



Hitachi Content Platform

Managing the Default Tenant and Namespace

© 2009–2015 Hitachi Data Systems Corporation. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi Data Systems Corporation (hereinafter referred to as “Hitachi Data Systems”).

Hitachi Data Systems reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. This document contains the most current information available at the time of publication. When new and/or revised information becomes available, this entire document will be updated and distributed to all registered users.

Some of the features described in this document may not be currently available. Refer to the most recent product announcement or contact Hitachi Data Systems for information about feature and product availability.

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

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

- a) Acquiring the relevant consents as may be required under local privacy laws or otherwise from employees and other individuals to access relevant data; and
- b) Ensuring that data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

Archivas, Essential NAS Platform, HiCommand, Hi-Track, ShadowImage, Tagmaserve, Tagmasoft, Tagmasolve, Tagmastore, TrueCopy, Universal Star Network, and Universal Storage Platform are registered trademarks of Hitachi Data Systems Corporation.

AIX, AS/400, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, ESCON, FICON, FlashCopy, IBM, Lotus, MVS, OS/390, RS6000, S/390, System z9, System z10, Tivoli, VM/ESA, z/OS, z9, z10, zSeries, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

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

Microsoft product screen shots reprinted with permission from Microsoft Corporation.

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.



Contents

| | |
|------------------------------------------------------------------------|-----------|
| Preface | ix |
| Intended audience | ix |
| Product version | ix |
| Syntax notation | x |
| Related documents | x |
| Getting help | xiii |
| Comments | xiii |
| 1 Introduction to Hitachi Content Platform | 1 |
| About Hitachi Content Platform | 2 |
| Object-based storage | 2 |
| Namespaces and tenants | 3 |
| Object representation | 5 |
| Namespace access | 5 |
| Namespace access protocols | 5 |
| HCP metadata query API | 6 |
| HCP Search Console | 6 |
| HCP Data Migrator | 8 |
| 2 Tenant and namespace properties | 9 |
| Data protection level | 10 |
| Cryptographic hash algorithm | 10 |
| Retention mode | 11 |
| Index setting | 12 |
| Retention-related properties | 12 |
| Ownership and permission changes for objects under retention | 12 |
| Custom metadata operations for objects under retention | 13 |
| XML checking for custom metadata | 13 |
| Compatibility properties | 14 |

| | |
|---------------------------------------------------|-----------|
| Disposition | 14 |
| Data access permission masks | 15 |
| Replication | 17 |
| Replication benefits | 18 |
| Replication implementation | 19 |
| Replication collision handling | 19 |
| Object content collisions | 20 |
| System metadata collisions | 21 |
| Custom metadata collisions | 24 |
| Retention class collisions | 25 |
| Service plans | 26 |
| 3 General administrative information..... | 27 |
| Tenant Management Console | 28 |
| Console access | 28 |
| Console sessions | 31 |
| Tenant Management Console URL | 32 |
| Logging in | 33 |
| Using the Tenant Management Console | 35 |
| Refreshing pages | 35 |
| Submitting changes | 36 |
| Viewing HCP documentation | 36 |
| Changing your password | 36 |
| Logging out | 37 |
| Administrative responsibilities | 38 |
| 4 Managing the tenant..... | 43 |
| About the tenant Overview page | 44 |
| Tenant statistics | 44 |
| Major tenant events | 44 |
| Tenant alerts | 45 |
| Tenant contact information | 45 |
| Tenant permission mask | 46 |
| Tenant description | 46 |
| Configuring the tenant | 46 |
| Changing the tenant contact information | 47 |
| Changing the tenant permission mask | 48 |
| Changing the tenant description | 49 |
| Monitoring the tenant | 50 |
| Viewing the complete tenant event log | 51 |
| Viewing the tenant security log | 52 |
| Viewing the tenant compliance log | 52 |

| | |
|---------------------------------------------------------------------------|-----------|
| Understanding log messages | 53 |
| Managing the message list | 54 |
| Enabling syslog logging | 55 |
| Enabling SNMP logging | 55 |
| Configuring email notification | 56 |
| Enabling email notification. | 57 |
| Testing email notification. | 57 |
| Constructing the email message template | 58 |
| Specifying email recipients. | 61 |
| Monitoring replication | 63 |
| High-level view of replication | 63 |
| Detailed view of replication | 64 |
| Up-to-date-as-of time | 64 |
| Data transmission rate | 64 |
| Operation rate | 65 |
| 5 Managing the namespace | 67 |
| About the namespace Overview page. | 68 |
| Namespace URL | 68 |
| Objects section. | 68 |
| Usage section | 69 |
| Major namespace events. | 70 |
| Namespace alerts | 71 |
| Namespace services, service plan, retention mode, hash algorithm, and DPL | 71 |
| Namespace permission mask | 72 |
| Namespace description | 72 |
| Configuring the namespace | 72 |
| Changing the namespace permission mask | 73 |
| Changing the namespace description | 74 |
| Changing retention-related settings | 75 |
| Enabling or disabling XML checking for custom metadata | 76 |
| Changing compatibility settings | 77 |
| Changing disposition settings. | 78 |
| Changing replication options | 78 |
| Changing the service plan | 80 |
| Changing the retention mode | 80 |
| Monitoring the namespace | 81 |
| Viewing the complete namespace event log | 82 |
| Viewing the namespace compliance log | 82 |
| Working with irreparable objects | 83 |
| 6 Configuring the namespace access protocols | 85 |

| | |
|---------------------------------------------------------------------------|-----|
| Namespace access protocol configuration | 86 |
| Specifying IP addresses in Allow and Deny lists | 86 |
| Adding and removing entries in Allow and Deny lists | 87 |
| Valid Allow and Deny list entries | 87 |
| Allow and Deny list handling | 87 |
| Specifying default ownership and permissions | 89 |
| Initial values for object owners, groups, and permissions | 90 |
| Specifying owner, group, and permission defaults | 90 |
| Octal permission values. | 91 |
| Configuring the HTTP and WebDAV protocols | 92 |
| HTTP and WebDAV protocol configuration | 92 |
| Enabling HTTP and WebDAV access to the namespace. | 92 |
| Configuring the CIFS protocol | 96 |
| CIFS protocol configuration | 96 |
| CIFS access to the namespace. | 97 |
| UID and GID for objects stored through CIFS | 98 |
| CIFS case sensitivity | 99 |
| Enabling CIFS access to the namespace | 101 |
| Working with username mapping files | 103 |
| Configuring the NFS protocol. | 104 |
| NFS protocol configuration | 104 |
| Enabling NFS access to the namespace. | 105 |
| Configuring the SMTP protocol. | 105 |
| SMTP protocol configuration | 105 |
| Email retention setting | 106 |
| Enabling SMTP access to the namespace | 107 |
| Configuring Microsoft Exchange for email archiving through SMTP | 109 |
| Configuring Microsoft Exchange 2003. | 110 |
| Configuring Microsoft Exchange 2007 | 111 |
| Configuring Microsoft Exchange 2010. | 112 |
| Configuring the NDMP protocol | 112 |
| Backup and restore operations. | 112 |
| OpenPGP format. | 113 |
| Signing and encryption keys | 114 |
| Backup performance considerations | 114 |
| NDMP protocol configuration | 115 |
| Enabling NDMP access to the namespace | 115 |
| Working with signing and encryption keys. | 118 |
| Downloading a key | 118 |
| Uploading a signing key | 119 |
| Uploading an encryption key | 119 |
| Deleting a signing key. | 120 |
| Using third-party applications with NDMP | 120 |

| | | |
|----------|------------------------------------------------------------------------------|------------|
| 7 | Managing search and indexing..... | 121 |
| | About search and indexing | 122 |
| | Content classes and content properties | 123 |
| | Metadata query engine indexing of custom metadata. | 124 |
| | Content class and content property workflow | 125 |
| | Content property definitions | 127 |
| | Content property names | 128 |
| | Content property expressions | 129 |
| | Content property data types | 131 |
| | Formats for the integer and float data types | 132 |
| | Datetime data type formats | 135 |
| | Multivalued content properties. | 137 |
| | Content properties extracted from sample XML | 137 |
| | Content property files | 139 |
| | About the Search page | 142 |
| | Managing the content class list | 142 |
| | Understanding the content property list for a content class. | 143 |
| | Creating a content class | 144 |
| | Managing content properties for a content class | 144 |
| | Adding, modifying, and deleting content properties | 144 |
| | Adding content properties individually | 145 |
| | Extracting content properties from sample XML | 146 |
| | Importing content properties from a content property file. | 146 |
| | Testing content properties. | 147 |
| | Exporting content properties | 147 |
| | Changing associations between the default namespace and content classes. . . | 148 |
| | Reindexing the default namespace from the Search page. | 149 |
| | Renaming a content class | 150 |
| | Deleting a content class | 150 |
| | Managing search and indexing for the default namespace | 151 |
| | Setting search and indexing options | 152 |
| | Reindexing the default namespace from the Search panel | 153 |
| 8 | Working with retention classes..... | 155 |
| | About retention classes. | 156 |
| | Understanding the retention class list. | 157 |
| | Creating a retention class | 158 |
| | Modifying a retention class | 159 |
| | Deleting a retention class | 159 |
| 9 | Using privileged delete | 161 |

| | |
|-----------------------------------------------------------------------------------|------------|
| About privileged delete | 162 |
| Object specification | 162 |
| Performing a privileged delete | 163 |
| 10 Downloading HCP Data Migrator | 165 |
| HCP-DM system requirements | 166 |
| Downloading the HCP-DM installation file | 166 |
| A Tenant Management Console alerts | 167 |
| Tenant Overview page alerts | 168 |
| Namespace Overview page alerts. | 169 |
| Search page alert | 170 |
| B Tenant log messages | 171 |
| C Browser configuration for single sign-on with Active Directory | 177 |
| Configuring Windows Internet Explorer for single sign-on. | 178 |
| Configuring Mozilla Firefox for single sign-on. | 179 |
| Glossary | 181 |
| Index | 199 |



Preface

This book explains how to use **Hitachi Content Platform (HCP)** to monitor and manage the default tenant and namespace in an HCP system. It presents the concepts and instructions you need to configure the tenant and namespace, set up the namespace access protocols, manage the search feature, and download the HCP Data Migrator installation files. It also covers activities you can perform to keep the namespace in compliance with local regulations.

This book does not address HCP tenants and namespaces. For information on these entities, see *Managing a Tenant and Its Namespaces*.



Note: Throughout this book, the word *Unix* is used to represent all UNIX[®]-like operating systems (such as UNIX itself or Linux[®]).

Intended audience

This book is intended for HCP administrators who configure, monitor, and manage the default tenant and namespace. It assumes you are familiar with your client operating system and the web browser you use to run the HCP Tenant Management Console.

Product version

This book applies to release 7.1 of HCP.

Syntax notation

The table below describes the conventions used for the syntax of commands, expressions, URLs, and object names in this book.

| Notation | Meaning | Example |
|-----------------|------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| boldface | Type exactly as it appears in the syntax (if the context is case insensitive, you can vary the case of the letters you type) | This book shows: https://default.hcp-domain-name:8000 You enter: https://default.hcp-ma.example.com:8000 |
| <i>italics</i> | Replace with a value of the indicated type | |

Related documents

The following documents contain additional information about Hitachi Content Platform:

- *Administering HCP* — This book explains how to use an HCP system to monitor and manage a digital object repository. It discusses the capabilities of the system, as well as its hardware and software components. The book presents both the concepts and instructions you need to configure the system, including creating the tenants that administer access to the repository. It also covers the processes that maintain the integrity and security of the repository contents.
- *Managing a Tenant and Its Namespaces* — This book contains complete information for managing the HCP tenants and namespaces created in an HCP system. It provides instructions for creating namespaces, setting up user accounts, configuring the protocols that allow access to namespaces, managing search and indexing, and downloading installation files for HCP Data Migrator. It also explains how to work with retention classes and the privileged delete functionality.
- *Replicating Tenants and Namespaces* — This book covers all aspects of tenant and namespace replication. Replication is the process of keeping selected tenants and namespaces in two or more HCP systems in sync with each other to ensure data availability and enable disaster recovery. The book describes how replication works, contains instructions for working with replication links, and explains how to manage and monitor the replication process.

- *HCP Management API Reference* — This book contains the information you need to use the HCP management API. This RESTful HTTP API enables you to create and manage tenants and namespaces programmatically. The book explains how to use the API to access an HCP system, specify resources, and update and retrieve resource properties.
- *Using a Namespace* — This book describes the properties of objects in HCP namespaces. It provides instructions for accessing namespaces by using the HTTP, WebDAV, CIFS, and NFS protocols for the purpose of storing, retrieving, and deleting objects, as well as changing object metadata such as retention and shred settings. It also explains how to manage namespace content and view namespace information in the Namespace Browser.
- *Using the HCP HS3 API* — This book contains the information you need to use the HCP HS3 API. This S3™-compatible, RESTful, HTTP-based API enables you to work with buckets and objects in HCP. The book introduces the HCP concepts you need to understand in order to use HS3 effectively and contains instructions and examples for each of the bucket and object operations you can perform with HS3.
- *Using the HCP OpenStack Swift API* — This book contains the information you need to use the HCP OpenStack Swift API. This S3™-compatible, RESTful, HTTP-based API enables you to work with containers and objects in HCP. The book introduces the HCP concepts you need to understand in order to use HSwift effectively and contains instructions and examples for each of the container and object operations you can perform with HSwift.
- *Using the Default Namespace* — This book describes the file system HCP uses to present the contents of the default namespace. It provides instructions for accessing the namespace by using the HCP-supported protocols for the purpose of storing, retrieving, and deleting objects, as well as changing object metadata such as retention and shred settings.
- *HCP Metadata Query API Reference* — This book describes the HCP metadata query API. This RESTful HTTP API enables you to query namespaces for objects that satisfy criteria you specify. The book

explains how to construct and perform queries and describes query results. It also contains several examples, which you can use as models for your own queries.

- *Searching Namespaces* — This book describes the HCP Search Console (also called the Metadata Query Engine Console). It explains how to use the Console to search namespaces for objects that satisfy criteria you specify. It also explains how to manage and manipulate queries and search results. The book contains many examples, which you can use as models for your own searches.
- *Using HCP Data Migrator* — This book contains the information you need to install and use HCP Data Migrator (HCP-DM), a utility that works with HCP. This utility enables you to copy data between local file systems, namespaces in HCP, and earlier HCAP archives. It also supports bulk delete operations and bulk operations to change object metadata. Additionally, it supports associating custom metadata and ACLs with individual objects. The book describes both the interactive window-based interface and the set of command-line tools included in HCP-DM.
- *Installing an HCP System* — This book provides the information you need to install the software for a new HCP system. It explains what you need to know to successfully configure the system and contains step-by-step instructions for the installation procedure.
- *Deploying an HCP-VM System* — This book contains all the information you need to install and configure an HCP-VM system. The book also includes requirements and guidelines for configuring the VMWare[®] environment in which the system is installed.
- *Third-Party Licenses and Copyrights* — This book contains copyright and license information for third-party software distributed with or embedded in HCP.
- *HCP-DM Third-Party Licenses and Copyrights* — This book contains copyright and license information for third-party software distributed with or embedded in HCP Data Migrator.

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

Getting help

The Hitachi Data Systems[®] customer support staff is available 24 hours a day, seven days a week. If you need technical support, call:

- United States: (800) 446-0744
- Outside the United States: (858) 547-4526



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

Comments

Please send us your comments on this document:

HCPDocumentationFeedback@hds.com

Include the document title, number, and revision, and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Data Systems.

Comments

Thank you!

Introduction to Hitachi Content Platform

Hitachi Content Platform (HCP) is a distributed storage system designed to support large, growing repositories of fixed-content data.

HCP stores objects that include both data and metadata that describes that data. HCP presents these objects as files in a standard directory structure.

An HCP repository is partitioned into namespaces. Each namespace consists of a distinct logical grouping of objects with its own directory structure. Namespaces are owned and managed by tenants.

HCP provides access to objects through a variety of industry-standard protocols, as well as through various HCP-specific interfaces.

This chapter contains an overview of Hitachi Content Platform.

About Hitachi Content Platform

Hitachi Content Platform is the distributed, fixed-content, data storage system from Hitachi Data Systems®. HCP provides a cost-effective, scalable, easy-to-use repository that can accommodate all types of data, from simple text files to medical images to multigigabyte database images.

A **fixed-content storage system** is one in which the data cannot be modified. HCP uses write-once, read-many (WORM) storage technology and a variety of policies and internal processes to ensure the integrity of the stored data and the efficient use of storage capacity. HCP also provides easy access to the repository for adding, retrieving, and deleting or shredding data.

Object-based storage

HCP stores **objects** in a repository. Each object permanently associates data HCP receives (for example, a document, an image, or a movie) with information about that data, called **metadata**.

An object encapsulates:

- **Fixed-content data** — An exact digital reproduction of data as it existed before it was stored in HCP. Once it's in the repository, this fixed-content data cannot be modified.
- **System metadata** — System-managed properties that describe the fixed-content data (for example, its size and creation date). System metadata includes policies, such as retention and data protection level, that influence how transactions and internal processes affect the object.
- **Custom metadata** — Optional metadata that a user or application provides to further describe the object. Custom metadata is typically specified in XML format.

You can use custom metadata to create self-describing objects. Users and applications can use this metadata to understand and repurpose object content.

HCP also stores directories and symbolic links. These items have system metadata but no fixed-content data or custom metadata.

HCP supports appendable objects. An **appendable object** is one to which data can be added after it has been successfully stored. Appending data to an object does not modify the original fixed-content data. Once the new data is added to the object, that data also cannot be modified.

Namespaces and tenants

An HCP repository is partitioned into namespaces. A **namespace** is a logical grouping of objects such that the objects in one namespace are not visible in any other namespace.

Namespaces provide a mechanism for separating the data stored for different applications, business units, or customers. For example, you could have one namespace for accounts receivable and another for accounts payable.

Namespaces also enable operations to work against selected subsets of objects. For example, you could perform a query that targets the accounts receivable and accounts payable namespaces but not the employees namespace.

HCP and default namespaces

An HCP system can have a maximum of 10,000 locally defined namespaces, including one special namespace called the **default namespace**. Applications are typically written against namespaces other than the default; these namespaces are called **HCP namespaces**. The default namespace is most often used with applications that existed before release 3.0 of HCP.



Note: Replication can cause an HCP system to have more than 10,000 namespaces. For information on replication, see ["Replication"](#) on page 17.

The table below outlines the major differences between HCP namespaces and the default namespaces.

| Feature | HCP namespaces | Default namespace |
|---------------------------------------------|----------------|-------------------|
| Storage usage quotas | ✓ | |
| Object ownership (not related to POSIX UID) | ✓ | |
| Access control lists (ACLs) for objects | ✓ | |
| Object versioning | ✓ | |
| Multiple custom metadata annotations | ✓ | |

(Continued)

| Feature | HCP namespaces | Default namespace |
|--------------------------------------------------------------------------|----------------|-------------------|
| Namespace ownership by users | ✓ | |
| RESTful HTTP/HTTPS API for data access | ✓ | |
| Non-RESTful HTTP/HTTPS protocol for data access | | ✓ |
| Data access authentication with HTTP/HTTPS | ✓ | |
| RESTful HS3 API for data access (compatible with Amazon [®] S3) | ✓ | |
| NDMP protocol for backup and restore | | ✓ |

Tenants

Namespaces are owned and managed by administrative entities called **tenants**. A tenant typically corresponds to an organization, such as a company or a division or department within a company.

HCP supports two types of tenants:

- The **default tenant**, which owns the default namespace and only that namespace. An HCP system can have only one default tenant.
- **HCP tenants**, which own HCP namespaces. An HCP system can have multiple HCP tenants, each of which can own multiple namespaces.

An HCP system can have a maximum of 1,000 locally defined tenants, including the default tenant.



Note: Replication can cause an HCP system to have more than 1,000 tenants. For information on replication, see [“Replication”](#) on page 17.

An HCP system has both system-level and tenant-level administrators:

- **System-level administrators** are concerned with monitoring the HCP system hardware and software, monitoring overall repository usage, configuring features that apply across the HCP system, and managing system-level users.
- **Tenant-level administrators** are concerned with monitoring namespace usage at the tenant and namespace level, configuring individual tenants and namespaces, and controlling access to namespaces.

System-level administrators create tenants. Tenant-level administrators create HCP namespaces. The default namespace is created automatically when the default tenant is created.

Object representation

HCP includes a standard POSIX file system called HCP-FS that represents each object in the default namespace as a set of files. One of these files has the same name as the object. This file contains the fixed-content data. When downloaded or opened, this file has the same content as the originally stored item.

The other files that HCP-FS presents contain object metadata. These files, which are either plain text or XML, are called **metafiles**.

All files containing fixed-content data are in a directory hierarchy headed by the `fcfs_data` directory. All metafiles are in a directory hierarchy headed by the `fcfs_metadata` directory. With this view of objects as conventional files and directories, HCP supports routine file-level calls and enables users and applications to find fixed-content data in familiar ways.

For more information on HCP-FS and object metadata, see *Using the Default Namespace*.

Namespace access

HCP supports access to the default namespace through:

- Several industry-standard protocols
- The HCP metadata query API
- The HCP Search Console
- HCP Data Migrator

Namespace access protocols

HCP supports access to the default namespace through several industry-standard protocols. The HTTP, WebDAV, CIFS, and NFS protocols support various operations: storing data, creating directories, viewing object data and metadata, viewing directories, modifying certain metadata, and deleting objects. You can use these protocols to access the namespace with a web browser, third-party applications, Windows[®] Explorer, and other native Windows and Unix tools.

HCP allows special-purpose access to the default namespace through the SMTP protocol. This protocol is used only for storing email.

HCP supports the NDMP protocol for backing up and restoring objects.

For more information on these protocols, see [Chapter 6, “Configuring the namespace access protocols,”](#) on page 85.

HCP metadata query API

The **HCP metadata query API** lets you search HCP for objects that meet specified criteria. The API supports two types of queries:

- **Object-based queries** search for objects based on object metadata. This includes both system metadata and the content of custom metadata. The query criteria can also include the object location (that is, the namespace and/or directory that contains the object). These queries use a robust query language that lets you combine search criteria in multiple ways.

Object-based queries search only for objects that currently exist in the repository.

- **Operation-based queries** search not only for objects currently in the repository but also for information about objects that have been deleted by a user or application or deleted through disposition. Criteria for operation-based queries can include object status (for example, created or deleted), change time, index setting, and location.

The metadata query API returns object metadata only, not object data. The metadata is returned either in XML format, with each object represented by a separate element, or in JSON format, with each object represented by a separate name/value pair. For queries that return large numbers of objects, you can use paged requests.

For information on using the metadata query API, see *HCP Metadata Query API Reference*.

HCP Search Console

The **HCP Search Console** is an easy-to-use web application that lets you search for and manage objects based on specified criteria. For example, you can search for objects that were stored before a certain date or that are larger than a specified size. You can then delete the objects listed in

the search results or prevent those objects from being deleted. Similar to the metadata query API, the Search Console returns only object metadata, not object data.

By offering a structured environment for performing searches, the Search Console facilitates e-discovery, namespace analysis, and other activities that require the user to examine the contents of namespaces. From the Search Console, you can:

- Open objects
- Perform bulk operations on objects
- Export search results in standard file formats for use as input to other applications
- Publish feeds to make search results available to web users

The Search Console works with either of these two search facilities:

- The **HCP metadata query engine** — This facility is integrated with HCP and works internally to perform searches and return results to the Search Console. The metadata query engine is also used by the metadata query API.



Note: When working with the metadata query engine, the Search Console is called the **Metadata Query Engine Console**.

- The **Hitachi Data Discovery Suite (DDS) search facility** — This facility interacts with HDDS, which performs searches and returns results to the HCP Search Console. HDDS is a separate product from HCP.

The Search Console can use only one search facility at any given time. The search facility is selected at the HCP system level. If no facility is selected, the HCP system does not support use of the Search Console to search namespaces.

Each search facility maintains its own index of objects in each search-enabled namespace and uses this index for fast retrieval of search results. The search facilities automatically update their indexes to account for new and deleted objects and changes to object metadata.

To learn which search facility is enabled for the HCP system, contact your HCP system administrator. For more information on the search indexes, see [“About search and indexing”](#) on page 122. For information on using the Search Console, see *Searching Namespaces*.

HCP Data Migrator

HCP Data Migrator (HCP-DM) is a high-performance, multithreaded, client-side utility for viewing, copying, and deleting data. With HCP-DM, you can:

- Copy objects, files, and directories between the local file system, HCP namespaces, default namespaces, and earlier HCAP archives
- Delete individual objects, files, and directories and perform bulk delete operations
- View the content of objects and files
- Rename files and directories on the local file system
- View object, file, and directory properties
- Change system metadata for multiple objects in a single operation
- Add, replace, or delete custom metadata for objects
- Create empty directories

HCP-DM has both a graphical user interface (GUI) and a command-line interface (CLI).

For information on downloading HCP-DM, see [Chapter 10, “Downloading HCP Data Migrator.”](#) on page 165. For information on installing and using HCP-DM, see *Using HCP Data Migrator*.

Tenant and namespace properties

Tenants and namespaces have certain properties that affect how they operate. Some of these properties are set when the tenant or namespace is created. Others are set after creation. Some can be modified after they are initially set; others cannot.

This chapter addresses these properties of tenants and namespaces:

- Data protection level
- Cryptographic hash algorithm
- Retention mode
- Index setting
- Whether ownership and permission changes are allowed for objects under retention
- Which custom metadata operations are allowed for objects under retention
- XML checking for custom metadata
- Compatibility with other storage products
- Disposition
- Data access permission masks
- Replication

- Service plans

For information on the search and indexing properties of tenants and namespaces, see [Chapter 7, “Managing search and indexing.”](#) on page 121.



Note: In this chapter, in the context of namespace access, the term *users* means both people and applications.

Data protection level

Each namespace has a **data protection level (DPL)** that specifies how many copies of each object HCP must maintain. HCP stores each copy of an object in a different location. All but one of these copies can become unavailable (for example, due to a hardware outage) without affecting access to the object.

The DPL for a namespace can be one (if allowed by the HCP system configuration), two, three, four, or dynamic. The highest allowed DPL is also determined by the HCP system configuration.

When the DPL is set to dynamic, the namespace uses the current HCP system-level DPL setting. Typically, this is the optimal DPL for the system configuration. If the system-level DPL setting changes, HCP adjusts the number of copies of objects in the namespace to match the new setting.

The DPL affects the amount of storage used when data is added to the namespace. With a DPL of one, HCP stores only one copy of the object. With a DPL of two, HCP stores two copies, thereby using twice as much storage.

The DPL for the default namespace is set when the service plan is configured. You can change this setting at any time.

Cryptographic hash algorithm

At object creation, HCP uses a cryptographic hash algorithm to calculate a hash value for the object from the object data. HCP then uses this hash value to ensure the integrity of the object over time.

HCP supports these cryptographic hash algorithms:

MD5
SHA-1
SHA-256
SHA-384
SHA-512
RIPEMD-160

The cryptographic hash algorithm HCP uses for the default namespace is set when the namespace is created. Once set, it cannot be changed.

Retention mode

Each object has a **retention setting** that specifies how long the object must remain in the namespace before it can be deleted; this duration is called the **retention period**. While an object cannot be deleted due to its retention setting, it is said to be **under retention**.

Retention mode is a property of a namespace that affects which operations are allowed on objects under retention. A namespace can be in either of two retention modes:

- In **compliance mode**, objects that are under retention cannot be deleted through any mechanism. Additionally, the duration of a retention class cannot be shortened, and retention classes cannot be deleted.
- In **enterprise mode**, you can delete objects under retention if your user account includes the compliance role. This is called **privileged delete**.

Also in enterprise mode, you can shorten the duration of a retention class, and you can delete retention classes.

The retention mode for the default namespace is set when the namespace is created. You can change the retention mode from enterprise to compliance but cannot do the reverse.

For information on:

- Changing the retention mode of the namespace, see [“Changing the retention mode”](#) on page 80
- Retention classes, see [Chapter 8, “Working with retention classes.”](#) on page 155

- Privileged delete, see [Chapter 9, “Using privileged delete,”](#) on page 161

Index setting

Each object in the repository has an index setting that is either **true** or **false**. This setting is present even if the namespace containing the object is not search-enabled or indexed.

The metadata query engine uses the index setting for an object to determine whether to index custom metadata for that object. Metadata query API requests can use the index setting as a search criterion. Additionally, third-party applications can use this setting for their own purposes.

Retention-related properties

The default namespace has two properties that affect objects under retention:

- Whether ownership and permission changes are allowed
- Which custom metadata operations are allowed

Ownership and permission changes for objects under retention

Each object in the default namespace has:

- An owner that's identified by a POSIX user ID
- An owning group that's identified by a POSIX group ID
- A set of POSIX permissions that determine which operations (read, write, or execute) the owner, members of the owning group, and others can perform on the object

A namespace-level setting determines whether users can change the owner, owning group, and permissions for objects under retention. When the namespace is created, these changes are not allowed. You can change this setting at any time.

For information on changing the setting for ownership and permission changes for objects under retention, see [“Changing retention-related settings”](#) on page 75. For more information on object ownership and permissions in general, see *Using the Default Namespace*.

Custom metadata operations for objects under retention

Custom metadata is user-supplied information that describes an object in a namespace. For objects that are not under retention, users can add, replace, and delete custom metadata as needed. For objects that are under retention, the operations allowed for custom metadata are determined by a namespace-level setting.

You can configure the default namespace to:

- Allow custom metadata to be added, replaced, and deleted for objects under retention
- Allow custom metadata to be added for objects under retention, but not replaced or deleted
- Disallow all custom metadata operations for objects under retention

When the default namespace is created, only the addition of custom metadata is allowed for objects under retention. You can change this setting at any time.

For information on changing custom metadata handling, see [“Changing retention-related settings”](#) on page 75. For more information on custom metadata in general, see *Using the Default Namespace*.

XML checking for custom metadata

By default, when custom metadata is added to or replaced in the default namespace, HCP checks whether it’s well-formed XML. If the XML is not well-formed, HCP rejects it

You can choose to allow users to provide custom metadata in non-XML formats (for example, as thumbnail images to accompany large objects with image content). In this case, you need to disable custom metadata XML checking for the namespace so that HCP accepts non-XML custom metadata.

Regardless of how custom metadata is specified, it is represented in the metadata file structure for an object by the `custom-metadata.xml` metafile.

XML checking applies only when custom metadata is added to or replaced in the namespace. It does not apply to custom metadata already in the namespace.

You can enable or disable custom metadata XML checking for the default namespace at any time. For instructions on doing this, see [“Enabling or disabling XML checking for custom metadata”](#) on page 76. For more information on custom metadata in general, see *Using the Default Namespace*.



Note: If the custom metadata XML for an object has a very large number of different elements and attributes, HCP may determine that the XML is not well-formed, even if it is, and reject the file. If this happens, the user can try restructuring the XML so that it includes fewer different elements and attributes. Alternatively, you can temporarily disable custom metadata XML checking to allow that XML to be stored in the namespace.

Compatibility properties

The default namespace has two features that support HCP compatibility with other storage products:

- Synchronization of POSIX **atime** values with retention settings. For information on the effects of this option, see *Using the Default Namespace*.
- Creation of appendable objects.



Note: Users can create and add data to appendable objects only through the CIFS and NFS protocols.

When the default namespace is created, both of these features are disabled. You can enable or disable these features at any time.

For information on enabling and disabling these features, see [“Changing compatibility settings”](#) on page 77.

Disposition

Disposition is the automatic deletion of objects. Disposition can be enabled for:

- Objects with expired retention periods. To be eligible for disposition, an object must have a retention setting that’s either:
 - A date in the past

- A retention class with automatic deletion enabled that results in a calculated expiration date in the past
- Objects flagged as replication collisions.

Disposition has the benefit of automatically freeing HCP storage space for the creation of more objects. Without disposition, users need to explicitly delete qualified objects to free the occupied space.

Disposition is enabled on a per-namespace basis. When the default namespace is created, this feature is disabled. You can change this setting at any time.

For information on enabling and disabling disposition, see [“Changing disposition settings”](#) on page 78. For information on retention classes, see [Chapter 8, “Working with retention classes.”](#) on page 155. [Chapter 9, “Working with retention classes.”](#) on page 241. For information on objects flagged as replication collisions, see [“Object content collisions”](#) on page 20.



Note: HCP system-level administrators can enable or disable disposition for the repository as a whole. While disposition is disabled for the repository, enabling it for the namespace has no effect.

Data access permission masks

A **data access permission mask** determines which of these operations is allowed in a namespace:

- **Read** — Lets users:
 - Read and retrieve objects, including object metadata (system metadata and custom metadata)
 - List directory contents
- **Write** — Lets users:
 - Add objects to the namespace
 - Modify system metadata
 - Add or replace custom metadata
- **Delete** — Lets users delete objects and custom metadata from the namespace.

- **Privileged** — Lets users delete objects that are under retention. For users to perform this operation, delete operations must also be allowed.
- **Search** — Lets users use the HCP metadata query API and the HCP Search Console to query or search the namespace. For users to query or search a namespace, read operations must also be allowed.

Data access permission masks are set at the system, tenant, and namespace levels:

- The system-level mask applies across all namespaces (that is, systemwide).
- The tenant-level mask is set individually for each tenant. This mask applies only to the namespaces owned by that tenant.
- The namespace-level mask is set individually for each namespace and applies only to that namespace.

The effective permissions for a tenant are the operations allowed by both the system-level and tenant-level permission masks. That is, to be in effect for a tenant, a permission must be included in the system-level permission mask *and* in the tenant-level permission mask.

The effective permissions for a namespace are the operations that are allowed by the masks at all three levels. That is, to be in effect for a namespace, a permission must be included in the system-level permission mask, the tenant-level permission mask, *and* the namespace-level permission mask.

The table below shows an example of the effective permissions for a namespace given a set of data access permission masks.

| Permission mask | Permissions | | | | |
|----------------------------------|-------------|-------|--------|--------------|--------|
| | Read | Write | Delete | Priv. delete | Search |
| Systemwide permission mask | ✓ | ✓ | ✓ | | ✓ |
| Tenant permission mask | ✓ | ✓ | ✓ | ✓ | |
| Namespace permission mask | ✓ | ✓ | ✓ | ✓ | ✓ |
| Effective permission mask | ✓ | ✓ | ✓ | | |

Both the default tenant and the default namespace initially have all permission in their data access permission masks.

HCP system administrators can change the systemwide permission mask at any time. Tenant administrators can change the tenant and namespace permission masks at any time.

You can make a namespace effectively read-only by removing all operations except read from its data access permission mask.

For information on setting the default tenant and namespace permission masks, see [“Changing the tenant permission mask”](#) on page 48 and [“Changing the namespace permission mask”](#) on page 73.

Replication

Replication is the process of keeping selected HCP tenants and namespaces and default-namespace directories in two or more HCP systems in sync with each other. Basically, this entails copying object creations, deletions, and metadata changes between systems. HCP also replicates retention classes, content classes, and all compliance log messages.

A **replication topology** is a configuration of HCP systems that are related to each other through replication. Typically, the systems in a replication topology are in separate geographic locations and are connected by a high-speed wide area network.

Clients can read from namespaces on all systems to which those namespaces are replicated. The replication configuration set at the system level determines on which systems clients can write to namespaces.



Note: Not all HCP systems support replication.

Replication benefits

Replication has several purposes:

- If one system in a replication topology becomes unavailable (for example, due to network issues), another system in the topology can provide continued data availability.
- If one system in a replication topology suffers irreparable damage, another system in the topology can serve as a source for disaster recovery.
- If multiple HCP systems are widely separated geographically, each system may be able to provide faster data access for some applications than the other systems can, depending on where the applications are running.
- If an enterprise has several satellite offices, an HCP system at a central facility can consolidate data from the HCP systems at those outlying locations.
- If an object cannot be read from one system in a replication topology (for example, because a server is unavailable), HCP can try to read it from another system in the topology. HCP tries to do this only if:
 - The directory that contains the object is being replicated.
 - The default namespace has the read-from-remote-system feature enabled.
 - The object has already been replicated. Users can check object metadata to determine whether an object has been replicated.

For more information on this feature, see [“Changing replication options”](#) on page 78.

- If a system that participates in a replication topology is unavailable, HTTP requests to that system can be automatically serviced by another system in the topology without the client needing to modify the target URL. The other system can service a request only if:
 - The directory named in the URL is replicated to the other system
 - The default namespace is configured to accept requests redirected from other HCP systems
 - The HCP systems involved use DNS for system addressing

For more information on this feature, see [“Changing replication options”](#) on page 78.

Replication implementation

Replication is configured at the system level. For the default namespace, the HCP system administrator selects the top-level directories to be replicated. HCP then replicates objects only from those directories and all subdirectories of those directories, recursively. Retention classes and compliance log messages are replicated only if at least one directory is selected for replication.

Depending on the replication topology, users may not be able to make any changes to namespace content on one or more systems in the topology.

Replication is asynchronous with other HCP activity. You can monitor its progress in the Tenant Management Console. For information on this, see [“Monitoring replication”](#) on page 63.

Replication collision handling

If clients can write to multiple systems in a replication topology, collisions can occur when different changes are made to the same objects on different systems and those changes are then replicated. Similarly, if you make changes to content classes defined for the default tenant or retention classes defined for the default namespace on multiple systems in a replication topology, configuration collisions can occur.

The way HCP handles collisions that occur due to replication depends on the type of collision. However, the general rule is that more recent changes have priority over conflicting less recent changes.

Object content collisions

An object content collision occurs when these events occur in the order shown:

1. An object is created with the same name in the same directory on two systems in a replication topology, but the object has different content on the two systems.
2. The object on one of the systems is replicated to the other system.

When an object content collision occurs, the more recently created object keeps its name and location. The other object is either moved to the `.lost+found` directory or renamed, depending on the namespace configuration.

When HCP moves an object to the `.lost+found` directory, the full object path becomes `.lost+found/replication/system-generated-directory/old-object-path`.

When renaming an object due to a content collision, HCP changes the object name to `object-name.collision`. If the new name is already in use, HCP changes the object name to `object-name.1.collision`. If that name is already in use, HCP successively increments the middle integer by one until a unique name is formed.

Objects that have been relocated or renamed due to content collisions are flagged as replication collisions in their system metadata. Clients can use the metadata query API to search for objects that are flagged as replication collisions.

If an object that's flagged as a replication collision changes (for example, if its retention period is extended), its collision flag is removed. If a client creates a copy of a flagged object with a new name, the collision flag is not set on the copy.

You can configure the default namespace to have the disposition service automatically delete objects that are flagged as replication collisions. When selecting this option, you specify the number of days the disposition service should wait before deleting such an object. The days are counted from the time the collision flag is set. If the collision flag is removed from an object, the object is no longer eligible for deletion by the disposition service.

For information on configuring the method HCP should use to handle object name collisions in a namespace, see [“Changing replication options”](#) on page 78. For information the metadata query API, see *HCP Metadata Query API Reference*.

System metadata collisions

A system metadata collision occurs when these events occur in the order shown:

1. Different changes are made to the system metadata for a given object on each of two systems in a replication topology.
2. The changed system metadata on one of the systems is replicated to the other system.

For example, suppose a user on one system changes the shred setting for an object while a user on the other system changes the index setting for the same object. When the object on either system is replicated to the other system, a system metadata collision occurs.

If a collision occurs when changed system metadata for a given object is replicated from one system (system A) in a replication topology to another system (system B) in the topology:

- For changed system metadata other than the retention setting and hold status:
 - If the last change made on system A is more recent than the last change made on system B, HCP changes the system metadata on system B to match the system metadata on system A.
 - If the last change on system B is more recent than the last change on system A, HCP does not change the system metadata on system B.
- For a changed retention setting:
 - If the retention setting on system A specifies a longer retention period than does the retention setting on system B, HCP changes the retention setting on system B to match the retention setting on system A.
 - If the retention setting on system B specifies a longer retention period than does the retention setting on system A, HCP does not change the retention setting on system B.

- For a changed hold status:
 - If the object is on hold on system A but not on system B, HCP places the object on hold on system B.
 - If the object is on hold on system B but not on system A, HCP leaves the object on hold on system B.

Here are some examples of how HCP handles collisions when changed system metadata for a given object is replicated from one system (system A) in a replication topology to another system (system B) in the topology.

Example 1

The object starts out on both system A and system B with these system metadata settings:

Shred: false
Index: false

The table below shows a sequence of events in which the system metadata for the object is changed and the changes are then replicated.

| Sequence | Event |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | On system A, a client changes the shred setting to true. |
| 2 | On system B, a client changes the index setting to true. |
| 3 | The changes on system A are replicated to system B. The resulting settings for the object on system B are: Shred: false Index: true |

Example 2

The object starts out on both system A and system B with these system metadata settings:

Retention: Initial Unspecified
Shred: false
Index: false

The table below shows a sequence of events in which the system metadata for the object is changed and the changes are then replicated.

| Sequence | Event |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | On system A, a client changes the retention setting to Deletion Prohibited. |
| 2 | On system B, a client changes the retention setting to Deletion Allowed. |
| 3 | On system B, a client changes the index setting to true. |
| 4 | On system A, a client changes the shred setting to true. |
| 5 | The changes on system A are replicated to system B. The resulting settings for the object on system B are: Retention: Deletion Prohibited Shred: true Index: false |

Example 3

The object starts out on both system A and system B with these system metadata settings:

Retention: Initial Unspecified
Hold: true
Shred: false
Index: false

The table below shows a sequence of events in which the system metadata for the object is changed and the changes are then replicated.

| Sequence | Event |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | On system A, a client changes the retention setting to Deletion Allowed. |
| 2 | On system B, a client changes the retention setting to Deletion Prohibited. |
| 3 | On system B, a client changes the index setting to true. |
| 4 | On system A, a client changes the shred setting to true. |
| 5 | On system A, a client releases the object from hold. |
| 6 | The changes on system A are replicated to system B. The resulting settings for the object on system B are: Retention: Deletion Prohibited Hold: true Shred: true Index: false |

(Continued)

| Sequence | Event |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 7 | <p>The changes on system B are replicated to system A. The resulting settings for the object on system A are:</p> <p>Retention: Deletion Prohibited Hold: true Shred: true Index: false</p> |

Custom metadata collisions

A custom metadata collision occurs when these events occur in the order shown:

1. One of these changes occurs:
 - Custom metadata is added to a given object on each of two systems in a replication topology, but the added custom metadata is different on the two systems.

The addition of custom metadata to an object on only one of the systems does not result in a custom metadata collision. Instead, the new custom metadata is replicated from that system to the other system without conflict.

- The custom metadata for a given object is replaced on each of two systems in a replication topology, but the replacement custom metadata is different on the two systems.
 - The custom metadata for a given object is replaced on one system in a replication topology, and the same custom metadata is deleted on another system in the topology.
2. The change made on one of the systems is replicated to the other system.

If a collision occurs when a custom metadata change for a given object is replicated from one system (system A) in a replication topology to another system (system B) in the topology:

- If the last change on system A is more recent than the last change on system B, HCP applies the change from system A to the custom metadata on system B

- If the last change on system B is more recent than the last change on system A, HCP does not change the custom metadata on system B

For example, suppose a given object starts out with the same custom metadata on system A and system B. The table below shows a sequence of events in which the custom metadata for the object is changed and the change is then replicated.

| Sequence | Event |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | On system B, a client replaces the custom metadata for the object with new custom metadata. |
| 2 | On system A, a client replaces the custom metadata for the object with different custom metadata from the custom metadata used on system B. |
| 3 | The change on system A is replicated to system B. The resulting custom metadata for the object on system B is the new custom metadata from system A. |

Retention class collisions

A retention class collision occurs when these events occur in the order shown:

1. Different changes are made to the same retention class on each of two systems in a replication topology.
2. The changed retention class on one of the systems is replicated to the other system.

If a collision occurs when a change to a retention class is replicated from one system (system A) in a replication topology to another system (system B) involved in the topology:

- **If the last change to the retention class on system A is more recent than the last change to the class on system B and:**
 - The value of the class on system A is greater than the value of the class on system B, HCP changes the value of the class on system B to the value of the class on system A

- The value of the class on system A is less than the value of the class on system B and:
 - System B is in enterprise mode, HCP changes the value of the class on system B to the value of the class on system A



Note: An exception to this rule is when the value of the class on system A is -2 (Initial Unspecified) and the value of the class on system B is *not* 0 (Deletion Allowed). In this case, the value of the class on system B does not change.

- System B is in compliance mode, HCP does not change the value of the class on system B
- **If the last change to the retention class on system B is more recent than the last change to the class on system A**, HCP does not change the value of the class on system B



Note: A retention class value of -1 (Deletion Prohibited) is greater than a value that's a specific duration. A retention class value of 0 (Deletion Allowed) or -2 (Initial Unspecified) is less than a value that's a specific duration.

Service plans

A **service plan** is a named option that can be associated with a namespace. This option determines how HCP manages the objects in that namespace. Service plan names are system specific.

When creating the default tenant, the HCP system administrator selects a service plan for the default namespace. You can then choose a different service plan for that namespace at any time.

The service plan you select for the default namespace should match the expected usage pattern and properties for the namespace. For example, if the purpose of the namespace is to store objects that most likely will not be accessed again, you would choose a service plan with a description indicating that the plan is intended for archiving.

General administrative information

As an administrator of the default tenant for an HCP system, you are responsible for managing that tenant and the namespace it owns. Your primary job is to ensure that users and applications have the access they need to the default namespace.

The tool you use for this purpose is a web application called the **Tenant Management Console**. Depending on the permissions you have, you can use the Console to configure and monitor the default tenant and namespace, as well as perform compliance activities.

This chapter:

- Presents basic information about using the Tenant Management Console
- Describes the responsibilities of default-tenant administrators

Tenant Management Console

The Tenant Management Console is a tenant-specific web application that lets you manage tenants and namespaces. The Console shows you tenant and namespace status in real time, so you can effectively monitor activity and take action as needed.

Using the Console, you can modify tenant and namespace settings and perform compliance activities. Changes you make through the Console take effect immediately.

Access to the Tenant Management Console is available only through HTTP with SSL security (HTTPS).



Note: As an alternative to the Tenant Management Console, you can use the HCP management API to configure the default tenant and namespace and to create, view, modify, and delete retention classes. For information on this API, see *HCP Management API Reference*.

Console access

To use the Tenant Management Console, you need either:

- A user account defined in HCP (either locally authenticated or RADIUS authenticated).
- A Windows Active Directory® (AD) user account for a user that belongs to one or more AD groups for which corresponding group accounts are defined in HCP. In this book, such an Active Directory user account is referred to as a **recognized AD user account**.

The HCP user account or group accounts specify what you have permission to do in the Console. The menu options, pages, and panels you see in the Console depend on your permissions.

User accounts, group accounts, and roles

Your permissions are determined by the roles associated with your HCP user account or HCP group accounts. Each role represents a set of permissions. Roles generally correspond to job functions.

If an AD user belongs to multiple AD groups for which HCP group accounts exist, that user has all the roles associated with all those group accounts.

User and group accounts are defined at the HCP system level. The same accounts that let you log into the HCP System Management Console are also used for the Tenant Management Console for the default tenant. However, to log into the Tenant Management Console, you need the monitor, administrator, security, or compliance role.

When you log into the System Management Console, you become an **HCP system administrator**. When you log into the Tenant Management Console, you become a **tenant administrator**.

For more information on user accounts and group accounts and to find out which roles you have, contact your HCP system administrator.

Permissions granted by roles

The table below lists the permissions that apply to the Tenant Management Console. Checkmarks indicate the permissions granted by each role.

| Permission | Role | | | |
|---------------------------------------------------------------------------------|---------|---------------|----------|------------|
| | Monitor | Administrator | Security | Compliance |
| View the tenant overview | ✓ | ✓ | ✓ | ✓ |
| Modify the tenant contact information, permission mask, and description | | ✓ | | |
| View tenant log messages about all events except compliance and security events | ✓ | ✓ | ✓ | ✓ |
| View tenant log messages about compliance events | | | | ✓ |
| View tenant log messages about security events | | | ✓ | |
| View syslog and SNMP logging options | ✓ | ✓ | | |
| Enable or disable syslog and SNMP logging | | ✓ | | |
| View email notification settings | ✓ | ✓ | | |
| Modify email notification settings | | ✓ | | |
| View the namespace overview | ✓ | ✓ | | ✓ |
| View the namespace permission mask and description | ✓ | ✓ | | ✓ |
| Modify the namespace permission mask and description | | ✓ | | |
| View namespace retention-related settings | ✓ | ✓ | | ✓ |
| Modify namespace retention-related settings | | | | ✓ |

(Continued)

| Permission | Role | | | |
|-------------------------------------------------------------------------|---------|---------------|----------|------------|
| | Monitor | Administrator | Security | Compliance |
| View the custom metadata XML checking setting for the namespace | ✓ | ✓ | | |
| Modify the custom metadata XML checking setting for the namespace | | ✓ | | |
| View namespace compatibility settings | ✓ | ✓ | | |
| Modify namespace compatibility settings | | ✓ | | |
| View the namespace disposition setting | ✓ | ✓ | | ✓ |
| Modify the namespace disposition setting | | | | ✓ |
| View namespace replication-related settings | ✓ | ✓ | | |
| Modify namespace replication-related settings | | ✓ | | |
| View the service plan associated with the namespace | ✓ | ✓ | | |
| Associate a service plan with the namespace | | ✓ | | |
| View the namespace DPL setting | ✓ | ✓ | | |
| Modify the namespace DPL setting | | ✓ | | |
| View the namespace retention mode | ✓ | ✓ | | |
| Modify the namespace retention mode | | ✓ | | |
| View namespace search settings | ✓ | ✓ | | |
| Enable or disable search for the namespace | | ✓ | | |
| Manage namespace indexing | | ✓ | | |
| View namespace access protocol configurations | ✓ | ✓ | | |
| Configure namespace access protocols for namespaces | | ✓ | | |
| Monitor replication | ✓ | ✓ | | |
| View all namespace log messages except messages about compliance events | ✓ | ✓ | ✓ | ✓ |
| View namespace log messages about compliance events | | | | ✓ |
| View the list of irreparable objects | ✓ | ✓ | | |
| Acknowledge irreparable objects | | ✓ | | |

(Continued)

| Permission | Role | | | |
|---------------------------------------------------------------------------------|---------|---------------|----------|------------|
| | Monitor | Administrator | Security | Compliance |
| Create, modify, and delete retention classes | | | | ✓ |
| View the list of retention classes | ✓ | ✓ | | ✓ |
| View individual retention classes | ✓ | ✓ | | ✓ |
| Perform privileged delete operations | | | | ✓ |
| Download HCP Data Migrator | ✓ | ✓ | ✓ | ✓ |
| Change your own locally authenticated password in the Tenant Management Console | ✓ | ✓ | ✓ | ✓ |
| View HCP documentation from the Tenant Management Console | ✓ | ✓ | ✓ | ✓ |

Console sessions

A Tenant Management Console session begins when you do one of these:

- Log into the Console using an HCP user account or recognized AD user account.
- Access a Console page while logged into Windows with a recognized AD user account. This is called **single sign-on**. With single sign-on, you don't need to explicitly log into the Console.

For single sign-on to work, your web browser must be configured to support it. For more information on this, see [Appendix C, "Browser configuration for single sign-on with Active Directory,"](#) on page 177.

A session ends when you log out. During a session, you can perform any actions for which you have permission.

During a session, if you don't take any action for a certain amount of time, the Console displays the **Idle Timeout** page. If you explicitly logged into the session, the Console automatically logs you out and, when you click on any tab on the **Idle Timeout** page, displays the login page. If you started the session by using single sign-on, when you click on any tab, the Console displays the requested page. The exact amount of idle time allowed is configured at the HCP system level.

You can access the Tenant Management Console directly from the HCP System Management Console. Doing so does not start a Tenant Management Console session. Rather, it continues the current System Management Console session.

Tenant Management Console URL

The URL for the Tenant Management Console has this format:

```
https://default.hcp-domain-name:8000
```

For example, to access the Tenant Management Console for the default tenant in the HCP system with the domain name hcp-ma.example.com, you would use this URL:

```
https://default.hcp-ma.example.com:8000
```

Using a hosts file

Typically, the HCP system is included as a subdomain in the corporate DNS. If this is not the case, you need to provide mappings of the tenant hostname to one or more IP addresses for the HCP system.

You specify hostname mappings in the `hosts` file on the client. The location of this file depends on the client operating system:

- On Windows, by default: `c:\windows\system32\drivers\etc\hosts`
- On Unix: `/etc/hosts`
- On Mac OS[®] X: `/private/etc/host`

Each entry in a `hosts` file maps one or more fully qualified hostnames to a single IP address. So, for example, if one of the IP addresses for the HCP system is 192.168.210.16, you would add this line to the `hosts` file on the client to enable access to the Tenant Management Console for the default tenant:

```
192.168.210.16    default.hcp-ma.example.com
```

The following considerations apply to `hosts` file entries:

- Each entry must appear on a separate line.
- Multiple hostnames in a single line must be separated by white space.
- Each hostname can map to multiple IP addresses.

You can include comments in a `hosts` file either on separate lines or following a mapping on the same line. Each comment must start with a number sign (`#`). Blank lines are ignored.

For the IP addresses for the HCP system, contact your HCP system administrator.

Hostname mapping considerations

An HCP system has multiple IP addresses. You can map the tenant hostname to more than one of these IP addresses in the `hosts` file. The way multiple mappings are used depends on the client platform. For information on how your client handles multiple mappings in a `hosts` file, see your client documentation.

If any of the IP addresses listed in the `hosts` file are unavailable, timeouts may occur when you use a `hosts` file to access the Tenant Management Console.

Logging in

Depending on the HCP system configuration, you can log into the Tenant Management Console with an HCP user account or a recognized AD user account. To log into the Console:

1. Open a web browser.
2. In the address field, enter the URL for your Tenant Management Console.



Note: If you inadvertently use *http* instead of *https* in the URL, the browser returns an error. Enter the URL again, this time using *https*.

One of these happens:

- If all of these are true, you are automatically logged into the Tenant Management Console, and the tenant **Overview** page appears:
 - You are currently logged into Windows with a recognized AD user account.
 - HCP is configured to support AD.
 - Your web browser is configured to support single sign-on with AD. For information on this, see [Appendix C, “Browser configuration for single sign-on with Active Directory.”](#) on page 177.

This is single sign-on. No further action is required.

- If HCP is configured to support AD but any of the following apply, a message appears indicating that single sign-on was not possible:
 - Your web browser is not configured to support single sign-on.
 - You are not currently logged into Windows with a recognized AD user account.
 - You are not on a Windows computer.

In these cases, you need to click on the **Console login page** link in the message to display the Tenant Management Console login page.

- If HCP is not configured to support AD, the Tenant Management Console login page appears.



Note: The Tenant Management Console login page shows the specific version of the HCP release. Once you enter the Console, the version number appears at the bottom of each page.

3. In the **Username** field, type your username.
4. In the **Password** field, type your case-sensitive password.

When using an HCP user account, if you try to log in with an invalid password multiple times in a row, you are locked out of the Console. The exact number of times is configured at the HCP system level.



Note: AD can also be configured to disable user accounts after a given number of authentication attempts with an invalid password.



Important: If you're using a locally authenticated HCP user account, you should change your password as soon as possible the first time you log into the Tenant Management Console.

5. If HCP is configured to support AD, do either of these in the **Domain** field:
 - If you're using an HCP user account, select the fully qualified name of the HCP system.
 - If you're using a recognized AD user account, select the AD domain in which your user account is defined.

If HCP is not configured to support AD, the login page does not display the **Domain** field.

6. Click on the **Log In** button.

The Console displays the tenant **Overview** page or, if you're using an HCP user account and are required to change your password, the **Change Password** page.

For information on the tenant **Overview** page, see ["About the tenant Overview page"](#) on page 44. For information on changing your password, see ["Changing your password"](#) on page 36.

Using the Tenant Management Console

Tenant Management Console pages display information about the default tenant and namespace. Some pages also let you configure various aspects of the tenant and namespace.

Console pages have menus and hyperlinks for navigation. Each page shows a horizontal menu at the top. Some of the menu options display a secondary menu when you mouse over them. To navigate to a page, you click on the corresponding menu option.

You can also use shortcut keys to navigate to pages in the Tenant Management Console. Each link that has a shortcut key has the applicable letter underlined. To use the shortcut key, follow the convention for the browser you're using.

Each page of the Tenant Management Console shows the username of the currently logged-in user in the upper right corner.



Note: While the HCP system is experiencing a heavy load, the Tenant Management Console may be slower to present certain information.

Refreshing pages

Tenant Management Console pages do not automatically refresh themselves while they remain open. To see the most recent values on a page, click again on the menu option that opens that page.



Note: Using the browser reload button to refresh a page that lets you create or modify an entity causes the Console to resubmit values you previously entered on the page.

Submitting changes

Tenant Management Console pages and panels on which you can modify information have action buttons (such as **Create Retention Class** and **Update Settings**) that submit your changes. Action buttons make the changes on a page permanent. These changes take effect immediately.


You need to submit the changes you make before switching to a different page or panel. If you switch without submitting those changes, the Console does not retain them.

For some checkbox options, selecting or deselecting the checkbox causes that change to take effect immediately.

After you submit changes, the Console displays a message indicating whether HCP successfully made the changes. To hide the message, click on **Dismiss** in the message area.

Viewing HCP documentation

HCP documentation is available online in PDF format. To view a document from the Tenant Management Console:

1. In the top right corner of the Tenant Management Console window, mouse over the documentation icon () to open a dropdown menu of the available documents.
2. In the dropdown menu, click on the document you want.

Changing your password

Depending on how your HCP user account is set up, HCP may authenticate your username and password locally or remotely when you log in. If your account is set up for local authentication, you can change your password in the Tenant Management Console. When you change your password in this Console, it also changes for any other HCP interfaces to which your user account gives you access.

If your account is set up for remote authentication or if you use an AD user account to access the Console, you use a method outside HCP to change your password.

For information on local and remote authentication, contact your HCP system administrator.

To change your locally authenticated password in the Tenant Management Console:

1. Log into the Tenant Management Console using your existing password.
2. In the top right corner of the Console window, click on the **Password** link.
3. On the **Change Password** page:
 - In the **Existing Password** field, type your current password.
 - In the **New Password** field, type your new password. Passwords can be up to 64 characters long, are case sensitive, and can contain any valid UTF-8 characters, including white space.

To be valid, a password must include at least one character from two of these three groups: alphabetic, numeric, and other.

The minimum password length is configured at the HCP system level. Typically, it's six or eight characters.

When changing your password, you cannot reuse your current password.
4. Click on the **Update Password** button.

Logging out

To log out of the Tenant Management Console:

1. In the top right corner of the Console window, click on the **Log Out** link.
2. If you explicitly logged in, close the browser window to ensure that other users cannot go back into the Tenant Management Console using the credentials you used to log in.



Tip: For extra security, clear the browser cache before closing the window.

Administrative responsibilities

Tenant-level administrative responsibilities consist of:

- Maintaining the default tenant
- Maintaining the default namespace
- Monitoring the tenant and namespace
- Managing namespace access
- Ensuring recovery
- Performing compliance activities

You perform these activities in the HCP Tenant Management Console. For information on using this Console, see [“Tenant Management Console”](#) on page 28.

Maintaining the default tenant

An HCP system-level administrator creates the default tenant and maintains certain aspects of its configuration. Other aspects of tenant configuration, however, are maintained at the tenant level.

As a tenant administrator, you are responsible for configuring:

- Tenant contact information
- The tenant description
- The tenant permission mask

For information on these activities, see [“Configuring the tenant”](#) on page 46.

Maintaining the default namespace

When creating the default tenant, an HCP system-level administrator also does the initial configuration of the default namespace. Once that’s done, you, as the tenant administrator, can change most of the namespace properties. You are also responsible for managing search and indexing for the namespace.

For information on these activities, see [“Configuring the namespace”](#) on page 72 and [Chapter 7, “Managing search and indexing,”](#) on page 121.

Monitoring the tenant and namespace

HCP is a self-monitoring, self-healing system that automatically alerts you to any issues you may need to address (such as a namespace running low on space). The Tenant Management Console enables you to review tenant and namespace activity and take action to address certain issues. Using the Console, you can view:

- Statistics and graphs showing namespace usage. For more information on this, see:
 - [“Tenant statistics”](#) on page 44
 - [“Objects section”](#) on page 68
 - [“Usage section”](#) on page 69
- Messages about tenant and namespace events, such as configuration changes and searches performed from the Search Console. For information on this, see:
 - [“Major tenant events”](#) on page 44
 - [“Viewing the complete tenant event log”](#) on page 51
 - [“Major namespace events”](#) on page 70
 - [“Viewing the complete namespace event log”](#) on page 82
- Messages about security events (that is, attempts to log into the Tenant Management Console with an invalid username). For information on this, see [“Viewing the tenant security log”](#) on page 52.
- Messages about compliance events (that is, retention class activity and privileged delete operations). For information on this, see:
 - [“Viewing the tenant compliance log”](#) on page 52
 - [“Viewing the namespace compliance log”](#) on page 82
- Reports of irreparable objects. For information on this, see [“Working with irreparable objects”](#) on page 83.
- Alerts warning you of conditions that may need your attention. For information on this, see:
 - [“Tenant alerts”](#) on page 45

- [“Namespace alerts”](#) on page 71
- The progress of replication activity. For information on this, see [“Monitoring replication”](#) on page 63.

You can also have HCP send tenant and namespace log messages to syslog servers, SNMP managers, and specified email addresses. For information on this see:

- [“Enabling syslog logging”](#) on page 55
- [“Enabling SNMP logging”](#) on page 55
- [“Configuring email notification”](#) on page 56

Managing namespace access

Managing namespace access entails:

- Setting the tenant and namespace data access permission masks (see [“Changing the tenant permission mask”](#) on page 48 and [“Changing the namespace permission mask”](#) on page 73)
- Configuring the protocols through which clients access the default namespace (see [Chapter 6, “Configuring the namespace access protocols,”](#) on page 85)
- Optionally, downloading HCP Data Migrator for installation on client computers (see [Chapter 10, “Downloading HCP Data Migrator,”](#) on page 165)

Ensuring recovery

With a DPL of two or higher or with replication in effect, stored data is well-protected, so you generally don't need to make backup copies. However, if your organization requires backups of the stored data, you are responsible for creating them.

You back up the default namespace through the NDMP protocol. Should the need ever arise, you also restore the data through that protocol. For more information on backing up and restoring the namespace with NDMP, see [“Configuring the NDMP protocol”](#) on page 112.

Performing compliance activities

The Tenant Management Console enables you to perform certain activities required for compliance with some local regulations. Using the Console, you can:

- Create, modify, and delete retention classes (see [Chapter 8, “Working with retention classes.”](#) on page 155)
- Perform privileged delete operations (see [Chapter 9, “Using privileged delete.”](#) on page 161)

Managing the tenant

In the Tenant Management Console, you can view all the available information about the default tenant, change certain of its properties, and monitor its activity.

This chapter:

- Describes the tenant **Overview** page
- Contains instructions for configuring the default tenant
- Explains how to monitor tenant activity, including replication

For an introduction to tenants, see [“Namespaces and tenants”](#) on page 3.

About the tenant Overview page

When you access the Tenant Management Console, the first page you see is the tenant **Overview** page. This page gives you a view of the tenant as a whole. It also shows the HCP system time.



Roles: To view the tenant **Overview** page, you need the monitor, administrator, security, or compliance role.

To return to the tenant **Overview** page from other Console pages, click on **Overview** in the top row of tabs.

Tenant statistics

The tenant **Overview** page shows these statistics:

- **Objects ingested** — The total number of objects currently in the default namespace.
- **Objects indexed** — The total number of indexed objects currently in the default namespace. This item appears only if the namespace is search enabled and a search facility is currently selected for use with the Search Console. In the case of the HDDS search facility, that facility must also be configured to show statistics.

Major tenant events

The **Major Events** section on the tenant **Overview** page lists log messages about major events related to the default tenant and namespace (for example, the tenant permission mask was changed). The list includes all such messages that have occurred since the tenant was created.

The list of messages in the **Major Events** section is a subset of the messages in the event log for the tenant. You can view all the messages in the tenant event log in the **All Events** panel on the **Tenant Events** page. For more information on the tenant **All Events** panel, see [“Viewing the complete tenant event log”](#) on page 51.

For a description of the information provided by each log message, see [“Understanding log messages”](#) on page 53. For information on the messages that can appear in the tenant log and how to respond to them, see [Appendix B, “Tenant log messages.”](#) on page 171.

By default, the messages in the **Major Events** section are listed ten at a time in reverse chronological order. For information on managing the message display, see [“Managing the message list”](#) on page 54.

If the **Overview** page shows alerts instead of log messages, click on the **Major Events** tab to display the log messages. For information on the alerts display, see [“Tenant alerts”](#) below.

Tenant alerts

The **Alerts** section on the tenant **Overview** page shows alerts that indicate tenant-related conditions that may require human intervention (for example, the default namespace contains irreparable objects). This section is visible only if alerts currently exist for the tenant.

Each alert is represented by an icon accompanied by descriptive text. For information on the alerts that can appear in the **Alerts** section, see [Appendix A, “Tenant Management Console alerts,”](#) on page 167.

If the **Overview** page shows log messages instead of alerts, click on the **Alerts** tab to display the alerts. For information on the log message display, see [“Major tenant events”](#) above.



Note: If the tenant is read-only due to metadata unavailability, a notice of the situation appears at the top of every Tenant Management Console page. In this case, you cannot make any configuration changes to the tenant or to the default namespace, nor can users and applications make any changes to namespace content. Additionally, in this situation, statistics that describe namespace content may be inaccurate.

Tenant contact information

You can view contact information for the current tenant from the tenant **Overview** page. To view this information, click on the **Contact Information** link.



Roles: To view the contact information for the tenant, you need the administrator role.

The contact information for a tenant is visible not only in the Tenant Management Console but also in the System Management Console. HCP system administrators can use this information to notify you about systemwide events (such as a system shutdown for maintenance). Therefore, you should keep this information up to date.

For information on updating the tenant contact information, see [“Changing the tenant contact information”](#) on page 47.

Tenant permission mask

The **Permissions** section on the tenant **Overview** page shows:

- The permissions included in the inherited permission mask, which is the systemwide permission mask. These permissions are indicated by gray dots.
- The permissions included in the tenant permission mask. These permissions are indicated by orange dots.
- The permissions included in the effective permission mask for the tenant. These permissions are indicated by checkmarks.

For an introduction to permission masks, see [“Data access permission masks”](#) on page 15. For information on modifying the tenant permission mask, see [“Changing the tenant permission mask”](#) on page 48.

Tenant description

The **Description** section on the tenant **Overview** page shows the description of the tenant, if a description exists. This description is visible only in the Tenant Management Console. It is not visible in the System Management Console.

The tenant description appears below the tenant name on the login page for the Tenant Management Console.

For information on providing a description of the tenant, see [“Changing the tenant description”](#) on page 49.

Configuring the tenant

In the Tenant Management Console, you can change these properties of the default tenant:

- The tenant contact information
- The tenant permission mask
- The tenant description

You can also specify whether HCP should send tenant log messages to syslog servers, SNMP managers, and/or specified email addresses if these features are enabled at the HCP system level. For information on these options, see:

- [“Enabling syslog logging”](#) on page 55
- [“Enabling SNMP logging”](#) on page 55
- [“Configuring email notification”](#) on page 56

Changing the tenant contact information

When creating the default tenant, the HCP system administrator has the option of entering contact information for it. Once the tenant exists, both system-level administrators and tenant-level administrators can modify this information.



Roles: To view or modify the tenant contact information, you need the administrator role.

To provide contact information for the default tenant or to update the existing contact information:

1. In the top-level menu in the Tenant Management Console, click on **Overview**.
2. On the tenant **Overview** page, click on the **Contact Information** link.
3. In the **Contact Information** window, fill in the contact information. The table below describes the values you can specify. Except as indicated, all fields are optional.

| Field | Description |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| First Name | First name of the tenant contact. First names can be up to 64 characters long and can contain any valid UTF-8 characters, including white space. |
| Last Name | The last name of the tenant contact. Last names can be up to 64 characters long and can contain any valid UTF-8 characters, including white space. |
| Email | An email address for the tenant contact. Email addresses cannot be more than 254 characters long. |
| Confirm Email | A repeat of the email address for the tenant contact. This field is required if you specify an email address in the Email field. |

(Continued)

| Field | Description |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Phone | A telephone number for the tenant contact. Do not include a telephone number extension. Telephone numbers can contain only numbers, parentheses, hyphens (-), periods (.), plus signs (+), and spaces and can be up to 24 characters long (for example, (800) 123-4567). |
| Extension | A telephone number extension for the tenant contact. Telephone number extensions can contain only numbers and can be up to five characters long. |
| Address Line 1 | The first line of an address for the tenant contact. Address lines can be up to 100 characters long and can contain any valid UTF-8 characters, including white space. |
| Address Line 2 | The second line of an address for the tenant contact. |
| City | The city for the tenant contact. City names can be up to 64 characters long and can contain any valid UTF-8 characters, including white space. |
| State/ Province | The state or province for the tenant contact. State and province names can be up to 64 characters long and can contain any valid UTF-8 characters, including white space. |
| Postal Code | The postal code for the tenant contact. Postal codes can be up to 64 characters long and can contain only alphanumeric characters and hyphens (-). |
| Country | The country for the tenant contact. Country names can be up to 64 characters long and can contain any valid UTF-8 characters, including white space. |

4. Click on the **Update Contact Info** button.

For more information on tenant contact information, see [“Tenant contact information”](#) on page 45.

Changing the tenant permission mask

When the default tenant is created, its data access permission mask includes all permissions. You can change this permission mask at any time.



Roles: To change the tenant permission mask, you need the administrator role.

To change the tenant permission mask:

1. In the top-level menu in the Tenant Management Console, click on **Overview**.
2. On the tenant **Overview** page, click on the **edit** link for the **Permissions** section.

The Console displays a set of checkboxes for the permissions. The permissions that are currently in the tenant permission mask are selected.

3. Select or deselect permissions as needed to modify the permission mask.

Selecting **Purge** automatically selects **Delete**. Selecting **Search** automatically selects **Read**.

4. Click on the **Submit** button.

For an introduction to permission masks, see [“Data access permission masks”](#) on page 15. For more information on the tenant permission mask, see [“Tenant permission mask”](#) on page 46.

Changing the tenant description

The tenant description is optional. You can enter a description or modify the existing description at any time.



Roles: To change the tenant description, you need the administrator role.

To change the tenant description:

1. In the top-level menu in the Tenant Management Console, click on **Overview**.
2. On the tenant **Overview** page, click on the **edit** link for the **Description** section.
3. In the edit area for the description, type the new description of the tenant. The description can be up to 1,024 characters long and can contain any valid UTF-8 characters, including white space.
4. Click on the **Submit** button.

For more information on the tenant description, see [“Tenant description”](#) on page 46.

Monitoring the tenant

While the tenant **Overview** page in the Tenant Management Console gives you a view of the tenant as a whole, the tenant log lets you monitor tenant and namespace activity on a more detailed level. The log records tenant and namespace events such as:

- Tenant Management Console logins
- Tenant and namespace configuration changes
- Creations, modifications, and deletions of retention classes
- Privileged delete operations

Each recorded entry about an event is called a **message**. The tenant log contains all the messages written to it since the tenant was created.

The Tenant Management Console provides several views of the log, as outlined in the table below.

| View | Shows | More information |
|--------------------------------|-----------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------|
| Tenant-level all events | All log messages recorded for the default tenant and namespace | “Viewing the complete tenant event log” below |
| Tenant-level major events | A subset of the log messages in the tenant-level all-events view | “Major tenant events” on page 44 |
| Tenant-level security events | Only log messages about attempts to log into the Tenant Management Console with an invalid username | “Viewing the tenant security log” on page 52 |
| Tenant-level compliance events | All log messages about default namespace events that require the compliance role | “Viewing the tenant compliance log” on page 52 |
| Namespace-level all events | All log messages for the default namespace | “Viewing the complete namespace event log” on page 82 |
| Namespace-level major events | A subset of the log messages in the namespace-level all-events view | “Major namespace events” on page 70 |

(Continued)

| View | Shows | More information |
|-----------------------------------|----------------------------------------------------------------------------------|-------------------------------------------------------------------|
| Namespace-level compliance events | All log messages about default namespace events that require the compliance role | “Viewing the namespace compliance log” on page 82 |

In addition to these views of the log, HCP gives you the option of sending log messages to syslog servers, SNMP managers, and/or specified email addresses. For more information on this, see:

- [“Enabling syslog logging”](#) on page 55
- [“Enabling SNMP logging”](#) on page 55
- [“Configuring email notification”](#) on page 56

For information on the messages that can appear in the tenant log and how to respond to them, see [Appendix B, “Tenant log messages.”](#) on page 171.

Viewing the complete tenant event log

The **All Events** panel on the **Tenant Events** page lists all messages recorded for the default tenant and namespace. By default, the panel displays ten messages at a time in reverse chronological order.



Roles: To view the tenant **All Events** panel, you need the monitor, administrator, security, or compliance role. However, only users with the compliance role can see messages about events that require the compliance role. Only users with the security role can see messages about attempts to log into the Tenant Management Console with an invalid username.

To display the **All Events** panel:

1. In the top-level menu in the Tenant Management Console, mouse over **Monitoring** to display a secondary menu.
2. In the secondary menu, click on **Tenant Events**.
3. On the left side of the **Tenant Events** page, click on **All Events**.

For a description of the information provided by each log message, see [“Understanding log messages”](#) on page 53. For information on managing the message display, see [“Managing the message list”](#) on page 54.

Viewing the tenant security log

The **Security Events** panel on the **Tenant Events** page lists all messages about attempts to log into the Tenant Management Console with an invalid username. By default, the panel displays ten messages at a time in reverse chronological order.



Roles: To view the tenant **Security Events** panel, you need the security role.

To display the **Security Events** panel:

1. In the top-level menu in the Tenant Management Console, mouse over **Monitoring** to display a secondary menu.
2. In the secondary menu, click on **Tenant Events**.
3. On the left side of the **Tenant Events** page, click on **Security Events**.

For a description of the information provided by each log message, see [“Understanding log messages”](#) on page 53. For information on managing the message display, see [“Managing the message list”](#) on page 54.

Viewing the tenant compliance log

The **Compliance Events** panel on the **Tenant Events** page lists all messages about events that require the compliance role. This includes all retention class activity and privileged delete operations. By default, the panel displays ten messages at a time in reverse chronological order.



Roles: To view the tenant **Compliance Events** panel, you need the compliance role.

To display the **Compliance Events** panel:

1. In the top-level menu in the Tenant Management Console, mouse over **Monitoring** to display a secondary menu.
2. In the secondary menu, click on **Tenant Events**.
3. On the left side of the **Tenant Events** page, click on **Compliance Events**.

For a description of the information provided by each log message, see [“Understanding log messages”](#) on page 53. For information on managing the message display, see [“Managing the message list”](#) on page 54.

For information on retention classes, see [Chapter 8, “Working with retention classes.”](#) on page 155. For information on privileged delete, see [Chapter 9, “Using privileged delete.”](#) on page 161.

Understanding log messages

Each message displayed in a view of the tenant log includes this information about an event:

- The username of the event initiator:
 - For user-initiated events, this is the username currently associated with the user account used by the user who initiated the event. These considerations apply:
 - For an HCP user account, if the account has been deleted, the username is followed by the letter D in parentheses.
 - For an AD user account, if the account has been deleted or if HCP currently cannot contact AD, the username for the message is blank.
 - For system-initiated events, the username is **[internal]**.
 - For events initiated by HCP service or support personnel by means other than the Tenant Management Console or HCP management API, the username is **[service]**.

Additionally, when directories in the default namespace are being replicated, messages for events initiated by a user who accessed the Tenant Management Console directly from the HCP System Management Console have a username of **[remote admin]** in the log messages on systems to which the directories are replicated.

- The severity of the event. Possible values are:
 - **Notice** — The event is normal and requires no special action. Events of this severity are informational only. Examples are:
 - Privileged delete requested
 - Object has been shredded

- **Warning** — The event is out of the ordinary and may require manual intervention. Examples are:
 - Account is disabled
 - Disposition service stopped without finishing
 - **Error** — The event is serious and most likely requires manual intervention. Examples are:
 - HCP found an irreparable object
 - Object did not replicate
- The date and time at which the event occurred, shown in the time zone of the HCP system.
 - A short description of the event.

To view more details about an event, click anywhere in the row containing the event message. To hide the details, click again in the row.



The details displayed for an event are:

- The user ID of the event initiator
- For user initiated events, the IP address from which the event request was sent
- For user-initiated events, the port through which HCP received the event request
- The message ID
- The full text of the event message

Managing the message list

You can take the following actions in any of the views of the tenant log:

- To display details for all the listed events, click on the **expand all** link. To hide all details, click on the **collapse all** link.
- To view a different number of messages per page, select the number you want in the **Items per page** field.

- To page forward or backward, click on the next () or back () control, respectively.

Enabling syslog logging

An HCP system can be configured to send system-level log messages to one or more specified syslog servers. You can choose to also send tenant log messages to those servers. The system-level configuration determines whether compliance and security messages are sent along with the other tenant log messages.

You use the **Syslog** page in the Tenant Management Console to enable or disable sending tenant log messages to the syslog servers. To display this page:

1. In the top-level menu in the Tenant Management Console, mouse over **Monitoring** to display a secondary menu.
2. In the secondary menu, click on **Syslog**.



Roles: To view the **Syslog** page, you need the monitor or administrator role. To enable or disable syslog logging, you need the administrator role.

To enable or disable the logging of tenant log messages to syslog servers:

1. On the **Syslog** page, select (to enable) or deselect (to disable) the **Enable syslog logging** option.

If the HCP system is not configured for syslog logging, selecting this option has no effect.

2. Click on the **Update Settings** button.

Enabling SNMP logging

An HCP system can be configured to send system-level log messages to one or more specified SNMP managers. You can choose to also send tenant log messages to those managers. The system-level configuration determines whether compliance and security messages are sent along with the other tenant log messages.

You use the **SNMP** page in the Tenant Management Console to enable or disable sending tenant log messages to the SNMP managers. To display this page:

1. In the top-level menu in the Tenant Management Console, mouse over **Monitoring** to display a secondary menu.
2. In the secondary menu, click on **SNMP**.



Roles: To view the **SNMP** page, you need the monitor or administrator role. To enable or disable SNMP logging, you need the administrator role.

To enable or disable the logging of tenant log messages to SNMP managers:

1. On the **SNMP** page, select (to enable) or deselect (to disable) the **Enable SNMP logging** option.

If the HCP system is not configured for SNMP logging, selecting this option has no effect.

2. Click on the **Update Settings** button.

Configuring email notification

HCP can be configured at the system level to support the use of email to notify recipients about messages added to the system-level log. If the HCP system supports email notification, you can configure HCP to send email about tenant log messages to recipients that you specify.

You can configure each email recipient to receive notification of only selected messages based on the message importance, severity, and type. Important messages are those that appear in the **Major Events** sections on the tenant **Overview** page and the namespace **Overview** page in the Tenant Management Console. Message severity levels are notice, warning, and error. Message types are general, security, and compliance. In all cases, HCP makes a best effort to send the applicable email in a timely manner.

Recipients are added to the blind carbon copy (bcc) list for each email, so the recipients of an email are not visible to one another. The To list remains empty.

You can configure the content of the email that HCP sends. For example, you could choose to have HCP send the full text, severity, and date and time for each log message. Or, if you're concerned about exposing tenant and namespace information in what is by nature an insecure medium, you could format the email to say only that a log message was recorded.

HCP writes messages to the tenant log about email that the email server fails to accept. The messages about failed email are not sent to email recipients.

You use the **Email** page in the Tenant Management Console to enable and configure email notification. To display the **Email** page:

1. In the top-level menu in the Tenant Management Console, mouse over **Monitoring** to display a secondary menu.
2. In the secondary menu, click on **Email**.



Roles: To view the **Email** page, you need the monitor or administrator role. To configure email notification, you need the administrator role.

Enabling email notification

To have HCP send email about log messages, on the **Email** page:

1. Optionally, test whether email notification is supported. For instructions on this, see ["Testing email notification"](#) below.
2. Select the **Enable email notification** option.
3. Optionally, change the format of the email to be sent. For instructions on this, see ["Constructing the email message template"](#) below.
4. Specify one or more recipients to receive email about log messages. For instructions on this, see ["Specifying email recipients"](#) on page 61.

Testing email notification

HCP email notification works only if the HCP system has been configured to enable support for this feature. At any time, you can test the HCP system to determine whether it has been configured to support email notification.

Testing support for email notification causes HCP to send an email to an address that you specify. This email comes from the email address specified in the **From** field in the **Message Settings** section on the **Email** page. The email subject is "Test email from HCP."

To test whether email notification is supported, on the **Email** page:

1. Click on the **Test** button.
2. In the **Test Email Notification** window, type the email address to which you want HCP to send the test email.
3. Click on the **Send** button.

If support for email notification is not configured at the system level, the Tenant Management Console displays an error message. If the Console displays a success message but the email does not arrive, ensure that you've correctly specified the email address to which you want the email sent. If the email still doesn't arrive, contact your HCP system administrator for help.

Constructing the email message template

The content of the email messages HCP sends is determined by the message template specified in the **Message Settings** section on the **Email** page. You can modify this template at any time. The **Message Preview** section shows a sample email that uses the current template.

The email template has three fields, each of which can be filled in with any combination of plain text and email template variables:

- The **From** field specifies the content of the email From line. This field must have a value. That value must have the form of a valid email address.

Some email servers require that the value in the From line be an email address that is already known to the server.

- The **Subject** field specifies the content of the email Subject line. This field must have a value.

For the email template subject, plain text can include spaces but not line breaks or tabs.

- The **Body** field specifies the body of the email. This field is optional.

For the email template body, plain text can include spaces and line breaks but not tabs. The character sequence consisting of a backslash (\) followed by a lowercase n creates a line break.

For a description of email template variables, see ["Email template variables"](#) below.

HCP comes with a default email template. At any time, you can change the email template back to the default. For instructions on this, see [“Restoring the default template”](#) on page 60.

To modify the template HCP uses for email notification about log messages, on the **Email** page:

1. In the **From**, **Subject**, and **Body** fields in the **Message Settings** section, specify the values that you want to use.
2. Optionally, click on the **Preview** button to preview the sample email with the specified format in the **Message Preview** field.
3. Click on the **Update Settings** button at the bottom of the page.

Email template variables

The values you specify in the **From**, **Subject**, and **Body** fields in the email template can include variables that correspond to the information available for each log message (for example, the severity of the event that triggered the message or the short description of the event). When sending email, HCP replaces the variables in the email message with the applicable information.

To include a variable in the email template, you specify the variable name preceded by the dollar sign (\$). A dollar sign followed by anything other than a variable name is displayed as a dollar sign in the email HCP sends.

The table below lists the variables you can use in the email template.

| Variable | Description |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$action | The action to take in response to the message |
| \$date | The date and time at which the event occurred (for example, Wed Feb 8 2012 3:15:57 PM EST) |
| \$fullText | The full text of the message |
| \$id | The message ID |
| \$location | The fully qualified name of the HCP system on which the event occurred (for example, hcp-ma.example.com) |
| \$origin | For user-initiated events, the IP address from which the event request was sent and the port through which HCP received the event request, separated by a colon (for example, 192.168.152.181:8000) |
| \$reason | The reason why HCP issued the message |
| \$scope | Either Tenant or Namespace |

(Continued)

| Variable | Description |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$severity | The severity of the event that triggered the message |
| \$shortText | A brief description of the event that triggered the message |
| \$type | The type of message (General, Security, or Compliance), preceded by Important and a comma if the message is important (for example, Important, Security) |
| \$user | The user ID and username of the event initiator (for example, 105ff38f-4770-4f98-b5b3-8371ab0af359 lgreen) |

For more information on log messages, see [“Understanding log messages”](#) on page 53 and [Appendix B, “Tenant log messages,”](#) on page 171.

Restoring the default template

The table below shows the format of the default email template.

| Field | Default value |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| From | log@\$location |
| Subject | [\$severity] \$shortText |
| Body | The following event occurred on \$date: \$fullText Reason: \$reason Action: \$action Details: User: \$user Origin: \$origin |

To change the email template back to the default, on the **Email** page:

1. Click on the **Reset** button.
2. Optionally, click on the **Preview** button to preview the sample email with the default format in the **Message Preview** field.
3. Click on the **Update Settings** button at the bottom of the page.

Specifying email recipients

You use the **Recipients** section on the **Email** page to specify the email addresses to which HCP sends email about log messages. HCP sends email as blind carbon copies, so email recipients are not visible to one another.



Each row in the **Recipients** section contains one or more email addresses and indicates which messages are sent to those addresses. The section can have at most 25 rows.


Because each row in the **Recipients** section can contain multiple email addresses, you can specify a total of more than 25 addresses in this section. However, HCP sends each email only to an arbitrary 25 of the addresses that are supposed to receive the email. For example, if 34 email addresses are supposed to receive email about log messages that are important and have a severity level of error and a type of general, HCP sends such email only to 25 of those addresses.

You can add, modify, and delete rows in the **Recipients** section at any time.

Understanding the recipients list

Each row in the **Recipients** section specifies:

- One or more email addresses.
- Whether to send email only about important log messages (**Major**) to the specified email addresses or to send email about all log messages (**All**).
- The severity of the log messages about which to send email:
 - **Notice** tells HCP to send email about log messages with a severity level of notice, warning, or error.
 - **Warning** tells HCP to send email about log messages with a severity level of warning or error.
 - **Error** tells HCP to send email only about log messages with a severity level of error.
- Whether to send email about general log messages (). General log messages are all messages that do not have a type of security or compliance.
- Whether to send email about log messages with a type of security ().

- Whether to send email about log messages with a type of compliance ().

Email recipients receive email only about log messages that have all the selected properties.

Adding, modifying, and deleting rows in the recipients list


To add, modify, and/or delete rows in the recipients list, on the **Email** page:

1. Take one or more of these actions:

– To add a row:

1. Optionally, in the **Recipients** field, type a comma-separated list of one or more well-formed email addresses.
2. Click on **Add**.


A new row appears in the recipients list with importance set to **Major**, severity set to **Error**, and only general selected as the type. The row is highlighted in green.

To remove the new row, click on the delete control () for the row.

– To modify a row:

- Optionally, in the **Address** field, type additional well-formed email addresses and/or modify or delete existing addresses. This field must contain at least one well-formed email address and no incorrectly formed addresses.
- Optionally, change the properties based on which HCP sends email to the specified addresses.

If you deselect all the types, no email is sent to the specified addresses.

– To delete a row, click on the delete control () for the row.

The row turns red. To undo the deletion, click again on the delete control.

2. Click on the **Update Settings** button at the bottom of the page.

Monitoring replication

The **Replication** page in the Tenant Management Console shows information about replication of the default namespace when at least one directory is selected for replication. You can use the statistics and graphs on this page to monitor replication progress.

If no directories are currently selected for replication, the **Replication** page shows a replication status of **Not Replicating**.

To display the **Replication** page:

1. In the top-level menu in the Tenant Management Console, mouse over **Services** to display a secondary menu.
2. In the secondary menu, click on **Replication**.



Roles: To view the **Replication** page, you need the monitor or administrator role.

For an introduction to replication, see ["Replication"](#) on page 17.

High-level view of replication

At the high level, the **Replication** page shows:

- The current rate of replication activity, expressed as operations per second and bytes per second. An operation is the replication or recovery of a single event, such as the creation of a new object or a metadata change.
- A progress bar that measures the up-to-date-as-of time for replication of the default namespace, along with this time as a numeric value. The length of the progress bar represents 30 days, with the right end representing the current time.

The up-to-date-as-of time is the difference between:

- The time before which objects and metadata changes are guaranteed to have been sent to other systems or received from other systems
- The current time

Detailed view of replication

To view more detailed information about the replication of the default namespace, click on the namespace name on the **Replication** page. The panel that opens shows:

- The date and time before which objects and metadata changes for the namespace are guaranteed to have been sent to other systems or received from other systems
- A graph of the history of the up-to-date-as-of time for replication of the namespace
- A graph of the history of the data transmission rate for replication of the namespace
- A graph of the history of the operation rate for replication of the namespace

Up-to-date-as-of time

The **Up to Date as Of** section in the namespace replication panel contains a graph that shows the history of the up-to-date-as-of time for replication of the namespace. The section heading shows the current up-to-date-as-of time. If the graph is not currently visible, click on **Up to date as of** to display it.

The x-axis in this graph marks the passage of time. It shows 30 days (or fewer if replication has been occurring for less than 30 days). The y-axis marks the up-to-date-as-of time in days, hours, or minutes. As the up-to-date-as-of time varies, the measurement unit for the y-axis grows or shrinks as needed (for example, from days to hours to minutes). The lower the up-to-date-as-of time, the closer replication or recovery is to being synchronized with current namespace activity.

Data transmission rate

The **Transfer Rate** section in the namespace replication panel contains a graph that shows the history of the rate of replication data transmissions for the namespace per second. The section heading shows the current data transmission rate. If the graph is not currently visible, click on **Transfer Rate** to display it.

The x-axis in this graph marks the passage of time. It shows 30 days (or fewer if replication or recovery has been occurring for less than 30 days). The y-axis marks the data transmission rate in KB, MB, or GB. As the transmission rate varies, the measurement unit for the y-axis grows or shrinks as needed (for example, from KB to MB to GB).

Operation rate

The **Operations per Second** section in the namespace replication panel contains a graph that shows the history of the rate of replication operations for the namespace per second. The section heading shows the current operation rate. If the graph is not currently visible, click on **Operations per Second** to display it.

The x-axis in this graph marks the passage of time. It shows 30 days (or fewer if replication or recovery has been occurring for less than 30 days). The y-axis marks the operation rate in tens, hundreds, or thousands. As the operation rate varies, the measurement unit on the y-axis grows or shrinks as needed (for example, from tens to hundreds to thousands).

Managing the namespace

In the Tenant Management Console, you can view all the available information about the default namespace, change certain of its properties, and monitor its activity.

This chapter describes these activities.

About the namespace Overview page

The namespace **Overview** page shows the current status of the default namespace. To display this page:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Overview**.



Roles: To view the namespace **Overview** page, you need the monitor, administrator, or compliance role.

Namespace URL

The top of the namespace **Overview** page shows the URL for access to the namespace content. If the HTTP or HTTPS protocol is enabled, this URL is a link for browsing that content.

The namespace URL has this format:

`http://default.default.hcp-domain-name`

For example, the URL for access to the content of the default namespace for the HCP system named hcp-ma.example.com is:

`http://default.default.hcp-ma.example.com`

If HTTPS is enabled, the URL contains *https* instead of *http*.

For information on enabling or disabling HTTP and HTTPS, see [“Configuring the HTTP and WebDAV protocols”](#) on page 92. For information on browsing the namespace content, see *Using the Default Namespace*.

Objects section

The **Objects** section on the namespace **Overview** page contains a graph showing the number of objects in the default namespace during the past 30 days (or since the namespace was created if that was less than 30 days ago).

While any of the search facilities is selected for use with the Search Console, the graph also shows the total number of indexed objects in the namespace during the past 30 days. For any point in time for which the indexed object count is shown, the count reflects the index maintained by the search facility that was selected for the Search Console at that time.



Notes:

- If the HDDS search facility is selected for use with the Search Console, the graph shows the number of indexed objects only if that facility is configured to show statistics.
 - For any period during which HCP cannot retrieve statistics from the HDDS server (for example, because the network connection is broken), the graph shows the number of indexed objects as zero.
-

The x-axis in the **Objects** graph marks the passage of time. The y-axis marks the number of objects. As the number of objects increases, the intervals on the y-axis get larger. The section heading indicates the current measurement unit (for example, thousands or millions).

The graph legend shows the most recent value for the number of objects stored and, if applicable, the number of indexed objects.

Below the **Usage** section (see below), the **Overview** page shows the date and time the **Objects** and **Usage** sections were last updated. To show the most current information in these sections, click on the **Refresh Now** link.

Usage section

The **Usage** section on the namespace **Overview** page contains a graph showing information about the namespace storage during the past 30 days (or since the namespace was created if that was less than 30 days ago).

The x-axis in the **Usage** graph marks the passage of time. The y-axis marks the amount of storage in gigabytes, terabytes, or petabytes, depending on the namespace size. The graph heading indicates the current measurement unit (gigabytes (GB), terabytes (TB), or petabytes (PB)).

The **Usage** graph shows:

- **Total system storage** — The total amount of storage space available for all data stored for all namespaces, including object data, metadata, the redundant data required to satisfy namespace DPL settings, and the metadata query engine index. For information on DPL, see [“Data protection level”](#) on page 10.
- **Used storage capacity** — The total amount of storage space currently occupied by all data stored in the default namespace, including object data, metadata, any redundant data required to satisfy namespace DPL settings.
- **Ingested volume** — The total size of the stored data and custom metadata before it was added to the default namespace. This value tells you how much data you have stored.

The graph legend shows the current value for each item.

Below the **Usage** section, the **Overview** page shows the date and time the **Objects** and **Usage** sections were last updated. To show the most current information in these sections, click on the **Refresh Now** link.

Major namespace events

The **Major Events** section on the namespace **Overview** page lists log messages about major events related to the namespace (for example, the namespace retention mode was changed). The list includes all such messages that have occurred since the namespace was created.

The list of messages in the **Major Events** section is a subset of the messages in the event log for the namespace. You can view all the messages in the namespace event log in the **All Events** panel for the namespace. For more information on the namespace **All Events** panel, see [“Viewing the complete namespace event log”](#) on page 82.

For a description of the information provided by each log message, see [“Understanding log messages”](#) on page 53. For information on the messages that can appear in the namespace log and how to respond to them, see [Appendix B, “Tenant log messages,”](#) on page 171.

By default, the messages in the **Major Events** section are listed ten at a time in reverse chronological order. For information on managing the message display, see [“Managing the message list”](#) on page 54.

If the **Overview** page shows alerts instead of log messages, click on the **Major Events** tab to display the log messages. For information on the alerts display, see [“Namespace alerts”](#) below.

Namespace alerts

The **Alerts** section on the namespace **Overview** page shows alerts that indicate namespace-related conditions that may require human intervention (for example, the existence of irreparable objects). This section is visible only if alerts currently exist for the namespace.

Each alert is represented by an icon accompanied by descriptive text. For information on the alerts that can appear in the **Alerts** section, see [Appendix A, “Tenant Management Console alerts,”](#) on page 167.

If the **Overview** page shows log messages instead of alerts, click on the **Alerts** tab to display the alerts. For information on the log message display, see [“Major namespace events”](#) above.

Namespace services, service plan, retention mode, hash algorithm, and DPL

The namespace **Overview** page shows which of these features are currently enabled for the namespace:

- **Replication** — If this feature is present, namespace content can be replicated. If this feature is not present, namespace content cannot be replicated.

For information on this feature, see [“Replication”](#) on page 17.

- **Search** — If this feature is present, search is currently enabled for the namespace. The namespace is searchable if its effective permission mask includes search. If this feature is not present, the namespace is not searchable.

For information on this feature, see [Chapter 7, “Managing search and indexing,”](#) on page 121. For information on the effective permission mask for a namespace, see [“Data access permission masks”](#) on page 15.

Additionally, the namespace **Overview** page shows:

- The **service plan** in effect for the namespace. For information on service plans, see [“Service plans”](#) on page 26.

- The **retention mode** of the namespace. This is either enterprise or compliance. For information on enterprise and compliance modes, see [“Retention mode”](#) on page 11.
- The **cryptographic hash algorithm** for the namespace. For more information on this, see [“Cryptographic hash algorithm”](#) on page 10.
- The **DPL** for the namespace. For more information on this, see [“Data protection level”](#) on page 10.

Namespace permission mask

The **Permissions** section on the namespace **Overview** page shows:

- The permissions included in the inherited mask, which is the effective tenant permission mask. These permissions are indicated by gray dots.
- The permissions included in the namespace permission mask. These permissions are indicated by orange dots.
- The permissions included in effective permission mask for the namespace. These permissions are indicated by checkmarks.

For an introduction to permission masks, see [“Data access permission masks”](#) on page 15. For information on modifying the namespace permission mask, see [“Changing the namespace permission mask”](#) on page 73.

Namespace description

The **Description** section in the namespace **Overview** page shows the description of the namespace, if a description exists.

For information on providing or modifying the description of the namespace, see [“Changing the namespace description”](#) on page 74.

Configuring the namespace

You can change these properties of the default namespace:

- Permission mask
- Description

- Retention-related settings
- Whether custom metadata XML checking is enabled
- Compatibility settings
- Whether disposition is enabled
- Replication options
- Associated service plan
- DPL
- Retention mode (only from enterprise to compliance)

You can also:

- Enable and configure the namespace access protocols. For information on this, see [Chapter 6, “Configuring the namespace access protocols.”](#) on page 85.
- Change search and indexing options for the namespace. For information on this, see [Chapter 7, “Managing search and indexing.”](#) on page 121.

Changing the namespace permission mask

When the default namespace is created, its data access permission mask includes all permissions. You can change the namespace permission mask at any time.



Roles: To change the namespace permission mask, you need the administrator role.

To change the permission mask for the namespace:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. On the secondary menu, click on **Overview**.
3. On the namespace **Overview** page, click on the **edit** link for the **Permissions** section.

The Console displays a set of checkboxes for the permissions. The permissions that are currently in the namespace permission mask are selected.

4. Select or deselect permissions as needed to modify the permission mask.

Selecting **Purge** automatically selects **Delete**. Selecting **Search** automatically selects **Read**.

5. Click on the **Submit** button.

For an introduction to permission masks, see [“Data access permission masks”](#) on page 15. For more information on the namespace permission mask, see [“Namespace permission mask”](#) on page 72.

Changing the namespace description

The namespace description is optional. You can enter a description or modify the existing description at any time.



Roles: To change the namespace description, you need the administrator role.

To change the description of the namespace:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. On the secondary menu, click on **Overview**.
3. On the namespace **Overview** page, click on the **edit** link for the **Description** section.
4. In the edit area for the description, type the new description of the namespace. The description can be up to 1,024 characters long and can contain any valid UTF-8 characters, including white space.
5. Click on the **Submit** button.

Changing retention-related settings

When the default namespace is created:

- Permission and ownership changes are not allowed for objects under retention
- Only add operations are allowed with custom metadata for objects under retention

You can change these settings at any time.



Roles: To view retention-related settings, you need the monitor, administrator, or compliance role. To change these settings, you need the compliance role.

To change retention-related settings for the namespace:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Policies**.
3. On the left side of the **Policies** page, click on **Retention**.
4. In the **Retention** panel:
 - Optionally, select (to allow) or deselect (to disallow) the **Allow permission and ownership changes for objects under retention** option.
 - Optionally, change the custom metadata setting:
 - To allow the addition, deletion, and replacement of custom metadata for objects under retention, select the **Add, delete, and replace** option.
 - To allow only the addition of custom metadata for objects under retention, select the **Add only** option.
 - To disallow all custom metadata operations for objects under retention, select the **None** option.
5. Click on the **Update Settings** button.

For more information on:

- Ownership and permission changes for objects under retention, see [“Ownership and permission changes for objects under retention”](#) on page 12
- Custom metadata handling, see [“Custom metadata operations for objects under retention”](#) on page 13
- Ownership, permissions, and custom metadata in general, see *Using the Default Namespace*

Enabling or disabling XML checking for custom metadata

When the default namespace is created, custom metadata XML checking is enabled. You can change this setting at any time.



Roles: To view the custom metadata XML checking setting for the namespace, you need the monitor or administrator role. To change the custom metadata XML checking setting for the namespace, you need the compliance role.

To change the custom metadata XML checking setting for the namespace:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Policies**.
3. On the left side of the **Policies** page, click on **Metadata**.
4. In the **Metadata** panel, do one of these:
 - To enable custom metadata XML checking, select the **Check on ingestion that XML in custom metadata files is well-formed** option.
 - To disable custom metadata XML checking, deselect the **Check on ingestion that XML in custom metadata files is well-formed** option.
5. Click on the **Update Settings** button.

For more information on custom metadata XML checking, see [“XML checking for custom metadata”](#) on page 13.

Changing compatibility settings

When the default namespace is created:

- **atime** synchronization is disabled
- The ability to create appendable objects is disabled

You can change these settings at any time.



Roles: To view the compatibility settings for the namespace, you need the monitor or administrator role. To change the compatibility settings for the namespace, you need the administrator role.

To change the compatibility settings for the namespace:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Settings**.
3. On the left side of the **Settings** page, click on **Compatibility**.
4. In the **Compatibility** panel:
 - Optionally, select or deselect the **Synchronize POSIX atime values and object retention settings** option.
 - Optionally, select or deselect the **Allow appendable objects with CIFS and NFS** option.



Note: If you enable appendable objects, you also need to disable disposition. For information on doing this, see [“Changing disposition settings”](#) above.



Note: If you enable both **atime** synchronization and appendable objects, you also need to enable permission and ownership changes for objects under retention. For information on this option, see [“Changing retention-related settings”](#) on page 75.

5. Click on the **Update Settings** button.

For more information on compatibility settings, see [“Compatibility properties”](#) on page 14.

Changing disposition settings

When the default namespace is created, disposition is disabled for both objects with expired retention periods and objects flagged as replication collisions. Once the namespace exists, you can change these settings at any time.



Roles: To view the disposition settings for the namespace, you need the monitor, administrator, or compliance role. To change the disposition settings for the namespace, you need the compliance role.

To change the disposition settings for the namespace:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Services**.
3. On the left side of the **Services** page, click on **Disposition**.
4. In the **Disposition** panel:
 - To enable or disable disposition for objects with expired retention periods, select or deselect the **Automatically delete objects with expired retention periods** option, respectively.
 - To enable or disable disposition for objects flagged as replication collisions, select or deselect the **Automatically delete replication collision objects after ... days** option, respectively.

If you select this option, in the option field, type the number of days objects flagged as replication collisions must remain in the namespace before they are automatically deleted. Valid values are integers in the range zero through 36,500 (that is, 100 years). A value of zero means delete immediately.

5. Click on the **Update Settings** button.

For more information on disposition, see [“Disposition”](#) on page 14.

Changing replication options

When the default namespace is created, the read-from-remote-system option and the option to service HTTP requests redirected from other HCP systems are both enabled, and the collision handling option is set to move

objects. If the HCP system includes the replication feature, you can change these settings at any time. If the HCP system doesn't include the replication feature, these options are not available.



Roles: To view the replication options for the namespace, you need the monitor or administrator role. To change the replication options for the namespace, you need the administrator role.

To change the replication options for the namespace:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Services**.
3. On the left side of the **Services** page, click on **Replication**.

If the HCP system doesn't support replication, the **Replication** option is hidden.

4. In the **Replication** panel:
 - To enable or disable the read-from-remote feature, select or deselect the **Enable read from remote system** option, respectively.
 - To allow or disallow HTTP requests that target the namespace to be redirected from other systems, select or deselect the **Accept requests redirected from other systems in the replication topology** option, respectively.
 - To change the collision handling option, click on **Collision Handling**. Then, in the **Collision Handling** section:
 - To have HCP move objects flagged as replication collisions to the **.lost+found** directory, select the **Move object to the .lost+found directory** option.
 - To have HCP rename objects flagged as replication collisions, select the **Rename object and store in the same location** option.
5. Click on the **Update Settings** button.

For more information on these options, see ["Replication"](#) on page 17.

Changing the service plan

By default, the service plan for the default namespace is Default. You can change this to a different service plan at any time.



Roles: To view the service plan for the default namespace, you need the monitor or administrator role. To change the service plan for the default namespace, you need the administrator role.

To change the service plan for the namespace:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Services**.
3. On the left side of the **Services** page, click on **Service Plan**.
4. In the list of service plans, select the service plan you want.



Note: HCP system administrators can delete service plans regardless of whether they're associated with any namespaces. In this case, the service plan name remains associated with the applicable namespaces, but the service plan is not available to be selected for any namespaces. HCP uses the Default service plan for the namespaces associated with that service plan name.

5. Click on the **Update Settings** button.

For more information on service plans, see ["Service plans"](#) on page 26.

Changing the retention mode

The retention mode of a namespace is either enterprise or compliance. You can change a namespace in enterprise mode to compliance mode, but you cannot do the reverse.

When you change the retention mode of a namespace from enterprise to compliance, you have no guarantee that objects that should have been retained were not already deleted.



Important: Changing the retention mode of a namespace may violate local regulations regarding data retention. Before taking this action, be sure you understand the implications.



Roles: To view the retention mode of the namespace, you need the monitor or administrator role. To change the retention mode of the namespace, you need the administrator role.

To change the retention mode of the default namespace from enterprise to compliance:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Settings**.
3. On the left side of the **Settings** page, click on **Retention Mode**.
4. In the **Retention Mode** panel, select the **Compliance** option.
5. Click on the **Update Settings** button.

A confirming message appears.

6. In the window with the confirming message, select **I understand** to confirm that you understand the consequences of your action. Then click on the **Update Settings** button.

Monitoring the namespace

While the namespace **Overview** page in the Tenant Management Console gives you a view of the default namespace as a whole, these namespace views of the tenant log let you monitor namespace activity on a more detailed level:

- The namespace-level all-events view shows all log messages for the namespace. For more information on this view, see [“Viewing the complete namespace event log”](#) below.
- The namespace-level compliance view shows all log messages about namespace events that require the compliance role. For more information on this view, see [“Viewing the namespace compliance log”](#) on page 82.

Although unlikely, if HCP finds a broken object it cannot repair, it reports the event in the tenant log. In the Tenant Management Console, you can see a list of the irreparable objects in the namespace. For more information on this, see [“Working with irreparable objects”](#) on page 83.

Viewing the complete namespace event log

The namespace **All Events** panel lists all namespace-specific log messages. By default, the panel displays ten messages at a time in reverse chronological order.



Roles: To view the namespace **All Events** panel, you need the monitor, administrator, security, or compliance role. However, only users with the compliance role can see messages about events that require the compliance role.

To display the **All Events** panel:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Monitoring**.
3. On the left side of the **Monitoring** panel, click on **All Events**.

For a description of the information provided by each log message, see [“Understanding log messages”](#) on page 53. For information on managing the message display, see [“Managing the message list”](#) on page 54.

Viewing the namespace compliance log

The namespace **Compliance Events** panel lists all log messages about namespace events that require the compliance role. This includes all retention class activity and privileged delete operations. By default, the panel displays ten messages at a time in reverse chronological order.



Roles: To view the namespace **Compliance** panel, you need the compliance role.

To display the **Compliance Events** panel:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Monitoring**.
3. On the left side of the **Monitoring** panel, click on **Compliance Events**.

For a description of the information provided by each log message, see [“Understanding log messages”](#) on page 53. For information on managing the message display, see [“Managing the message list”](#) on page 54.

For information on retention classes, see [Chapter 8, “Working with retention classes.”](#) on page 155. For information on privileged delete, see [Chapter 9, “Using privileged delete.”](#) on page 161.

Working with irreparable objects

The HCP system keeps track of the irreparable objects it finds. You can view a list of these objects for the default namespace in the Tenant Management Console. To display this list:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Monitoring**.
3. On the left side of the **Monitoring** panel, click on **Irreparable Objects**.



Roles: To view the **Irreparable Objects** panel, you need the monitor or administrator role. To acknowledge irreparable objects, you need the administrator role.

For each object it lists, the **Irreparable Objects** panel shows the full path to the object (starting after `fcfs_data`) and the date and time at which HCP discovered that the object was irreparable. If HCP subsequently repairs a listed object, the object is removed from the list.

You can acknowledge irreparable objects in the **Irreparable Objects** panel. Acknowledging an object leaves a checkmark in the object row. You can use this option to distinguish objects you’ve already seen from objects that have recently become irreparable.

You can delete irreparable objects from the namespace by using normal delete operations as long as the objects are not under retention (or by using privileged delete if the objects are under retention). When you delete an object, HCP removes it from the list of irreparable objects.



Note: Acknowledging that an object is irreparable does not delete the object from the namespace.

By default, the objects in the **Irreparable Objects** panel are listed ten at a time in reverse chronological order by discovery time:

- To view a different number of messages, select the number of messages you want in the **Items per page** field.
- To page forward or backward, click on the next (▶) or back (◀) control, respectively.

To acknowledge one or more objects in the list of irreparable objects:

1. In the **Irreparable Objects** panel, individually select the objects you want to acknowledge, or click on the **select all** link to select all the unacknowledged objects on the current page.

To undo all your selections, click on the **deselect all** link.

2. Click on the **Acknowledge Selected** button.



Tip: To acknowledge all objects, whether selected or not, on *all* pages in a single operation, click on the **Acknowledge All** button.

For information on deleting objects from the namespace, see *Using the Default Namespace*.

Configuring the namespace access protocols

HCP supports the HTTP, WebDAV, CIFS, NFS, SMTP, and NDMP protocols for access to the default namespace. You use these protocols to add, view, modify, delete, backup, and restore the contents of the namespace.

You enable each protocol separately (except for HTTP and WebDAV). You also configure each protocol separately, specifying information such as permission defaults and the IP addresses that can access HCP through it.

This chapter provides instructions for enabling and configuring each of the supported protocols.

For information on using the HTTP, WebDAV, CIFS, and NFS protocols to access the namespace and the SMTP protocol for sending individual emails to the namespace, see *Using the Default Namespace*.

Namespace access protocol configuration

Users and applications have access to the content stored in the default namespace through these industry-standard protocols: HTTP, WebDAV, CIFS, NFS, SMTP, and NDMP. When the default namespace is created, all these namespace access protocols are initially disabled. For any namespace access to occur, at least one protocol must be enabled.



Tip: For enhanced security, disable unused namespace access protocols.

When you enable a namespace access protocol, you also need to configure it. Each protocol, with the exception of HTTP and WebDAV, has its own set of configuration options; HTTP and WebDAV share a set of these options. Some configuration options are common to multiple protocols; others are protocol specific.

When you change the configuration of a protocol that's already enabled, access to the namespace through that protocol is briefly disrupted.

To enable and configure the protocols HCP supports, you use the **Protocols** page in the Tenant Management Console. This page has a separate panel for each protocol except HTTP and WebDAV, which share a panel.

To display the **Protocols** page:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Protocols**.



Roles: To view the **Protocols** page, you need the monitor or administrator role. To enable, disable, and configure namespace access protocols, you need the administrator role.

Specifying IP addresses in Allow and Deny lists

For each namespace access protocol, you have the option of allowing access only from specific IP addresses. For all but NFS, you can also deny access to the namespace from specific IP addresses.



Tip: For enhanced security, restrict access to the namespace to as few IP addresses as possible.


The Tenant Management Console panels for the namespace access protocols each contain an **Allow** list and, except for the **NFS Protocol** panel, a **Deny** list. Each list has an associated field in which you type entries for it.

Adding and removing entries in Allow and Deny lists

To add an entry to an **Allow** or **Deny** list:

1. In the field above the list, type the entry you want. For a description of valid entries, see [“Valid Allow and Deny list entries”](#) below.
2. Click on **Add**.

To remove entries from the **Allow** or **Deny** list:

- To remove a single entry, click on the delete control () for that entry.
- To remove all entries, click on **Delete All**.

Changes you make to either list of IP addresses take effect immediately.

Valid Allow and Deny list entries

Each entry in an **Allow** or **Deny** list can be one of:

- An IP address
- A comma-separated list of IP addresses
- A range of IP addresses specified as *ip-address/subnet-mask* (for example, 192.168.100.197/255.255.255.0) or in CIDR format (for example, 192.168.100.0/24)

The CIDR entry that matches all IP addresses is 0.0.0.0/0.

Allow and Deny list handling

IP addresses can be included in neither, one, or both of the **Allow** and **Deny** lists for HTTP, WebDAV, CIFS, SMTP, and NDMP. The way HCP handles this differs depending on the protocol.

Allow and Deny list handling for HTTP and WebDAV

For HTTP and WebDAV, you can choose how HCP handles **Allow** and **Deny** list entries by selecting or deselecting the **Allow request when same IP is used in both lists** option in the applicable panel or page. The table below describes the effects of selecting or deselecting this option. Either action takes effect immediately.

| List entries | Allow Requests When Same IP Is Used in Both Lists | |
|-------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| | Selected | Not selected |
| Allow list: empty Deny list: empty | All IP addresses can access the namespace through HTTP and WebDAV. | No IP addresses can access the namespace through HTTP or WebDAV. |
| Allow list: at least one entry Deny list: empty | All IP addresses can access the namespace through HTTP and WebDAV. | Only IP addresses in the Allow list can access the namespace through HTTP and WebDAV. |
| Allow list: empty Deny list: at least one entry | All IP addresses not in the Deny list can access the namespace through HTTP and WebDAV. IP addresses in the Deny list cannot. | No IP addresses can access the namespace through HTTP or WebDAV. |
| Allow list: at least one entry Deny list: at least one entry | IP addresses appearing in both or neither of the lists can access the namespace through HTTP and WebDAV. | Only IP addresses appearing in the Allow list and not in the Deny list can access the namespace through HTTP and WebDAV. |

Allow and Deny list handling for CIFS

For CIFS, HCP handles **Allow** and **Deny** list entries as described in the table below.

| List entries | Effect |
|-------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Allow list: empty Deny list: empty | All IP addresses can access the namespace through CIFS. |
| Allow list: at least one entry Deny list: empty | Only IP addresses in the Allow list can access the namespace through CIFS. |
| Allow list: empty Deny list: at least one entry | All IP addresses that are not in the Deny list can access the namespace through CIFS. IP addresses in the Deny list cannot. |
| Allow list: at least one entry Deny list: at least one entry | All IP addresses appearing in the Allow list and only those addresses can access the namespace through CIFS, regardless of whether those addresses also appear in the Deny list. |

Allow list handling for NFS

For NFS, if the **Allow** list in the **NFS Protocol** panel includes one or more IP addresses, those addresses have access to the namespace through NFS and all others don't. If the list is empty, all IP addresses can access the namespace through NFS.

Allow and Deny list handling for SMTP and NDMP

For SMTP and NDMP, HCP handles **Allow** and **Deny** list entries as described in the table below.

| List entries | Effect |
|-------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| Allow list: empty Deny list: empty | All IP addresses can access the namespace through SMTP and NDMP. |
| Allow list: at least one entry Deny list: empty | Only IP addresses in the Allow list can access the namespace through SMTP and NDMP. |
| Allow list: empty Deny list: at least one entry | No IP addresses can access the namespace through SMTP or NDMP. |
| Allow list: at least one entry Deny list: at least one entry | Only IP addresses appearing in the Allow list and not in the Deny list can access the namespace through SMTP and NDMP. |

Specifying default ownership and permissions

All objects in the default namespace have owners, groups, and permissions that follow POSIX standards:

- Each object is associated with one owner and one group, represented by an user ID (UID) and a group ID (GID), respectively.
- A permission is either read, write, or execute. A set of permissions is any combination of these, including none.
- Each object has three sets of permissions — one for its owner, one for its group, and one for all others.



Note: Even if an object has write permission, its data is secure because WORM semantics prevent it from being modified.

For information on the effects of ownership and permissions and on the relationship between the POSIX-style permissions HCP uses and Windows-style permissions, see *Using the Default Namespace*.

Initial values for object owners, groups, and permissions

The POSIX UID, GID, and permissions for an object are initially set when the object is added to the namespace. The values for these properties depend on the namespace access protocol through which the object is added.

For certain protocols, you set default ownership and permission values in the applicable protocol panel. The default values can differ for objects and directories.

The table below describes how HCP gets the UID, GID, and permissions for objects added to the namespace with each protocol (except NDMP, which is used only for backup and recovery).

| Protocol | UID and GID | Permissions |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|----------------------------------------------------------|
| HTTP ¹ , WebDAV, and SMTP | UID and GID defaults set in the applicable protocol panel | Permission defaults set in the applicable protocol panel |
| CIFS | See “UID and GID for objects stored through CIFS” on page 98 | Determined by the client |
| NFS | Determined by the client | Determined by the client |
| 1. For information on overriding UID, GID, and permission defaults in HTTP PUT and MKDIR requests, see <i>Using the Default Namespace</i> . 2. For information on the username mapping file, see “Working with username mapping files” on page 103. | | |



Note: If an object is immediately placed under retention when it's stored, its write permissions are automatically cleared.

Specifying owner, group, and permission defaults

To specify the default owner, group, and permissions for objects added through a particular protocol:

- In the applicable protocol panel, select either **Data Objects** or **Directories** to indicate what you're setting default permissions for. Then do either of the following:
 - Select or deselect the permission for each combination of permission type and owner, group, or other.

- In the applicable **Numeric permission code** field, type the three-digit octal value that corresponds to the permissions you want. For an explanation of these values, see [“Octal permission values”](#) below.



Tip: You can use different methods in the **Data Objects** and **Directories** sections.

- In the **User ID** field, type the UID of the default owner of each object added to the namespace through the protocol. Valid values are integers greater than or equal to zero.
- In the **Group ID** field, type the GID of the default group for each object added to the namespace through the protocol. Valid values are integers greater than or equal to zero.



Tip: Users and applications can change the UID, GID, and permissions for objects that are not under retention. You can allow or disallow these changes for objects that are under retention. For information on allowing this, see [“Changing retention-related settings”](#) on page 75.

Octal permission values

Each permission for owner, group, and other has a unique octal value, as shown in the table below.

| | Read | Write | Execute |
|-------|------|-------|---------|
| Owner | 400 | 200 | 100 |
| Group | 040 | 020 | 010 |
| Other | 004 | 002 | 001 |

You can represent permissions numerically by combining these values. For example, the octal value 755 represents these permissions:

- Owner has read, write, and execute permissions (700).
- Group has read and execute permissions (050).
- Other has read and execute permissions (005).

Configuring the HTTP and WebDAV protocols

With the HTTP and WebDAV protocols, users and applications can add, view, and, when allowed, delete objects and modify object metadata through familiar directory structures.

HTTP and WebDAV protocol configuration

You use the **HTTP & WebDAV** panel to enable and configure the HTTP and WebDAV protocols. To display this panel, on the left side of the **Protocols** page, click on **HTTP & WebDAV**.

The **HTTP & WebDAV** panel lets you:

- Enable the HTTP and WebDAV protocols
- Activate SSL security for HTTP and WebDAV
- Specify the client IP addresses that have access to the namespace through HTTP and WebDAV
- Specify the default UID, GID, and permissions for objects and directories added to the namespace through HTTP or WebDAV
- Control permission checking for namespace access through HTTP and WebDAV
- Allow or prevent ownership and permission overrides through HTTP and ownership and permission changes through HTTP and WebDAV
- For WebDAV only, control the use of basic authentication and the handling of dead properties

For information on using the HTTP and WebDAV protocols for namespace access, see *Using the Default Namespace*. For information on WebDAV basic authentication and dead properties, see the WebDAV specification at <http://www.webdav.org/specs/rfc2518.html>.

Enabling HTTP and WebDAV access to the namespace

The **HTTP & WebDAV** panel has four sections for enabling and configuring the HTTP and WebDAV protocols.

Settings section

To enable the HTTP and WebDAV protocols, in the **Settings** section:

1. Do either or both of these:
 - To enable the HTTP and WebDAV protocols without SSL security, select the **Enable HTTP and WebDAV protocols** option.
 - To enable and secure the HTTP and WebDAV protocols, select the **Enable SSL for HTTP and WebDAV protocols** option.

These two options are independent of each other. If you select both, users and applications can send both secure and unsecure data through the HTTP and WebDAV protocols.

2. Click on the **Update Settings** button in the **Settings** section.

Allow/Deny section

To set the IP addresses to be allowed or denied access to the namespace through HTTP and WebDAV:

- Optionally, in the **Allow/Deny** section, specify IP addresses to be allowed or denied access to the namespace through HTTP and WebDAV. For instructions on doing this, see [“Adding and removing entries in Allow and Deny lists”](#) on page 87.
- To specify how HCP should handle IP addresses that appear in both or neither of the **Allow** and **Deny** lists, select or deselect the **Allow request when same IP is used in both lists** option. For the effects of this option, see [“Allow and Deny list handling for HTTP and WebDAV”](#) on page 88.

Data Object and Directory Permissions section

To set owner, group, and permission options:

1. Click on **Data Object and Directory Permissions**.
2. In the **Data Object and Directory Permissions** section:
 - Set the default UID, GID, and permissions for objects and directories added to the namespace through HTTP or WebDAV. For instructions on doing this, see [“Specifying owner, group, and permission defaults”](#) on page 90.

- Optionally, select the **Allow UID, GID, and permission overrides and changes** option to allow:
 - Overrides of UID, GID, and permissions in HTTP **PUT** and **MKDIR** requests
 - Changes to UID, GID, and permissions with HTTP **CHMOD** and **CHOWN** requests and WebDAV **PROPPATCH** requests
- In the **Permission Enforcement** field:
 - To disable permission checking on HTTP and WebDAV requests, select **Do not check permissions**.
 - To enable permission checking only on the final object or directory appearing in each HTTP or WebDAV request, select **Check permissions only on first object or directory**. This includes permission checking on the target directory when a new object is added to the namespace.
 - To enable permission checking on each object and each directory appearing anywhere in a path in each HTTP or WebDAV request, select **Strict POSIX permission checking**.



Note: Strict permission checking enhances the security of stored data but results in slower performance. With no permission checking, performance is unaffected, but you gain no security benefit.



Tip: To learn the permission-checking requirements for your applications, contact the application vendors.

3. Click on the **Update Permissions** button.

Additional WebDAV Settings section

To set WebDAV-specific options:

1. Click on **Additional WebDAV Settings**.

2. In the **Additional WebDAV Settings** section:

- To use basic authentication for WebDAV access to the namespace, select the **Enable basic authentication** option. Then:
 - In the **Username** field, type the username to use for basic authentication. Usernames must be from one through 64 characters long and can contain any valid UTF-8 characters but cannot start with an opening square bracket ([). White space is allowed.

Usernames are not case sensitive.

- In the **Password** field, type the password to use for basic authentication. Passwords can up to 64 characters long, are case sensitive, and can contain any valid UTF-8 characters, including white space. The minimum password length is configured at the HCP system level. Typically, it's six or eight characters.

To be valid, a password must include at least one character from two of these three groups: alphabetic, numeric, and other.

If you're modifying settings in the **HTTP(S)** panel and you leave the **Password** field empty, the previously set password remains in effect.

- In the **Confirm Password** field, type the password again.

Selecting this option has no effect on the HTTP protocol.



Tip: Be sure to tell WebDAV users the username and password you specify.

- To have WebDAV store dead properties and other properties that don't correspond to HCP metadata in `custom-metadata.xml` files, select the **Use custom metadata to store WebDAV properties** option.

Selecting this option has no effect on the HTTP protocol.

For more information on storing dead properties with WebDAV, see *Using the Default Namespace*.

3. Click on the **Update Settings** button in the **Additional WebDAV Settings** section.

Configuring the CIFS protocol

With the CIFS protocol, users and applications can add, view, and, when allowed, delete objects and modify object metadata through familiar directory structures.

CIFS protocol configuration

You use the **CIFS** panel to enable and configure the CIFS protocol. To display this panel, on the left side of the **Protocols** page, click on **CIFS**.

The **CIFS** panel lets you:

- Enable the CIFS protocol
- Specify the client IP addresses that have access to the namespace through CIFS
- Specify whether the CIFS protocol requires user authentication for access to the namespace or allows anonymous or authenticated access
- Depending on the HCP system configuration, either specify the UID and GID for objects added to the namespace through CIFS or optionally provide a list of username mappings
- Change CIFS case sensitivity

When you reconfigure the CIFS protocol while it's already enabled, the changes you make don't affect current CIFS mounts of the namespace. This includes changing from allowing anonymous access to requiring authenticated access. To force changes to take effect, you can do either of these:

- Disable and then reenable the protocol. This causes all CIFS clients to lose their connections to the namespace. When they reconnect, the changes will be in effect.
- Direct all clients with current CIFS mounts to disconnect from the namespace and then to either reboot or wait five minutes for cached connections to be released before reconnecting.

For information on using the CIFS protocol for namespace access, see *Using the Default Namespace*.

CIFS access to the namespace

The way in which users access the namespace through CIFS depends on:

- Whether HCP support for Windows Active Directory (AD) is enabled at the system level.
- Whether the user has an account in AD.
- Whether the CIFS protocol requires user authentication for access to the namespace or allows anonymous or authenticated access

If AD support is enabled and:

- The CIFS protocol requires authenticated access, the user must have an account in AD and must provide valid credentials for that account when starting the CIFS session
- The CIFS protocol allows anonymous or authenticated access, HCP first checks whether the user has provided AD credentials for the CIFS session:
 - If the user has provided credentials and the credentials are valid, HCP gives the user access to the namespace.
 - If the user has provided credentials but the specified username does not exist in AD, HCP gives the user anonymous access to the namespace.
 - If the user has provided credentials but the specified password is invalid, HCP does not give the user access to the namespace.
 - If the user has not provided AD credentials, HCP gives the user anonymous access to the namespace.

If AD support is not enabled and:

- The CIFS protocol requires authenticated access, HCP does not give the user access to the namespace
- The CIFS protocol allows anonymous or authenticated access, HCP gives the user anonymous access to the namespace

UID and GID for objects stored through CIFS

If AD support is enabled at the system level, the CIFS protocol configuration can optionally include username mappings. A username mapping associates an AD user with a UID and GID. Only one username can map to any given UID. For more information on username mappings, see [“Working with username mapping files”](#) above.

If AD support is not enabled at the system level, the CIFS protocol configuration includes a UID and GID that are associated with each user who accesses the namespace anonymously.

The UID and GID for an object stored through the CIFS protocol depend on the way in which the user storing the object accessed the namespace, as described in [“CIFS access to the namespace”](#) above. They also depend on whether a username mapping exists for the user.

The table below shows the UID and GID assigned to objects stored through the CIFS protocol while AD support is enabled.

| User | Authenticated access only | Anonymous or authenticated access |
|-----------------------------------------------|-------------------------------------------------------------------|-------------------------------------------------------------------|
| User exists in AD and has a username mapping | UID and GID from username mapping | UID and GID from username mapping |
| User exists in AD but has no username mapping | UID and GID arbitrarily assigned to the user for the CIFS session | UID and GID arbitrarily assigned to the user for the CIFS session |
| User does not exist in AD | Access denied | UID 99 and GID 99 |

While AD support is not enabled and either anonymous or authenticated access is allowed, the object gets the UID and GID specified in the CIFS configuration.

These considerations apply to the UIDs and GIDs for objects stored through the CIFS protocol:

- When an AD user without a username mapping starts a CIFS session, HCP assigns that user an arbitrary UID and GID. HCP maintains the association between that user and the UID and GID for the duration of the session and for a limited amount of time after the session ends. If the user starts a new session after that time, HCP assigns a new UID and GID. The user, therefore, is no longer the owner of objects stored during the previous session and no longer has owner permissions for them.

If maintaining the connection between objects and the users who stored them is important, you should ensure that all AD users have username mappings.

- When an authenticated user with a username mapping adds an object to the namespace through the CIFS protocol, the UID in the mapping becomes the object owner. If you then map that username to a different UID, the user who added the object is no longer its owner and, therefore, no longer has owner permissions for it. To resolve this situation, run a script to change the ownership of each affected object.
- You can change the UID and GID specified for anonymous access at any time. If you do that, anonymous users will no longer have owner permissions for objects stored with the previous UID and GID. To resolve this situation, run a script to change the ownership of each affected object.
- If the access method changes from anonymous or authenticated to authenticated only, at least one AD user account needs to be mapped to a UID that matches the UID used for anonymous access. If that is not the case, no user will have owner permissions for objects stored using anonymous access. To resolve this situation, run a script to change the ownership of each affected object.
- With the UID set to 0 (zero) for anonymous access or the AD user account mapped to 0 (zero) for authenticated access, all objects stored by the user appear writable, even though HCP WORM semantics prevent their fixed-content data from being modified.
- HCP does not support Active Directory access control lists (ACLs).

CIFS case sensitivity

The Windows operating system is case preserving but not case sensitive. The HCP CIFS implementation, by default, is both case preserving and case sensitive. One result of this discrepancy is that Windows applications that do not observe differences in case may not be able to access HCP objects by name.

For example, suppose a Windows application adds a file named `File.txt` to the namespace by using the CIFS protocol. CIFS preserves case, so the namespace then contains an object named `File.txt`. Now suppose the application tries to retrieve that object using the name `file.txt`. CIFS is case sensitive, so it passes the request to HCP with only the name `file.txt`. It doesn't include any case variations on the name, such as `File.TXT`, `FILE.txt`, or `File.txt`. As a result, HCP cannot find the object.

If you have Windows applications that ignore case, you may want HCP to ignore case as well. You can change the CIFS protocol configuration in either of two ways to meet this need:

- **Make CIFS case forcing** — With this behavior, CIFS changes names to all upper- or lowercase in the requests it passes to HCP. To Windows applications, then, HCP appears to be case-insensitive. An application that stores `File.txt` and then retrieves `File.TXT` will get the right object back.

The drawback to this method is that applications using other namespace access protocols must accommodate this behavior. For example, suppose CIFS changes names to all uppercase. If an application using the CIFS protocol stores an object named `File.txt`, applications using the case-sensitive HTTP, WebDAV, and NFS protocols need to retrieve the object as `FILE.TXT`.

- **Make CIFS case insensitive** — With this behavior, CIFS preserves case as objects are stored in the namespace but passes through every case variation possible when applications make other requests for objects.

For example, suppose an application using the CIFS protocol requests an object named `FILE.txt`. CIFS passes the request through with the names `File.txt`, `FILE.txt`, `file.TXT`, and so on. HCP then returns the first object it finds with a name that matches any of these.

The major drawback to this method is that performance is slowed by the need to check for matches to multiple case variations. A second drawback is that if the namespace contains multiple objects with names that differ only in case, HCP may return the wrong object.

If you make CIFS both case forcing and case insensitive, it is case forcing when storing objects and case insensitive on requests for existing objects.

For more information on CIFS case sensitivity, see *Case Sensitivity versus Case Preservation in CIFS Server (Samba)* at http://www.manualshark.org/manualshark/files/28/pdf_27202.pdf.

Enabling CIFS access to the namespace

The **CIFS** panel has three sections for enabling and configuring the CIFS protocol.

Settings section

To enable the CIFS protocol, in the **Settings** section:

1. Select the **Enable CIFS protocol** option.
2. Click on the **Update Settings** button.

Allow/Deny section

Optionally, in the **Allow/Deny** section, specify IP addresses to be allowed or denied access to the namespace through CIFS. For instructions on doing this, see [“Adding and removing entries in Allow and Deny lists”](#) on page 87.

For information on how HCP handles IP addresses that appear in both or neither of the **Allow** and **Deny** lists, see [“Allow and Deny list handling for CIFS”](#) on page 88.

Authentication section

To configure the access method for the CIFS protocol:

1. Click on **Authentication**.
2. In the **Authentication** section:
 - Do either of these:
 - To require authentication with the CIFS protocol, select the **Authenticated access only** option.
 - To allow both anonymous and authenticated access to the namespace with CIFS, select the **Anonymous or authenticated access** option.

- Do either of these:
 - If the **Authentication** section contains the **Username Mapping File** field, optionally click on the **Browse** button to select a username mapping file for upload. For information on username mapping files, see [“Working with username mapping files”](#) below.



Important: Uploading a username mapping file while the CIFS protocol is enabled causes HCP to automatically disable and then reenables CIFS for all namespaces, including both the default namespace and HCP namespaces. This causes all CIFS clients to lose their connections to HCP. If you need to upload a new username mapping file, ask your HCP system administrator when would be the best time to do it.

- If the **Authentication** section contains the **User ID** and **Group ID** fields, specify the UID and GID for the CIFS protocol. For instructions on doing this, see [“Specifying default ownership and permissions”](#) on page 89.
3. Click on the **Update Settings** button in the **Authentication** section.

Case Sensitivity section

To change CIFS case sensitivity:

1. Click on **Case Sensitivity**.
2. In the **Case Sensitivity** section:
 - To make the CIFS protocol case insensitive, deselect the **Make CIFS case sensitive** option.



Note: Disabling CIFS case sensitivity has a significant negative impact on performance.

- To make the CIFS protocol case forcing, select **Make CIFS case forcing**. Then select either **Lowercase** or **Uppercase** to force object names to be lower- or uppercase, respectively.
3. Click on the **Update Settings** button in the **Case Sensitivity** section.

For more information on making the CIFS protocol case insensitive or case forcing, see [“CIFS case sensitivity”](#) on page 99.

Working with username mapping files

A username mapping file is a plain text file that maps AD usernames to UID and GID pairs. Each entry in the file must be on a separate line and have this format:

```
username→uid→gid
```

For each entry:

- *username* is a valid account in the Active Directory domain.
- *uid* associates a numeric user ID with *username*.
- *gid* associates a numeric group ID with *username*.

Creating a username mapping file

You can use any text or spreadsheet editor to create a username mapping file. You can give the file any name you want.

If you're using a spreadsheet editor, be sure to save the file in a tab-delimited format. For example, in Microsoft® Excel®, you would save the file as type **Text (Tab-delimited) (*.txt)**.

After creating a username mapping file, you need to upload it to HCP. For instructions on uploading the file, see ["Enabling CIFS access to the namespace"](#) above.

Viewing username mappings

To view the set of username mappings CIFS is currently using, click on the **View Username Mappings** link in **Authentication** section in the **CIFS** panel.



Tip: You can view username mappings even if your user account includes only the monitor role.

Adding, modifying, and deleting username mappings

To add, modify, and delete username mappings in the set CIFS is currently using:

1. In **Authentication** section in the **CIFS** panel, click on the **View Username Mappings** link.
2. In the **Username Mappings** window, click on **Download Username Mappings**.

3. Save and, optionally, rename the downloaded file. (The name of the file HCP downloads is always `user_mappings.csv`.)
4. In the text or spreadsheet editor of your choice, add, modify, and delete mappings in the saved file, as needed.
5. Save the file in a tab-delimited format. For example, in Microsoft Excel, you would save the file as type **Text (Tab-delimited) (*.txt)**.
6. Upload the edited file, as described in [“Enabling CIFS access to the namespace”](#) above.

Configuring the NFS protocol

With the NFS protocol, users and applications can add, view, and, when allowed, delete objects and modify object metadata through familiar directory structures.

NFS protocol configuration

You use the **NFS** panel to enable and configure the NFS protocol. To display this panel, on the left side of the **Protocols** page, click on **NFS**.

The **NFS** panel lets you:

- Enable the NFS protocol
- Specify the client IP addresses that have access to the namespace through NFS

When you reconfigure the NFS protocol while it's already enabled, the changes you make don't affect current NFS mounts of the namespace. As a result, clients with IP addresses that are no longer in the **Allow** list remain connected. To break those connections, disable and then reenabling the protocol. This causes all NFS clients to lose their connections to the namespace. Only those clients in the changed **Allow** list can then reconnect to the namespace.

For information on using the NFS protocol to access stored data, see *Using the Default Namespace*.

Enabling NFS access to the namespace

The **NFS** panel has two sections for enabling and configuring the NFS protocol.

Settings section

To enable the NFS protocol, in the **Settings** section:

1. Select the **Enable NFS protocol** option.
2. Click on the **Update Settings** button.

Allowed IP Addresses section

To set the IP addresses to be allowed access to the namespace through NFS, optionally, specify IP addresses in the **Allowed IP Addresses** section. For instructions on doing this, see [“Adding and removing entries in Allow and Deny lists”](#) on page 87.

Configuring the SMTP protocol

With the SMTP protocol, HCP can automatically archive emails forwarded by email servers. The protocol also enables users and applications to send individual emails to the default namespace.

For information on sending individual emails to the namespace, see *Using the Default Namespace*.

SMTP protocol configuration

You use the **SMTP** panel to enable and configure the SMTP protocol. To display this panel, on the left side of the **Protocols** page, click on **SMTP**.

The **SMTP** panel lets you:

- Enable the SMTP protocol
- Specify the email server IP addresses that have access to the namespace through SMTP
- Specify the index setting for email objects and directories added to the namespace through SMTP
- Specify the shred setting for email objects and directories added to the namespace through SMTP

- Specify the retention setting for email objects and directories added to the namespace through SMTP (see [“Email retention setting”](#) below)
- Specify the UID, GID, and permissions for email objects and directories added to the namespace through SMTP
- Specify where and in what format email objects are stored
- Specify whether to store email attachments separately

The SMTP protocol always stores attachments along with the email they accompany. The HDDS and HCP search facilities and the metadata query engine index the attachments along with the email.

You can choose to additionally store attachments separately from the email they accompany. With this option, searches return not only the original email with the attachments but also the attachments as separate objects.

When attachments are stored only with the email they accompany, searches return only the email objects. You then need to retrieve the email objects and separate the attachments yourself.

Storing attachments as separate objects can have a significant impact on performance and storage space.

Email retention setting

When configuring the SMTP protocol, you specify the retention setting to be assigned to emails added to the namespace through that protocol. Valid values for this setting are:

- A retention class. For information on retention classes, see [Chapter 8, “Working with retention classes,”](#) on page 155.
- One of these special values:
 - **Deletion Allowed** — The email object can be deleted at any time.
 - **Deletion Prohibited** — The email object can never be deleted by means of a normal delete operation. If the namespace is in enterprise mode, however, the object can be deleted by means of a privileged delete operation.

Once an object has this retention setting, its retention setting cannot be changed.

- **Initial Unspecified** — The email object cannot be deleted, but can have its retention setting changed to any other retention setting.
- A fixed date.



Tip: If you specify a fixed date, remember to change the default retention setting again before that date occurs. Otherwise, new objects will have a retention setting in the past and will immediately be deletable.

Enabling SMTP access to the namespace

The **SMTP** panel has five sections for enabling and configuring the SMTP protocol.

Settings section


To enable the SMTP protocol, in the **Settings** section:

1. Select the **Enable SMTP protocol** option.
2. Click on the **Update Settings** button in the **Settings** section.

Default SMTP Settings section

To set defaults for new emails added through the SMTP protocol, in the **Default SMTP Settings** section:

1. Do any of these:
 - Optionally, change the index setting for new emails added through SMTP by selecting (to index) or deselecting (not to index) the **Index objects** option.
 - Optionally, change the shred setting for new emails added through SMTP by selecting (to shred) or deselecting (not to shred) the **Shred on delete** option.
 - Optionally, change the retention setting for new emails added through SMTP:
 - To make the retention setting a retention class, in the **Retention Method** field, select **Retention Class**. Then, in the **Retention Class** field, select the retention class you want.

- To make the retention setting a special value, in the **Retention Method** field, select **Special Value**. Then, in the **Special Value** field, select the special value you want.
- To make the retention setting a fixed date, in the **Retention Method** field, select **Fixed Date**. Then, either type a date in the **Fixed Date** field or click on the calendar icon () to select a date. If you type a date, use this format: *mm/dd/yyyy*

If you specify certain invalid dates, HCP automatically adjusts the value to make a real date. For example, if you specify a retention setting of 11/33/2015, which is three days past the end of November, email objects added to the namespace get a retention setting of 12/03/2015.

2. Click on the **Update Settings** button in the **Default SMTP Settings** section.

Allow/Deny section

Optionally, in the **Allow/Deny** section, specify IP addresses of email servers to be allowed or denied access to the namespace through SMTP. For instructions on doing this, see [“Adding and removing entries in Allow and Deny lists”](#) on page 87.

For information on how HCP handles IP addresses that appear in both or neither of the **Allow** and **Deny** lists, see [“Allow and Deny list handling for CIFS”](#) on page 88.

Data Object and Directory Permissions section

To set owner, group, and permission options:

1. Click on **Data Object and Directory Permissions**.
2. In the **Data Object and Directory Permissions** section, set the default UID, GID, and permissions for emails and directories added to the namespace through SMTP. For instructions on doing this, see [“Specifying owner, group, and permission defaults”](#) on page 90.
3. Click on the **Update Permissions** button.

Emails section

To set options specific to email objects:

1. Click on **Emails**.

2. In the **Emails** section:

- In the **Email Location** field, type the path for the directory in which you want email objects stored. This path is appended to the `fcfs_data` directory. Be sure to start and end the path with a forward slash (`/`), like this:

```
/email/company_all/
```

If any part of the specified directory path doesn't exist, HCP creates it.



Important: If the email directory is being replicated and email is also being sent to other HCP systems in the replication topology, the email directories on the systems should have different names.

For information on the `fcfs_data` directory and examples of the complete path and object names for email added to the namespace through SMTP, see *Using the Default Namespace*.

- In the **Format** field, select either **.eml** or **.mbox**. The one you choose depends on the application you use to read the stored email.
- To store email attachments separately from the emails they're attached to, select the **Separate attachments from parent email** option. For more information on storing email attachments, see ["SMTP protocol configuration"](#) on page 105.



Important: Storing attachments separately from the email they accompany can have a significant impact on performance and storage space. Unless you have a specific reason to do so, do not enable this option.

3. Click on the **Update Settings** button in the **Emails** section.

Configuring Microsoft Exchange for email archiving through SMTP

HCP provides SMTP support for Microsoft Exchange 2003, 2007, and 2010 email servers. The following sections outline the procedures for configuring each supported version of Exchange for email archiving to the default namespace through SMTP. For information on configuring Exchange for email archiving to HCP namespaces, see *Managing a Tenant and Its Namespaces*.

For considerations that apply to configuring Exchange 2003 and 2010 in an environment that uses both versions, see “Exchange 2003 and Exchange 2010 Journaling Interoperability” (<http://technet.microsoft.com/en-us/library/aa997918.aspx#Exc>).



Note: HCP supports archiving email to multiple namespaces. However, with Microsoft Exchange 2003, the recommendation is to archive email to only one namespace.

Configuring Microsoft Exchange 2003

To configure Microsoft Exchange 2003 for archiving emails to the default namespace through SMTP:

1. Create a custom SMTP recipient. For the new email address, select **SMTP Address**. For details, see “How to Create a Custom SMTP Recipient for Exchange Server 2003 Journaling” (<http://technet.microsoft.com/en-us/library/bb124642%28EXCHG.65%29.aspx>).
2. Create journal recipient mailboxes. For details, see “Planning an Exchange Server 2003 Journaling Deployment” (<http://technet.microsoft.com/en-us/library/aa998762%28EXCHG.65%29.aspx>).



Tip: For better performance, create the journal recipient mailboxes on separate servers from the user mailbox servers.

3. Define a server-side forwarding rule for the journal recipient mailboxes. For details, see “How to Set a Server-Side Rule for Journal Recipient Mailboxes” (<http://technet.microsoft.com/en-us/library/bb124633%28EXCHG.65%29.aspx>).
4. Configure Exchange Mailbox Manager to clean out the journal recipient mailboxes after the journalized messages in them are transmitted. For details, see “How to Configure Mailbox Manager to Clean the Journal Recipient Mailbox” (<http://technet.microsoft.com/en-us/library/aa995756%28EXCHG.65%29.aspx>).
5. Configure a dedicated SMTP connector to transmit journalized messages to `smtp.hcp-domain-name`. For details, see “How to Create an SMTP Connector” (<http://technet.microsoft.com/en-us/library/aa996625%28EXCHG.65%29.aspx>).

6. For each mailbox store, enable standard journaling. For details, see “How to Enable Standard Journaling” (<http://technet.microsoft.com/en-us/library/bb124786%28EXCHG.65%29.aspx>).
7. Enable envelope journaling. For details, see “How to Enable Envelope Journaling” (<http://technet.microsoft.com/en-us/library/aa997541%28EXCHG.65%29.aspx>).

For additional information on this procedure, see “Implementing Exchange 2003 Message Journaling” (<http://www.msexchange.org/tutorials/Implementing-Exchange-2003-Message-Journaling.html>).

Configuring Microsoft Exchange 2007

To configure Microsoft Exchange 2007 for archiving emails to the default namespace through SMTP:

1. Create a custom SMTP recipient. For the email address, use *username@hcp-domain-name*, where *username* is any new or existing username. For details, see “How to Create a New Mail Contact” ([http://technet.microsoft.com/en-us/library/aa997220\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/aa997220(EXCHG.80).aspx)).



Tip: For *username*, use admin.

2. Create the mailboxes to be journaled. For details, see “How to Create a Mailbox for a New User” ([http://technet.microsoft.com/en-us/library/aa998197\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/aa998197(EXCHG.80).aspx)).



Tip: For better performance, create the journal mailboxes on separate servers from the user mailbox servers.

3. Create a custom Send connector for the Exchange server. For the address space, use *hcp-domain-name*. Choose smart host for email routing. For details, see “How to Create a New Send Connector” ([http://technet.microsoft.com/en-us/library/aa998814\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/aa998814(EXCHG.80).aspx)).
4. Create a journal rule to send journal reports to the custom SMTP recipient you created in [step 1](#) above. For details, see “How to Create a New Journal Rule” ([http://technet.microsoft.com/en-us/library/bb124723\(EXCHG.80\).aspx](http://technet.microsoft.com/en-us/library/bb124723(EXCHG.80).aspx)).

Configuring Microsoft Exchange 2010

To configure Microsoft Exchange 2010 for archiving emails to the default namespace through SMTP:

1. Create the user mailboxes to be journaled, along with the new user for each mailbox. For the user email address use, *username@hcp-domain-name*, where *username* is any new or existing username. For details, see "Create a Mailbox" (<http://technet.microsoft.com/en-us/library/bb123809.aspx>).



Tip: For *username*, use admin.

2. Create a new SMTP Send connector for the Exchange server. For the address space, use *hcp-domain-name*. For the smart host authentication settings, select **None**. For details, see "Create an SMTP Send Connector" (<http://technet.microsoft.com/en-us/library/aa997285.aspx>).
3. Optionally, if you have the Enterprise edition of Microsoft Exchange 2010, create a journal rule to selectively journal emails. For details, see "Create a Journal Rule" (<http://technet.microsoft.com/en-us/library/aa995915.aspx>).

Configuring the NDMP protocol

The NDMP protocol enables you to back up and restore objects in the default namespace using third-party backup applications. HCP supports NDMP v4.

For a list of backup applications that have been tested with HCP, please contact your HCP system administrator.

Backup and restore operations

In the Tenant Management Console, you can enable backup and restore operations separately from each other. This capability lets you disable restore operations until you actually need to perform them while leaving backup operations enabled at all times (for example, to facilitate daily backups of the stored data). By disabling restore operations, you prevent them from occurring except when you explicitly allow them.

HCP backs up objects in OpenPGP format, which uses a tar file to package files that represent an object. This standard format, which can be both signed and encrypted, allows backed-up objects to be restored to other storage systems.

OpenPGP format

With the OpenPGP format, you can back up or restore the entire namespace or selected objects in it, depending on the parameters supported by the backup application you're using. The target of a restore operation can be the original HCP system the objects came from or any other HCP system.

For each object you back up, HCP creates an OpenPGP **HCP object package (HOP)**. The HOP contains the object data, metadata, and, if applicable, custom metadata.

For each object to be restored, if the target directory already contains an object that:

- Is identical to the backup object in both data and metadata, no action is taken
- Differs from the backup object only in system metadata that can be modified (such as retention and permissions) or in custom metadata, the metadata for the backup object is used to update the metadata in the namespace, if such updates are allowed
- Differs from the backup object in data or metadata that cannot be modified, the restore of that object fails, and HCP sends a message to the backup application

Under no circumstances does HCP replace an existing object in the namespace. This is true even for objects that are not under retention.

If an object being restored has a retention setting that's a retention class and that class doesn't already exist in the namespace, HCP creates the class with a value of **Initial Unspecified**. For information on retention classes, see [Chapter 8, "Working with retention classes,"](#) on page 155.



Note: To ensure that objects restored through the NDMP protocol are included in the applicable search indexes, HCP sets the POSIX **ctime** value of each restored object to the current time. This also ensures that the objects are replicated if replication is in effect.

Signing and encryption keys

The first time HCP starts, it generates both a signing key pair and an encryption key pair for backup and restore operations. When you back up the default namespace, you can cryptographically sign and/or encrypt the backup data:

- When signing backup data, HCP uses its private signing key. To restore the data to a different HCP system, you need to download the public signing key from the original system and upload it into the target system. If the signing key is lost, you can choose not to verify signatures when restoring signed data, although this disables cryptographic verification of the restored data.

The default signing algorithm for HCP is SHA-256. To use a different algorithm, please contact your authorized HCP service provider.

- When encrypting backup data, HCP uses its public encryption key. To restore the encrypted data to a different HCP system, you need to download the private encryption key from the original system and upload it into the target system. (HCP actually downloads and uploads both portions of the encryption key together, enabling you to use the uploaded key as the encryption key for a backup operation on another system.)

Downloaded encryption keys are password protected.



Important: Be sure to keep a backup copy of the downloaded encryption key. If the encryption key is lost, the backup data is not recoverable.

For instructions on downloading and uploading signing and encryption keys, see [“Working with signing and encryption keys”](#) on page 118.

Backup performance considerations

How long a backup operation takes depends on several factors, including:

- The size of the HCP system.
- The number of objects in the namespace.
- The CPU speed of the HCP servers.

- Whether the data is being signed, encrypted, and/or compressed. Using signing, encryption, or compression results in slower performance for both backup and restore operations.
- The network bandwidth between HCP and the backup device.

With NDMP applications that support multiple backup streams, simultaneous backups can lessen the time it takes to back up the namespace. They are recommended for larger systems — the smaller the system, the less time you save by running multiple streams.

HCP supports up to 30 simultaneous backup streams. To enable simultaneous backups, you need to segment the namespace by assigning different streams to different directory trees.

NDMP protocol configuration

You use the **NDMP** panel to enable and configure the NDMP protocol. To display this panel, on the left side of the **Protocols** page, click on **NDMP**.

The **NDMP** panel lets you:

- Enable the NDMP protocol
- Specify the client IP addresses that have access to the namespace through NDMP
- Enable backup and restore operations separately and control the signing, compression, and encryption of objects in OpenPGP format
- Set the level of user authentication required for backup and restore operations
- Download and upload signing and encryption keys for backup and restore operations and set the active encryption key for backup operations

Enabling NDMP access to the namespace

The **NDMP** panel has four sections for enabling and configuring the NDMP protocol and an additional section for working with signing and encryption keys.

Settings section

To enable the NDMP protocol, in the **Settings** section:

1. Select the **Enable NDMP protocol** option.
2. Click on the **Update Settings** button.

Allow/Deny section

Optionally, in the **Allow/Deny** section, specify IP addresses to be allowed or denied access to the namespace through NDMP. For instructions on doing this, see [“Adding and removing entries in Allow and Deny lists”](#) on page 87.

For information on how HCP handles IP addresses that appear in both or neither of the **Allow** and **Deny** lists, see [“Allow and Deny list handling for SMTP and NDMP”](#) on page 89.

Backup and Restore section

To set backup and restore options:

1. Click on **Backup and Restore**.
2. In the **Backup and Restore** section:
 - To allow HCP backups through NDMP, select the **Enable backups** option. This option takes effect only if the **Enable NDMP protocol** option is also selected.
 - To allow HCP restores through NDMP, select the **Enable restores** option. This option takes effect only if the **Enable NDMP protocol** option is also selected.
 - To cryptographically sign each HOP when using the OpenPGP format, select the **Sign backup data** option.
 - To compress each HOP when using the OpenPGP format, select the **Compress backup data** option.
 - To encrypt each HOP when using the OpenPGP format, select the **Encrypt backup data** option.
 - To ignore signatures during restore operations when using the OpenPGP format, select the **Ignore signatures when restoring** option. When HCP ignores signatures, it accepts unsigned objects and objects with signatures for which it doesn't have the applicable key.
3. Click on the **Update Settings** button in the **Backup and Restore** section.

Authentication section

To set the level of user authentication required for backup and restore operations:

1. Click on **Authentication**.
2. In the **Authentication** section:
 - To specify the types of user authentication HCP accepts from backup applications, select one or more of:
 - **Allow unauthenticated operations** — The backup and restore application is not required to provide authentication to back up and restore data in the namespace.
 - **Allow username/pwd authenticated operations** — Only backup applications providing the username and password specified in this section can back up and restore data in the namespace. With this level of authentication, usernames and passwords are sent over the network as clear text.
 - **Allow digest-authenticated operations** — Only backup applications providing the username and password specified in this section can back up and restore data in the namespace. With this option, the username and password are handled in a more secure manner than with username/password authentication.

If you don't select any authentication options, backup and restore operations are not allowed.



Note: Some backup applications require specific types of user authentication.

- If you select **Allow username/pwd authenticated operations** or **Allow digest-authenticated operations** as an accepted type of authentication:
 - In the **Username** field, type the username that's authorized to perform backup and restore operations with HCP.
 - In the **Password** field, type the password associated with that username.
 - In the **Confirm Password** field, type the password again.
3. Click on the **Update Settings** button in the **Authentication** section.

Working with signing and encryption keys

You use the **Keys** section in the **NDMP** panel to:

- Download NDMP signing and encryption keys
- Upload NDMP signing and encryption keys
- Delete NDMP signing keys (you cannot delete NDMP encryption keys)

To open this section, click on **Keys**.


To see the list of currently installed signing keys, click on **Signing Keys**. To see the list of currently installed encryption keys, click on **Encryption Keys**.

For each key in the lists of signing and encryption keys, the **Keys** section shows:

- The name of the key. For the system-generated keys, this name includes the name of the HCP system.
- The date and time the key was created.
- The bit strength of the key.
- The fingerprint of the key.

Downloading a key

To download a signing or encryption key in the **Signing Keys** or **Encryption Keys** list:

1. Click on the download control () for the key.

For an encryption key, the download requires you to specify a password. HCP uses this password to encrypt the key before downloading it.

In the **Download Encryption Key** window that opens:

- In the **Secret Key Password** field, type a password for the key. This password can contain any valid UTF-8 characters, including white space.
- In the **Confirm Password** field, type the password again.

Then click on the **Download** button.

2. Save the key file in the location of your choice.

Uploading a signing key

To upload a signing key:

1. In the **Keys** section, click on **Signing Keys**.
2. Click on the **Browse** button for the **Upload Key** field and select the key file you want.
3. Click on the **Upload Key** button.

You can upload multiple signing keys that HCP can use for cryptographic verification of restored data from other HCP systems. However, the system always uses its own signing key when creating data backups.

Uploading an encryption key

To upload a signing key:

1. In the **Keys** section, click on **Encryption Keys**.
2. Click on the **Upload Key** button.
3. In the **Upload Encryption Key** window that opens:
 - In the **Secret Key Password** field, type the password for the key.
 - Click on the **Browse** button for the **Key File** field and select the key file you want.
 - To use the uploaded key instead of the original HCP encryption key to encrypt HOPs in subsequent backup operations, select the **Use for encryption** option.



Note: The list of encryption keys does not indicate which key HCP will use to encrypt namespace backups. The one that will be used is either the only one listed or the last one uploaded with the **Use for encryption** option selected. When you upload an encryption key with this option selected, you need to keep track of which key it is.

4. Click on the **Upload Key** button.

Deleting a signing key

To delete a signing key, in the **Signing Keys** list, click on the delete control () for that key.



Note: You cannot delete or replace the signing key HCP generated when it first started.

Using third-party applications with NDMP

Third-party applications differ in their implementation of backup and restore operations, but they all require certain information:

- The source of the data to be backed up or target of the restore operation. Typically, this is the HCP domain name.

To run multiple simultaneous backup streams, you need to segment the namespace by directory trees.

- The portion of data to be backed up or restored.
- The TCP port on which the source or target server listens for communications from the backup and restore application. The default TCP port for the NDMP protocol is 10000.
- For applications that require authentication, the username and password specified in the NDMP configuration.

For instructions on using a third-party application to back up or restore the data in the namespace, see the applicable product documentation.

Managing search and indexing

You manage search and indexing for the default namespace at both the tenant and namespace level. At the tenant level, you create content classes and content properties. At the namespace level, you enable search and indexing options.

This chapter contains:

- An overview of search and indexing
- An explanation of content classes and content properties and instructions for working with them
- Instructions for managing search and indexing for the default namespace



Note: The discussion of content properties in this chapter assumes a basic understanding of XML.

About search and indexing

For the default namespace to be searchable through either the metadata query API or the Search Console, it must be search enabled, and its effective permission mask must include the read and search permissions. Additionally, to get results from object-based queries through the metadata query API and from searches through the Search Console with any search facility, the namespace must be indexed by the applicable search facility.

Metadata query engine indexing

The metadata query engine indexes the default namespace when both search and indexing are enabled. The metadata query engine index is based on system metadata and, optionally, custom metadata that is well-formed XML. When you enable search for the namespace, indexing is enabled by default.

You can enable or disable indexing for the namespace at any time while search is enabled for the namespace. Disabling indexing prevents the metadata query engine from updating its index with new objects and metadata changes. When indexing is reenabled, the metadata query engine updates its index with the backlogged objects and changes and continues indexing from there.

You can enable or disable indexing of custom metadata for the namespace at any time while indexing is enabled for the namespace. Because indexing custom metadata can significantly increase the size of the indexes, you should enable it only if users need to perform searches based on custom metadata.

Indexing can be disabled for the metadata query engine at the HCP system level. If this indexing is disabled at the system level, enabling it at the namespace level has no effect.

Similarly, custom metadata indexing can be disabled for the metadata query engine at the HCP system level. If this indexing is disabled at the system level, enabling it at the namespace level has no effect.

You can reduce the amount of custom metadata that the metadata query engine indexes by creating content properties. Content properties not only decrease the size of the index but also enable users to query for objects more easily and intuitively. For information on content properties, see [“Content classes and content properties”](#) on page 123.

Content properties affect only metadata query engine indexing. They have no effect on HDDS indexing.

For an introduction to the metadata query engine, see [“HCP Search Console”](#) on page 6.

HDDS search facility indexing

The HDDS search facility indexes objects in the default namespace only while all of these are true:

- Search is enabled for the namespace.
- HTTP is enabled for the namespace.
- The effective permission mask for the namespace includes the read and search permissions.
- The HCP management API is enabled at both the system level and tenant levels.
- The HCP system uses DNS for system addressing.
- The namespace is known to HDDS.

For an introduction to the HDDS search facility, see [“HCP Search Console”](#) on page 6.

Disabling search

You can disable and reenable search for the default namespace at any time. When you disable search, indexing is automatically disabled for both the metadata query engine and the HDDS search facility.

Disabling search also removes objects in the namespace from the metadata query engine index. If you subsequently reenable search for the namespace, the namespace must be completely reindexed. The amount of time required to rebuild the indexes depends on the amount of data in the namespace. With a very large amount of data, this process can take several days.

Content classes and content properties

A **content class** is a named construct that is used to characterize objects in the default namespace. Content classes use object metadata to impose structure on unstructured namespace content. They do this through content properties.

A **content property** is a named construct used to extract an element or attribute value from custom metadata that's well-formed XML. Content properties use XPath expressions to identify the metadata of interest. When content properties are indexed, users can use them to find unstructured content that matches structured patterns.

For example, consider the following XML structure that could occur in the custom metadata for multiple objects if the namespace contains medical data:

```
<doctor>
  <name>doctor-name</name>
</doctor>
<patient>
  <name>patient-name</name>
</patient>
```

The information of interest in this custom metadata consists of the doctor's name and the patient's name.

Based on the metadata structure above, you could create content properties named `Doctor_Name` and `Patient_Name` that extract the doctor's name and patient's name from the custom metadata XML for each object. The metadata query engine could then index objects with this metadata structure by those property values. Using the metadata query API or the Metadata Query Engine Console, users could query for objects that have `Doctor_Name` or `Patient_Name` equal to a specific value.

Content properties belong to content classes. Both content classes and content properties are defined at the tenant level. Content classes are optionally associated with the default namespace. Through this association, content properties are associated with the namespace.

Metadata query engine indexing of custom metadata

By default, when custom metadata indexing is enabled for the default namespace, the metadata query engine indexes the content properties for the namespace and not the full text of custom metadata. If the namespace doesn't have any content properties, no custom metadata is indexed.

You can choose to have the metadata query engine index the full text of custom metadata. If you enable this option, the metadata query engine indexes both content properties, if any exist, and the full text of custom metadata.

With content properties, the metadata query engine indexes only the values that you determine are of interest. When indexing the full text of custom metadata, the metadata query engine indexes each word individually.

For example, suppose an object has this XML in its custom metadata:

```
<doctor>
  <name>Lee Green</name>
</doctor>
<patient>
  <name>Paris Black</name>
</patient>
```

If you've defined the `Doctor_Name` and `Patient_Name` properties, the metadata query engine index includes:

```
Lee Green
Paris Black
```

If full text indexing is enabled, the metadata query engine index includes:

```
doctor name Lee Green name doctor patient name Paris Black name patient
```

In this case, to use the metadata query API to find the objects that have a doctor named Lee Green, users would need to query for custom metadata containing "doctor.name.Lee Green.name.doctor". This kind of query can become very complex when elements are nested to deeper levels or when they have attributes.

Content class and content property workflow

Here's the basic procedure for working with content classes and content properties:

1. Create one or more content classes for the tenant. Give each class a meaningful name. For example, if the class will contain object properties that pertain to medical images, you could name the class DICOM. (DICOM is a standard for managing medical images.)

A tenant can have at most 25 content classes.

For more information, see ["Creating a content class"](#) on page 144.

2. Create content properties for each content class. Create only the content properties that will be useful to metadata query API and Search Console users. Creating content properties that won't be used unnecessarily increases the size of the metadata query engine index.

A content class can have at most 100 content properties.

For more information, see [“Content property definitions”](#) below and [“Managing content properties for a content class”](#) on page 144.

3. If custom metadata indexing isn't already enabled for the default namespace, enable it. For more information, see [“Setting search and indexing options”](#) on page 152.
4. Associate the default namespace with the applicable content classes. For clarity, associate the namespace with a content class only if the namespace contains objects that can be characterized by the content properties in the content class.

A namespace can be associated with any number of content classes.

For more information, see [“Changing associations between the default namespace and content classes”](#) on page 148.

5. Optionally, reindex the namespace. You would reindex the namespace if you want objects that were already in the namespace to be indexed by the new content properties.

You can reindex the namespace starting from the time it was created or starting from a specific date and time. When reindexing the namespace, the metadata query engine reindexes all objects with a change time that's equal to or later than the time you specify.



Tip: Because reindexing can take a long time, before reindexing the namespace:

- Create all the content properties you want for the namespace
 - Associate all the content classes containing those properties with the namespace
-

For more information, see [“Reindexing the default namespace from the Search page”](#) on page 149 and [“Reindexing the default namespace from the Search panel”](#) on page 153.

Content property definitions

The definition of a content property consists of:

- A name for the property.
- The XPath expression that identifies the property values.
- The data type of the property values.
- For numeric and datetime data types, the format of the property values.
- An indication of whether the property is single-valued or multivalued. A multivalued property can have multiple values for any given object.

The examples of content property definitions in the following sections are based on this sample custom metadata XML:

```
<?xml version="1.0" ?>
<dicom_image>
  <image type="MRI">
    <date>09/27/2012</date>
    <technician>Morgan Grey</technician>
  </image>
  <doctor>
    <name>Lee Green</name>
    <office>ABC Oncology</office>
    <address>
      <address1>Anytown Medical Building</address1>
      <address2>1 Main Street</address2>
      <city>Anytown</city>
      <state>MA</state>
      <zip>02000</zip>
    </address>
    <specialties>
      <specialty primary="true">Oncology</specialty>
      <specialty>Internal Medicine</specialty>
    </specialties>
  </doctor>
</patient>
  <id>243789</id>
  <name>Paris Black</name>
  <address>
    <address1>10 Elm Street</address1>
    <address2/>
    <city>Anytown</city>
```

```
        <state>MA</state>
        <zip>02000</zip>
    </address>
</patient>
<followup_needed>true</followup_needed>
</dicom_image>
```

Content property names

When you define a content property, you specify a name for it. Content property names must be from one through 25 characters long, can contain only alphanumeric characters and underscores (`_`), and are case sensitive. White space is not allowed.

Content property names should be intuitive for users of the metadata query API and Metadata Query Engine Console. For example, for the property that extracts the name of the doctor from the sample custom metadata, you should use a name like `Doctor_Name` rather than a name like `dname`.

Content properties with the same name

You can use the same name for multiple content properties as long as those properties have the same data type. For example, suppose the custom metadata for some objects includes a **physician** element instead of a **doctor** element, like this:

```
<physician>
  <name>Lee Green</name>
  <office>ABC Oncology</office>
  <address>
    <address1>Anytown Medical Building</address1>
    <address2>1 Main Street</address2>
    <city>Anytown</city>
    <state>MA</state>
    <zip>02000</zip>
  </address>
  <specialties>
    <specialty primary="true">Oncology</specialty>
    <specialty>Internal Medicine</specialty>
  </specialties>
</physician>
```

You could define two could define two content properties named `Doctor_Name`, one with an XPath expression that includes the **doctor** element, the other with an XPath expression that includes the **physician** element.

Within a content class, content properties with the same name must have the same data type. For information on data types, see [“Content property data types”](#) on page 131.

Reserved words

The following words are reserved and cannot be used as content property names:

- accessTime
- accessTimeString
- acl
- changeTimeMilliseconds
- changeTimeString
- customMetadata
- customMetadataAnnotation
- dpl
- gid
- hash
- hashScheme
- hold
- index
- ingestTime
- ingestTimeString
- namespace
- objectPath
- operation
- owner
- permission
- replicated
- retention
- retentionClass
- retentionString
- shred
- size
- type
- uid
- urlName
- updateTime
- updateTimeString
- utf8Name
- version

Content property expressions

For each content property you create, you specify an XPath expression. An XPath expression is an instruction for navigating an XML document to find an element or attribute value.

XPath expressions use the XPath language. HCP supports the full syntax of this language. The examples in this section illustrate only a small part of the XPath syntax.

You can learn more about XPath expressions at:

<http://www.w3schools.com/xpath>

XPath expressions that find element values

Here's a simple XPath expression that finds the value of the **followup_needed** element:

```
/dicom_image/followup_needed
```

The forward slash (/) at the beginning of the expression means that the first element is the root element in the XML. The element after the second forward slash is a child of the root element.

Here's another simple XPath expression:

```
//name
```

This expression is probably not very useful. The double slash at the beginning means find the value of any **name** element, regardless of whether that element is a child of the **doctor** element or the **patient** element.

A more useful XPath expression specifies a path to the **name** element:

```
/dicom_image/doctor/name
```

This expression means start at the root element, find the **doctor** element that's the child of the root element, and then find the **name** element that's the child of the **doctor** element. A content property with this expression finds only the name of the doctor, not the name of the patient.

A different content property with this Xpath expression finds only the patient's name:

```
/dicom_image/patient/name
```

The element path in an XPath expression can go deeper than the three levels shown above. Here's an XPath expression that's four levels deep and finds the city in which the doctor's office is located:

```
/dicom_image/doctor/address/city
```


XPath expressions that find attribute values

To find the value of an attribute, you include an at sign (@) followed by the attribute name at the end of the XPath expression. For example, here's an XPath expression that finds the value of the **type** attribute of the **image** element:

```
/dicom_image/image@type
```

Complex XPath expressions

XPath expressions can be much more complex than the ones shown so far. For example, an XPath expression can navigate XML based on the values of elements and attributes. Here's an expression that finds the name of a doctor whose primary specialty is oncology:

```
/dicom_image/doctor/specialties/specialty[@primary='true' and text()='Oncology']/
ancestor::doctor/name
```

This expression navigates down from the **doctor** element to the **specialty** elements and finds the one that has both a value of Oncology and a **primary** attribute with a value of true. The expression then navigates back up to the same **doctor** element and from there down to the **name** element that's the child of the **doctor** element.

Content property data types

Each content property has a data type that determines how the property values are treated by the metadata query engine. The possible data types are:

- **String** — The metadata query engine indexes the value as a text string. The value is handled as a single unit, even if it contains white space. Users cannot base queries on individual terms within a string value.
- **Tokenized** — The metadata query engine indexes the value as a text string after breaking it into tokens. A token is a string of either alphabetic or numeric characters. For example, the value *SSN12345789* becomes this string of two tokens: *ssn 123456789*. Tokens are not case sensitive.

The metadata query engine treats white space and special characters as token separators. For example, the value *12A Elm Street, apt. 2D* becomes this string of seven tokens: *12 a elm street apt 2 d*.

Users can base queries on any individual token or sequence of tokens within a tokenized string.

- **Boolean** — The metadata query engine indexes the value as true or false. Values that start with *1*, *t*, or *T* are treated as true. Any other values are treated as false.
- **Integer** — The metadata query engine indexes the value as an integer. Users can base queries on comparative numeric values.

The metadata query engine indexes values for a content property with a data type of integer only if the values conform to the format for the property. For more information, see [“Formats for the integer and float data types”](#) below.

- **Float** — The metadata query engine indexes the value as a decimal number with or without an exponent, depending on the value. Users can base queries on comparative numeric values.

The metadata query engine indexes values for a content property with a data type of integer only if the values conform to the format for the property. For more information, see [“Datetime data type formats”](#) on page 135.

- **Datetime** — The metadata query engine indexes the value as a date and time. Users can base queries on comparative datetime values.

The metadata query engine indexes values for a content property with a data type of date only if the values conform to the format for the property. For more information, see [“Datetime data type formats”](#) on page 135.

Formats for the integer and float data types

For a content property with the integer or float data type, you can specify a format that values needs to match in order to be indexed. The following sections include basic information about these formats. You can find more information at:

<http://docs.oracle.com/javase/6/docs/api/java/text/DecimalFormat.html>

Integer data type formats

The basic format for a content property with the integer data type is:

optional-prefix number-pattern optional-suffix

A number pattern for the integer data type consists of any number of number signs (#), followed by any number of zeroes. Both number signs and zeroes represent any number of digits, including none. The metadata query engine does not consider the length of the number pattern when matching values.

A number pattern can include a thousands separator. With the integer data type, the metadata query engine recognizes either commas (,) or periods (.) as the thousands separator.

For example, a value of *1234* matches any of these number patterns:

```
0
000
##
###0000
0,0
##,000
```

If a content property value contains a thousands separator, the value matches only number patterns that contain the same thousands separator. For example, the value *1,234* matches the last two patterns above, but not the first four. It also does not match *0.0* or *##.000*.

The prefix or suffix in the format for the integer data type can be any character string, with a few exceptions. For example, a prefix or suffix cannot include a period (.) or percent sign (%). The format must include white space between the integer pattern and the suffix, if used.

For example, for the metadata query engine to index the value *\$1234* as an integer, the format for the content property must have a dollar sign (\$) in front of the integer pattern, with no space between them.

Here are some examples of integer formats with examples of values that match them:

| Format | Example |
|---------------|----------------|
| \$ 0,0 | \$ 1,234 |
| ###0 AD | 2012 AD |
| ~# mph | ~55 mph |

If you don't specify a format for a content property with the integer data type, the metadata query engine indexes only sequences of digits with no special characters.

Float data type formats

For the format for a content property with the float data type, you can use any of the formats for the integer data type. However, with the float data type, the thousands separator, if used, must be a comma (,).

You can include a period as a decimal separator in the number pattern for the float data type, although this is not required. If you do include it, any number signs (#) must come after any zeroes in the part following separator.

For example, a value of *1234.5* matches any of these number patterns:

```
0
00.0
.0
#0.0#
##,000
0,0
#,0.0#
```

You can also include an exponent character (E) followed by one or more zeroes in the number pattern for the float data type. However, values with an exponent character also match patterns that don't include the exponent character, and values without an exponent character also match patterns with an exponent character.

For example, a value of *1234E5* matches any of these number patterns:

```
0
00.0
.0E0
#0.0#E000
##,000E0
0,0
#,0.0E00
```

You can use a percent sign (%) by itself as the prefix or suffix in the format for the float data type. Before indexing values with a matching percent sign, the metadata query engine converts them to their decimal equivalents. For example, a value of *1234%* matches a format of *0%* and is indexed as *12.34*.

White space is not required between the number pattern and a suffix that's a percent sign.

If you don't specify a format for a float data type, the metadata query engine indexes only sequences of digits that optionally include one decimal point.

Datetime data type formats

For a content property with the datetime data type, you can specify a format that values needs to match in order to be indexed. The format consists of a pattern of letters, optional separators, and optional quoted text. The letters represent date or time components, as outlined in the table below. Letters can be repeated, which can affect their meaning.

| Letter | Description |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| G | Represents a valid era indicator, such as AD, BC, or BCE. Repetition has no effect. If a datetime pattern doesn't include any occurrences of G, the metadata query engine assumes an era of AD for matching values. |
| y | Represents a year. For matching values with a two-digit year, a pattern that includes y more than twice in a row causes the metadata query engine to interpret the two digits as being preceded by two zeroes rather than by the number that indicates the current century. If a datetime pattern doesn't include any occurrences of y, the metadata query engine assumes a year of 1970 for matching values. |
| M | Represents a month. Values that include the month as a number match a pattern that includes M or MM. Values that include the name of the month, either in full or as a three-letter abbreviation, match a pattern that includes three or more occurrences of M in a row. If a datetime pattern doesn't include any occurrences of M, the metadata query engine assumes a month of January for matching values. |
| w | Represents the number of the week into the year. Repetition has no effect. |
| W | Represents the number of the week into the month, where the first week is the week that includes the first day of the month. Repetition has no effect. |
| D | Represents the number of the day into the year. Repetition has no effect. |
| d | Represents the number of the day into the month. Repetition has no effect. |
| F | Represents the number of the week into the month, where the first week starts with the first Sunday in the month. Repetition has no effect. |
| E | Represents the day of the week. Matching values include the name of the day in full or as a three-letter abbreviation. Repetition has no effect. |
| a | Represents a valid morning or afternoon indicator, such as AM or pm. Repetition has no effect. |

(Continued)

| Letter | Description |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| H | Represents the hour on a 24-hour clock, where midnight is represented by zero. Repetition has no effect. |
| k | Represents the hour on a 24-hour clock, where midnight is represented by 24. Repetition has no effect. |
| K | Represents the hour on a 12-hour clock, where midnight and noon are represented by zero. Repetition has no effect. |
| h | Represents the hour on a 12-hour clock, where midnight and noon are represented by 12. Repetition has no effect. |
| m | Represents the minute into the hour. Repetition has no effect. If a datetime pattern doesn't include any occurrences of m, the metadata query engine assumes that the number of minutes is zero for matching values. |
| s | Represents the second into the minute. Repetition has no effect. If a datetime pattern doesn't include any occurrences of s, the metadata query engine assumes that the number of seconds is zero for matching values. |
| S | Represents a number of milliseconds past the applicable second. Repetition has no effect. |
| z | Represents a valid time zone specified as text, such Eastern Standard Time, EDT, or GMT. Repetition has no effect. |
| Z | Represents a valid time zone specified as an offset from GMT, formatted as (+ -)nnnn, such as +0500 or -0200. Repetition has no effect. |

If a datetime format doesn't include a representation for:

- A day, the metadata query engine assumes that the day is the first day of the applicable month for matching values
- An hour, the metadata query engine assumes that the hour is midnight
- A time zone, the metadata query engine assumes that the time is in the HCP system time zone

The separators in a datetime format can be any of several different special characters, including forward slashes (/), hyphens (-), colons (:), semicolons (;), at signs (@), and spaces.

To include text in a datetime format, enclose the text in single quotation marks ('). To include a single quotation mark, specify two single quotation marks in a row.

Here are some examples of datetime formats with examples of values that match them:

| Format | Example |
|----------------------------|-------------------------------------|
| MM/dd/yy HH:mm:ss z | 03/19/12 14:35:27 EST |
| hh 'o'clock' a, zzz | 2 o'clock PM, Eastern Standard Time |
| yyyy-MM-dd'T'HH:mm:ss.SSSZ | 2012-03-19T14:35:27.236-0400 |
| E., MMM d, yyyy 'at' k:s | Mon., March 19, 2012 at 14:35 |

If you don't specify a format for a content property with the datetime data type, the metadata query engine indexes values that match patterns such as MM/dd/yyyy, MM-dd-yyyy, yyyy-MM-dd, or yyyy-MM-dd'T'HH:mm:ssZ.

You can find more information about datetime formats at:

<http://docs.oracle.com/javase/6/docs/api/java/text/SimpleDateFormat.html>

Multivalued content properties

A content property is defined as either single-valued or multivalued:

- If you define a content property as single-valued, the metadata query engine indexes only one occurrence of it for any given object, regardless of how many times it occurs in the custom metadata XML for that object.
- If you define a content property as multivalued, the metadata query engine indexes all occurrences of it in the custom metadata XML for an object.

For example, based on the sample custom metadata XML, you would define as multivalued a content property that extracts the value of the **specialty** element.

With the metadata query API, users can sort query results based on single-valued content properties but not on multivalued properties.

Content properties extracted from sample XML

When working with content properties in the Tenant Management Console, you can supply sample well-formed XML and have HCP extract content properties from that XML. You can then select which of those properties you want to add to a content class.

HCP extracts only content properties for XPath expressions that follow a straight path from the root element. These conventions apply to the content property definitions:

- The XPath expression always starts from the root element.
- The name of a content property that extracts an element value is the name of the element preceded by the name of the parent element.
- The name of a content property that extracts an attribute value is the name of the attribute preceded by the name of the element the attribute applies to.
- Content property names that would exceed 25 characters in length are truncated to 25 characters, starting from the beginning.
- The definitions do not include formats.
- The definitions are listed alphabetically by XPath expression.

When adding extracted content properties to a content class, you can change any parts of their definitions.

The table below shows the definitions of the content properties HCP extracts from the sample custom metadata XML.

| XPath expression | Name | Data Type | Multivalued |
|----------------------------------------------------|---------------------------|-----------|-------------|
| /dicom_image/doctor/address/address1 | addressAddress1 | String | No |
| /dicom_image/doctor/address/address2 | addressAddress2 | String | No |
| /dicom_image/doctor/address/city | addressCity | String | No |
| /dicom_image/doctor/address/state | addressState | String | No |
| /dicom_image/doctor/address/zip | addressZip | Integer | No |
| /dicom_image/doctor/name | doctorName | String | No |
| /dicom_image/doctor/office | doctorOffice | String | No |
| /dicom_image/doctor/specialties/specialty | specialtiesSpecialty | String | Yes |
| /dicom_image/doctor/specialties/specialty/@primary | specialtyPrimary | Boolean | No |
| /dicom_image/followup_needed | icom_imageFollowup_needed | Boolean | No |
| /dicom_image/image/@type | imageType | String | No |
| /dicom_image/image/date | imageDate | String | No |

(Continued)

| XPath expression | Name | Data Type | Multivalued |
|---------------------------------------|-----------------|-----------|-------------|
| /dicom_image/image/technician | imageTechnician | String | No |
| /dicom_image/patient/address/address1 | addressAddress1 | String | No |
| /dicom_image/patient/address/address2 | addressAddress2 | String | No |
| /dicom_image/patient/address/city | addressCity | String | No |
| /dicom_image/patient/address/state | addressState | String | No |
| /dicom_image/patient/address/zip | addressZip | Integer | No |
| /dicom_image/patient/id | patientId | Integer | No |
| /dicom_image/patient/name | patientName | String | No |

Content property files

You can export the content properties for a content class to a file that you can then use to import the properties to another class. The exported file contains XML definitions of the content properties in this format:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes" ?>
<contentClass>
  <contentProperties>
    <contentProperty>
      <name>property-name</name>
      <expression>xpath-expression</expression>
      <type>data-type</type>
      <multivalued>>true-or-false</multivalued>
      <format>format</format>
    </contentProperty>
    .
    .
    .
  </contentProperties>
</contentClass>
```

Using the same format, you can also create content property files yourself.

Here's an example of XML that defines some content properties based on the sample custom metadata XML:

```
<?xml version="1.0" encoding="UTF-8" standalone="yes" ?>
<contentClass>
  <contentProperties>
    <contentProperty>
      <name>Doctor_City</name>
      <expression>/dicom_image/doctor/address/city</expression>
      <type>STRING</type>
      <multivalued>>false</multivalued>
      <format />
    </contentProperty>
    <contentProperty>
      <name>Doctor_State</name>
      <expression>/dicom_image/doctor/address/state</expression>
      <type>STRING</type>
      <multivalued>>false</multivalued>
      <format />
    </contentProperty>
    <contentProperty>
      <name>Doctor_Name</name>
      <expression>/dicom_image/doctor/name</expression>
      <type>STRING</type>
      <multivalued>>false</multivalued>
      <format />
    </contentProperty>
    <contentProperty>
      <name>Doctor_Office</name>
      <expression>/dicom_image/doctor/office</expression>
      <type>STRING</type>
      <multivalued>>false</multivalued>
      <format />
    </contentProperty>
    <contentProperty>
      <name>Doctor_Specialty</name>
      <expression>/dicom_image/doctor/specialties/specialty</expression>
      <type>STRING</type>
      <multivalued>>true</multivalued>
      <format />
    </contentProperty>
    <contentProperty>
      <name>Followup_Needed</name>
      <expression>/dicom_image/followup_needed</expression>
      <type>BOOLEAN</type>
      <multivalued>>false</multivalued>
      <format />
    </contentProperty>
  </contentProperties>
</contentClass>
```

```

    <name>Image_Type</name>
    <expression>/dicom_image/image/@type</expression>
    <type>STRING</type>
    <multivalued>>false</multivalued>
    <format />
  </contentProperty>
  <contentProperty>
    <name>Image_Date</name>
    <expression>/dicom_image/image/date</expression>
    <type>STRING</type>
    <multivalued>>false</multivalued>
    <format />
  </contentProperty>
  <contentProperty>
    <name>Patient_City</name>
    <expression>/dicom_image/patient/address/city</expression>
    <type>STRING</type>
    <multivalued>>false</multivalued>
    <format />
  </contentProperty>
  <contentProperty>
    <name>Patient_State</name>
    <expression>/dicom_image/patient/address/state</expression>
    <type>STRING</type>
    <multivalued>>false</multivalued>
    <format />
  </contentProperty>
  <contentProperty>
    <name>Patient_ID</name>
    <expression>/dicom_image/patient/id</expression>
    <type>INTEGER</type>
    <multivalued>>false</multivalued>
    <format />
  </contentProperty>
  <contentProperty>
    <name>Patient_Name</name>
    <expression>/dicom_image/patient/name</expression>
    <type>STRING</type>
    <multivalued>>false</multivalued>
    <format />
  </contentProperty>
</contentProperties>
</contentClass>

```

About the Search page

To manage search and indexing for the default tenant, including viewing, creating, and managing content classes and content properties, you use the **Search** page in the Tenant Management Console. To display this page:

1. In the top-level menu, mouse over **Services** to display a secondary menu.
2. In the secondary menu, click on **Search**.



Roles: To view existing content classes and content properties, you need the monitor or administrator role. To create, modify, and delete content classes and content properties and reindex the default namespace, you need the administrator role.

Managing the content class list

The **Search** page lists existing content classes. For each content class, the list shows the content class name.

By default, the content class list includes all existing content classes. The content classes are listed 20 at a time in ascending order by name.

You can page through, sort, and filter the list of content classes. The **Search** page indicates which content classes are shown out of the total number of content classes in the current list.

Paging

You can change the number of content classes shown at a time on the **Search** page. To do this, in the **Items per page** field, select the number of content classes you want. The options are 10, 20, and 50.

To page forward or backward through the content class list, click on the next () or back () control, respectively.

To jump to a specific page in the content class list:

1. In the **Page** field, type the page number you want.
2. Press Enter.


Sorting


You can sort the content class list in ascending or descending order by content class name. To change the sort order, click on the **Name** column heading. Each time you click on the column heading, the sort order switches between ascending and descending.

Filtering

You can filter the content class list by content class name. The filtered list includes only those content classes with a name that begins with or is the same as a specified text string.

To filter the content class list:


1. In the entry field above the list, type the text string you want to use as a filter. This string can be up to 64 characters long and can contain any valid UTF-8 characters, including white space. It is not case sensitive.
2. Click on the find control ().

To redisplay the entire list of storage pools after filtering it, click on the clear filter control ().

Understanding the content property list for a content class

To view the content properties defined for a content class, click on the content class name in the content class list. The panel that opens (the **Settings** panel) contains a list of the content properties in that content class. The properties are listed in alphabetical order by name.

For each content property, the list shows:

- **Name** — The content property name
- **Expression** — The XPath expression for the content property, with an annotation prefix if applicable
- **Type** — The data type of the content property
- **Format** — The format for the content property
-  — An indication of whether the content property is multivalued

Creating a content class

When you create a content class, you can create content properties for it at the same time. Alternatively, you can create a content class with no properties and add properties to it later.

To create a content class, on the **Search** page:

1. Click on **Create Content Class**.
2. In the **Name** field, type a name for the content class. Content class names must be from one through 64 characters long, can contain any valid UTF-8 characters, including white space, and are not case sensitive.
3. Optionally, define one or more content properties to be added to the content class. For instructions on doing this, see [“Adding, modifying, and deleting content properties”](#) below.
4. If you defined content properties in [step 2](#) above, optionally test or export them. For instructions for these actions, see [“Testing content properties”](#) on page 147 and [“Exporting content properties”](#) on page 147.
5. Click on the **Create Content Class** button.

Managing content properties for a content class

You can add, modify, and delete content properties for a content class at any time. You can also test the properties against XML that you supply. Additionally, you can export the properties to a content property file.

Adding, modifying, and deleting content properties

To add, modify, or delete content properties for a content class, on the **Search** page:


1. In the list of content classes in the **Settings** panel, click on the name of the content class for which you want to add, modify, or delete content properties.

2. Do one or more of these:

- To add content properties, do one or more of these:
 - Add one or more properties individually (see [“Adding content properties individually”](#) below)
 - Extract one or more properties from XML that you supply (see [“Extracting content properties from sample XML”](#) on page 146)
 - Import properties from a content property file (see [“Importing content properties from a content property file”](#) on page 146)

A row for each new content property appears in the list of content properties for the content class. The row is highlighted in green.

To remove a new row, click on the delete control () for the row.

- To modify a content property, in the row for the property, make the changes you want.
- To delete an existing content property, click on the delete control () for the row containing the property.

The row turns red. To undo the deletion, click again on the delete control.

3. Optionally, test or export all the listed content properties. For instructions for these actions, see [“Testing content properties”](#) on page 147 and [“Exporting content properties”](#) on page 147.4. Click on the **Update Settings** button.

If you also typed a new name for the content class in the **Name** field, clicking on the **Update Settings** button changes the content class name.

Adding content properties individually

To add an individual content property to the content property list for a content class:

1. Above the content property list, click on **Add**.
2. In the new row that appears in the content properties list, fill in the definition for the new content property.

Extracting content properties from sample XML

To extract content properties from sample XML and add them to the content property list for a content class:

1. In the **Sample Custom Metadata** field, type or paste the well-formed XML from which you want to extract properties.
2. Above the content property list, click on **Extract**.

The **Extract Content Properties** window opens. The window lists all the content properties HCP was able to extract from the sample XML.

3. Select each content property you want to add to the content class.

To select all the listed content properties, click in the checkbox at the top of the list. To deselect all the properties after selecting all, click in the checkbox again.

4. Click on the **Add Selected** button.

A row for each content property you selected appears in the list of content properties for the content class.

5. Optionally, modify the definitions of the new content properties.

Importing content properties from a content property file

To import content properties from a content property file and add them to the content property list for a content class:

1. Above the content property list, click on **Import**.
2. In the **Import Content Properties** window, click on the **Browse** button. Then select the content property file you want.
3. Click on the **Import** button.

A row for each content property defined in the content property file appears in the list of content properties for the content class.

4. Optionally, modify the definitions of the new content properties.

For more information on content property files, see [“Content property files”](#) on page 139.

Testing content properties

You can test the definitions of the content properties for a content class at any time. You do this by having HCP extract content property values from sample XML that you supply. The test applies to all the content properties in the content property list, including those that have not yet been committed.

To test the definitions of the content properties in the content property list for a content class, on the **Search** page:

1. In the list of content classes, click on the name of the content class with the content properties you want to test.
2. In the **Sample Custom Metadata** field in the **Settings** panel, type or paste the well-formed XML you want to use for the test.
3. Above the content property list, click on **Test**.

HCP extracts the values it can find for the content properties in the content property list. The extracted values appear in the **Test Value** column in the content property list.

Exporting content properties

You can export the definitions of the content properties for a content class to a content property file at any time. The resulting file includes the definition of each property in the list, including those that have not yet been committed.

To export the definitions of the content properties in the content property list for a content class to a content property file, on the **Search** page:

1. In the list of content classes, click on the name of the content class with the content properties you want to export.
2. Above the content property list in the **Settings** panel, click on **Export**.
3. When prompted, save the content property file to the location of your choice.

For more information on content property files, see [“Content property files”](#) on page 139.

Changing associations between the default namespace and content classes

You can change the associations between the default namespace and content classes at any time.


When you change the association between the namespace and a content class, you have the option of reindexing the namespace. If you choose to do this, the metadata query engine reindexes the namespace starting from the time the namespace was created.

To change the association between the namespace and a content class, on the **Search** page:

1. In the list of content classes, click on the name of the content class with which you want to associate namespaces.
2. In the row of tabs below the content class name, click on **Namespaces**.


In the **Namespaces** panel, the **Associate Namespaces with Content Class** section lists the default namespace if it's enabled for search and is not already associated with the content class.

If the namespace is already associated with the content class, it is listed in the **Content Class Namespaces** section.

3. Do either of these:
 - To associate the namespace with the content class, click on the add control () for the namespace in the **Associate Namespaces with Content Class** section.

The row containing the namespace turns green.

To deselect the namespace, click on the remove control () for it.

- To dissociate the namespace from the content class, click on the remove control () for the namespace in the **Content Class Namespaces** section.

The row containing the namespace turns red.

To deselect the namespace, click on the add control () for it.

4. Optionally, to reindex the namespace, select the **Reindex all objects in added/removed namespaces** option.

5. Click on the **Update Content Class** button.

Reindexing the default namespace from the Search page


You can reindex the default namespace at any time. You might do this for example, if you add a new content property to the content class.


To reindex the namespace from the **Search** page:

1. In the list of content classes, click on the name of the content class with which the namespace is associated.


2. In the row of tabs below the content class name, click on **Reindex**.

The **Content Class Namespaces** section in the **Reindex** panel is empty if the namespace is not associated with the content class.

3. Click on the add control () to select the namespace. The namespace row turns green.

To deselect a selected namespace, click on the remove control () for the namespace.

4. In the **Reindex Selected Namespaces** section, select either the **All objects** option or the **Objects modified after** option.

If you select the **Objects modified after** option, either type a date in the associated field or click on the calendar control () to select a date. If you type a date, use this format: *mm/dd/yyyy*

If you enter an invalid date using the correct date format, HCP tries to convert it to a real date. For example, if you enter 11/31/2012, HCP converts it to 12/01/2012.

5. Click on the **Reindex Namespaces** button.

If you selected the **All objects** option, a confirming message appears.

In the window with the confirming message, select **I understand** to confirm that you understand the consequences of your action. Then click on the **Reindex Objects** button.



Note: If indexing is disabled for the namespace, the above procedure by itself does not start the reindexing process. To start reindexing the namespace, you need to reenable indexing for it. For information on doing that, see [“Setting search and indexing options”](#) on page 152.

For more information on reindexing the default namespace, see [“Content class and content property workflow”](#) on page 125. For information on reindexing the namespace from the namespace **Search** panel, see [“Reindexing the default namespace from the Search panel”](#) on page 153.

Renaming a content class

You can rename a content class at any time. Renaming a content class has no effect on the associations between the default namespace and that class.

To rename a content class, on the **Search** page:


1. In the list of content classes, click on the name of the content class you want to rename.
2. In the **Name** field in the **Settings** panel, type a new name for the content class.
3. Click on the **Update Settings** button.

If you also modified the list of content properties for the content class, clicking on the **Update Settings** button commits those changes as well.

Deleting a content class

You can delete a content class at any time. When you delete a content class, you have the option of reindexing the default namespace if the namespace is associated with the class. If you choose to do this, the metadata query engine reindexes the namespace starting from the time the namespace was created.

To delete a content class, on the **Search** page:

1. In the list of content classes in the **Settings** panel, click on the delete control () for the content class you want to delete.
2. In response to the confirming message:
 - a. Optionally, to reindex the default namespace, select the **Reindex all objects in associated namespaces** option.
 - b. Click on the **Delete** button.

Managing search and indexing for the default namespace

You can change the search and indexing settings for the default namespace at any time. You can:

- Enable or disable search. Disabling search also:
 - Disables all indexing for the namespace
 - Removes the objects in the namespace from the metadata query engine index
 - Deletes all associations the namespace has with any content classes
- If search is enabled:
 - Enable or disable indexing. Disabling indexing also disables custom metadata indexing for the namespace. However, it does not remove objects in the namespace from the indexes.
 - Add or delete associations between content classes and the namespace.
- If indexing is enabled, enable or disable indexing of custom metadata. Disabling custom metadata indexing also disables metadata query engine indexing of the full text of custom metadata.

If you disable custom metadata indexing after it has been enabled, custom metadata that has already been indexed is not removed from the index.

- If indexing of custom metadata is enabled, enable or disable metadata query engine indexing of the full text of custom metadata. If you disable this option after it has been enabled, custom metadata text that has already been indexed is not removed from the index.



Note: If the default namespace existed before the HCP system was upgraded from a release earlier than 6.0 and had custom metadata indexing enabled before the upgrade, it has full-text custom metadata indexing enabled after the upgrade.

You can also reindex the namespace at any time. You might do this, for example, if you associate additional content classes with it.

To manage search and indexing for the namespace, you use the namespace **Search** panel in the Tenant Management Console. To display this panel:

1. In the top-level menu, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Services**.
3. On the left side of the **Services** panel, click on **Search**.

While the metadata query engine is selected for use with the Search Console, the **Search** panel shows the date and time before which eligible objects in the namespace are guaranteed to be included in the applicable index.



Roles: To view the search and indexing settings for the namespace, you need the monitor or administrator role. To modify the search and indexing settings for the namespace or to reindex the namespace, you need the administrator role.

Setting search and indexing options

To change the search and indexing options for the default namespace, in the **Search** panel for the namespace:

1. Optionally, select or deselect the **Enable search** option to enable or disable search, respectively.
2. Optionally, if the **Enable search** option is selected, select or deselect the **Enable indexing** option to enable or disable indexing, respectively.
3. Optionally, if the **Enable indexing** option is selected, select or deselect the **Enable indexing of custom metadata** option to enable or disable indexing of custom metadata, respectively.
4. Optionally, if the **Enable indexing of custom metadata** option is enabled:
 - Optionally, select or deselect the **Enable full custom metadata indexing** option to enable or disable metadata query engine indexing of the full text of custom metadata, respectively.
 - Optionally, in the **Content Classes** list, select or deselect content classes to associate them with or dissociate them from the namespace, respectively.

To select all the listed content classes, click in the checkbox at the top of the list. To deselect all the content classes after selecting all, click in the checkbox again.

5. Click on the **Update Settings** button.


If you deselected the **Enable search** option, HCP displays a confirming message.

In the window with the confirming message, select **I understand** to confirm that you understand the consequences of your action. Then click on the **Update Settings** button.

Reindexing the default namespace from the Search panel

To reindex the default namespace from the **Search** panel:

1. In the **Reindex** section of the **Search** panel, select either the **All objects** option or the **Objects modified after** option.

If you select the **Objects modified after** option, either type a date in the associated field or click on the calendar control () to select a date. If you type a date, use this format: *mm/dd/yyyy*

If you enter an invalid date using the correct date format, HCP tries to convert it to a real date. For example, if you enter 11/31/2015, HCP converts it to 12/01/2015.

2. Click on the **Reindex Objects** button.

If you selected the **All objects** option, a confirming message appears.

In the window with the confirming message, select **I understand** to confirm that you understand the consequences of your action. Then click on the **Reindex Objects** button.



Note: If indexing is disabled for the namespace, the procedure above does not by itself start the reindexing process. To start reindexing the namespace, you need to reenable indexing. For information on doing that, see ["Setting search and indexing options"](#) above.

For more information on reindexing namespaces, see ["Content class and content property workflow"](#) on page 125. For information on reindexing the namespace from the namespace **Search** page, see ["Reindexing the default namespace from the Search page"](#) on page 149.



Working with retention classes

Retention classes provide a means to consistently manage data that must remain in the namespace for a specific amount of time. For example, if local law requires that medical records be kept for a specified number of years, you can use a retention class to enforce that requirement.

Retention classes are defined on a per-namespace basis. The classes you create for the default namespace are not visible in any other namespace.

You create retention classes in the Tenant Management Console. Users and applications can then use those classes as retention settings for objects.

This chapter describes retention classes and explains how to create, modify, and delete them.

For information on how users and applications use retention classes, see *Using the Default Namespace*.

About retention classes

A **retention class** is a named value that, when used as the retention setting for an object, specifies how long the object must remain in the repository. This value can be:

- An offset from the time the object is created. You specify an offset as numbers of years, months, and/or days.

For example, you could create a retention class named HlthReg-107 with an offset of 21 years. Then, all objects assigned HlthReg-107 as their retention setting could not be deleted for 21 years after they're created.

- One of these special values:
 - **Deletion Allowed** — The object can be deleted at any time.
 - **Deletion Prohibited** — The object can never be deleted by means of a normal delete operation. If the namespace is in enterprise mode, however, the object can be deleted by means of a privileged delete operation.
 - **Initial Unspecified** — The retention period for the object is unspecified. The object cannot be deleted by means of a normal delete operation while it has this retention setting. If the namespace is in enterprise mode, however, the object can be deleted by means of a privileged delete operation.

The retention period for an object assigned to a retention class is calculated from the value of that class. So, for example, if an object stored on October 8, 2011, is assigned to a retention class with an offset value of seven years, the retention period for that object expires on October 8, 2018.

You can choose to have HCP automatically delete the objects assigned to a retention class when they expire. This applies only to retention classes with a value that's an offset. It does not apply to retention classes with special values.

Automatic deletion of expired objects in retention classes occurs only if disposition is enabled. For information on disposition, see ["Disposition"](#) on page 14.

To view, create, and manage retention classes for a namespace, you use the **Retention Classes** panel for that namespace. To display this panel:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Compliance**.
3. On the left side of the **Compliance** page, click on **Retention Classes**.



Roles: To view retention classes, you need the monitor, administrator, or compliance role. To create, modify, and delete retention classes, you need the compliance role.

Understanding the retention class list

The **Retention Classes** panel lists the retention classes defined for a namespace. For each class, the list shows:

- The retention class name.
- Whether the retention class value is an offset or a special value.
- The retention class value.

A value that's an offset has this format:

$$A+n\gamma+rm+nd$$

In this format, y is the number of years, m is the number of months, and d is the number of days. Only the measurement units with nonzero values are shown. For example, an offset of two years and 5 days looks like this:

$$A+2y+5d$$

- Whether expired objects in the retention class are automatically deleted (**Allow Disposition**).

To view the description of a listed retention class, click on the class name.

Creating a retention class

To create a retention class in the default namespace, in the **Retention Classes** panel:

1. Click on **Create Retention Class**.
2. In the **Create Retention Class** panel:
 - In the **Retention class name** field, type a name for the retention class. Retention class names must be from one through 64 characters long, can contain only alphanumeric characters, hyphens (-), and underscores (_), and are not case sensitive.
 - Do one of these:
 - To make the retention class value an offset, in the **Retention Method** field, select **Offset**. Then do one or more of:
 - In the **Years** field, type a number of years. Valid values are integers in the range zero through 9,999.
 - In the **Months** field, type a number of months. Valid values are integers in the range zero through 9,999.
 - In the **Days** field, type a number of days. Valid values are integers in the range zero through 9,999.
 - To make the retention class value a special value:
 1. In the **Retention method** field, select **Special Value**.
 2. In the field below the **Retention method** field, select the special value you want.
 - Optionally, in the **Description** field, type a description of the retention class. The description can be up to 1,024 characters long and can contain any valid UTF-8 characters, including white space.
 - Optionally, select **Allow Disposition service to delete objects when expired** to have HCP automatically delete expired objects in the retention class.
3. Click on the **Create Retention Class** button.

Modifying a retention class

You can increase the value of a retention class at any time. These changes increase the value:

- From an offset to a larger offset
- From **Deletion Allowed** to an offset or **Deletion Prohibited**


If the namespace is in enterprise mode, you can decrease the value of a retention class at any time. These changes decrease the value:

- From an offset to a smaller offset
- From an offset to **Deletion Allowed**
- From **Deletion Prohibited** to an offset or **Deletion Allowed**

If the value of a retention class is **Initial Unspecified**, you can change it to any other value at any time. You can change the value to **Initial Unspecified** only from **Deletion Allowed**.

You can enable or disable the automatic deletion feature at any time.

To modify an existing retention class, in the **Retention Classes** panel:


1. In the list of retention classes, click on the edit control () for the retention class you want to modify.
2. In the **Edit Retention Class** window, make the changes you want. For information on the fields and options in this panel, see ["Creating a retention class"](#) above.
3. Click on the **Update Settings** button.

Deleting a retention class

You can delete a retention class only if the namespace is in enterprise mode. You cannot delete a retention class if the namespace is in compliance mode.

When you delete a retention class, the retention setting of each object in the class changes to **Deletion Prohibited**.

To delete a retention class, in the **Retention Classes** panel:

1. In the list of retention classes, click on the delete control () for the retention class you want to delete.
2. In response to the confirming message, click on the **Delete Retention Class** button.

Using privileged delete

When created in enterprise mode, the default namespace supports privileged delete operations. If you have the compliance role through your HCP user account or group accounts, you can perform these operations through the Tenant Management Console and the HCP Search Console. The namespace access protocols and HCP Data Migrator do not support privileged delete operations in the default namespace.

This chapter describes the privileged delete feature and explains how to use the Tenant Management Console to perform privileged delete operations.

For information on performing privileged delete operations through the Search Console, see *Searching Namespaces*.

About privileged delete

Privileged delete is an HCP feature that enables you to delete objects even if they are under retention. This feature is available only if the namespace is in enterprise mode. If the namespace is in compliance mode, you *cannot* delete objects that are under retention.

Privileged delete supports government regulations that require the destruction of certain types of data in response to changing circumstances. For example, companies may be required to destroy particular information about employees who leave. If that data is under retention, it cannot be deleted through normal delete operations.

When using privileged delete to delete an object, you need to specify a reason for the deletion. The tenant log records all privileged delete operations, including the specified reasons, thereby creating an audit trail.

Using privileged delete, you can also delete objects that are not under retention. You would do this, for example, if you wanted to record the reason for an object deletion.

You cannot use privileged delete to delete objects that are on hold, regardless of their retention settings.



Roles: To perform a privileged delete operation, you need the compliance role.

Object specification

With privileged delete, you can delete only one object at a time. To specify the object, you need to include the full path to it in the namespace (starting after `fcfs_data`). The path must begin with a forward slash (/).

For example, to delete the `Lee_Green_1254` object from the `Corporate/Employees` directory, you would specify:

```
/Corporate/Employees/Lee_Green_1254
```

Directory and object names are case sensitive. The separator is the forward slash (/).

Non-UTF-8 characters in directory and object names must be percent encoded. To avoid ambiguity, you should also percent-encode the characters listed in the table below.

| Character | Percent-encoded value |
|-----------------|-----------------------|
| Space | %20 |
| Tab | %09 |
| New line | %0A |
| Carriage return | %0D |
| + | %2B |
| % | %25 |
| # | %23 |
| ? | %3F |
| & | %26 |

Percent-encoded values are not case sensitive.

Performing a privileged delete

To use privileged delete to delete an object:

1. In the top-level menu in the Tenant Management Console, mouse over **Default Namespace** to display a secondary menu.
2. In the secondary menu, click on **Compliance**.
3. On the left side of the **Compliance** page click on **Privileged Delete**.



Note: This option is present only if the namespace is in enterprise mode and your user account includes the compliance role.

4. In the **Privileged Delete** panel:
 - In the **Object to Delete** field, type the path to and name of the object you want to delete. For information on identifying objects for deletion, see [“Object specification”](#) above.
 - In the **Reason for Deletion** field, type the reason why you’re deleting the object. This text must be from one through 1,024 characters long and can contain any valid UTF-8 characters, including white space.

5. Click on the **Delete This Object** button.
6. In response to the confirming message, click on the **Delete Object** button.

Downloading HCP Data Migrator

HCP Data Migrator runs on both Windows and Unix clients. You download the applicable HCP-DM installation file from the Tenant Management Console. That file is then used to install HCP-DM on the client computers.

This chapter describes the system requirements for running HCP-DM and contains instructions for downloading the applicable installation file.

For an introduction to HCP-DM, see [“HCP Data Migrator”](#) on page 8. For information on installing and using HCP-DM, see *Using HCP Data Migrator*.

HCP-DM system requirements

HCP-DM runs on any Windows or Unix client that supports the Oracle Java Runtime Environment (JRE) version 7 update 6 or later. The computer that runs HCP-DM must meet these minimum requirements:

- 1.6 Ghz processor
- 2 Gb RAM
- 100 Mbps Ethernet interface

Windows clients should have the most recent applicable Microsoft Windows Service Pack installed.

Downloading the HCP-DM installation file


The HCP-DM installation file is:

- For Windows, either `hcpdm.exe` or `hcpdm.zip`
- For Unix, `hcpdm.tgz`



Roles: You can download an HCP-DM installation file while logged into the Tenant Management Console with any user account.

To download the HCP-DM installation file you want to use, in the Tenant Management Console:

1. In the top right corner of the Console window, mouse over the tools icon () to open a dropdown menu.
2. In the dropdown menu, click on the option for the HCP-DM installation file you want:
 - **HCP-DM (Windows Installer)** for `hcpdm.exe`
 - **HCP-DM (zip)** for `hcpdm.zip`
 - **HCP-DM (tgz)** for `hcpdm.tgz`
3. Save the installation file in the location of your choice.








Tenant Management Console alerts

The Tenant Management Console uses icons to report tenant and namespace status on the tenant and default namespace **Overview** pages and on the **Search** page. These icons, called **alerts**, are accompanied by text. Each alert also has text that's displayed when you mouse over the alert.






This appendix describes the alerts that can appear on these Console pages and shows the mouse-over text for each alert. The appendix also tells you how to respond to the alerts.

The alerts in this appendix are listed alphabetically by their mouse-over text.


Tenant Overview page alerts

| Icon | Mouse-over text | Description |
|-------------------------------------------------------------------------------------|-------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Default namespace may be under-replicated | <p>The default namespace has a DPL of 1 (one) and at least one of the directories in it is not being replicated. When the default namespace has a DPL of 1, all directories in it should be replicated to ensure that the stored data is protected. If replication is not available, increase the namespace DPL.</p> <p>This alert appears only in HCP systems in which setting the DPL to one for the namespace can leave objects unprotected.</p> |
|  | Irreparable and unavailable objects | <p>The default namespace contains irreparable objects and unavailable objects. Please contact your HCP system administrator.</p> <p>To see which objects are irreparable, go to the Irreparable Objects panel. For information on this panel, see “Working with irreparable objects” on page 83.</p> |
|  | Irreparable objects | <p>The default namespace contains irreparable objects. Please contact your HCP system administrator.</p> <p>To see which objects are irreparable, go to the Irreparable Objects panel. For information on this panel, see “Working with irreparable objects” on page 83.</p> |
|  | Tenant with DPL 1 namespaces that are not replicating | <p>The default namespace has a DPL of 1 (one) and none of the directories in it are being replicated. When the default namespace has a DPL of 1, all directories in it should be replicated to ensure that the stored data is protected. If replication is not available, increase the namespace DPL.</p> <p>This alert appears only in HCP systems in which setting the DPL to one for the namespace can leave objects unprotected.</p> |
|  | Unavailable objects | <p>The default namespace contains unavailable objects. Please contact your HCP system administrator.</p> |

Namespace Overview page alerts

| Icon | Mouse-over text | Description |
|-------------------------------------------------------------------------------------|-------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Default namespace may be under-replicated | <p>The default namespace has a DPL of 1 (one) and at least one of the directories in it is not being replicated. When the default namespace has a DPL of 1, all directories in it should be replicated to ensure that the stored data is protected. If replication is not available, increase the namespace DPL.</p> <p>This alert appears only in HCP systems in which setting the DPL to one for the namespace can leave objects unprotected.</p> |
|  | DPL1 namespace is not replicating | <p>The default namespace has a DPL of 1 (one) and none of the directories in it are being replicated. When the default namespace has a DPL of 1, all directories in it should be replicated to ensure that the stored data is protected. If replication is not available, increase the namespace DPL.</p> <p>This alert appears only in HCP systems in which setting the DPL to one for the namespace can leave objects unprotected.</p> |
|  | Irreparable and unavailable objects | <p>The default namespace contains irreparable objects and unavailable objects. Please contact your HCP system administrator.</p> <p>To see which objects are irreparable, go to the Irreparable Objects panel. For information on this panel, see "Working with irreparable objects" on page 83.</p> |
|  | Irreparable objects | <p>The default namespace contains irreparable objects. Please contact your HCP system administrator.</p> <p>To see which objects are irreparable, go to the Irreparable Objects panel. For information on this panel, see "Working with irreparable objects" on page 83.</p> |
|  | Unavailable objects | <p>The default namespace contains unavailable objects. Please contact your HCP system administrator.</p> |

Search page alert

| Icon | Mouse-over text | Description |
|-----------------------------------------------------------------------------------|---------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Namespaces without custom metadata indexing | The content class is associated with one or more namespaces for which custom metadata indexing is disabled. The metadata query engine is not indexing content properties for objects in those namespace. |

Tenant log messages

The tenant log contains messages about events that happen at the tenant and namespace levels. The table in this appendix lists the messages HCP can write to the tenant log. The messages are listed in order by event ID.

For each message, the table shows:

- The message ID
- The short form of the message, which identifies the event to which the message applies
- An explanation of the message
- The action, if any, you should take in response to the message
- The message severity

For more information on the system log, see [“Monitoring the tenant”](#) on page 50.

| ID | Event | Explanation | Action | Severity |
|------|-----------------------------------------------------------|------------------------------------------------------------------------------------------------------------------|----------------------------------------|----------|
| 2005 | HCP found an unavailable object | HCP could not repair an object because the object was unavailable. | Contact your HCP system administrator. | Warning |
| 2006 | HCP found an irreparable object | HCP found a broken object it could not repair. | Contact your HCP system administrator. | Error |
| 2028 | HCP found an irreparable object | HCP was unable to repair the object. The repair may be retried at a later time. | Contact your HCP system administrator. | Error |
| 2044 | HCP found an unavailable object | HCP could not repair an object because the object was unavailable. | Contact your HCP system administrator. | Warning |
| 2046 | HCP found an irreparable object | HCP found a broken object it could not repair. | Contact your HCP system administrator. | Error |
| 2070 | Object has been shredded | An object was shredded. | No action is required. | Notice |
| 2087 | Disposition service stopped: run complete | The disposition service finished successfully. | No action is required. | Notice |
| 2088 | Disposition service stopped without finishing | The disposition service stopped without completing its run. The service will resume at some point in the future. | No action is required. | Warning |
| 2089 | Disposition service stopped: run complete | The disposition service finished successfully. | No action is required. | Notice |
| 2090 | Disposition service stopped without finishing | The disposition service stopped without completing its run. The service will resume at some point in the future. | No action is required. | Warning |
| 2092 | Disposition service stopped without finishing | The disposition service stopped without completing its run. The service will resume at some point in the future. | No action is required. | Notice |
| 2093 | Disposition service stopped without finishing | The disposition service stopped without completing its run. The service will resume at some point in the future. | No action is required. | Notice |
| 2158 | Additional top-level directories selected for replication | The indicated top-level directories were selected for replication. | No action is required. | Notice |

(Continued)

| ID | Event | Explanation | Action | Severity |
|------|------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|----------|
| 2159 | One or more top-level directories removed from replication | The indicated top-level directories were removed from replication. | No action is required. | Notice |
| 2900 | Privileged delete requested | A user requested a privileged delete operation. | No action is required. | Notice |
| 2901 | Privileged delete succeeded | A privileged delete operation succeeded. | No action is required. | Notice |
| 2902 | Privileged delete failed | A privileged delete operation failed. | No action is required. | Notice |
| 2903 | Retention class created | A user created a retention class. | No action is required. | Notice |
| 2904 | Retention class updated | A user updated a retention class. | No action is required. | Notice |
| 2905 | Retention class deleted | A user deleted a retention class. | No action is required. | Notice |
| 2906 | Retention mode set | The namespace retention mode has been changed. | No action is required. | Notice |
| 3005 | Namespace updated | A user updated a namespace. | No action is required. | Notice |
| 3019 | Failed login to the Search Console | A user tried to log into the Search Console with a username and password that are not valid for any data access account. | Have the user log into the Search Console with a username and password for that are valid for a data access account. | Warning |
| 3028 | Failed data access account login attempt | An attempt to log into a tenant failed because the data access account could not be authenticated. | No action is required. | Warning |
| 3508 | HCP found an irreparable object | HCP found a broken object it could not repair. | Contact your HCP system administrator. | Error |
| 3998 | Search enabled or disabled | A user enabled or disabled search. | No action is required. | Notice |
| 3999 | Search indexing enabled or disabled | A user enabled or disabled search indexing. | No action is required. | Notice |
| 4005 | User authenticated | A user login to the Tenant Management Console was successfully authenticated. | No action is required. | Notice |

(Continued)

| ID | Event | Explanation | Action | Severity |
|------|---------------------------------------------|--------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|----------|
| 4006 | Authentication attempt by unknown user | A user tried to log in with an unknown username. | Have the user log in with a valid username and password. | Warning |
| 4008 | Account is disabled | A user tried to log in with a disabled account. | Reenable the user account to allow the user to log in. | Warning |
| 4009 | Account has been inactive for too long | A user tried to log in with an account that was disabled due to inactivity. | Reenable the user account to allow the user to log in. | Warning |
| 4010 | Account does not include the required roles | A user tried to log in with an account that does not include a required role. | Update the account to include the required role to allow the user to log in. | Warning |
| 4011 | Password is invalid | A user tried to log in with an invalid password. | Have the user log in with a valid username and password. | Warning |
| 4012 | Remote authentication server error | The login for a remotely authenticated user failed due to an error communicating with a RADIUS server. | Contact your HCP system administrator. | Warning |
| 4013 | Password changed | A user changed the password for a user account. | No action is required. | Notice |
| 4019 | Configuration changed | A user changed a configuration value of an HCP component. | No action is required. | Notice |
| 4020 | Configuration changed | A user changed a configuration value of an HCP component. | No action is required. | Notice |
| 4021 | Irreparable object acknowledged | A user acknowledged an irreparable object. | No action is required. | Warning |
| 4022 | All irreparable objects acknowledged | A user acknowledged all irreparable objects. | No action is required. | Warning |
| 4023 | Unauthorized action | A user has requested an operation that is not authorized for the user account. | If the user should be allowed to perform this operation, add the required role to the user account. | Warning |

(Continued)

| ID | Event | Explanation | Action | Severity |
|------|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| 4100 | Object replicated with collisions | An object being replicated conflicts with an existing object on the target system. The object has been stored in the .lost+found directory on the target system. | No action is required. | Warning |
| 4101 | Object did not replicate | An object was not replicated. | Contact your HCP system administrator. | Error |
| 4102 | Object did not replicate; will retry later | An object was not replicated. Replication of the object will be retried later. | Monitor the replica to see whether this object is eventually replicated. If the object does not replicate within one week, contact your HCP system administrator. | Warning |
| 4114 | Metadata query engine indexing failure | The metadata query engine encountered an object it could not index. | If this situation persists, contact your HCP system administrator. | Warning |
| 4115 | Metadata query engine checkpoint reset | A user reset the metadata query engine checkpoint for the indicated namespace. | No action is required. | Warning |
| 4124 | Content class created | A user created a content class. | No action is required. | Notice |
| 4125 | Content class updated | A user updated a content class. | No action is required. | Notice |
| 4126 | Content class deleted | A user deleted a content class. | No action is required. | Notice |
| 4127 | Namespaces associated with content class | A user associated namespaces with a content class. | No action is required. | Notice |
| 4128 | Content classes associated with namespace | A user associated content classes with a namespace. | No action is required. | Notice |



Browser configuration for single sign-on with Active Directory

If HCP is configured to support AD, you can use a recognized AD user account to access the Tenant Management Console with single sign-on. However, for this to work, the web browser you use to access the Console must be configured to support single sign-on.

This appendix contains instructions for configuring Windows Internet Explorer[®] and Mozilla[®] Firefox[®] to support single sign-on.

Configuring Windows Internet Explorer for single sign-on

To configure Windows Internet Explorer for single sign-on with Active Directory:

1. Open Internet Explorer.
2. On the **Tools** menu, click on **Internet Options**.
3. In the **Internet Options** window, click on the **Security** tab.
4. On the **Security** page, select **Local intranet**.
5. Click on the **Sites** button.
6. In the **Local intranet** window, ensure that all the options are selected.
7. Click on the **Advanced** button.
8. In the **Add this website to the zone** field, do either of these:
 - To enable single sign-on with HTTP, type:
`http://*.hcp-name.domain-name`

For example:
`http://*.hcp.example.com`
 - To enable single sign-on with HTTPS, type:
`https://*.hcp-name.domain-name`

For example:
`https://*.hcp.example.com`
9. Click on the **Add** button.
10. Click on the **Close** button.
11. In the **Local intranet** window, click on the **OK** button.
12. In the **Internet Options** window, click on the **Advanced** tab.

13. In the **Settings** list, under **Security**, select **Enable Integrated Windows Authentication**.
14. Click on the **OK** button.
15. Close Internet Explorer.

Configuring Mozilla Firefox for single sign-on

To configure Mozilla Firefox for single sign-on with Active Directory:

1. Open Firefox.
2. In the address field in the Firefox window, enter:

`about:config`
3. In response to the warning message, click on the **I'll be careful, I promise!** button.
4. In the **Preference Name** list, double-click on **network.negotiate-auth.delegation-uris**.
5. In the **Enter string value** window, type:

`http://*.hcp-name.domain-name,https://*.hcp-name.domain-name`

For example:

`http://*.hcp.example.com,https://*.hcp.example.com`
6. Click on the **OK** button.
7. In the **Preference Name** list, double-click on **network.negotiate-auth.trusted-uris**.
8. In the **Enter string value** window, type:

`http://*.hcp-name.domain-name,https://*.hcp-name.domain-name`
9. Click on the **OK** button.
10. Close Firefox.



Glossary

A

access protocol

See [namespace access protocol](#).

Active Directory (AD)

A Microsoft product that, among other features, provides user authentication services.

AD

See [Active Directory \(AD\)](#).

alert

A graphic that indicates the status of some particular element of an HCP system in the Tenant Management Console.

allow list

A list of IP addresses that are allowed access to the HCP system when using a particular external interface (such as a namespace access protocol).

anonymous access

A method of access to a namespace wherein the user or application gains access without presenting any credentials. *See also* [authenticated access](#).

appendable object

An object to which data can be added after it has been successfully stored. Appending data to an object does not modify the original fixed-content data, nor does it create a new version of the object. Once the new data is added to the object, that data also cannot be modified.

Appendable objects are supported only with the CIFS and NFS protocols.

atime

In POSIX file systems, metadata that specifies the date and time a file was last accessed. In HCP, POSIX metadata that initially specifies the date and time at which an object was ingested. HCP does not automatically change the **atime** value when the object is accessed.

Users and applications can change this metadata, thereby causing it to no longer reflect the actual storage time. Additionally, HCP can be configured to synchronize **atime** values with retention settings.

authenticated access

A method of access to the HCP system or a namespace wherein the user or application presents credentials to gain access. *See also* [anonymous access](#).

C

capacity

The total amount of primary storage space in HCP, excluding the space required for system overhead and the operating system. This is the amount of space available for all data to be stored in primary running storage and primary spindown storage, including the fixed-content data, metadata, any redundant data required to satisfy service plans, and the metadata query engine index.

CIFS

Common Internet File System. One of the namespace access protocols supported by HCP. CIFS lets Windows clients access files on a remote computer as if the files were part of the local file system.

compliance mode

The retention mode in which objects under retention cannot be deleted through any mechanism. This is the more restrictive retention mode.

cryptographic hash value

A system-generated metadata value calculated by a cryptographic hash algorithm from object data. This value is used to verify that the content of an object has not changed.

ctime

POSIX metadata that specifies the date and time of the last change to the metadata for an object. For a directory, this is the time of the last change to the metadata for any object in the directory.

custom metadata

User-supplied information about an HCP object. Users and applications can use custom metadata to understand and repurpose object content.

D**data access permission mask**

A set of permissions that determine which of these operations are allowed in a namespace: read, write, delete, privileged operations, and search. Data access permission masks are defined at the system, tenant, and namespace level. The effective permissions for a namespace are those that are allowed at all three levels.

Data Migrator

See [HCP Data Migrator \(HCP-DM\)](#).

data protection level (DPL)

The number of copies of the data for an object HCP must maintain in the repository. The DPL for an object is determined by the service plan that applies to the namespace containing the object.

dead properties

For WebDAV only, arbitrary name/value pairs that the server stores but does not use or modify in any way.

default namespace

A namespace that supports only anonymous access through the HTTP protocol. An HCP system can have at most one default namespace. The default namespace is used mostly with applications that existed before release 3.0 of HCP.

default tenant

The tenant that manages the default namespace.

deny list

A list of IP addresses that are denied access to the HCP system when using a particular external interface (such as a namespace access protocol).

disposition

The automatic deletion of an expired object by HCP.

DNS

See [domain name system \(DNS\)](#).

domain

A group of computers and devices on a network that are administered as a unit.

domain name system (DNS)

A network service that resolves domain names into IP addresses for client access.

DPL

See [data protection level \(DPL\)](#).

dynamic DPL

A namespace data protection level that, at any given time, matches the system-level DPL setting.

E

economy storage

See [HCP S Series Node](#).

effective permissions

For a tenant, the permissions that are included in both the system-level *and* tenant-level permission masks.

For a namespace, the permissions that are included in all three of the system-level, tenant-level, *and* namespace-level permission masks.

enterprise mode

The retention mode in which these operations are allowed:

- Privileged delete
- Changing the retention class of an object to one with a shorter duration
- Reducing retention class duration
- Deleting retention classes

This is the less restrictive retention mode.

expired object

An object that is no longer under retention.

F**fixed-content data**

A digital asset ingested into HCP and preserved in its original form as the core part of an object. Once stored, fixed-content data cannot be modified.

G**GID**

POSIX group identifier.

group account

A representation of an Active Directory group in HCP. A group account enables Active Directory users in the Active Directory group to access one or more HCP interfaces.

H

hash value

See [cryptographic hash value](#).

HCP

See [Hitachi Content Platform \(HCP\)](#).

HCP Data Migrator (HCP-DM)

An HCP utility that can transfer data from one location to another, delete data from a location, and change object metadata in a namespace. Each location can be a local file system, an HCP namespace, a default namespace, or an HCAP 2.x archive.

HCP-DM

See [HCP Data Migrator \(HCP-DM\)](#).

HCP-FS

See [HCP file system \(HCP-FS\)](#).

HCP file system (HCP-FS)

The HCP runtime component that represents each object in a namespace as a set of files. One of these files contains the object data. The others contain the object metadata.

HCP management API

A RESTful HTTP interface to a subset of the administrative functions of an HCP system. Using this API, you can manage tenants, namespaces, content classes, retention classes, and tenant-level user and group accounts.

HCP metadata query API

See [metadata query API](#).

HCP namespace

A namespace that supports user authentication for data access through the HTTP, HS3, and CIFS protocols. HCP namespaces also support storage usage quotas, access control lists, and versioning. An HCP system can have multiple HCP namespaces.

HCP object package (HOP)

The backup format for an HCP object. An HOP is an OpenPGP-wrapped tar file that unites multiple files containing the data, if applicable, and metadata for an object.

HCP service

See [service](#).

HCP tenant

A tenant created to manage HCP namespaces.

HDDS

See [Hitachi Data Discovery Suite \(HDDS\)](#).

HDDS search facility

One of the search facilities available for use with the HCP Search Console. This facility interacts with Hitachi Data Discovery Suite.

Hitachi Content Platform (HCP)

A distributed object-based storage system designed to support large, growing repositories of fixed-content data. HCP provides a single scalable environment that can be used for archiving, business continuity, content depots, disaster recovery, e-discovery, and other services. With its support for multitenancy, HCP securely segregates data among various constituents in a shared infrastructure. Clients can use a variety of industry-standard protocols and various HCP-specific interfaces to access and manipulate objects in an HCP repository.

Hitachi Data Discovery Suite (HDDS)

A Hitachi product that enables federated searches across multiple HCP systems and other supported systems.

hold

A condition that prevents an object from being deleted by any means and from having its metadata modified, regardless of its retention setting, until it is explicitly released.

HOP

See [HCP object package \(HOP\)](#).

HS3 API

One of the namespace access protocols supported by HCP. HS3 is a RESTful, HTTP-based API that is compatible with Amazon S3. Using HS3, users and applications can create and manage buckets and bucket contents.

HSwift API

One of the namespace access protocols supported by HCP. HSwift is a RESTful, HTTP-based API that is compatible with OpenStack Swift. Using HSwift, users and applications can create and manage containers and container contents.

HTTP

HyperText Transfer Protocol. One of the namespace access protocols supported by HCP.

HCP also uses HTTP for client communication with the Tenant Management and Search Consoles, for client access through the HCP management API, and for access to namespace content through the metadata query API.

HTTPS

HTTP with SSL security. *See [HTTP](#) and [SSL](#).*

I

index

An index of the objects in namespaces that is used to support search operations. Each of the two search facilities, the metadata query engine and the HDDS search facility, creates and maintains its own separate index.

index setting

The property of an object that determines whether the metadata query engine indexes the custom metadata associated with the object.

M

management API

See [HCP management API](#).

metadata

System-generated and user-supplied information about an object. Metadata is stored as an integral part of the object it describes, thereby making the object self-describing.

metadata query API

A RESTful HTTP interface that lets you search HCP for objects that meet specified metadata-based or operation-based criteria. With this API, you can search not only for objects currently in the repository but also for information about objects that are no longer in the repository.

metadata query engine

One of the search facilities available for use with HCP. The metadata query engine works internally to perform searches and return results either through the metadata query API or to the HCP Metadata Query Engine Console (also known as the HCP Search Console).

Metadata Query Engine Console

The web application that provides interactive access to the HCP search functionality provided by the metadata query engine.

N**namespace**

A logical partition of the objects stored in an HCP system. A namespace consists of a grouping of objects such that the objects in one namespace are not visible in any other namespace. Namespaces are configured independently of each other and, therefore, can have different properties.

namespace access protocol

A protocol that can be used to transfer data to and from namespaces in an HCP system. HCP supports the HTTP, HS3, WebDAV, CIFS, NFS, and SMTP protocols for access to HCP namespaces and the default namespace. HCP also supports the NDMP protocol for access to the default namespace.

NDMP

Network Data Management Protocol. The namespace access protocol HCP supports for backing up and restoring objects in the default namespace.

NFS

Network File System. One of the namespace access protocols supported by HCP. NFS lets clients access files on a remote computer as if the files were part of the local file system.

O

object

An exact digital representation of data as it existed before it was ingested into HCP, together with the system and custom metadata that describes that data. An object is handled as a single unit by all transactions and internal processes, including shredding, indexing, and replication.

object-based query

In the metadata query API, a query that searches for objects based on object metadata. This includes both system metadata and the content of custom metadata. The query criteria can also include the object location (that is, the namespace and/or directory that contains the object).

Object-based queries search only for objects that currently exist in the repository.

operation-based query

In the metadata query API, a query that searches not only for objects currently in the repository but also for information about objects that have been deleted by a user or application or deleted through disposition.

Criteria for operation-based queries can include object status (for example, created or deleted), change time, index setting, and location (that is, the namespace and/or directory that contains the object).

P

permission

One of these:

- In POSIX permissions, the ability granted to the owner, the members of a group, or other users to access an object, directory, or symbolic link. A POSIX permission can be read, write, or execute.

- In a data access permission mask, the condition of allowing a specific type of operation to be performed in a namespace.
- The granted ability to access the HCP System Management Console, Tenant Management, or HCP Search Console and to perform a specific activity or set of activities in that Console. Permissions of this type are granted by roles associated with the user account.

permission mask

See [data access permission mask](#).

policy

One or more settings that influence how transactions and internal processes work on objects.

POSIX

Portable Operating System Interface for UNIX. A set of standards that define an application programming interface (API) for software designed to run under heterogeneous operating systems. HCP-FS is a POSIX-compliant file system, with minor variations.

privileged delete

A delete operation that works on an object regardless of whether the object is under retention, except if the object is on hold. This operation is available only to users and applications with explicit permission to perform it.

Privileged delete operations work only in namespaces in enterprise mode.

protocol

See [namespace access protocol](#).

Q**query**

A request submitted to HCP to return metadata for objects that satisfy a specified set of criteria. Also, to submit such a request.

query API

See [metadata query API](#).

R

recognized Active Directory user account

An Active Directory user account for a user that belongs to one or more Active Directory groups for which corresponding group accounts are defined in HCP.

remote authentication

Authentication wherein HCP uses a remote service to check the validity of the specified username and password.

replica

The HCP system to which the replication service copies objects and other information from the primary system during normal replication.

replication

The process of keeping selected HCP tenants and namespaces and selected default-namespace directories in two HCP systems in sync with each other. Basically, this entails copying object creations, deletions, and metadata changes from each system to the other or from one system to the other. HCP also replicates retention classes and all compliance log messages.

repository

The aggregate of the namespaces defined for an HCP system.

REST

Representational State Transfer. A software architectural style that defines a set of rules (called constraints) for client/server communication. In a REST architecture:

- Resources (where a resource can be any coherent and meaningful concept) must be uniquely addressable.
- Representations of resources (for example, in XML format) are transferred between clients and servers. Each representation communicates the current or intended state of a resource.
- Clients communicate with servers through a uniform interface (that is, a set of methods that resources respond to) such as HTTP.

retention class

A named retention setting. The value of a retention class can be a duration, Deletion Allowed, Deletion Prohibited, or Initial Unspecified.

retention hold

See [hold](#).

retention mode

A namespace property that affects which operations are allowed on objects under retention. A namespace can be in either of two retention modes: compliance or enterprise.

retention period

The period of time during which an object cannot be deleted (except by means of a privileged delete).

retention setting

The property that determines the retention period for an object.

role

A named collection of permissions that can be associated with an HCP user account, where each permission allows the user to perform some specific interaction or set of interactions with the HCP Tenant Management Console, the HCP management API, the metadata query API, or the HCP Search Console. Roles generally correspond to job functions.

S**Search Console**

The web application that provides interactive access to HCP search functionality. When the Search Console uses the HCP metadata query engine for search functionality, it is called the Metadata Query Engine Console.

search facility

An interface between the HCP Search Console and the search functionality provided by the metadata query engine or HDDS. Only one search facility can be selected for use with the Search Console at any given time.

service

A background process that performs a specific function that contributes to the continuous tuning of the HCP system. In particular, services are responsible for optimizing the use of system resources and maintaining the integrity and availability of the data stored in the HCP repository.

service plan

A named option that can be associated with a namespace and that determines how HCP manages the objects in that namespace. Service plan names are system specific.

shred setting

The property that determines whether an object will be shredded or simply removed when it's deleted from HCP.

shredding

The process of deleting an object and overwriting the locations where all its copies were stored in such a way that none of its data or metadata can be reconstructed. Also called **secure deletion**.

single sign-on

In a Windows environment, the use of an already authenticated Active Directory user account to access the System Management Console, Tenant Management Console, or HCP Search Console without the need to explicitly log in.

SMTP

Simple Mail Transfer Protocol. The namespace access protocol HCP uses to receive and store email data directly from email servers.

SNMP

Simple Network Management Protocol. A protocol HCP uses to facilitate monitoring and management of the system through an external interface.

SSL

Secure Sockets Layer. A key-based Internet protocol for transmitting documents through an encrypted link.

syslog

A protocol used for forwarding log messages in an IP network. HCP uses syslog to facilitate system monitoring through an external interface.

System Management Console

The system-specific web application that lets you monitor and manage HCP.

system metadata

System-managed properties that describe the content of an object. System metadata includes policies, such as retention and data protection level, that influence how transactions and services affect the object.

systemwide permission mask

The data access permission mask defined at the HCP system level. The systemwide permission mask applies across all tenants and namespaces.

T**tenant**

An administrative entity created for the purpose of owning and managing namespaces. Tenants typically correspond to customers or business units.

Tenant Management Console

The tenant-specific web application that lets you monitor and manage tenants and namespaces.

U**UID**

POSIX user ID.

Unix

Any UNIX-like operating system (such as UNIX itself or Linux).

user account

A set of credentials that gives a user access to one or more of the System Management Console, the Tenant Management Console, the HCP management API, the HCP Search Console, namespace content through the namespace access protocols, the metadata query API, and HCP Data Migrator.

user authentication

The process of checking that the combination of a specified username and password is valid when a user tries to log into the Tenant Management Console or the HCP Search Console, to access the HCP system through the management API, or to access a namespace.

V

versioning

An optional namespace feature that enables the creation and management of multiple versions of an object.

W

WebDAV

Web-based Distributed Authoring and Versioning. One of the namespace access protocols supported by HCP. WebDAV is an extension of HTTP.

Windows workgroup

A named collection of computers on a LAN that share resources such as printers and file servers.

workgroup

See [Windows workgroup](#).

WORM

Write once, read many. A data storage property that protects the stored data from being modified or overwritten.

X

XML

Extensible Markup Language. A standard for describing data content using structural tags called elements.

Index

Symbols

- [internal] event initiator 53
- [remote admin] event initiator 53
- [service] event initiator 53

A

- access methods, CIFS 97
- access protocols
 - See [namespace access protocols](#)
- acknowledging irreparable objects 83–84
- Active Directory
 - authentication, about 97
 - authentication, configuring 103–104
 - single sign-on 31
 - user permissions 28
- adding
 - See also [creating](#)
 - content properties to content classes 144–145
 - content properties to content classes individually 145
 - email recipients 62
 - IP addresses to Allow/Deny lists 87
 - username mappings 103–104
- alerts
 - content classes 170
 - namespace 71, 169
 - Overview page, namespace 169
 - Overview page, tenant 168
 - tenant 45, 168
- All Events panel
 - namespace 82
 - tenant 51
- Allow lists
 - about 86–87
 - adding IP addresses 87
 - CIFS protocol handling of 88
 - deleting IP addresses 87
 - HTTP protocol handling of 88

- NDMP protocol handling of 89
- NFS protocol handling of 89
- SMTP protocol handling of 89
- valid entries 87
- WebDAV protocol handling of 88
- allowing namespace access by IP addresses 86–89
 - See also [enabling](#)
- anonymous access, CIFS
 - about 97
 - configuring 101–102
- appendable objects
 - about 3, 14
 - enabling/disabling 77
- associating
 - content classes with default namespace 151, 152–153
 - default namespace with content classes 148–149
 - service plan with default namespace 80
- atime synchronization
 - about 14
 - enabling/disabling 77
- authenticated access, CIFS
 - about 97
 - configuring 101–102
- automatic deletion
 - See [disposition](#)

B

- backing up default namespace
 - about 112–113
 - encryption/signing keys for 114
 - NDMP protocol 112–115
 - need for 40
 - performance considerations 114–115
 - third-party applications 120
- basic authentication, WebDAV 95
- Boolean data type 132

capacity

C

- capacity 70
- case forcing, CIFS protocol
 - about 100
 - changing 102
- case sensitivity, CIFS protocol
 - about 99–100
 - changing 102
- Change Password page 37
- changing
 - See also* [configuring](#); [modifying](#)
 - appendable objects setting 77
 - atime synchronization setting 77
 - CIFS case forcing 102
 - CIFS case sensitivity 102
 - collision handling option 78–79
 - compatibility properties 77
 - custom metadata operations allowed for
 - objects under retention 75–76
 - namespace description 74
 - namespace permission mask 73–74
 - passwords 36–37
 - retention class values 159
 - retention mode 80–81
 - retention-related properties 75–76
 - service plan for default namespace 80
 - tenant contact information 47–48
 - tenant description 49–50
 - tenant permission mask 48–49
- CIFS panel 96, 101–102
- CIFS protocol
 - about 5
 - access methods 97
 - Active Directory authentication 103–104
 - Allow/Deny list handling 88
 - Allow/Deny lists 86–87
 - anonymous access 97
 - authenticated access 97
 - case forcing 100, 102
 - case sensitivity 99–100, 102
 - configuring 96, 101–102
 - default ownership/permissions 90
 - UIDs/GIDs 98–99
- collision handling
 - See* [replication collisions](#)
- Compatibility panel (Settings page) 77
- compatibility properties
 - about 14
 - changing 77
- compliance activities, performing 41
- Compliance Events panel
 - namespace 82
 - tenant 52
- compliance mode 11
 - See also* [retention mode](#)
- Compliance page 157, 163
 - Create Retention Class panel 158
 - Privileged Delete panel 163–164
 - Retention Classes panel 157
- configuring
 - See also* [changing](#); [modifying](#)
 - CIFS protocol 96
 - default namespace 72–73
 - default tenant 46–47
 - email notification 56–62
 - Firefox for single sign-on 179
 - HTTP protocol 92–94
 - Internet Explorer for single sign-on 178–179
 - Microsoft Exchange 2003 for SMTP 110–111
 - Microsoft Exchange 2007 for SMTP 111
 - Microsoft Exchange 2010 for SMTP 112
 - namespace access protocols 86
 - NDMP protocol 112–120
 - NFS protocol 104–105
 - SMTP protocol 105–109
 - WebDAV protocol 92–95
- Console pages
 - See also* [Tenant Management Console](#)
 - Change Password 37
 - Compliance 157, 163
 - Email 57–62
 - Idle Timeout 31
 - login 34–35
 - Monitoring 82–84
 - Overview, default namespace 68–72
 - Overview, tenant 44–46
 - Policies (Metadata panel) 76
 - Policies (Retention panel) 75
 - Protocols 86
 - refreshing 35
 - Replication 63–65
 - Services (Disposition panel) 78
 - Services (Replication panel) 79
 - Services (Search panel) 152
 - Services (Service Plan panel) 80
 - Settings (Compatibility panel) 77
 - Settings (Retention Mode panel) 81
 - SNMP 56
 - Syslog 55
 - Tenant Events 51–53
- contact information
 - changing 47–48
 - viewing 45–46
- Contact Information window 47–48
- content classes
 - about 123

- alerts 170
 - associating/dissociating with default namespace 148–149, 151, 152–153
 - creating 144
 - deleting 150
 - list 142–143
 - managing content properties 144–147
 - names 125, 144
 - reindexing default namespace 149–150
 - renaming 150
 - workflow 125–126
 - content properties
 - about 124
 - adding 144–145
 - adding individually 145
 - data types 131–132
 - definitions 127–128
 - definitions in content property files 139–141
 - deleting 144–145
 - exporting 147
 - expressions 129–131
 - extracted from XML 137–139
 - extracting from XML 146
 - importing 146
 - indexing 124
 - list 143
 - modifying 144–145
 - multivalued 137
 - names 128–129
 - testing 147
 - workflow 125–126
 - content property files
 - about 139–141
 - exporting content properties to 147
 - Create Retention Class panel 158
 - creating
 - See also* [adding](#)
 - content classes 144
 - retention classes 158
 - username mapping files 103
 - cryptographic hash algorithms
 - about 10–11
 - viewing 72
 - cryptographic hash values 10
 - custom metadata
 - about 2
 - dead properties as 95
 - enabling/disabling full indexing 151, 152–153
 - enabling/disabling indexing 151, 151, 152–153
 - indexing by metadata query engine 124–125
 - indexing, about 122–123
 - operations allowed for objects under retention, about 13
 - operations allowed for objects under retention, changing 75–76
 - sample 127–128
 - XML checking, about 13–14
 - XML checking, enabling/disabling 76
 - custom metadata collisions 24–25
- D**
- data
 - fixed content 2
 - security 94
 - transmission rate, replication 64–65
 - data access permission masks 15–17
 - See also* [namespace permission mask](#); [tenant permission mask](#)
 - Data Migrator
 - See* [HCP Data Migrator](#)
 - data protection level
 - See* [DPL](#)
 - data types for content properties
 - about 131–132
 - formats 132–137
 - datetime data type
 - about 132
 - formats 135–137
 - dead properties 95
 - default email message template 60
 - default namespace
 - See also* [namespaces](#)
 - about 3–4
 - access to 5–8
 - alerts 71, 169
 - all events 82
 - allowing/denying access by IP addresses 86–89
 - associating/dissociating content classes with 151, 152–153
 - associating/dissociating with content classes 148–149
 - backing up 40, 112–115
 - collision handling option, changing 78–79
 - compatibility properties, changing 77
 - compliance events 82–83
 - configuring 72–73
 - cryptographic hash algorithm, viewing 72
 - custom metadata operations allowed for objects under retention, changing 75–76
 - custom metadata XML checking, enabling/disabling 76
 - description, changing 74

default namespace, DPL, changing

- description, viewing 72
 - disposition, enabling/disabling 78
 - DPL, viewing 72
 - enabling/disabling search 152–153
 - ensuring recovery 40
 - features 71
 - indexed object count 69
 - indexing 122–123
 - maintaining 38
 - major events 70–71
 - managing access to 40
 - monitoring 39–40, 81
 - object count 68–69
 - permission mask, changing 73–74
 - permission mask, viewing 72
 - read from remote system, enabling/
disabling 78–79
 - reindexing from Search page 148, 149–150
 - reindexing from Search panel 151, 153
 - reindexing, about 126
 - restoring 112–115
 - retention mode, changing 80–81
 - retention mode, viewing 72
 - search and indexing options 151
 - search enabled 122
 - search feature, viewing 71
 - segmenting for backup 115
 - service by remote systems, enabling/
disabling 78–79
 - service plan, changing 80
 - service plan, viewing 71
 - storage usage 69–70
 - URL 68
- default ownership/permissions 90
- default tenant
- See also* [tenants](#)
 - about 4
 - administrator responsibilities 38–41
 - alerts 45, 168
 - all events 51
 - compliance events 52
 - configuring 46–47
 - contact information, changing 47–48
 - contact information, viewing 45–46
 - description, changing 49–50
 - description, viewing 46
 - maintaining 38
 - major events 44–45
 - monitoring 39–40, 50–51
 - permission mask, changing 48–49
 - permission mask, viewing 46
 - read-only 45
 - security events 52
 - statistics 44
- delete permission 15
- deleting
- content classes 150
 - content properties 144–145
 - email recipients 62
 - IP addresses from Allow/Deny lists 87
 - irreparable objects 83
 - NDMP signing keys 120
 - objects under retention 162–164
 - retention classes 159–160
 - username mappings 103–104
- Deletion Allowed 106, 156
- Deletion Prohibited 106, 156
- Deny lists
- about 86–87
 - adding IP addresses 87
 - CIFS protocol handling of 88
 - deleting IP addresses 87
 - HTTP protocol handling of 88
 - NDMP protocol handling of 89
 - SMTP protocol handling of 89
 - valid entries 87
 - WebDAV protocol handling of 88
- denying namespace access by IP addresses 86–89
- See also* [disabling](#)
- descriptions
- namespace 72, 74
 - retention classes 158
 - tenant 49–50
- detailed view of replication 64–65
- directories
- default permissions 90–91
 - fcfs_data 5
 - fcfs_metadata 5
 - replicated 19
 - storing 2
- disabling
- appendable objects 77
 - atime synchronization 77
 - custom metadata indexing 151, 151, 152–153
 - custom metadata XML checking 76
 - disposition 78
 - full custom metadata indexing 151, 152–153
 - indexing 151, 152–153
 - metadata overrides 94
 - ownership/permission changes for objects
under retention 75–76
 - read from remote system 78–79
 - search 123, 151, 152–153
 - service by remote systems 78–79
- disposition
- about 14–15

- enabling/disabling 78
 - with retention classes 156
 - Disposition panel (Services page) 78
 - dissociating
 - content classes from default namespace 151, 152–153
 - default namespace from content classes 148–149
 - documentation, viewing 36
 - downloading
 - encryption/signing keys for NDMP 118–119
 - HCP Data Migrator 166
 - HCP documentation 36
 - DPL
 - about 10–??
 - dynamic 10
 - viewing 72
 - dynamic DPL 10
- E**
- Edit Retention Class window 159
 - effective permission mask
 - namespace 72
 - tenant 46
 - email
 - See also* [email notification](#)
 - archiving 105
 - retention settings 106–107
 - settings 107–108
 - storing attachments 106
 - email message template
 - See also* [email notification](#)
 - about 58
 - default 60
 - modifying 59
 - variables 59–60
 - email notification
 - about 56–57
 - enabling 57
 - message template 58–60
 - recipients 61–62
 - testing 57–58
 - Email page 57–62
 - email recipients 61–62
 - See also* [email notification](#)
 - enabling
 - appendable objects 77
 - atime synchronization 77
 - CIFS protocol 101–102
 - custom metadata indexing 151, 151, 152–153
 - custom metadata XML checking 76
 - disposition 78
 - email notification 57
 - full custom metadata indexing 151, 152–153
 - HTTP protocol 92–94
 - indexing 151, 152–153
 - metadata overrides 94
 - NDMP protocol 115–118
 - NFS protocol 105
 - ownership/permission changes for objects
 - under retention 75–76
 - read from remote system 78–79
 - search 123, 151, 152–153
 - service by remote systems 78–79
 - SMTP protocol 107–109
 - WebDAV protocol 92–95
 - encryption keys, NDMP
 - about 114, 118
 - downloading 118–119
 - uploading 119–120
 - enforcing permissions, HTTP/WebDAV protocols 94
 - enterprise mode 11
 - See also* [retention mode](#)
 - Error (event severity) 54
 - events
 - See also* [log messages](#)
 - initiator 53, 54
 - severity 53–54
 - Exchange
 - See* [Microsoft Exchange, configuring for SMTP](#)
 - expired objects, automatic deletion 14–15
 - exported content properties 139–141
 - exporting content properties 147
 - expressions for content properties 129–131
 - extracted content properties 137–139
 - extracting content properties from XML 146
- F**
- fcfs_data directory 5
 - fcfs_metadata directory 5
 - features, namespace 71
 - files
 - hosts 32–33
 - objects stored for 2–3
 - username mapping 103
 - filtering content class list 143
 - Firefox, configuring for single sign-on 179
 - fixed content 2
 - float data type
 - about 132
 - formats 134–135
 - formats for content properties
 - datetime data type 135–137

formats for content properties, float data type

float data type 134–135
integer data type 132–133

full custom metadata indexing
about 124–125
enabling/disabling 151, 152–153

G

GIDs
defaults 90
objects added through CIFS 98–99
valid values 91
group accounts 28–29

H

hash algorithms
See [cryptographic hash algorithms](#)
hash values 10
HCP
about 1–2
documentation 36
email server access to 105
file system 5
service by remote systems 19, 78–79
storage capacity 70
version 34
HCP Data Migrator
about 8
downloading 166
system requirements 166
HCP management API 28
HCP namespaces 3–4
See also [namespaces](#)
HCP object packages 113
HCP Search Console 6–8
HCP search facility
indexing 122–123
HCP system administrator 29
HCP tenants 4
See also [tenants](#)
HCP-DM
See [HCP Data Migrator](#)
HCP-FS 5
HDDS search facility
about 7
indexing 123
high-level view of replication 63
Hitachi Content Platform
See [HCP](#)
Hitachi Data Discovery Suite
See [HDDS search facility](#)
HOPs 113
hostname mappings 33

hosts file 32–33
HTTP & WebDAV panel 92–95
HTTP protocol
about 5
Allow/Deny list handling 88
Allow/Deny lists 86–87
configuring 92–94
default ownership/permissions 90
metadata overrides, enabling/disabling 94
permissions enforcement 94
secure namespace access 93
HTTPS access to default namespace 93
See also [HTTP protocol](#)

I

Idle Timeout page 31
importing content properties 146
index setting 12
indexes
See [indexing](#); [search](#)
indexing
See also [search](#)
about 122–123
custom metadata by metadata query
engine 124–125
custom metadata, about 122–123
enabling/disabling 151, 152–153
ingested volume 70
Initial Unspecified 107, 156
integer data type
about 132
formats 132–133
Internet Explorer, configuring for single sign-
on 178–179
IP addresses
in Allow/Deny lists 87
for HCP system 33
irreparable objects
acknowledging 83–84
deleting 83
viewing 83
Irreparable Objects panel 83–84

L

log messages
about 53–54
descriptions 171–175
details 54
event severity 53–54
importance 56
managing list of 54–55
sending through email 56–57

- sending to SNMP managers 55–56
 - sending to syslog servers 55
 - types 56
 - viewing 50–51
- logging into Tenant Management Console 33–35
- logging out of Tenant Management Console 37

M

- Mac OS X hosts file 32
- maintaining
 - namespace 38
 - tenant 38
- major events
 - See also* [log messages](#)
 - namespace 70–71
 - tenant 44–45
- management API 28
- managing
 - content class list 142–143
 - content properties 144–147
 - namespace access 40
 - tenant log message list 54–55
- mappings, hostname 33
- maximum content classes per tenant 125
- metadata
 - about 2
 - custom 2
 - enabling/disabling overrides 94
 - system 2
- Metadata panel (Policies page) 76
- metadata query API 6
- metadata query engine
 - about 7
 - custom metadata indexing 124–125
 - indexing 122–123
- Metadata Query Engine Console 7
- metafiles 5
- Microsoft Exchange, configuring for SMTP
 - considerations for mixed Exchange 2003 and 2010 environments 110
 - Exchange 2003 110–111
 - Exchange 2007 111
 - Exchange 2010 112
- modifying
 - See also* [changing](#); [configuring](#)
 - content properties 144–145
 - email recipients 62
 - retention classes 159
 - username mappings 103–104
- monitoring
 - namespace 39–40, 81
 - replication 63–65
 - tenant 39–40, 50–51

- Monitoring page, default namespace
 - All Events panel 82
 - Compliance Events panel 82
 - Irreparable Objects panel 83–84
- Mozilla Firefox, configuring for single sign-on 179
- multivalued content properties 137

N

- names
 - content classes 125, 144
 - content properties 128–129
- names, retention class 158
- namespace access protocols
 - about 5–6
 - Allow/Deny lists 86–87
 - configuring 86
 - default ownership/permissions 90
- namespace permission mask
 - See also* [data access permission masks](#)
 - about 16
 - changing 73–74
 - viewing 72
- namespaces
 - See also* [default namespace](#)
 - about 3–4
 - compatibility properties 14
 - cryptographic hash algorithms 10–11
 - custom metadata operations allowed for
 - objects under retention 13
 - custom metadata XML checking 13–14
 - data access permission masks 15–17
 - default 3–4
 - disposition 15
 - DPL 10–??
 - effective permissions 16–17
 - enabling/disabling search 151
 - HCP 3–4
 - object ownership/permission changes 12
 - read from remote system, about 18
 - read-only 17
 - retention mode 11
 - service by remote systems 19
 - service by remote systems, about 19
- Namespaces panel (Search page) 148–149
- NDMP panel 115–120
- NDMP protocol
 - about 6
 - Allow/Deny list handling 89
 - Allow/Deny lists 86–87
 - backup/restore operations 112–113
 - configuring 112, 115–118
 - encryption keys 114, 118–120
 - signing keys 114, 118–120

NFS panel 104–105

NFS protocol

about 5

Allow list handling 89

Allow lists 86–87

configuring 104–105

default ownership/permissions 90

Notice (event severity) 53

O

objects

about 2–3

appendable 3

automatic deletion 14–15, 78

content collisions, about 20–21

content collisions, handling 78–79

default ownership 90–91

default permissions 90–91

deleting under retention 162–164

hash value 10

index setting 12

indexed in default namespace 44

list of irreparable 83

number of in default namespace 44

number of indexed over time 69

number of over time 68–69

ownership 12

ownership/permission changes under

retention, enabling/disabling 75–76

permissions 12

privileged delete 162–163

representation 5

UID/GID through CIFS 98–99

write permission, POSIX 89

octal values for permissions 91

operation rate, replication 65

operations allowed for objects under retention,

custom metadata

about 13

changing 75–76

Overview page, namespace

about 68–72

alerts 169

Overview page, tenant

about 44–46

alerts 168

ownership

about 12

defaults by protocol 90

enabling/disabling changes for objects under

retention 75–76

enabling/disabling changes to 94

P

paging through content class list 142

passwords, changing 34, 36–37

percent encoding for privileged delete 163

performance considerations, namespace

backup 114–115

performing

compliance activities 41

privileged delete operations 163–164

permissions

Active Directory users 28

granted by roles 29–31

permissions, object

about 12, 89

defaults by protocol 90

enabling/disabling changes for objects under

retention 75–76

enabling/disabling changes to 94

HTTP enforcement of 94

octal values for 91

selecting defaults from dropdown lists 90

WebDAV enforcement of 94

Policies page 75, 76

privileged delete

about 161–162

performing 163–164

Privileged Delete panel 163–164

privileged permission 16

protocols

See [namespace access protocols](#)

Protocols page

CIFS panel 96, 101–102

HTTP & WebDAV panel 92–95

NDMP panel 115–120

NFS panel 104–105

SMTP panel 105–106, 107–109

R

read from remote system

about 18

enabling/disabling 78–79

read permission 15

read-only tenants 45

recipients, email 61–62

recognized Active Directory user account 28

recovery, ensuring 40

refreshing Console pages 35

Reindex panel (Search page) 149–150

reindexing

about 126

default namespace from Search page 148,

149–150

- default namespace from Search panel 151, 153
 - removing
 - See [deleting](#)
 - renaming content classes 150
 - replication
 - about 17–18
 - benefits 18–19
 - collision handling 19–26, 78–79
 - data transmission rate 64–65
 - detailed view 64
 - high-level view 63
 - implementation 19
 - monitoring 63–65
 - operation rate 65
 - read from remote system 78–79
 - service by remote systems 19, 78–79
 - topologies 17
 - up-to-date-as-of time 63, 64
 - replication collisions
 - about 19
 - automatic deletion 15
 - custom metadata 24–25
 - object content, about 20–21
 - object content, handling 78–79
 - retention classes 25–26
 - system metadata 21–24
 - Replication page 63–65
 - Replication panel (Services page) 79
 - reserved words, content property names 129
 - responsibilities, tenant administrator 38–41
 - restoring default email message template 60
 - restoring default namespace
 - about 112–113
 - encryption/signing keys for 114
 - NDMP protocol 112–115
 - third-party applications 120
 - retention
 - deleting objects under 162–164
 - email settings 106–107
 - periods 11
 - settings 11
 - retention classes
 - about 155–156
 - automatically deleting expired objects in 156
 - collisions 25–26
 - creating 158
 - deleting 159–160
 - description 158
 - list 157
 - modifying 159
 - names 158
 - Retention Classes panel 157
 - retention mode
 - about 11
 - changing 80–81
 - viewing 72
 - Retention Mode panel (Settings page) 81
 - Retention panel (Policies page) 75
 - retention-related properties
 - about 12–13
 - changing 75–76
 - roles 28–31
- S**
- sample custom metadata
 - about 127–128
 - content properties extracted from 138–139
 - exported content property definitions 140–141
 - search
 - See also [indexing](#)
 - about 122–123
 - enabling/disabling 151, 152–153
 - namespace feature 71
 - permission 16
 - Search Console 6–8
 - search facilities
 - See also [HCP search facility](#); [HDDS search facility](#); [metadata query engine](#)
 - about 7
 - indexes 7
 - search indexes
 - See [indexes](#)
 - Search page
 - about 142–143
 - adding content properties 145
 - adding content properties individually 145
 - alerts 170
 - associating/dissociating default namespace
 - with content classes 148–149
 - creating content classes 144
 - deleting content classes 150
 - deleting content properties 145
 - exporting content properties 147
 - extracting content properties from XML 146
 - importing content properties 146
 - modifying content properties 145
 - reindexing default namespace 149
 - renaming content classes 150
 - testing content properties 147
 - Search panel (Services page)
 - displaying 152
 - reindexing default namespace 153
 - setting search and indexing options 152–153
 - Security Events panel 52

service plans

segmenting default namespace for backup 115

service by remote systems

about 19

enabling/disabling 78–79

Service Plan panel (Services page) 80

service plans

about 26

associating with default namespace 80

viewing 71

Services page, default namespace

Disposition panel 78

Replication panel 79

Search panel 152–153

Service Plan panel 80

Settings page

Compatibility panel 77

Retention Mode panel 81

Settings panel (Search page)

about 143

adding content properties 144–145

adding content properties individually 145

deleting content properties 144–145

exporting content properties 147

extracting content properties from XML 146

importing content properties 146

modifying content properties 144–145

renaming content classes 150

testing content properties 147

severity, events 53–54

shortcut keys 35

signing keys, NDMP

about 114, 118

deleting 120

downloading 118–119

uploading 119

simultaneous backups 115

single sign-on

about 31

configuring Firefox for 179

configuring Internet Explorer for 178–179

SMTP panel 105–106, 107–109

SMTP protocol

about 6

Allow/Deny list handling 89

Allow/Deny lists 86–87

configuring 105–109

default ownership/permissions 90

email retention settings 106–107

with Microsoft Exchange 109–112

SNMP logging 55–56

SNMP page 56

sorting content class list 143

statistics, tenant 44

storage

amount used per object 10

capacity 70

namespace usage 69–70

object based 2–3

string data type 131

submitting changes in Tenant Management

Console 36

symbolic links 2

syslog logging 55

Syslog page 55

system administrator 29

system metadata 2

system metadata collisions 21–24

systemwide permission mask 16

T

template, email message 58–60

tenant administrator responsibilities 38–41

Tenant Events page

All Events panel 51

Compliance Events panel 52

Security Events panel 52

tenant log 50–51

See also [log messages](#); [Monitoring page](#);

[Tenant Events page](#)

Tenant Management Console

See also [Console pages](#)

about 28

logging in 33–35

logging out 37

permissions granted by roles 29–31

sessions 31

submitting changes 36

tenant description 46

URL 32

using 35

tenant permission mask

See also [data access permission masks](#)

about 16

changing 48–49

viewing 46

tenants

See also [default tenant](#)

about 4–5

default 4

effective permissions 16

HCP 4

testing content properties 147

testing email notification 57–58

third-party backup/restore applications 120

tokenized data type 131
transfer rate, replication 64–65

U

UIDs
 defaults 90
 objects added through CIFS 98–99
 valid values 91
Unix hosts file 32
uploading
 encryption keys for NDMP 119–120
 signing keys for NDMP 119
up-to-date-as-of time 63, 64
URL encoding for privileged delete 163
URLs
 namespace 68
 Tenant Management Console 32
user accounts
 about 28–29
 recognized Active Directory 28
username mapping files
 about 103
 creating 103
 modifying 103–104
 viewing 103

V

variables, email message template 59–60
version, HCP 34
viewing
 all namespace events 82
 all tenant events 51
 HCP documentation 36
 irreparable objects 83
 namespace compliance events 82–83
 namespace description 72
 namespace permission mask 72
 namespace retention mode 72
 namespace service plan 71
 tenant compliance events 52
 tenant contact information 45–46
 tenant description 46
 tenant permission mask 46
 tenant security events 52
 username mapping files 103
volume, ingested 70

W

Warning (event severity) 54
WebDAV protocol
 about 5
 Allow/Deny list handling 88

Allow/Deny lists 86–87
basic authentication 95
configuring 92–95
dead properties 95
default ownership/permissions 90
permissions enforcement 94
secure namespace access 93
Windows case sensitivity 99–100
Windows hosts file 32
Windows Internet Explorer, configuring for single
 sign-on 178–179
WORM 2, 89
write permission
 in data access permission masks 15
 POSIX 89

X

XML checking for custom metadata
 about 13–14
 enabling/disabling 76
XML, extracting content properties from 146
XPath expressions 129–131

Hitachi Data Systems

Corporate Headquarters

2845 Lafayette Street
Santa Clara, California 95050-2627
U.S.A.

www.hds.com

Regional Contact Information

Americas

+1 408 970 1000

info@hds.com

Europe, Middle East, and Africa

+44 (0) 1753 618000

info.emea@hds.com

Asia Pacific

+852 3189 7900

hds.marketing.apac@hds.com



MK-99ARC024-13