cloud scale system remains available and accessible. These services cannot be moved or reconfigured.			
Admin App The system management application.	Service Options <ul style="list-style-type: none">▪ Max heap size: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 512 MB.	RAM needed per instance:	N/A
		Number of instances:	N/A
		Persistent or floating?	Persistent
		Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	No
Cluster Coordination Manages hardware resource allocation.	N/A	RAM needed per instance:	N/A
		Number of instances:	N/A
		Persistent or floating?	Persistent
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	No
Cluster Worker Receives and performs work from other services.	N/A	RAM needed per instance:	N/A
		Number of instances:	N/A
		Service unit cost:	5

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
		Persistent or floating?	Persistent
		Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	No
Network Proxy Network request load balancer.	Security Protocol: Select which Transport Layer Security (TLS) versions to use: <ul style="list-style-type: none"> ▪ TLS 1.2 ▪ TLS 1.3 	RAM needed per instance:	N/A
		Number of instances:	N/A
	SSL Ciphers: To use your own cipher suite, type it here.	Service unit cost:	1
	Custom Global Configuration: Select Enable Advanced Global Configuration to enable adding your own parameters to the HAProxy "global" section.	Persistent or floating?	Persistent
	Custom Defaults Configuration: Select Enable Defaults Configuration to enable adding your own parameters to the HAProxy "global" section.	Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	No
Sentinel Runs internal system processes and monitors the health of the other services.	Service Options <ul style="list-style-type: none"> ▪ Max heap size: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 1024 MB. 	RAM needed per instance:	N/A
		Number of instances:	N/A
		Persistent or floating?	Persistent

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
		Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	No
Service Deployment Handles deployment of high-level services (that is, the services that you can configure).	N/A	RAM needed per instance:	N/A
		Number of instances:	N/A
		Persistent or floating?	Persistent
		Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	No
Synchronization Coordinates service configuration settings and other information across instances.	Service Options <ul style="list-style-type: none"> ▪ Max heap size: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 1024 MB. 	RAM needed per instance:	N/A
		Number of instances:	N/A
		Persistent or floating?	Persistent
		Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	No
Watchdog	Service Options	RAM needed per instance:	N/A

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
Monitors the other System services and restarts them if necessary. Also responsible for initial system startup.	<ul style="list-style-type: none"> ▪ Max heap size: Maximum amount of memory to allocate to the Java heap for each instance of the service. Type an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 1024 MB. ▪ Heap new size: Heap size for the Java virtual machine. Valid values are integers representing a number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 512 MB. 	Number of instances:	N/A
		Service unit cost:	5
		Persistent or floating?	Persistent
		Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	No
		Service unit cost:	5

Viewing services

You can use Admin App, CLI, and REST API to view the status of all services for the system.

Viewing all services

Admin App instructions

Procedure

1. To view the status of all services, in the Admin App, click **Services**.

For each service, the page shows:

- The service name
- The service state:
 - **Healthy:** The service is running normally.
 - **Unconfigured:** The service has yet to be configured and deployed.
 - **Deploying:** The system is currently starting or restarting the service. This can happen when:
 - You move the service to run on a completely different set of instances.
 - You repair a service.

For information on viewing the status service functions, see [Monitoring service operations \(on page 99\)](#).

- **Balancing:** The service is running normally, but performing background maintenance.
- **Under-protected:** In a multi-instance system, one or more of the instances on which a service is configured to run are offline.
- **Failed:** The service is not running or the system cannot communicate with the service.
- **CPU Usage:** The current percentage CPU usage for the service across all instances on which it's running.
- **Memory:** The current RAM usage for the service across all instances on which it's running.
- **Disk Used:** The current total amount of disk space that the service is using across all instances on which it's running.

Viewing individual service status

Procedure

1. To view the detailed status for an individual service, select the service in the **Services** window.

In addition to status information, the window shows:

- **Instances:** A list of all instances on which the service is running.
- **Volumes:** To view a list of volumes used by the service, click the row for an instance in the **Instances** section.
- **Network: [Internal | External]:** Which network type this service uses to receive communications.

This section also displays a list of the ports that the service uses:

- **Configuration settings:** The settings you can configure for the service.

- **Service Units:** The total number of service units currently being spent to run this service. This value is equal to the service's service unit cost times the number of instances on which the service is running.
- **Service unit cost:** The number of service units required to run the service on one instance.
- **Service Instance Types:** For services that have multiple types, the types that are currently running.
- **Instance Pool:** For floating services, the instances that this service is eligible to run on.
- **Events:** A list of all system events for the service.

Related CLI commands

getService

listServices

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

GET /services

GET /services/{id}

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Listing service ports

You can list service port information for ports available for customer use.

You can list public service ports using an API without an access token.

Related API method

```
POST /public/discovery/get_service_port
```

For information about specific API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

Managing services

This section describes how you can reconfigure, restart, and otherwise manage the services running on your system.

Moving and scaling services

You can change a service to run on:

- Additional instances (for example, to improve service performance and availability)
- Fewer instances (for example, to free up resources on an instance for running other services)
- A different set of instances (for example, to retire the piece of hardware on which an instance is installed)

Moving and scaling floating services

For floating services, instead of specifying the specific instances on which the service runs, you can specify a pool of eligible instances, any of which can run the service.

Moving and scaling services with multiple types

When moving or scaling a service that has multiple types, you can simultaneously configure separate rebalancing for each type.

Best practices

- Avoid running multiple services with high service unit costs together on the same instance.
- On master instances, avoid running any services besides those classified as System services.
- To use your instances evenly, try to deploy a comparable number of service units on each instance.

Considerations

- You cannot remove a service from an instance if doing so would cause or risk causing data loss.
- Service relocations can take a long time to complete and can impact system performance while they are running.
- Instance needs vary from service to service. Each service defines the minimum and maximum number of instances on which it can run.



Tip: Use the Available Instances option to make a floating service eligible to run on any instance in the system, including any new instances added in the future.

Relocating services

To manually relocate a service, in the Admin App:

Procedure

1. Select **Services**.

The **Services** page opens, displaying the services and system services.

2. Click on the service that you want to scale or move.
Configuration information for the service is displayed.
3. Click **Scale**, and if the service has more than one type, select the instance type that you want to scale.

The next step depends on whether the service is floating or persistent (non-floating).

4. If the service is a floating service, you are presented with options for configuring an instance pool. For example:

The screenshot displays the configuration page for a service named 'MAPI-Gateway'. At the top, there are tabs for 'INSTANCES', 'VOLUMES', 'NETWORK', 'CONFIGURATION', 'SCALE' (selected), and 'EVENTS'. Below the tabs, there are statistics: 'Average CPU Usage' at 1.02%, 'Memory Used' at 531.2 MB of 768.0 MB, and 'Disk Used' at 40.0 kB. A 'Service Unit Cost' section shows 'Total: 5 Per Instance: 5' and 'Service Units In Use' at '574 of Unlimited'. The 'Service Instance Configuration' section shows 'Service Instances' set to '1' with a note 'The number of service instances that should be run in the pool. Minimum: 1 Maximum: 3'. There is a checkbox 'All Available Instances' which is checked, with a note 'Enabling this will allow the service to run on any of the instances in the system.' Below this is the 'Instance Pool' section, which lists three instances with their IP addresses, service counts, allocated service units, and load averages.

Instance IP	Services	Allocated Service Units	Load Average
172.18.46.50	18	200	0.58
172.18.46.51	16	197	0.93
172.18.46.52	13	177	0.30

- a. In the field **Service Instances**, specify the number of instances on which the service should be running at any time.
- b. Configure the instance pool:
 - For the service to run on any instance in the system, select **All Available Instances**.
With this option, the service can be restarted on any instance in the instance pool, including instances that were added to the system after the service was configured.
 - For the service to run on a specific set of instances, deselect **All Available Instances**. Then:
 - To remove an instance from the pool, select it from the list **Instance Pool**, on the left, and then click **Remove Instances**.
 - To add an instance to the pool, select it from the list **Available Instances**, on the right, and then click **Add Instances**.
5. If the service is a persistent (non-floating) service, you are presented with options for selecting the specific instances that the service should run on. Do one or both of

these, then click **Next**:

- To remove the service from the instances it's currently on, select one or more instances from the list **Selected Instances**, on the left, and then click **Remove Instances**.
- To add the service to other instances, select one or more instances from the list **Available Instances**, on the right, and then click **Add Instances**.

6. Click **Update.**

The **Processes** page opens, and the **Service Operations** tab displays the progress of the service update as "Running." When the update finishes, the service shows "Complete."

Next steps

After reconfiguration, the service may take a few minutes to appear on the Services page.

Related CLI commands

updateServiceConfig

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /services/configure

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Scaling Metadata Gateway instances

The HCP for cloud scale software lets you deploy an instance of the Metadata Gateway service on every node in your system. You can scale the number of instances up or down as needed.

The Metadata Coordination service manages Metadata Gateway scaling. The service does the following:

- Constantly monitors the Metadata Gateway service and balances data among Metadata Gateway instances as needed
- Moves data into new Metadata Gateway instances
- Moves data out of a Metadata Gateway instance set for removal

Scaling up

Use the System Management application to add new Metadata Gateway instances. You can add more than one instance at a time.

Scaling down

Use the System Management application to remove a Metadata Gateway instance. Before you scale down Metadata Gateway instances, consider the following:

- You can only remove a Metadata Gateway instance from the system when there is one or zero Metadata Gateway instances down.



Note: If more than one instance is down, call Support to remove a Metadata Gateway instance.

- You cannot remove a Metadata Gateway instance when there are only three instances. You first need to add a new Metadata Gateway instance.
- You can only remove one Metadata Gateway instance at a time.

Failure recovery

If a Metadata Gateway instance is down, the data in this instance becomes underprotected. To solve this situation, remove the Metadata Gateway instance that is down so that the Metadata Gateway service can recover the data protection. You should first add a new Metadata Gateway instance before removing the instance that is down. This ensures that the system keeps the same performance and capacity usage and also that there is a suitable target instance to recover the data protection. When removing the Metadata Gateway instance, the considerations on scaling down services apply.

A snapshot shows the current state of the state machine from a leader node (service instance) to any follower service instance that is out of synch. If a leader node runs out of space to store snapshots and can't send out its latest snapshot, the follower node cannot resynchronize. If this happens, bring down the leader service instance, increase its storage space, and restart the service.

Configuring service settings

You can configure settings for some of the services that the system runs.



Note: If you make an unwanted change to a service configuration, wait for the configuration to finish before correcting the error.

Configuring service settings

Procedure

1. Select the **Services** window.
2. Select the service you want to configure.
3. On the **Configuration** tab, configure the service.
4. Click **Update**.

Related CLI commands

updateServiceConfig

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /services/configure

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Repairing services

If a service becomes slow, unresponsive, or shows a status of Failed, you can repair it. If you change the configuration of a service you use the same process to restart it.

Repairing a service stops and restarts the service on each instance on which it's running.

If you change the cluster name (cluster hostname), you must repair the S3 Gateway services for the change to take effect.

If you regenerate or upload an SSL certificate, you must repair the S3 Gateway and MAPI services for the change to take effect.

If you upload an SSL certificate for access to a remote system for bucket synchronization, you must repair the Policy Engine and MAPI services for the change to take effect.



Important: Depending on which service you're repairing, parts of the system will be unavailable until the repair finishes.

To repair a service:

Procedure

1. Click the **Services** window.
2. Select the service you want to repair.
3. Click **Repair**.
The **Processes** window opens, displaying a progress bar for the repair process.

Configuring TLS cipher suite

By default, the S3 Gateway service supports the Transport Layer Security (TLS) 1.2 and 1.3 protocols to establish secure links to client nodes. TLS uses a suite of cryptographic algorithms to provide security for HTTPS traffic. The default cipher suite used by the S3 Gateway service is the Open Web Application Security Project (OWASP) security recommendation for Tomcat. You can provide your own cipher suite.



Note: Contact Support if you need to use the TLS 1.0 or 1.1 protocols.

To configure cipher suites:

Procedure

1. Select the **Services** window.
2. Select the service **S3-Gateway**.
3. On the **Configuration** tab, in the **HTTPS Options** section, enter the new cipher suite in the field **SSL Ciphers**.
4. Click **Update**.
The service redeploys.

Avoiding Message Queue shutdown

If two of the three Message Queue service instances fail, the service shuts down. To avoid the possible loss of queued messages, resolve any situation in which only two service instances are running.

To protect messaging consistency, the Message Queue service always has three service instances. To prevent being split into disconnected parts, the service shuts down if half of the service instances fail. In practice, messaging stops if two of the three instances fail.

Do not let the service run with only two instances, because in that scenario if one of the remaining instances fails, the service shuts down. However, when one of the failed instances restarts, messaging services recover and resume.

To protect the Message Queue service, immediately address a node failure where an instance cannot be restarted, because if two service instances are lost and cannot be recovered, the service cannot recover its previous state. You can still add new instances to form a new cluster, but messages that were queued are lost.

In case of such a multi-node failure, the best practice is to restart the Policy Engine service instances one at a time after the Message Queue service cluster is re-formed. This forces the service instances to recover configurations that might have been missed while the Message Queue service was down. Additionally, after the Message Queue service cluster is re-formed, bucket sync-to events that were in the messaging queues are lost, so you might need to regenerate bucket sync-to events for such objects.

The cluster forms based on instance names, including the IP address of the node on which an instance runs. Therefore, changing node configurations such as IP addresses can cause nodes to be permanently removed from the cluster, possibly triggering a shutdown. If this happens, first add instances to the messaging service. Ensure the instances synchronize with the cluster before taking nodes offline or changing node configurations such as IP addresses. This way, the cluster can keep over half of its instances running at all times.

Chapter 4: Monitoring

Your system gives a number of mechanisms that allow you to monitor the health and performance of the system and all of its instances and services.


Monitoring instances

You can use the Admin App, CLI, and REST API to view a list of all instances in the system.


Viewing all instances

To view all instances, in the Admin App, click Dashboard > Instances.

The page shows all instances in the system. Each instance is identified by its IP address.



This table describes the information shown for each instance.

Property	Description
State	<ul style="list-style-type: none">▪ Up: The instance is reachable by other instances in the system.▪ Down: The instance cannot be reached by other instances in the system.
Services	The number of services running on the instance.
Service Units	<p>The total number of service units for all services and job types running on the instance, out of the best-practice service unit limit for the instance.</p> <p>An instance with a higher number of service units is likely to be more heavily used by the system than an instance with a lower number of service units.</p> <p>The Instances page displays a blue bar for instances running less than the best-practice service unit limit.</p> <p>The Instances page displays a red bar for instances running more than the best-practice service unit limit.</p> <p>Service Units  138 of 100 recommended Service Units</p>

Property	Description
Load Average	The load averages for the instance for the past one, five, and ten minutes.
CPU	The sum of the percentage utilization for each CPU core in the instance.
Memory Allocated	This section shows both: <ul style="list-style-type: none"> ▪ The amount of RAM on the instance that's allocated to all services running on that instance. ▪ The percentage of this allocated RAM to the total RAM for the instance.
Memory Total	The total amount of RAM for the instance.
Disk Used	The current amount of disk space that your system is using in the partition on which it is installed.
Disk Free	The amount of free disk space in the partition in which your system is installed.

Viewing the services running on an instance

To view the services running on an individual instance, in the Admin App:

Procedure

1. Click **Dashboard > Instances**.
2. Select the instance you want.

The page lists all services running on the instance.

For each service, the page shows:

- The service name
 - The service state:
 - **Healthy:** The service is running normally.
 - **Unconfigured:** The service has yet to be configured and deployed.
 - **Deploying:** The system is currently starting or restarting the service. This can happen when:
 - You move the service to run on a completely different set of instances.
 - You repair a service.
- For information on viewing the status service operations, see [Monitoring service operations \(on page 99\)](#).
- **Balancing:** The service is running normally, but performing background maintenance.
 - **Under-protected:** In a multi-instance system, one or more of the instances on which a service is configured to run are offline.
 - **Failed:** The service is not running or the system cannot communicate with the service.
 - **CPU Usage:** The current percentage CPU usage for the service across all instances on which it's running.
 - **Memory:** The current RAM usage for the service across all instances on which it's running.
 - **Disk Used:** The current total amount of disk space that the service is using across all instances on which it's running.

Related CLI commands

`getInstance`

`listInstances`

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

`GET /instances`

`GET /instances/{uuid}`

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Monitoring services

You can use Admin App, CLI, and REST API to view the status of all services for the system.

Viewing all services

To view the status of all services, in the Admin App, click Services.

For each service, the page shows:

- The service name
- The service state:
 - Healthy: The service is running normally.
 - Unconfigured: The service has yet to be configured and deployed.
 - Deploying: The system is currently starting or restarting the service. This can happen when:
 - You move the service to run on a completely different set of instances.
 - You repair a service.

For information on viewing the status service operations, see [Monitoring service operations. \(on page 99\)](#)

- Balancing: The service is running normally, but performing some background maintenance operations.
- Under-protected: In a multi-instance system, one or more of the instances on which a service is configured to run are offline.
- Failed: The service is not running or the system cannot communicate with the service.
- CPU Usage: The current percentage CPU usage for the service across all instances on which it's running.
- Memory: The current RAM usage for the service across all instances on which it's running.
- Disk Used: The current total amount of disk space that the service is using across all instances on which it's running.

Viewing individual service status

To view the detailed status for an individual service, select the service on the Services page.

In addition to the status information, the page shows:

- Instances: A list of all instances on which the service is running.
- Volumes: To view a list of volumes used by the service, select the row for an instance in the Instances section.
- Network: [Internal | External]: Which network type this service uses to receive communications.

This section also displays a list of the ports that the service uses.

- Configuration settings: The settings you can configure for the service.

- **Service Units:** The total number of service units currently being spent to run this service. This value is equal to the service's service unit cost times the number of instances on which the service is running.
- **Service unit cost:** The number of service units required to run the service on one instance.
- **Service Instance Types:** For services that have multiple types, the types that are currently running.
- **Instance Pool:** For floating services, the instances that this service is eligible to run on.
- **Events:** A list of all system events for the service.

Related CLI commands

getService

listServices

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

GET /services

GET /services/{id}

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Monitoring processes

The Processes page lets you view information about what the system is doing. This includes any service operations you started and any internal maintenance processes the system needs to run.

Monitoring service operations

You can use the Admin App, CLI, and REST API to monitor all service operations. These include:

- The initial deployments of services when the system was installed.
- Service relocations that you begin.

For each one, the system shows:

- The name of the service involved
- The status of the operation
- The number of steps completed out of the total number of steps

Admin App instructions

Procedure

1. Select the **Processes** window.

Result

The Service Operations tab shows information about in-progress and completed service operations.

Related CLI commands

listSystemTasks

getSystemTask

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

GET /tasks/system

GET /tasks/system/{uuid}

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Monitoring system processes

You can use Admin App, REST API, and CLI to view the progress of internal system processes. These include package installation tasks and regularly scheduled system maintenance activities such as log rotation.

For each process, your system shows:

- The process name
- The process state
- The times at which each step in the process run occurred



Note: System processes have a type of SCHEDULED or ONE-TIME.

Admin App instructions

Procedure

1. In the Admin App, select **Processes**.
2. To view the currently running processes, select the **System** tab.
3. To view the scheduled processes, select the **Scheduled** tab.

Related CLI commands

listSystemTasks

getSystemTask

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

GET /tasks/system

GET /tasks/system/{uuid}

You can get help on specific REST API methods for the Admin App at REST API - Admin.

System events

Your system maintains a log of system events that you can view through the Admin App, CLI, and REST API.

Admin App instructions

Procedure

1. To view all system events, in the Admin App, click **Events**.

Related CLI commands

queryEvents

To view events through the CLI, your requests need to specify which events you want to retrieve.

For example, this JSON request body searches the event log for all events that have a severity level of `warning`:

```
{
  "severities": [
    "warning"
  ]
}
```

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /events

To view events through the REST API, your requests need to specify which events you want to retrieve.

For example, this JSON request body searches the event log for all events that have a severity level of `warning`:

```
{
  "severities": [
    "warning"
  ]
}
```

You can get help on specific REST API methods for the Admin App at [REST API - Admin](#).

HCP for cloud scale events

Most events are generated by and reported through the System Management application.

Events specific to HCP for cloud scale are reported with the IDs 6006 (informational), 6007 (warning), and 6008 (severe). These events are:

ID	Severity	Message	Description
6006	INFO	Service Information: Job Configuration 'id' of type 'job_type' updated with status 'status'	The policy configuration has changed and the status of job <i>id</i> of type 'job_type' is one of the following: <ul style="list-style-type: none"> ▪ ENABLED ▪ DISABLED
6006	INFO	Service Information: Job Configuration 'id' started	Policy configuration for job <i>id</i> has started.
6006	INFO	Service Information: Lifecycle policy {CREATE UPDATE DELETE} bucket 'bucket_name'	The lifecycle policy for bucket <i>bucket_name</i> has been either created, updated, or deleted.
6006	INFO	Service Information: Lifecycle policy deleted for bucket 'bucket_name'	The lifecycle policy for bucket <i>bucket_name</i> has been deleted.

ID	Severity	Message	Description
6006	INFO	Service Information: Serial number updated to <i>value</i>	The HCP for cloud scale serial number has been changed to <i>value</i> .
6006	INFO	Service Information: <i>setting</i> was set to <i>value</i>	The S3 setting <i>setting</i> has been changed to <i>value</i> .
6006	INFO	Service Warning: Storage component ' <i>id</i> ' created	The storage component <i>id</i> has been created.
6006	INFO	Service Warning: Storage component ' <i>id</i> ' is now <i>state</i>	The storage component <i>id</i> is in one of the following:states <ul style="list-style-type: none"> ▪ ACTIVE ▪ INACTIVE ▪ UNVERIFIED
6006	INFO	Service Warning: Storage component ' <i>id</i> ' updated: <i>configuration</i>	The storage component <i>id</i> has been updated. <i>configuration</i> lists the changes.
6007	WARNING	Service Warning: Certificate for Storage component ' <i>id</i> ' is about to expire in ' <i>n</i> ' days	The SSL certificate for the storage component <i>id</i> is set to expire in <i>n</i> days. If the certificate expires, HCP for cloud scale will not be able to read from or write to the storage component.
6007	WARNING	Service Warning: Storage component ' <i>id</i> ' is now INACCESSIBLE	The storage component <i>id</i> is inaccessible.HCP for cloud scale cannot read from or write to the storage component.

ID	Severity	Message	Description
6008	SEVERE	Service Error: Storage Component Certificate Expired.	The SSL certificate for a storage component has expired. HCP for cloud scale cannot read from or write to the storage component.
6008	SEVERE	Service Error: Vault Connection Issue	The active vault node can't be reached.
6008	SEVERE	Service Error: Vault Service Completely Sealed	The vault service (Key Management Server service) is completely sealed. Unseal it using the unseal keys you obtained when you turned on encryption.
6008	SEVERE	Service Error: Vault Service Node Error. IP: <i>ip_address</i>	One of the vault nodes can't be reached. If other active nodes are available service continues, but attend to this issue immediately.
6008	SEVERE	Service Error: Vault Service Node Sealed. IP: <i>ip_address</i>	One of the vault nodes is sealed. If other active nodes are available service continues, but attend to this issue immediately. Unseal it using the unseal keys you obtained when you turned on encryption.

Alerts

Alert messages notify you of situations that need your attention. Alerts can have a severity of Info, Warning, Severe, or Critical. You can view system alerts through the Admin App, CLI, or REST API, and storage component alerts through the Object Storage Management app.

Each alert corresponds to a system event.

System alerts

Severity	Alert Description	Action
Severe	Instance <i>ip-address</i> disk usage severe threshold	The specified instance has less than 10% free disk space. Add additional storage to the instance. Important: If an instance runs out of disk space, the system can become unresponsive.
Severe	Master Instance <i>ip-address</i> is down	Do one of these: <ul style="list-style-type: none"> Restart the instance hardware or virtual machine. Restart the script <code>run</code> on the instance. This script is located in the folder <code>bin</code> in the installation folder.
Severe	Service is down	Verify the health of your instances. If one is down, do one of these: <ul style="list-style-type: none"> Restart the instance hardware or virtual machine. Restart the script <code>run</code> on the instance. This script is located in the folder <code>bin</code> in the installation folder. Otherwise, if your instances are healthy and the problem persists, contact Support.
Severe	Worker Instance <i>ip-address</i> is down	Do one of these: <ul style="list-style-type: none"> Restart the instance hardware or virtual machine. Restart the script <code>run</code> on the instance. This script is located in the folder <code>bin</code> in the installation folder.

Severity	Alert Description	Action
Warning	Instance <i>ip-address</i> disk usage warning threshold	The specified instance has less than 25% free disk space. Add additional storage to the instance. Important: If an instance runs out of disk space, the system can become unresponsive.
Warning	Package installation failed	Your system failed to install a package that you uploaded.
Warning	Service below recommendation	The service is currently running on fewer than the minimum number of instances. Configure this service to run on additional instances.
Warning	Service under-protected	A service has lost redundancy; that is, one or more instances on which that service is running are unresponsive. Verify the health of your instances. If one is down, do one of these: <ul style="list-style-type: none"> Restart the instance hardware or virtual machine. Restart the script <code>run</code> on the instance. This script is located in the folder <code>bin</code> in the installation folder. Otherwise, if your instances are healthy and the problem persists, contact Support.
Warning	SSL server certificate chain expires soon	A certificate in the SSL server certificate chain for this system expires soon. If the certificate chain expires, users can't access the system.
Warning	SSL server certificate chain expired	The SSL server certificate chain for this system contains an expired certificate. Users cannot access the system until the certificate chain is replaced.
Info	Package installation in progress	Your system is currently installing a package that you uploaded. Depending on the contents of the package, this might take a while.
Warning	The certificate for the storage component (<i>storage-id</i>) is about to expire in n days	Renew the storage component certificate.
Info	The storage component (<i>storage-id</i>) is unavailable	Verify that the storage component ID is correct and valid and that the storage component is active.

Storage component alerts

Severity	Message	Description
Warning	Certificate for Storage component <i>id</i> is about to expire in <i>n</i> days	The SSL certificate for the storage component <i>id</i> is set to expire in <i>n</i> days. If the certificate expires, HCP for cloud scale will not be able to read from or write to the storage component.
Warning	Storage component <i>id</i> is now inaccessible	The storage component <i>id</i> is in the state INACCESSIBLE. HCP for cloud scale cannot read from or write to the storage component.
Severe	Certificate for Storage component <i>id</i> expired	The SSL certificate for the storage component <i>id</i> has expired. HCP for cloud scale cannot read from or write to the storage component. Install a new certificate.
Severe	Error communicating with a vault node. Node IP: <i>ip_address</i>	One of the vault nodes can't be reached. If other active nodes are available service continues, but attend to this issue immediately. Examine the vault instance logs to determine the cause of this issue.

Severity	Message	Description
Severe	Failed to connect to KMS server	<p>One of the vault nodes can't be reached. If other active nodes are available service continues, but attend to this issue immediately.</p> <p>If ingest is halted, then investigate why the KMS service is failing to run on all nodes. If ingest is still working, the original active node has failed over. Examine the vault instance logs to determine the cause of the failure.</p>
Severe	Failed to connect to KMS server as it is completely sealed	<p>The vault service (Key Management Server service) is completely sealed.</p> <p>Unseal it using the unseal keys you obtained when you turned on encryption.</p>
Severe	Service error: There is a critical issue with the Metadata Gateway database. Shutting down the Metadata Gateway Service.	<p>A Metadata Gateway instance has encountered an issue and shut down. Use the System Management Services function Repair to restart it.</p> <p>If restarting the service doesn't resolve the issue, contact Support.</p>
Severe	Vault node is sealed. Node IP: <i>ip_address</i>	<p>One of the vault nodes is sealed. If other active nodes are available service continues, but attend to this issue immediately.</p> <p>Unseal it using the unseal keys you obtained when you turned on encryption.</p>

Severity	Message	Description
Critical	Failed to connect to KMS server	<p>The Key Management System service is not available. Until the service is available, data on encrypted storage components can't be read or written.</p> <p>When KMS service restarts, if there is only one active instance log in to HCP for cloud scale on port 8200 and provide unseal keys to reopen the vault.</p>
Critical	Failed verification for retrieved encryption key for StorageComponent_ID{uuid= <i>uuid</i> }	<p>The encryption key returned from the Key Management System server doesn't match the key for the storage component <i>uuid</i>.</p> <p>Verify that the KMS service is available. If the service is available, verify that you have provided the service with a quorum of unseal keys. If objects on the storage component still can't be read, contact Support.</p>
Critical	Metadata-Coordination cannot communicate with Sentinel service to get state information	<p>The Sentinel service is not responding to requests for state information. Using the System Management application, immediately review the health of the Metadata-Coordination and Sentinel services and ensure that the Sentinel container has adequate heap size for the configuration of the cluster.</p>

Viewing alerts

Admin App instructions

Procedure

1. To view alerts, click the user icon (👤) in the top right corner of each Admin App page and then click **Notifications**.

Object Storage Management application instructions

The Object Storage Management application displays alerts about storage components. If an alert is raised the alert icon turns red and displays a badge with the number of active alerts. For example:



Click the icon to display a window listing alert text.

Related CLI commands

listAlerts

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

GET /alerts

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Related REST API methods (HCP for cloud scale)

```
POST /alert/list
```

For information about specific HCP for cloud scale API methods, see the *MAPI Reference* or, in the Object Storage Management application, click the profile icon and select REST API.

Email notification rules

For the system to send email notifications, you need to create a rule that specifies who to email, what email server to use, what events to send emails about, and what information to include in email messages.

SMTP settings

- Enable: Turns on email notifications.
- Host: The hostname or IP address of the email server.
- Port: The port on which the email server listens for email messages.

- **Security:** The security protocol used by the email server (SSL or STARTTLS) or None if the email server doesn't use a security protocol.
- **Authenticated:** Enable this if the email server needs authentication, then specify:
 - In the Username field, the username for an email account that's authorized to establish the connection between the system and the email server.
 - In the Password field, the password for the email account.

Message settings

You use the email notification message settings to configure a template for formatting all email notifications sent by the system.

- **From:** The email address from which you want email notifications to be sent.
- **Subject:** The email subject.
- **Body:** The email message body.

Message variables

This table lists the variables you can use to make the email notification template. When the system sends an email notification, it replaces the variables in the notification with event-specific information.

Variable	Description
\$severity	Event severity: INFO, WARNING, or SEVERITY.
\$subject	A short description of the event.
\$message	Event message text.
\$userName	Name of the user responsible for the event.
\$objectId	Unique identifier for component affected by the event.
\$subsystem	Category for the component affected by the event.
\$objectSourceId	Unique identifier of the internal system component or process that was the source of the event. Value is [unknown] for most events.

Recipient settings

- **Email addresses:** A comma-separated list of email addresses to send notification emails to.
- **Severity Filter:** The event severities about which to send email notifications. Can be one or more of these: INFO, WARNING, SEVERITY.

Creating email notification rules

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Notifications**.
3. Click **Create**.
4. In the **Type** field, select **Email**.
5. Type a name for the notification rule.
6. Configure the SMTP settings and message settings for the notification rule.
7. Specify a comma-separated list of emails to send notifications to.
8. Specify a comma-separated list of emails to send notifications to.
9. Click **Create**.

Related CLI commands

createNotificationRule

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /notifications

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Creating syslog notification rules

When you create a syslog notification rule, the system sends log messages to your syslog server for each applicable system event.

Syslog settings

- **Enable:** Turns on syslog notifications
- **Host:** The hostname or IP address of the syslog server
- **Port:** The port on which the syslog server listens for log messages
- **Facility:** Category for the messages sent by this notification rule

Message settings

You use the syslog notification message settings to configure a template for formatting all syslog notifications sent by this notification rule.

- **Message:** The message to send. You can use these variables as part of the message:

Variable	Description
\$severity	Event severity: INFO, WARNING, or SEVERITY
\$subject	A short description of the event
\$message	Event message text
\$time	Time at which the event occurred
\$userName	Name of the user responsible for the event
\$subsystem	Category for the component affected by the event
\$objectId	Unique identifier for component affected by the event
\$objectType	The type of the component affected by the event.
\$objectSourceId	Unique identifier of the internal system component or process that was the source of the event. Value is [unknown] for most events.
\$objectSourceType	Type of the internal system component or process that was the source of the event. Value is [unknown] for most events.

- **Sender Identity:** Identity of the sender for the event. Sent with every syslog message.

Severity filter

The event severities about which to send email notifications. Can be one or more of these: INFO, WARNING, or SEVERITY.

Creating syslog notification rules

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Notifications**.
3. Click **Create**.
4. In the **Type** field, select **Syslog**.
5. Type a name for the notification rule.
6. Configure the settings for the notification rule.
7. Specify a severity filter for the notification rule.
8. Click **Create**.

Related CLI commands

createNotificationRule

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /notifications

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Logs and diagnostic information

Each service maintains its own set of logs. By default, log files are maintained in the folder `install_path/log` on each instance in the system. During installation, you can configure each service to store its log files in a different, non-default location.

Log levels

This table lists the available log levels.



Note: Raising the log level (for example, from WARN to INFO) results in writing more data to the log file, but the file size increases more rapidly. Lowering the log level (for example, from WARN to ERROR) results in the file size increasing more slowly, but results in writing less data to the log file.

Level	Levels included
ALL	FATAL, ERROR, WARN, INFO, DEBUG, TRACE
TRACE	FATAL, ERROR, WARN, INFO, DEBUG, TRACE
DEBUG	FATAL, ERROR, WARN, INFO, DEBUG
INFO	FATAL, ERROR, WARN, INFO
WARN	FATAL, ERROR, WARN (default)
ERROR	FATAL, ERROR
FATAL	FATAL
OFF	None

Log management

You can manage any of the log files yourself. That is, you can delete or archive them as necessary.



Caution: Deleting log files can make it more difficult for support personnel to resolve issues you might encounter.

System logs are managed automatically in one or more of these ways:

- All log files are periodically added to a compressed file and moved to `install_path/retired/`. This occurs at least one time per day, but can also occur at other times. For example:
 - Whenever you run the `log_download` script.
 - Hourly, if the system instance's disk space is more than 60% full.
 - At the optimum time for a specific service.
- When a log file grows larger than an optimum size, the system stops writing to that file, renames it, and begins writing to a new file. For example, if the file `exampleService.log.0` grows to 10MB, it is renamed to `exampleService.log.1` and the system creates a new file named `exampleService.log.0` to write to.
- When an optimum number of log files for a specific service is reached, the system can overwrite the oldest file. For example, if a service is limited to 20 log files, when the file `exampleService.log.19` is filled, the system overwrites the file named `exampleService.log.0`.

Retrieving logs and diagnostic information

The tool `log_download` lets you easily retrieve logs and diagnostic information from all instances in the system. This tool is located at this path on each instance:

```
install_path/bin/log_download
```

For information about running the tool, use this command:

```
install_path/bin/log_download -h
```



Note:

- When using the tool `log_download`, if you specify the option `--output`, do not specify an output path that contains colons, spaces, or symbolic links. If you omit the option `--output`, you cannot run the script from within a folder path that contains colons, spaces, or symbolic links.
- When you run the script `log_download`, all log files are automatically compressed and moved to the folder `install_path/retired/`.
- If an instance is down, you need to specify the option `--offline` to collect the logs from that instance. If your whole system is down, you need to run the `log_download` script with the `--offline` option on each instance.

Default log locations

By default, each service stores its logs in its own folder at this path:

```
install_path/log
```

This table shows the default log folder names for each service. Depending on how your system was configured when first deployed, your system's logs might not be stored in these directories.

Default log folder name	Related service	Contains information about
com.hds.ensemble.plugins.service.adminApp	Admin-App	The System Management application.
com.hds.ensemble.plugins.service.cassandra	Database	<ul style="list-style-type: none"> System configuration data. Document fields and values.
com.hds.ensemble.plugins.service.chronos	Scheduling	Workflow task scheduling.
com.hds.ensemble.plugins.service.elasticsearch	Metrics	The storage and indexing of: <ul style="list-style-type: none"> System events Performance and failure metrics for workflow tasks
com.hds.ensemble.plugins.service.haproxy	Network-Proxy	Network requests between instances.
com.hds.ensemble.plugins.service.logstash	Logging	The transport of system events and workflow task metrics to the Metrics service.
com.hds.ensemble.plugins.service.marathon	Service-Deployment	The deployment of high-level services across system instances. High-level services are the ones that you can move and configure, not the services grouped under System Services.

Default log folder name	Related service	Contains information about
com.hds.ensemble.plugins.service.mesosAgent	Cluster-Worker	Work ordered by the Cluster-Coordination service.
com.hds.ensemble.plugins.service.mesosMaster	Cluster-Coordination	Hardware resource allocation.
com.hds.ensemble.plugins.service.remoteAction	Watchdog	Internal system processes.
com.hds.ensemble.plugins.service.sentinel	Sentinel	Internal system processes.
com.hds.ensemble.plugins.service.solr	Index	Index collections and search indexes.
com.hds.ensemble.plugins.service.watchdog	Watchdog	General diagnostic information.
com.hds.ensemble.plugins.service.zookeeper	Synchronization	Coordination of actions and database operations across instances.
com.hitachi.aspen.foundry.service.mapi.gateway	MAPI-Gateway	Management of API requests.
com.hitachi.aspen.foundry.service.rabbitmq.server	Message Queue	Transmission of data between instances.
com.hitachi.aspen.foundry.service.metadata.coordination	Metadata-Coordination	Transmission of system metadata between instances.
com.hitachi.aspen.foundry.service.metadata.gateway	Metadata-Gateway	Transmission of system metadata between instances.
com.hitachi.aspen.foundry.service.metadata.async.policy.engine	Policy-Engine	Updates.
com.hitachi.aspen.foundry.service.clientaccess.data	S3-Gateway	Transmission of S3 requests to endpoints.
com.hitachi.aspen.foundry.service.jaeger.agent	Tracing-Agent	Usage of traces.

Default log folder name	Related service	Contains information about
com.hitachi.aspen.foundry.service.jaeger.collector	Tracing-Collector	Usage of tracing collections.
com.hitachi.aspen.foundry.service.jaeger.query	Tracing-Query	Usage of tracing queries.

Chapter 5: Security

The System Management application supports configuring system security features, including user authentication.

Granting access to users

These are the general steps you need to take to grant users access to the system:

1. Add one or more identity providers to the system.
For information, see [Adding identity providers \(on page 121\)](#).
2. Add one or more groups from your identity providers to the system.
For information, see [Adding groups \(on page 126\)](#).
3. Create a role that contains the system permissions you want to associate with a group of users.
For information, see [Creating roles \(on page 129\)](#).
4. Associate roles with groups.
For information, see [Assigning roles to groups \(on page 128\)](#).

Setting the session timeout limit

You can use the Admin App, REST API, or CLI to set the system session timeout limit. This limit affects user sessions in all applications that your system runs and also affects the length of time that REST API authorization tokens are valid.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. On the **Settings** tab, type a number of minutes in the **Session Timeout** field.
4. Click **Update**.

Related CLI commands

`editSecuritySettings`

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

PUT /security/settings

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Setting the refresh token timeout limit

You can use the Admin App, REST API, or CLI to set the refresh token timeout limit. The refresh token timeout limit must be greater than the session timeout limit so that if the access token expires, the refresh token will still be active and you can request a new session token. When your refresh token expires, you will need to resubmit your credentials to access your system.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. On the **Settings** tab, type a number of minutes in the **Refresh Token Timeout** field.
4. Click **Update**.

Related CLI commands

editSecuritySettings

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

PUT /security/settings

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Setting the CORS Allowed Origins

You can use the Admin App, REST API, or CLI to set CORS (cross-origin resource sharing) origins that are allowed on your system. Specifying multiple origins allows you to access restricted resources.

Admin App instructions

Procedure

1. Select the **Configuration** window.

2. Click **Security**.
3. On the **Settings** tab, type a list of origins in the **CORS Allowed Origins** field.
4. Click **Update**.

Related CLI commands

editSecuritySettings

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

PUT /security/settings

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Identity providers

The system supports these identity provider types for user authentication:

- Active Directory (AD)
- OpenLDAP
- 389 Directory Server
- LDAP Compatible: Other LDAP-compatible identity providers not listed above.

To use one of these systems to authenticate users with your system, you need to first add your identity provider to the system.

Adding identity providers

For information on the types of identity providers you can add, see [Identity provider configuration settings \(on page 122\)](#).

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. On the **Identity Providers** tab, click **Create**.
4. Select and configure an identity provider type. For information, see [Identity provider configuration settings \(on page 122\)](#).
5. Click **Create**.

Related CLI commands

createIdentityProvider

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /security/identityProviders

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Identity provider configuration settings

These sections describe the configuration settings for each type of identity provider that your system supports.

All types

Security Realm Name: The name by which to identify this identity provider in the system. This name appears as an option in the Security Realm list on Admin App login pages.



Tip: To ensure that your users can easily log into the system, pick security realm names that your users will recognize and understand.

Active Directory

- **Identity Provider Hostname:** Hostname or IP address for the identity provider.
- **Transport Security:** The protocol to use for securing communications between the system and the identity provider. Options are:
 - None
 - TLS Security (Transport Layer Security)
 - SSL (Secure Sockets Layer)
- **Identity Provider Host Port:** Network port used to communicate with the identity provider. The default value depends on the Transport Security setting:
 - For None or TLS Security (Transport Layer Security), 389
 - For SSL (Secure Sockets Layer), 636
- **User Name:** A user account on the identity provider. Your system uses this user account to read information from the identity provider.
- **Password:** The user account password.
- **Domain:** The AD domain in which the user account is defined.



Note: Use the short name for the AD domain. For example, use `MYACTIVEDIRECTORY` instead of `MYACTIVEDIRECTORY.local`.

- **Search Base DN:** The distinguished name (DN) of the identity provider location where you want your system to begin its searches for users and groups.

For example, if you specify a value of OU=Users,DC=corp,DC=example,DC=com, the system searches for users and groups in the organization unit called Users in the corp.example.com domain.

- **Default Domain Name:** The default domain for users logging into the Admin App and Search App. For example, if you specify a default domain name of east.example.com, the user jdoe@east.example.com needs to specify only jdoe when logging into either app.

LDAP Compatible

- **Identity Provider Hostname:** Hostname or IP address for the identity provider.
- **Transport Security:** The protocol to use for securing communications between the system and the identity provider. Options are:
 - None
 - TLS Security (Transport Layer Security)
 - SSL (Secure Sockets Layer)
- **Identity Provider Host Port:** Network port used to communicate with the identity provider. The default value depends on the Transport Security setting:
 - For None or TLS Security (Transport Layer Security), 389
 - For SSL (Secure Sockets Layer), 636
- **User Name:** A user account on the identity provider. Your system uses this account to read information from the identity provider.
- **Password:** The user account password.
- **User DN Template:** A template on the LDAP server. When a user logs into their system, the provided username is inserted into this template to determine the user's LDAP distinguished name (DN).
- **Unique ID:** The unique identifier for the specified LDAP server.
- **Member Name Attribute:** The name of the attribute that each group on the identity provider uses to list its members.
- **Search Base DN:** The distinguished name (DN) of the identity provider location where you want your system to begin searching for users and groups.

For example, if you specify a value of OU=Users,DC=corp,DC=example,DC=com, your system searches for users and groups in the organization unit called Users in the corp.example.com domain.
- **Group Object Class:** The objectClass value for groups on the LDAP server.

OpenLDAP and 389 Directory Server

- Identity Provider Hostname: Hostname or IP address for the identity provider.
- Transport Security: The protocol to use for securing communications between the system and the identity provider. Options are:
 - None
 - TLS Security (Transport Layer Security)
 - SSL (Secure Sockets Layer)
- Identity Provider Host Port: Network port used to communicate with the identity provider. The default value depends on the Transport Security setting:
 - For None or TLS Security (Transport Layer Security), 389
 - For SSL (Secure Sockets Layer), 636
- User Name: A user account on the identity provider. Your system uses this account to read information from the identity provider.
- Password: The user account password.
- User DN Template: A template on the LDAP server. When a user logs into their system, the provided username is inserted into this template to determine the user's LDAP distinguished name (DN).
- Unique ID: The unique identifier for the specified LDAP server.
- Member Name Attribute: The name of the attribute that each group on the identity provider uses to list its members.
- Search Base DN: The distinguished name (DN) of the identity provider location where you want your system to begin searching for users and groups.

For example, if you specify a value of OU=Users,DC=corp,DC=example,DC=com, your system searches for users and groups in the organization unit called Users in the corp.example.com domain.

User information caching

The system caches the following information from each of your identity providers:

- The names of users who access the system.
- The groups that each user belongs to.

As long as this information is in the system's cache, your users can perform any activities for which they have permissions, without the system needing to reconnect to the identity provider.

LDAP user information remains in the cache for four hours.

Clearing the cache

Any changes that you make on the identity provider are not reflected in the system until the information is removed from the cache. For example, if you add a user to an LDAP identity provider, that user cannot access the system for up to four hours, or until the cache is cleared. If you delete a user from an LDAP identity provider, that user will be able to access the system for up to four hours, or until the cache is cleared.

To ensure that a change is reflected immediately, use the `clearCache` command or API.

Related CLI commands

`clearCache`

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /security/clearCache

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Viewing identity providers

You can use the Admin App, REST API, and CLI to view the identity providers that have been added to your system.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. Select the **Identity Providers** tab.

Related CLI commands

`getIdentityProvider`

`listIdentityProviders`

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

GET /security/identityProviders/{uuid}

GET /security/identityProviders


You can get help on specific REST API methods for the Admin App at REST API - Admin.

Deleting identity providers

When you delete an identity provider from your system, all users from that provider lose access to the system.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. On the **Identity Providers** tab, click the delete icon () for the server you want to remove.

Related CLI commands

```
deleteIdentityProvider
```

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

DELETE /security/identityProviders/{uuid}

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Groups

To allow user access to your system, you need to add groups to your system. These groups are defined on your organization's identity providers. After you add a group to your system, you can specify what roles its members have.

For information on:

- Adding identity providers to your system, see [Adding identity providers \(on page 121\)](#).
- Roles, see [Roles \(on page 129\)](#).

Adding groups

You use the REST API, Admin App, or CLI to add groups from your identity providers to your system.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.

3. On the **Groups** tab, click **Create**.
4. Select an identity provider and type a string on which to query the identity provider for groups.
5. Click **Discover Groups**.
6. Click **Continue**.
7. Select one or more roles to associate with the group.
8. Click **Create**.

Related CLI commands

createGroup

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /security/groups

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Viewing groups

You use the REST API, CLI, or Admin App to view all the groups that have been created for your system.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. Select the **Groups** tab.

Related CLI commands

getGroup

listGroups

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

GET /security/groups/{uuid}

GET /security/groups

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Assigning roles to groups

You use the REST API, Admin App, and CLI to assign roles to the groups that you've added your system.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. On the **Groups** tab, select the group you want to edit.
4. On the **Roles** tab, select one or more roles to enable for the group.
5. Click **Update**.

Related CLI commands

editGroup

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

PUT /security/groups/{uuid}


You can get help on specific REST API methods for the Admin App at REST API - Admin.

Deleting groups

When you delete a group, all users in the group lose access to your system.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. Select the **Groups** tab.
4. Click the delete icon () for the group you want to remove.

Related CLI commands

deleteGroup

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

DELETE /security/groups/{uuid}

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Roles

Roles determine what actions a group of users can perform. You create your own roles, each of which can grant permission to perform any combination of actions.

For information on associating a role with a group of users, see [Assigning roles to groups \(on page 128\)](#).

Creating roles

You can use the REST API, Admin App, and CLI to create roles and select which permissions the roles contain.

About permissions

Each permission in a role grants a user the ability to perform an action in some area of the system. For example, the `admin:services:read` permission grants the ability to view services through the Admin App.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. On the **Roles** tab, click **Create**.
4. Specify a name and, optionally, a description for the role.
5. Use the **Individual** and **Wildcard** tabs to edit the permissions for the role.
On the **Individual** tab, you can enable individual permissions or categories of permissions:

- Select a category of permissions and select one or more individual permissions within the category.

For example, with the permissions selected in this image, a user can read, create, and update certificates, but cannot delete them.

Permissions Group - Certificates			BACK
<input checked="" type="checkbox"/>	Name ↑	Description	
<input checked="" type="checkbox"/>	workflowcertificatescreate	Create certificate	
<input type="checkbox"/>	workflowcertificatesdelete	Delete certificates	
<input checked="" type="checkbox"/>	workflowcertificatesread	Read certificate(s)	
<input checked="" type="checkbox"/>	workflowcertificatesupdate	Generate certificate	

- On the **Wildcard** tab, you can enable permissions for multiple categories at the same time. To do this:
 - a. Click **Add Permission**.
 - b. Use the menus to select a category of permissions.
 - c. Leave the last menu set to the wildcard character (*).
6. Click **Create**.

7. Click **Update**.

Related CLI commands

createRole

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /security/roles

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Viewing roles

You can use the REST API, CLI, and Admin App to view all the roles that have been created for your system.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. Select the **Roles** tab.

Related CLI commands

getRole

listRoles

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

GET /security/roles/{uuid}

GET /security/roles

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Editing roles

You can use the REST API, Admin App, and CLI to change the permissions that a role contains.

About permissions

Each permission in a role grants a user the ability to perform an action in some area of the system. For example, the admin:services:read permission grants the ability to view services through the Admin App.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. Select the **Roles** tab.
4. Select the role you want to edit.
5. Use the **Individual** and **Wildcard** tabs to edit the permissions for the role.

On the **Individual** tab, you can enable individual permissions or categories of permissions:

- Select a category of permissions and select one or more individual permissions within the category.

For example, with the permissions selected in this image, a user can read, create, and update certificates, but cannot delete them.



Permissions Group - Certificates		BACK
<input checked="" type="checkbox"/>	Name ↑	Description
<input checked="" type="checkbox"/>	workflowcertificatescreate	Create certificate
<input type="checkbox"/>	workflowcertificatesdelete	Delete certificates
<input checked="" type="checkbox"/>	workflowcertificatesread	Read certificate(s)
<input checked="" type="checkbox"/>	workflowcertificatesupdate	Generate certificate

- On the **Wildcard** tab, you can enable permissions for multiple categories at the same time. To do this:
 - a. Click **Add Permission**.
 - b. Use the menus to select a category of permissions.
 - c. Leave the last menu set to the wildcard character (*).
- 6. Click **Create**.
- 7. Click **Update**.

Related CLI commands

editRole

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

PUT /security/roles/{uuid}


You can get help on specific REST API methods for the Admin App at REST API - Admin.

Deleting roles

When you delete a role, all groups associated with that role lose the permissions that the role granted.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. Select the **Roles** tab.
4. Click the delete icon () for the role you want to remove.

Related CLI commands

deleteRole

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

DELETE /security/roles/{uuid}

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Permissions

The following tables list the HCP for cloud scale permissions available for system roles. The words *Yes* and *No* indicate whether or not the permission is assigned for a default role.

Data Service		
Permission name	Description	Default admin role permission?
data:bucket:notification:get	View bucket notification configuration	Yes
data:bucket:notification:set	Configure bucket notification	Yes
data:bucket:sync:from:set	Create bucket sync-from rules for buckets the user owns or has access to	Yes
data:bucket:sync:get	View bucket sync-from and sync-to rules for buckets the user owns or has access to	Yes
data:bucket:sync:to:set	Create bucket sync-to rules for buckets the user owns or has access to	Yes

License		
Permission name	Description	Default admin role permission?
mapi:license:add	Add licensed feature	Yes
mapi:license:list	List all licensed feature	Yes

MAPI Alerts		
Permission name	Description	Default admin role permission?
mapi:alert:list	List all active alerts	Yes

MAPI Job Configurations		
Permission name	Description	Default admin role permission?
mapi:job_configuration:list	List all job configurations	Yes
mapi:job_configuration:run	Run a job configuration immediately	Yes
mapi:job_configuration:update	Modify a job configuration	Yes

MAPI S3 Settings		
Permission name	Description	Default admin role permission?
mapi:s3_settings:get	Read S3 settings	Yes
mapi:s3_settings:set	Modify S3 settings	Yes

MAPI Storage Component		
Permission name	Description	Default admin role permission?
mapi:storage_component:activate	Activate a storage component	Yes
mapi:storage_component:create	Create a storage component	Yes
mapi:storage_component:list	List storage component(s)	Yes
mapi:storage_component:test	Test a storage component	Yes
mapi:storage_component:update	Modify a storage component	Yes
mapi:storage_component:update_state	Modify state of a storage component	Yes

MAPI Stored Objects		
Permission name	Description	Default admin role permission?
mapi:client_object:lookup	List stored objects	Yes

MAPI System		
Permission name	Description	Default admin role permission?
mapi:certificates:refresh	Refresh SSL certificates	
mapi:system:info	List system information	Yes

MAPI User		
Permission name	Description	Default admin role permission?
mapi:user:list	List existing users	Yes

MAPI User		
Permission name	Description	Default admin role permission?
mapi:user:list_buckets	List user's buckets	Yes
mapi:user:revoke_credentials	Revoke S3 credentials	Yes
mapi:user:revoke_tokens	Revoke OAuth tokens	Yes

S3 Encryption Setting		
Permission name	Description	Default admin role permission?
mapi:s3_encryption:get	Read S3 encryption settings	Yes
mapi:s3_encryption:set	Enable global encryption	Yes

S3 User		
Permission name	Description	Default admin role permission?
s3:user:generate_credentials	Generate S3 credentials	Yes


Serial Number		
Permission name	Description	Default admin role permission?
mapi:serial_number:get	Read serial number	Yes
mapi:serial_number:set	Modify serial number	Yes

Changing the admin account password

Your system includes a single local user account called admin, which is available when you first install the system. You can use the REST API, Admin App, or CLI to change the password for this account.

Admin App instructions

Procedure

1. When logged into the Admin App with the admin user account, click the user icon () in the top right corner of the window.
2. Click **Change Password**.
3. Confirm your current password and specify a new password.
4. Click **Change Password**.

Related CLI commands

updateCurrentUserPassword

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /setup/password

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Certificates - HCPCS

Your system uses SSL to provide security for the System Management application and the S3 Console application. To enable SSL security, you need a valid SSL server certificate or chain of certificates.

Your system comes with its own self-signed SSL server certificate, which is generated and installed automatically when the system is installed. This certificate is not automatically trusted by web browsers.

You can choose to trust this self-signed certificate to replace the certificate with either one from a certificate authority (CA) or one that you create yourself. You can also have the system generate and install a new self-signed SSL server certificate. You do this, for example, if the current certificate is close to expiring and you are waiting to retrieve a new one from your CA.

Viewing installed certificates

You can use the REST API, CLI, and Admin App to view information about:

- The system certificate. This is the certificate used to secure communications for your system's applications, CLIs, and REST APIs.
- Data source certificates. These are the certificates retrieved from the systems that your system has connected to using a data connection.

For each certificate, you can view:

- The distinguished name of the certificate
- The date and time when the certificate goes (or went) into effect
- The date and time when the certificate expires (or expired)

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Certificates**.
The **System** tab displays the currently active system certificate.
3. To view the data source certificates, select the **Client** tab.

Related CLI commands

listCertificates

getCertificate

getSystemCertificate

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

GET /certificates

GET /certificates/system

GET /certificates/{subjectDn}

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Adding data source certificates

For your system to retrieve documents from a data source that uses SSL-protected communication, it must accept the certificate from the data source. Your system prompts you to accept a data source certificate when testing the connection to the data source. You can also upload data source certificates manually.

Uploading data source certificates manually

Admin App instructions

Procedure

1. Retrieve the SSL certificate from your data source.
2. In the Admin App, click **Configuration**.
3. Click **Certificates**.
4. On the **Client** tab, click **Upload Client Certificate**.
5. Drag the certificate file into the **Upload Certificate** box.

Related CLI commands

testDataSource

createCertificate

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /datasources/test

POST /certificates

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Changing the system certificate

By default, your system includes a self-signed certificate when the system is first installed.

You cannot delete the currently installed certificate. However, you can replace the certificate by:

- Installing a new PKCS12 certificate (for instructions, see [Installing a certificate you created \(on page 139\)](#)).
- Generating and installing a new self-signed certificate (for instructions, see [Installing a new self-signed certificate \(on page 140\)](#)).
- Generating a certificate signing request (CSR) and installing the certificate you receive in response to this request (for instructions, see [Creating a CSR and installing the returned certificate \(on page 141\)](#)).

System certificate considerations

Keep the following in mind when configuring SSL certificates for your system, especially if you are configuring the system to use one or more certificates that you create yourself:

- Do not allow any of the SSL certificates to expire.
- Adhere to the established best practices for setting up SSL certificates. For example, if you are using wildcards to identify hostnames in an SSL certificate, a wildcard should appear only at the beginning of the hostname, not in the middle.

For information on SSL best practices, see <http://tools.ietf.org/html/rfc5280> and <http://tools.ietf.org/html/rfc6125>.

- Ensure that the DNS name for the system matches the name defined in the certificate. If you rename the system you need a new certificate.
- When configuring a certificate chain, ensure that all intermediate issuers have the appropriate signing authority permissions so that the entire chain is signed.
- If you regenerate or upload an SSL certificate, you must repair (that is, restart) the S3 Gateway, Key Management Server, and MAPI services for the change to take effect.
- If you upload an SSL certificate for access to a remote system for bucket synchronization, you must repair the Policy Engine and MAPI services for the change to take effect.

Installing a certificate you created

You can create an SSL server certificate by using a third-party tool such as OpenSSL. When creating the certificate, you specify two passwords, one for the PKCS12 object containing the certificate and one for the private key for the certificate. To use the certificate with your system, these passwords must be the same.

When you create your own SSL server certificate, you can choose to have that certificate signed by a certificate authority (CA). In this case, the CA you use may provide you with one or more intermediate certificates. These certificates are used in conjunction with the SSL server certificate you created to establish a *certificate chain*, which is an ordered list of certificates in which each certificate is trusted by the next.

To preserve the chain of trust among the certificates, you need to upload the certificates in the correct order. That is, each certificate you upload must be immediately followed by the certificate that signs it. For information on the correct order for the certificate chain, see your CA.



Important: Read and understand the topic [System certificate considerations \(on page 138\)](#) before creating your own SSL certificates, especially if you are using an in-house CA.

To install your certificates:

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Certificates**.

3. Click **Update System Certificate**.
4. On the **PKCS12** window, drag your certificate into the **Upload Certificate Chain** box.
5. In the **PKCS12 Password** field, type the password for your certificate.
6. Drag the certificate into the **Upload Certificate Chain** box.
7. Click **Continue**.
8. Click **Accept**.

Related CLI commands

uploadPKCS12Certificate

applyCertificateChanges

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /certificates/system/pkcs12

POST /certificates/system

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Installing a new self-signed certificate

Your system can generate and install a new self-signed SSL server certificate. The new certificate is good for five years.



Note:

- If the system is using a self-signed certificate, when you change the hostname name of the system, you need to generate a new SSL certificate.
- When you install the new certificate, if you are using encryption, you must repair the Vault service and then unseal the vault.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Certificates**.
3. Click **Update System Certificate**.
4. Select the **Self-Signed** window.
5. Click **Continue**.
Your system generates a new self-signed server certificate.
6. Click **Accept**.
Your system installs the new certificate.
7. To continue using the Admin App, log out and then log back in.

Related CLI commands

generateSelfSignedCertificate

applyCertificateChanges

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /certificates/system/selfsigned

POST /certificates/system

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Creating a CSR and installing the returned certificate

SSL server certificates are available from several trusted sources. To obtain a certificate created by a certificate authority (CA), you need to create a certificate signing request (CSR) and give it to the CA. The CA then generates the requested certificate and makes it available to you.

Creating a certificate signing request

You can create a CSR using the Admin App or a third-party tool. When you use the Admin App, the system securely stores the private key needed for installing the returned certificate, so you don't need to save the key yourself.

To know exactly what information is required, check with the certificate authority (CA) that you plan to use.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Certificates**.
3. Select the **System** tab.
4. Click **Update System Certificate**.
5. Select the **CSR** window.
6. Choose **Generate a new certificate signing request** and click **Continue**.
7. Fill in the fields as needed:
 - In the field **Common Name (CN)**, type the DNS name of the system preceded by an asterisk (*) and a period (.) (for example, *.system.example.com).
The field **Common Name (CN)** is required.
 - In the field **Organizational Unit (OU)**, type the name of the organizational unit that uses the system (for example, the name of a division or a name under which your company does business).
 - In the field **Organization (O)**, type the full legal name of your organization.

- In the field **Location (L)**, type the name of the city in which your organization's headquarters are located.
 - In the field **State/Province (ST)**, type the full name of the state or province in which your organization's headquarters are located.
 - In the field **Country (C)**, type the two-letter ISO 3166-1 abbreviation for the country in which your organization's headquarters are located (for example, US for the United States).
 - In the field **Subject Alternate Name**, type `*.hcpcs_cluster_name`, where `hcpcs_cluster_name` is the name of your HCP for cloud scale cluster.
8. Click **Generate CSR**.
The page displays the generated certificate request.
 9. Copy and paste the request text into a file and send that file to your CA.
 10. Continue to [Installing the certificates returned for a system-generated CSR \(on page 143\)](#).

Creating a certificate signing request

You can create a CSR using the Admin App or a third-party tool. When you use the Admin App, the system securely stores the private key needed for installing the returned certificate, so you don't need to save the key yourself.

To know exactly what information is required, check with the CA you plan to use.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Certificates**.
3. Select the **System** tab.
4. Click **Update System Certificate**.
5. Select the **CSR** window.
6. Choose **Generate a new certificate signing request** and click **Continue**.
7. Fill in the fields as needed:
 - In the **Common Name (CN)** field, type the DNS name of the system preceded by an asterisk (*) and a period (.) (for example, *.system.example.com).

The **Common Name (CN)** field is required.

 - In the **Organizational Unit (OU)** field, type the name of the organizational unit that uses the system (for example, the name of a division or a name under which your company does business).
 - In the **Organization (O)** field, type the full legal name of your organization.
 - In the **Location (L)** field, type the name of the city in which your organization's headquarters are located.

- In the **State/Province (ST)** field, type the full name of the state or province in which your organization's headquarters are located.
 - In the **Country (C)** field, type the two-letter ISO 3166-1 abbreviation for the country in which your organization's headquarters are located (for example, US for the United States).
8. Click **Generate CSR**.
The page displays the generated certificate request.
 9. Copy and paste the request text into a file and send that file to your CA.
 10. Continue to [Installing the certificates returned for a system-generated CSR \(on page 143\)](#).

Related CLI commands

generateCSR

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

PUT /certificates/system/csr

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Installing the certificates returned for a system-generated CSR

In response to a CSR, your CA gives you an SSL server certificate and any required *intermediate certificates*. These certificates are used in conjunction with the SSL server certificate to establish a *certificate chain*, an ordered list of certificates in which each certificate is trusted by the next. You need to upload and install these certificates on your system.

To preserve the chain of trust among the certificates, you need to upload the certificates in the correct order. That is, each certificate you upload must be immediately followed by the certificate that signs it. For information on the correct order for the certificate chain, see your CA.

Admin App instructions**Procedure**

1. Select the **Configuration** window.
2. Click **Certificates**.
3. Select the **System** tab.
4. Click **Update System Certificate**.
5. Select the **CSR** window.
6. Select the **I already generated a CSR and obtained a signed certificate** option and then click **Continue**.
7. Drag the certificate into the **Upload certificate obtained from Certificate Authority** box.
8. Click **Accept**.

Related CLI commands

uploadCSR

applyCertificateChanges

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /certificates/system/csr

POST /certificates/system

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Chapter 6: System management

As an administrator, you play a role in ensuring the continued accessibility and performance of the system. You can use the Admin App, command line, or REST API to manage the system.

Your responsibilities for administering the system include:

- Managing and monitoring system performance and resource usage by configuring how instances are deployed in your infrastructure.
- Expanding functionality by writing and installing plugins.
- Setting up email notifications.
- Upgrading the system.

Setting the system hostname

After installing your system, you need to configure it with the hostname that you've assigned to it in your corporate DNS environment.

Admin App instructions

Procedure

1. Select the **Configuration** window.
2. Click **Security**.
3. On the **Settings** tab, specify the hostname in the **Cluster Hostname** field.
4. Click **Update**.

Changing system hostname

If you change the cluster name (cluster hostname), you must restart the S3 Gateway services for the change to take effect.

For information about restarting a service, see [Repairing services \(on page 92\)](#).

Related CLI commands

`editSecuritySettings`

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

PUT /security/settings

You can get help on specific REST API methods for the Admin App at REST API - Admin.

System scaling

You manage how the system scales by adding or removing instances to the system and also by specifying which services run on those instances.

Instances

An *instance* is a server or virtual machine on which the software is running. A system can have either a single instance or multiple instances. Multi-instance systems have a minimum of four instances.

A system with multiple instances maintains higher availability if instances fail. Additionally, a system with more instances can run tasks concurrently and can typically process tasks faster than a system with fewer or only one instance.

A multi-instance system has two types of instances: *master instances*, which run an essential set of services, and non-master instances, which are called *workers*.

Services

Each instance runs a configurable set of services, each of which performs a specific function. For example, the Metadata Gateway service stores metadata persistently.

In a single-instance system, that instance runs all services. In a multi-instance system, services can be distributed across all instances.

Networking

This topic describes the network usage by, and requirements for, both system instances and services.



Note:

- You can configure the network settings for each service when you install the system. You cannot change these settings after the system is up and running.
- If your networking environment changes such that the system can no longer function with its current networking configuration, you need to reinstall the system.

Site hostname

The HCP for cloud scale site hostname is configured during installation. This hostname is very important because it's needed for access to:

- The HCP for cloud scale user interface
- The S3 API

Instance IP address requirements

All instance IP addresses must be static. This includes both internal and external network IP addresses, if applicable to your system.



Important: If the IP address of any instance changes, you must reinstall the system.

Network types

Each of the HCP for cloud scale services can bind to one type of network, either **internal** or **external**, for receiving incoming traffic. If your network infrastructure supports having two networks, you might want to isolate the traffic for most system services to a secured internal network that has limited access. You can then leave the following services on your external network for user access:

- Admin-App
- Message Queue
- Metadata-Cache
- Metadata-Coordination
- Metadata-Gateway
- Policy-Engine
- Metrics
- S3-Gateway
- Tracing-Agent
- Tracing-Collector
- Tracing-Query
- MAPI-Gateway

You can use either a single network type for all services or a mix of both types. To use both types, every instance in your system must be addressable by two IP addresses, one on your internal network and one on your external network. If you use only one network type, each instance needs only one IP address.

Allowing access to external resources

Regardless of whether you're using a single network type or a mix of types, you need to configure your network environment to ensure that all instances have outgoing access to the external resources you want to use, such as:

- The storage components where your object data is stored
- Identity providers for user authentication
- Email servers that you want to use for sending email notifications

Ports

Each service binds to a number of ports for receiving incoming traffic.

Before installing HCP for cloud scale, you can configure services to use different ports, or use the default values shown in the following tables.

The following services must be deployed with their default port values:

- Message Queue
- Metadata-Cache
- Tracing Agent
- Tracing Collector
- Tracing Query

External ports

The following table contains information about the service ports that users use to interact with the system.

On every instance in the system, each of these ports:

- Must be accessible from any network that needs administrative or data access to the system
- Must be accessible from every other instance in the system

Default Port Value	Used by Service	Purpose
80 (S3 HTTP port, if enabled)	S3 Gateway	Object persistence and access
443 (S3 HTTPS port)	S3 Gateway	Object persistence and access
8000	Admin App	System Management application GUI
9099	MAPI Gateway	Object Storage Management application GUI

Load balancing

The supported options for load balancing S3 traffic affect performance.

The S3 Gateway service processes S3 traffic and can serve as an SSL termination point. It can listen on port 80, port 443 (the standard SSL port) or port 8443. The Network Proxy service balances the flow of S3 traffic to S3 Gateway instances. The Network Proxy service listens only on port 443. By default, Network Proxy passes S3 SSL traffic through to the S3 Gateway service.

To improve performance, you can configure an external load balancer and bypass Network Proxy. If your load balancer supports SSL termination, you can configure S3 Gateway instances to accept HTTP traffic on port 80.

If you want your load balancer to pass through SSL S3 traffic and your firewall rules permit traffic on port 8443, configure your load balancer to point to port 8443.

If you want your load balancer to pass through SSL S3 traffic but your firewall rules block traffic on port 8443, you can use IP tables to redirect the traffic from port 8443 to port 443.

HCP for cloud scale provides scripts to enable and disable IPtable redirection of S3 traffic. An additional script lists the IP addresses of affected instances. For more information, see [Script to enable S3 traffic redirection \(on page 149\)](#), [Script to disable S3 traffic redirection \(on page 150\)](#), and [Script to list S3 traffic redirection \(on page 151\)](#).

Script to enable S3 traffic redirection

A script is included to enable redirection of S3 traffic from port 443 to port 8443.

The script is written in Python and located in the folder `install_path/product/bin` (for example, `/opt/hcps/bin`).

The script redirects S3 traffic from port 443 to port 8443 using the file `iptables`.

Syntax

```
enable_s3_redirect.py
```

Options and parameters

None

Example

```
$ enable_s3_redirect.py
```

This example can produce the following output:

```
*** PREROUTING chain in NAT table before adding Redirect
Chain PREROUTING (policy ACCEPT 135 packets, 8100 bytes)
num  pkts bytes target    prot opt in     out     source
destination
1      14M  845M DOCKER    all  --  *      *       0.0.0.0/0
0.0.0.0/0          ADDRTYPE match dst-type LOCAL
-----
*** PREROUTING chain in NAT table after adding Redirect
```

```
Chain PREROUTING (policy ACCEPT 0 packets, 0 bytes)
num  pkts bytes target    prot opt in     out     source
destination
1      14M  845M DOCKER     all  --  *      *       0.0.0.0/0
0.0.0.0/0          ADDRTYPE match dst-type LOCAL
2        0      0 REDIRECT  tcp  --  *      *       0.0.0.0/0
0.0.0.0/0          tcp dpt:443 redir ports 8443
```

Script to disable S3 traffic redirection

A script is included to disable redirection of S3 traffic from port 443 to port 8443.

The script is written in Python and located in the folder *install_path/product/bin* (for example, */opt/hcps/bin*).

The script redirects S3 traffic from port 8443 to port 443 using the file *iptables*.

Syntax

```
disable_s3_redirect.py
```

Options and parameters

None

Example

```
$ disable_s3_redirect.py
```

This example can produce the following output:

```
*** PREROUTING chain in NAT table before deleting Redirect
Chain PREROUTING (policy ACCEPT 3227 packets, 195K bytes)
num  pkts bytes target    prot opt in     out     source
destination
1      14M  845M DOCKER     all  --  *      *       0.0.0.0/0
0.0.0.0/0          ADDRTYPE match dst-type LOCAL
2        0      0 REDIRECT  tcp  --  *      *       0.0.0.0/0
0.0.0.0/0          tcp dpt:443 redir ports 8443
```

```
-----
*** PREROUTING chain in NAT table after deleting Redirect
Chain PREROUTING (policy ACCEPT 0 packets, 0 bytes)
num  pkts bytes target    prot opt in     out     source
destination
1      14M  845M DOCKER     all  --  *      *       0.0.0.0/0
0.0.0.0/0          ADDRTYPE match dst-type LOCAL
```

Script to list S3 traffic redirection

A script is included to list redirection of S3 traffic.

The script is written in Python and located in the folder *install_path/product/bin* (for example, */opt/hcpcs/bin*).

The script lists the instances running S3 Gateway system instances and writes the output to the file *s3NodeIPs.txt*.

Syntax

```
list_s3_node_ips.py username password
```

Options and parameters

username

User name of an HCP for cloud scale user with administrative privileges.

password

Password for the administrative user.

Example

```
$ list_s3_node_ips.py username password
```

This example can produce the following output:

```
INSTALL_DIR: /opt/hcpcs
ADMIN_CLI: /opt/hcpcs/cli/admin/admincli
IPs of S3 Gateway nodes:
172.10.24.195
172.10.24.196
172.10.24.197
172.10.24.198
172.10.24.199
Output file is located at /opt/hcpcs/s3NodeIPs.txt
```

Handling network changes

After your system is deployed, its network infrastructure and configuration should not change. Specifically:

- All instance IP addresses should not change.
- All services should continue to use the same ports.
- All services and instances should continue to use the same network types.

If any of these things change, you will need to reinstall the system.

Safely changing an instance IP address

If you need to change the IP addresses for one or more instances in the system, use this procedure to manually change the IP addresses without risk of data loss.

For each instance whose IP address you need to change:

Procedure

1. Move all services off of the instance. Distribute those services among all the other instances.
2. On the instance from step 1, stop the `run` script using whatever tool or process you used to run it. For example, with `systemd`, run:
`systemctl stop <service-name>`
3. Remove the instance from the system.
4. Delete the installation folder from the instance.
5. Add the instance back to the system.

After a network change

If a network infrastructure or configuration change occurs that prevents your system from functioning with its current network settings, you need to reinstall all instances in the system.

Procedure

1. If the Admin App is accessible, back up your system components by exporting a package.
2. On each instance in the system:
 - a. Navigate to the installation folder.
 - b. Stop the run script using whatever tool or process you used to run it. For example, with `systemd`, run:
`systemctl stop <service-name>`
 - c. Run `bin/stop`
 - d. Run the setup script, including the list of master instances:

```
sudo bin/setup -i <ip-address-for-this-instance> -m
<comma-separated-list-of-master-instance-IP-addresses>
```

- e. Run the run script using whatever methods you usually use to run scripts.
3. Log into Admin App and use the wizard to set up the system.
4. After the system has been set up, upload your package.

Volumes

Volumes are properties of services that specify where and how a service stores its data.

You can use volumes to configure services to store their data in external storage systems, outside of the system instances. This allows data to be more easily backed up or migrated.


Volumes can also allow services to store different types of data in different locations. For example, a service might use two separate volumes, one for storing its logs and the other for storing all other data.



Note: Some functions described here are not used with HCP for cloud scale. They are not visible in the System Management application, or have no effect when used.

Example

In this example, service A runs on instance 101. The service's Log volume stores data in a folder on the system instance and the service's Data volume stores data in an NFS mount.



Creating and managing volumes

Volumes are separated into these groups, depending on how they are created and managed:

- *System-managed volumes* are created and managed by the system. When you deploy the system, you can specify the volume driver and options that the system should use when creating these volumes.

After the system is deployed, you cannot change the configuration settings for these volumes.

- *User-managed volumes* can be added to services and job types after the system has been deployed. These are volumes that you manage; you need to create them on your system instances before you can configure a service or job to use them.



Note: As of v1.3.0, none of the built-in services support adding user-managed volumes.

Volume drivers

When configuring a volume, you specify the volume driver that it uses. The volume driver determines how and where data is stored.

Because services run in Docker containers on instances in the system, volume drivers are provided by Docker and other third-party developers, not by the system itself. For information about volume drivers you can use, see the applicable Docker or third-party developer's documentation.

By default, all services do not use volume drivers but instead use the *bind-mount* setting. With this setting, data for each service is stored within the system installation folder on each instance where the service runs.

For more information on volume drivers, see the Docker documentation.

Viewing volumes

The System Management application shows this information about the Docker volumes used by jobs and services:

- Name: The unique identifier for the volume.
- Type: Either of these:
 - System: The volume is managed automatically for you by the system.
 - User: You need to manage the volume yourself.
- Capacity: Total storage space available in the volume.
- Used: Space used by the job or service.
- Pool: The volume category, as defined by the service or job that uses the volume.



Note: Some functions described here are not used with HCP for cloud scale. They are not visible in the System Management application, or have no effect when used.

For each volume, you can also view this information about the volume driver that controls how the volume stores data:

- Volume driver: The name of the volume driver.
- Option/Value: The command-line options used to create the volume and their corresponding values. The available options and valid values for those options are determined by the volume driver.

Viewing job volumes

To view the volumes being used by a job:

Procedure

1. In the Admin App, click the **Jobs** window.
2. On the **Job Type** page, click the job you want.
3. Click the **Volumes** tab.

Viewing service volumes

To view the volumes being used by a service:

Procedure

1. In the Admin App, click the **Services** window.
2. Click the service you want.

3. Click the **Volumes** tab.

Instances

A system is made up of one or more instances of the software. This section includes information on adding and removing instances to the system.

About master and worker instances

Master instances are special instances that run an essential set of services, including:

- Admin-App service
- Cluster-Coordination service
- Synchronization service
- Service-Deployment service

Non-master instances are called workers. Workers can run any services except for those listed previously.

Single-instance systems have one master instance while multi-instance systems have either one or three master instances.



Important: You cannot add master instances to a system after it's installed. You can, however, add any number of worker instances.

Single-instance systems versus multi-instance systems

A system can have a single instance or can have multiple instances (four or more).



Note: Every instance must meet the minimum RAM, CPU, and disk space requirements.

Single instance

A single-instance system is useful for testing and demonstration purposes. A single-instance system requires only a single server or virtual machine and can perform all product functionality.

However, a single-instance system has these drawbacks:

- It has a single point of failure. If the instance hardware fails, you lose access to the system.
- With no additional instances, you cannot choose where to run services. All services run on the single instance.

Multiple instances

A multi-instance system is suitable for use in a production environment because it offers these advantages over a single-instance system:

- You can control how services are distributed across the multiple instances, providing improved service redundancy, scale out, and availability.
- A multi-instance system can survive instance outages. For example, with a four-instance system running the default distribution of services, the system can lose one instance and still remain available.
- Performance is improved as work can be performed in parallel across instances.
- You can add additional instances to the system at any time.



Note: You cannot change a single-instance system into a production-ready multi-instance system by adding new instances. This is because you cannot add master instances. Master instances are special instances that run a particular set of HCP for cloud scale services. Single-instance systems have one master instance. Multi-instance systems have at least three.

By adding additional instances to a single-instance system, your system still has only one master instance, meaning there is still a single point of failure for the essential services that only a master instance can run.

Three-instance system considerations

Three-instance systems should have only a single master instance. If you deploy a three-instance system where all three instances are masters, the system may not have enough resources to do much beyond running the master services.

Requirements for running system instances

This section lists the hardware and operating system requirements for running system instances.

Hardware requirements

To install HCP for cloud scale on on-premises hardware for production use, you must provision at least four instances (nodes) with sufficient CPU, RAM, disk space, and networking capabilities. This table shows the minimum and best-practice hardware requirements for each instance in an HCP for cloud scale system.

Resource	Minimum	Best
RAM	32 GB	128 GB
CPU	8-core	24-core
Available disk space	500 GB 10k SAS RAID	2000 GB 15k SAS RAID

Resource	Minimum	Best
Network interface controller (NIC)	(1) 10 Gb Ethernet	(2) 10 Gb Ethernet
IP addresses	(1) static	(2) static
Firewall Port Access	Port 443 for S3 API and S3 Console application Port 8000 for System Management application GUI Port 8443 for S3 Gateway API Port 9084 for MAPI and Object Storage Management application GUI	Same
Internal IP Ports	See Networking (on page 146)	Same
Network Time	IP address of time service (NTP)	Same



Important: Each instance uses all available RAM and CPU resources on the server or virtual machine on which it's installed.

Operating system and Docker minimum requirements

Each server or virtual machine you provide must have the following:

- 64-bit Linux distribution
- Docker version 1.13.1 or later installed
- IP and DNS addresses configured

Additionally, you should install all relevant patches on the operating system and perform appropriate security hardening tasks.



Important: The system cannot run with Docker versions before 1.13.1.

Operating system and Docker qualified versions

This table shows the operating system, Docker, and SELinux configurations with which the HCP for cloud scale system has been qualified.



Important: Docker versions after 18.09 are not supported.

Operating system	Docker version	Docker storage configuration	SELinux setting
Fedora 31	Docker 18.09.0-ce	direct-lvm	Enforcing
Red Hat Enterprise Linux 7.4	Docker 18.09.0-ce	direct-lvm	Enforcing
Ubuntu 16.04-LTS	Docker 17.03.0-ce	aufs	N/A
CentOS 7.4	Docker 18.09.1-ce	overlay2	Enforcing

Docker considerations

The Docker installation folder on each instance must have at least 20 GB available for storing the Docker images.

Make sure that the Docker storage driver is configured correctly on each instance before installing the product. After you install the product, to change the Docker storage driver you must reinstall the product. To view the current Docker storage driver on an instance, run:

```
docker info
```

Core dumps can fill a host's file system, which can result in host or container instability. Also, if your system uses the data at rest encryption (DARE) feature, encryption keys are written to the dump file. It's best to disable core dumps.

To enable SELinux on the system instances, you need to use a Docker storage driver that SELinux supports. The storage drivers that SELinux supports differ depending on the Linux distribution you're using. For more information, see the Docker documentation.

If you are using the Docker `devicemapper` storage driver:

- Make sure that there's at least 40 GB of Docker metadata storage space available on each instance. The product needs 20 GB to install successfully and an additional 20 GB to successfully update to a later version.

To view Docker metadata storage usage on an instance, run:

```
docker info
```

- On a production system, do not run `devicemapper` in `loop-lvm` mode. This can cause slow performance or, on certain Linux distributions, the product might not have enough space to run.

SELinux considerations

- You should decide whether you want to run SELinux on system instances and enable or disable it before installing additional software on the instance.

Enabling or disabling SELinux on an instance needs a restart of the instance.

To view whether SELinux is enabled on an instance, run: `sestatus`

- To enable SELinux on the system instances, you need to use a Docker storage driver that SELinux supports.

The storage drivers that SELinux supports differ depending on the Linux distribution you're using. For more information, see the Docker documentation.

Supported browsers

The HCP for cloud scale web applications support these web browsers:

- Google Chrome latest
- Mozilla Firefox latest

Time source

If you are installing a multi-instance system, each instance should run NTP (network time protocol) and use the same external time source. For information, see support.ntp.org.

Adding new instances

You might want to add additional instances to the system if:

- You want to improve system performance.
- You are running out of disk space on one or more instances.



Important: You cannot add new master instances, only new worker instances.

However, these situations might also be improved by adding additional CPU, RAM, or disks to the instances you already have.

Before adding a new instance

- Obtain the product installation file. When adding an instance, you unpack and deploy this file on a bare-metal server or a pre-existing Linux virtual machine.
- Record the IP address of at least one of the master instances in the system.

If your system uses internal and external networks, you need to record both the internal and external IP addresses for the master instances.

You can view instance IP addresses on the **Instances** page in the Admin App.

- Ensure that the new instances you are adding meet the minimum hardware, OS, and networking requirements. For information, see [Requirements for running system instances \(on page 156\)](#).
- Record the Docker volume drivers currently used by services and jobs across all existing instances. You need to install all of these volume drivers on the new instance that you're adding.

To find the volume drivers currently in use by your system, run this command on each system instance:

```
docker volume ls
```

Take note of each value for the DRIVER field.

Install Docker on each server or virtual machine

On each server or virtual machine that is to be an HCP for cloud scale instance:

Procedure

1. In a terminal window, verify whether Docker 1.13.1 or later is installed:
`docker --version`
2. If Docker is not installed or if you have a version before 1.13.1, install the current Docker version suggested by your operating system.
The installation method you use depends on your operating system. See the [Docker website](#) for instructions.

Configure Docker on each server or virtual machine

Before installing the product, configure Docker with settings suitable for your environment. For guidance on configuring and running Docker, see the applicable Docker documentation.

Procedure

1. Ensure that the Docker installation folder on each instance has at least 20 GB available for storing the product Docker images.
2. Ensure that the Docker storage driver is configured correctly on each instance. After installation, changing the Docker storage driver needs reinstallation of the product.
To view the current Docker storage driver on an instance, run: `docker info`.
3. To enable SELinux on the system instances, use a Docker storage driver that SELinux supports.
The storage drivers that SELinux supports differ depending on the Linux distribution you're using. For more information, see the Docker documentation.
4. If you are using the Docker `devicemapper` storage driver, ensure that there's at least 40 GB of Docker metadata storage space available on each instance.
The product needs 20 GB to install successfully and an additional 20 GB to successfully update to a later version.
To view Docker metadata storage usage on an instance, run: `docker info`

Next steps

On a production system, do not run `devicemapper` in `loop-lvm` mode. This can cause slow performance or, on certain Linux distributions, the product might not have enough space to run.

(Optional) Configure Docker volume drivers

If any services or jobs on your system are using Docker volume drivers (that is, not the `bind-mount` setting) for storing data, you need to install those volume drivers on the new instance that you are adding. If you don't, jobs and services might fail to run on the new instance.

Volume drivers are provided by Docker and other third-party developers, not by the system itself. For information on volume drivers, their capabilities, and their valid configuration settings, see the applicable Docker or third-party developer's documentation.

Configure maximum map count setting

You need to configure a value in the file `sysctl.conf`.

Procedure

1. On each server or virtual machine that is to be a system instance, open the file `/etc/sysctl.conf`.
2. Append this line: `vm.max_map_count = 262144`
If the line already exists, ensure that the value is greater than or equal to 262144.
3. Save and close the file.

(Optional) Enable or disable SELinux on each server or virtual machine

You should decide whether you want to run SELinux on system instances before installation.

Procedure

1. Enable or disable SELinux on each instance.
2. Restart the instance.

Configure the firewall rules on each server or virtual machine

Before you begin

Determine the port values currently used by your system. To do this, on any instance, view the file `install_path/config/network.config`.

On each server or virtual machine that is to be a system instance:

Procedure

1. Edit the firewall rules to allow communication over all network ports that you want your system to use. You do this using a firewall management tool such as `firewalld`.
2. Restart the server or virtual machine.

Install and configure NTP

Install NTP (Network Time Protocol) on the new server or virtual machine and configure it to use the same time source as the other system instances. For information, see <http://support.ntp.org>.

Run Docker on each server or virtual machine

On each server or virtual machine that is to be a system instance, you need to start Docker and keep it running. You can use whatever tools you typically use for keeping services running in your environment.

For example, to run Docker using `systemd`:

Procedure

1. Verify that Docker is running:
`systemctl status docker`
2. If Docker is not running, start the `docker` service:
`sudo systemctl start docker`
3. (Optional) Configure the Docker service to start automatically when you restart the server or virtual machine:
`sudo systemctl enable docker`

Unpack the installation package

On each server or virtual machine that is to be a system instance:

Procedure

1. Download the product installation package and MD5 checksum file and store both in a folder on the server or virtual machine.
2. Verify the integrity of the installation package:
`md5sum -c product-version_number.tgz.md5`
If the package integrity is verified, the command displays `OK`.
3. In the largest disk partition on the server or virtual machine, create a product installation folder.
`mkdir install_path/product`
4. Move the installation package from the folder where you stored it to the product installation folder.
`mv product-version_number.tgz install_path/product/product-version_number.tgz`
5. Navigate to the installation folder.
`cd install_path/product`

6. Unpack the installation package:

```
tar -zxvf hcpcs-version_number.tgz
```

A number of directories are created within the installation folder.



Note:

If you encounter problems unpacking the installation file (for example, the error message "tar: This does not look like a tar archive"), the file might have been packed more than one time during download. Use the following commands to fully extract the file:

```
$ gunzip product-version_number.tgz
$ mv product-version_number.tar product-
version_number.tgz
$ tar -zxvf product-version_number.tgz
```

7. Run the installation script `install`, located within a folder matching the version number of system software used by the product software.

```
sudo ./cluster/sys_ver_num/bin/install
```

This version number is different from the product version number. The subfolder is the only one in the folder `cluster`.

For example:

```
sudo ./cluster/1.4.0.260/bin/install
```



Note:

- Don't change directories after running the installation script. The following tasks are performed in your current folder.
- The installation script can be run only one time on each instance. You cannot rerun this script to try to repair or upgrade a system instance.

Set up networking

On each server or virtual machine that is to be a system instance, edit the file `installation-folder/config/network.config` file to be identical to the copies of the same file on the existing system instances.

Run the setup script on each server or virtual machine

Before you begin



Note:

- When installing a multi-instance system, make sure you specify the same list of master instance IP addresses on every instance that you are installing.
- When entering IP address lists, do not separate IP addresses with spaces. For example, the following is correct:

```
sudo install_path/hcpcs/bin/setup -i 192.0.2.4
-m 192.0.2.0,192.0.2.1,192.0.2.3
```

On each server or virtual machine that is to be a system instance:

Procedure

1. Run the script `setup` with the applicable options:

Option	Description
<code>-i</code>	The external network IP address for the instance on which you're running the script.
<code>-I</code>	The internal network IP address for the instance on which you're running the script.
<code>-m</code>	Comma-separated list of external network IP addresses of each master instance.
<code>-M</code>	Comma-separated list of internal network IP addresses of each master instance.

Use this table to determine which options you need to use:

Number of instances in the system	Network type usage	Options to use
Multiple	Single network type for all services	Either: <code>-i</code> and <code>-m</code> or <code>-I</code> and <code>-M</code>
Multiple	Internal for some services, external for others	All of these: <code>-i</code> , <code>-I</code> , <code>-m</code> , <code>-M</code>
Single	Single network type for all services	Either <code>-i</code> or <code>-I</code>
Single	Internal for some services, external for others	Both <code>-i</code> and <code>-I</code>

Result



Note: If the terminal displays Docker errors when you run the `setup` script, ensure that Docker is running.

Example

This example sets up a single-instance system that uses only one network type for all services:

```
sudo install_path/hcpcs/bin/setup -i 192.0.2.4
```

To set up a multi-instance system that uses both internal and external networks, type the command in this format:

```
sudo install_path/hcpcs/bin/setup -i external_instance_ip -I
internal_instance_ip -m external_master_ips_list -M
internal_master_ips_list
```

For example:

```
sudo install_path/hcpcs/bin/setup -i 192.0.2.4 -I 10.236.1.0 -m
192.0.2.0,192.0.2.1,192.0.2.3 -M 10.236.1.1,10.236.1.2,10.236.1.3
```

This table shows sample commands to create a four-instance system. Each command is entered on a different server or virtual machine that is to be a system instance. The resulting system contains three master instances and one worker instance and uses both internal and external networks.

Instance internal IP	Instance external IP	Master or worker	Command
192.0.2.1	10.236.1.1	Master	<code>sudo install_path/hcpcs/bin/setup -I 192.0.2.1 -i 10.236.1.1 -M 192.0.2.1,192.0.2.2,192.0.2.3 -m 10.236.1.1,10.236.1.2,10.236.1.3</code>
192.0.2.2	10.236.1.2	Master	<code>sudo install_path/hcpcs/bin/setup -I 192.0.2.2 -i 10.236.1.2 -M 192.0.2.1,192.0.2.2,192.0.2.3 -m 10.236.1.1,10.236.1.2,10.236.1.3</code>
192.0.2.3	10.236.1.3	Master	<code>sudo install_path/hcpcs/bin/setup -I 192.0.2.3 -i 10.236.1.3 -M 192.0.2.1,192.0.2.2,192.0.2.3 -m 10.236.1.1,10.236.1.2,10.236.1.3</code>
192.0.2.4	10.236.1.4	Worker	<code>sudo install_path/hcpcs/bin/setup -I 192.0.2.4 -i 10.236.1.4 -M 192.0.2.1,192.0.2.2,192.0.2.3 -m 10.236.1.1,10.236.1.2,10.236.1.3</code>

Start the application on each server or virtual machine

On each server or virtual machine that is to be a system instance:

Procedure

1. Start the application script `run` using whatever methods you usually use to run scripts.



Important: Ensure that the method you use can keep the `run` script running and can automatically restart it if a server restarts or there is other availability event.

Result

When the service starts, the server or virtual machine automatically joins the system as a new instance.

Example

Here are some examples of how you can start the script:

- You can run the script in the foreground:

```
sudo install_path/product/bin/run
```

When you run the `run` script this way, the script does not automatically complete, but instead remains running in the foreground.

- You can run the script as a service using `systemd`:

- Copy the product `.service` file to the appropriate location for your OS. For example:

```
cp install_path/product/bin/product.service /etc/systemd/system
```

- Enable and start the `product.service` service:

```
sudo systemctl enable product.service
sudo systemctl start product.service
```

Configure services and jobs on the new instances

The system does not automatically begin running services on the instances you've added. You need to manually configure services to run on those new instances.

Also, depending on how your jobs are configured, jobs might not run on the new instances that you've added. You need to manually configure jobs to run on the instances.

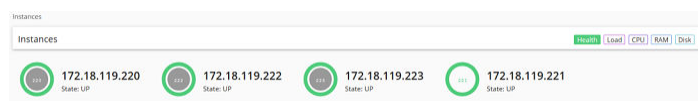
Viewing instances

You can use the Admin App, CLI, and REST API to view a list of all instances in the system.

Viewing all instances


To view all instances, in the Admin App, click Dashboard > Instances.

The page shows all instances in the system. Each instance is identified by its IP address.



IP Address	State
172.18.119.220	UP
172.18.119.222	UP
172.18.119.223	UP
172.18.119.221	UP

This table describes the information shown for each instance.

Property	Description
State	<ul style="list-style-type: none"> ▪ Up: The instance is reachable by other instances in the system. ▪ Down: The instance cannot be reached by other instances in the system.
Services	The number of services running on the instance.
Service Units	<p>The total number of service units for all services and job types running on the instance, out of the best-practice service unit limit for the instance.</p> <p>An instance with a higher number of service units is likely to be more heavily used by the system than an instance with a lower number of service units.</p> <p>The Instances page displays a blue bar for instances running less than the best-practice service unit limit.</p> <p>The Instances page displays a red bar for instances running more than the best-practice service unit limit.</p> 
Load Average	The load averages for the instance for the past one, five, and ten minutes.
CPU	The sum of the percentage utilization for each CPU core in the instance.
Memory Allocated	<p>This section shows both:</p> <ul style="list-style-type: none"> ▪ The amount of RAM on the instance that's allocated to all services running on that instance. ▪ The percentage of this allocated RAM to the total RAM for the instance.
Memory Total	The total amount of RAM for the instance.
Disk Used	The current amount of disk space that your system is using in the partition on which it is installed.
Disk Free	The amount of free disk space in the partition in which your system is installed.

Viewing the services running on an instance

To view the services running on an individual instance, in the Admin App:

Procedure

1. Click **Dashboard > Instances**.

2. Select the instance you want.

The page lists all services running on the instance.

For each service, the page shows:

- The service name
 - The service state:
 - **Healthy:** The service is running normally.
 - **Unconfigured:** The service has yet to be configured and deployed.
 - **Deploying:** The system is currently starting or restarting the service. This can happen when:
 - You move the service to run on a completely different set of instances.
 - You repair a service.
- For information on viewing the status service operations, see [Monitoring service operations \(on page 99\)](#).
- **Balancing:** The service is running normally, but performing background maintenance.
 - **Under-protected:** In a multi-instance system, one or more of the instances on which a service is configured to run are offline.
 - **Failed:** The service is not running or the system cannot communicate with the service.
 - **CPU Usage:** The current percentage CPU usage for the service across all instances on which it's running.
 - **Memory:** The current RAM usage for the service across all instances on which it's running.
 - **Disk Used:** The current total amount of disk space that the service is using across all instances on which it's running.

Related CLI commands

`getInstance`

`listInstances`

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

`GET /instances`

`GET /instances/{uuid}`

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Removing instances

You typically remove an instance from your system in these situations:

- You are retiring the hardware on which the instance runs.
- The instance is in the Down state and cannot be recovered.
- You want to run a system with fewer instances.

(Optional) Shut down the instance you want to remove

If the instance has already shut down because of a failure, the instance is in the Down state. Your system automatically tries to move all services from that instance to other instances in the system. After all services have been moved, the instance is eligible for removal. Continue to the next step ([Remove the shut down instance from the system \(on page 170\)](#)).

If the instance that you want to remove is in the Up state, you need to shut the instance down yourself before you can remove it from the system.

Procedure

1. Move all the services that the instance is currently running to the other instances in the system.



Important: Shutting down an instance without first moving its services can cause data loss.

2. If the system has jobs configured to run on only the failed instance, configure those jobs to run on other instances.
3. Stop the `run` script from running. You do this using whatever method you're currently using to run the script.
4. Run this command to stop all system Docker containers on the instance:


```
sudo <installation-folder>/bin/stop
```
5. Delete the system Docker containers:
 - a. List all Docker containers:


```
sudo docker ps
```
 - b. Note the container IDs for all containers that use a `com.hds.ensemble` or `com.hitachi.aspen` image.
 - c. Delete each of those containers:


```
sudo docker rm <container-id>
```
6. Delete the system Docker images:
 - a. List all Docker images:


```
sudo docker images
```
 - b. Note the image IDs for all images that use a `com.hds.ensemble` or `com.hitachi.aspen` repository.
 - c. Delete each of those images:


```
sudo docker rmi <image-id>
```

7. Delete the system installation folder:

```
rm -rf /<installation-folder>
```

Remove the shut down instance from the system

Admin App instructions

Procedure

1. Select the **Instances** window.
2. Click the instance you want to remove.
3. Click **Remove Instance**.

Related CLI commands

deleteInstance

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

DELETE /instances/{uuid}

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Replacing a failed instance

If an instance suffers an unrecoverable failure, you need to replace that instance with a new one.

Procedure

1. In the Admin App, view the **Instances** page to determine whether the failed instance was a master instance.
2. Select a new server or virtual machine to add as a new instance to the system.
3. Remove the failed instance from the system.



WARNING: If the failed instance was a master, after you remove the instance, you have only two master instances remaining. If any other instance fails while you are in this state, the system becomes completely unavailable until you add a third master back to the system by completing this procedure.

4. Add the replacement instance to the system.



Important: If the instance you are replacing was a master instance, when you run setup on the replacement instance, the list of masters that you specify for the **-m** option needs to include:

- The IP addresses of the two remaining healthy master instances.
- The IP address of the new instance that you're adding.

For example, in a system with master instance IPs ranging from 192.0.2.1 to 192.0.2.3 and you are replacing instance 192.0.2.3 with 192.0.2.5, run setup with these options:

```
sudo bin/setup -i 192.0.2.5 -m
192.0.2.1,192.0.2.2,192.0.2.5
```

This does not apply when you're replacing a worker instance. In that case, specify the IP addresses of the three existing masters.

Plugins

Plugins are modular pieces of code that allow your system to perform specific activities.

Plugins are organized in groups called plugin bundles. When adding or removing plugins from your system, you work with plugin bundles, not individual plugins.

Viewing installed plugins

Use the Admin App, REST API, and CLI to view all plugin bundles and individual plugins that have been installed. You can view all individual plugins at the same time or filter the list based on plugin type.

Admin App instructions

Procedure

1. Click the **Configuration** window.
2. Click **Plugins**.

The **Plugin Bundles** tab shows all installed plugin bundles.

3. To view all individual plugins, click the **All Plugins** tab.

Related CLI commands

listPlugins

Related REST API methods

GET /plugins

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Upgrading plugin bundles

To upgrade plugins, you upload a new version of the bundle that contains those plugins.

You can select which version of the plugin bundle is the active one (that is, the one that connectors or stages will use). If you select the new version, all connectors and stages immediately begin using the new versions of the plugins in the bundle.

You can change the active plugin bundle version at any time.

Admin App instructions

Procedure

1. Click the **Configuration** window.
2. Click **Plugins**.
3. Click **Upload Bundle**.
4. In the **Upload Plugins** window, drag and drop the new version of the plugin bundle.
5. In the list of plugin bundles, click the row for the plugin bundle version that you want.
If the bundle you uploaded isn't listed, click **Reload Plugins**.
6. Click **Set Active**.

Related CLI commands

uploadPlugin

setPluginBundleActive

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /plugins/upload

POST /plugins/bundles/{name}/{bundleVersion}/active

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Setting the active plugin bundle version

If you've uploaded multiple versions of a plugin bundle, only one version can be active at a time. The active plugin bundle version is the one that the system uses.

When you change the active version of a plugin bundle, any workflow tasks that contain connectors and stages that use the bundle immediately begin using the new active version.

Admin App instructions

Procedure

1. Click the **Configuration** window.

2. Click **Plugins**.
3. Click the row for the plugin bundle version that you want.
4. Click **Set Active**.

Related CLI commands

setPluginBundleActive

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /plugins/bundles/{name}/{bundleVersion}/active

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Deleting plugin bundles


To delete plugins from your system, you delete plugin bundles from the system. You cannot delete individual plugins.



Note: You cannot delete a plugin bundle, or any of its versions, if any of that bundle's plugins are currently in use by the system.

Admin App instructions

Procedure

1. Click the **Configuration** window.
2. Click **Plugins**.
3. Click the delete icon () for the plugin bundle you want to remove.

Related CLI commands

deletePluginBundle

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

DELETE /plugins/bundles/{name}/{bundleVersion}

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Packages

You can back up all of your system configuration by exporting packages. You can back up these package files and use them to restore your configurations in the event of a system failure.

Exporting packages

You can export the configurations for system components as package files. You can back up these package files and use them to restore your configurations in the event of a system failure.

After exporting a package, you can store it in one of your data sources. When you want to import the package, your system can retrieve it directly from the data source.

Admin App instructions

Procedure

1. Click the **Configuration** window.
2. Click **Packages**.
3. Click **Export**.
4. Under **Customize Package Description**, give your package a name and an optional description.
5. Under **Configuration**, select any configuration items to export.
6. Under **Plugins**, select any plugin bundles to export.
7. Under **Components**, select any available components to export.

If you select one component but not the components it depends on, the Admin App prompts you to add those missing components to the package.
8. Under **Validate**, make sure your package is valid and then click **Download Package**.
9. When your package downloads, click **Download Package** to download it again, or click **Finish** to exit.

Related CLI commands

buildPackage

downloadPackage

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /package/build

POST /package/download

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Importing packages

To import a package, you can upload it from your computer or have your system retrieve it from one of your data sources. After you import the package, your system runs a system task to synchronize the package components across all instances in your system.

The system can have only one imported package at a time.

**Note:**

- Importing a component that already exists on your system might cause conflicts and should be avoided.
- You need to manually resolve conflicts with Components, while conflicts with Configuration are handled automatically by the system.

Admin App instructions**Procedure**

1. Select the **Configuration** window.
2. Click **Packages**.
3. Click **Import**.
4. Do one of these:
 - If the package you want to import is stored on your computer, click and drag the package file into the **Upload Package** window.
 - If the package you want to import is stored in one of your data sources, click the **Click to Upload** window and then browse for the package file.
5. Under **Package Description**, review the description and then click **Continue**.
6. Under **Configuration**, select any configuration items to import.
7. Under **Plugins**, select any plugin bundles to import.
8. Under **Components**, select any available components to import.
9. Under **Validate**, make sure your package is valid and then click **Install Package**.
Your system starts a system task to install the package components on all instances in the system.
You can monitor the task from the current page or from the **Processes** page.
10. When the task has completed and all package components have been installed, clicking **Complete Install** deletes the package from the system.

Related CLI commands

uploadPackage

loadPackage: loads a package from a data connection

installPackage

getPackageStatus

deletePackage

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

POST /package (Uploads a package)

POST /package/load (Loads a package from a data connection)

POST /package/install

GET /package (Gets the status of the imported package)

DELETE /package

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Setting a login welcome message

You can use the Admin App, REST API, and CLI to set a welcome message for the Admin App. The message appears on the app's login page.

Admin App instructions

Procedure

1. Click the **Configuration** window.
2. Click **Security**.
3. On the **Settings** tab, type a message in the **Single Sign-on Welcome Message** field.
4. Click **Update**.

Related CLI commands

editSecuritySettings

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

PUT /security/settings

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Updating the system

You can update system software by uploading new update packages.



Important: Hitachi Vantara does not provide updates or security fixes for the host operating systems running on system instances.

Before updating

In order for a system to be updated:

- All instances and services must be healthy.
- Each service must be running on its best-practice number of instances.
- Each instance must have enough disk space for the update.

- All required network ports must be available on each instance.
- There can be no in-progress package uploads or installations.

During an update

- All running jobs are paused.
- System availability considerations:
 - Instances shut down and restart one at a time during the upgrade. Other instances remain online and able to service requests.
 - The Admin App remains available but is in a read-only state. You can monitor the progress of the update, but you cannot make any other changes to the system.



Note: Systems with two instances are more susceptible to availability outages during an update than systems with three or more instances.

Verifying update status

As an update runs, you can view its progress on the Configuration > Update page. Also on this page, you can view all system events related to system updates.

Results of an update

After an update, the system runs a new version of the software. Additionally:

- If any of the built-in plugins were updated, your system automatically uses the latest versions of those plugins.
- If an existing service is replaced with a new service, the system automatically runs that new, replacement service.
- If any new services were added, you might need to manually configure those services to run on the system instances.

Update errors

If errors occur during an update, the **Update** page displays information about each error and also displays a Retry button for starting the update over again. Some errors might not be resolved by restarting the update.

If you encounter errors during an update, contact your authorized service provider.

New services and components added during an update

A system update might add new services or plugins. You need to manually configure your system to start using these new components; your system does not start using them automatically.

Applying a system update

Admin App instructions

Procedure

1. Click the **Configuration** window.
2. Click **Update**.
3. Click the **Install** tab.
4. Click and drag the file into the **Upload** window.
The update file is uploaded and the system verifies that the file is valid. This might take several minutes.
5. On the **Update** page, click **View** in the **Update Status** window.
The **Verify & Apply Update** page displays information about the contents of the update.
6. To start the update, click **Apply Update**.

Result

The system verifies that it is ready to be updated. If it isn't, the update stops. In this case, you need to correct the problems before the update can continue.

Related CLI commands

getUpdateStatus

installUpdate

deleteUpdate

loadUpdate

uploadUpdate

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

GET /update

POST /update/install

DELETE /update/package

POST /update/package

POST /update/package/load — (Retrieves update package from a data connection)

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Viewing update history

You can view a list of all updates that have previously been applied to your system. For each update, you can view the corresponding version number and the date on which it was installed.

Admin App instructions

Procedure

1. Click the **Configuration** window.
2. Click **Update**.

Result

The History tab lists previously installed versions and when each was installed.

Related CLI commands

getUpdateHistory

For information on running CLI commands, see [CLI reference \(on page 183\)](#).

Related REST API methods

GET /update/history

You can get help on specific REST API methods for the Admin App at REST API - Admin.

Removing the system

To completely remove your system, do the following on all instances:

Procedure

1. Stop the `run` script from running. You do this using whatever method you're currently using to run the script.
2. Run this command to stop all system Docker containers on the instance:


```
sudo <installation-folder>/bin/stop
```
3. Delete the system Docker containers:
 - a. List all Docker containers:


```
sudo docker ps
```
 - b. Note the container IDs for all containers that use a `com.hds.ensemble` or `com.hitachi.aspen` image.
 - c. Delete each of those containers:


```
sudo docker rm <container-id>
```
4. Delete the system Docker images:

- a. List all Docker images:
`sudo docker images`
 - b. Note the image IDs for all images that use a `com.hds.ensemble` or `com.hitachi.aspen` repository.
 - c. Delete each of those images:
`sudo docker rmi <image-id>`
- 5. Delete the system installation folder:**
`rm -rf /<installation-folder>`

Chapter 7: Best practices

This section contains topics that describe some best practices for administering your system.

Best practices for distributing services

Each service has an associated service unit cost. These costs indicate how computationally expensive one service is to run compared to another. You can use these costs as a guide for how to distribute services across the instances in your system.

Best practices

- Avoid running multiple services with high service unit costs together on the same instance.
- On master instances, avoid running any services besides those classified as System services.
- To use your instances evenly, try to deploy a comparable number of service units on each instance.

Best practices for maintaining system availability

Run master instances on separate physical hardware

For a multi-instance system, master instances should run on separate physical hardware. If your instances run on virtual machines, run the master instances on separate physical hosts.

Run services on more than one instance

In a multi-instance system, you can choose which and how many instances that each service can run on. For redundancy, you should run each service on more than one instance.

Chapter 8: Reference

This section contains information about troubleshooting, information on using the command-line interface, and information on using the REST API.

Troubleshooting

System event issues

For information on viewing system events, see [System events \(on page 101\)](#).

Issue	Description/Resolution
The event log contains instances of event id 6007 with this message: <code>Service <i>service-name</i> health check against HTTP /health port cannot succeed because it is on a different network than service Cluster-Worker. Health check ignored to prevent service interruption.</code>	Your system uses both internal and external networks, but the service specified by the event is on a different network type from the Cluster-Worker service. Because of this, the system cannot perform an additional health verification on the specified service. You can ignore this event. The additional health verification is not required and does not affect the other health verifications used to display the service status on the Monitoring > Services page in the Admin App.
The MAPI Gateway service becomes unresponsive.	If you're working with a storage component that is configured with multiple retries and long timeouts, and if the endpoint for the storage component is unreachable, and if as a result you send multiple verification or activation requests to the endpoint, the MAPI Gateway service can become unresponsive. If the MAPI Gateway service becomes unresponsive, use the System Management application Services function Repair on it.

Issue	Description/Resolution
Users can't write to storage components, objects read from storage components are encrypted, and you receive the critical alert <code>Failed to connect to KMS server.</code>	The Key Management System service is down. The service should restart automatically. When it restarts, log in to HCP for Cloud Scale on port 8200 and provide a quorum of unseal keys.

CLI reference

The system includes an administrative CLI for system management. This interface allows you to perform all tasks relating to system setup and configuration. Any administrative activity that you can perform in the Admin App or REST API can be performed through the CLI.

Your system can also include additional, product-specific CLI tools.

Accessing the CLI tools on a system instance

You can access the CLI tools from any instance. To do this:

Procedure

1. Log in or SSH into a system instance.
2. Navigate to the CLI tool:
`cd <installation-folder>/cli/`
3. Navigate to the folder for the CLI tool you want. For example:
`cd admin`

Installing CLI tools on your computer

You can install your system's CLI tools on your Linux computer. To do this, you must have version 1.8 of the Java Runtime Environment (JRE) installed.

Your system's CLI tools are distributed in .tgz files along with the software installation package.

To install a CLI tool:

Procedure

1. Store the .tgz file in a folder on your computer.
2. Unpack the file:
`tar zxvf filename`

Syntax

The CLI tools have this syntax:

```
<tool-name> [options] [command] [command-specific-options]
```

Options

```
-c, --command <command|category>
```

Specifies the command you want to run. When used with the `--help` option, displays information about the specified command.

You can also use this option to specify a category of commands when using the `--help` option. Doing this displays information about all commands within the specified category.

```
-d, --model-definition <name>
```

Returns information about the specified request model. See [Viewing request models \(on page 185\)](#).

```
--debug
```

Includes verbose debug output for troubleshooting purposes.

```
-h, --help <all>
```

Displays help information. If you specify the `all` argument, displays information on all commands. If you specify the `--help` option, displays information about commands in the specified category.

```
-k, --check-ssl-cert <true|false>
```

Whether to enable SSL security verification. When false, insecure connections are allowed.

```
-m, --model-schema <ModelName>
```

Returns the JSON-formatted schema for the specified request model. See [Viewing request models \(on page 185\)](#).

```
-p, --password <password>
```

Password for the specified user account.

```
--port
```

The port for the system application that supports the CLI tool.

```
-r, --realm <realm>
```

Security realm where your user account is defined. For information, see [Adding identity providers \(on page 121\)](#).

```
-s, --server <server>
```

The hostname or IP address of a system instance.

```
-u, --username <username>
```

Username for an account that has permission to access system.

```
-V, --version
```


Displays the CLI version.

Viewing available commands

- To view all available commands, run:
`<cli-tool-name> --help all`
- To view all command categories, run:
`<cli-tool-name> --help`
- To view all commands within a category, run:
`<cli-tool-name> --help -c <category>`

For example:

```
admincli --help -c instances
```

- To view all information about a single command, run:

```
<cli-tool-name> --help -c <command>
```

For example:

```
admincli --help -c listInstances
```

Viewing request models

Some commands need a JSON-formatted request body along with the command. The command's request model determines how you need to format the request body.

The help command for an individual command indicates what request model it needs.

For example, this help command output indicates that a command to update a service needs a ServiceUpdateModel request:

```
# ./admincli -c updateServiceConfig -h
usage: updateServiceConfig
Name:
  updateServiceConfig
Description:
  Configure service instances
Added:
  1.0
Usage:
  admincli -c updateServiceConfig <options>
Options:
  --service-update-model <ServiceUpdateModel>
File containing JSON text representing a ServiceUpdateModel for the
command
updateServiceConfig. Use the -m and -d options to retrieve information on
request and
response models.
```

Viewing request model information

To view detailed information about the contents of a request model, run:

```
<cli-tool-name> -d <ModelName>
```

For example:

```
admincli -d ServiceUpdateModel
```

Viewing request model formatting

To view the JSON format for the request model, run:

```
<cli-tool-name> -m <ModelName>
```

Editing configuration preferences

You can use the CLI tool's `.conf` file to specify settings to use every time you run a CLI command.

The CLI configuration file has this format:

```
{
  "defaultSettings": {
    "checkSSLCert": "[false|true]", (optional)
    "server": "<hostname>", (optional)
    "realm": "[local|<security-realm-name>]", (optional)
    "username": "<your-username>", (optional)
    "password": "<your-password>" (optional)
  }
}
```

For example, with the following configuration, all commands:

- Are run against the `system.example.com` system
- Verify the SSL certificate for the system before connecting
- Uses the `exampleUsersEast` security realm to authenticate the specified username and password

```
{
  "defaultSettings": {
    "checkSSLCert": "true",
    "server": "system.example.com",
    "realm": "exampleUsersEast"
  }
}
```

File location

You can configure CLI preference by editing the existing `.conf` files in the CLI installation folder.

The options you specify explicitly in a CLI command override the options specified in the `.conf` file.

System error responses

If a CLI request reaches the system and the system returns an error, the CLI response contains:

- An HTTP status code
- Conditionally, a product-specific error code
- A JSON-formatted error response body

HTTP status codes

This table describes the typical reasons why these HTTP status codes are returned.

Status code	Meaning	Description
400	Bad Request	The request body contains one or more of these: <ul style="list-style-type: none"> ▪ An entry that isn't valid ▪ A value for an entry that isn't valid ▪ Incorrectly formatted JSON If the request includes a UUID, the UUID might be incorrectly formatted.
401	Unauthorized	The provided credentials are incorrect.
403	Forbidden	You do not have permission to perform the request.
404	No Found	The resource you are trying to retrieve or edit cannot be found.
409	Conflict	The resource you are trying to create already exists.
500	Internal Server Error	The system experienced an error.
599	Network Connection Timeout Error	The CLI request timed out while trying to connect to the system or one of its instances.

Product-specific error codes

Some CLI requests return product-specific error codes in addition to an HTTP status code. These error codes are listed in the `errorCodes` field in the JSON response body. This table describes these error codes.

Status code	Description
4000	SSL certificate not trusted.

JSON response body

Error response bodies have this format:

```
{
  "statusCode": <HTTP-status-code>,
  "errorCode": <product-specific-error-code>,
  "errorMessage": <message>,
  "errorProperties": [
    {
      "name": <error-property>,
      "message": <error-property-message>
    }
  ]
}
```

REST API Reference

Your system includes a RESTful API that you can use for writing applications that manage the system. Anything you can do in either the Object Storage Management application or the System Management application can also be performed using REST APIs.

Input and output formats

The REST API accepts and returns JSON.

Access and authentication

To use the REST API, you need a user account that has permission to perform the actions you want.

Requesting an access token

After you have a user account, you need to request an authentication token from the system. To do this, you send an HTTP POST request to the `/auth/oauth` method.

When you generate a new access token, a refresh token also gets generated automatically.

Here's an example using the cURL command-line tool:

```
curl -ik -X POST https://mysystem.example.com:8000/auth/oauth/ \
-d grant_type=password \
-d username=user1 \
```

```
-d password=password1 \
-d scope=* \
-d client_secret=my-client \
-d client_id=my-client \
-d realm=marketingUsers
```

In response to this request, you receive a JSON response body containing an `access_token` field. The value for this field is your token. For example:

```
{"access_token" : "eyJr287bjle..."}
```



Note:

- To get a list of security realms for the system, send an HTTP GET request to the `/setup` method. For example, to do this with cURL:

```
curl -k -X GET --header 'Accept: application/json'
'https://mysystem.example.com:<admin-app-port>/api/admin/setup'
```

For information on configuring security realms, see [Adding identity providers \(on page 121\)](#).

- To get an access token for the local admin user account, you can omit the realm option for the request, or specify a realm value of Local.

Submitting your access token

You need to specify your access token as part of all REST API requests that you make. You do this by submitting an Authorization header along with your request. Here's an example that uses cURL:

```
curl -X GET --header "Accept:application/json" https://
mysystem.example.com:<admin-app-port>
/api/admin/instances --header "Authorization: Bearer eyJr287bjle..."
```

Changing a password

You can use the REST API to change your system's password using the following cURL command.

```
$1=<server-name>
```

```
$2=<current-password>
```

```
$3=<new-password>
```

```
TOKEN=$(curl -ik -X POST https://$1.mysystem.com:8000/auth/oauth/ -d
grant_type=password
-d username=admin -d password=$2 -d scope=* -d client_secret=hci-client -
```

```
d client_id=hci-client
-d realm=local 2>&1 | grep access_token | awk -F: '{print $2}' | awk -F
\" '{print $2}')
```



```
curl -v -X POST --header 'Content-Type: application/json' --header
"Authorization: Bearer $TOKEN"
https://$1.mysystem.com:8000/api/admin/setup/password -d '{"password":
"$$$3$$$"}'
```

Viewing and using REST API methods

Your system includes web-based documentation pages where you can view all supported REST API methods, including the request bodies, request URLs, response bodies, and return codes for each. You can also use these pages to run each REST API method.

Viewing the API documentation

To view the REST API documentation:

- In the Object Storage Management application, click the user icon () and then click REST API.
- In the System Management application, click the user icon () and then click REST API - Admin.

Making requests

You can use the REST API documentation pages to experiment with the REST API.



Note: Any requests you submit on the REST API page take effect on the system.

To use the REST API page to run a method:

Procedure

1. Select the row for the method you want.
2. If the method you want needs a UUID:
 - a. Select the row for the GET method for the resource type that you want.
 - b. Click **Try it Out!**.
 - c. In the JSON response body, copy the value for the **uuid** field for the resource that you want.



Note: If you specify UUIDs when creating resources, the UUIDs are ignored.

3. If the method you want needs a request body:

- a. In the **Parameters** section, under **Model Schema**, click inside the JSON text box.

The JSON text is added to the **Value** field.

- b. Edit the JSON in the **Value** field.



Note: Some methods might need other information in addition to or instead of UUIDs or JSON-formatted text. Some need particular string values or need you to browse for and select a file to upload.

4. Click **Try it out!**.

Error responses

When an API request fails, the API returns:

- An HTTP status code
- Conditionally, a system-specific error code
- A JSON-formatted error response body

HTTP status codes

This table describes the typical reasons why these HTTP status codes are returned. For information on the status codes for a particular method, view the system REST API web interface.

Status code	Meaning	Description
400	Bad Request	The request body contains one or more of these: <ul style="list-style-type: none"> ▪ An entry that's not valid ▪ A value for an entry that's not valid ▪ JSON code that's not validly formatted If the request includes a UUID, the UUID might not be validly formatted.
403	Forbidden	You do not have permission to perform the request.
404	Not Found	The resource you are trying to retrieve or edit cannot be found.
409	Conflict	The resource you are trying to create already exists.
500	Internal Server Error	The system experienced an error.

System-specific error codes

Some API requests return system-specific error codes in addition to an HTTP status code. These error codes are listed in the `errorCodes` field in the JSON response body. This table describes these error codes.

Status code	Description
4000	SSL certificate not trusted.

JSON response body

REST API error responses have this format:

```
{
  "statusCode": <HTTP-status-code>,
  "errorCode": <system-specific-error-code>,
  "errorMessage": <message>,
  "errorProperties": [
    {
      "name": <error-property>,
      "message": <error-property-message>
    }
  ]
}
```


Hitachi Vantara



Corporate Headquarters

2535 Augustine Drive

Santa Clara, CA 95054 USA

HitachiVantara.com | community.HitachiVantara.com

Contact Information

USA: 1-800-446-0744

Global: 1-858-547-4526

HitachiVantara.com/contact