HITACHI
Inspire the Next

Hitachi Content Platform

Replicating Tenants and Namespaces

Hitachi Data Systems

MK-98ARC015-14

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.

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

This book is your guide to data replication with **Hitachi Content Platform** (**HCP**). Replication is an HCP service that copies selected tenants and namespaces from one HCP system to another.

This book describes how replication works and contains complete instructions for setting up replication between HCP systems. It also provides information on monitoring and managing replication activity, including failover and recovery.

Intended audience

This book is intended for HCP system administrators who are responsible for configuring replication links, monitoring replication, and managing failover and recovery with a replica. It assumes you are familiar with HCP concepts and functionality and have some experience with the HCP System Management Console.

Product version

This book applies to release 7.1.1 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: **replication.admin.***hcp-domain-name* You enter: replication.admin.hcp-ca.example.com
italics	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.

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

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

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

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

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

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

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

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

- *Using HCP Data Migrator* — This book contains the information you need to install and use HCP Data Migrator (HCP-DM), a utility that works with HCP. This utility enables you to copy 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.

Thank you!

Comments

1

Replication overview

Hitachi Content Platform (**HCP**) is the distributed, fixed-content, data storage system from Hitachi Data Systems®. An HCP system maintains a repository of objects that encapsulate fixed-content data and metadata that describes the data. You can create a configuration wherein selected tenants and namespaces are copied between two or more HCP systems. This copying process is carried out by the replication service through a replication link.

This chapter provides an overview of how replication works and how it supports business continuity and disaster recovery. For an introduction to HCP systems and how they work, see *Administering HCP*.

Note: Replication is an add-on feature to HCP. If your HCP system doesn't include this feature and you would like to add it, please contact your HCP sales representative.

About replication

Replication, an HCP service, is the process of keeping selected tenants and namespaces in two HCP systems in sync with each other. This entails copying object creations, object deletions, metadata changes, and other information between the two systems. Typically, the two systems are in separate geographic locations and are connected by a high-speed wide area network.

A **replication topology** is a configuration of HCP systems wherein each system is related to each other system through a path of one or more replication relationships. The HCP system from which you're viewing the replication topology is called the **local system**. Any other system in the replication topology is called a **remote system**.

Replication has several purposes:

- If the local system becomes unavailable (for example, due to network issues), a remote system can provide continued data availability.

- If the local system suffers irreparable damage, a remote system 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 the content verification or protection service discovers objects it cannot repair on the local system, the service can use object data and metadata from a remote system to make the needed repairs. For more information on these services, see *Administering HCP*.

- If an object cannot be read from the local system (for example, because a node is unavailable), HCP can try to read it from a remote system. HCP tries to do this only if:

  - The namespace that contains the object is involved in replication.

  - The 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.

  - The replication networks for the local and remote systems can both use the same IP mode (IPv4 or IPv6).

  For more information on the read-from-remote-system feature, see *Managing a Tenant and Its Namespaces* or *Managing the Default Tenant and Namespace*.

- If the local system is unavailable, HTTP requests to that system can be automatically serviced by a remote system without the client needing to modify the target URL. The remote system can service a request only if:

  - The namespace named in the URL is replicated to the remote system

  - The namespace named in the URL is configured to accept requests redirected from other HCP systems

  - The HCP systems involved use a shared DNS for system addressing

  - The replication networks for the two systems can both use the same IP mode (IPv4 or IPv6)

  Additionally, unless you're using a load balancer that can target the requests to the remote system:

  - The DNS must be configured to support redirection between the two systems

  - The HCP system that's unavailable must have been configured to allow DNS failover to other systems

  For more information on the service-by-remote-system feature, see *Managing a Tenant and Its Namespaces* or *Managing the Default Tenant and Namespace*.

For information on:

- Replication networks, see "Selecting the network for replication" on page 132

- Configuring HCP in the DNS, see *Administering HCP*

- Configuring HCP to support DNS failover, see “Managing DNS failover” on page 136

How replication works

To start the replication service, you create a **replication link** between the local system and another HCP system. A replication link is a secure trust relationship between the two systems. It determines what is replicated and how data is transmitted between the systems.

Link types

A replication link can be an active/active or active/passive:

- With an **active/active link**, the HCP tenants and namespaces and default-namespace directories being replicated are read-write on both systems, and data is replicated in both directions between the two systems. Active/active links are designed for use in a cloud storage environment, where applications need seamless access to namespaces regardless of where those applications are located.

  With an active/active link, HCP gives preference to synchronizing object metadata between the two systems, followed by data synchronization. This preferential treatment enables the fastest access to objects regardless of which system they were created on. If the data for an object has not yet been replicated to the system targeted by a client request, HCP can read the data from a remote system. Additionally, synchronizing metadata first ensures that object names are reserved as quickly as possible.

  Note: Until the data for an object is replicated to a given system, the object is metadata-only on that system.

- With an **active/passive link**, the HCP tenants and namespaces and default-namespace directories being replicated are read-write on one system, called the **primary system**, and read-only on the other system, called the **replica**. In this case, data is replicated in one direction only — from the primary system to the replica. Active/passive links are designed to support disaster recovery.

  With an active/passive link, HCP replicates object data and metadata together, thereby ensuring that objects are intact on the replica and can be recovered as whole objects should that become necessary.

However, this approach to object replication means that the replica cannot service requests for any given object until that object is fully replicated.

To a primary system, an active/passive link is an **outbound link**; that is, the system is sending data through the link. To a replica, an active/passive link is an **inbound link**; that is, the system is receiving data through the link. To either system involved in an active/active link, the link functions as both an outbound link and an inbound link.

Link security

Replication relies on SSL to ensure the integrity and security of transmitted data. Before you can create a replication link, each system involved in the link must each have a valid SSL server certificate. Each system must also have the server certificate from the other system installed as a **trusted replication server certificate**.

Replication priority

Replication is asynchronous with other tenant and namespace activity on both systems involved in a link. An HCP system can, therefore, develop a backlog of objects to be replicated.

You can choose to have the replication service work on objects with the oldest changes first, regardless of which namespaces they're in. Alternatively, you can have the service balance its processing time evenly across the namespaces being replicated. In this case, the service may replicate recent changes in some namespaces before older changes in others.

What is replicated

Replication works by tenant. For each replication link, you select the HCP tenants to replicate on that link and, if you have administrative access to the tenant, which namespaces to include in replication of that tenant. For the default tenant, you select the directories to replicate.

Regardless of the system on which you're creating the link:

- When creating an active/active link, you can select items to replicate from both systems involved in the link

- When creating an active/passive link, you can select items to replicate only from the primary system

What is replicated for each tenant differs between HCP tenants and the default tenant.

Replication of HCP tenants

Part of the configuration of an HCP tenant is whether it's eligible to be replicated. A replication link can include only HCP tenants that are eligible to be replicated.

Replicated tenants and namespaces

When an HCP tenant is replicated, the replication service works only on selected namespaces owned by that tenant. The tenant administrator can select and deselect namespaces for replication at any time. If the tenant is being replicated, its selected namespaces are replicated along with it on each link that includes the tenant. If the tenant is not being replicated, the selected namespaces are not replicated either.

If an HCP tenant has granted system-level users administrative access to itself, you can select and deselect its namespaces for replication when you create the replication link or by modifying the link at a later time. You should coordinate namespace selection with the tenant administrator to ensure that you both have the same understanding of what should be replicated.

Tenants can be replicated from one HCP system to another regardless of the number of tenants that already exist in the other system. However, you can create a new tenant in a system only if the total number of tenants in the system is less than one thousand.

Similarly, namespaces can be replicated from one HCP system to another regardless of the number of namespaces that already exist in the other system. However, you can create a new namespace in a system only if the total number of namespaces in the system is less than ten thousand.

HCP tenants can be created in both of the systems involved in a replication link, so the two systems can have tenants with the same name. You cannot, however, replicate a tenant created in one of the systems if a locally created tenant with the same name already exists in the other system. To replicate the tenant in this case, you first need to change the name of one of the tenants involved.

Other replicated items

In addition to object creations, object deletions, and metadata changes, the replication service replicates these items for an HCP tenant:

- Old versions of objects in the namespaces being replicated

- Retention class creations, modifications, and deletions for the namespaces being replicated

- Content class and content property creations, modifications, and deletions
- All tenant-level log messages
- All namespace-level log messages for the namespaces that are being replicated
- The tenant configuration
- The configuration of each namespace being replicated
- User accounts defined for the tenant
- Group accounts defined for the tenant

Replicated changes

With an active/active link, replicated HCP tenants are read-write on both systems involved in the link. Administrators can make changes to the configuration of the replicated tenants and namespaces on either system, and users and applications can make changes to the content of the replicated namespaces on either system. All changes made on each system are replicated to the other system.

With an active/passive link, replicated HCP tenants are read-write on the primary system and read-only on the replica. Administrators cannot make any configuration changes to the replicated tenants or to any of their namespaces on the replica, nor can users and applications make any changes to the content of the replicated namespaces on the replica. However, all changes made in the primary system are replicated to the replica.

Note: Although you cannot make any changes to the configuration of a replicated namespace on a replica, you can reindex the namespace.

Replication of the default tenant

For the default tenant, the replication service works on one or more selected directories directly under `fcfs_data` in the default namespace. The service replicates the entire directory hierarchy under each selected directory.

Default namespace properties

Directories in the default namespace on one system involved in a replication link can be replicated only to the default namespace on the other system involved in the link. They cannot be replicated to HCP namespaces. Additionally:

- Before you can include any default-namespace directories in a replication link, the default namespace must exist on both systems involved in the link.
- The default namespaces in the two systems must have the same cryptographic hash algorithm and retention mode.
- If the default namespace on one system supports appendable objects, the default namespace on the other system should, too. For an active/active link, this ensures that users and applications can add data to replicated appendable objects on either system. For an active/passive link, this ensures that, in case of failover, users and applications can continue to add data to the replicated appendable objects on the replica. For information on failover, see "Failover and failback" on page 37.

Note: For an active/passive link, regardless of whether the default namespace on the replica supports appendable objects, when these objects are recovered to the primary system, they are still appendable.

The default namespace can be configured to prevent writes and/or deletes. This configuration can differ between the default namespace in the two systems involved in a link. However, even if the system receiving data is configured to prevent these actions, changes to namespace content on the sending system are replicated.

Replicated directories

You select directories for replication when you create the replication link. You can select additional directories or deselect directories by modifying the link at a later time. A replication link can include at most 256 directories.

Note: In releases of HCP earlier than 6.0, a replication link could include more than 256 directories. In a system that was originally upgraded from a release earlier than 6.0, the maximum number of directories that can be included in a preexisting link in which the system participates is equal to either 256 or the number of directories included in the link at the time of the upgrade to release 7.0, whichever is greater.

Directories can be created in the default namespace in both of the systems involved in a replication link, so the two systems can have directories with the same name. You cannot, however, replicate a directory created in one of the systems if a locally created directory with the same name already exists in the other system.

The one exception to this is the directory used for email archiving through SMTP. You can replicate the email directory if the email directory name is the same on both systems.

Replicated changes

With an active/active link, replicated directories are read-write on both systems involved in the link. All changes made to the content of those directories on either system are replicated to the other system.

With an active/passive link, the replicated directories are read-write on the primary system and read-only on the replica. Users and applications cannot make any changes to the content of the replicated directories on the replica. However, all changes made on the primary system are replicated to the replica.

Other replicated items

In addition to object creations, object deletions, and metadata changes, the replication service replicates:

- Retention class creations, modifications, and deletions for the default namespace
- Content class and content property creations, modifications, and deletions
- All log messages relating to retention classes and privileged delete operations

Replication service processing

With an active/active link, the replication service runs on both HCP systems involved in the link. With an active/passive link, the replication service runs only on the primary system. The replication service starts on the applicable HCP systems or system the first time you add HCP tenants and/or default-namespace directories to a replication link that involves that system.

When you first add HCP tenants and namespaces to a replication link, the replication service immediately replicates the configuration of those items. The service then begins replicating objects in those namespaces and in the selected default-namespace directories. The service starts with the objects with the oldest metadata changes either across all namespaces or within each namespace, depending on the link configuration.

Whenever a new item is added to a replication link or a change is made to the configuration of an item already included in the link, the replication service immediately replicates the configuration of the new item or the configuration change. This ensures the correct behavior of objects on the receiving system.

Replicated systems

Replication occurs between two separate HCP systems, each of which is complete in its own right. Because replication is a software function, the two systems can have entirely different hardware configurations, including differing amounts of storage.

The two systems in a replicated pair are connected through the front-end network infrastructure, as shown in the figure below. Replication traffic from each system must be routable to the network selected for replication on the other system.

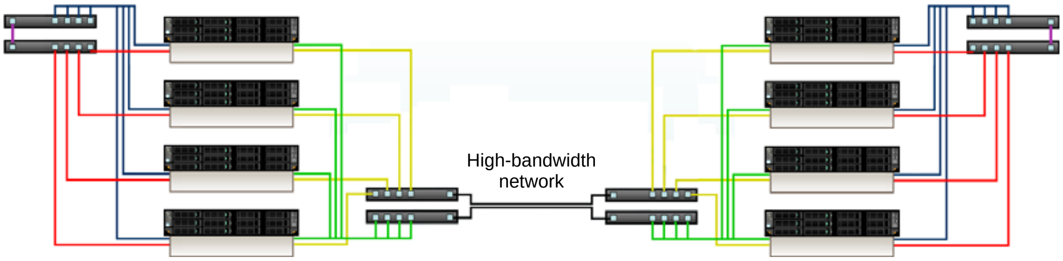

Replication is a network bandwidth-intensive process. To ensure sufficient front-end network capacity when replicating, you need to give extra consideration to planning the front-end infrastructure.

For information on selecting a network for replication, see “Selecting the network for replication” on page 132.

Replication and Active Directory authentication

An HCP system can be configured to support Windows® Active Directory® (AD) for user authentication. Part of this configuration is the specification of an AD domain. The domain determines the AD groups from which HCP group accounts can be created.

If HCP is configured to support AD, HCP tenants can be configured to allow access by users authenticated by AD. For this access to work, the AD user must belong to one or more AD groups for which corresponding group accounts are defined for the tenant.

For the same AD users to be able to access a given tenant on both systems involved in a replication link, the group accounts on each system must correspond to the same AD groups as they do on the other system. To make this happen, support for AD must be enabled on both systems, and either of these must be true:

- The same domain is specified in the AD configuration on both systems.
- The domain specified in the AD configuration on one system is trusted by the domain specified in the AD configuration on the other system.

Similarly, for the same AD users to be able to access the default namespace on two systems involved in a replication link, the system-level group accounts on each system must correspond to the same AD groups as they do on the other system.

Replication and RADIUS authentication

An HCP system can be configured to support RADIUS for user authentication. If HCP is configured this way, HCP tenants can be configured to allow access by RADIUS-authenticated users. Also, if HCP is configured this way, RADIUS-authenticated users with system-level user accounts can access the default namespace.

For the same RADIUS-authenticated users to have access to both systems involved in a replication link, the HCP RADIUS configuration on the two systems must specify the same RADIUS servers.

DPL, service plans, and cross-release replication

In release 6.*x* of HCP, each namespace has a data protection level (DPL) setting that applies to each object in the namespace. This setting determines the number of copies of each object HCP should maintain in the repository.

In release 7.0 of HCP, namespaces no longer have a DPL setting. Instead, the service plan that applies to a namespace specifies the number of copies of each object that should exist at any given time during the life of the object.

When a namespace is replicated from a release 7.0 primary system to a release 6.*x* replica, the namespace DPL is set to **Dynamic** on the 6.*x* replica, regardless of the definition of the service plan on the 7.0 system.

When a release 6.*x* system is upgraded to release 7.0, HCP redefines existing service plans and automatically generates new service plans as needed to accommodate all combinations of:

- The existing service plans that are directly associated with namespaces

- The DPL settings of those namespaces

HCP associates a new service plan with each of those namespaces, as necessary, except if the namespace is being replicated to the upgraded system from a 6.*x* system. In that case, the service plan that was associated with the namespace before the upgrade remains associated with the namespace after the upgrade, even if the service plan definition has changed.

For more information on service plans, see *Administering HCP*.

Replication topologies

An HCP system can have up to five outbound links and up to five inbound links. An active/active link counts as both an outbound link and an inbound link. The ability for a single system to participate in multiple links enables HCP to support various replication topologies.

Replication can occur:

- In two directions on a single link between two HCP systems (active/active replication)

- In one direction on a single link between two HCP systems (active/passive replication)
- From multiple HCP systems to a single HCP system (many-to-one replication)
- From one HCP system to a second HCP system and from that second system to a third HCP system, such that the same HCP tenants and namespaces and default-namespace directories that are replicated to the second system are then replicated to the third system (chained replication)
- From one HCP system to multiple other HCP systems (one-to-many replication)

These configurations can be combined to form complex replication topologies.

Note: A replication topology can include one or more HCP systems in which the occurrences of a given namespace contain metadata-only objects. If, in this topology, the namespace contains fully hydrated objects on only one system and the topology is being used for production purposes, the topology should include a disaster recovery system in which the namespace also contains fully hydrated objects.

For information on metadata-only objects, see *Administering HCP*.

Simple active/active replication

In an **simple active/active replication topology**, two HCP systems replicate the same HCP tenants and namespaces and default-namespace directories to each other over an active/active link. The items being replicated can be originally created on either system. The items are read-write on the system in which they were created and, after being replicated, are also read-write on the other system.

With active/active replication, client requests can be directed to either system. All changes made on each system, including both configuration changes and changes to namespace content, are replicated to the other system.

Note: HCP supports active/active links only between systems that are both at release 7.0 or higher.

What this looks like

The figure below shows a simple active/active replication topology in which two HCP systems (A and B) are replicating to each other.

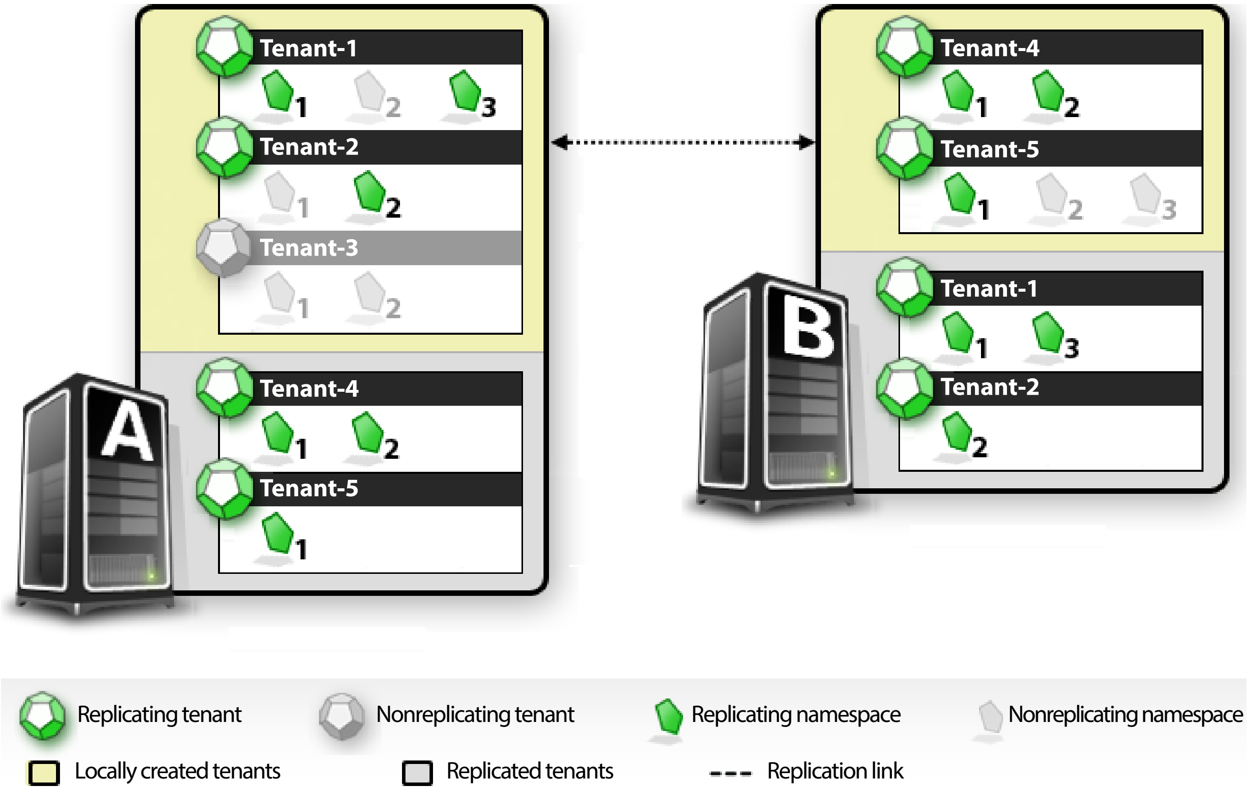

In this figure:

- Two of three HCP tenants created locally in system A are being replicated between system A and system B. In the first tenant being replicated, two of three namespaces are selected for replication. In the second tenant being replicated, one of two namespaces is selected for replication.

- Two tenants created locally in system B are being replicated between system A and system B. In the first tenant, two namespaces are selected for replication. In the second tenant, one of three namespaces is selected for replication.

Uses

The active/active replication topology supports a cloud storage model, where any type of client request can be serviced equally by either HCP system involved in the replication link. With this topology, if one system becomes unavailable (for example, either unexpectedly or for scheduled maintenance), the other system can still provide all the required functionality.

With an active/active replication topology, the processing load can be distributed between the two systems. This distribution can be accomplished in two ways:

- By having some applications direct requests to one system and other applications direct requests to the other system. Typically, the assignment of applications to systems is based on geographic location.

  With this configuration, a single shared DNS can enable requests to be redirected automatically from one system to the other in case one of the systems becomes unavailable. For more information on this, see “Managing DNS failover” on page 136.

- By using a load balancer to divide requests between the two systems. For each client request, the decision as to which system to target can be based on several factors, such as the current load on each system and the geographical distance between the client and each system.

  With this configuration, if one system becomes unavailable, the load balancer simply directs all requests to the other system.

Active/active replication ring

In an active/active replication ring topology, multiple HCP systems are connected by active/active links where:

- Each system has exactly two active/active links that connect that system to two other systems in such a way that all the systems are linked in the form of a ring

- The same HCP tenants and namespaces and default-namespace directories are replicated on each link

An active/active replication ring can include any number of HCP systems. Changes made on one system travel around the ring in both directions from that system until they reach a system in which those changes have already occurred.

To configure an active/active replication ring topology:

1. Create an active/active link between each pair of systems in the ring.

2. For each system that has locally created tenants, namespaces, and directories you want to replicate, add those items to one of the links connected to that system.

3. After those tenants, namespaces, and directories have been replicated, add them to the second link that's connected to the system to which the items have been replicated.

4. After the items have been replicated to that system, add them to the next link in the ring.

5. Continue adding the replicated items to successive links until you have added them to the second link that connects to the system on which the items were originally created.

What this looks like

The figure below shows an active/active replication ring topology that includes three systems (A, B, C, and D).

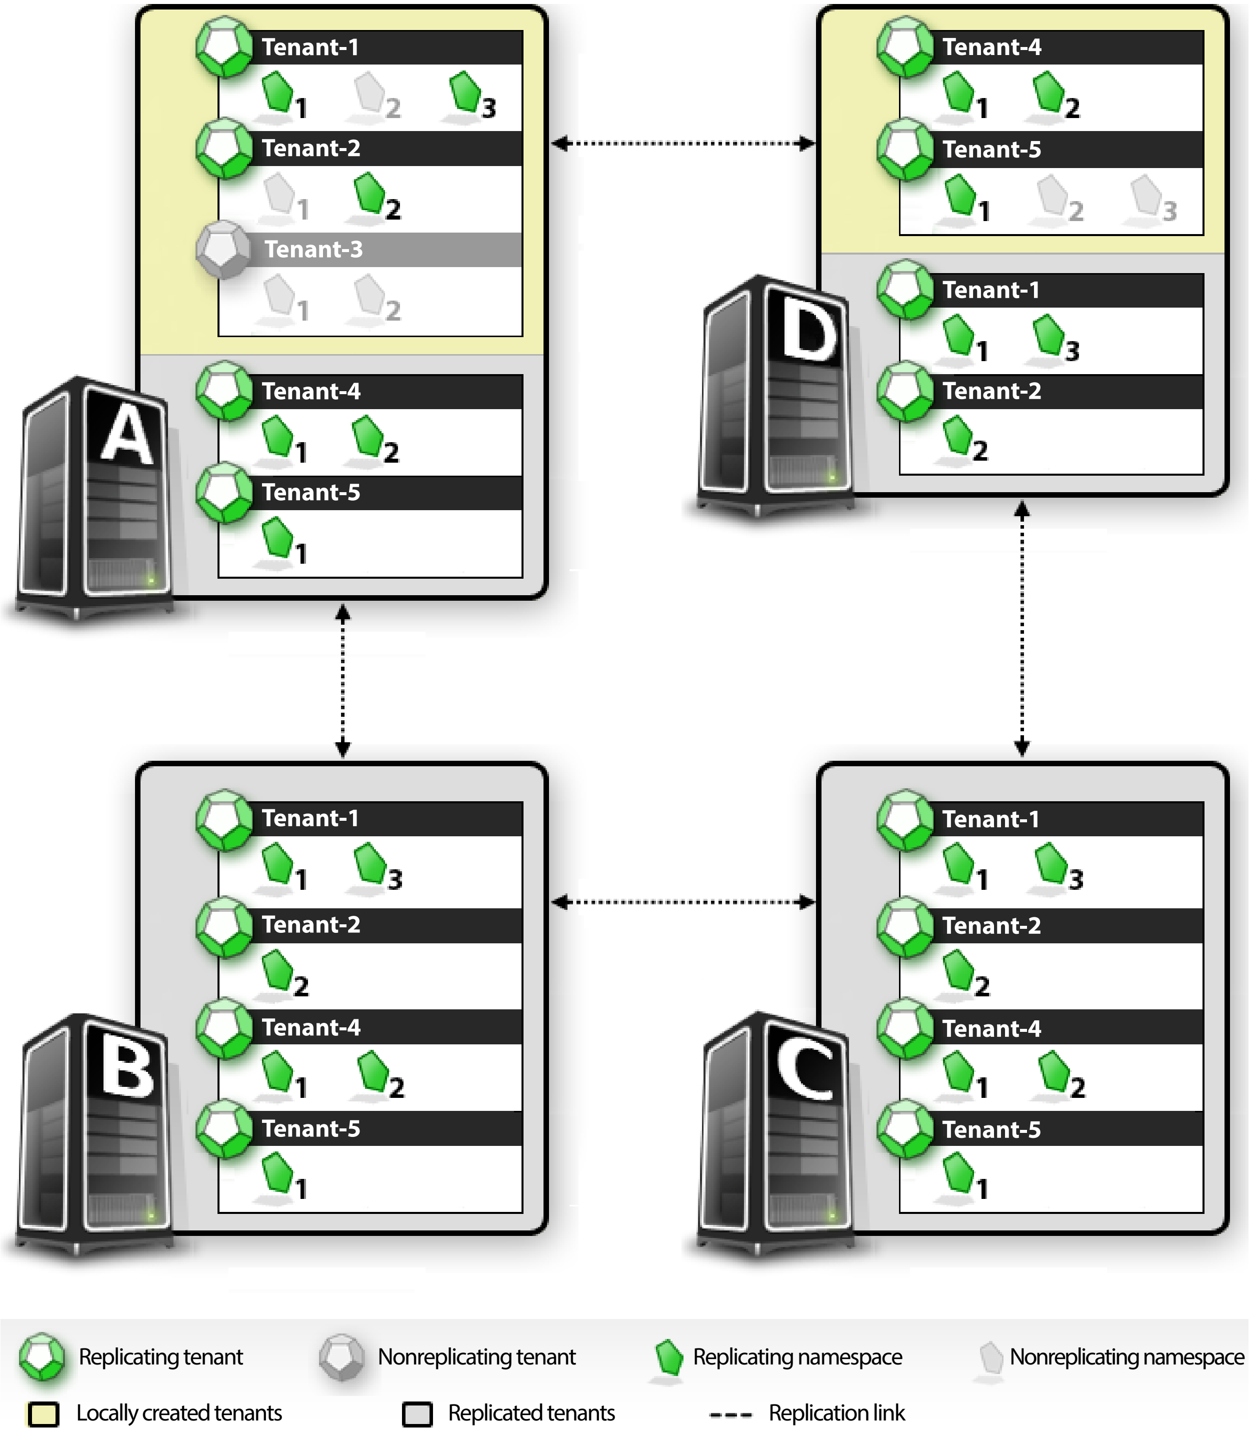

In this figure:

- Two of three HCP tenants created locally in system A are being replicated around the ring. In the first tenant being replicated, two of three namespaces are selected for replication. In the second tenant being replicated, one of two namespaces is selected for replication.

- Two tenants created locally in system D are being replicated around the ring. In the first tenant, two namespaces are selected for replication. In the second tenant, one of three namespaces is selected for replication.

Uses

The active/active replication ring topology supports increasingly large storage clouds. The HCP systems included in the ring are typically in geographically disperse locations, with each system providing the same functionality as all the others. Changes made on any one system are propagated to all the other systems, effectively enabling shared read-write access to the same HCP tenants and namespaces and default-namespace directories regardless of where the client applications are located.

New HCP systems can be added to an active/active replication ring topology at any time. To add a new system (for example, system X) to the topology:

1. Create a new active/active link between any system already in the topology (for example, system B) and system X.

2. Add the tenants, namespaces, and directories already being replicated around the ring to the new link.

3. Create a new link between system X and one of the other systems with which system B participates in a link (for example, system C).

4. Add the tenants, namespaces, and directories already being replicated around the ring to that new link.

5. Delete the link between system B and system C.

Simple active/passive replication

In a **simple active/passive replication topology**, one HCP system, the primary system, replicates HCP tenants and namespaces and default-namespace directories to another HCP system, the replica, over an active/passive link. The items being replicated are originally created only in the primary system. The items are read-write on the primary system and, after being replicated, read-only on the replica.

What this looks like

The figure below shows a simple active/passive replication topology in which one HCP system (A) is replicating to a second HCP system (B).

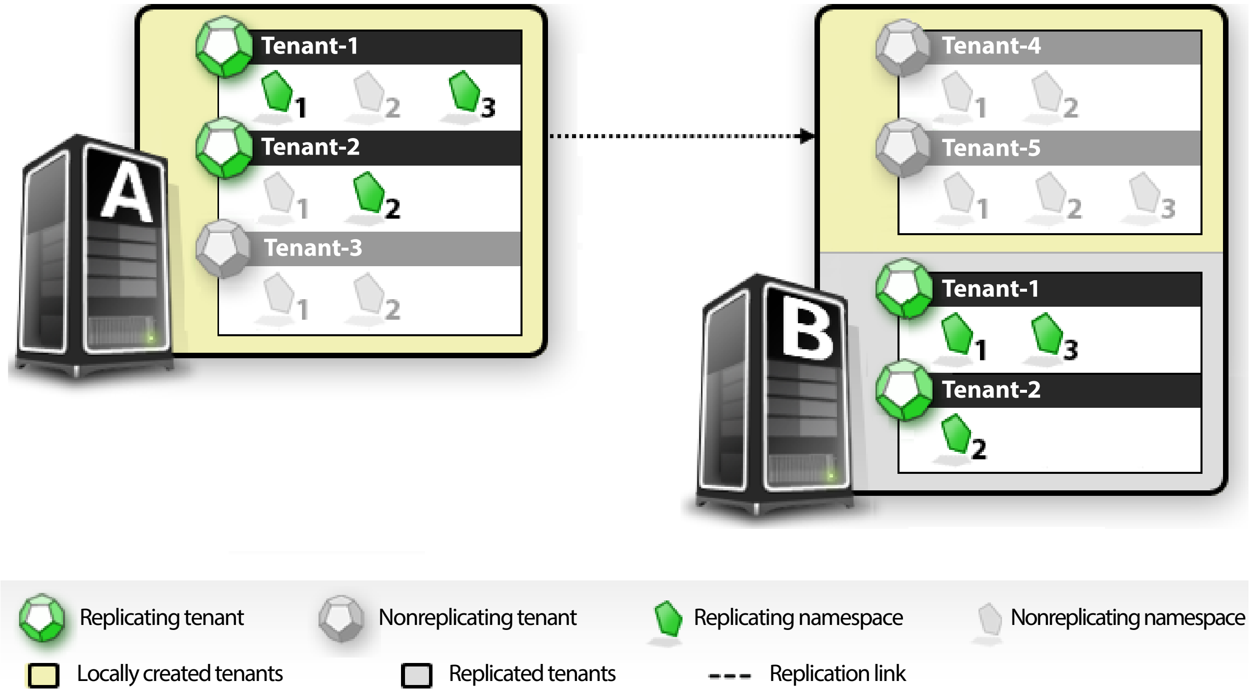

In this figure:

- From system A, two of three locally created tenants are being replicated to system B. In the first tenant being replicated, two of three namespaces are selected for replication. In the second tenant being replicated, one of two namespaces is selected for replication.
- In system B, no locally created tenants are being replicated.

Uses

The main purpose of a simple active/passive replication topology is to have a backup of the primary system for use in disaster recovery. However, this topology also supports a scenario in which some applications need read-write access to the replicated items while other applications need only read access. By directing requests from the latter to the replica, you can reduce the load on the primary system.

Active/passive many-to-one replication

In an **active/passive many-to-one replication topology**, multiple HCP systems replicate to a single other HCP system. Each of the replicating systems is the primary system for an active/passive replication link. The other system is the replica for each of those links.

In this topology, you can think of the replica as being a hub, with the primary systems being at the ends of spokes connected to that hub. The spokes themselves are the replication links. The hub can have at most five spokes.

Each HCP tenant and default-namespace directory selected for replication on a primary system must be unique on the replica. For example, if two of the primary systems each have a tenant named Finance, the Finance tenant can be replicated from only one of the systems. For both tenants to be replicated, they must have different names.

If you want to replicate the email directory in the default namespace on multiple primary systems, the SMTP protocol on each system should specify a different email directory name.

Note: You can also create a many-to-one replication topology that consists of a combination of active/passive and active/active links or even only active/active links.

What this looks like

The figure below shows an active/passive many-to-one replication topology in which three primary systems (A, B, and C) are replicating to a single replica (D).

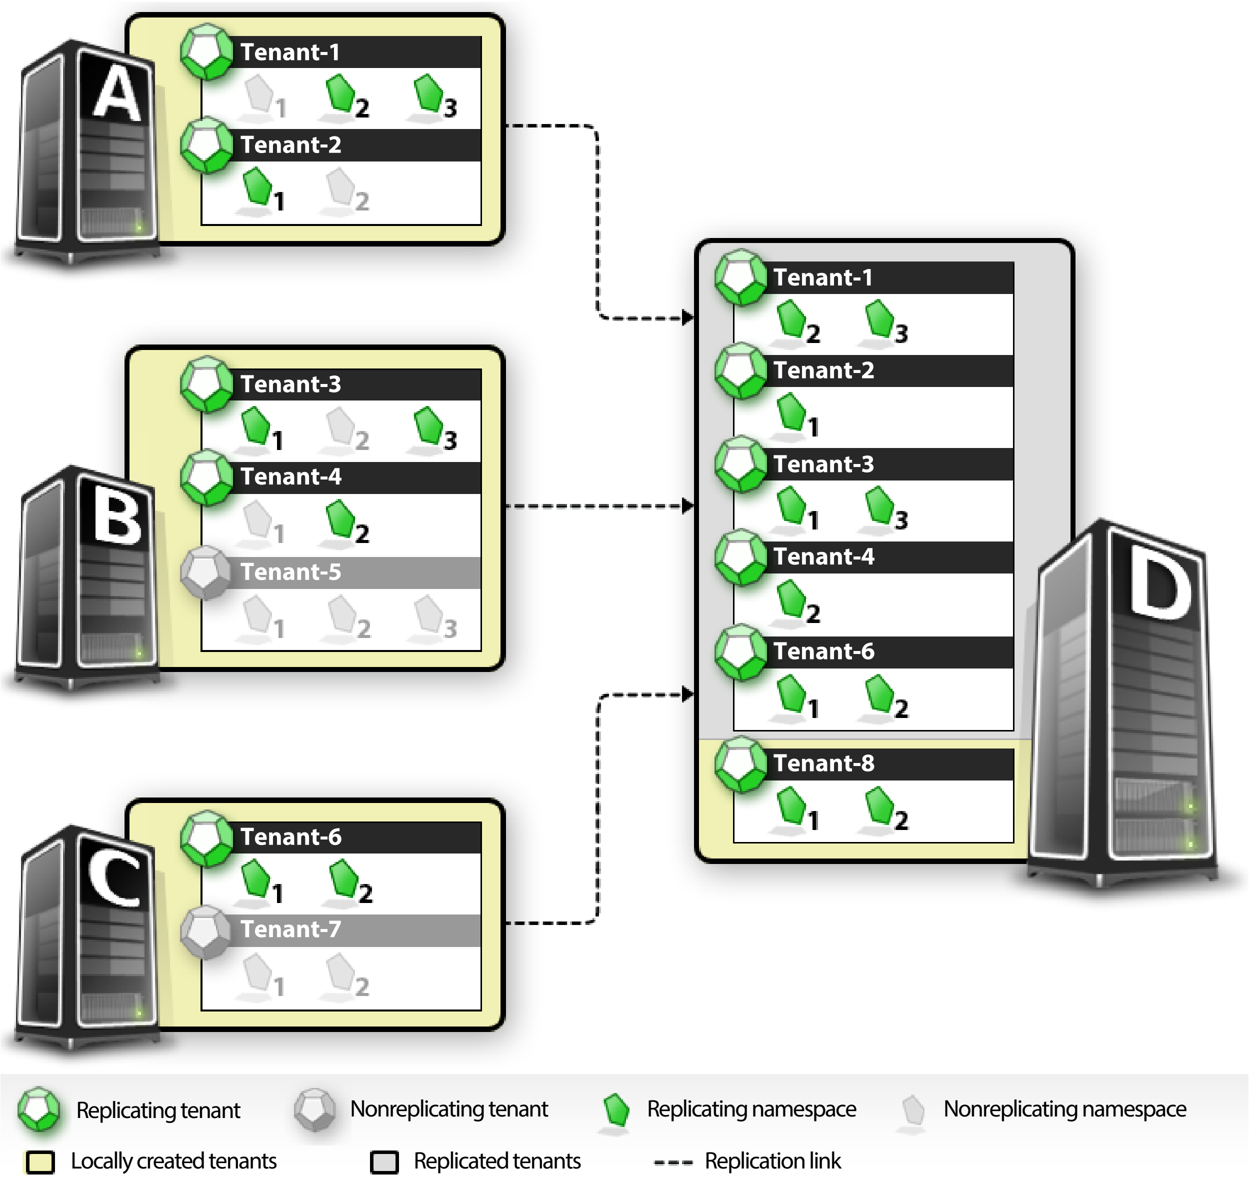

In this figure:

- From system A, two HCP tenants are being replicated to system D. In the first tenant, two of three namespaces are selected for replication. In the second tenant, one of two namespaces is selected for replication.

- From system B, two of three HCP tenants are being replicated to system D. In the first tenant being replicated, two of three namespaces are selected for replication. In the second tenant being replicated, one of two namespaces is selected for replication.

- From system C, one of two HCP tenants is being replicated to system D. In that tenant, both namespaces are selected for replication.

- System D has one HCP tenant of its own that is not the result of replication. This tenant has two namespaces.

Uses

In an active/passive many-to-one replication topology, the replica is typically the largest HCP system. The data center in which it resides is usually fully staffed with IT personnel. The primary systems are typically smaller. Because replication links can be created and managed from either side, the data centers where these systems reside can be minimally staffed with less experienced people.

Active/passive many-to-one replication supports a scenario in which the outlying sites are branch offices of an enterprise that has an HCP system at a primary data center. Applications running at each of those offices connect over a local area network (LAN) to an HCP system at the local data center. Because they are close to the storage they use, the applications perform better than they would if they were using a wide area network (WAN) to connect to the HCP system at the hub. Also, the distribution of the network load among the outlying sites further enhances application performance.

With active/passive many-to-one replication, you can consolidate similar data from multiple sources in a single HCP system. Using a search application, you can then perform federated queries across the replicated namespaces on that system. For example, suppose each of the branch offices of an enterprise stores completed orders in a namespace in its own HCP system. If each of those namespaces is replicated to a single HCP system, you could query that system to find all orders placed for a specific product at any of the branch offices.

Active/passive many-to-one replication enables a single HCP system to be used for backup and recovery of multiple other HCP systems, so the outlying sites don’t need backup systems of their own. For additional data protection, the hub system can be backed up, for example, to tape or to another HCP system (see “Many-to-one replication with disaster recovery support” on page 29).

For information on data recovery in an active/passive many-to-one replication topology, see “Failover and failback in an active/passive many-to-one topology” on page 153.

Active/passive chained replication

In an **active/passive chained replication topology**, HCP tenants and namespaces and default-namespace directories in one HCP system (A) are replicated to a second HCP system (B) and then from the second system to a third HCP system (C). Both the replication link from A to B and the replication link from B to C are active/passive links.

To configure an active/passive chained replication topology:

1. When you create the replication link from A to B, select the tenants, namespaces, and directories you want to replicate. In this link, A is the primary system and B is the replica.

2. When you create the replication link from B to C, select the link from A to B to be included in that link. This causes the HCP tenants, namespaces, and directories replicated on the link from A to B to be replicated again from B to C. In this link, B is the primary system and C is the replica.

   The link from B to C can also include tenants, namespaces, and directories that were originally created on B.

   You cannot individually select the replicated tenants, namespaces, and directories on B for inclusion in the link from B to C.

In an active/passive chained replication topology with the configuration A ▸ B ▸ C, B is the replica for the link from A to B. Therefore, in B, the tenants and directories that were replicated from A are read-only, even though B is also the primary system for the link from B to C.

HCP supports replication chains that consist of three HCP systems. Longer chains are not supported.

Notes:

- You can also create a chained replication topology in which the first link is active/active and the second link is active/passive. You cannot, however, create a chained replication topology in which the first link is active/passive and the second link is active/active or in which both links are active/active.
- In a replication chain A ▸ B ▸ C in which the objects in a namespace on system B are supposed to be metadata-only, HCP does not remove data from those objects until that data has been replicated to system C. For information on metadata-only objects, see *Administering HCP*.

What this looks like

The figure below shows an active/passive chained replication topology in which system A is replicating to system B and system B is replicating to system C.

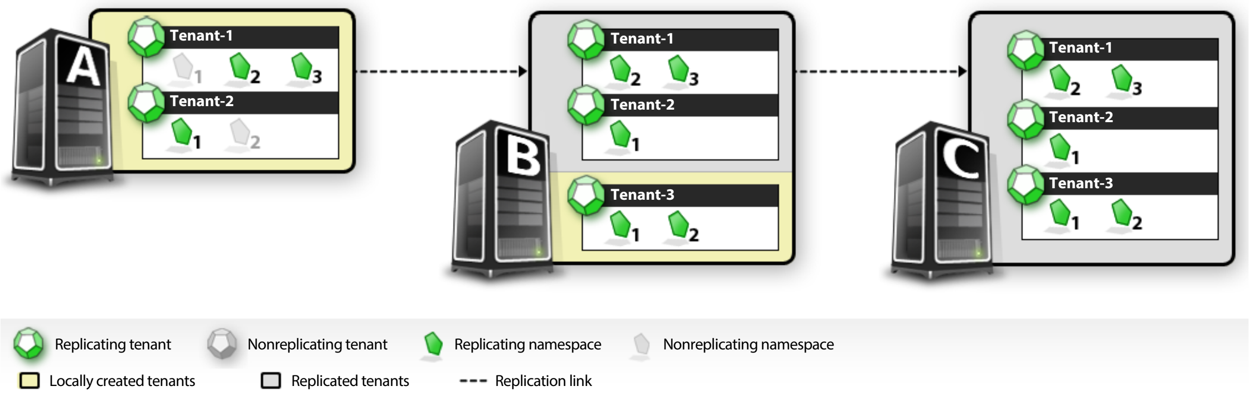

In this figure:

- From system A, two HCP tenants are being replicated to system B. In the first tenant, two of three namespaces are selected for replication. In the second tenant, one of two namespaces is selected for replication.
- From system B, the tenants and namespaces created by replication from system A are being replicated to system C because the link from A to B is included in the link from B to C.
- One tenant that was originally created on system B is also being replicated to system C. Both the namespaces in that tenant are selected for replication.

Uses

In an active/passive chained replication topology, all three HCP systems are typically located at full-scale data centers. For example, suppose a corporation has three major locations — New York, Los Angeles, and Tokyo. Each of these locations runs an application that is vital to the corporation, but only New York generates the data for these applications. Replicating the data from New York to Los Angeles and from Los Angeles to Tokyo enables applications at each location to access the data through a local area network. Because those applications are not accessing the data on the HCP system in New York, this topology also reduces the load on that system.

Active/passive chained replication helps ensure continuous data availability. If any one or even two of the HCP systems in the chain become unavailable, applications still have access to the stored data.

In an active/passive replication chain, the third HCP system provides disaster recovery functionality for the second system, and the second system provides disaster recovery functionality for the first system. For information on data recovery in an active/passive chained replication topology, see “Failover and failback in an active/passive chained topology” on page 154.

In an active/passive chained replication topology, you can create and manage both replication links from the HCP system in the middle of the chain. To this system, the link from the first system is an inbound link, and the link to the third system is an outbound link.

Active/passive one-to-many replication

In an **active/passive one-to-many replication topology**, one HCP system replicates to two or more other HCP systems using separate active/passive replication links. The replicating system is the primary system for each link. The other systems are replicas.

Typically, in an active/passive one-to-many replication topology, the replication links include different HCP tenants and namespaces and default-namespace directories.

If three systems need to have the same tenants, namespaces, and directories as each other, consider using a chained topology for them rather than a one-to-many topology. The one-to-many topology puts greater load on the single primary system than does a chained topology, which splits the load between the first and second systems in the chain.

In some cases, however, replicating the same tenants and namespaces to two different replicas in a one-to-many topology may be appropriate. For an example of this, see the description of the second use case in "Uses" on page 27.

Note: You can also create a one-to-many replication topology that consists of a combination of active/passive and active/active links or even only active/active links.

What this looks like

The figure below shows an active/passive one-to-many replication topology in which one primary system (A) replicates to two replicas (B and C) over two separate replication links.

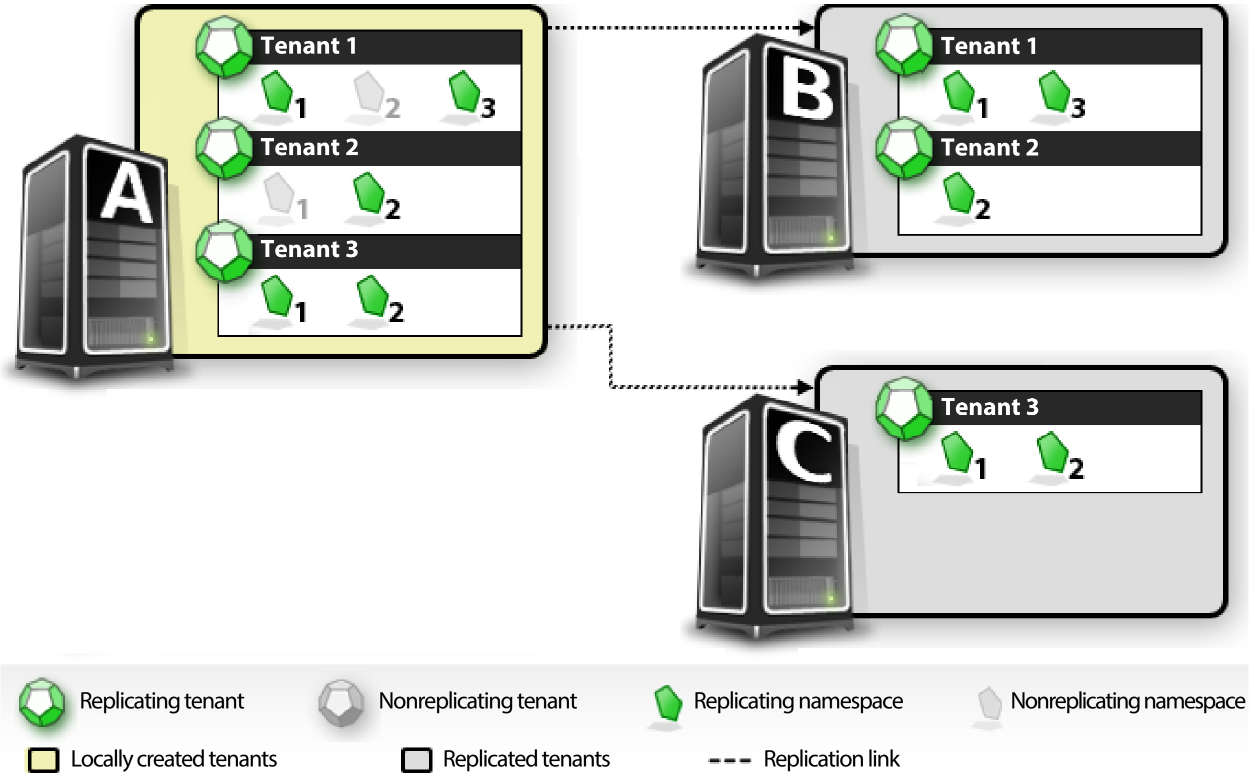

In this figure:

- Two of three HCP tenants in system A are being replicated to system B. In the first tenant being replicated, two of three namespaces are selected for replication. In the second tenant being replicated, one of two namespaces is selected for replication.

- The third HCP tenant in system A is being replicated to system C. This tenant has two namespaces, both of which are selected for replication.

Uses

An active/passive one-to-many replication topology enables you to replicate different tenants and default-namespace directories from a single primary system to multiple different replicas. You may choose to do this, for example, if the replicas are at sites where users need read access to only a limited number of namespaces. Those sites can then have HCP systems that use smaller amounts of storage.

You can also use an active/passive one-to-many replication topology as a means to upgrade your HCP systems from RAIN to SAIN. This scenario assumes you currently have a RAIN primary system replicating to a RAIN replica on an active/passive link. You want to have a SAIN primary system replicating to a SAIN replica. To make this happen:

1. Create a second active/passive link from the RAIN primary system to the SAIN system that will be the primary system after the upgrade. Use this link to replicate all the HCP tenants, HCP namespaces, and default-namespace directories in the RAIN system.

2. Create a replication link from the new SAIN primary system to the second SAIN system. Include the link created in step 1 above in this link to create a replication chain.

3. When replication is completely up to date, fail over the link from the RAIN system to the SAIN system. To enable replication to become completely up to date, you may need to stop clients from writing to the RAIN system before you fail over the link.

4. Delete the link from the RAIN system to the SAIN system.

5. Decommission the two RAIN systems.

For information on failover, see “Failover and failback” on page 37. For information on data recovery in an active/passive one-to-many replication topology, see “Failover and failback in an active/passive one-to-many topology” on page 159.

Active/active replication with disaster recovery support

You can combine active/active replication with active/passive replication to create a topology in which a third system provides disaster recovery support for the two systems that participate in an active/active link. In this topology, a set of HCP tenants and namespaces and default-namespace directories is replicated on an active/active link between two systems (A and B). Some of those items are also replicated on an active/passive link from A to a third system (C). The rest of the items are also replicated on an active/passive link from system B to system C.

The choice of which items to replicate on each active/passive link is independent of which system the items were initially created on. For example, an HCP tenant created locally on system A could be included in the active/passive link from B to C.

The figure below shows a simple active/active replication topology with disaster recovery support.

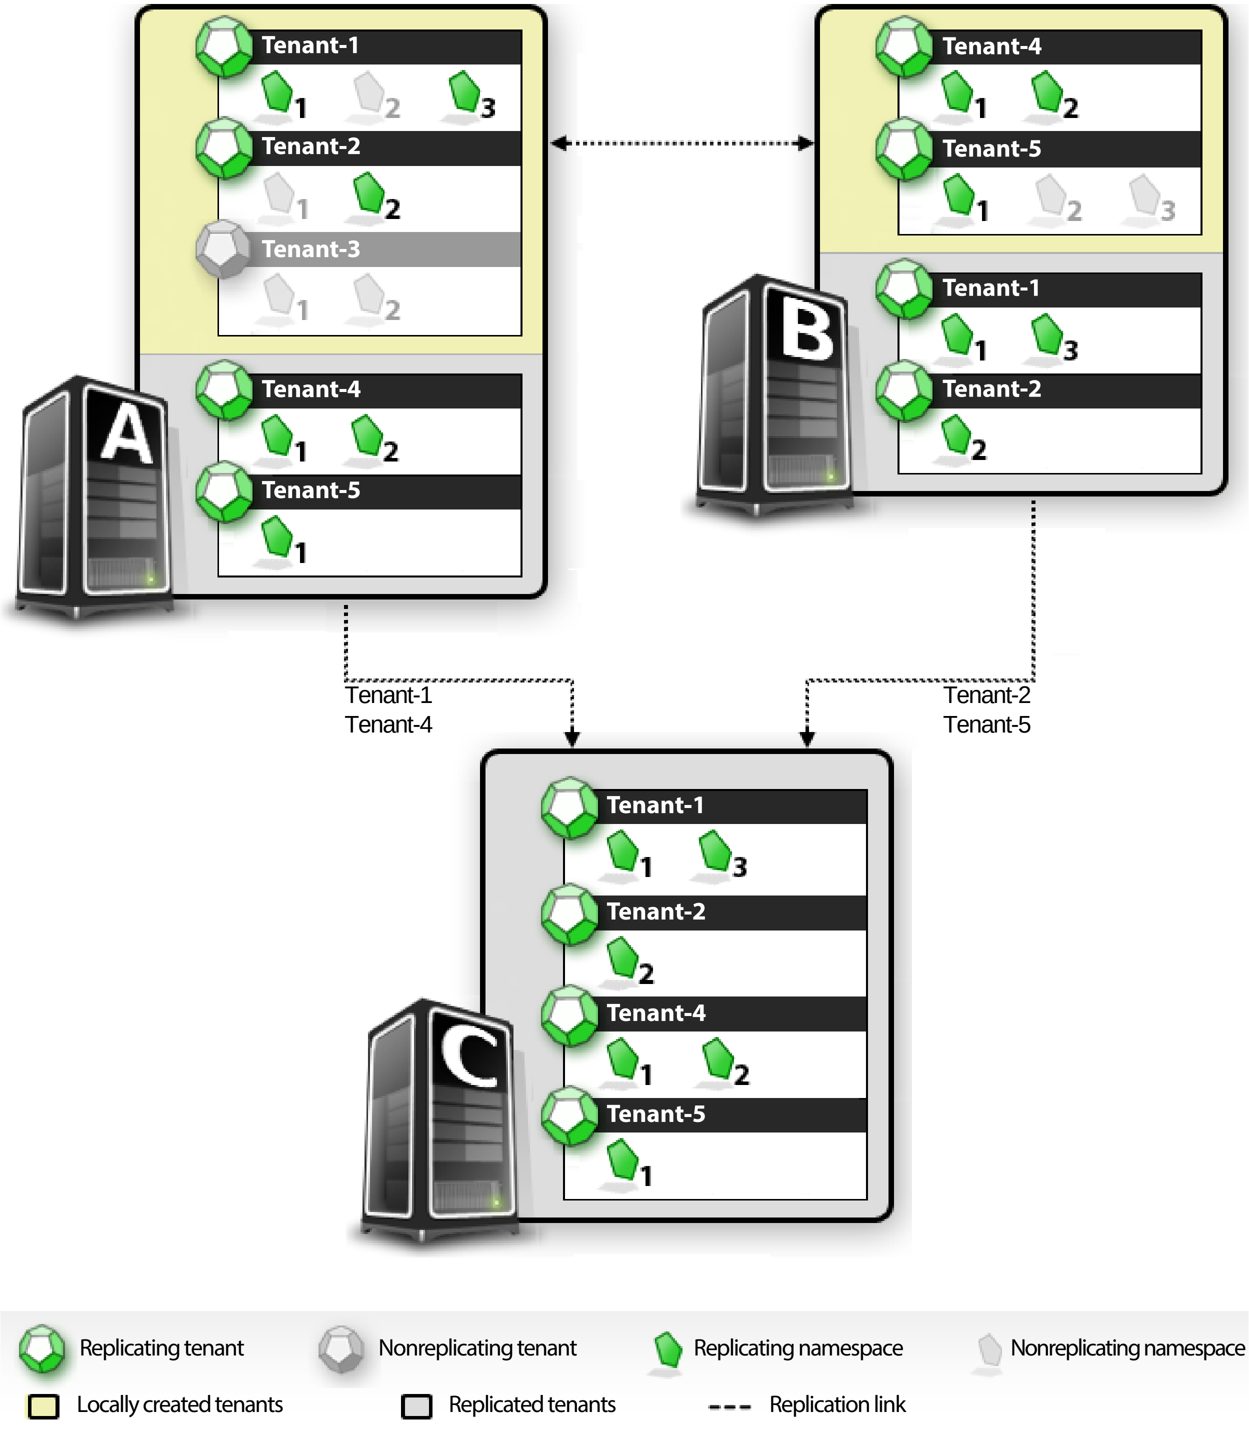

In this figure:

- Two of three HCP tenants created locally in system A are being replicated between system A and system B. In the first tenant being replicated, two of three namespaces are selected for replication. In the second tenant being replicated, one of two namespaces is selected for replication.

- Two HCP tenants created locally in system B are being replicated between system A and system B. In the first tenant, two namespaces are selected for replication. In the second tenant, one of three namespaces is selected for replication.

- From system A, one HCP tenant initially created on system A and one HCP tenant initially created on system B are being replicated to system C.

- From system B, one HCP tenant initially created on system A and one HCP tenant initially created on system B are being replicated to system C.

Many-to-one replication with disaster recovery support

You can combine the many-to-one and chained replication topologies to create a configuration in which multiple HCP systems replicate to a single HCP system (many-to-one), which, in turn, replicates to another HCP system (chained). The last system provides additional assurance of continuous data availability for all the HCP tenants and namespaces and default-namespace directories originally selected for replication. It also provides disaster recovery functionality for the HCP system to which those tenants, namespaces, and directories are initially replicated.

For example, suppose HCP systems A, B, and C replicate to system D, which replicates to system E. To create this combined topology:

1. When you create the replication links from A to D, from B to D, and from C to D, select the tenants, namespaces, and directories you want to replicate.

2. When you create the replication link from D to E, select the links from A to D, from B to D, and from C to D to be included in that link. This causes the tenants, namespaces, and directories replicated on all three of those links to be replicated again from D to E.

In effect, you’re creating three replication chains: A ▸ D ▸ E, B ▸ D ▸ E, and C ▸ D ▸ E.

The figure below shows a many-to-one replication topology with disaster recovery support.

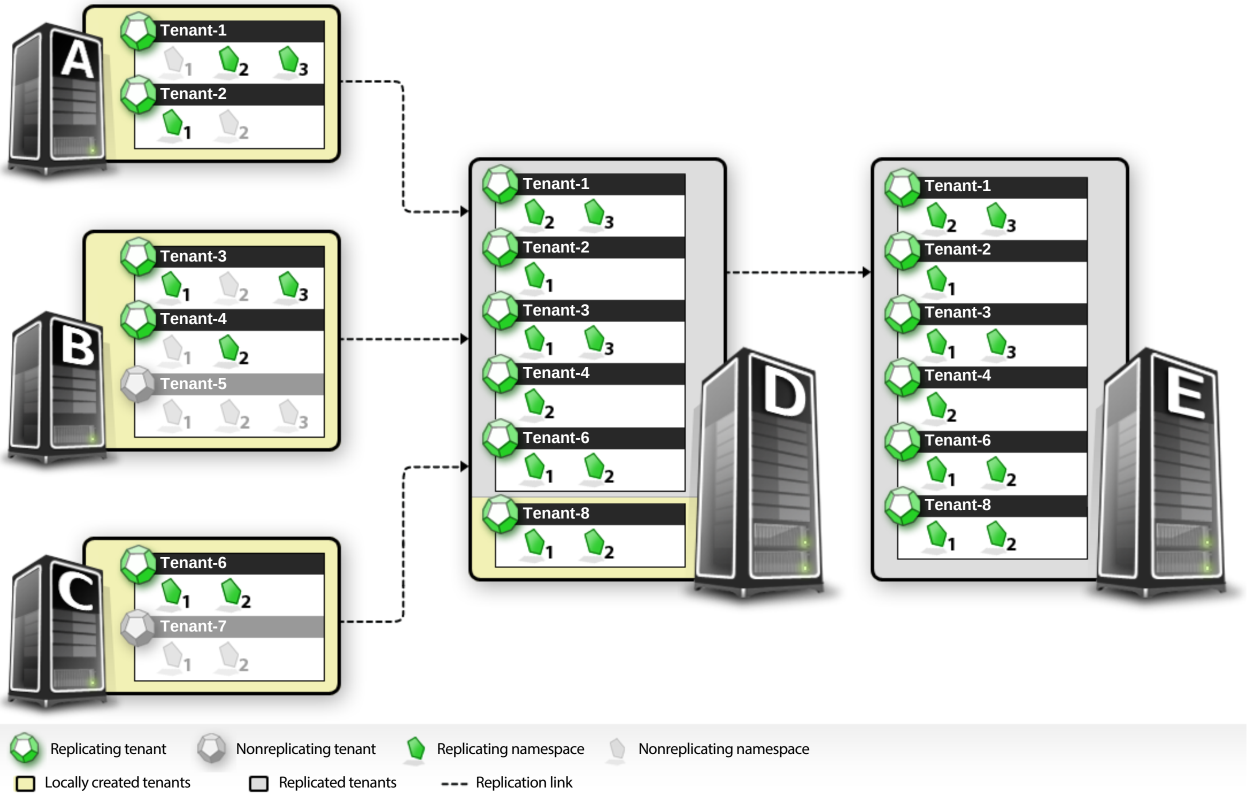

For information on data recovery in an active/passive many-to-one replication topology with disaster recovery support, see "Failover and failback in an active/passive many-to-one topology with disaster recovery support" on page 160.

Bidirectional active/passive replication

For any given pair of HCP systems, each system can serve as both a primary system and a replica for the other system, as long as different HCP tenants and different default-namespace directories are being replicated in each direction. A system that's serving as both a primary system and a replica at the same time has two active/passive links — one outbound and one inbound. This topology is called **bidirectional active/passive replication**.

What this looks like

The figure below shows a bidirectional active/passive replication topology in which two systems (A and B) are each replicating to the other system.

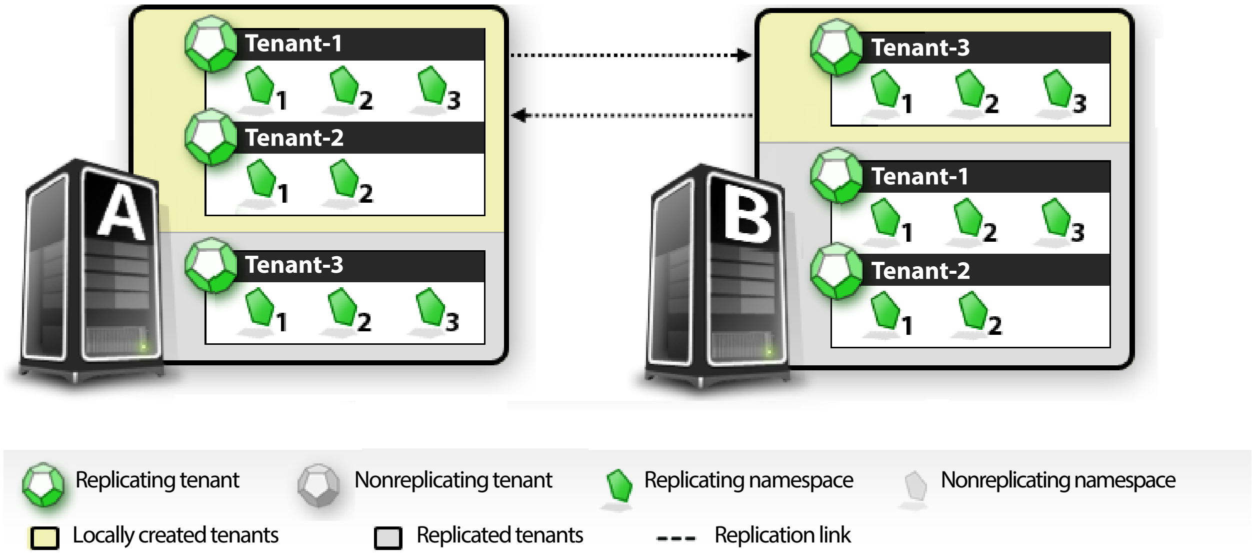

In this figure:

- From system A, two locally created tenants are being replicated to system B. In both tenants, all namespaces are selected for replication.

- From system B, one locally created tenant is being replicated to system A. In this tenant, all namespaces are selected for replication.

Uses

For each link in a bidirectional active/passive replication topology, the tenants, namespaces, and directories being replicated are read-write on the primary system and read-only on the replica. A bidirectional active/passive replication topology thus supports a scenario like this:

- Application 1 running in a data center in New York needs read-write access to Tenant-1, which was created locally in the HCP system (A) in New York.

- Application 2 running in a data center in Tokyo needs read-write access to Tenant-2, which was created locally in the HCP system (B) in Tokyo.

- Application 1 needs only read access to Tenant-2.

- Application 2 needs only read access to Tenant-1

To meet these needs, you could:

- Create an active/passive link from system A to system B that includes Tenant-1. System A sends Tenant-1 to system B, where the tenant is in read-only mode. Application 2 can access the tenant more efficiently on system B than on system A.

- Create an active/passive link from system B to system A that includes Tenant-2. System B sends Tenant-2 to system A, where the tenant is in read-only mode. Application 1 can access the tenant more efficiently on system A than on system B.

Bidirectional active/passive chained replication

In a **bidirectional active/passive chained replication** topology, three HCP systems participate in two active/passive replication chains, with each chain going in a different direction. That is, given three HCP systems A, B, and C:

- In one chain, system A replicates to system B, which replicates to system C.

- In the other chain, system C replicates to system B, which replicates to system A.

What this looks like

The figure below shows a bidirectional active/passive chained replication topology in which the replication chains are A ▸ B ▸ C and C ▸ B ▸ A.

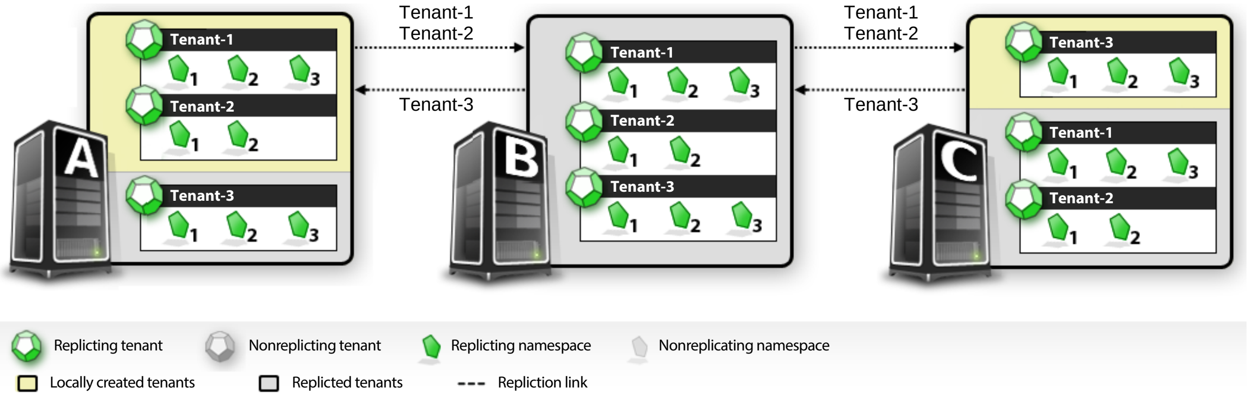

In this figure:

- From system A, Tenant-1 and Tenant-2 are being replicated to system B. In both tenants, all namespaces are selected for replication.

- From system C, Tenant-3 is being replicated to system B. This tenant has three namespaces, all of which are selected for replication.

- From system B:

  - The tenants and namespaces created by replication from system A (that is, Tenant-1 and Tenant-2 and their namespaces) are being replicated to system C because the link from A to B is included in the link from B to C.

  - The tenant and namespaces created by replication from system C (that is, Tenant-3 and its namespaces) are being replicated to system A because the link from C to B in included in the link from B to A.

Uses

A bidirectional active/passive chained replication topology allows clients to write to selected HCP tenants and namespaces and default-namespace directories on three different systems while reading from all those tenants, namespaces, and directories on all those systems. For example, suppose you have data centers in New York, Los Angeles, and Tokyo, and:

- The HCP system in New York has a locally created HCP tenant named Tenant-1

- The HCP system in Los Angeles has a locally created HCP tenant named Tenant-2

- The HCP system in Tokyo has a locally created HCP tenant named Tenant-3

- At each data center, a locally running application needs write access to the namespaces owned by the locally created tenant

- Clients at all three locations need read access to the namespaces owned by all three HCP tenants

To meet these needs, you could:

- Create an active/passive link (Link-1) from the New York system to the Los Angeles system that includes Tenant-1, which is writable in New York. The New York system sends Tenant-1 to the Los Angeles system, where the tenant is in read-only mode.

- Create an active/passive link (Link-2) from the Los Angeles system to the Tokyo system that includes Link-1 and Tenant-2, which is writable in Los Angeles. The Los Angeles system sends both Tenant-1 and Tenant-2 to the Tokyo system, where the tenants are in read-only mode.

- Create an active/passive (Link-3) from the Tokyo system to the Los Angeles system that includes Tenant-3, which is writable in Tokyo. The Tokyo system sends Tenant-3 to the Los Angeles system, where the tenant is in read-only mode.

- Create an active/passive link (Link-4) from the Los Angeles system to the New York system that includes Link-3 and Tenant-2. The Los Angeles system sends both Tenant-3 and Tenant-2 to the New York system, where the tenants are in read-only mode.

Bidirectional active/passive replication with disaster recovery support

You can combine the bidirectional and one-to-many replication topologies to create a configuration in which each of two systems replicates its locally created HCP tenants and namespaces and default-namespace directories to both the other system and a third system. This third system then provides each of the first two systems with both additional assurance of continuous data availability and disaster recovery functionality.

For example, suppose HCP systems A and B each replicate their locally created HCP tenants and namespaces and default-namespace directories to each other, and you want a third system, C, to provide disaster recovery functionality for each of those systems. To create this combined topology:

- Create a link from A to B that includes all the HCP tenants and namespaces and default-namespace directories that were originally created on system A.

- Create a link from A to C that includes the same tenants, namespaces, and directories as the link from A to B.

- Create a link from B to A that includes all the HCP tenants and namespaces and default-namespace directories that were originally created on system B.

- Create a link from B to C that includes the same tenants, namespaces, and directories as the link from B to A.

The figure below shows a bidirectional replication topology with disaster recovery support.

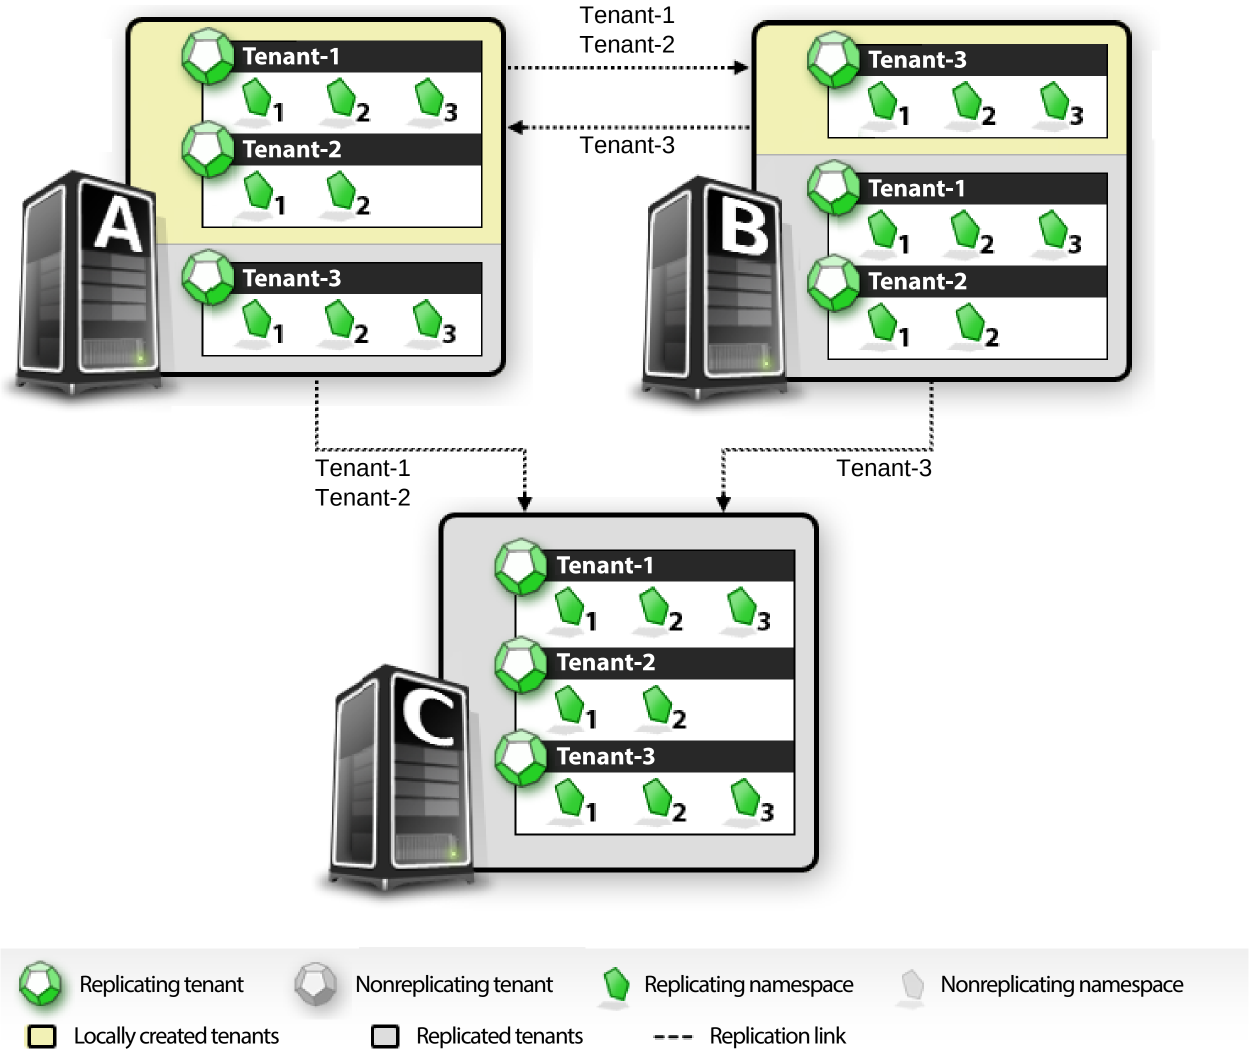

You can combine the bidirectional and chained replication topologies to have the same result as the topology described above, which involves the bidirectional and one-to-many topologies. That is, the replicated items in the link from A to B also end up on C, and the replicated items in the link from B to A also end up on C.

In this case, however, instead of the link from A to C replicating the same items as the link from A to B, the link from A to C would include the link from B to A. Likewise, instead of the link from B to C replicating the same items as the link from B to A, the link from B to C would include the link from A to B.

The figure below shows this topology.

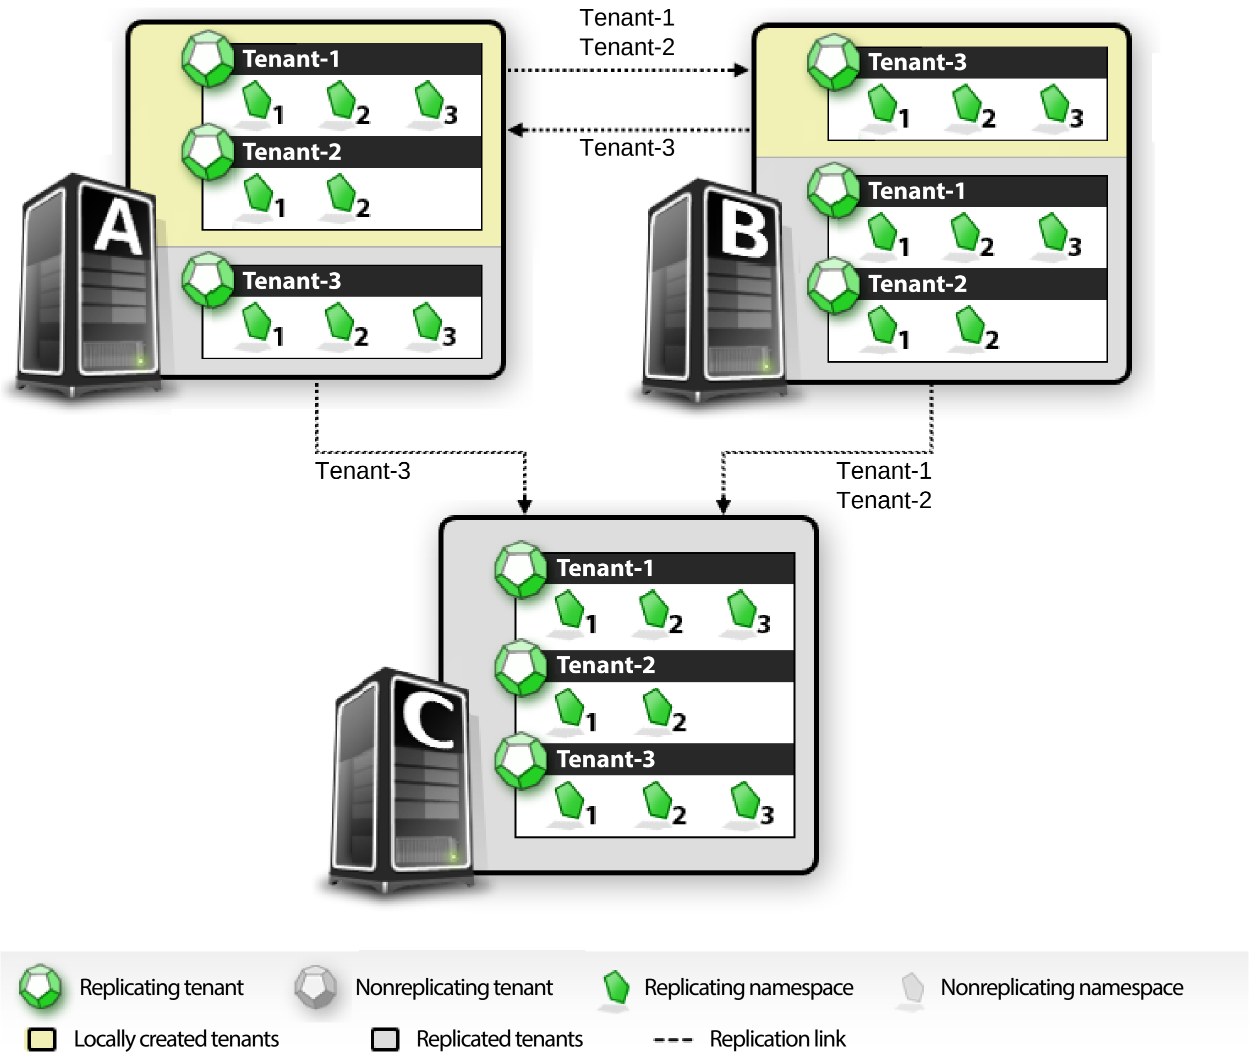

This topology provides better performance than does the first topology described above. Additionally, if you need to add a tenant to the link from A to B, for example, you don't need to also add it to the link from A to C. Instead, the tenant is automatically replicated to C because the link from A to B is included in the link from B to C.

On the other hand, the first topology described above is better for assuring continuous data availability in the case of the failure of system A or B. For example, if system A is replicating to systems B and C and system B fails, the items being replicated from system A are still making it to system C. In the second topology, if system B fails, the items from system A are not replicated to system C because the replication chain is broken.

Failover and failback

Failover is a process that stops replication on a link and results in a situation in which, for read-write access:

- For an active/active link, applications should use only the HCP system to which the link was failed over

- For an active/passive link, applications should use only the replica

Typically, you fail over a link when one of the systems involved in the link becomes unavailable. With an active/passive link, this system must be the primary system. You don't need to fail over the link when the replica fails.

You can fail over a link while both systems are available. You might do this, for example, if you need to shut down one of the systems for maintenance.

Depending on the link configuration, failover either is a manual procedure or occurs automatically. When automatic failover is enabled for a link, the link fails over automatically after the applicable system is unavailable for a specified amount of time.

You enable or disable automatic failover separately for each system involved in an active/active link. For an active/passive link, you enable or disable automatic failover only for the replica.

Failback is the process that restarts replication on a link that has been failed over and returns the HCP systems involved in the link to normal operation. Typically, you fail back a link when an unavailable system becomes available again.

If connectivity was lost between the two systems involved in a failed-over link, before failback can occur, connectivity must be restored. Connectivity exists when the network infrastructure through which the two systems communicate is healthy and the applicable SSL server certificates have been shared between the two systems.

In a disaster recovery situation in which the system that became unavailable has been rebuilt, the link no longer exists on that system. In this case, before failback can occur, the link must be restored to the rebuilt system.

With an active/active link, failback is a manual procedure. With an active/passive link, the failback procedure can be partially automated.

Note: Automatic failover and failback are not supported for cross-release links (that is, links between 7.0 system and 6.*x* systems).

For information on sharing SSL server certificates, see Chapter 2, "Configuring SSL for replication," on page 57. For instructions on enabling automatic failover and failback, see "Changing automatic failover and failback settings" on page 93. For instructions on performing failover and failback, see Chapter 6, "Managing failover and failback," on page 143.

Failover and failback with active/active links

The effects of failing over and failing back an active/active link differ depending on whether DNS failover was enabled for the system that became unavailable. In all cases, however, when the link fails over, replication on that link stops. When the link fails back, normal replication restarts.

Failover with an active/active link

With an active/active link, failover can occur in either direction between the two systems involved in the link. While the link is failed over, the replicated HCP tenants and namespaces and default-namespace directories remain read-write on both systems. However, because failover normally occurs when one system is unavailable, to avoid wasting resources, neither system tries to read or repair objects from the other system.

With DNS failover enabled, when an active/active link fails over from one of the HCP systems involved in the link (system A) to the other system involved in the link (system B), system B broadcasts a new configuration to the DNS. This new configuration causes client requests targeted to system A by domain name to be redirected to system B when the request is for an HCP namespace or default-namespace directory that's included in the failed-over link.

Note: System B can service redirected namespace access requests only if the applicable namespace is configured to allow service by remote systems.

If a client request targeted to system A is for an HCP namespace or default-namespace directory that is not included in the failed-over link, the request is not redirected to system B. Client requests that target system A by IP address are also not redirected to system B.

Client requests that use a domain name to target the Tenant Management Console for a replicated HCP tenant on system A are redirected to system B, but system B cannot process such requests. Instead, system B returns a 403 error code.

While the link is failed over, system A does not broadcast any configuration information to the DNS.

With DNS failover disabled, failing over an active/active link stops replication on the link but does not cause any other changes. Clients can still access system A by domain name (if system A is available).

For more information on DNS failover, see "Managing DNS failover" on page 136.

Failback with an active/active link

Failing back an active/active link entails a single action, fail back. When an active/active link fails back:

- Replication immediately restarts in both directions on the link.
- With DNS failover enabled, each HCP system involved in the link broadcasts its own configuration to the DNS. From that point on, client requests that target either system by domain name are directed to the specified system.

Failover and failback with active/passive links

The effects of failing over and failing back an active/passive link differ depending on whether DNS failover was enabled for the system that became unavailable. In all cases, however, when an active/passive link fails over, replication on that link stops. When the link fails back, normal replication restarts.

Failover with an active/passive link

Failover with an active/passive link occurs only from the primary system to the replica. When an active/passive link fails over to the replica, the replicated HCP tenants and namespaces and default-namespace directories become read-write on the replica.

With DNS failover enabled, when an active/passive link fails over, the replica broadcasts a new configuration to the DNS. This new configuration causes client requests targeted to the primary system by domain name to be redirected to the replica when the request is for an HCP namespace or default-namespace directory that's included in the failed-over link.

Note: A replica can service redirected namespace access requests only if the applicable namespace is configured to allow service by remote systems.

If a client request targeted to the primary system is for an HCP namespace or default-namespace directory that is not included in the failed-over link, the request is not redirected to the replica. Client requests that target the primary system by IP address are also not redirected to the replica.

Client requests that use a domain name to target the Tenant Management Console for a replicated HCP tenant on the primary system are redirected to the replica, but the replica cannot process such requests. Instead, the replica returns a 403 error code.

If the primary system is still available when an active/passive link is failed over and the primary system and the replica can communicate with each other, the replicated HCP tenants and namespaces and default-namespace directories become read-only on the primary system. Also, while the link is failed over, the primary system does not broadcast any configuration information to the DNS.

If the two systems cannot communicate with each other when the link is failed over, the replicated items remain read-write on the primary system, and the primary system continues to broadcast its configuration information to the DNS. If DNS failover is disabled, clients can still access the primary system by domain name.

If clients are allowed to write to both the primary system and the replica while an active/passive link is failed-over, configuration conflicts and conflicts in namespace content may occur when the link is failed back. Although HCP resolves such conflicts in a predictable manner, the recommended practice is to avoid them in the first place. Therefore, when you fail over an active/passive link without DNS failover enabled, you should tell the applicable tenant administrators to direct all client access requests to the replica.

For information on how HCP resolves conflicts resulting from writing to both systems involved in a failed-over active/passive link, see "Replication collisions" on page 41. For more information on DNS failover, see "Managing DNS failover" on page 136.

Failover with bidirectional active/passive links

With bidirectional active/passive links, failover is a independent process for each of the two links. If one of the HCP systems involved in the links becomes unavailable, failover needs to occur only on the link for which that system is the primary system. The link for which that system is the replica does not need to be failed over, and the status of the HCP tenants and namespaces and default-namespace directories included in that link does not change.

Failback with an active/passive link

Failing back an active/passive link has two phases, begin recovery and complete recovery. The begin recovery phase is always started manually. The complete recovery phase can be started manually or automatically.

When recovery begins on a link, the replica starts sending configuration changes and changes to namespace content back to the primary system. The replicated HCP tenants and namespaces and default-namespace directories remain read-write on the replica and read-only on the primary system.

The complete recovery phase is designed to allow data recovery to catch up to the current time before normal replication restarts. At the beginning of this phase, the replicated items change to read-only on the replica and remain read-only on the primary system. The replica continues to send data to the primary system until the two HCP systems involved in the link are in sync with each other. Within a minute after that point:

- Normal replication restarts on the link.
- With DNS failover enabled, each system involved in the link broadcasts its own configuration to the DNS. From that point on, client requests that target either system by domain name are directed to the specified system.

Before starting the complete recovery phase manually, you should wait until data recovery is caught up with the current time. When you configure automatic failback for an active/passive link, you specify how up to date data recovery must be before the complete recovery phase begins.

Replication collisions

With an active/active link, clients can make configuration changes and changes to namespace content on both HCP systems involved in the link. This situation can result in configuration and content collisions.

With an active/passive link, if clients are allowed to make changes on both HCP systems involved in the link while the link is failed over to the replica, configuration and content collisions can occur on failback. Collisions can also occur on failback if changes made on the replica while a link is failed over conflict with configuration or content that was not yet replicated at the time of failover.

The way HCP handles collisions that occur due to replication depends on the type of collision. The general rule for namespace content is that more recent changes have priority over conflicting less recent changes. If conflicting changes occur at exactly the same time, HCP gives priority to the change that occurred on the system on which the link was created.

When configuration changes result in a conflict, the replication service pauses replication or recovery of the tenant, as applicable. For more information on conflicting configuration changes, see "Automatically paused tenant replication or recovery" on page 123.

Note: If the system time is not in sync on the two systems involved in a replication link, replication collision handling may have unexpected results.

Object content collisions

An object content collision occurs when:

- For an HCP namespace without versioning enabled, these events occur in the order shown:

  1. An object is created with the same name in that namespace on both systems involved in a replication link, but the object has different content on the two systems.

  2. The object on one of the systems is replicated to the other system.

  If versioning is enabled for the namespace, no collision occurs. Instead, the less recently created of the two objects becomes an old version of the more recently created object. If the two objects were created at exactly the same time, the object that was created on the system on which the link was created is treated as the newer version.

- For a default-namespace directory, these events occur in the order shown:

    1. An object is created with the same name in that directory on both systems involved in a replication link, 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 in the same namespace or renamed, depending on the namespace configuration. If the two objects were created at exactly the same time, the object that was created on the system on which the link was created keeps its name and location.

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

When renaming an object due to a content collision, HCP changes the object name to *`object-name`*`.collision` or *`object-name`*`.`*`version-id`*`.collision`, where *`version-id`* is the version ID of the object. HCP uses the second format only if versioning has ever been enabled for the namespace that contains the object but is not currently enabled.

If the new name is already in use, HCP changes the object name to *`object-name`*`.1.collision` or *`object-name`*`.`*`version-id`*`.1.collision`, as applicable. 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.

Namespaces can be configured to have the disposition service automatically delete objects that are flagged as replication collisions. When selecting this option for a namespace, the tenant administrator specifies 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 *Using a Namespace* or *Using the Default Namespace*. 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 the two systems involved in a replication link.

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) involved in a replication link to the other system (system B) involved in the link:

- 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) involved in a replication link to the other system (system B) involved in the link.

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	The changes on system B are replicated to system A. The resulting settings for the object on system A are: Retention: Deletion Prohibited Hold: true Shred: true Index: false

Custom metadata collisions

A custom metadata collision occurs when:

- For an HCP namespace, these events occur in the order shown:

  1. One of these changes occurs:

     – An annotation is added with the same name to a given object on both systems involved in a replication link, but the annotation has different content on the two systems.

       The addition of an annotation to a given object on only one of the systems does not result in a custom metadata collision if the object does not have an annotation with the same name on the other system. In this case, the new annotation is replicated without conflict.

     – Different changes are made to the content of a given annotation for a given object on each of the two systems involved in a replication link.

     – A change is made to the content of a given annotation for a given object on one of the systems involved in a replication link, and the same annotation is deleted on the other system involved in the link.

  2. The change made on one of the systems is replicated to the other system.

- For a default-namespace directory, these events occur in the order shown:

  1. One of these changes occurs:

     - Custom metadata is added to a given object on both systems involved in a replication link, 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 involved in a replication link does not result in a custom metadata collision. Instead, the new custom metadata is replicated from that system to the other system involved in the link without conflict.

     - The custom metadata for a given object is replaced on both systems involved in a replication link, but the replacement custom metadata is different on the two systems.

     - The custom metadata for a given object is replaced on one of the systems involved in a replication link, and the same custom metadata is deleted on the other system involved in the link.

  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) involved in a replication link to the other system (system B) involved in the link:

- 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

Here are some examples of how HCP handles collisions when custom metadata changes for a given object are replicated from one system (system A) involved in a replication link to the other system (system B) involved in the link.

Example 1

In an HCP namespace, the object starts out with annotations named a1 and a2 on both system A and system B.

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

Sequence	Event
1	On system B, a client changes the content of a1.
2	On system A, a client makes a different change to the content of a1.
3	On system A, a client adds annotation a3 to the object.
4	On system B, a client adds annotation a3 with different content from the a3 added on system A.
5	The changes on system A are replicated to system B. The resulting annotations for the object on system B are: a1 with the changed content from system A a2 (unchanged) a3 with the content added on system B
6	The changes on system B are replicated to system A. The resulting annotations for the object on system A are: a1 with the changed content from system A a2 (unchanged) a3 with the content added on system B

Example 2

In an HCP namespace, the object starts out with the annotations named a1, a2, and a3 on both system A and system B.

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

Sequence	Event
1	On system B, a client changes the content of a1.
2	On system A, a client deletes a1.
3	On system A, a client changes the content of a2.
4	On system B, a client changes the content of a2.
5	On system A, a client deletes a3.
6	On system B, a client changes the content of a3.
7	The changes on system A are replicated to system B. The resulting annotations for the object on system B are: a2 with the changed content from system B a3 with the changed content from system B

(Continued)

Sequence	Event
8	The changes on system B are replicated to system A, the resulting annotations for the object on system A are: a2 with the changed content from system B a3 with the changed content from system B

Example 3

In a default-namespace directory, the object starts out with 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.

Access control list collisions

An access control list (ACL) collision occurs when these events occur in the order shown:

1. Different changes are made to the ACL for a given object on each of the two systems involved in a replication link.

2. The changed ACL on one of the systems is replicated to the other system.

An ACL is treated as a single unit. If a collision occurs when a changed ACL for a given object is replicated from one system (system A) involved in a replication link to the other system (system B) involved in the link:

- If the last change to the ACL on system A is more recent than the last change to the ACL on system B, HCP changes the ACL on system B to match the changed ACL on system A

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

For example, suppose the ACL for a given object starts out with these grants on both system A and system B:

All users: read
User lgreen: write
User mwhite: write, delete

The table below shows a sequence of events in which the ACL for the object is changed and the change is then replicated.

Sequence	Event
1	On system B, a client changes the grants in the ACL to: All users: read User lgreen: write, delete User mwhite: write, delete, read ACL
2	On system A, a client changes the grants in the ACL to: All users: read User mwhite: write User pdgrey: write
3	The changed ACL on system A is replicated to system B. The resulting ACL for the object on system B contains these grants: All users: read User mwhite: write User pdgrey: write

Configuration collisions

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

1. Different changes are made to the same configuration property on each of the two systems involved in a replication link.

2. The changed property on one of the systems is replicated to the other system.

Examples of configuration properties are:

The namespace quota for a tenant
The data access permission mask for a tenant
The versioning setting for a namespace
The default shred setting for a namespace
The roles for a user account
The data access permissions a group account has for a namespace
The protocol optimization setting on a namespace

Certain groups of properties are treated as a single unit. Generally, these groups consist of properties that are updated by a single submission in the System or Tenant Management Console. Two notable exceptions to this rule are data access permissions for user accounts and content properties for content classes. In these cases, each set of data access permissions for a namespace and each content property is treated as an individual property.

If a collision occurs when a configuration change is replicated from one system (system A) involved in a replication link to the other system (system B) involved in the link:

- If the last change to the configuration on system A is more recent than the last change to the configuration on system B, HCP changes the configuration on system B to match the configuration on system A

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

The rules above apply to all configuration collisions except collisions that occur when retention class properties are changed. For information on how HCP handles this type of collision, see "Retention class collisions" below.

Here are two examples of how HCP handles collisions when configuration changes are replicated from one system (system A) involved in a replication link to the other system (system B) involved in the link.

Example 1

A given tenant starts out on both system A and system B with these properties:

Namespace quota: 5
Versioning: disabled

The table below shows a sequence of events in which the tenant configuration is changed and the change is then replicated.

Sequence	Event
1	In the System Management Console on system B, an administrator changes the namespace quota for the tenant to ten.
2	In the System Management Console on system A, an administrator enables versioning for the tenant.
3	The change on system A is replicated to system B. Because namespace quota and versioning are properties in the same submission group, the resulting properties for the tenant on system B are: Namespace quota: 5 Versioning: enabled

Example 2

A given tenant-level user account starts out on both system A and system B with these roles and data access permissions:

Roles: monitor, compliance
Namespace-1 permissions: browse, read, write, delete

The table below shows a sequence of events in which the user account is changed and the changes are then replicated.

Sequence	Event
1	In the Tenant Management Console on system B, a security administrator adds the administrator role to the user account.
2	In the Tenant Management Console on system A, a security administrator removes the monitor role from the account.
3	In the Tenant Management Console on system B, an administrator removes the write permission from Namespace-1 and adds the purge permission.
4	In the Tenant Management Console on system A, an administrator adds the privileged permission to Namespace-1.
5	In the Tenant Management Console on system B, an administrator gives the user account browse and read permissions for Namespace-2.
6	In the Tenant Management Console on system A, an administrator gives the user account browse, read, and write permissions for Namespace-3.

(Continued)

Sequence	Event
7	The changes on system A are replicated to system B. Because roles and data access permissions are in separate submission groups, the resulting roles and data access permissions for the user account on system B are: Roles: compliance Namespace-1 permissions: browse, read, write, delete, privileged Namespace-2 permissions: browse, read Namespace-3 permissions: browse, read, write

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 the two systems involved in a replication link.

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) involved in a replication link to the other system (system B) involved in the link:

- **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.

2

Configuring SSL for replication

Before replication can occur between two HCP systems, the systems involved must have a trust relationship with each other. This trust is based on shared SSL server certificates.

This chapter provides instructions for downloading, uploading, and deleting SSL server certificates for replication. For general information on SSL server certificates for HCP, see *Administering HCP*.

Roles: To view the SSL configuration for replication, you need the monitor or administrator role. To configure SSL for replication, you need the administrator role.

Sharing SSL server certificates

Before replication can occur on a replication link, the two systems involved in the link each need to have installed at least one valid replication SSL server certificate. Each system also needs to have installed at least one valid replication SSL server certificate from the other system as a trusted replication server certificate.

A valid replication SSL server certificate on any given HCP system is one that:

- Is associated with the domain that's associated with the network that's selected for replication on that system
- Has already reached its start date
- Has not expired

To share a certificate, the HCP administrator for the system in which the certificate is installed needs to download the certificate and give it to the administrator for the other system. That administrator then needs to upload the certificate as a trusted replication server certificate on the other system. Because what's downloaded is only the public portion of each server certificate, you can transfer the certificate unsecured.

For any given replication link, the two systems directly involved must share certificates. In a replication chain, for example, from system A to system B to system C, systems A and B must share certificates, and systems B and C must share certificates, but systems A and C don't need to do this.

If you take any of these actions on one of the systems in a replication pair, the replication certificate on that system automatically changes:

- Delete the certificate that's currently being used for replication
- Associate a different domain with the network that's selected for replication
- Select a different network for replication, where that network is associated with a different domain from the previously selected network
- Install a new valid certificate for the applicable domain where that certificate has an earlier start date than the certificate that's currently being used for replication

The replication certificate also automatically changes if:

- The certificate that's currently being used for replication expires and the applicable domain has at least one other certificate that's valid
- A future certificate for the applicable domain becomes valid and the domain has no other valid certificates

In any of these cases, when the certificate changes, HCP automatically suspends replication or recovery activity on all links in which the system participates. At that point, you need to download the new certificate and upload it to the other system for each link. After you upload the new certificate, activity on the link resumes automatically.

For more information on:

- Domains and SSL server certificates, see *Administering HCP*
- Selecting the network to use for replication, see "Selecting the network for replication" on page 132
- Suspended replication links, see "Suspending and resuming activity on an individual link" on page 122

Downloading an SSL server certificate

To download an installed replication SSL server certificate:

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

   On the replication **Certificates** page, the **Replication Server** section is displayed. The section shows this information about each currently installed replication SSL server certificate, regardless of whether it's expired:

   - **Server Certificate Domain** — The distinguished name for the certificate
   - **Valid On** — The date and time at which the certificate goes (or went) into effect

   - **Expires On** — The date and time at which the certificate expires (or expired)

4. Click on the download control (⤓) for the certificate you want to download. Then save the downloaded certificate in the location of your choice.

Uploading a trusted replication server certificate

To upload a downloaded replication SSL server certificate as a trusted replication server certificate:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the left side of the **Replication** page, click on **Certificates**.

4. On the replication **Certificates** page, click on **Trusted Replication**.

   The **Trusted Replication** section shows this information about each trusted replication server certificate:

   - **Server Certificate Domain** — The distinguished name for the certificate

   - **Valid On** — The date and time at which the certificate goes (or went) into effect

   - **Expires On** — The date and time at which the certificate expires (or expired)

5. In the **Trusted Replication** section, click on the **Browse** button. Then select the file containing the downloaded SSL server certificate.

6. Click on the **Upload Certificate** button.

   The **Trusted Replication** section displays the uploaded certificate.

Tip: You can also download a trusted replication server certificate. To do this, click on the download control (⤓) for the certificate in the **Trusted Replication** section.

Deleting a trusted replication server certificate

You can delete trusted replication server certificates. You may want to do this, for example, with certificates that become invalid.

You can delete a trusted replication server certificate only while all links in which the system participates are suspended. For instructions on suspending a link, see "Suspending and resuming activity on an individual link" on page 122.

To delete a trusted replication server certificate:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the left side of the **Replication** page, click on **Certificates**.

4. On the replication **Certificates** page, click on **Trusted Replication**.

5. In the **Trusted Replication** section, click on the delete control (🗑) for the certificate you want to delete.

6. In response to the confirming message, click on the **Remove Certificate** button.

3

Working with replication links

A replication link provides all the information the replication service needs to perform its functions. You can create a link in the HCP System Management Console for either of the systems involved in the link.

When you create a link, you specify whether it is active/active, the outbound side of an active/passive link (that is, the primary system), or the inbound side of an active/passive link (that is, the replica). Once a link exists, you can select HCP tenants, default-namespace directories, and, for active/passive links only, other links to replicate on the link.

In most cases, replication on a link starts automatically when you submit your first selection of items to be included in the link. However, if you create an active/passive link on the primary system for a link between a release 6.*x* HCP system and a release 7.0 HCP system, the primary system sends a link request to the replica. The replica administrator must explicitly accept the link before replication can begin on the link.

You can modify replication links at any time. You can also delete replication links. However, this action should be taken only with a full understanding of the consequences.

This chapter contains considerations and instructions for performing the activities mentioned above. For information on changing the time periods and intensity at which activity occurs on a link, see “Scheduling activity on a replication link” on page 118.

Roles: To view a replication link, you need the monitor or administrator role. To create, accept, reject, modify, or delete a replication link, you need the administrator role.

Note: You can also use the HCP management API to create, modify, and delete replication links. For information on doing this, see *HCP Management API Reference*.

Link properties

You create a replication link on either one of the systems you want to participate in the link. When you first create the link, you set these properties:

- The link name.
- Optionally, a description of the link.
- The link type. This can be active/active, outbound (for an active/passive link you're creating on the primary system), or inbound (for an active/passive link you're creating on the replica).
- Whether HCP should compress and/or encrypt data before transmitting it to the other system. For an active/active link, these settings apply to data being replicated in both directions on the link.
- Whether priority should be given to objects with the oldest changes, regardless of namespace, or processing should be balanced across namespaces. For an active/active link, this setting applies to both systems involved in the link.
- The identity of the other system you want to participate in the link. You can identify the other system either by domain name or by the IP addresses of storage nodes in that system.
- The port on which the other system will listen for transmissions from the system on which you're creating the link. Typically, you accept the default of 5748.
- If the system on which you're creating the link uses network address translation (NAT)
    - The domain name or IP addresses the other system should use to communicate with the system on which you're creating the link
    - The port on which the system on which you're creating the link will listen for transmissions from the other system

Note: If either or both of the systems involved in a link use a NAT, the NAT or NATs must be configured for port forwarding.

When you create a link, the automatic failover and failback properties default to disabled. You cannot modify these properties during link creation. For instructions on modifying them for an existing link, see "Changing automatic failover and failback settings" on page 93.

Considerations for creating a replication link

These considerations apply to creating replication links:

- You can create a replication link only after you have installed the required SSL certificates, as described in Chapter 2, "Configuring SSL for replication," on page 57.

- HCP supports replication between two release 7.0 systems and between release 7.0 and 6.*x* systems but not between 7.0 systems and systems earlier than release 6.0. A link between a 7.0 system and a 6.*x* system is called a **cross-release link**.

- You cannot create an active/active link between a release 7.0 HCP system and a release 6.*x* HCP system.

- When you create an active/passive link on the primary system for a link between a release 7.0 HCP system and a release 6.*x* HCP system, the primary system sends a link request to the replica. The system administrator for the replica can either accept of reject the request. If the system administrator rejects the request, the link is deleted.

- You can create multiple active/passive links from a single primary system to the same replica. However, the links cannot include the same HCP tenants and default-namespace directories.

- If you don't want HCP tenant-level users to see replication details, you should deselect the option that allows this before you create the replication link. For more information on this, see "Controlling the Tenant Management Console replication display" on page 112.

Creating a replication link

To create a replication link:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the replication **Overview** page, click on the **Create Link** button.

   The link creation wizard opens.

4. On the wizard **Name** page:

   - In the **Name** field, type a name for the link. Link names must be from one through 64 characters long and can contain any valid UTF-8 characters, including white space. Link names are not case sensitive.

     Link names must be unique within a replication topology.

   - Optionally, in the Description field, type a description for the link. This text can be up to 1,024 characters long and can contain any valid UTF-8 characters, including white space.

5. Click on the **Next** button.

6. On the wizard **Settings** page:

   - In the **Link Type** section:

     - To create an active/active link, select **Active/Active**.

     - To create an active/passive link for which the system on which you're creating the link is the primary system, select **Outbound**.

     - To create an active/passive link for which the system on which you're creating the link is the replica, select **Inbound**.

   - In the **Data Transfer** section:

     - Optionally, to compress data before it's transmitted from one system to the other, select the **Compress data** option.

       Compressing data increases network throughput but also increases processing time. Therefore, because compression applies to all transmitted data, you should enable it only if most of the data being replicated is compressible.

     - Optionally, to encrypt data before it's transmitted from one system to the other, select the **Encrypt data** option. Encrypting data keeps it confidential during transmission. The transmitted data is automatically decrypted on the target system.

If you don't select this option, data is transmitted unencrypted, even if the data is encrypted on disk in the system sending it. However, regardless of whether the data is encrypted before transmission, it is still secured by SSL during transmission.

- In the **Replication Priority** section:
  - To have the replication service balance its processing time evenly across the namespaces being replicated, select the **Balanced across namespaces** option.
  - To replicate objects with the oldest changes first, regardless of which namespaces they're in, select the **Oldest object first** option.

7. Click on the **Next** button.

8. On the wizard **Connection** page:

- In the **Remote Hostname or IP Addresses** field, type either the domain name of the other system you want to participate in the link or one or more comma-separated IP addresses of storage nodes in that system. For the domain name, use either of these formats:

  replication.*hcp-domain-name*

  replication.admin.*hcp-domain-name*

  In these formats, *hcp-domain-name* must be the name of the domain associated with the network that's selected for replication on the other system. For example, if the domain name is hcp-ca.example.com, you specify this in the **Remote Hostname or IP Addresses** field when using the second format shown above:

  replication.admin.hcp-ca.example.com

  For information on when the second format is required, see "Displaying the zone definition for the replication network domain" on page 135.

If you specify IP addresses for one or more nodes, they must be IP addresses in the network that's selected for replication on the other system.

Note: If you specify a domain name, the system on which you're creating the link transmits data to all storage nodes in the other system. If you specify IP addresses, the system on which you're creating the link transmits data to only those nodes. Therefore, you should specify IP addresses only if you have a compelling reason to do so (for example, HCP is not using DNS, or you need to reduce the processing load on some number of nodes).

- Optionally, to specify a different port for the other system to listen on and/or network address translation (NAT) information for the system on which you're creating the link, click on **Advanced Configuration**. Then, in the **Advanced Configuration** section, as applicable:

  - In the **Remote Port** field, type the number of the port on which the other system will listen for data from the system on which you're creating the link. The default port is 5748. Typically, you specify a different port only if other port usage makes it necessary.

  - If the system on which you're creating the link uses NAT for communication with the other system, specify the target for data transmissions from the other system:

    - In the **Local Hostname or IP Addresses** field, type either the name of the domain associated with the network that's selected for replication on the system on which you're creating the link (as that name is known to the other system) or one or more comma-separated IP addresses of storage nodes in that network (as those addresses are known to the other system). For the domain name, use either of these formats:

      `replication.`*`hcp-domain-name`*

      `replication.admin.`*`hcp-domain-name`*

      For information on when the second format is required, see "Displaying the zone definition for the replication network domain" on page 135.

Make sure the other system can resolve the domain name you specify.

Note: The other system will transmit data only to the nodes identified by the domain name or IP addresses you specify. Therefore, if you're using IP addresses, you should specify the addresses of all the storage nodes in the system.

- In the **Local Port** field, type the number of the port on which the system on which you're creating the link will listen for data from the other system. This is the port that's exposed to the other system.

Note: For network configurations that use NAT, multiple IP addresses must be configured to expose the HCP nodes. You cannot use multiple ports on the same IP address.

9. Click on the **Next** button.

10. On the wizard **Review** page, review the link configuration.

11. If the link configuration is what you want, click on the **Finish** button to create the link.

    If you're creating a cross-release active/passive link on the primary system, HCP sends a link request to the replica. On the primary system, the link has a status of **Pending remote reply**. On the replica, the link has a status of **Pending**. For information on how to handle a pending link request, see "Accepting or rejecting a pending cross-release link on a 7.0 replica" on page 72.

    If the link configuration is not what you want, you can use the **Previous** button in the wizard to return to pages on which you want to make corrections. Alternatively, you can press the **Cancel** button to leave the wizard without creating the link.

Tip: If HCP cannot create the configured link, check that the other system is healthy and that you've correctly shared SSL server certificates between the two systems.

Handling cross-release link requests

When you create an outbound active/passive link between a release 7.0 HCP system and a release 6.*x* HCP system, the primary system sends a link request to the replica. The administrator for the replica must accept the request before replication can start on the link.

Before a link request is accepted:

- On the primary system, the link has a status of **Pending remote reply**
- On the replica, the link has a status of **Pending**
- In the HCP System Management Console for the replica, the **Overview** page displays an alert indicating that a link request is pending

At this point, any of these can happen:

- The system administrator for the primary system can update the link properties or specify content for the link. For information on doing this on a release 7.0 primary system, see "Modifying a pending cross-release link on a 7.0 primary system" below.
- The system administrator for the primary system can cancel the link request. For information on doing this on a release 7.0 primary system, see "Canceling a cross-release link request on a 7.0 primary system" on page 71
- The system administrator for the replica can accept or reject the link. For information on doing this on a release 7.0 primary system, see "Accepting or rejecting a pending cross-release link on a 7.0 replica" on page 72.

For information on taking the above actions on a release 6.*x* system, see *Replicating Tenants and Namespaces* for the applicable release.

Modifying a pending cross-release link on a 7.0 primary system

To modify a pending cross-release link on a release 7.0 primary system:

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

3. In the list of links on the replication **Overview** page, click on the link you want to modify.

4. Take either or both of these actions:

   ○ To modify the link properties:

      1. In the top row of tabs, click on **Link**.

      2. In the link **Settings** panel, make the changes you want. For information on link property settings, see "Creating a replication link" on page 65.

      3. Click on the **Update Pending Link** button.

   ○ To modify the link content:

      1. In the top row of tabs, click on **Content**.

      2. In the **Content** panel, make the changes you want. For information on changing the content for a link, see "Specifying link content" on page 78.

      3. Click on the **Update Pending Link** button.

Canceling a cross-release link request on a 7.0 primary system

To cancel a cross-release link request on a release 7.0 primary system before the request is accepted:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. In the list of links on the replication **Overview** page, click on the link for which you want to cancel the request.

4. In the panel that opens, click on **Link**.

5. In the link **Settings** panel, click on the **Delete Link Request** button.

The requested link is deleted from both the primary system and the replica.

Accepting or rejecting a pending cross-release link on a 7.0 replica

To accept or reject a pending cross-release link on a release 7.0 replica:

1. On the **Replication** page in the HCP System Management Console, display the **Settings** panel for the link you want to accept or reject. You can display this panel in either of these ways:

   - In the alerts section on the System Management Console **Overview** page, click on the link name below the alert about the pending request.

   - In the top-level menu in the System Management Console:

     1. Mouse over **Services** to display a secondary menu.

     2. In the secondary menu, click on **Replication**.

     3. In the list of links on the replication **Overview** page, click on the link you want to accept or reject.

     4. In the panel that opens, click on **Link**.

2. Optionally, review the link settings.

   You can also click on the **Content** tab to review the link content.

3. Take one of these actions:

   - **To accept** the replication link, click on the **Accept Link** button.

     HCP sends an acceptance message to the primary system, which immediately starts replicating.

   - **To reject** the replication link, click on the **Reject Link** button.

     HCP sends a rejection message to the primary system, and both systems delete the pending link.

Considerations for specifying link content

These considerations apply to specifying the content for replication links:

- The content for a replication link (that is, what is replicated on the link) consists of any number of:
  - HCP tenants
  - HCP namespaces
  - Directories defined in the default namespace
  - For active/passive links only, chained links
- You can choose which namespaces to replicate for any given HCP tenant only if that tenant has granted system-level users administrative access to itself. Namespaces that are selected for replication are replicated on every link that includes the owning tenant.
- You can add and remove HCP tenants, default-namespace directories, and chained links from a replication link at any time, except during an upgrade of either system involved in the link, while the link is failed over, while data recovery is occurring on the link, or while the two systems involved in the link cannot communicate with each other.
- To select and deselect HCP namespaces for replication, you can use any link that includes the owning tenant.
- You can select HCP namespaces for replication at any time, except during an upgrade of either system involved in the link you started from. You can deselect HCP namespaces from replication at any time, except during an upgrade of either system involved in the link you started from, while a link that includes the owning tenant is failed over, or while data recovery is occurring on a link that includes the owning tenant.
- When you add one or more default-namespace directories to a link, some objects in directories that were already included in the link may be reprocessed. This can increase the replication backlog for the default namespace.
- You cannot add an HCP tenant to an active/passive link if the same tenant already exists on the replica. Two tenants are the same as each other if they have the same internal ID.

The only way the same tenant can exist on two systems is if it was created on one system and then replicated to the other system. Changing the name of the tenant on either of the systems does not make the two tenants different from each other.

- You cannot add an HCP tenant to any type of replication link if a different tenant with the same name already exists on the other system involved in the link. In this case, however, if you change the name of the tenant on either of the systems involved, you can add the tenant to the link.

- You cannot add a default-namespace directory to an active/passive link if a directory with the same name already exists on the replica. If the directory is empty on the primary system or the replica, you can change the directory name on that system. Then you can add the directory to the link.

- You cannot chain a replication link between system A and system B into an active/passive link between system B and system C if an HCP tenant included in the first link has the same name as a different tenant that already exists on system C. However, if you change the name of the tenant on system A or system C, you chain add first link into the second link.

- You cannot chain a replication link between system A and system B into an active/passive link between system B and system C if an HCP tenant included in the first link is the same as a tenant that already exists on system C, regardless of whether the tenant names are the same.

- You cannot chain a replication link between system A and system B into an active/passive link between system B and system C if a default-namespace directory included in the first link has the same name as a default-namespace directory that already exists on system C.

- In a replication chain with the configuration A ▸ B ▸ C, if you add an HCP tenant to the link from system A to system B and the same tenant already exists on system C, replication of that tenant from system B to system C is automatically paused. To recover from this situation:

  1. Either delete the tenant on system C, or remove the tenant from the link from system A to system B.

  2. Resume replication of the tenant on the link from system B to system C.

  For information on resuming replication of a tenant, see "Pausing and resuming replication or recovery of a tenant" on page 122.

- In a replication chain with the configuration A ▸ B ▸ C, if you add an HCP tenant to the link from system A to system B and a different tenant with the same name already exists on system C, replication of that tenant from system B to system C is automatically paused. To recover from this situation:

  1. Either rename the tenant on system C, delete the tenant on system C, or remove the tenant from the link from system A to system B.

  2. Resume replication of the tenant on the link from system B to system C.

- In a replication chain with the configuration A ▸ B ▸ C, if you add a default-namespace directory to the link from system A to system B and a default-namespace directory with the same name already exists on system C, replication of the default tenant from system B to system C is automatically paused. To recover from this situation:

  1. Either rename the directory on system C, delete the directory on system C, or remove the directory from the link from system A to system B. The first two options are possible only if the directory is empty on system C.

  2. Resume replication of the default tenant on the link from system B to system C.

- In a replication chain with the configuration A ▸ B ▸ C, if you add a default-namespace directory to the link from system A to system B and either of these is true, replication of the default tenant from system B to system C is automatically paused:

  - The default namespace doesn’t exist on system C. To recover from this situation:

    1. Create the default namespace on system C with the same retention mode and cryptographic hash algorithm as the default namespaces on systems A and B have.

    2. Resume replication of the default tenant on the link from system B to system C.

  - The default namespace exists on system C but has a different retention mode or cryptographic hash algorithm from the default namespaces on systems A and B. To recover from this situation:

    1. Delete the default-namespace directory from the link from system A to system B.

2. Resume replication of the default tenant on the link from system B to system C.

- Once you remove an HCP tenant or directory from an active/passive link, you cannot add it back to that link unless you first delete it from the replica. This is because the tenant or directory now already exists on the replica.

- On the **Replication** page in the System Management Console, each panel in which you can add and remove items of a specific type to and from a replication link includes a list of the items of the applicable type that are already included in the link. In those lists, rows containing items that have been deleted from HCP are highlighted in red and have a trash can icon (🗑) on the right.

  HCP automatically removes each deleted namespace from a link after the deletion has been replicated. You need to remove deleted tenants and directories yourself.

Important: Do not remove a deleted item from a replication link until after the deletion has been replicated.

- If you deselect a namespace from a replication, any further action on that namespace, including deletion of the namespace, is not replicated. This enables a situation in which the owning tenant can end up with more namespaces than are allowed by its namespace quota and/or using more storage than is allowed by its hard quota.

  For example, consider the following scenario for a tenant has reached its namespace quota of five:

  1. The tenant is selected for replication on an active/passive link, along with all five of its namespaces.

     The tenant now has five namespaces on both the primary system and the replica.

  2. One of the namespaces is removed from the replication link.

  3. On the primary system, the tenant empties and deletes the removed namespace.

     The tenant now has four namespaces on the primary system and five namespaces on the replica.

  4. On the primary system, the tenant creates a new namespace.

   The tenant now has five namespaces on the primary system and five namespaces on the replica, but the namespaces on the two systems are not the same ones.

5. The new namespace is added to the replication link.

   The tenant now has five namespaces on the primary system and six namespaces on the replica.

- A namespace can contain metadata-only objects on one or more HCP systems in a replication topology. If such a system is involved in only one replication link that includes the tenant that owns the namespace, the following actions can result in the data for those metadata-only objects becoming inaccessible from that system, possibly permanently:

  - You remove the tenant that owns the namespace from the link. In this case, to make the data accessible again:

    - If the link is active/active, add the tenant back into the link.

    - If the link is active/passive, either change the link to active/active, or create a new active/active link between the same two systems. Then add the tenant back to the changed link or to the new link, as applicable.

  - You deselect the namespace from replication. In this case, to make the data accessible again, reselect the namespace.

  - The tenant that owns the namespace is included in the first link in a replication chain, the namespace with the metadata-only objects is on the third HCP system in the chain, and you remove the first link from the link to the third system. In this case, to make the data accessible again, take either of these actions:

    1. If the first link in the chain is active/passive, change the link to active/active.

    2. Either change the link to the third system to active/active, or create a new active/active link between the second and third systems in the chain.

    3. Add the tenant back to the changed link to the third system or to the new link, as applicable.

  - ○ The namespace with the metadata-only objects is on the third HCP system in a replication chain, and you remove the tenant that owns the namespace from the first replication link in the chain. In this case, to make the data accessible again:

    1. Create a new active/active link between the first and second systems in the replication chain.

    2. Add the tenant to the new link.

  HCP displays a warning about metadata-only objects when you submit a request to delete an active/passive link or to remove a tenant from an active/passive link while any of the namespaces being replicated on the link contain metadata-only objects.

  For more information on metadata-only objects, see *Administering HCP*.

- You can add an HCP tenant to a replication link on one of the systems involved in the link even if the management and data access networks associated with that tenant aren't defined on the other system involved in the link. However, if these networks are not defined on the other system, the tenant and its namespaces will be inaccessible on that system.

  If a network associated with a tenant is undefined in a given system, the tenant list on the **Tenants** page in the System Management Console for that system shows this alert for the tenant:

  For more information on networks, see *Administering HCP*.

Specifying link content

To specify the content for a replication link:

1. In the top-level menu in the HCP System Management Console for either system involved in the link, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. In the list of links on the replication **Overview** page, click on the link for which you want to specify content.

4. In the panel that opens, click on **Content**.

5. In the **Content** panel, take either of the following actions:

   - To specify content for an active/active link, follow the instructions in "Specifying link content for an active/active link" on page 79.

   - To specify content for an active/passive link, follow the instructions in "Specifying link content for an active/passive link" on page 84.

Specifying link content for an active/active link

To specify content for an active/active link, take one or more of these actions in the **Content** panel for the link:

- To add and remove HCP tenants to and from the link, click on the **Tenants** tab. Then follow the instructions in "Adding and removing HCP tenants in an active/active link" below.

- To select and deselect namespaces for replication, follow the instructions in "Selecting and deselecting HCP namespaces for tenants in an active/active link" on page 81.

  Selected namespaces are replicated on every link that includes the owning tenant.

- To add and remove default-namespace directories to and from the link, click on the **Default Tenant** tab. Then follow the instructions in "Adding and removing default-namespace directories in an active/active link" on page 82.

  The **Default Tenant** tab is present only if the default namespace exists on the local system.

Adding and removing HCP tenants in an active/active link

The link content **Tenants** panel for an active/active link has these sections:

- **Tenants Replicating** — Lists the HCP tenants already included in the link.

- **Local Tenants** — Lists the HCP tenants on the local system that are eligible for replication but not currently included in the link. This list does not include tenants that are read-only on the local system due to replication.

- **Remote Tenants** — Lists the HCP tenants on the other system involved in the link that are eligible for replication but not currently included in the link. This list does not include tenants that are read-only on the other system due to replication.

If a tenant has granted system-level users administrative access to itself, the row for the tenant in the applicable list shows the number of namespaces owned by the tenant that are selected for replication. On the primary system for an active/passive link, the row shows this number out of the total number of namespaces owned by the tenant.

To add HCP tenants to an active/active link, in the link content **Tenants** panel:

1. Optionally, filter the list from which you want to select one or more tenants to add to the link. For instructions on doing this, see "Filtering item selection lists" on page 91.

2. For each listed tenant you want to add to the link, click in the tenant row.

   The tenant is selected, and the tenant row changes color.

   To select all the tenants in the list, click on the **Select All** button above the list.

   To deselect a selected tenant, click in the tenant row.

   To deselect all the selected tenants, click on the **Clear** button.

3. Click on the **Add Selected Tenants** button below the list.

To remove HCP tenants from an active/active link, in the link content **Tenants** panel:

1. Optionally, in the **Tenants Replicating** section, filter the list of tenants that are currently included in the link. For instructions on doing this, see "Filtering item selection lists" on page 91.

2. For each listed tenant you want to remove from the link, click in the tenant row.

   The tenant is selected, and the tenant row changes color.

   To select all the tenants in the list, click on the **Select All** button above the list.

To deselect a selected tenant, click in the tenant row.

To deselect all the selected tenants, click on the **Clear** button.

3. Click on the **Remove Selected Tenants** button below the list.

Selecting and deselecting HCP namespaces for tenants in an active/active link

You select and deselect HCP namespaces for replication in the **Manage Namespaces** window for the tenant that owns the namespaces. To open this window, in the link content **Tenants** panel, click on the manage namespaces control ( ) for the owning tenant. This control is present only if the tenant has granted system-level users administrative access to itself.

When you start from an active/active link, the **Manage Namespaces** window has these sections:

- **Namespaces Selected for Replication** — Lists the namespaces that are owned by the tenant and that are currently selected for replication

- **Local Namespaces** — Lists the namespaces that are owned by the tenant on the local system and that are not currently selected for replication

- **Remote Namespaces** — Lists the namespaces that are owned by the tenant on the other system involved in the link and that are not currently selected for replication

To close the **Manage Namespaces** window without taking any action, click on the close control (×) in the upper right corner.

To select HCP namespaces for replication when you start from an active/active link, in the **Manage Namespaces** window:

1. Optionally, filter the list from which you want to select one or more namespaces for replication. For instructions on doing this, see “Filtering item selection lists” on page 91.

2. For each listed namespace you want to select for replication, click in the namespace row.

   The namespace is selected, and the namespace row changes color.

To select all the namespaces in the list, click on the **Select All** button above the list.

To deselect a selected namespace, click in the namespace row.

To deselect all the selected namespaces, click on the **Clear** button.

3. Click on the **Add Selected Namespaces** button below the list.

To deselect HCP namespaces from replication when you start from an active/active link, in the **Manage Namespaces** window:

1. Optionally, in the **Namespaces Selected for Replication** section, filter the list of namespaces that are currently selected for replication. For instructions on doing this, see "Filtering item selection lists" on page 91.

2. For each listed namespace you want to deselect from replication, click in the namespace row.

   The namespace is selected, and the namespace row changes color.

   To select all the namespaces in the list, click on the **Select All** button above the list.

   To deselect a selected namespace, click in the namespace row.

   To deselect all the selected namespaces, click on the **Clear** button.

3. Click on the **Remove Selected Namespaces** button below the list.

Adding and removing default-namespace directories in an active/active link

The link content **Default Tenant** panel for an active/active link has these sections:

- **Directories Replicating** — Lists the HCP default-namespace directories already included in the link.

- **Local Directories** — Lists the default-namespace directories on the local system that are not currently included in the link. This list does not include directories that are read-only on the local system due to replication.

- **Remote Directories** — Lists the default-namespace directories on the other system involved in the link that are not currently included in the link. This list does not include directories that are read-only on the other system due to replication.

To add default-namespace directories to an active/active link, in the link content **Default Tenant** panel:

1. Optionally, filter the list from which you want to select one or more directories to add to the link. For instructions on doing this, see "Filtering item selection lists" on page 91.

2. For each listed directory you want to add to the link, click in the directory row.

   The directory is selected, and the directory row changes color.

   To select all the directories in the list, click on the **Select All** button above the list.

   To deselect a selected directory, click in the directory row.

   To deselect all the selected directories, click on the **Clear** button.

3. Click on the **Add Selected Directories** button below the list.

To remove default-namespace directories from an active/active link, in the link content **Default Tenant** panel:

1. Optionally, in the **Directories Replicating** section, filter the list of directories that are currently included in the link. For instructions on doing this, see "Filtering item selection lists" on page 91.

2. For each listed directory you want to remove from the link, click in the directory row.

   The directory is selected, and the directory row changes color.

   To select all the directories in the list, click on the **Select All** button above the list.

   To deselect a selected directory, click in the directory row.

   To deselect all the selected directories, click on the **Clear** button.

3. Click on the **Remove Selected Directories** button below the list.

Specifying link content for an active/passive link

To specify content for an active/passive link, take one or more of these actions in the **Content** panel for the link:

- To add and remove HCP tenants to and from the link, click on the **Tenants** tab. Then follow the instructions in "Adding and removing HCP tenants in an active/passive link" below.

- To select and deselect namespaces for replication, follow the instructions in "Selecting and deselecting HCP namespaces for tenants in an active/passive link" on page 86.

  Selected namespaces are replicated on every link that includes the owning tenant.

- To add and remove default-namespace directories to and from the link, click on the **Default Tenant** tab. Then follow the instructions in "Adding and removing default-namespace directories in an active/passive link" on page 88.

  The **Default Tenant** tab is present only if the default namespace exists on the local system.

- To add and remove chained links to and from the link, click on the **Chained Links** tab. Then follow the instructions in "Adding and removing chained links in an active/passive link" on page 89.

Adding and removing HCP tenants in an active/passive link

The link content **Tenants** panel for an active/passive link has these sections:

- **Tenants Replicating** — Lists the HCP tenants already included in the link.

- **Local Tenants** — For outbound links only, lists the HCP tenants on the local system that are eligible for replication but not currently included in the link. This list does not include tenants that are either:

  - Read-only on the local system due to replication

  - Included in a link that is chained into the current link

  This section is not present for inbound links.

- **Remote Tenants** — For inbound links only, lists the HCP tenants on the other system involved in the link that are eligible for replication but not currently included in the link. This list does not include tenants that are either:

  - Read-only on the other system due to replication

  - Included in a link that is chained into the current link

  This section is not present for outbound links.

If a tenant has granted system-level users administrative access to itself, the row for the tenant shows the number of namespaces owned by the tenant that are selected for replication.

To add or remove HCP tenants in an active/passive link, in the link content **Tenants** panel:

1. Take either or both of these actions:

   - **To add HCP tenants to the link:**

     1. Optionally, filter the list of tenants in the **Local Tenants** or **Remote Tenants** section, as applicable. For instructions on doing this, see “Filtering item selection lists” on page 91.

     2. For each listed tenant you want to add to the link, click in the tenant row.

        The tenant is selected, and the tenant row changes color.

        To select all the tenants in the list, click on the **Select All** button above the list.

        To deselect a selected tenant, click in the tenant row.

        To deselect all the selected tenants, click on the **Clear** button.

   - **To remove HCP tenants from the link:**

     1. Optionally, filter the list of tenants in the **Tenants Replicating** section. For instructions on doing this, see “Filtering item selection lists” on page 91.

     2. For each listed tenant you want to remove from the link, click in the tenant row.

The tenant is selected, and the tenant row changes color.

To select all the tenants in the list, click on the **Select All** button above the list.

To deselect a selected tenant, click in the tenant row.

To deselect all the selected tenants, click on the **Clear** button.

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

   If you're removing tenants from the link, 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 **Update Link** button.

Selecting and deselecting HCP namespaces for tenants in an active/passive link

You select and deselect HCP namespaces for replication in the **Manage Namespaces** window for the tenant that owns the namespaces. To open this window, in the link content **Tenants** panel, click on the manage namespaces control ( ) for the owning tenant. This control is present only if the tenant has granted system-level users administrative access to itself.

When you start from an active/passive link, the **Manage Namespaces** window has these sections:

- **Namespaces Selected for Replication** — Lists the namespaces that are owned by the tenant and that are currently selected for replication.

- **Select Namespaces for Replication** — Lists the namespaces that are owned by the tenant and that are not currently selected for replication.

To close the **Manage Namespaces** window without taking any action, click on the close control (×) in the upper right corner.

To select or deselect HCP namespaces for replication when you start with an active/passive link, in the **Manage Namespaces** window:

1. Take either or both of these actions:

   - **To select HCP namespaces for replication:**

     1. Optionally, filter the list of namespaces in the **Select Namespaces for Replication** section. For instructions on doing this, see “Filtering item selection lists” on page 91.

     2. For each listed namespace you want to select for replication, click in the namespace row.

        The namespace is selected, and the namespace row changes color.

        To select all the namespaces in the list, click on the **Select All** button above the list.

        To deselect a selected namespace, click in the namespace row.

        To deselect all the selected namespaces, click on the **Clear** button.

   - **To deselect HCP namespaces from replication:**

     1. Optionally, filter the list of namespaces in the **Namespaces Selected for Replication** section. For instructions on doing this, see “Filtering item selection lists” on page 91.

     2. For each listed namespace you want to deselect from replication, click in the namespace row.

        The namespace is selected, and the namespace row changes color.

        To select all the namespaces in the list, click on the **Select All** button above the list.

        To deselect a selected namespace, click in the namespace row.

        To deselect all the selected namespaces, click on the **Clear** button.

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

Adding and removing default-namespace directories in an active/passive link

The link content **Default Tenant** panel for an active/passive link has these sections:

- **Directories Replicating** — Lists the default-namespace directories already included in the link.
- **Local Directories** — For outbound links only, lists the default-namespace directories on the local system that are not currently included in the link. This list does not include directories that are either:
  - Read-only on the local system due to replication
  - Included in a link that is chained into the current link

  This section is not present for inbound links.
- **Remote Directories** — For inbound links only, lists the default-namespace directories on the other system involved in the link that are not currently included in the link. This list does not include directories that are either:
  - Read-only on the other system due to replication
  - Included in a link that is chained into the current link

  This section is not present for outbound links.

To add or remove default-namespace directories in an active/passive link, in the link content **Default Tenant** panel:

1. Take either or both of these actions:
   - **To add default-namespace directories to the link:**
     1. Optionally, filter the list of directories in the **Local Directories** or **Remote Directories** section, as applicable. For instructions on doing this, see “Filtering item selection lists” on page 91.
     2. For each listed directory you want to add to the link, click in the directory row.

        The directory is selected, and the directory row changes color.

To select all the directories in the list, click on the **Select All** button above the list.

To deselect a selected directory, click in the directory row.

To deselect all the selected directories, click on the **Clear** button.

- ○ **To remove default-namespace directories from the link:**

  1. Optionally, filter the list of directories in the **Directories Replicating** section. For instructions on doing this, see "Filtering item selection lists" on page 91.

  2. For each listed directory you want to remove from the link, click in the directory row.

     The directory is selected, and the directory row changes color.

     To select all the directories in the list, click on the **Select All** button above the list.

     To deselect a directory tenant, click in the directory row.

     To deselect all the selected directories, click on the **Clear** button.

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

   If you're removing directories from the link, 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 **Update Link** button.

Adding and removing chained links in an active/passive link

The link content **Chained Links** panel for an active/passive link has these sections:

- **Links Replicating** — Lists the links already included in the current link.

- **Local Links** — For outbound links only, lists the active/active links in which the current link participates and the active/passive links for which the current system is the replica.

  This section is not present for inbound links.

- **Remote Links** — For inbound links only, lists the active/active links in which the other system involved in the current link participates and the active/passive links for which the other system involved in the current link is the replica.

  This section is not present for outbound links.

To view the content of a link in any of the lists described above:

1. Click on the view content control ( ) for the link.

2. In the **Link Content** window:

   - To see a list of the HCP tenants included in the link, click on the **Tenants** tab.

   - To see a list of the default-namespace directories included in the link, click on the **Default Tenant** tab.

   - To see a list of the chained links included in the link, click on the **Chained Links** tab.

   You can filter any of these lists to display a subset of the listed items. For instructions on doing this, see “Filtering item selection lists” on page 91.

3. When you’ve finished viewing the link content, click on the **Close** button.

To add or remove chained links in an active/passive link, in the link content **Chained Links** panel:

1. Take either or both of these actions:

   - **To add chained links to the current link:**

     1. Optionally, filter the list of links in the **Local Links** or **Remote Links** section, as applicable. For instructions on doing this, see “Filtering item selection lists” on page 91.

     2. For each listed link you want to add to the current link, click in the link row.

        The link is selected, and the link row changes color.

        To select all the links in the list, click on the **Select All** button above the list.

To deselect a selected link, click in the link row.

To deselect all the selected links, click on the **Clear** button.

- **To remove chained links from the current link:**

  1. Optionally, filter the list of links in the **Links Replicating** section. For instructions on doing this, see “Filtering item selection lists” on page 91.

  2. For each listed link you want to remove from the current link, click in the link row.

     The link is selected, and the link row changes color.

     To select all the links in the list, click on the **Select All** button above the list.

     To deselect a selected link, click in the link row.

     To deselect all the selected links, click on the **Clear** button.

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

   If you’re removing chained links from the current link, 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 **Update Link** button.

For a discussion of chained replication, see “Active/passive chained replication” on page 23.

Filtering item selection lists

By default, the lists you use to add and remove items in replication links contain all of the applicable items. You can filter any of these lists to display a subset of the applicable items. A filtered list contains only those items with a name that begins with or is the same as a specified text string.

You cannot filter a list while any of the items in it are selected.

To filter a list:

1. In the field above the list, type the text string you want to use as a filter. This string can be up to 64 characters long, can contain any valid UTF-8 characters, and is not case sensitive. White space is allowed.

2. Click on the find control (🔍).

To redisplay the entire list, click on the clear filter control (☒).

Considerations for modifying replication link settings

These considerations apply to modifying the settings for replication links:

- You can modify the settings for a replication link in the HCP System Management Console for either system involved in the link.

- You can modify the settings for a replication link at any time.

- You can change an active/passive link to an active/active link, but you cannot change an active/active link to an active/passive link.

- When changing an active/passive link to active/active, you need to ensure that replication checkpoints are up to date before converting, otherwise you risk losing directory metadata.

- After an upgrade of an HCP system from release 6.x to release 7.0, you can change preexisting active/passive links to active/active links. However, if such a link includes one or more default-namespace directories and metadata changes were made to those directories before the upgrade, those metadata changes may be lost. To prevent this loss, wait until replication is up to date as of the time of the upgrade before you change the link type.

- You cannot change the direction of an active/passive link.

- If you change the domain name or IP addresses for either system involved in a link, the new domain name or IP addresses must identify the same system. The exception to this is when you're reconfiguring the link before restoring it to a new or rebuilt system.

- When you change an active/passive link to an active/active link, HCP suspends activity on the link. You need to resume activity on the link manually. For information on resuming activity on link, see "Suspending and resuming activity on an individual link" on page 122.

- While updating link properties other than the link type, HCP temporarily suspends activity on the link.

- With a cross-release link, you can enable automatic failover and failback. However, enabling automatic failover has no effect, and automatic failback works only when the release 7.0 system is the replica. Enabling automatic failback when the release 7.0 system is the primary system has no effect.

Modifying replication link settings

You use separate procedures to modify the link properties you can set when you create a link and to modify the automatic failover and failback settings for a link.

Changing link settings

To change replication link properties that you can set when you create a link:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the **Replication** page, click on the name of the link you want to modify.

4. In the panel that opens, click on the **Link** tab.

5. In the replication **Link** panel, click on the **Settings** tab.

6. In the link **Settings** panel, make the changes you want:

   For information on link property settings, see “Creating a replication link” on page 65.

7. Click on the **Update Link** button.

Changing automatic failover and failback settings

To change the automatic failover and failback settings for a replication link:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the **Replication** page, click on the name of the link you want to modify.

4. In the panel that opens, click on the **Link** tab.

5. In the replication **Link** panel, click on the **Failover** tab.

6. In the **Failover** panel, click on **Automatic Failover Settings**.

7. In the **Automatic Failover Settings** section:

   - For an active/active link or an inbound active/passive link, in the **Local System** section:

     - To enable automatic failover to the current system, select the **Fail over automatically** option. Then, in the **Automatically fail over link after loss of connectivity for** field, type the number of minutes HCP should wait before automatically failing over the link after a loss of connectivity to the other system involved in the link. Valid values are integers in the seven through 9,999.

     - To disable automatic failover to the current system, deselect the **Fail over automatically** option.

     - For active/passive links only, to enable automatic failback to the other system involved in the link, select the **Complete recovery automatically** option. Then, in the **Automatically complete recovery when synchronization is less than** field, type the number of minutes the up-to-date-as-of time for the link must be less than before HCP should automatically fail back the link. Valid values are integers in the one through 9,999.

     - For active/passive links only, to disable automatic failback to the other system involved in the link, deselect the **Complete recovery automatically** option.

   - For an active/active link or an outbound active/passive link, in the **Remote System** section:

     - To enable automatic failover to the other system involved in the link, select the **Fail over automatically** option. Then, in the **Automatically fail over link after loss of connectivity for** field, type the number of minutes HCP should wait before automatically failing over the link after a loss of connectivity to the current system. Valid values are integers in the seven through 9,999.

- To disable automatic failover to the other system involved in the link, deselect the **Fail over automatically** option.
- For active/passive links only, to enable automatic failback to the current system, select the **Complete recovery automatically** option. Then, in the **Automatically complete recovery when synchronization is less than** field, type the number of minutes the up-to-date-as-of time for the link must be less than before HCP should automatically fail back the link. Valid values are integers in the one through 9,999.
- For active/passive links only, to disable automatic failback to the current system, deselect the **Complete recovery automatically** option.

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

For information on how automatic failover and failback work, see “Failover and failback” on page 37.

Considerations for deleting a replication link

These considerations apply to deleting replication links:

- You can delete a replication link in the HCP System Management Console for either system involved in the link.
- You can delete a replication link only while activity on the link is suspended. For information on suspending activity on a link, see “Suspending and resuming activity on an individual link” on page 122.
- When you delete a replication link, replication on that link immediately stops. The HCP tenants and default-namespace directories included in the link remain or become read-write on both systems involved in the link.
- You can delete a replication link that’s included as a chained link in a second replication link. When you do this, the HCP tenants and default-namespace directories that were included in the deleted link are automatically included individually in the second link.

- After deleting an active/passive link, you can no longer replicate the HCP tenants and default-namespace directories included in it to the same replica on an active/passive link unless you first delete them on the replica. If you don't delete them, when you create a new active/passive link with the same replica, HCP doesn't let you include them because they now exist on both systems.

  If the possibility exists that you will want to replicate those tenants and directories again to the same replica on an active/passive link, you should either pause replication for the individual tenants or suspend activity on the link instead of deleting the link.

- After deleting an active/passive link, you can include the HCP tenants and default-namespace directories that were included in the deleted link in a new or existing active/active link. In this case, you do not need to delete those tenants and directories from the system that was the replica for the deleted link.

- A namespace can contain metadata-only objects on one or more systems in a replication topology. If you delete the only replication link that involves such a system, the data for those metadata-only objects becomes inaccessible from that system. In this case, to make the data accessible again:

  - If the deleted link was active/active, recreate the link.

  - If the deleted link was active/passive, recreate the link as an active/active link.

  For information on metadata-only objects, see *Administering HCP*.

Deleting a replication link

To delete a replication link:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the **Replication** page, click on the name of the link you want to delete.

4. In the panel that opens, click on the **Management** tab.

5. In the **Management** panel, click on **Delete Replication Link**.

6. In the **Delete Replication Link** section, click on the **Delete** button.

7. In response to the confirming message:

   - If any namespace included in the link contains metadata-only objects on either of the systems involved in the link or if the two systems involved in the link cannot communicate with each other:

     1. In the field in the message window, type *YES*. This value is case sensitive.

     2. Click on the **Delete Link** button.

   - If no namespaces included in the link contain metadata-only objects on either of the systems involved in the link and the two systems involved in the link can communicate with each other, click on the **Delete Link** button.

4

Monitoring replication

You can monitor the status of each currently configured replication link, as well as replication and recovery activity on the link, in the HCP System Management Console for either system involved in the link. By periodically reviewing link activity, you can decide whether to change the replication rate in the link configuration to accommodate other loads on system performance. You can also determine whether you need to add more storage capacity to either system.

Under normal circumstances, the replication service works without any intervention required. System Management Console alerts inform you of conditions, such as network connection problems, that may require action on your part to allow the service to continue processing.

For HCP tenants, you can control whether information about replication activity is displayed in the Tenant Management Console. For the default tenant, this information is always displayed.

This chapter describes the information available to you for the links you've created.

Roles: To monitor replication, you need the monitor or administrator role.

Note: You can also use the HCP management API to monitor replication links. For information on doing this, see *HCP Management API Reference*.

Replication overview

The replication **Overview** page in the HCP System Management Console lists the currently defined replication links in which the current system participates. To display this page:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

For each listed link, the replication **Overview** page shows:

- **Name** — The link name.

- **Type** — The link type. The link type can be **Active/Active**, **Outbound**, or **Inbound**.

- **Status** — The link status, represented by an icon and text. The icon indicates the general health of the link. The text specifies the status of the link. The possible icons are

    - ✓ — The link is healthy. The statuses that can accompany this icon are:

        - **Synchronizing data** — On either system for an active/active link, the local system is sending data to the other system involved in the link.

        - **Sending data** — On the primary system for an active/passive link during replication, the primary system is sending data to the replica.

        - **Receiving data** — On the replica for an active/passive link during replication, the primary system is sending data to the replica.

    - ! — The link is healthy, but normal replication is not occurring on the link. The statuses that can accompany this icon are:

        - **Recovering data** — On the replica for an active/passive link during the first part of data recovery, the replica is sending data to the primary system.

  - **Completing recovery** — On the replica for an active/passive link during final data recovery, the replica is sending data to the primary system.

  - **Receiving data** — On the primary system for an active/passive link during data recovery, the replica is sending data to the primary system.

  - **Scheduled off period:**

    - On the local system for an active/active link, the local system is not sending data to the other system involved in the link because the scheduled performance level for the current time period for the local system is **Off**.

    - On both the primary system and the replica for an active/passive link, no replication or recovery activity is occurring on the link because the scheduled performance level for the current time period is **Off**.

  - **Suspended by user** — On both systems for an active/active or active/passive link, activity on the link has been manually suspended. No replication or recovery is occurring.

  - **Pending remote reply** — On the primary system for a cross-release link, the primary system has sent a link request to the replica, and the administrator for the replica has not yet responded to the request.

  - **Pending**:

    - On the replica for a cross-release link, the primary system has sent a link request to the replica, and the administrator for the replica has not yet responded to the request.

    - On the primary system for a cross-release link, the replica has sent a request to restore a broken link, and the administrator for the primary system has not yet responded to the request.

- ⚠ — The link is unhealthy. The statuses that can accompany this icon are:

  - **Failed over** — On both systems for an active/active or active/passive link, the link is currently failed over.

- **Remote storage full, link suspended** — On the local system for an active/active link, on the primary system for an active/passive link during replication, or on the replica for an active/passive link during data recovery, the other system involved in the link does not have enough free space in primary running storage to accept any more data transmissions. In any case, HCP has automatically suspended activity on the link.

- **Local storage full, link suspended** — On the local system for an active/active link, on the replica for an active/passive link during replication, or on the primary system for an active/passive link during data recovery, the local system does not have enough free space in primary running storage to accept any more data transmissions. In any case, HCP has automatically suspended activity on the link.

- **High error rate** — On both systems for an active/active or active/passive link, errors are occurring at a high rate. If you see this status for:

  - An active/active link, check the health of both systems involved in the link

  - An active/passive link during replication, check the health of the replica.

  - An active/passive link during data recovery, check the health of the primary system

  In any case, also check the health of the network connection between the two systems involved in the link.

  If you cannot find the problem, contact your authorized HCP service provider for help.

- **Stalled link** — On both systems for an active/active or active/passive link, activity on the link has stopped. Check the network connection between the two systems. If the connection appears to be working properly, contact your authorized HCP service provider for help.

- **Unrecognized link** — On either system for an active/active or active/passive link, the other system involved in the link doesn't recognize the link. Restore the link. Then, if applicable, start data recovery. For instructions on these activities, see "Recovering from a failure" on page 148.

   – **Broken link** — On both systems for an active/active or active/passive link, the replication service cannot contact any nodes on the other system involved in the link. If you see this status for:

     - An active/active link, check the health of both systems involved in the link

     - An active/passive link during replication, check the health of the replica

     - An active/passive link during data recovery, check the health of the primary system

     In any case, also check the health of the network connection between the two systems involved in the link.

     If you cannot find the problem, contact your authorized HCP service provider for help.

  If more than one status applies to a link, the page shows the icon and text for only one of the statuses.

  For additional information on the status of a link, mouse over the status icon.

Note: The System Management Console may not immediately reflect certain changes in link status.

- **Alerts** — If replication or recovery activity is paused for one or more tenants included in the link, either of these icons:

  ○ ❗ — Replication or recovery has been manually paused for one or more tenants included in the link.

  ○ ⚠ — Replication or recovery has been automatically paused for one or more tenants included in the link. To view text describing the condition that's causing this alert, mouse over the alert icon.

  For information on events that can cause replication or recovery of a tenant to be paused automatically, see “Automatically paused tenant replication or recovery” on page 123.

  If replication or recovery is paused manually for one or more tenants and automatically for one or more other tenants, the page shows only the icon indicating an automatic pause.

- The number of nodes on the other system that can receive replication transmissions from the current system. If the current system can determine the total number of storage nodes on the other system, the page shows that number as well.

General link information

The list of links on the replication **Overview** page shows the current status of each replication link. To further monitor the status of a link and to manage the link, click on the name of the link in the list.

The top part of the page that opens shows this information for the link:

- The link name.

- The name of the domain associated with the [hcp_system] network for the local system (the system on which you're viewing the link).

- The name of the domain associated with the [hcp_system] network for the other system involved in the link (the remote system for the link).

- An large arrow showing the direction of the link, overlaid with the link status. For an active/active link, this arrow has two arrowheads, one at each end.

- For active/active links, the date and time before which configuration changes and changes to namespace content are guaranteed to be synchronized in both directions between the two systems involved in the link. For active/passive links, the date and time before which configuration changes and changes to namespace content are guaranteed to have been replicated or recovered on the link, as applicable.

Link overview

To monitor the status of and activity on a replication link, you use the link status **Overview** panel in the HCP System Management Console. To open this panel:

1. In the list of links on the replication **Overview** page, click on the name of the link you want to monitor.

2. On the page for the individual link, click on **Status**.

3. In the **Status** panel, click on the **Overview** tab.

The top part of the link status **Overview** panel shows the current status of activity on the link:

- For an active/active link, a section labeled **Sending** shows the status of send activity on the local system. A section labeled **Receiving** shows the status of send activity on the remote system.
- For an outbound active/passive link during replication and an inbound active/passive link during data recovery, a section labeled **Sending** shows the status of send activity on the local system.
- For an outbound active/passive link during data recovery and an inbound active/passive link during replication, a section labeled **Receiving** shows the status of send activity on the remote system.

The **Sending** and **Receiving** sections show:

- **Data pending** — The approximate amount of data currently waiting to be sent to the other system. This is the sum of the amounts of data waiting to be sent in each HCP namespace included in the link. This value does not include data in the default namespace.

  This information is not available in the **Receiving** section on a replica during replication or on a primary system during data recovery.

- **Objects pending** — The approximate number of objects currently waiting to be sent to the other system. This is the sum of the numbers of objects waiting to be sent in each HCP namespace included in the link. This number does not include objects in the default namespace.

  This information is not available in the **Receiving** section on a replica during replication or on a primary system during data recovery.

- **Up to date as of** — The amount of time that is the difference between:
  - The date and time before which configuration changes and changes to namespace content are guaranteed to have been sent to the other system
  - The current date and time
- A graph showing the history of the rate of data transmissions on the link (see “Data transmission rate” below).
- A graph showing the history of the rate of operations on the link (see “Operation rate” on page 106).

- For active/active links and outbound active/passive links, one or two pie charts showing the amount of used and free storage space on the remote system (see “Remote storage space” on page 107).

The statistics for replication links are updated every ten minutes. To see updated statistics, redisplay the **Replication** page.

Data transmission rate

The **Transfer Rate** graph in the link status **Overview** panel for a replication link shows the history of the rate of data transmissions on the link per second. For an active/active link, this is the rate of data transmission from the local system to the remote system. For an active/passive link, this is the rate of data transmission during replication or recovery, whichever is happening at the time. In any case, the data transmission rate is cumulative for all the HCP tenants and default-tenant directories included in the link.

If the **Transfer Rate** graph is not currently visible, click on **Transfer Rate** to display it.

The x-axis in the **Transfer Rate** graph marks the passage of time. It shows 30 days (or fewer if the replication link was created less than 30 days ago). 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).

The graph heading displays the current data transmission rate. Factors that affect this rate include the amount of other traffic on the network, the current load on both systems involved in the link, and whether data is being compressed or encrypted. Also, larger objects have higher rates of throughput.

Operation rate

The **Operations per Second** graph in the link status **Overview** panel for a replication link shows the history of the rate of operations on the link per second. An operation is the replication of any of these:

- An object, directory, symbolic link, metadata change, or object deletion

- An HCP tenant or HCP namespace or the modification or deletion of an HCP tenant or HCP namespace

- For HCP tenants only, the creation, modification, or deletion of a user account

- The creation, modification, or deletion of a retention class
- The creation, modification, or deletion of a content class
- A tenant log message

For an active/active link, the **Operations per Second** graph shows the operation rate for replication from the local system to the remote system. For an active/passive link, the graph shows the operation rate during replication or recovery, whichever is happening at the time. In any case, the operation rate is cumulative for all the tenants being replicated or recovered on the link.

If the **Operations per Second** graph is not currently visible, click on **Operations per Second** to display it.

The x-axis in the **Operations per Second** graph marks the passage of time. It shows 30 days (or fewer if the replication link was created less than 30 days ago). 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).

The graph heading displays the current rate of replication operations. Factors that affect the operation rate are the same as those that affect the data transmission rate.

Remote storage space

The **Remote Total Primary Storage** pie chart in the link status **Overview** panel for an active/active link or an outbound active/passive link during replication shows the amounts of free and used space in primary storage for the other system involved in the link. Each amount is also shown as a number of bytes.

In a SAIN system with spindown storage, the storage being measured in this chart includes both primary running storage and primary spindown storage. In this case, a second pie chart to the right, labeled **Remote Primary Running Storage**, shows the amounts of free and used space in only primary running storage.

The two systems involved in a replication link do not necessarily have the same storage capacity or storage tiering strategies. If a system receiving replicated data starts to run low on free space in primary running storage, you need to evaluate whether to take action on that system to address the issue (for example, by increasing the primary running storage capacity or by tiering more data to other types of storage).

When primary running storage in the remote system for a replication link is almost full, the **Overview** page in the HCP System Management Console for the local system displays an alert. When the storage is full, the Console displays a different alert, and HCP automatically suspends replication on the link. For more information on these alerts, see "System Management Console alerts" on page 113.

For information on different types of storage, see *Administering HCP*.

Tenants view

To monitor the status of replication or recovery activity for the individual tenants included in a given replication link, you use the link status **Tenants** panel in the HCP System Management Console. To open this panel:

1. In the list of links on the replication **Overview** page, click on the name of the link.

2. On the page for the individual link, click on **Status**.

3. In the **Status** panel, click on the **Tenants** tab.

Understanding the tenant list

The link status **Tenants** panel lists the tenants included in a replication link, including the default tenant if the link includes any default-namespace directories. For each tenant, the list shows:

- **Name** — The tenant name.

- **Sent as Of** — For an active/active link, an outbound active/passive link during replication, or an inbound active/passive link during data recovery, the average of the sent-as-of times for the namespaces owned by the tenant. The sent-as-of time for a namespace is the amount of time that is the difference between:

    - The date and time before which the namespace content is guaranteed to have been sent from the local system to the remote system

- The current date and time

For the default tenant, this is the average sent-as-of time among the directories included in the link.

Tips:

- If some tenants are significantly less up to date than others, set the replication priority for the link to **Oldest Object First** to reduce the gap.
- If the sent-as-of time is consistently increasing for one or more tenants, even with the replication performance level set to **High**, you may need to add more processing power (for example, additional nodes) to your system or increase the bandwidth between the two systems involved in the link.

- **Received as Of** — For an active/active link, an outbound active/passive link during data recovery, or an inbound active/passive link during replication, the average of the received-as-of times for the namespaces owned by the tenant. The received-as-of time for a namespace is the amount of time that is the difference between:

  - The date and time before which the namespace content is guaranteed to have been sent from the remote system to the local system
  - The current date and time

  For the default tenant, this is the average received-as-of time among the directories included in the link.

- **Status** — The current status of replication or recovery activity for the tenant, represented by one or more of these icons, as applicable:

  - ••• — If the icon is animated, replication or recovery of the tenant is proceeding normally. If the icon is static, no replication or recovery is occurring on the replication link.
  - ⏸ — A user paused replication or recovery of the tenant.
  - ⏸ — The replication service automatically paused replication or recovery of the tenant. To view text describing the condition that's causing this status, mouse over the status icon.

For information on events that can cause replication or recovery of a tenant to be paused automatically, see “Automatically paused tenant replication or recovery” on page 123.

- 🔗 — The tenant is included in a link that’s chained into the link you’re viewing.

Managing the tenant list

By default, the tenant list in the **Tenants** panel includes all tenants in the link. The tenants are listed 20 at a time in ascending order by tenant name.

You can page through, sort, and filter the list of tenants. The **Tenants** panel indicates which tenants are shown out of the total number of tenants in the current list.

Paging

You can change the number of tenants shown at a time in the **Tenants** panel. To do this, in the **Items per page** field, select the number of tenants you want. The options are 10, 20, and 50.

To page forward or backward through the tenant list, click on the next (▶) or back (◀) control, respectively.

To jump to a specific page in the tenant list:

1. In the **Page** field, type the page number you want.

2. Press Enter.

Sorting

You can sort the tenants in the tenant list by tenant name, outbound or inbound synchronization time, or status. To sort the list:

- By name, click on the **Name** column heading. Each time you click on the column heading, the sort order switches between ascending and descending.

- By outbound or inbound synchronization time, click on the **Outbound Synchronization** or **Inbound Synchronization** column heading, as applicable. The first time you click on the column heading, the tenants are sorted in ascending order by the applicable synchronization time. Each time you click on the column heading after that, the sort order switches between ascending and descending.

- By status, click on the **Status** column heading. The first time you click on the column heading, tenants for which replication has been automatically paused are listed first, followed by tenants for which replication has been manually paused, followed by the remaining tenants. Each time you click on the column heading after that, the order reverses from what it currently is.

Filtering

You can filter the tenant list by tenant name or tag. The filtered list includes only those tenants with a name or tag, as applicable, that begins with or is the same as a specified text string.

To filter the tenant list:

1. In the field above the **Name** column, select **Name** to filter by name or **Tag** to filter by tag.

2. In the next field, type the text string you want to use as a filter. This string can be up to 64 characters long, can contain any valid UTF-8 characters except commas (,), and is not case sensitive. White space is allowed.

3. Click on the find control (🔍).

To redisplay the entire list of tenants after filtering it, click on the clear filter control (✖).

Tenant replication details

To see more details about the replication of a tenant, mouse over the row for the tenant in the list of tenants in the link status **Tenants** panel. The information shown is:

- **Most up-to-date namespace**:

    ○ For an active/active link, an outbound active/passive link during replication, or an inbound active/passive link during data recovery, the smallest sent-as-of time among the namespaces owned by the tenant

    ○ For an outbound active/passive link during data recovery or an inbound active/passive link during replication, the smallest received-as-of time among the namespaces owned by the tenant

For the default tenant, this is the smallest sent-as-of time or received-as-of time, as applicable, among the directories included in the replication link.

- **Average up-to-date time**:
  - For an active/active link, an outbound active/passive link during replication, or an inbound active/passive link during data recovery, the average of the sent-as-of times for the namespaces owned by the tenant
  - For an outbound active/passive link during data recovery or an inbound active/passive link during replication, the average of the received-as-of times for the namespaces owned by the tenant

  For the default tenant, this is the average of the sent-as-of times or received-as-of times, as applicable, for the directories included in the replication link.

- **Objects pending** — For an active/active link, an outbound active/passive link during replication, or an inbound active/passive link during data recovery, the approximate number of objects currently waiting to be sent to the remote system in the namespaces owned by the tenant.

  This information is not available for an outbound link during data recovery or for an inbound link during replication. This information is also not available for the default tenant.

- **Data pending** — For an active/active link, an outbound active/passive link during replication, or an inbound active/passive link during data recovery, the approximate amount of data currently waiting to be sent to the remote system in the namespaces owned by the tenant.

  This information is not available for an outbound link during data recovery or for an inbound link during replication. This information is also not available for the default tenant.

Controlling the Tenant Management Console replication display

For HCP tenants that are included in at least one replication link, you control whether the Tenant Management Console displays information about the status of replication of the tenant and its namespaces. The Tenant Management Console for the default tenant always displays this information.

To control whether the Tenant Management Console displays replication status information for an HCP tenant and its namespaces:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the left side of the **Replication** page, click on **Settings**.

4. On the replication **Settings** page, take either of these actions:

   - To have the Tenant Management Console display replication status information for all HCP tenants, select the **Allow tenants to monitor replication of their namespaces** option.

   - To have the Tenant Management Console hide replication status information for all HCP tenants, deselect the **Allow tenants to monitor replication of their namespaces** option.

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

System Management Console alerts

The **Overview** page in the HCP System Management Console has a section in which it displays alerts about abnormal conditions. Each alert consists of an icon and accompanying text that identifies the problem. Each alert also has text that's displayed when you mouse over the icon.

The table below describes the alerts that relate to replication. The alerts are listed alphabetically by their mouse-over text.

Icon	Mouse-over text	Description
	Expired replication certificates	One or more trusted replication server certificates have expired. Replication with the systems from which the applicable certificates were obtained has stopped. When the other system in any of the applicable replication pairs installs a new SSL server certificate, download that certificate and upload it to this system as a trusted replication server certificate. For ease of maintenance, delete expired certificates from the list of trusted replication server certificates.

(Continued)

Icon	Mouse-over text	Description
	Front-end connection error	All front-end connections to one or more nodes in the local system are unavailable. If the system is currently a source for replication, some objects will not be replicated. If the system is a replication target, replication performance is degraded.
	Partial network assigned to replication	The network selected for replication is a partial network. The load from replication activity is not being fully distributed across the nodes in the system. Check the replication **Settings** page to see which network is selected for replication. Then either assign IP addresses in those networks to all nodes, or select a different network for replication.
	Remote storage almost full	Primary running storage space on the remote system for a replication link is at least 90% used. Consider adding more primary running storage capacity to the remote system or changing storage tiering strategies on that system to create more space in the existing primary running storage. For an active/active link, this alert appears only in the System Management Console for the local system. For an active/passive link, this alert appears only in the System Management Console for the primary system.
	Remote storage full (link suspended)	Primary running storage space on the remote system for a replication link is 94% used. The remote system cannot act on any more replication data transmissions from the local system. HCP has automatically suspended activity on the link. Either add more primary running storage capacity to the remote system, or change storage tiering strategies on that system to create more space in the existing primary running storage. Then resume activity on the link. For an active/active link, this alert appears only in the System Management Console for the local system. For an active/passive link, this alert appears only in the System Management Console for the primary system.

(Continued)

Icon	Mouse-over text	Description
✓	Replication certificates expire soon	One or more trusted replication server certificates will expire within 90 days. If a certificate expires, replication with the system from which the certificate was obtained will fail. To ensure that replication is not disrupted, when the other system in the replication pair installs a new SSL server certificate, download that certificate and upload it to this system as a trusted replication server certificate.
✓	Replication certificates expire soon	One or more trusted replication server certificates will expire within 30 days. If a certificate expires, replication with the system from which the certificate was obtained will fail. To ensure that replication is not disrupted, when the other system in the replication pair installs a new SSL server certificate, download that certificate and upload it to this system as a trusted replication server certificate.
X	Replication link failure	An active/active or outbound active/passive replication link is not working as expected. Check the network connection between the two systems involved in the link. If the connection appears to be working properly, contact your authorized HCP service provider for help.
II	Replication link has autopaused tenants	HCP has automatically paused replication of one or more HCP tenants. For each tenant, correct the situation that caused replication to be paused. Then resume replication of the tenant.
▶	Replication link pending	Either the system has sent a request for a replication link to another system and is waiting for a response, or the system has received a request for a replication link and has not yet responded.
?	Replication link stalled	Replication has unexpectedly stopped on a replication link. Check the network connection between the two systems involved in the link. If the connection appears to be working properly, contact your authorized HCP service provider for help.
!	Replication links shut down	All replication links in which this system is involved are shut down. No replication or recovery activity is occurring on these links, and the links cannot be used for read-from-remote or object-repair purposes.

(Continued)

Icon	Mouse-over text	Description
🕒	Time out of sync between replicating systems	The system time on this system is more than one minute out of sync with the system time on one or more other systems with which this system participates in a replication link. The recommended practice is to have all HCP systems in a replication topology use the same external time server.

For more information on alerts, see *Administering HCP*.

5

Managing replication

To help manage the load on an HCP system and the replication network, you can control the performance level for send activity on each replication link in which the system participates. You can use a schedule to change the performance level for a link automatically at specific times on a weekly basis. Alternatively, you can choose a single performance level for the entire week.

Occasionally, you may need to temporarily stop all send activity on an individual link. Or, you may want to temporarily stop send activity only for particular tenants.

Rarely, you may need to stop all activity on all links in which the HCP system participates. This action stops not only replication and recovery but also read and repair from remote. It also prevents you from changing which items are included in a link.

If the HCP system uses virtual networking, you can select the network to use for communications through any replication link. You typically do this only once.

You can control whether the HCP system allows DNS failover to other systems in the replication topology. Disallowing DNS failover prevents the system from servicing requests redirected from remote systems regardless of whether this is allowed by the targeted namespace.

You can configure an HCP system to automatically share its domains and SSL server certificates with other systems with which it participates in replication links. This ensures that SSL works for access to replicated namespaces on those systems.

This chapter explains how to perform the tasks outlined above.

Roles: To view replication link schedules and global replication settings, you need the monitor role. To modify replication link schedules, manage replication at the link and tenant levels, and set global replication options, you need the administrator role.

Note: You can also use the HCP management API to manage replication. For information on doing this, see *HCP Management API Reference*.

Scheduling activity on a replication link

The amount of send activity that can occur on a replication link at any given time is controlled by the performance level that applies to the sending system. An active/active link has two performance levels at any given time, one for each system involved in the link. An active/passive link has only one performance level at any given time. That level applies to the primary system during replication and to the replica during data recovery.

The performance levels for different replication links are independent of each other, as are the performance levels for the two systems involved in an active/active link. In each case, the performance level can be low, medium, high, custom, or off. Off means that no send activity is occurring for the applicable link on the applicable system.

For each system involved in an active/active link and for an active/passive link, you can schedule the performance level to change automatically over the course of a week. At any time, if you don’t want to use the schedule, you can override it by selecting a single performance level to apply until you cancel the override.

Overriding a replication schedule lets you change the performance level without changing the configured schedule. You might do this, for example, on a holiday when the load on the applicable HCP system is expected to be light.

Note: When data recovery begins on an active/passive link, a schedule override with a performance level of high takes effect. When data recovery is complete, the performance level automatically returns to the currently scheduled level if the schedule was in effect when data recovery began or to the override level that was in effect when data recovery began, as applicable.

The performance level for a sending system determines the amount of load replication or recovery processing puts on that system and on the network connecting that system to the other system involved in the applicable replication link. If a system is sending data on multiple links, the total load on the system is determined by the performance levels for all of those links together. To minimize the load, you should schedule the links to be active at different times.

Tip: If the system load from other activities is light, consider raising the performance level for one or more links to reduce any replication backlog. If the system load is heavy, consider lowering the performance level for one or more links to free resources for other system activity.

For more detailed information on performance levels, see “Changing the custom performance level” on page 166.

About replication schedules

You use the **Schedule** panel on the **Replication** page in the HCP System Management Console to set the schedule for replication activity on a replication link. For an active/active link, this panel has two tabs on each system involved in the link — one labeled **Local** for the local system and one labeled **Remote** for the remote system.

For an active/passive link, you set the schedule only for the primary system. Therefore, on the primary system the **Schedule** panel has a tab labeled **Local**. On the replica, the **Schedule** panel has a tab labeled **Remote**.

Each **Local** and **Remote** schedule panel contains a grid in which the weekdays from Sunday through Saturday are each broken out into 24 hours. To set a schedule, you assign performance levels to time periods in the grid. The Console makes it easy for you to do this for the whole week, individual days, individual hours, or ranges of hours within a day.

In the schedule grid, the top of each time period with a given performance level displays the start and end times for that period (for example, **8 am to 6 pm**). These times are in the time zone of the primary system.

The top of each time period also displays the performance level for that period. Additionally, each time period is color coded in the schedule grid to indicate the performance level so you can easily see which levels are assigned to which periods:

- Low: 🟩 (light green)
- Medium: 🟩 (green)

- High: 🟩 (dark green)
- Custom: 🟦 (blue)
- Off: ⬜ (gray)

While a schedule is overridden, the words **Schedule Overridden** appear on the schedule grid.

You cannot set a schedule to off for the entire week. Instead, to disable activity on a replication link for an extended period of time, suspend activity on the link. For information on suspending activity on a link, see “Suspending and resuming activity on an individual link” on page 122.

Modifying a replication schedule

By default, the schedules for an active/active link and the schedule for an active/passive link specify a performance level of medium for the entire week. To change a schedule for a replication link:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.
2. In the secondary menu, click on **Replication**.
3. On the replication **Overview** page, click on the link name in the list of links.
4. In the panel that opens, click on the **Schedule** tab.
5. In the **Schedule** panel, click on the **Local** or **Remote** tab, as applicable.
6. In the **Local** or **Remote** panel, as applicable, take any of these actions as many times as needed to set the schedule you want:
   - To set the performance level for the entire week:
     1. Mouse over **All** to display the list of performance levels.
     2. Click on the performance level you want.
   - To set the performance level for an individual day:
     1. Mouse over the name of the day to display the list of performance levels.
     2. Click on the performance level you want.

- To set the performance level for a single hour or a range of hours:

  1. Either click on an hour, or click and drag from one hour to another in the same day.

     The **Set Performance Level** window opens.

  2. Optionally, select different start and end times in the **Start time** and **End time** fields, respectively.

  3. In the **Level** field, select the performance level you want.

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

7. Click on the **Update Schedule** button.

Overriding a replication schedule

To override a schedule for a replication link or to cancel a schedule override:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the replication **Overview** page, click on the link name in the list of links.

4. In the panel that opens, click on the **Schedule** tab.

5. In the **Schedule** panel, click on the **Local** or **Remote** tab, as applicable.

6. In the **Local** or **Remote** panel, as applicable, click on **Schedule Override**.

7. In the **Schedule Override** section, take either of these actions:

   - To override the schedule:

     1. Select the **Override schedule** option.

     2. Under **Performance Level**, select the performance level you want.

   - To cancel the schedule override, deselect the **Override schedule** option.

8. Click on the **Update Schedule** button.

Suspending and resuming activity on an individual link

You can suspend and resume replication or recovery activity on individual replication links. When you suspend activity on a link, HCP stops all send activity on the link. You might suspend activity on a link, for example, before making changes to system hardware or to the network over which the two systems involved in the link communicate with each other.

While send activity is suspended on a link, the applicable HCP tenants and namespaces and default-namespace directories on each system remain read-write or read-only, as applicable. Additionally, the link continues to support other functions such as read and repair from remote.

The replication service periodically checkpoints its progress. When you suspend activity on a link, no special checkpoint occurs. When you resume link activity, therefore, processing starts from the last checkpoint before the suspension.

To suspend or resume activity on a replication link:

1. In the top-level menu in the System Management Console for either the primary system or the replica, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the replication **Overview** page, click on the name of the link on which you want to suspend or resume activity.

4. In the panel that opens, click on the **Link** tab.

5. In the replication **Link** panel, click on the **Management** tab.

6. In the link **Management** panel, click on the **Suspend** or **Resume** button, as applicable.

Pausing and resuming replication or recovery of a tenant

You can pause and resume replication or recovery for an individual tenant that's included in a replication link. You might pause replication for some tenants, for example, to give more processing time to other tenants with greater backlogs. When you pause replication for a tenant on an active/active link, replication of the tenant stops in both directions on the link.

Replication or recovery of a tenant can also be paused automatically due to certain events. After replication or recovery of a tenant is paused automatically, you need to resume replication or recovery manually. However, you cannot do this until the issue that caused replication or recovery to be paused is resolved. For information on events that can cause tenant replication or recovery to be paused automatically, see "Automatically paused tenant replication or recovery" on page 123.

The replication service periodically checkpoints its progress. When you pause replication or recovery for a tenant, no special checkpoint occurs. When you resume processing, therefore, processing starts from the last checkpoint before the pause.

Pausing and resuming tenant replication or recovery

To pause or resume replication or recovery for an individual tenant:

1. In the top-level menu in the System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the replication **Overview** page, click on the name of the link that includes the tenant for which you want to pause or resume replication or recovery.

4. In the link **Status** panel that opens, click on the **Tenants** tab.

5. In the list of tenants, click on the pause control (⏸) or resume control (▶), as applicable, for the tenant for which you want to pause or resume activity.

Automatically paused tenant replication or recovery

Certain events can cause the replication service to automatically pause replication or recovery of a tenant on a replication link.

Tip: To avoid situations that can cause the replication service to automatically pause replication of a tenant, do not try to create the same tenants, namespaces, content classes, user accounts, and group accounts on both of the systems that are replicating to each other on an active/active link. Instead, allow the items you create on each system to replicate to the other system.

Tenant name collisions

A tenant name collision occurs when the replication service tries to replicate an HCP tenant from one HCP system to another HCP system that already has a different tenant with the same name. To recover from a tenant name collision, you need to either rename the tenant on one of the systems involved in the link or delete the tenant on the receiving system.

Here are two scenarios that show how a tenant name collision can cause the replication service to pause replication of a tenant.

Scenario 1

In this scenario:

- System A replicates to system B on link AB. Link AB can be either active/active or active/passive.
- System A has a tenant named T1 that is not included in link AB.

These events occur in the order shown:

1. On system A, you add T1 to link AB.
2. Before T1 is replicated to system B, you create a tenant named T1 on system B.
3. The replication service tries to replicate T1 to system B. The replication is unsuccessful because a different tenant named T1 already exists on system B. As a result, the service automatically pauses replication of T1 on link AB.

Scenario 2

In this scenario:

- System A replicates to system B on link AB, which replicates to system C on link BC, where link AB is chained into link BC. Link AB can be either active/active or active/passive. Link BC is active/passive.
- System A and system C each have a tenant named T1, where T1 was created independently on each system.

These events occur in the order shown:

1. On system A, you add T1 to link AB.
2. T1 is replicated to system B.

3. Because link BC includes link AB, the replication service tries to replicate T1 to system C. The replication is unsuccessful because a different tenant named T1 already exists on system C. As a result, the service automatically pauses replication of T1 on link BC.

Namespace name collisions

Each HCP namespace you create in an HCP system has an internal ID that uniquely identifies it. As a result, two namespaces created on different systems are different from each other, even if they have the same name and are owned by the same tenant.

A namespace name collision occurs when the replication service tries to replicate a namespace from one system to another system that already has a different namespace with the same name, where both namespaces are owned by the same tenant.

Here's a scenario that shows how a namespace name collision can cause the replication service to pause replication of a tenant. In this scenario:

- System A and system B replicate to each other over active/active link AB.
- Link AB includes tenant T1, so T1 exists on both systems.
- On system A, T1 owns namespace NS1, which is not selected for replication.

These events occur in the order shown:

1. On system A, you select NS1 for replication.
2. Before NS1 is replicated to system B, you create a namespace named NS1 for T1 on system B.
3. The replication service tries to replicate NS1 to system B. The replication is unsuccessful because a different namespace named NS1 already exists on system B. As a result, the service automatically pauses replication of T1 on link AB.

To recover from a namespace name collision, you can take any of these actions:

- Rename the namespace on one of the systems involved in the link.
- Deselect the namespace from replication.

- Delete the namespace on the receiving system.

Namespace compliance issues

Different HCP systems can have different definitions for service plans with the same name. When an HCP tenant or namespace is replicated, the name of its associated service plan, not the service plan itself, is replicated with it. As a result, the service plan that applies to a namespace can differ on the two HCP systems involved in a link on which the namespace is being replicated.

Service plans can be compliant or noncompliant. However, the service plan that applies to a namespace in compliance mode must be compliant. A namespace compliance issue occurs when replication of a namespace in compliance mode would cause a noncompliant service plan to apply to the namespace on the receiving system.

To recover from namespace compliance issue, you can take any of these actions:

- Redefine the noncompliant service plan on the receiving system to be compliant.

- If the service plan is assigned to the tenant that owns the namespace, assign a different service plan to the tenant on the sending system, where that service plan is complaint on both systems involved in the link.

- If the service plan is assigned to the namespace, have the tenant administrator assign a different service plan to the namespace on the sending system, where that service plan is complaint on both systems involved in the link.

- Deselect the namespace from replication.

Release 6.*x* systems do not have the concept of compliant or noncompliant service plans. As a result, namespace compliance issues do not occur with cross-release replication where the replica is the 6.*x* system.

Here are two scenarios that show how a namespace compliance issue can cause the replication service to pause replication of a tenant.

Scenario 1

In this scenario:

- System A replicates to system B on link AB. Link AB can be either active/active or active/passive.
- Link AB includes tenant T1, so T1 exists on both systems.
- On system A, T1 owns namespace NS1, which is in compliance mode. NS1 not selected for replication.
- The service plan that applies to NS1 is named SP1. SP1 is compliant on system A and noncompliant on system B.

These events occur in the order shown:

1. On system A, you select NS1 for replication.
2. The replication service tries to replicate NS1 to system B. The replication is unsuccessful because it would cause NS1, which is in compliance mode, to have a noncompliant service plan on system B. As a result, the service automatically pauses replication of T1 on link AB.

Scenario 2

In this scenario:

- System A replicates to system B on link AB. Link AB can be either active/active or active/passive.
- Link AB includes namespace NS1, which is owned by tenant T1, so NS1 exists on both systems.
- NS1 is in enterprise mode, not compliance mode.
- The service plan that applies to NS1 is named SP1. SP1 is compliant on system A and noncompliant on system B.

These events occur in the order shown:

1. On system A, you change NS1 to be in compliance mode.
2. The replication service tries to replicate the change to system B. The replication is unsuccessful because it would cause NS1 to be in compliance mode with a noncompliant service plan on system B. As a result, the service automatically pauses replication of T1 on link AB.

Content class collisions

Each content class you create in an HCP system has an internal ID that uniquely identifies it. As a result, two content classes created on different HCP systems are different from each other, even if they have the same name and are defined for the same tenant.

A content class collision occurs when the replication service tries to replicate a content class from one system to another system that already has a different content class with the same name, where both content classes are defined for the same tenant.

Here’s a scenario that shows how a content class collision can cause the replication service to pause replication of a tenant. In this scenario:

- System A and system B replicate to each other over active/active link AB.
- Link AB includes tenant T1, so T1 exists on both systems.

These events occur in the order shown:

1. On system A, you create a content class named CC1 for T1.
2. Before CC1 is replicated to system B, you create a content class named CC1 for T1 on system B.
3. The replication service tries to replicate CC1 to system B. The replication is unsuccessful because a different content class named CC1 already exists on system B. As a result, the service automatically pauses replication of T1 on link AB.

To recover from a content class collision, you can take any of these actions:

- Rename the content class on either of the systems involved in the link.
- Delete the content class on either of the systems involved in the link.

User account collisions

Each user account you create in an HCP system has an internal ID that uniquely identifies it. As a result, two user accounts created on different systems are different from each other, even if they have the same username and are defined for the same HCP tenant.

A user account collision occurs when the replication service tries to replicate a user account from one system to another system that already has a different user account with the same username, where both user accounts are defined for the same tenant.

Here's a scenario that shows how a user account collision can cause the replication service to pause replication of a tenant. In this scenario:

- System A and system B replicate to each other over active/active link AB.
- Link AB includes tenant T1, so T1 exists on both systems.

These events occur in the order shown:

1. On system A, you create a user account with username U1 for T1.
2. Before U1 is replicated to system B, you create a user account with username U1 for T1 on system B.
3. The replication service tries to replicate U1 to system B. The replication is unsuccessful because a different user account with username U1 already exists on system B. As a result, the service automatically pauses replication of T1 on link AB.

To recover from a user account collision, you can take either of these actions:

- Change the username for the user account on either of the systems involved in the link.
- Delete the user account on either of the systems involved in the link.

Group account collisions

Each HCP group account you create in an HCP system has an internal ID that uniquely identifies it. As a result, two group accounts created on different systems are different from each other, even if they are created from the same Active Directory group and are defined for the same HCP tenant. (An HCP group account always has the same name as the AD group it's created from, so group accounts created from the same AD group on two different systems have the same name as each other.)

A group account collision occurs when the replication service tries to replicate a group account from one system to another system that already has a different group account created from the same AD group, where both group accounts are defined for the same tenant.

Here's a scenario that shows how a group account collision can cause the replication service to pause replication of a tenant. In this scenario:

- System A and system B replicate to each other over active/active link AB.
- Link AB includes tenant T1, so T1 exists on both systems.

These events occur in the order shown:

1. On system A, you create an HCP group account for T1 from the AD group named AD1. The name of the group account you create is AD1.
2. Before AD1 is replicated to system B, you create an HCP group account for T1 from the AD group named AD1 on system B. The name of the group account you create is AD1.
3. The replication service tries to replicate the HCP group account named AD1 to system B. The replication is unsuccessful because a different group account named AD1 already exists on system B. As a result, the service automatically pauses replication of T1 on link AB.

To recover from a group account collision, you can delete the group account on either of the systems involved in the link.

Shutting down and reestablishing all replication links

Shutting down replication links on an HCP system stops all activity on all links in which the system participates. You cannot shut down an individual link.

While links are shut down:

- No replication or recovery activity occurs on the links.
- The system on which the links are shut down cannot read or repair objects from other HCP systems.
- Other HCP systems cannot read or repair objects from the system on which the links are shut down.
- The system on which the links are shut down can still service requests that are redirected from other HCP systems in the replication topology.

- You cannot change which HCP tenants and namespaces, default-namespace directories, and inbound links are included in the links.

- An alert indicating that all links are shut down appears on the **Overview** page in the System Management Console for the system on which the links are shut down.

- A message indicating that all links are shut down appears at the top of the **Replication** page in the System Management Console for the system on which the links are shut down.

You may want to shut down replication links, for example, if you need to temporarily dedicate as much network bandwidth as possible to applications that are unrelated to HCP. Additionally, your authorized HCP service provider may ask you to shut down all links for certain system upgrade scenarios.

When you shut down all replication links, you are required to specify a reason for the action.

To restart activity on replication links after you've shut them down, you need to reestablish the links. Reestablishing the links returns each link to the state it was in before you shut down the links. If a link was active when you shut down all links, it becomes active again when you reestablish the links. If a link was suspended when you shut down all links, it remains suspended when you reestablish the links.

To shut down or reestablish all replication links on an HCP system:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the left side of the **Replication** page, click on **Settings**.

4. On the replication **Settings** page, take either of these actions:

   - To shut down all replication links:

     1. Click on the **Shut Down All Links** button.

        The **Shut Down All Replication Links** window appears.

     2. In the **Reason** field, type the reason why you're shutting down all links. This text can be up to 1,024 characters long and can contain any valid UTF-8 characters, including white space.

3. Click on the **Shut Down All Links** button.

- To reestablish all replication links, click on the **Reestablish All Links** button.

Selecting the network for replication

In an HCP system that uses virtual networking, you can define multiple networks and then select the networks to use for various purposes, including replication. The network you select for replication on a given system is used for both incoming and outgoing replication traffic.

For information on virtual networking with HCP, see *Administering HCP*.

Networking infrastructure

Different HCP systems can use different networks for the purpose of replication. When you select a replication network for a given system, you need to ensure that your networking infrastructure is configured to allow communications to be routed between that system and all other systems with which that system participates in a replication link.

IP mode

The two systems involved in a replication link must be able to use the same IP mode to communicate with each other. That is, either both networks must be configured with IPv4 addresses, or both networks must be configured with IPv6 addresses.

One or both of the networks involved can have both IPv4 and IPv6 addresses. If the replication networks for the two systems involved in a replication link have both types of addresses, HCP uses the IP addresses for the first mode in which it can establish communication between the two systems, with preference given to IPv6.

Within a replication topology, different pairs of systems can use different IP modes for communication. The figure below shows a replication chain in which different IP modes are used for communication over each link.

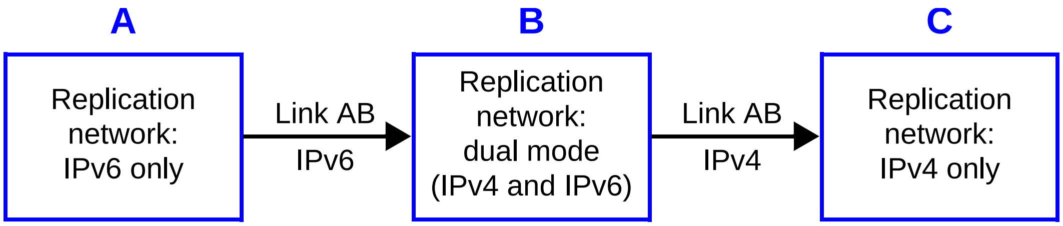
A
Replication network: IPv6 only
Link AB
IPv6
B
Replication network: dual mode (IPv4 and IPv6)
Link AB
IPv4
C
Replication network: IPv4 only

In the topology shown above:

- System A replicates to system B on link AB.
- System B replicates to system C on link BC.
- The replication network for system A has only IPv6 IP addresses.
- The replication network for system B has both IPv4 IP address and IPv6 addresses.
- The replication network for system C has only IPv4 IP addresses.
- System A and system B use IPv6 addresses to communicate with each other.
- System B and system C use IPv4 addresses to communicate with each other.

Loss of connectivity

You can select a different network for replication at any time. Selecting a different replication network on any given system can result in loss of connectivity to other systems with which that system is directly involved in replication.

When connectivity is lost between the two systems that participate in a replication link, the replication service automatically suspends activity on that link. After restoring connectivity, you need to manually resume activity on the link.

For any given replication link, connectivity is lost between the two systems involved when you select a different network on one of the systems and either of these apply:

- The networking infrastructure is not configured to route communications between the new network and the network selected for replication on the other system.
- The new network is associated with a different domain from the previously selected network. In this case, the SSL server certificate used for replication changes, and you need to share the new certificate with the other system involved in the link.

  Additionally, in this case, if the system on which you selected the new network is identified by domain name in the link configuration, you need to update the domain name in the link configuration.

In any case, if you select a different network on one system and the link identifies that system by IP addresses, you need to update those IP addresses in the link configuration.

If you need to both share new SSL server certificates and update the identification of one of the systems in the link configuration, you should share the certificates first. If you update the link first and then share the certificates, the link status changes to broken. To recover from this situation, you need to click again on the **Update Link** button in the **Settings** panel for the link.

For information on the SSL certificate change that can occur when you select a different network for replication, see "Sharing SSL server certificates" on page 58.

Shared domain name

When you create a replication link, if both of the following are true, you need to use the **replication.admin.***hcp-domain-name* format to identify the other system involved in the link:

- You are identifying the other system by domain name.
- The domain associated with the network selected for replication on the other system is also associated with another network.

In the configuration of an existing replication link, the domain name used to identify a system can have a format other than **replication.admin.***hcp-domain-name*. In this case, you need to update the system identification to use the **replication.admin.***hcp-domain-name* format if any of these happens on that system:

- You change the domain associated with the replication network to a domain that's also associated with another network.
- You select a new replication network, and the new network is associated with a domain that's also associated with another network.
- You change the domain associated with another network to the domain that's associated with the replication network.

Selecting the replication network

By default, HCP uses the [hcp_system] network for replication. To select a different network to be used for replication traffic:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the left side of the **Replication** page, click on **Settings**.

4. On the replication **Settings** page, in the **Replication Network** field, select the network you want to use for replication traffic. The dropdown list of networks does not include empty, degraded, or partial networks.

   The **Replication Network** field is present only while virtual network management is enabled for the HCP system.

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

Displaying the zone definition for the replication network domain

Each network in HCP is associated with a domain. Multiple networks can be associated with the same domain. If you're using DNS for domain name resolution, the DNS needs to include a zone definition for each combination of network and domain.

If an HCP system is involved in replication, the zone definition for the replication network domain for that system needs to be added to the upstream DNS servers for the other system on the replication link. An **upstream DNS server** is a DNS server to which HCP routes the outbound communications it initiates (for example, for sending log messages to syslog servers or for communicating with Active Directory).

If the domain for the replication network is shared with other networks, the domain name in the zone definition you add for the replication network domain must be:

replication.admin.*replicaton-network-domain-name*

From the **Replication** page, you can display the zone definition for the replication network domain, formatted for Unix DNS servers. You can then copy that definition to the upstream DNS servers.

Note: You can display the zone definition only while virtual network management is enabled for the HCP system and only if at least one user-defined network exists.

To display the zone definition:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the left side of the **Replication** page, click on **Settings**.

4. On the replication **Settings** page, click on **Zone Definition**.

   The **Zone Definition** section opens. This section shows the zone definition for the network (with a type of slave or stub, as applicable), formatted as shown in this example:

   ```
   # net1 Special-purpose HCP system network
   zone "replication.admin.system_1.example.com" IN {
       type slave;
       file "/var/named/slave/replication.admin.system_1.example.com";
       masters {
           172.25.147.11;
           172.25.147.12;
           172.25.147.13;
           172.25.147.14;
       };
   };
   ```

For more information on defining zones for HCP domains in the DNS, see *Administering HCP*.

Managing DNS failover

HCP replication helps ensure continuous namespace availability. If a system in a replication topology fails, client requests for access to a replicated namespace can be serviced by other systems in the topology.

You can use the HCP DNS failover feature to automate the process of redirecting client requests from a failed system to a healthy one. For active/active links, this process can be automated in other ways as well. Automating redirection means that applications do not need to be modified to explicitly direct requests to another system when the normally targeted system fails.

The DNS failover feature requires that the HCP systems involved use a shared DNS for system addressing. If the systems don't use a shared DNS, DNS failover is not an option for automating redirection of client requests.

An HCP system can service automatically redirected requests only if the target namespace is configured to support service by remote systems. If the namespace is not configured this way, requests to access the namespace on a failed system fail.

About DNS failover

DNS failover is an HCP system configuration option that, when enabled on the one system involved in a replication link, forces requests to that system to be automatically redirected to the other system involved in the link while the link is failed over to the other system. This redirection occurs only when the request identifies the target system by domain name, not by IP address.

In effect, DNS failover causes the domain name for the failed-over system to be associated with the IP addresses for the nodes in the other system. Therefore, all types of requests that specify that domain name are redirected to the other system. This includes not only requests for namespace access but also requests for access to HCP interfaces such as the Tenant Management Console and HCP management API.

An HCP system can service redirected requests only if they come in through a namespace access protocol. This means that requests for access to the failed-over system that are made through other interfaces fail.

With an active/active link, failover can occur in either direction between the two systems involved in the link. Therefore, if you are using DNS failover for automatic redirection of client requests, you should enable it on both systems.

With an active/passive link, failover can occur only from the primary system to the replica. In this case, therefore, you need to enable DNS failover only on the primary system. However, if the replica is also the primary system for another link, you need to enable DNS failover on the replica as well.

For DNS failover to work for the system on which it's enabled, the HCP domains for that system in the DNS must be configured to support service by remote systems. If DNS failover is not enabled, the HCP domains should not be configured that way.

DNS failover is intended to address cases of catastrophic failure of the HCP system on which it's enabled. However, DNS failover also applies if you fail over a link while the system is healthy. In this case, the method used to access nonreplicated items on that system depends on the data access network for the tenant that owns the target namespace.

For example, suppose:

- Tenants ten1 and ten2 both use the network named net1 for data access.

- Tenant ten1 and its namespace ns1 are in a replication link that is failed over from system A to system B.

- Tenant ten2 and its namespace ns2 are not in the failed-over replication link.

Client requests for access to ns1 on system A, where the request URL specifies the name of the domain associated with net1, are redirected to system B. Because they come in on the same network, client requests for access to ns2 on system A, where the request URL specifies the domain name, are also redirected to the system B and, therefore, fail. For those requests to succeed, they need to access system A by using an IP address assigned to a node in net1 on that system instead of by using the domain name.

The same consideration applies to access to other HCP interfaces. For example, if the data access network for a tenant in a link that's failed over from system A to system B is [hcp_system], you need to use an IP address to access the HCP System Management Console on system A.

DNS failover also affects replication between the failed-over system and any other system with which that system participates in a replication link. If the other system identifies the failed-over system by domain name, all replication activity on the link between the two systems stops.

For an introduction to failover with HCP replication, see “Failover and failback” on page 37. For information on configuring namespaces to accept redirected requests, see *Managing a Tenant and Its Namespaces* or *Managing the Default Tenant and Namespace*. For information on networks and configuring DNS for HCP, see *Administering HCP*.

Enabling or disabling DNS failover

To enable or disable DNS failover for an HCP system:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the left side of the **Replication** page, click on **Settings**.

4. On the replication **Settings** page, take either of these actions:

   - To enable DNS failover, select the **Enable DNS failover to other systems in the replication topology** option.

   - To disable DNS failover, deselect the **Enable DNS failover to other systems in the replication topology** option.

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

Alternatives to DNS failover

DNS failover is an HCP-specific method for automatically managing service by remote systems. With active/active replication, other options exist:

- In an environment in which load balancers are used to spread client requests among multiple HCP systems, if one of the systems fails, the load balancers can ensure that the requests go to other systems.

- In a cloud storage environment, the networking and DNS infrastructure can be configured to support multiple HCP systems that use the same domain names. In this configuration, client requests are normally handled by the local DNS, but if an HCP system fails, the request can be passed on to another DNS that’s local to another HCP system.

Automatically sharing domains and SSL server certificates

With DNS failover, client requests to a failed system that identify the system by domain name are automatically redirected to another system in the replication topology. If the client request uses HTTP with SSL security (HTTPS), the system to which the request is redirected must have an SSL server certificate for the domain specified in the request. Because the fully qualified domain name for a replicated namespace is different on different systems in a replication topology that uses DNS failover, the system to which the request is redirected would not normally have such a certificate.

An HCP system can be configured to periodically send its domains and SSL server certificates to each other system with which it participates as a sending system in a replication link. If the system targeted by an HTTPS request has shared its certificates with the system to which the request is redirected, that second system can service the request.

When a system is configured to share its domains and SSL server certificates, it sends all the domains and certificates it has, including those that were sent to it by another system. If you configure all the systems in a replication topology to share their domains and certificates, any system in the topology can service an HTTPS request redirected from any other system in the topology.

Sharing domains and SSL server certificates has an additional benefit. If an HCP system is rebuilt after a catastrophic failure, the domains and certificates originally created on that system can be recovered from another system with which they were shared before the failure.

Note: Automatically sharing domains and certificates is not supported for cross-release links (that is, links between 7.0 system and 6.*x* systems).

To enable or disable automatic sharing of domains and SSL server certificates for an HCP system:

1. In the top-level menu in the HCP System Management Console, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the left side of the **Replication** page, click on **Settings**.

4. On the replication **Settings** page, take either of these actions:

    - To enable automatic sharing of domains and certificates, select the **Send local domains and certificates to other systems in the replication topology** option.

    - To disable automatic sharing of domains and certificates, deselect the **Send local domains and certificates to other systems in the replication topology** option.

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

6

Managing failover and failback

When one system involved in an active/active link fails or when the primary system for an active/passive link fails, you can fail over the link to the other system involved in the link. Failing over the link stops replication on the link and, if DNS failover is enabled, allows client requests that target the failed system to be redirected to the healthy system.

If a system failure results in a replication link being broken (for example, due to the system being rebuilt after a catastrophic failure), you need to restore the link before replication can restart or data recovery can occur on that link. This applies regardless of the link type and, for an active/passive link, regardless of whether the failed system is the primary system or the replica.

To restart replication after failing over a link, you need to fail back the link. Failing back the link restarts replication on the link and returns the HCP systems involved to normal operation. For an active/passive link, failing back includes recovering data from the replica to the primary system.

Failover can be automated for both active/active and active/passive links. Failback can be partially automated for active/passive links.

This chapter provides instructions and considerations for managing failover and failback manually. For an overview of manual and automatic failover and failback with HCP replication, see "Failover and failback" on page 37.

Roles: To fail over a replication link, restore a link, or recover replicated data, you need the administrator role.

Note: You can also use the HCP management API to manage failover and failback of replication links. For information on doing this, see *HCP Management API Reference*.

Failover and failback workflows

Three failover and failback scenarios are possible, depending on which of these HCP systems fails:

- One of the systems involved in an active/active link
- The primary system for an active/passive link
- The replica for an active/passive link

The following sections describe the basic workflows for these scenarios. For information on failover and recovery workflows in more complex replication topologies, see [“Failover and failback considerations”](#) on page 152.

System failure workflow with an active/active link

The table below outlines what happens when one of the systems involved in an active/active link fails, where the system that fails is system A and the system that remains healthy is system B.

Step	What you do	What happens
System A fails		
1	On system B, fail over the link	If DNS failover is enabled, system B broadcasts new DNS configuration
2	If DNS failover is disabled, direct clients to write only to system B	
System A comes back online		
3	If system A has been rebuilt: • On system A, upload the replication SSL server certificate from system B • On system B, upload the replication SSL server certificate from system A	
4	On system B, update the link configuration as needed	
5	If the link is broken, on system B, send a request to restore the link	Replication link is recreated

(Continued)

Step	What you do	What happens
6	On system B, fail back the link	System A and system B broadcast original DNS configurations; replication restarts in both directions on the link

Primary system failure workflow

The table below outlines what happens when the primary system for an active/passive link fails.

Step	What you do	What happens
Primary system fails		
1	On the replica, fail over the link	Applicable tenants and directories on the replica become read-write; applicable tenants and directories on the primary system either remain read-write or become read-only depending on whether the two systems can communicate with each other; if DNS failover is enabled, the replica broadcasts new DNS configuration
2	If DNS failover is disabled, direct clients to write only to the replica	
Primary system comes back online		
3	If the primary system has been rebuilt: • On the primary system, upload the replication SSL server certificate from the replica • On the replica, upload the replication SSL server certificate from the primary system	
4	On the replica, update the link configuration as needed	
5	If the link is broken, on the replica, send a request to restore the link	For a cross-release link, pending link request appears on the primary system Otherwise, replication link is recreated

(Continued)

Step	What you do	What happens
6	For a cross-release link, on the primary system, accept the restored link	For a cross-release link, replication link is recreated
7	On the replica, begin data recovery	Applicable tenants and directories on the replica remain read-write; applicable tenants and directories on the primary system remain or become read-only; data recovery from the replica to the primary system begins
8	Wait for data recovery to come close to being up to date	
9	On the replica, complete data recovery	Applicable tenants and directories on the replica become read-only; applicable tenants and directories on the primary system remain read-only; data recovery from the replica to the primary system continues to completion
Data recovery finishes		
10	Nothing	Applicable tenants and directories on the replica remain read-only; applicable tenants and directories on the primary system become read-write; the primary system and the replica broadcast original DNS configurations; replication from the primary system to the replica restarts
11	If DNS failover is disabled, after you see this message in the system log, direct clients to write only to the primary system: *Replication data recovery completed*	

Replica failure workflow

The table below outlines what happens when the replica for an active/passive link fails.

Step	What you do	What happens
Replica fails		
1	On the primary system, suspend the link	

(Continued)

Step	What you do	What happens
Replica comes back online		
2	If the replica has been rebuilt: • On the replica, upload the replication SSL server certificate from the primary system • On the primary system, upload the replication SSL server certificate from the replica	
3	On the primary system, update the link configuration as needed	
4	If the link is broken, on the primary system, send a request to restore the link	For a cross-release link, pending link request appears on the replica Otherwise, replication link is recreated; applicable tenants and directories on the primary system remain read-write; applicable tenants and directories on the replica are read-only
5	For a cross-release link, on the replica, accept the restored link	For a cross-release link, replication link is recreated; applicable tenants and directories on the primary system remain read-write; applicable tenants and directories on the replica are read-only
6	On the primary system or the replica, resume the link	Replication from the primary system to the replica restarts from the beginning

Failing over

To fail over a replication link:

1. In the top-level menu in the HCP System Management Console for the system you want to fail over to, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the replication **Overview** page, click on the name of the link you want to fail over.

4. In the panel that opens, click on the **Link** tab.

5. In the replication **Link** panel, click on the **Failover** tab.

6. In the link **Failover** panel, click on the **Fail Over** button.

   A confirming message appears.

7. In the window with the confirming message, select **I understand** to confirm that you understand the consequences of your action. Then click on the **Fail Over Link** button.

8. For an active/passive link, if DNS failover is disabled, tell the applicable tenant administrators to direct all client access requests to the replica.

Recovering from a failure

During a catastrophic failure, an HCP system can lose all configuration information. In this case, if the system participates in a replication link, you need to restore the link configuration after the system is rebuilt. If the failure did not cause the system to lose the link configuration, you don't need to restore the link.

Once the link configuration exists on both systems involved in the link:

- For an active/active link, you need to perform the failback procedure.
- For an active/passive link:
  - If the primary system failed, you need to recover namespace content and other applicable information from the replica.

    Note: After you fail over an active/passive link, the only way to return to normal replication is by going through the data recovery procedure. You need to perform this procedure even if you don't need to restore the link and even if no changes have been made to the configuration or content of the replicated items. Even when nothing has change, the data recovery process can take more than five minutes.

  - If the replica failed, replication automatically restarts, beginning again with the objects with the oldest metadata changes either across all namespaces or within each namespace, depending on the link configuration.

Restoring a link

Before you can restore a replication link, both the primary system and the replica must have the required SSL certificates installed, as described in Chapter 2, "Configuring SSL for replication," on page 57.

To restore a replication link after a system failure:

1. In the top-level menu in the HCP System Management Console for the system on which the link configuration still exists, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the replication **Overview** page, click on the name of the link you want to restore.

4. In the panel that opens, click on the **Link** tab.

5. If necessary, update the link configuration. You need to do this only if the domain name or IP addresses of the other system have changed.

   If the domain name of the other system has not changed, do not replace the displayed IP addresses with the domain name.

   For instructions on updating the link configuration, see "Modifying replication link settings" on page 93.

6. In the replication **Link** panel, click on the **Failover** tab.

7. In the **Failover** panel, click on the **Restore Link** button.

   Note: If you don't see the **Restore Link** button, wait a few minutes and then redisplay the **Replication** page. Before the System Management Console can display this button, HCP needs to recognize that the other system is available, has the necessary SSL certificates installed, and isn't aware of the existing link.

   For a cross-release link, a pending link request appears on the other system involve in the link. You respond to this request the same way you do to an initial request to create a link. For information on responding to a link request, see "Handling cross-release link requests" on page 70.

Failing back an active/active link

To fail back an active/active link:

1. In the top-level menu in the HCP System Management Console for the system you failed over to, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the replication **Overview** page, click on the name of the link you want to fail back.

4. In the panel that opens, click on the **Link** tab.

5. In the replication **Link** panel, click on the **Failover** tab.

6. In the link **Failover** panel, click on the **Fail Back** button.

Recovering the data after a primary system failure

To recover namespace content and other applicable information from a replica to a primary system:

1. In the top-level menu in the HCP System Management Console for the replica, mouse over **Services** to display a secondary menu.

2. In the secondary menu, click on **Replication**.

3. On the replication **Overview** page, click on the name of the link on which you want to recover data.

4. In the panel that opens, click on the **Link** tab.

5. In the replication **Link** panel, click on the **Failover** tab.

6. In the link **Failover** panel, click on the **Begin Recovery** button.

Note: After uploading new trusted replication server certificates, you may need to wait more than ten minutes for the **Begin Recovery** button to become active.

The applicable HCP tenants and default-namespace directories become read-only on the primary system, and the replication service starts copying the applicable objects and configuration information from the replica to the primary system. As with replication from the primary system to the replica, the service starts with the objects with the oldest metadata changes either across all namespaces or within each namespace, depending on the link configuration.

Note: If the primary system cannot communicate with Active Directory and either of these is true for a tenant, recovery of that tenant is automatically paused:

- The tenant to be recovered supports AD authentication.
- A namespace owned by the tenant supports AD single sign-on.

When communication between the primary system and AD is restored, you can resume recovery of the tenant.

7. Monitor the recovery process by periodically reviewing the information in the status **Overview** and status **Tenants** panels for the link.

8. When data recovery is almost synchronized with current tenant and namespace activities on the replica, return to the **Failover** panel for the link. Synchronization is nearing completion when the up-to-date-as-of time for the link is close to zero.

Note: As long as clients continue writing to the replica, synchronization won’t reach one hundred percent. Synchronization doesn’t need to be completely up to date for you to start the complete recovery phase.

9. In the link **Failover** panel, click on the **Complete Recovery** button.

   The applicable tenants and directories on the replica immediately become read-only. The tenants and directories on both systems then remain read-only until the replication service finishes the data recovery. The amount of time this takes depends on how much data is left to recover.

   When recovery is complete, the tenants and directories on the primary system become read-write, those on the replica remain read-only, and the replication service on the primary system starts copying objects to the replica again.

Tips:

- You can schedule completion of the data recovery process for a time when client usage of the repository is low.
- If, before final recovery is complete, you need to allow clients to write to the applicable tenants and directories on the replica again, click on the **Cancel Final Recovery** button in the link **Failover** panel. The recovery process continues, but the applicable tenants and directories become read-write on the replica and remain read-only on the primary system until you click on the **Complete Recovery** button again.

10. If DNS failover is disabled:

    a. Wait for this message to appear in the system log:

    ```
    Replication data recovery completed
    ```

    b. Tell the applicable tenant administrators to redirect all client access requests to the primary system.

Failover and failback considerations

These basic rules apply to replication, failover, and failback on replication links, regardless of the replication topology:

- Multiple failed-over active/active links in a replication topology can be failed back in any order.
- With failback of multiple failed-over active/passive links in a replication topology, order matters.
- If a topology includes both failed-over active/active links and failed-over active/passive links, order matters for failing back the active/passive links, but the active/active links can be failed back in any order and at any time.
- With active/passive links, failover occurs from the primary system to the replica for the same link. Failover cannot occur from the primary system for one link to the replica for a different link.
- With active/passive links, failback occurs from the replica to the primary system for the same link. Failback cannot occur from the replica for one link to the primary system for another link.

- With an active/passive link, when a link fails over to the replica, only the HCP tenants and namespaces and default-namespace directories that were read-write on the primary system become read-write on the replica.

- In a complex replication topology that includes only active/passive links in many-to-one and/or chained relationships, the HCP tenants and namespaces and default-namespace directories being replicated in the topology are read-write on at most one HCP system at a time. This is true during regardless of the type of activity on the links.

Additional considerations apply to failover and failback in replication topologies that include multiple links.

Failover and failback in an active/passive many-to-one topology

In an active/passive many-to-one replication topology, multiple HCP systems replicate to a single other HCP system. For an explanation of this topology, see "Active/passive many-to-one replication" on page 20.

To recover from a single primary system failure in an active/passive many-to-one replication topology, you follow the normal pattern:

1. Fail over the link between the failed primary system and the replica.

2. When the primary system becomes available again, restore the link from the replica.

3. Begin and complete data recovery from the replica to the primary system.

If more than one primary system fails in an active/passive many-to-one topology, you need to fail the link from each failed system over to the replica. Multiple inbound links on the replica can be in the failed-over state at the same time.

When the failed systems become available again, you can restore the links at any time. However, you can perform data recovery on only one link at a time. (You don't need to wait for all the failed systems to become available before recovering data to the first one.)

For example, suppose systems A, B, and C all replicate to system D. If both A and B fail, you can return to normal replication with these steps after the failed systems become available again:

1. On system D, restore the link from A to D.

2. On system D, begin and complete data recovery from D to A.

3. On system D, restore the link from B to D.

4. After the recovery of data from D to A is complete and replication from A to D has restarted, on system D, begin and complete data recovery from D to B.

Failover and failback in an active/passive chained topology

In an active/passive chained replication topology, one HCP system replicates to a second HCP system, which replicates to a third HCP system. For an explanation of this topology, see "Active/passive chained replication" on page 23.

The way you manage failover and failback in an active/passive chained replication topology depends on which system or systems have failed. The links from the first system in the chain to the second system and from the second system to the third system function independently of each other, but order matters when you fail them over or restore and perform data recovery on them.

The following sections outline the steps you need to take to return an active/passive replication chain to normal operation after the failure of any one or two of the systems in the chain. These sections assume a replication topology in which, when all three systems are healthy:

- System A replicates to system B on link AB.
- Link AB includes HCP tenant T1 and default-namespace directory D1, both of which were originally created on system A.
- System B replicates to system C on link BC.
- Link BC includes link AB and HCP tenant T2, which was originally created on system B.
- T1 and D1 are read-write on system A and read-only on systems B and C.
- T2 is read-write on system B and read-only on system C.

The figure below shows this topology.

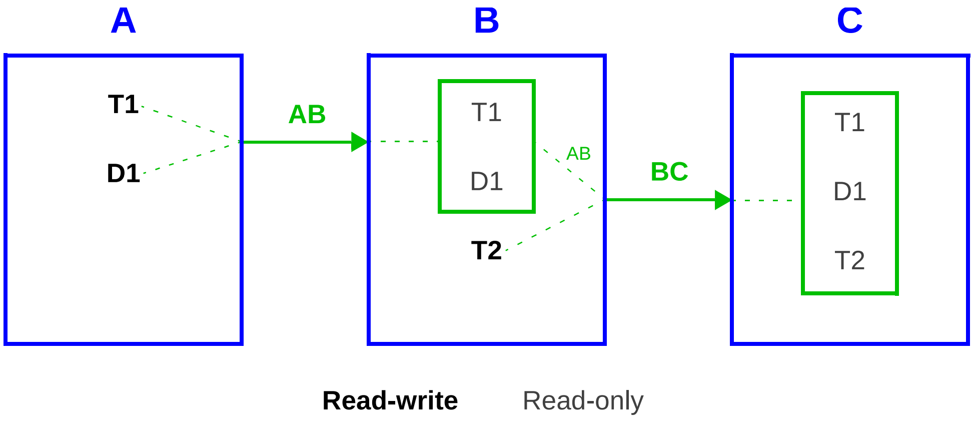

Scenario: System A fails

To return to normal operation after system A fails:

1. On system B, fail link AB over to B.

   T1 and D1 become read-write on B.

2. When system A becomes available again, on system B, restore link AB.

3. On system A, accept the restored link.

4. On system B, begin and complete data recovery on link AB.

   When data recovery is complete, T1 and D1 become read-write on A and read-only on B, and replication resumes on link AB.

Scenario: System B fails

To return to normal operation after system B fails:

1. On system C, fail link BC over to C.

   T1 and D1 remain read-write on A and read-only on C. T2 becomes read-write on C.

2. When system B becomes available again, on system A, restore link AB.

3. On system B, accept the restored link.

   T1 and D1 are read-write on A and read-only on B and C, and replication restarts on link AB.

4. On system C, restore link BC.

5. On system B, accept the restored link.

6. On system C, begin and complete data recovery on link BC.

   When data recovery is complete, T1 and D1 are read-only on B and C. T2 becomes read-write on B and read-only on C, and replication resumes on link BC.

Scenario: System C fails

To return to normal operation after system C fails:

1. When system C becomes available again, on system B, restore link BC.

2. On system C, accept the restored link.

   T1 and D1 remain read-only on B and C. T2 remains read-write on B and read-only on C, and replication restarts on link BC.

Scenario: Systems A and B fail

To return to normal operation after systems A and B fail:

1. On system C, fail link BC over to C.

   T1 and D1 remain read-only on C. T2 becomes read-write on C.

2. When system B becomes available again, on system C, restore link BC.

3. On system B, accept the restored link.

4. On system C, begin and complete data recovery on link BC.

   When data recovery is complete, T1 and D1 are read-only on B and C. T2 becomes read-write on B and read-only on C, and replication resumes on link BC.

   System B automatically recreates link AB without the primary system IP addresses or hostname from the inbound link AB in link BC.

5. When system A becomes available again, take either of these actions:

   - **If link AB still exists on A:**

     1. On system A, restore link AB.

2. On system B, accept the restored link.

3. On system B, begin and complete data recovery on link AB.

   When data recovery is complete, T1 and D1 become read-write on A and remain read-only on B, and replication resumes on link AB.

- **If link AB no longer exists on A,** follow the steps for option one or option two in the table below.

Option one	Option two
1. On system B, update the configuration of the automatically recreated link AB to include the IP addresses or hostname for system A.	1. On system B, suspend and then delete link AB. T1 and D1 become directly included in link BC, which still includes T2, and become read-write on B.
2. On system B, restore link AB.	2. Optionally, create a new replication chain B ▸ C ▸ A: a. Reinstall HCP on A. b. On system C, create outbound link CA, including link BC as an inbound link. c. On system A, accept the new link. T1, D1, and T2 are read-write on B and read-only on C and A, and replication starts on link CA.
3. On system A, accept the restored link.	N/A
4. On system B, begin and complete data recovery on link AB. When data recovery is complete, T1 and D1 are read-write on A and read-only on B and C, and replication restarts on link AB.	

Scenario: Systems A and C fail

To return to normal operation after systems A and C fail:

1. On system B, fail link AB over to B.

   T1 and D1 become read-write on B. T2 remains read-write on B.

2. When system C becomes available again, on system B, restore link BC.

3. On system C, accept the restored link.

   T1, D1, and T2 are read-write on B and read-only on C, and replication restarts on link BC.

4. When system A becomes available again, on system B, restore link AB.

5. On system A, accept the restored link.

6. On system B, begin and complete data recovery on link AB.

   When data recovery is complete, T1 and D1 become read-write on A and read-only on B, and replication resumes on link AB.

Scenario: Systems B and C fail

To return to normal operation after systems B and C fail:

1. When system B becomes available again, on system A, restore link AB.

2. On system B, accept the restored link.

   T1 and D1 are read-write on A and read-only on B, and replication restarts.

3. When system C becomes available again, if T1 and D1 still exist on system C, delete them. (This requires that all objects in the namespaces owned by T1 and in D1 be deleted first.)

   Note: If T2 still exists on systems B and C, delete it from C. If T2 still exists on system C and not on system B, you can recover it to B by creating a link from C to B. However, you cannot then replicate T2 from B to C unless you first delete it from C.

4. On system B, create outbound link BC, including link AB as an inbound link. If T2 still exists on B, also include T2 in the link.

5. On system C, accept the new link.

   T1 and D1 are read-only on B and C. If T2 is included in link BC, T2 is read-write on B and read-only on C. Replication restarts on link BC.

Failover and failback in an active/passive one-to-many topology

In an active/passive one-to many replication topology, one HCP system replicates to two or more other systems. For an explanation of this topology, see “Active/passive one-to-many replication” on page 25.

If one of the replicas fails, you follow the normal pattern for recovering from a replica failure. If more than one replica fails, you follow the normal recovery pattern for each replication link individually. The order in which you perform the recovery procedures doesn’t matter.

Throughout these failure scenarios, the HCP tenants and namespaces and default-namespace directories included in each link remain read-write on the primary system and read-only on the replicas. Therefore, even if the links include the same items, no conflicts can occur.

However, if the two or more links include the same HCP tenants and namespaces and default-namespace directories and the primary system fails, these items can be read-write on multiple systems at the same time. This can lead to conflicts during data recovery. For information on how HCP handles conflicts that occur during data recovery, see “Replication collisions” on page 41.

For example, assume a replication topology in which, when all three systems are healthy:

- System A replicates to system B on link AB. Link AB includes HCP tenant T1.
- System A replicates to system C on link AC. Link AC also includes HCP tenant T1.

To return to normal operation after system A fails:

1. On system B, fail link AB over to B.

   T1 becomes read-write on B and read-only on A and remains read-only on C.

2. On system C, fail link AC over to C.

T1 becomes read-write on C and remains read-only on A and read-write on B. T1 is now read-write on two systems.

Tip: To prevent recovery conflicts, ensure that clients write to only system B or only system C while both systems are read-write.

3. When system A becomes available again, on system B, restore link AB. (You could restore link AC first. The order in which you restore the links doesn't matter.)

4. On system A, accept the restored link.

5. On system B, begin and complete data recovery on link AB.

   When data recovery is complete, T1 remains read-only on A because link AC is stilled failed over to C. It becomes read-only on B and remains read-write on C. Replication resumes on link AB.

6. When data recovery on link AB is complete, on system C, restore link AC.

7. On system A, accept the restored link.

8. On system C, begin and complete data recovery on link AC.

   T1 becomes read-write on A and read-only on C and remains read-only on B. Replication resumes on link AC.

Failover and failback in an active/passive many-to-one topology with disaster recovery support

When HCP systems fail in an active/passive many-to-one topology with disaster recovery support, you need to combine the failback patterns for the many-to-one and chained topologies. For an explanation of the active/passive many-to-one topology with disaster recovery support, see “Many-to-one replication with disaster recovery support” on page 29.

For example, assume a replication topology in which, when all five systems are healthy:

- System A replicates to system D on link AD. Link AD includes HCP tenant T1, which was originally created on system A.

- System B replicates to system D on link BD. Link BD includes HCP tenant T2, which was originally created on system B.

- System C replicates to system D on link CD. Link CD includes HCP tenant T3, which was originally created on system C.
- System D replicates to system E on link DE. Link DE includes links AD, BD, and CD and tenant T4, which was originally created on system D.
- T1, T2, and T3 are read-write on systems A, B, and C, respectively, and read-only on systems D and E.
- T4 is read-write on system D and read-only on system E.

The figure below shows this topology.

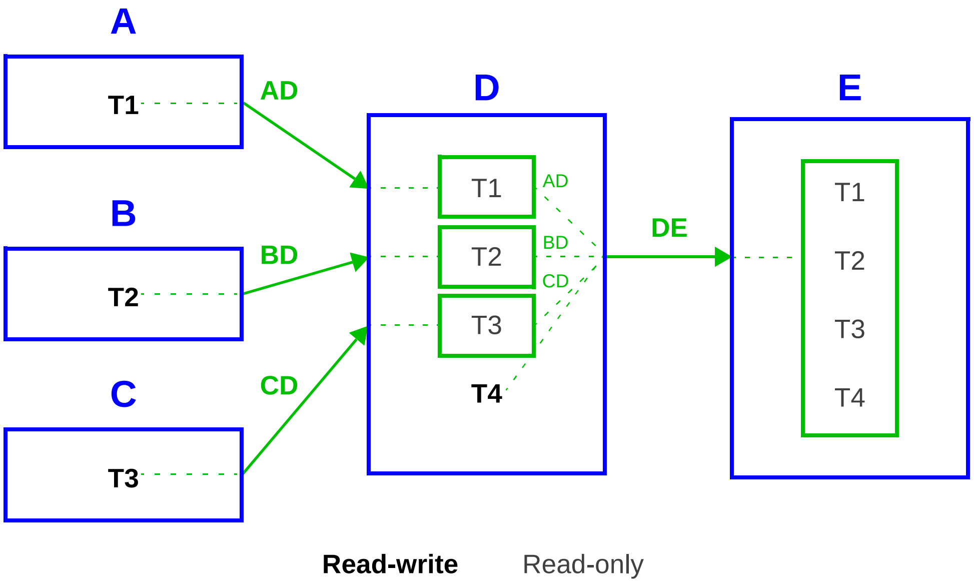

To return to normal operation after systems A, B, and D fail:

1. On system E, fail link DE over to E.

   T1, T2, and T3 are read-only on E. T4 becomes read-write on E. T3 remains read-write on C.

2. When system D becomes available again, on system E, restore link DE.
3. On system D, accept the restored link.
4. On system E, begin and complete data recovery on link DE.

   When data recovery is complete, T1, T2, and T3 are read-only on D and E. T4 becomes read-write on D and read-only on E, and replication resumes on link DE.

5. On system C, restore link CD.

6. On system D, accept the restored link.

   T3 is read-write on C and read-only on D, and replication restarts on link CD.

7. When system A becomes available again, on system D, update the configuration of the automatically recreated link AD.

8. On system D, restore link AD.

9. On system A, accept the restored link.

10. On system D, begin and complete data recovery on link AD.

    When data recovery is complete, T1 becomes read-write on A and remains read-only on D. Replication resumes on link AD.

11. When system B becomes available again and data recovery on link AD is complete, on system D, update the configuration of the automatically recreated link BD.

12. On system D, restore link BD.

13. On system B, accept the restored link.

14. On system D, begin and complete data recovery on link BD.

    When data recovery is complete, T2 becomes read-write on B and remains read-only on D. Replication resumes on link BD.

A

Reenabling user accounts disabled on the replica

For HCP tenants on an active/passive link, tenant-level user accounts are replicated and are available for use on the replica. If an account becomes disabled on the replica due to consecutive failed login attempts, it cannot be reenabled on the replica because the tenant is read-only.

This appendix explains how to reenable a tenant-level user account that has been disabled on a replica.

To reenable a tenant-level user account that's disabled on the replica for an active/passive link, in the Tenant Management Console for the applicable HCP tenant on the primary system:

1. In the top-level menu, mouse over **Security** to display a secondary menu.

2. In the secondary menu, click on **Users**.

3. In the list of user accounts on the **Users** page, click on the username for the account you want to modify.

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

5. Wait for the updated account to be replicated.

   When the account is replicated, it becomes enabled on the replica.

B

Setting advanced replication options

Two advanced options are available for replication:

- Setting the number of threads for the custom performance level
- Setting the interval at which HCP sends SNMP notifications about broken links

This appendix describes these options.

Roles: To perform the activities listed above, you need the service role.

Changing the custom performance level

At any given time, a sending system on a replication link has a performance level that's set in accordance with a schedule or a schedule override. The performance level controls the number of replication threads that can run on each node in the system for that link. For the low, medium, and high performance levels, the numbers of threads are one, five, and ten, respectively. For the custom performance level, the default number of threads is 25, but you can change this number.

The performance level is a maximum. On any given node, fewer threads can be running than the number indicated by the performance level.

If a system is sending data on more than one link, the maximum number of threads per node on that system is the total allowed by all those links.

The number of threads for the custom performance level is a systemwide setting that applies to all links in which the system participates.

To change the number of threads for the custom performance level:

1. In the top-level menu in the HCP System Management Console, mouse over **Configuration** to display a secondary menu.

2. In the secondary menu, click on **Miscellaneous**.

3. In the **Replication Threads for Custom Performance Level** field, type the new number of threads for the custom performance level. Valid values are integers in the range one through 25.

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

Note: If a sending system is already using the custom performance level when you change the number of threads, you need to suspend and resume the link for the change to take effect. For information on suspending and resuming a replication link, see “Suspending and resuming activity on an individual link” on page 122.

For information on replication link schedules and schedule overrides, see “Scheduling activity on a replication link” on page 118.

Changing the broken-link reporting interval

HCP supports the use of SNMP for reporting certain events and conditions. When this capability is enabled, HCP reports broken links on a periodic basis. That is, while a link is broken, HCP sends an SNMP notification about it every *n* minutes. By default, *n* equals 60.

You can change the SNMP reporting interval for broken links. To do this:

1. In the top-level menu in the HCP System Management Console, mouse over **Configuration** to display a secondary menu.

2. In the secondary menu, click on **Miscellaneous**.

3. In the **SNMP Broken-link Reporting Interval** field, type the number of minutes you want HCP to wait between SNMP notifications about broken links. Valid values are integers greater than or equal to one.

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

Glossary

A

access control list (ACL)

Optional metadata consisting of a set of grants of permissions to perform various operations on an object. Permissions can be granted to individual users or to groups of users.

access protocol

See namespace access protocol.

ACL

See access control list (ACL).

Active Directory (AD)

A Microsoft product that, among other features, provides user authentication services.

Active Directory domain

A structural unit within Active Directory that serves as a container for objects such as users and groups.

active/active link

A replication link on which data is replicated in both directions between the two HCP systems. The HCP tenants and namespaces and default-namespace directories included in the link are read-write on both systems.

active/passive link

A replication link on which data is replicated in one direction between the two HCP systems. The HCP tenants and namespaces and default-namespace directories included the link are read-write on only one system at a time.

AD

See Active Directory (AD).

alert

A graphic that indicates the status of some particular element of an HCP system in the System Management Console.

annotation

A discrete unit of custom metadata.

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.

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.

content verification service

The HCP service that ensures the integrity of each object by checking that the object data still matches its cryptographic hash value.

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.

custom metadata

User-supplied information about an HCP object. Custom metadata is specified as one or more annotations, where each annotation is a discrete unit of information about the object. Users and applications can use custom metadata to understand and repurpose object content.

D

data recovery

For an active/passive replication link, the process of copying data from the replica back to the primary system after the link has been failed over.

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.

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.

E

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.

F

failback

The process that restarts replication on a link that has been failed over and returns the HCP systems involved in the link to normal operation. Typically, you fail back a link when an unavailable system becomes available again.

failover

The process that stops replication on a replication link. Typically, you fail over a link when one of the systems involved in the link becomes unavailable.

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

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 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 node

See node.

HCP service

See service.

HCP tenant

A tenant created to manage HCP namespaces.

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.

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.

HTTP

HyperText Transfer Protocol. One of the namespace access protocols supported by HCP.

I

inbound link

An active/passive replication link from the perspective of the replica for the link.

L

link

See replication link.

M

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.

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.

NAT

See network address translation (NAT).

network address translation (NAT)

The translation of a set of IP addresses used within a local area network to a different set of IP addresses used within another network.

node

A server running HCP software and networked with other such servers to form an HCP 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. Objects can also include ACLs that give users and groups permission to perform certain operations on the object.

An object is handled as a single unit by all transactions and internal processes, including shredding, indexing, versioning, and replication.

outbound link

An active/passive replication link from the perspective of the primary system for the link.

P

primary system

For an active/passive replication link, the HCP system from which the replication service objects and other information to the replica during normal replication.

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.

protection service

The HCP service that ensures the stability of the repository by maintaining a set level of data redundancy within each namespace, as specified by the service plan for the namespace.

protocol

See namespace access protocol.

R

RADIUS

Remote Authentication Dial-In User Service. A protocol for authenticating credentials that authorize access to an IP network.

recovery

See data recovery.

replica

For an active/passive link, 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 tenant and namespace configuration, tenant-level user and group accounts, retention classes, content classes, all compliance log messages, and all HCP tenant log messages.

replication link

A configurable, secure trust relationship between two HCP systems that determines what is replicated between the systems and how data is transmitted between the systems.

replication service

The HCP service that performs replication.

repository

The aggregate of the namespaces defined for an HCP system.

retention class

A named retention setting. The value of a retention class can be a duration, Deletion Allowed, Deletion Prohibited, or Initial Unspecified.

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 System Management Console. Roles generally correspond to job functions.

S

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 specification of an HCP service behavior that determines how HCP manages objects in a namespace. Service plans enable you to tailor service activity to specific namespace usage patterns or properties.

SMTP

Simple Mail Transfer Protocol. The namespace access protocol HCP uses to receive and store email data directly from email servers.

SSL

Secure Sockets Layer. A key-based Internet protocol for transmitting documents through an encrypted link.

SSL server certificate

A file containing cryptographic keys and signatures. When used with the HTTP protocol, an SSL server certificate helps verify that the web site holding the certificate is authentic. An SSL server certificate also helps protect data sent to or from that site.

storage node

An HCP node that manages the objects that are added to HCP and can be used for object storage. Each storage node runs the complete HCP software.

System Management Console

The system-specific web application that lets you monitor and manage HCP.

T

tag

An arbitrary text string associated with an HCP tenant or namespace. Tags can be used to group tenants or namespaces and to filter tenant or namespace lists.

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

upstream DNS server

A DNS server to which HCP routes the outbound communications it initiates (for example, for sending log messages to syslog servers or for communicating with Active Directory).

user account

A set of credentials that gives a user access to one or more of the System Management Console, the Tenant Management Console, and namespace content.

Index

Hitachi Data Systems

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

Regional Contact Information

Americas
+1 408 970 1000
info@hds.com

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

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

Hitachi Data Systems

MK-98ARC015-14