

Hitachi Content Platform for Cloud Scale

v1.2.0

Object Storage Management Guide

This document describes the Object Storage Management application, one of the applications available as part of Hitachi Content Platform for cloud scale.

© 2019 Hitachi, Ltd. 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 Corporation (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 Corporation 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 Corporation.

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>.

Contents

- Preface..... 7**
 - About this document..... 7
 - Intended audience..... 7
 - Product version..... 7
 - Document conventions..... 7
 - Release notes..... 9
 - Related documents..... 9
 - Accessing product downloads..... 9
 - Accessing product documentation..... 10
 - Getting help..... 10
 - Comments..... 10

- Chapter 1: Introducing Hitachi Content Platform for cloud scale.... 11**
 - Data access..... 11
 - Storage components, buckets, and objects..... 11
 - High availability..... 12
 - Site availability..... 12
 - Service availability..... 13
 - Metadata availability..... 13
 - Object data availability..... 13
 - Network availability..... 13
 - Failure recovery..... 14
 - Instance failure recovery..... 14
 - Service failure recovery..... 14
 - Storage component failure recovery..... 15
 - Security and authentication..... 15
 - User accounts..... 15
 - Data access control..... 16
 - API access..... 16
 - Data security..... 17
 - Network isolation and port mapping..... 18
 - Scalability of instances, service instances, and storage components..... 19
 - Bucket synchronization..... 20
 - Supported limits..... 20

Logging in.....	21
HCP for cloud scale applications.....	22
Switching between applications.....	23
Chapter 2: Dashboard.....	24
Serial number.....	24
Entering your serial number.....	24
Object Storage Management application instructions.....	24
Related API method.....	24
Displaying your serial number.....	25
Object Storage Management application instructions.....	25
Related API method.....	25
System reports.....	25
Displaying the active object count.....	25
Object Storage Management application instructions.....	25
Displaying the alert count.....	25
Object Storage Management application instructions.....	25
Related API method.....	26
Metrics.....	26
Displaying metrics.....	26
Object Storage Management application instructions.....	26
Available metrics.....	26
Tracing requests and operations.....	31
Displaying traces.....	31
Traceable operations.....	31
Chapter 3: Managing storage components.....	38
Displaying storage component analytics.....	38
Displaying counts of storage components.....	38
Object Storage Management application instructions.....	38
Related API method.....	39
Viewing storage components.....	39
Object Storage Management application instructions.....	39
Related API method.....	39
Adding a storage component.....	40
Object Storage Management application instructions.....	40
Related API method.....	41
Modifying a storage component.....	42
Object Storage Management application instructions.....	42
Related API method.....	42
Activating a storage component.....	42
Object Storage Management application instructions.....	43

Related API method.....	43
Deactivating a storage component.....	43
Object Storage Management application instructions.....	43
Related API method.....	43
Marking a storage component as read-only.....	44
Object Storage Management application instructions.....	44
Related API methods.....	44
Marking a storage component as read-write.....	44
Object Storage Management application instructions.....	44
Related API methods.....	45
Chapter 4: Managing bucket synchronization.....	46
About bucket synchronization.....	46
Bucket synchronization configuration.....	49
Configure bucket synchronization (PUT bucket replication).....	49
Get bucket synchronization rules (GET bucket replication).....	54
Get object synchronization status.....	56
Delete bucket synchronization rules (DELETE bucket replication).....	57
Chapter 5: Notifications and user profiles.....	59
Viewing alerts.....	59
Object Storage Management application instructions.....	59
Related API method.....	59
HCP for cloud scale alerts.....	60
HCP for cloud scale events.....	61
User profiles.....	63
Chapter 6: Services.....	64
Service categories.....	64
HCP for cloud scale services.....	64
Listing service ports.....	80
Scaling Metadata Gateway instances.....	80
Chapter 7: S3 User Credentials.....	82
Obtaining S3 credentials.....	82
S3 application instructions.....	82
Related API method.....	83
Revoking S3 credentials.....	83
Chapter 8: HCP for cloud scale APIs.....	84
Object Storage Management APIs.....	84
System Management APIs.....	85
Support for Amazon S3 APIs.....	85

Preface

This document contains information about using the Object Storage Management application, which is available as part of the Hitachi Content Platform for cloud scale (HCP for cloud scale) software.

This document matches the information in the online Help available in the Object Storage Management application.

About this document

This document describes the Object Storage Management application, one of the applications available as part of Hitachi Content Platform for cloud scale.

Intended audience

This book is intended for those who use the Object Storage Management application, one of the applications available as part of Hitachi Content Platform for cloud scale.

Product version

This document revision applies to HCP for cloud scale v1.2.0.




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.

Convention	Description
<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: <pre>pairdisplay -g group</pre> <p>(For exceptions to this convention for variables, see the entry for angle brackets.)</p>
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairdisplay -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).

Icon	Label	Description
	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

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

Referenced documents

- *Installing Hitachi Content Platform for Cloud Scale* (MK-HCPCS002-03): This document provides the information you need to install the HCP for cloud scale software.
- *Hitachi Content Platform for Cloud Scale System Management Application Help* (MK-HCPCS001-03): This Help system contains the instructions for using the HCP for cloud scale System Management application to 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 Management API Reference* (MK-HCPCS007-01): This document is for customers, and describes the management application programming interface (API) endpoints available for customer use.
- *Hitachi Content Platform for Cloud Scale Third-party Copyrights and Licenses* (MK-HCPCS003-03): This document contains copyright and license information for third-party software distributed with or embedded in the HCP for cloud scale operating system, core software, and applications.
- *Hitachi Content Platform for Cloud Scale Release Notes* (RN-HCPCS004-04): 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.

Accessing product downloads

Product software, drivers, and firmware downloads are available on Hitachi Vantara Support Connect: <https://support.hitachivantara.com/>.

Log in and select Product Downloads to access the most current downloads, including important updates that may have been made after the release of the product.

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](#) 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](#) 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 Corporation.

Thank you!

Chapter 1: Introducing Hitachi Content Platform for cloud scale

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

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

HCP for cloud scale has the ability to federate S3-compatible storage from virtually any private or public source, and present the combined capacity in a single, centrally managed, global namespace.

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 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 and repair. The HCP for cloud scale system provides functions to send notification of alerts, track and monitor throughput and performance, and trace actions through the system.

Data access

HCP for cloud scale supports the Amazon Simple Storage Service (S3) application programming interface (API), which allows client applications to store and retrieve unlimited amounts of data from configured storage services.

Storage components, buckets, and objects

A storage component is an Amazon S3-compatible storage system, running independently, that is manageable by HCP for cloud scale 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 Nodes
- Any Amazon S3-compatible storage service

An HCP for cloud scale bucket is modeled on a storage service bucket. A bucket is a logical collection of secure data objects that is created and managed by a client application. HCP for cloud scale uses buckets to manage storage components, and an HCP for cloud scale site can be thought of 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 ACL support using S3 APIs. 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. If you want to empty a bucket and reuse it, 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 the bucket and keep it.

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](#).

High availability

HCP for cloud scale provides high availability for multi-instance sites. High availability requires 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.

Site availability

An HCP for cloud scale site has three master instances, and can tolerate the failure of one master instance without interruption of service. Even if two or all three master instances fail, HCP for cloud scale services may be functional (but 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 up, the nodes "elect a leader" using the raft consensus algorithm. The leader processes all GET and PUT requests. If the followers cannot identify the leader, they elect a new leader. The Metadata Gateway service can tolerate service instance failure, and service remains available without loss of data, so 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, performance decreases.

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 only has one service instance fails, HCP for cloud scale will automatically start a new service instance.

Metadata availability

Metadata is available as long as two services are available:

- S3 Gateway
- Metadata Gateway

Object data availability

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

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

The availability of object data depends on the storage component. For high availability of object data, you should use a storage component with high availability, such as HCP, HCP-S, and AWS S3. This is true as well for data protection.

Network availability

You can install each HCP for cloud scale instance with an internal and an external network interface. If you want to avoid networking single points of 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 (that is, as Active-Passive) into one virtual network interface
- 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, provides real-time visual health representations, issues alert messages when needed, and can automatically take action to recover from the following types of failures:

- 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 recommended minimum number of service instances. Data on the failed instance is not lost and remains consistent. However, while the instance is down, data redundancy may degrade.

HCP for cloud scale only adds new service instances automatically for floating services. Depending on the remaining number of instances and service instances running, you may 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 brings up another one 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 one of those service instances 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 failure, it sets the storage component state to INACCESSIBLE, so that HCP for cloud scale will not try to write new objects to it. HCP for cloud scale can send an alert when this event happens. While a storage component is down, the data in it is not accessible.

HCP for cloud scale keeps checking 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 can send an alert when this event happens as well. Once the storage component is repaired and brought back online, the data its contains is again accessible, and HCP for cloud scale can write new objects to it.

Security and authentication

HCP for cloud scale controls access to system functions by means of user accounts, roles, permissions, and OAuth tokens, where user accounts reside in an external identity provider. HCP for cloud scale controls access to data (S3 APIs) by means of 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 functions. After the initial user account is created, you can change its password any time, but you cannot disable it and you cannot change its permissions.

The initial user is the only local account allowed, and is only intended to let you configure an identity provider (IdP). HCP for cloud scale can communicate to 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 any 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 S3 APIs

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, and you can switch between applications without logging in again.

Data access control

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

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 only list the buckets that user owns

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

API access

The Object Storage Management application APIs require a valid OAuth access token for a user account with suitable permissions; otherwise, the requests are rejected. With one exception, the System Management application APIs also require a valid OAuth access token for a user account with suitable permissions; otherwise, the requests are rejected. (The API call to generate an OAuth token 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 the request. OAuth tokens are valid for five hours.



Note: A user can revoke all OAuth tokens for any other HCP for cloud scale user. You would do this 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 operations 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.

A valid user account with 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 unless and until revoked.

A valid user account with 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.

Data security

HCP for cloud scale supports encryption of data sent between systems ("in flight") and data stored persistently within the system ("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 certificates are required:

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

For a system certificate, HCP for cloud scale 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 or replace it by using one of three options:

1. Upload a PKCS12 certificate chain and password and apply it as the active system certificate.
2. Download a certificate signing request (CSR), 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 a client certificate, you need to upload the certificate of the clients HCP for cloud scale needs to access using SSL.

You can manage certificates, as well as view the installed certificates and their details, 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 App user interface (GUI)
- Object Storage Management App GUI

You can enable or disable data-in-flight encryption for these data paths:

- Between HCP for cloud scale and an identity provider (IDP) server
- Between HCP for cloud scale and each application using TLS or SSL
- Between HCP for cloud scale and each managed storage component
- Between HCP for cloud scale and each SMTP server using SSL or STARTTLS

Communication among HCP for cloud scale instances are without data-in-flight encryption. Depending on your security requirements, you may need to set up an isolated internal network for your HCP for cloud scale site.

Data-at-rest encryption

HCP for cloud scale stores three kinds of data persistently:

1. HCP for cloud scale services data
2. HCP for cloud scale metadata and user-defined metadata
3. 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. Data of the last kind is handled by storage components. If needed, you can use storage components that support data-at-rest encryption. Storage components can self-manage their keys, or HCP for cloud scale can facilitate customer-supplied keys following the S3 API specification.

Network isolation and port mapping

When you install HCP for cloud scale, you can achieve network isolation by configuring it with one external network and one internal network.

HCP for cloud scale software creates a cluster using commodity x86 Linux servers that are networked using Ethernet. The software uses two networks constructed on the operating system hosting the HCP for cloud scale software. These networks may additionally employ link aggregation defined by the OS administrator. While two networks provide optimal traffic isolation, it is possible to deploy the software using a single network. These networking decisions are made by the OS administrator. These network topology decisions must be completed and already in place when you install HCP for cloud scale. HCP for cloud scale uses a variety of network ports identified during the installation process. You will have this one opportunity to adjust or alter the default ports used by any service.

When you install HCP for cloud scale, you can also configure it to use specific ports instead of the default ports.

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

Scalability of instances, service instances, and storage components

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

- Instances: Additional 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 additional instances if you want to improve system performance or if you are running out of disk 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 a site with 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 instances of the S3-Gateway service to improve throughput of S3 API transactions without having to add a compute instance.
- You can configure it run on fewer instances. For example, you can free up 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 it.
- 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. These include:

- Metadata-Coordination
- Metadata-Cache

You might add additional 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 require
- The existing storage components do not provide the functionality you require

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.

Bucket synchronization is implemented using cross-region replication. An external storage system can be another HCP for cloud scale system, AWS, or any S3-compatible system.

Bucket sync-to provides 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 provides the following advantages:

- **Data consolidation:** Transformed data can be stored on an HCP for cloud scale system. An HCP for cloud scale system can synch 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 require using 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 managing bucket synchronization, see [Managing bucket synchronization \(on page 46\)](#).

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	

Entity	Minimum	Maximum	Notes
Users (external)	None	Unlimited	The local user can do all operations including MAPI calls and S3 API calls. However, it is recommended that HCP for cloud scale be configured 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 operation is 5 GB.
Storage components	1	Unlimited	

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. Enter 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.

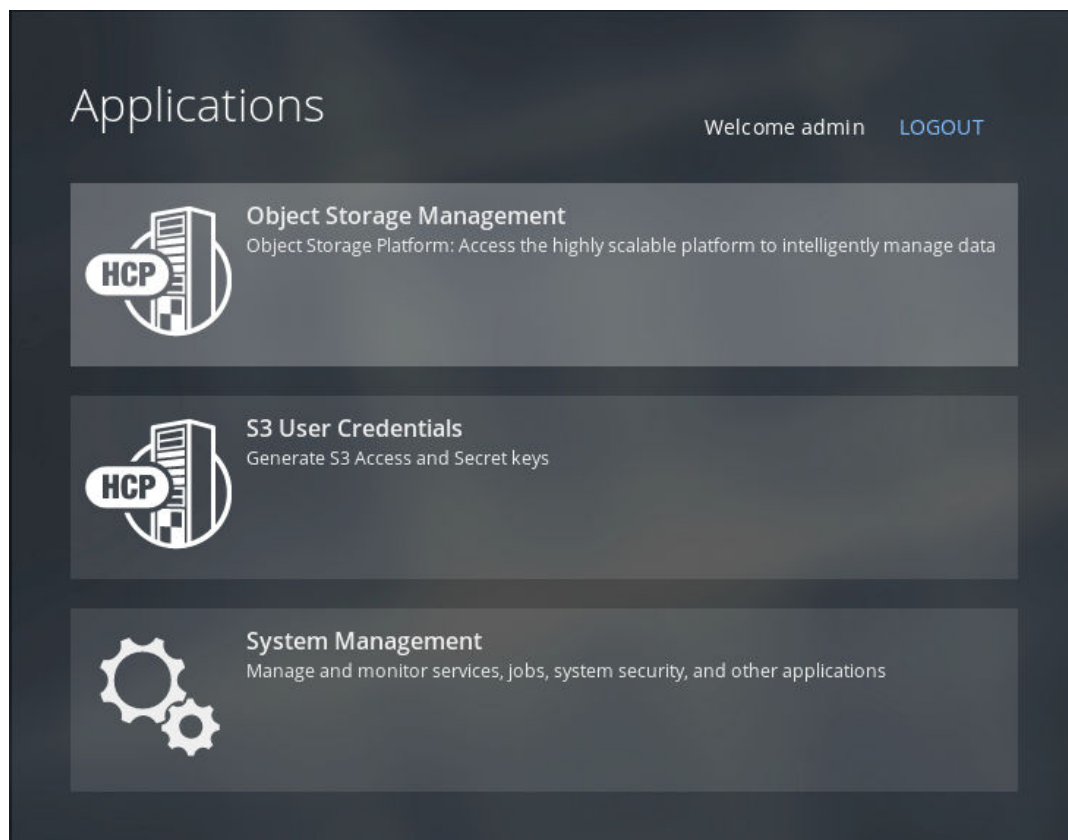


Note: When a new user is created and added to a group, that user might not have immediate access to HCP for cloud scale. Instead, login fails with the message "Not authorized. Please contact your system administrator." Verify the credentials. If the condition persists, the system administrator can use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

HCP for cloud scale applications

After you log in, HCP for cloud scale presents you with applications you can launch:

- Object Storage Management: Manage and monitor storage components, data objects, alerts, and regions
- S3 User Credentials: Generate S3 access and secret keys
- 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




You can return to the Applications page to switch back and forth between these 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, and you can switch between applications without logging in again.

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.

Chapter 2: Dashboard

Hitachi Content Platform for cloud scale (HCP for cloud scale) provides functions that let you monitor the activity and performance of the system, storage components, and objects stored on storage components in real time.

When you select the Object Storage Management application, your landing point is the **Dashboard** page.



Note: If the browser window is reduced horizontally until the side navigation bar is collapsed, there is no function to toggle it open again. If this happens, expand the browser window horizontally until the bar reappears.

Serial number

You can use the Object Storage Management application or APIs 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 proceeding further.

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 **Dashboard** and click on the Edit icon next to the **Serial Number** field. The **Add Serial Number** window opens.
2. Enter your serial number and click **Add**.

Related API method

```
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 APIs 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 API method

```
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.

System reports

The **Dashboard** page includes a System Reports section that displays the current counts of active objects and alerts in the system.

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 the alert count

You can use the Object Storage Management application or APIs to display a count of active alerts.

Object Storage Management application instructions

To display the Alert Count report, select Dashboard.

The report displays the number of active alerts, if any. If there are no active alerts, this infographic is not displayed.

In addition, the alert icon () in the upper right corner of the page, displays a badge with the current count of active alerts, if any.

Related API method

```
POST /alert/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
- Metadata Policy Engine
- Metadata Cache
- 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 the operation of HCP for cloud scale instances and services are collected and provided by the System Management application. Collection of these metrics cannot be disabled. For information about these metrics, see the Help available in that application.

Displaying metrics

You can use the metrics service to display or graph metrics, or use the service APIs 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 click the Metrics panel. 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 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>	Count of the total number of requests made to the health check API
<code>http_monitoring_requests_total</code>	Count of the total number of requests made to the monitoring API
<code>scrape_duration_seconds</code>	Duration in seconds of the scrape (collection interval)
<code>scrape_samples_post_metric_relabeling</code>	Number of samples remaining after metric relabeling was applied
<code>scrape_samples_scraped</code>	Number 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

S3 Gateway

The following metrics are available from the S3 Gateway service.

Metric	Description
<code>http_s3_servlet_errors_total</code>	Count of total number of errors returned by the s3 servlet, grouped by error
<code>http_s3_servlet_get_object_response_bytes_total</code>	Count of total bytes in the body of S3 GET object responses
<code>http_s3_servlet_ingest_object_bytes_per_bucket</code>	Count of total objects ingested to S3 for each bucket, grouped by bucket
<code>http_s3_servlet_operations_total</code>	Count of total number of S3 operations made to the s3 servlet for each endpoint, grouped by operation
<code>http_s3_servlet_post_object_bytes_total</code>	Count of total bytes of objects posted to S3

Metric	Description
http_s3_servlet_put_copied_bytes_total	Count of total bytes of objects PUT copied to S3
http_s3_servlet_put_object_bytes_total	Count of total bytes of objects put to S3
http_s3_servlet_put_object_part_bytes_total	Count of total bytes of PUT part operations to S3
http_s3_servlet_requests_histogram_latency_seconds	Latency in seconds as measured by a histogram timer, grouped by operation
http_s3_servlet_requests_histogram_latency_seconds_count	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; used with count to determine average
http_s3_servlet_requests_latency_seconds	Latency in seconds as measured by a summary timer, grouped by operation
http_s3_servlet_requests_latency_seconds:hour_average	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	
http_s3_servlet_requests_per_bucket	Count of total number of requests made to the s3 servlet per bucket, grouped by bucket
http_s3_servlet_requests_total	Count of total number of requests made to the s3 servlet, grouped by method
http_s3_servlet_unimplemented_api_request_total	Count of total number of requests made for unimplemented S3 APIs
http_s3_servlet_unimplemented_bucket_api_request_total	Count of total number of requests made for unimplemented S3 APIs per bucket, grouped by API
http_s3_servlet_unimplemented_object_api_request_total	Count of total number of requests made for unimplemented S3 APIs per object, grouped by API
http_s3_servlet_unimplemented_service_api_request_total	Count of total number of requests made for unimplemented S3 APIs per service, grouped by API

Metric	Description
http_s3_servlet_unknown_api_requests_to tal	Count of total number of requests made for unknown S3 APIs, grouped by API

Metadata Policy Engine

The following metrics are available from the Metadata Policy Engine service.

Metric	Description
duq_query_latency	Time to get a response from a Get DUQ query
duq_query_latency_count	Number of times the DUQ is queried (for determining the average)
duq_query_latency_sum	Aggregate sum of latencies for DUQ latencies (for determining the average)

Metadata Coordination

The following metrics are available from the Metadata Coordination service.

Metric	Description
mcs_copies_per_partition	Count of number of copies of each metadata partition of each key space (to verify protection)
mcs_disk_usage_per_instance	Total disk usage of each metadata instance
mcs_disk_usage_per_partition	Disk usage of each metadata partition of each key space
mcs_partitions_per_instance	Count of total number of metadata partitions per metadata instance (to verify balance)

Metadata Gateway

The following metrics are available from the Metadata Gateway service.

**Note:**

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

Metric	Description
metadata_clientobject_active_count	Number of client objects in metadata that are in the ACTIVE state
metadata_clientobject_and_part_active_space	Space occupied by client objects and parts in metadata that are in the ACTIVE state
metadata_clientobject_part_active_count	Number of client object parts in metadata that are in the ACTIVE state
metadata_storedObject_active_space	Space occupied by stored objects on the back-end storage components
sync_from_object_count_failed	Number of objects that failed to synchronize from external storage (sync-from) by this instance
sync_from_object_count_succeeded	Number 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_to_object_count_failed	Number of objects that failed to synchronize to external storage (sync-to) by this instance
sync_to_object_count_succeeded	Number of objects synchronized to external storage (sync-to) by this instance
sync_to_object_size_total	Total size of object data synchronized to external storage (sync-to) by this instance
update_queue_inprogress	Number of update queue entries in progress
update_queue_size	Size of the update queue

Tracing requests and operations

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

The Tracing service provides end-to-end, distributed tracing of S3 requests and operations by HCP for cloud scale services. By tracing requests and operations you can 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 panel. 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 click on 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
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: Managing storage components

Hitachi Content Platform for cloud scale (HCP for cloud scale) provides functions to let you manage and monitor storage components.

Your starting point is the **Storage Component** page.



Note: 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 Services function Repair on it.

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 API method

```
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.

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 issues (for example, network, authentication, or certificate issues) accessing it
- **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 API method

```
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.

Adding a storage component

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



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

Before adding a storage component to HCP for cloud scale, you must have created an S3 bucket on it.

To define a storage component, it must be operating, 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 proxy user name and password
- API port
- S3 credentials (the access key and secret key you use for access to the storage component bucket)

Object Storage Management application instructions

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


The storage component must include 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. Enter the following information:
 - a. **Storage Component Name** (optional): The display name you choose for the storage component. Enter up to 1024 alphanumeric characters.
 - b. **Storage Type**: Select **AMAZON_S3**, **HCP_S3**, **HCPS_S3**, or **GENERIC_S3**.
 - c. **Region** (optional): Enter a region name of up to 1024 characters.
HCP for cloud scale doesn't validate this field except for its length.
 - d. **Host**: Enter the host name of the storage component. Use ASCII characters only.
4. Click **Next**.
The **Connection Advanced** page opens.
5. Enter the following information:
 - a. Select the protocol **HTTPS** (the default) or **HTTP**.


- b. If **Use Default** is selected, the appropriate default port number is filled in. If you deselect **Use Default**, enter the **Port** number.
 - c. If you select **Proxy**, enter values for the fields **Proxy Host** and **Proxy Port**, and if the proxy requires authentication, enter the **Proxy User Name** and **Proxy Password**.
6. Click **Next**.
The **Activation** page opens.
7. Enter the following information:
 - a. **Bucket Name**: The name of the bucket on the storage component. Enter a name from 3 to 63 characters long containing 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: Select **V2** or **V4**.
 - d. Enter your **Access Key**.
 - e. Enter 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 with the information to correct.
 - If the information is correct, click **Create**.

Result

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

 **Note:** After you define the storage component, if its state is UNVERIFIED, go back and check the parameters you used when adding it.

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 provides additional details that you can use to troubleshoot the problem.

Related API method

POST /storage_component/create



Note: After you define the storage component, if its state is UNVERIFIED, check the parameters you used when adding 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 API method

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 API method

```
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.

Once 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 API method

```
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 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, deselect 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 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.

Chapter 4: Managing bucket synchronization

Hitachi Content Platform for cloud scale (HCP for cloud scale) provides functions to let you configure and manage bucket synchronization.

About bucket synchronization

HCP for cloud scale can synchronize the following kinds of data in buckets:

- Object data
- All user metadata (that is, anything that can be returned in the header `x-amz-meta-*`)
- Tags
- `Content-Type` system metadata
- Objects that the owner of the source bucket doesn't have permission to read

**Note:**

Objects that existed before the functions are configured are not synchronized.

HCP for cloud scale checks the rules that are valid at the time an object is synchronized, not at the time the object is ingested.

Objects that are marked as deleted are not synchronized.

Most system metadata is not synchronized, specifically:

- Owner ID and Name
- Timestamps (when last modified)
- Metadata returned in `x-amz-grant-*`
- Metadata returned in `x-amz-acl`
- Metadata returned in `x-amz-grant-*`
- Metadata returned in `x-amz-acl`
- Metadata returned in `x-amz-storage-class`
- Metadata returned in `x-amz-replication-status`
- Metadata returned in `x-amz-server-side-encryption-*`
- Metadata returned in `x-amz-restore-*`
- Metadata returned in `x-amz-version-id-*`
- Metadata returned in `x-amz-website-redirect-location`
- Metadata returned in `x-amz-object-lock-*`

Comparing synchronization to replication

Unlike AWS replication, HCP for cloud scale can synchronize with buckets on storage systems outside of AWS.

AWS determines the destination bucket using rules, but only applies one rule to each new object. In contrast, HCP for cloud scale can apply multiple rules to each new object so long as the destination buckets are different. This is how one-to-many synchronization is implemented.

AWS does not replicate, but HCP for cloud scale synchronizes, objects that the owner of the source bucket doesn't have permission to read.

In contrast with AWS replication, HCP for cloud scale does not synchronize the following:

- Access control lists (ACLs)
- Lock retention information
- Objects that are encrypted using Amazon S3 managed keys (SSE-S3) and AWS KMS managed keys (SSE-KMS)

If an object being synchronized has the same name as an object in the target bucket, the result depends on whether or not the target bucket uses versioning:

- If versioning is used, the old object is kept as an old version.
- If versioning is not used, the old object is replaced by the new object.

HCP for cloud scale buckets always use versioning. The best practice is to use versioning in all target buckets.

Best-effort ordering

HCP for cloud scale guarantees that operations are applied in the order of their arrival (*strong consistency*). However, synchronizing multiple operations applied in a short period of time to the same object presents the following difficulties:

- In a distributed system, especially when many systems are involved, synchronizing all operations in correct order is complex.
- Even if HCP for cloud scale synchronizes all operations in correct order to an external storage component, that component might not guarantee that the operations are applied with strong consistency. In particular, AWS guarantees only "eventual consistency."
- For bucket sync-from, the external queue service might not guarantee that messages are provided in correct order. In particular, AWS Simple Queue Service (SQS) does not support first-in, first-out (FIFO) queues for S3 notifications.

Therefore, HCP for cloud scale makes its best effort to synchronize only the latest state of an object, not each version or operation for the object. For example:

- Assume that a client sends three operations to an object, and that they are all committed: (1) PUT, (2) PUT, (3) DEL. The latest state of the object is (3) DEL. HCP for cloud scale only synchronizes DEL.
- Assume that a client sends three operations to an object, and that they are all committed: (1) PUT, (2) DEL, (3) PUT. The latest state of the object is (3) PUT. HCP for cloud scale only synchronizes (3) PUT.

This approach does not guarantee that the latest state of an object will be in the external storage for all situations. Partly because of the "eventual consistency" provided by AWS S3 API, corner cases still exist.

Overview of tasks

These are the high-level steps involved in setting up bucket synchronization:

1. If appropriate, the HCP for cloud scale administrator assigns a sync-to or sync-from role to tenant administrators.
2. The administrator creates buckets.
3. The administrator configures synchronization rules.
4. Users can now use synchronization when writing objects to buckets.



Note: If you use the AWS command-line interface to configure bucket synchronization, use at least `aws-cli v1.16.211`.

Bucket synchronization configuration

Bucket synchronization is configured using S3 `PUT bucket replication` API requests that define rules. Each bucket can have up to 1,000 rules, but all rules must be sync-to or sync-from rules. Each rule defines the following:

- External bucket settings
- A set of one or more prefixes; an object with one of the prefixes is mirrored
- A set of one or more tags; an object with all, or any, of the tags is mirrored
- For sync-from, external queue settings

Because you can configure multiple rules with multiple tags, you have flexibility in selecting objects to mirror. For example:

- To mirror all objects that contain `Tag1` and `Tag2`, you can configure one rule that includes both tags.
- To mirror all objects that contain `Tag1` or `Tag2`, you can configure two rules, one for each tag.

For information on `PUT bucket replication` see [Configure bucket synchronization \(PUT bucket replication\) \(on page 49\)](#).

Rule collisions

HCP for cloud scale can apply multiple bucket synchronization rules to each new object so long as the destination buckets are different. This is how one-to-many synchronization is implemented.

A rule collision is when two or more rules that apply to an object have the same destination (that is, the same external host, port, and bucket). HCP for cloud scale does not allow rule collisions, so `PUT bucket replication` requests are rejected if they contain rule collisions. To avoid rule collisions, you can define as many tags in a rule as necessary, so that multiple rules with the same destination are not needed.

Effect of configuration changes

When bucket synchronization rules are created, updated, or deleted, the changes only apply to new objects or new S3 API operations. Objects that existed before the rules are configured are not synchronized. If an object exists in the state `PENDING` when a rule is created, updated, or deleted, the rule might not be applied to it, because the object might be in the midst of copying.

Configure bucket synchronization (PUT bucket replication)

You can configure S3 bucket sync-to and sync-from settings.

HTTP request syntax (URI)

```
aws --endpoint -url https://host_ip s3api put-bucket-replication --bucket
"bucket" --replication-configuration '{body}'
```

Request structure

A rule consists of up to 1000 prefixes and tag-value pairs. You can configure up to 1000 rules per bucket. Separate tag-value pairs in the rule using the keywords "And" : or "Or" :.

The request body is shown below:

```
'{
  "Role": "",
  "Rules": [{
    "ID": "string",
    "Filter": {
      "Prefix": "string",
      "Tag": {
        "Key": "string",
        "Value": "string"
      }
    }
  },
  "Status": "boolean",
  "Destination": {
    "Bucket": "json",
    "Account": "B64_key, B64_key",
    "StorageClass": ""
  }
}
.
.
.
}]
}'
```



Note: S3 parameters not shown are not required, not supported, and if specified should be left empty.

Parameter	Required	Type	Description
Role	Yes	N/A	Not supported; leave empty.
ID	No	String	Unique identifier for rule, up to 255 characters. All rules must specify the same bucket.

Parameter	Required	Type	Description
Priority	Yes	Integer	Not supported; ignored.
DeleteMarkerReplication.Status	No	String	Not supported; if provided, leave as <code>Disabled</code> .
Prefix	No	String	Prefix (one per rule). Up to 1024 characters.
Key	No	String	Tag key (up to 1000 per rule). Up to 128 characters.
Value	No	String	Tag value. Up to 256 characters.
Rules.Status	Yes	Boolean	Enter <code>Enabled</code> or <code>Disabled</code> . If <code>Disabled</code> , rule is ignored.
Bucket	Yes	Base64-encoded JSON	External S3 bucket access settings. <ul style="list-style-type: none"> For bucket synch-to, the settings to access the external bucket. For bucket sync-from, the settings to access the external bucket and the SQS queue settings.
Account	No	Base64 encoded string	The S3 access key and secret key credentials to the external S3 bucket.
StorageClass	No	Enum	Optional destination storage class override. If provided, leave empty.

Bucket sync-to structure

Bucket sync-to settings are defined by a set of parameters and passed in the value of `Destination.Bucket` as a Base64-encoded string.

The syntax of a bucket sync-to setting is shown below:

```
arn::sync-
to::version::host:port>::region::bucket_name::auth_version::path_style_always
```

Parameter	Required	Type	Description
version	Yes	String	Enter 1 . 0.
host	Yes	IP address	Host IP address.
port	Yes	integer	Host port.
region	Yes	String	The S3 region.
bucket_name	Yes	String	The name of the bucket. Enter a name from 3 to 63 characters long containing only lowercase characters (a-z), numbers (0-9), periods (.), or hyphens (-). The bucket must already exist.
auth_version	Yes	String	AWS Signature version: enter v2 or v4.
path_style_always	Yes	Boolean	Path-style URLs for bucket access: enter true or false.

Bucket sync-from structure

Bucket sync-from settings include both a bucket address and a notification queue. The settings are defined by a set of parameters and passed in the value of `Destination.Bucket` as a Base64-encoded string.

The syntax of a bucket sync-from setting is shown below:

```
arn::sync-
to::version::host:port>::S3_region::bucket_name::auth_version::path_style_a
lways::AWS_SQS::SQS_region::SQS_queue::SQS_access_key::SQS_secret_key
```

Parameter	Required	Type	Description
version	Yes	String	Enter 1 . 0.
host	Yes	IP address	Host IP address.
port	Yes	integer	Host port.
S3_region	Yes	String	The S3 region.
bucket_name	Yes	String	The name of the bucket. Enter a name from 3 to 63 characters long containing only lowercase characters (a-z), numbers (0-9), periods (.), or hyphens (-). The bucket must already exist.
auth_version	Yes	String	AWS Signature version: enter v2 or v4.

Parameter	Required	Type	Description
path_style_always	Yes	Boolean	Path-style URLs for bucket access: enter <code>true</code> or <code>false</code> .
SQS_region	Yes	String	The SQS region.
SQS_queue	Yes	String	The name of the notification queue.
SQS_access_key	Yes	Base64-encoded string	The access key of the S3 credentials for access to the notification queue.
SQS_secret_key	Yes	Base64-encoded string	The secret key of the S3 credentials for access to the notification queue.

Response structure

None.

Example

Request example:

```
aws --endpoint-url https://10.08.1019 s3api put-bucket-replication --
bucket "hcpcs_bucket" --replication-configuration '{body}'
```

JSON request:

```
{
  "Role": "",
  "Rules": [{
    "ID": "sync_rule1_for_images",
    "Filter": {
      "Prefix": "/images/september/",
      "Tag": {
        "Key": "target",
        "Value": "cloud"
      }
    },
    "Status": "Enabled",
    "Destination": {
      "Bucket": "arn::sync-to::1.0::s3.amazonaws.com:443::us-east-1::redbucket::v4::true",
      "Account": "access_key, secret_key",
      "StorageClass": "STANDARD_IA"
    }
  },
  {
    {
```

```

    "ID": "sync_rule2_for_music",
    "Filter": {
      "Prefix": "/music/october/",
      "Tag": {
        "Key": "target",
        "Value": "cloud"
      }
    },
    "Status": "Enabled",
    "Destination": {
      "Bucket": "arn::sync-from::1.0::s3.amazonaws.com:443::us-east-1::bluebucket::v4::true::AWS_SQS::us-east-1::blackqueue::MTIzNA==:Njc4OQ==",
      "Account": "access_key, secret_key",
      "StorageClass": "STANDARD_IA"
    }
  ]
}'

```

Get bucket synchronization rules (GET bucket replication)

You can retrieve the synchronization rules for a bucket.

HTTP request syntax (URI)

```
aws --endpoint -url https://host_ip s3api get-bucket-replication --bucket "bucket"
```

Request structure

Not applicable.

Response structure

The response body is shown below:

```

{
  "ReplicationConfiguration": {
    "Role": "",
    "Rules": [
      {
        "Filter": {
          "And": {
            "Prefix": "string",
            "Tags": [
              {
                "Key": "string",
                "Value": "string"
              }
            ]
          }
        }
      }
    ]
  }
}

```

```

    .
    .
    .
  },
  "Status": "boolean",
  "Destination": {
    "Bucket": "access_settings",
  },
  "ID": "string",
}
],
}
}

```

Parameter	Required	Type	Description
Role	Yes	N/A	Not supported; empty.
Prefix	No	String	Prefix.
Key	No	String	Tag key.
Value	No	String	Tag value. Sets of prefixes and key-value pairs.
Status	Yes	Boolean	If <code>false</code> , rule is ignored.
Bucket	Yes	Base64-encoded JSON	Bucket access settings. S3 access and secret keys are masked.
ID	No	String	Unique identifier for rule, up to 255 characters.

Return codes

Status code	HTTP name	Description
200	OK	The request was executed successfully.
401	Unauthorized	Access was denied due to invalid credentials.

Example

Request example:

```
aws --endpoint-url https://10.08.1019 s3api get-bucket-replication --
bucket "hcpcs_bucket"
```

JSON response:

```
{
  "ReplicationConfiguration": {
    "Role": "",
    "Rules": [
      {
        "Filter": {
          "And": {
            "Prefix": "SQS",
            "Tags": [
              {
                "Value": "cloud",
                "Key": "target"
              }
            ]
          }
        },
        "Status": "Enabled",
        "Destination": {
          "Bucket": "arn::sync-from::1.0::s3.amazonaws.com:443::<AWS-Region>::hcpcs_bucket::V4::true::AWS_SQS::<SQS-Region>::<SQS-QUEUE-TopicName>",
          "ID": "mirrorBack_rule_for_images"
        }
      }
    ]
  }
}
```

Get object synchronization status

The synchronization status of an object is returned in metadata as part of the response to a GET object or HEAD object request.

For a GET object or HEAD object request, the synchronization functions return a replication status header in addition to the standard response metadata. This information is useful before deletion from a source bucket to verify synchronization.

Response header	Description
x-amz-replication-status	<p>Status of synchronization:</p> <ul style="list-style-type: none"> ▪ COMPLETED: For sync-to, all rules were successfully executed and the object was successfully synchronized. <p>Note: This status is also returned for objects that match a sync-to rule but were skipped because they are not the most recent version.</p> <ul style="list-style-type: none"> ▪ PENDING: For sync-to, one of the following: (1) a check is pending to see if the object needs synchronization; (2) the object needs synchronization, but the process is not complete. ▪ FAILED: For sync-to, the process has failed multiple times. To be synchronized, the object must be reloaded. ▪ REPLICA: For sync-from, the object is a replica created by Amazon S3.
(Header not in response)	The object did not match any rules.

Delete bucket synchronization rules (DELETE bucket replication)

You can delete S3 synchronization settings for buckets. This function is the same as in AWS S3.

HTTP request syntax (URI)

```
aws --endpoint -url https://host_ip s3api delete-bucket-replication --
bucket "bucket"
```

Request structure

None.

Response structure

Example

Request example:

```
aws --endpoint-url https://10.08.1019 s3api delete-bucket-replication --  
bucket "hcpcs_bucket"
```

Chapter 5: Notifications and user profiles

Within Hitachi Content Platform for cloud scale (HCP for cloud scale), the Object Storage Management application provides functions to display notifications, user profiles, online help, and API reference information.

HCP for cloud scale actively monitors the health and performance of the system and its resources and collects data for statistical reports. It provides:

- Alerts available through the GUI and REST APIs, as well as syslog notifications, so that you can respond quickly to events that require your attention
- Logs, metrics, and tracing to provide statistics and aid troubleshooting to improve performance and solve issues

Viewing alerts

You can use the Object Storage Management application or an API method to list active alerts.

An alert is a message to notify you of an event that may require your attention. Alerts are triggered by events, and remain active until the condition that caused the event is resolved. Once the condition is resolved, the alert is cleared.



Note: System alerts are generated by the System Management application to help you monitor overall system health and status of your HCP for cloud scale system. For information about system alerts and how to configure email notifications, see the Help in the System Management application.

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 panel listing alert text.

Related API method

```
POST /alert/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.

HCP for cloud scale alerts

An alert is a message that notifies you of a situation that requires your attention. Each alert corresponds to an event. Alerts have the severity INFO, WARNING, SEVERE, or CRITICAL.

Most alerts are generated by and reported through the System Management application. For information about those alerts, see the online help for the System Management application.

This table lists the alerts specific to HCP for cloud scale.

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	Service error: There is a critical issue with the Metadata Gateway database. Shutting down the Metadata Gateway Service.	A Metadata Gateway instance has encountered an issue and shut down. Use the System Management Services function Repair to restart it. If restarting the service doesn't resolve the issue, contact Support.

Severity	Message	Description
CRITICAL	Metadata-Coordination cannot communicate with Sentinel service to get state information	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.

HCP for cloud scale events

Most events are generated by and reported through the System Management application. For information about those events, see the online help for 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 ' <i>id</i> ' of type ' <i>job_type</i> ' updated with status ' <i>status</i> '	The policy configuration has changed, and the status of job <i>id</i> of type ' <i>job_type</i> ' is one of the following: <ul style="list-style-type: none"> ▪ ENABLED ▪ DISABLED
6006	INFO	Service Information: Job Configuration ' <i>id</i> ' started	Policy configuration for job <i>id</i> has started.
6006	INFO	Service Information: Lifecycle policy {CREATE UPDATE DELETE} bucket ' <i>bucket_name</i> '	The lifecycle policy for bucket <i>bucket_name</i> has been either created, updated, or deleted.

ID	Severity	Message	Description
6006	INFO	Service Information: Lifecycle policy deleted for bucket ' <i>bucket_name</i> '	The lifecycle policy for bucket <i>bucket_name</i> has been deleted.
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.

ID	Severity	Message	Description
6007	WARNING	Service Warning: Storage component 'id' is now INACCESSIBLE	The storage component <i>id</i> is inaccessible.HCP for cloud scale cannot read from or write to the storage component.
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.

User profiles

The profile icon, on the right end of the top navigation bar, provides access to these functions:

- Help: information about the Object Storage Management application
- REST API: information about the Object Storage Management APIs
- Logout: Log out of the Object Storage Management application and return to the **Login** page

Chapter 6: Services

Services perform functions essential to the health and function of the Hitachi Content Platform for cloud scale (HCP for cloud scale) system.

For example, the S3 Gateway service serves S3 API endpoints 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 these categories depending on what actions they perform:

- Product services enable HCP for cloud scale functions. For example, the S3 Gateway service serves S3 API endpoints 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

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

- **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 required number of instances on which a service must run to function properly.
 - The recommended number of instances on which a service should run. These are recommended minimums; if your system includes more instances, you should take advantage of them 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 **Container Memory** size larger than the **Max Heap Size** setting. For other services, you should not set the **Container Memory** size larger than the **Max Heap Size** setting.

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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter 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. Enter an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 1800m. 	Number of instances:	Required: 3 Recommended: 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 units:	10
	<p>Advanced Options:</p> <p>Compaction Frequency: How often database compaction operations are run. The options are Weekly (default) and Daily.</p> <p>Caution: Changing this setting can negatively affect the service. Use with caution.</p>	Persistent or floating?	Persistent
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes
<p>Chronos Job scheduling.</p>	<p>Container Options: Default</p> <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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter a decimal number. The default is 0.1. 	RAM needed per instance:	712 MB
		Service units:	1
	<p>Service Options</p> <ul style="list-style-type: none"> ▪ Heap settings: Maximum amount of memory to allocate to the Java heap for each instance of the service. Enter an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 512 MB. 	Number of instances:	<ul style="list-style-type: none"> ▪ 1 required ▪ 1 recommended
		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
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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter a decimal number. The default is 0.1. 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. Enter 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. 	RAM needed per instance:	2 GB
		Service units:	25
		Number of instances:	Required: 3 Recommended: All
		Persistent or floating?	Persistent
		Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	Yes
Kafka Stream processing platform for handling real-time data streams.	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. 	RAM needed per instance:	2 GB

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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter a decimal number. The default is 0.1. 	Service units:	5
	<p>Service Options</p> <ul style="list-style-type: none"> ▪ Heap settings: Maximum amount of memory to allocate to the Java heap for each instance of the service. Enter an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is -Xmx1800m-Xms512m. 	Number of instances:	Required: 3 Recommended: All
		Persistent or floating?	Persistent
		Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	Yes
<p>Logstash Logging.</p>	<p>Container Options: Default</p> <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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter a decimal number. The default is 0.1. 	RAM needed per instance:	700 MB
	<p>Service Options</p> <ul style="list-style-type: none"> ▪ Max heap: Maximum amount of memory to allocate to the Java heap for each instance of the service. Enter an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 512 MB. 	Number of instances:	Required: 1 Recommended: 1
		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
MAPI Gateway Serves MAPI endpoints.	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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter 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. Enter an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 64 MB. 	Service units:	5
		Number of instances:	Required: 1 Recommended: All
		Persistent or floating?	Floating
		Supports volume configuration?	No
		Single or multiple types?	Single
Metadata Cache Cache for system 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 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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter a decimal number. The default is 0.1. 	Service units:	10
	<p>Service Options</p> <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. Enter an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 512 MB. 	Number of instances:	Required: 1 Recommended: 1
		Persistent or floating?	Persistent
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	No
<p>Metadata Coordination</p> <p>Coordinates metadata gateway service instances, and coordinates scaling and balancing.</p>	<p>Container Options: Default</p> <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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter a decimal number. The default is 0.1. 	RAM needed per instance:	768 MB
	<p>Service Options</p> <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. Enter an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 64 MB. 	Number of instances:	Required: 1 Recommended: 1
		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
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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter a decimal number. The default is 0.1. 	RAM needed per instance:	768 MB
		Service units:	50
	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. Enter 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:	Required: 3 Recommended: All
		Persistent or floating?	Persistent
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes
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. 	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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter a decimal number. The default is 0.1. <p>Service Options</p> <ul style="list-style-type: none"> ▪ Prometheus Scrape Interval: The time interval between runs of the metrics collection task. Enter 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. Enter a path. The default is tsdb/. ▪ Prometheus Database Retention: The number of days to retain files. Enter an integer number of days. You can optionally specify the suffix d (days). The default is 15 days. 	Service units:	10
		Number of instances:	Required: 1 Recommended: All
		Persistent or floating?	Persistent
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes
<p>Metadata Policy Engine</p> <p>Executes asynchronous metadata updates.</p>	<p>Container Options: Default</p> <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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter a decimal number. The default is 0.1. <p>Service Options</p> <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. Enter 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:	768 MB
		Service units:	25
		Number of instances:	Required: 1 Recommended: 1
		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
S3 Gateway Serves S3 API endpoints 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. ▪ CPU: The relative CPU usage weight for the service's Docker container. Generally, a higher value means that the container received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter 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. Enter an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 512 MB. 	Number of instances:	Required: 1 Recommended: 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. Enter an integer. The default is 100 request headers. 	Persistent or floating?	Floating
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
<p>Tracing Agent</p> <p>Provides end-to-end distributed tracing for S3 API calls and MAPI calls.</p>	<p>Container Options: Default</p> <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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter a decimal number. The default is 0.1. 	RAM needed per instance:	768 MB
	<p>Service Options</p> <ul style="list-style-type: none"> ▪ Collector TChannel Hostname: Enter a host name. The default is localhost. ▪ Collector TChannel Port: Enter a port number. The default is 14267. 	Service units:	1
		Number of instances:	Required: All
		Persistent or floating?	Floating
		Supports volume configuration?	No
		Single or multiple types?	Single
Scalable?	Yes		
<p>Tracing Collector</p> <p>Provides end-to-end distributed tracing for S3 API calls and MAPI calls.</p>	<p>Container Options: Default</p> <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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter a decimal number. The default is 0.1. 	RAM needed per instance:	768 MB
	<p>Service Options</p> <ul style="list-style-type: none"> ▪ ElasticSearch Hostname: Enter a host name. The default is localhost. 	Number of instances:	Required: 1 Recommended: All

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
	<ul style="list-style-type: none"> ▪ ElasticSearch Port: Enter a port number. The default is 9200. ▪ Sampling Rate: The sampling rate for all clients implementing remote sampling. Enter 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. Enter 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. Enter a value from 1 to 365 days inclusive. The default is 60 days. 	Persistent or floating?	Floating
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes
<p>Tracing Query</p> <p>Provides end-to-end distributed tracing for S3 API calls and MAPI calls.</p>	<p>Container Options: Default</p> <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 received more CPU resources than other processes (including other service Docker containers) running on the instance. Enter a decimal number. The default is 0.1. 	RAM needed per instance:	768 MB
	<p>Service Options</p> <ul style="list-style-type: none"> ▪ ElasticSearch Hostname: Enter a host name. The default is localhost. ▪ ElasticSearch Port: Enter a port number. The default is 9200. 	Number of instances:	Required: 1 Recommended: All
		Persistent or floating?	Floating
		Supports volume configuration?	No
		Single or multiple types?	Single
		Scalable?	Yes

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.			
<p>Admin App</p> <p>The system management application.</p>	<p>Service Options</p> <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. Enter 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
<p>Cluster Coordination</p> <p>Manages hardware resource allocation.</p>	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

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
		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 units:	5
		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 TLS versions to use: <ul style="list-style-type: none"> ▪ TLS 1.0 ▪ TLS 1.1 ▪ TLS 1.2 ▪ TLS 1.3 	RAM needed per instance:	N/A
	SSL Ciphers: If you want to provide your own cipher suite, enter it here.	Number of instances:	N/A
		Service units:	1
	Custom Global Configuration: Select Enable Advanced Global Configuration to enable adding your own parameters to the HAProxy "global" section.	Persistent or floating?	Persistent

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
	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. Enter an integer number of bytes. You can optionally specify the suffixes k (kilobytes), m (megabytes), or g (gigabytes). The default is 256 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
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

Service name and description	Configuration settings (changes cause the service to redeploy)	Properties	
		Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	No
<p>Synchronization</p> <p>Coordinates service configuration settings and other information across instances.</p>	<p>Service Options</p> <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. Enter 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
<p>Watchdog</p> <p>Monitors the other System services and restarts them if necessary. Also responsible for initial system startup.</p>	<p>Service Options</p> <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. Enter 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
		Service units:	5

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. 	Persistent or floating?	Persistent
		Supports volume configuration?	Yes
		Single or multiple types?	Single
		Scalable?	No

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.

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. For information on scaling up services, see the System Management application online help.

Scaling down

Use the System Management application to remove a Metadata Gateway instance. For information on scaling down services, see the System Management application online help. 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.

Chapter 7: S3 User Credentials

HCP for cloud scale provides an application to obtain S3 user credentials.

Amazon Web Services uses security credentials, called S3 credentials, to authenticate and authorize data requests. The credentials consist of an access key and a secret key. Client applications that post S3 requests, such as uploading documents, reading documents, and adding buckets, to Hitachi Content Platform for cloud scale (HCP for cloud scale) also need these credentials. HCP for cloud scale provides a simple application, S3 User Credentials, to obtain these credentials for registered users of the system. It obtains an OAuth token from system services once you log in, and provides credentials on demand.

Obtaining S3 credentials

You can use the S3 User Credentials application or APIs to obtain S3 credentials.

The S3 User Credentials application retrieves credentials (access key and secret key) to access Amazon S3 bucket services. These credentials are linked to the username and password supplied in the API call. Thus, each unique user will retrieve a unique set of credentials.

If a user makes multiple, repeated API calls, only the last set of credentials remain active. Previously retrieved credentials become invalidated and will no longer work. Credentials expire automatically if a user changes his or her password held in the identity provider.

S3 application instructions

Use the S3 User Credentials application to obtain S3 access credentials.

Obtaining credentials nullifies any pre-existing S3 credentials you may already have.

To obtain S3 credentials:

Procedure

1. From the **Applications** page, select the application **S3 User Credentials**.
2. Click **Generate S3 Credentials**.
You are warned that any existing credentials for the logged-in user will be nullified.
3. Click **Generate**.



Note: If this step fails, your session may have timed out.

The application generates and displays an **Access Key** and a **Secret Key**.

4. Click **Copy**, next to the **Access Key** field, and paste the credential into the client application that you use to post S3 requests to HCP for cloud scale.
5. Click **Copy**, next to the **Secret Key** field, and paste the credential into the client application that you use to post S3 requests to HCP for cloud scale.

Related API method

```
POST /s3/user/generate_credentials
```

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.

Revoking S3 credentials

Amazon S3 credentials can be revoked by the associated user or by other users with appropriate permissions. If you have permissions you can revoke all Amazon S3 credentials belonging to a specific user. Use the endpoint `/user/list` to look up the ID of the user for whom you want to revoke credentials.

Related API methods

```
POST /user/list
```

```
POST /user/revoke_credentials
```

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.

Chapter 8: HCP for cloud scale APIs

The Hitachi Content Platform for cloud scale (HCP for cloud scale) system includes a set of 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 APIs

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 APIs

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 APIs

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

Object Storage Management APIs

The Object Storage Management application provides a RESTful HTTPS interface for the following functions:

- Managing storage components and Amazon Simple Storage Service (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 APIs are 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 RESTful APIs.



Note: The system configuration, management, and monitoring functions provided through the System Management application can be performed using the System Management APIs.

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

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

System Management APIs

The System Management application provides a RESTful HTTPS interface 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 RESTful APIs.

For information on the System Management APIs, see the System Management online help.

Support for Amazon S3 APIs

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 and has its own individual settings, such as ownership and lifecycle. Using HCP for cloud scale, users can perform common read and write operations 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 [S3 User Credentials \(on page 82\)](#).

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

Authentication and addressing operations

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.
Pre-signed URL	Fully implemented.

Service operations

Feature	Implementation differences
GET service (list buckets)	Fully implemented.

Bucket operations

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 an invalid bucket name, Amazon S3 performs an access check first and returns 403. HCP for cloud
DELETE Bucket	

Feature	Implementation differences
HEAD Bucket	scale returns 400 if the bucket name validation fails.
PUT Bucket ACL GET Bucket 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> <p>HCP for cloud scale does not mirror or mirror back ACLs or policies.</p>
List Multipart Uploads	Fully implemented.
GET Bucket Lifecycle (except transition action)	HCP for cloud scale supports the latest API for bucket lifecycle management. Old and deprecated V1.0 APIs are not supported.
PUT Bucket Lifecycle (except transition action)	HCP for cloud scale does not support Object Transition actions. Including these actions causes a Malformed XML exception.
DELETE Bucket Lifecycle (except transition action)	
PUT Bucket Replication	<p>Amazon S3 allows only one-to-one replication. HCP for cloud scale supports one-to-many mirroring and many-to-one mirroring back.</p> <p>For mirroring back, HCP for cloud scale supports one queue server, <code>AMAZON_SQS</code>.</p>

Feature	Implementation differences
GET Bucket Versioning	Version Listing Requests do not strictly comply to documented behavior for NextKeyMarker/NextVersionIdMarker. S3 documentation currently states that these values "specifies 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 (opaque to the caller); internally, HCP for cloud scale shares the same listing logic across all three listing types.
GET Bucket Object Versions	Fully implemented.
GET Bucket Location	You must be the bucket owner. The only location supported is <code>us-west-2</code> .

Object operations

Feature	Implementation differences
GET Object	If a lifecycle policy is configured for a bucket, HCP for cloud scale displays the expiration date of an object (in the <code>x-amz-expiration</code> header) fetched using the <code>?versionId</code> subresource. Amazon only displays this when performing unversioned GET requests.
HEAD Object	If a lifecycle policy is configured for a bucket, HCP for cloud scale displays the expiration date of an object (in the <code>x-amz-expiration</code> header) fetched using the <code>?versionId</code> subresource. Amazon only displays this when performing unversioned HEAD requests.
PUT Object	Amazon S3 limits the maximum file size for a single PUT or POST object operation to 5 GB. In HCP for cloud scale, this value is configurable, and the default is 5 GB.

Feature	Implementation differences
	<p>Content-Type Validations: Amazon S3 is extremely liberal in what is accepted for the Content-Type of an object. HCP for cloud scale adds additional checks for what is allowed.</p> <p>Mirroring and mirroring back is supported.</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>
Object and Version Encoding	<p>Amazon S3 Object and Version listing documentation mentions the ability to pass an encoding parameter (url). This is so the object name XML in the response to the client can be escaped to avoid names containing invalid XML characters. This encoding is only documented as applied to object names and not Owner/DisplayNames. Additionally, there is no mention of escaping for Bucket Listing requests. The Owner/DisplayName is a concern as there is a possibility that user display names may not be able to contain characters that could cause XML parsing issues. Amazon may be able to restrict this, though it does not currently return a display name for all regions. HCP for cloud scale uses Foundry IDPs, thus controlling restriction is not in the realm of HCP for cloud scale. 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 it in XML responses.</p>

Feature	Implementation differences
Object tagging	<p>Amazon S3 wraps eTags in double-quotes. For XML listings (v1 object, v2 object, version) it escapes these, for example:</p> <pre><ETag>&quot;32c81604d07395b1aa39a7e206c3af06\$quot;</ETag></pre> <p>It's not necessary for HCP for cloud scale to perform this because double-quotes do not need to be escaped within content, only attributes.</p>
	<p>Expiration Date URL Encoding (x-amz-expiration header)</p> <p>The <code>RuleID</code> portion of the x-amz-expiration header is URL-encoded by HCP for cloud scale using the same encoding strategy that Amazon suggests for V4 authentication. This may result in encoded strings that do not exactly match how Amazon encodes RuleIDs in general. However, decoding them should always return the original string.</p>
	<p>HCP for cloud scale mirrors and mirrors back object tagging and tag updates.</p>
GET Object ACL	<p>Mirroring and mirroring back is not supported.</p>
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>

Feature	Implementation differences
DELETE Object	Mirroring and mirroring back of deletion of an object or a specific version of an object is not supported.
DELETE Multiple Objects	Fully implemented. Mirroring and mirroring back is not supported.
POST Object	Fully implemented. Amazon S3 limits the maximum file size for a single PUT or POST object operation to 5 GB. In HCP for cloud scale, this value is configurable, and the default is 5 GB. Mirroring and mirroring back is supported.
Initiate/Complete/Abort Multipart Upload	Fully implemented. Mirroring and mirroring back is supported.
Upload Part	Fully implemented.
List Multipart Uploads	Fully implemented.

Unsupported S3 APIs

The following lists are the unsupported Amazon S3 API features.

Authentication API

- Authentication v2 (deprecated by AWS)

Bucket APIs

- 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 (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 APIs

- Options Object
- GET/POST Object Torrent
- SELECT Object Content (SQL)

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