Hitachi Content Platform Anywhere Enterprise

v8.3

Portal Installation Guide for ESXi

This document describes how to install and configure an HCP Anywhere Enterprise Portal in a VMware ESXi environment.

MK-HCPAE000-09
June 2026

© 2023, 2026 Hitachi Vantara. All rights reserved.

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

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

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

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

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

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

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

AIX, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, OS/390, PowerHA, PowerPC, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z15, z16, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

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

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

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

Contents

Preface 3

Chapter 1. Planning Your Installation 4

Scalability, Sizing, and Load Balancing 6
Data Replication and Failover 6
Security 7
Requirements 7
Port Requirements 10
Inbound Ports 10
Outbound Ports 10
Additional Ports Not Requiring Internet Access 11

Chapter 2. Installing HCP Anywhere Enterprise Portal Instances 12

Creating a HCP Anywhere Enterprise Portal 12
Create the Archive Storage 14
Troubleshooting the Installation 14
Configuring Network Settings 15
Changing the HCP Anywhere Enterprise Portal Server's Hostname 15
Configuring a Network Interface 15
Configuring Static Routes 17
Configuring a Default Gateway 18
Additional Installation Instructions for Customers Without Internet Access 18

Chapter 3. Configuring the Primary Server 19

Performing Initial HCP Anywhere Enterprise Portal Setup 19
Installing the HCP Anywhere Enterprise Portal License 22
Installing a TLS Certificate 24
Note the HCP Anywhere Enterprise Portal DNS Suffix 24
Obtain a TLS Certificate 25
Generate a Certificate Signing Request 26
Sign the Certificate Request 28
Validate and Prepare Certificates for Upload 29
Install the Signed Certificate on HCP Anywhere Enterprise Portal 29
Importing a TLS Certificate 31
Creating DNS Records 32
Configuring a Public NAT Address 33
Setting Up the Time Zone and Configuring NTP 34
Restarting a Server from the User Interface 36

Chapter 4. Installing Additional HCP Anywhere Enterprise Portal Servers 37

Configuring the Replication Server 39
Configuring Additional Servers 39

Chapter 5. Configuring HCP Anywhere Enterprise Portal Database for Backup and Restore 41

Backing Up the HCP Anywhere Enterprise Servers and Storage 41
Backing Up the HCP Anywhere Enterprise Portal Database 41
Calculating the Minimum Space Required for the Database Backup 41
Configuring PostgreSQL for TLS 42
Using PostgreSQL Continuous Archiving 43
Using PostgreSQL Streaming Replication 45
Monitoring the Database Backup and Streaming Replication 48
Reverting the Primary Database to a Snapshot 55

Chapter 6. Additional Functionality for HCP Anywhere Enterprise Portal Servers 57

Enabling Federal Information Processing Standard (FIPS) 57
Enabling/Disabling Remote Support 57
Load Balancing HCP Anywhere Enterprise Portal Servers 57
General Load Balancing Best Practices 57
Using F5 Load Balancer 57

Chapter 7. ESXi Specific Management 59

Increasing the Data or Archive Pool Size 59
Protecting the HCP Anywhere Enterprise Portal Main Database Using vSphere HA 60
vSphere HA Hardware and Software Requirements 60
Configuring vSphere HA 60
Testing vSphere HA Failover 61

Chapter 8. Upgrading HCP Anywhere Enterprise Portal 62

Upgrading a HCP Anywhere Enterprise Portal to 8.3.3300.x 62
Upgrading the HCP Anywhere Enterprise Portal Image and Software 63
Upgrading a HCP Anywhere Enterprise Portal with the HCP Anywhere Enterprise Messaging Service Implemented 64
Upgrading a HCP Anywhere Enterprise Portal with Subfolder Quotas Configured 64
Upgrading the HCP Anywhere Enterprise Portal Software 64

Preface

About this document

This book describes how to install and initial configuration of an Hitachi Content Platform Anywhere Enterprise Portal in a VMware ESXi environment.

Document conventions

This document uses the following typographic convention:

Convention	Description
Bold	• Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: **Click OK**. • Indicates emphasized words in list items.
Italic	Indicates a document title or emphasized words in text.
`Monospace`	Indicates text that is displayed on screen or entered by the user. Example: `pairdisplay -g oradb`

Intended audience

This document is intended for personnel who will install an HCP Anywhere Enterprise Portal.

Accessing product downloads

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

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

Getting Help

Hitachi Vantara Support Connect is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to Hitachi Vantara Support Connect for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

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

Chapter 1. Planning Your Installation

HCP Anywhere Enterprise Portal is a scalable cloud service delivery platform that you install at your own data center or in a cloud environment and use to create, deliver and manage cloud storage applications, including a Global File System, file access via stubbing/caching, backup, and mobile collaboration. HCP Anywhere Enterprise Portal is compatible with cloud storage infrastructure from multiple vendors and cloud storage providers such as AWS.

HCP Anywhere Enterprise Portal facilitates access to cloud storage services; handles data protection and file sync & share services; used for provisioning and monitoring global file services. This is the beating heart of the system and is the component that will run in the customer's Datacenter or VPC. The portal ensures data consistency, maintains version history, and facilitates file sharing among users, regardless of their access method. Both global source-based deduplication and data compression are used to ensure that only incremental data changes are transferred for storage in the cloud, and that data blocks are stored only once. This dramatically reduces storage capacity needs and overall network traffic. You extend the global file system to users, via HCP Anywhere Enterprise Edge Filers, HCP Anywhere Enterprise Drive Share, and HCP Anywhere Enterprise Drive Connect. The portal ensures data consistency, maintains version history, and facilitates file sharing among users, regardless of their access method.

HCP Anywhere Enterprise Portal enables you to create one or more tenants, or Virtual Portals. These tenants are accessed by end-users and management staff via web-based interfaces. HCP Anywhere Enterprise Edge Filers and endpoint HCP Anywhere Enterprise Agents are centrally managed from HCP Anywhere Enterprise Portal using a single web-based console. Template-based management, centralized monitoring, customized alerting and remote software and firmware upgrade capabilities make it easy to manage gateways of various types and sizes as well as individual endpoints – up to hundreds of thousands of connected devices – with no need for on-site IT presence in remote locations.

Together, these components allow Hitachi Vantara to offer true global file services: files are centrally stored and protected, while users can easily access them everywhere. On top of the all-in-one global namespace/file system approach, Hitachi Vantara allows its customers to achieve the following goals:

Military-grade security – A private and secure architecture powered by end-to-end encryption, advanced authentication, anti-virus, and behind-the-firewall deployment.

Global deduplication – Most modern storage solutions apply deduplication only to centrally stored files. Hitachi Vantara has taken deduplication to the next level, applying the algorithms at both cloud and edge. Not only does the Portal support global deduplication, but HCP Anywhere Enterprise Edge Filers and Drive clients offer source-based deduplication, greatly reducing the size of files being sent to the cloud and lowering storage costs substantially.

WAN optimization – To overcome bandwidth and latency limitations, a slew of optimization techniques are used in order to reduce file sizes and transfer times to/from any access point.

Intelligent caching – Every HCP Anywhere Enterprise Edge Filer and HCP Anywhere Enterprise Drive application comes with a built-in file cache. Caching accelerates remote access, plus it enables access points to "view" the full file storage space, and have on-demand access to every available file.

Managing a large global file system, with thousands of access points and tens of thousands of users, can be quite challenging. To simplify the process and support scale, to tens of thousands of users and sites, the HCP Anywhere Enterprise Portal comes with advanced management tools, including template-based automation. In addition, HCP Anywhere Enterprise Portal comes with rich activity dashboards and analytics, allowing administrators to observe, monitor and troubleshoot every aspect of their global file system.

A HCP Anywhere Enterprise Portal installation comprises a cluster of one or more VMs (servers). Each server can host any combination of the following services:

- **Main database**. Only one server can host the main database. The server that hosts the main database is called the primary server.
- **Secondary replication server**. A passive database service set to replicate the primary server. During server installation, you can turn on the replication service and select the primary server from which to replicate.
- **Application server**. This service accepts connections and handles requests from Web and CTTP clients. Application servers are added to the cluster to increase client handling capacity. Any servers that are enabled as application servers automatically balance the connected clients between them, allowing for maximized capacity and availability.

  Note: CTTP clients are HCP Anywhere Enterprise Edge Filers and Agents that communicate with the HCP Anywhere Enterprise Portal application server using a proprietary secure, WAN optimized transport protocol.

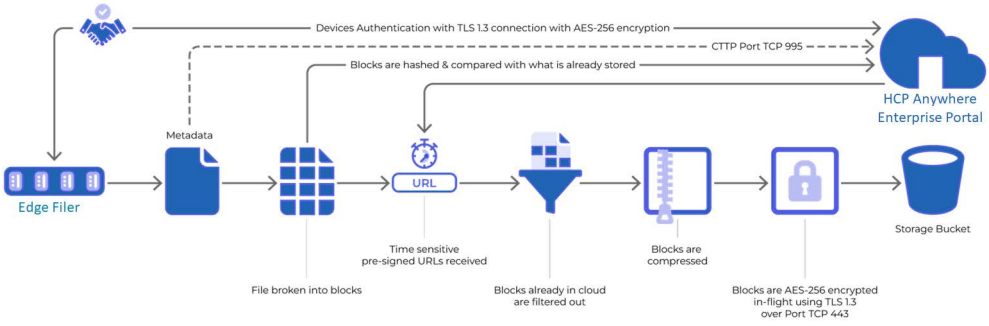

- **Messaging server**. This server enables sending notifications from the portal to various consumers, for example the Varonis Data Security Platform, which is a connector running on top of the HCP Anywhere Enterprise Messaging Service.
- **Document preview server**. This server is used to process document preview requests. The

document preview server supports high availability. You can install one or more servers, in order to ensure uninterrupted document preview generation and redundancy in the event of a server failure.

By default, the first installed server is the primary server, hosting the main database and application server. You can install any number of additional servers. The same procedure is used to install all the server instances.

Scalability, Sizing, and Load Balancing

HCP Anywhere Enterprise Portal is horizontally scalable. Additional servers can be added:

- As application servers, to increase client handling capacity. Any servers that are enabled as application servers automatically balance the connected clients between them, allowing for maximized capacity and availability. The number of application servers deployed depends on the use case:
  - ROBO (remote office, branch office) use case – The users connect to the local HCP Anywhere Enterprise Edge Filers, each edge filer connection to a virtual portal is one connection, even if there are thousands of users connected to each edge filer. You require one application server for every 100 edge filers.
  - FSS (file sync and share) use case – The users connect directly to a virtual portal. You require one application server for every 10,000 HCP Anywhere Enterprise Drive Share users or 1250 HCP Anywhere Enterprise Drive Connect users and a minimum of one virtual portal for every 100,000 users.
- As messaging servers, to enable sending notifications from the portal to various consumers, for example the Varonis Data Security Platform, which is a connector running on top of the HCP Anywhere Enterprise Messaging Service. In production environments that use the messaging service, the HCP Anywhere Enterprise Portal must include three application servers defined as messaging servers.
- As document preview servers. The document preview servers supports high availability. You can install one or more servers, in order to ensure uninterrupted document preview generation and redundancy in the event of a server failure.

Data Replication and Failover

The main database is stateful and contains critical data. You must replicate all such servers to maintain the availability of critical data. The application service is stateless, and therefore, any dedicated application servers do not require replication or backup. Failover between application servers is automatic.

Replicating the database is described as part of the installation for a HCP Anywhere Enterprise Portal.

HCP Anywhere Enterprise Portal includes a built-in replication function for achieving higher level of availability. Replication can be achieved using other platform dependent replication methods (such as SAN or VMWare-level replication).

Security

All internal communication between HCP Anywhere Enterprise Portal servers is authenticated to prevent unauthorized access. Nevertheless, to follow the defense in-depth security philosophy, the primary database server, which stores sensitive data, should be placed in its own firewalled, isolated network, and only the application servers should be allowed to face the Internet.

Requirements

The HCP Anywhere Enterprise Portal image, obtainable from Hitachi Vantara support. The HCP Anywhere Enterprise Portal can be managed in VMware vCenter and in VMware vCloud Director.

HCP Anywhere Enterprise Portal must be installed on a machine that meets the following requirements:

- VMware ESXi 7.0.x and later.
- **Production Deployment Blueprint**
  A minimal production installation of HCP Anywhere Enterprise Portal comprises of four 64-bit virtual machines: Two database servers (primary and secondary) and two application servers. The minimum two application servers are required for high availability and load balancing. Additional application servers may be deployed for further load balancing.
  Note: Either only one or only three application servers function as messaging servers.
  Optionally, one or more preview servers can be deployed for document previews.
  The following table details the requirements per HCP Anywhere Enterprise Portal Server in a production environment.

Server	Minimum Requirements	Notes
Primary Database Server	8 vCPUs, 32GB RAM, 200GB data pool (SSD), 200GB-400GB archive pool (Magnetic)	The data pool should have at least 2000 IOPS and should be sized around 1% of the expected global file system size. Each server is deployed with a 250GB data pool. The archive pool size should be around 2% of the expected global file system size.
Secondary, Replication, Database Server	The replication database server must have the same configuration as the primary database server.	–

Server	Minimum Requirements	Notes
Application Server	8 vCPUs, 32GB RAM, 100GB data pool (Magnetic) **or**, with the HCP Anywhere Enterprise Messaging service: 8 vCPU, 32GB RAM, 250GB data pool (Magnetic)	An application server can handle up to 10,000 clients. When the number of expected clients will be near 10,000, 8 vCPUs and an additional 16GB is recommended. When the HCP Anywhere Enterprise Messaging service is enabled, the data pool size should be a minimum of 250GB but a more accurate amount can be calculated as follows: `data_pool = number_of_daily_audit_events × 4.3` where 4.3 is 3 retention days multiplied by 1.2 for the message size and then a 20%buffer. Hitachi Vantara recommends separate servers to use as messaging servers. Each server is deployed with a 250GB data pool.
Preview Server	8 vCPUs, 32GB RAM, 60GB data pool (SSD)	–

Note: All resources allocated to a server must be dedicated to that server and not shared with other servers. You must not run non-HCP Anywhere Enterprise applications on any of the HCP Anywhere Enterprise Portal servers.
Hitachi Vantara recommends seeking guidance from Hitachi Vantara support for a more accurate estimation of the required sizing.

- **Test Deployment Blueprint**
  The following details the minimal requirements in a test configuration, with a single 64-bit virtual machine deployment.

  Warning: Do not use this setup for production.

  If the HCP Anywhere Enterprise Messaging service will not be part of the test deployment:
  - Single server, 4 vCPU, 8GB RAM, 100GB SSD storage.

  If the HCP Anywhere Enterprise Messaging service will be part of the test deployment:
  - Single server, 4 vCPU, 24GB RAM, 250GB SSD storage.

  Note: In a test deployment with messaging, you specify the primary database server as the messaging server. If you defined more than one server for the test deployment, you cannot set the primary database server as a messaging server.

- **Other requirements**
  - Two networks, one for external client-facing traffic and one for internal communication. You **require** two networks for each portal server.
    External NIC – For client-facing traffic.
    Internal NIC – For inter-server/cluster communication and platform services.
    In order not to compromise security, the internal NIC must **not** be accessible from the external network.
    Both NICs can be on the same VLAN but the internal NIC must have a different IP than the external NIC. Hitachi Vantara strongly recommends a separate VLAN or proper firewall rules.

By default, the NICs are eth0 for the external interface and eth1 for the internal interface. If different names are used, `/etc/ctera/portal.cfg` on each server must be edited to reflect the different names, where:
`CTERA_EXTERNAL_INTERFACE=<external_iface_name>`
`CTERA_INTERNAL_INTERFACE=<internal_iface_name>`
Upgrading a portal automatically assigns the correct name to `CTERA_INTERNAL_INTERFACE`. If the external NIC is not eth0, `/etc/ctera/portal.cfg` must be edited on each server with `CTERA_EXTERNAL_INTERFACE=<external_iface_name>`.
In a multi-datacenter (multi-DC) architecture with separate VLANs, the internal IPs of all portal servers must be routable to each other across the datacenters.

- Web browser: The latest two releases of Apple Safari, Google Chrome, Microsoft Edge, and Mozilla Firefox.
- SSH and SCP clients. For example, the freeware PuTTY.
- Access from the virtual machine to a Storage Area Network (SAN) or directly attached hard drive.
- Specific ports must be open. For details, see Port Requirements.
- A DNS name for the HCP Anywhere Enterprise Portal installation. This can be changed after the installation.
- An ICAP Server and license if the antivirus feature will be used.
  Best Practices:
  At least two ICAP antivirus servers are required for high availability (HA) and load balancing. If no ICAP server is available, scanning will stop, and users will receive error messages when attempting to open files which were not scanned. Antivirus server sizing and the number of servers in the cluster should be determined in consultation with the ICAP antivirus vendor.
  At least 2 background scanning servers are required for HA.
  A single server from the designated background scanners is elected active. If this server fails, another eligible server is elected.
- A mail server address and port for sending notifications.

Port Requirements

To allow access to and from the Internet on the firewall on each machine that will operate as an application server or database server, ensure the relevant network ports are open. In addition, portal servers communicate with each other over specific ports. These ports also need to be opened but do not need Internet access.

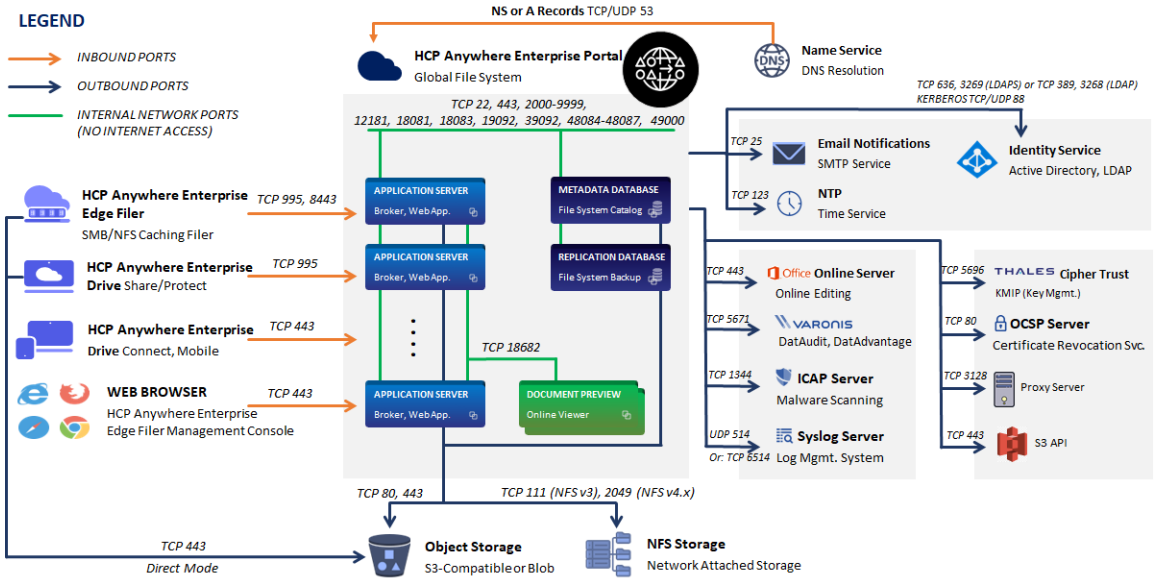

Inbound Ports

Port	Protocol	Notes
22	TCP	SSH. Hitachi Vantara recommends limiting SSH access to specific IP addresses that may require access to the HCP Anywhere Enterprise application servers, for example to perform scheduled maintenance and support related work.
53	UDP	DNS resolution server (If portal internal DNS server is registered in DNS)
80	TCP	HTTP (redirects to port 443)
443	TCP	HTTPS
995	TCP	CTTP protocol communications with HCP Anywhere Enterprise Edge Filers and agents.
8443	TCP	Communications with HCP Anywhere Enterprise Edge Filers for log collection

Outbound Ports

Port	Protocol	Notes
25	TCP	Default SMTP port. This port can be configured on the SMTP server and specified in the portal Web interface
80	TCP	HTTP
88	TCP & UDP	If Kerberos is used

Port	Protocol	Notes
111	TCP	NFS, only required if NFS storage is used
123	UDP	NTP (Network Time Protocol)
389	TCP & UDP	LDAP/LDAP GC (Global Catalog)
443	TCP	HTTPS
514	UDP	Default Syslog port. This port can be configured on the Syslog server
636	TCP	LDAP and LDAP GC with TLS (Hitachi Vantara recommends using LDAPS and LDAPS GC instead of LDAP and LDAP GC)
1344	TCP	If using an antivirus server
2049	TCP	NFS, only required if NFS storage is used
3128	TCP	Default Proxy server port, only required if a proxy server is defined in the global administration, where a different port can be configured
3268	TCP & UDP	LDAP/LDAP GC (Global Catalog)
3269	TCP	LDAPS and LDAPS GC (Hitachi Vantara recommends using LDAPS and LDAPS GC instead of LDAP and LDAP GC)
5671	TCP	Only required when using the Varonis service. This port can be configured in Varonis Data Security Platform
5696	TCP	Only required when using the Key Management service to connect to the Key Management Interoperability Protocol (KMIP) server
6514	TDP	Default Syslog port over TCP/TLS, can be configured on the Syslog server

Additional Ports Not Requiring Internet Access

The following ports must be opened between the HCP Anywhere Enterprise Portal servers.

Port	Protocol	Notes
22, 443	TCP	SSH, internal messaging between HCP Anywhere Enterprise Portal servers.
2181, 2888, 3888, 4646, 4647, 4648, 5044, 5045, 5432, 8081, 8082, 8083, 8088, 8443, 9092, 9093, 9094, 9095, 9096, 9600, 12181, 18081, 18083, 19092, 39092, 48084, 48085, 48086, 40487, 49000	TCP	Internal communication between HCP Anywhere Enterprise Portal server micro services. Port 5432 is used for PostgreSQL access from all cluster members to the main and replication databases. Ports 5045 and 8443 are required when the Thumbnail service is implemented.
4648, 8300, 8301, 8500, 8600	TCP & UDP	Internal communication between HCP Anywhere Enterprise Portal server micro services.
18682	TCP	Only required when a Preview server is used.

Warning: HCP Anywhere Enterprise Portal operates behind a firewall, and it is important to leave all other ports closed.

Chapter 2. Installing HCP Anywhere Enterprise Portal Instances

Creating a HCP Anywhere Enterprise Portal

Contact Hitachi Vantara and request the latest ESXi HCP Anywhere Enterprise Portal OVA file.

Note: The following procedure uses the vSphere Client. You can also use the vSphere Host Client. When using the vSphere Host Client, because the OVA file is larger than 2GB, you must unpack the OVA file, which includes the OVF file, VMDK and MF files. Use the OVF and VMDK files to deploy the HCP Anywhere Enterprise Portal.
For the primary database server and secondary, replication, server, the HCP Anywhere Enterprise Portal instance is created with a fixed size data pool. If you require a larger data pool, which should be approximately 1% of the expected global file system size, you can extend the data pool.

To create the HCP Anywhere Enterprise Portal instance:

Note: The following procedure is based on vSphere Client 7.0.3. The order of actions might be different in different versions.

1. In the vSphere Client console click **File > Deploy OVF Template**.
2. The **Deploy OVF Template** wizard is displayed.
3. Browse to the HCP Anywhere Enterprise Portal OVA file and choose it.
4. Click **NEXT**.
5. Continue through the wizard specifying the following information, as required for your configuration:
   - A name to identify the HCP Anywhere Enterprise Portal in vCenter.
   - The location for the HCP Anywhere Enterprise Portal: either a datacenter or a folder in a datacenter.
   - The compute resource to run the HCP Anywhere Enterprise Portal.
6. Click **NEXT** to review the configuration details.

   Note: Click **Ignore** in the warning to be able to proceed.
7. Click **NEXT**.
8. Select the virtual disk format for the HCP Anywhere Enterprise Portal software and the storage to use for this software. Refer to VMware documentation for a full explanation of the disk provisioning formats. For **Select virtual disk format** select either **Thick Provision Lazy Zeroed** or **Thick Provision Eager Zeroed** according to your preference.
   - **Thick Provision Lazy Zeroed** – Creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the virtual disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out on demand at a later time on first write from the virtual machine. Using the default flat virtual disk format does not zero out or eliminate the possibility of recovering deleted files or restoring old data that might be present on this allocated space.

- **Thick Provision Eager Zeroed** – Creates a virtual disk that supports clustering features such as Fault Tolerance. Space required for the virtual disk is allocated at creation time. In contrast to the flat format, the data remaining on the physical device is zeroed out when the virtual disk is created. It might take much longer to create disks in this format than to create lazy zeroed disks.

9. Click **NEXT**.
10. Select the **Destination Network** that the HCP Anywhere Enterprise Portal will use.
11. Click **NEXT** to customize the template.

    You can customize the following:
    - The name of the portal.
    - The IP address, netmask, gateway and DNS server.

    If you do not enter values, the default values for the portal are applied.
12. Click **NEXT** to review the configuration before creating the VM and then click **FINISH**.

    The HCP Anywhere Enterprise Portal is created and powered off.
13. Power on the HCP Anywhere Enterprise Portal virtual machine.

    The virtual machine starts up and on the first start up a script is run to create a data pool from the data disk and then to load portal dockers on to this data pool. Loading the dockers can take a few minutes.

    Note: You need a data pool on every server. The data pool is created automatically when you first start the virtual machine.
    If the script to create the data pool does not successfully run, it will start on every boot until it completes. The script has a timeout which means it will exit if the data pool is not created within the timeout after boot time. If the data pool is not created, dockers required by the portal are not loaded to the data pool.
    To make sure that the script completed successfully, **before continuing**, run docker images to check that the docker images are available, including zookeeper, which is the last docker to load to the data pool.
    If all the dockers do not load you need to run the script
    `/usr/bin/ctera_firstboot.sh`
    Also, refer to Troubleshooting the Installation if the script does not complete successfully.
14. For the primary database server and secondary, replication database, server, right-click the VM in the navigation page and choose **Edit Settings**.

    The **Hard disk 2** and **Hard disk 3** entries together make up the data pool.

    a) Click **ADD NEW DEVICE > Hard Disk**.
    A new hard disk is added to the list of hard disks.

    b) Enter a size for the new hard disk.

    Note: The minimum archive pool should be 200GB but it should be sized around 2% of the expected global file system size.

    c) Expand the **New Hard disk** item and for **Location** browse to a the datastore you want for the archive pool.

    d) Select the disk to use for the archive pool.

    e) Click **OK**.
15. For every server, right click the VM in the navigation page and choose **Edit Settings**.

    a) Click **ADD NEW DEVICE > Network Adapter**.
    A new network is added to the list of networks.

    b) Click **OK**.
16. Restart the VM.

17. Log in as **root**, using SSH or through the console.

    The default password is `%1Change:Me0`.
    You are prompted to change the password on your first login.
18. For the primary database server and the secondary, replication database, server, continue with Create the Archive Storage.
19. If an IP is not assigned, configure network settings before starting the HCP Anywhere Enterprise Portal services in the next step.
20. Start HCP Anywhere Enterprise Portal services, by running the following command: `portal-manage.sh start`

    Note: Do not start the portal until both the **sdconv** and **envoy** dockers have been loaded to the data pool. You can check that these dockers have loaded in by running `docker images`.

Create the Archive Storage

You need to create an archive pool on the primary database server, and when PostgreSQL streaming replication is required, also on the secondary, replication database, server.

To create the archive pool:

1. Log in as `root`, using SSH or through the console.
2. Run `lsblk` or `fdisk -l` to identify the disk to use for the archive pool.
3. Run the following command to create the archive pool: `portal-storage-util.sh create_db_archive_pool` *`Device`*

   where *Device* is the **device name** of the disk to use for the archive pool.
   For example: `portal-storage-util.sh create_db_archive_pool sdc`
   This command creates both a logical volume and an LVM volume group using the specified device. Therefore, multiple devices can be specified if desired. For example: `portal-storage-util.sh create_db_archive_pool sdc sdd sde`.

   Note: When using NFS storage, before you can create the database archive pool, you have to disable the root squashing security setting for NFS export while setting up the database replication. In the NFS implementation use the `disable root squash` setting. After disabling root squash, run the following command to create the database archive pool: `portal-storage-util.sh create_db_archive_pool -nfs <`*`NFS_IP`*`>:/export/db_archive_dir`
   where *NFS_IP* is the IP address of the NFS mount point.
   Hitachi Vantara recommends re-enabling root squash for the NFS export after the database replication is set up and verified to be working.

Troubleshooting the Installation

You can check on the progress of the docker loads by running `docker images` to check that all the dockers are loaded: The last docker to load is called *zookeeper*:

If all the dockers do not load you need to run the script `/usr/bin/ctera_firstboot.sh`

Configuring Network Settings

To use nmtui:

1. Log in as `root`, using SSH or through the console.
2. Run the following command: `nmtui`

   The **NetworkManager TUI** screen is displayed.
3. Use your keyboard arrow keys to navigate between options.

Changing the HCP Anywhere Enterprise Portal Server's Hostname

To change the HCP Anywhere Enterprise Portal server's hostname:

1. In nmtui, navigate to **Set system hostname** and press **Enter**.
   The **Set Hostname** screen opens, displaying the current HCP Anywhere Enterprise Portal hostname.
2. In the field provided, enter the server hostname.
3. Navigate to **OK** and press **Enter**.

   A confirmation message is displayed.
4. Press **Enter**.

   The new hostname is configured.
5. Navigate to **Quit** and press **Enter** to exit nmtui.
6. You need to reboot the system for the change to take effect. You can reboot the system by entering the command: `reboot`

Configuring a Network Interface

Listing Network Interfaces

To list all network interfaces:

- Run the following command: `ifconfig`

Configuring a Static IP Address for a Network Interface

To configure a static IP address for a network interface:

1. In nmtui, navigate to **Edit a connection** and press **Enter**.
   A window opens, displaying all network adapters attached to the HCP Anywhere Enterprise Portal server.
2. Navigate to the network adapter for which you want to set a static IP address and press **Enter**.

   The **Edit connection** window is displayed.
3. Navigate to **Automatic** next to **IPv4 CONFIGURATION**, press **Enter**, and then select **Manual**.
4. Press **Enter**.
5. Navigate to **Show** next to **IPv4 CONFIGURATION** and press **Enter**.

   Additional fields are displayed.
6. Navigate to **Add** next to **Addresses** and press **Enter**.
7. Type the static IP address.

   To specify a subnet mask, use the classless inter-domain routing (CIDR) notation. For example:
   - To set a class C subnet mask [255.255.255.0], use: *`IP_Address`*`/24`, for example, 192.168.93.204/24

- To set a class B subnet mask [255.255.0.0], use: *IP_Address*/16, for example, 192.168.93.204/16

You can refer to the following link for a full IPv4 CIDR reference:
https://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing#IPv4_CIDR_blocks

8. To configure a default gateway for the current network interface, navigate to **Gateway**, and then type the IP address of the default gateway.
9. To configure a DNS server, navigate to **Add** next to **DNS servers**, press **Enter**, and then enter the IP address of the DNS server.

   Note: You can add multiple DNS servers if desired, by repeating this step.
10. Navigate to **OK** and press **Enter**.
11. Navigate to **Activate a connection** and press **Enter**.
12. Navigate to **Quit** and press **Enter** to exit nmtui.

Your changes take effect.

Note: If using SSH to access nmtui, navigating to **Activate a connection** and pressing **Enter** can close the SSH session. To prevent this from happening, enter the following command:
```
sudo systemctl restart NetworkManager
```

Enabling DHCP for a Network Interface

If you want to use DHCP, for example, for a demo, and you are configured to use a static IP, you can change to DHCP using nmtui.

To enable DHCP for a network interface:

1. In nmtui, navigate to **Edit a connection** and press **Enter**.
   A window opens, displaying all network adapters attached to the HCP Anywhere Enterprise Portal server.
2. Navigate to the network adapter for which you want to enable DHCP, and then press **Enter**.
   The **Edit connection** screen is displayed.
3. Navigate to **Manual** next to **IPv4 CONFIGURATION**, press **Enter**, and then select **Automatic**.
4. Navigate to **OK** and press **Enter**.
5. Navigate to **Activate a connection** and press **Enter**.
6. Navigate to **Quit** and press **Enter** to exit nmtui.

Your changes take effect.

Note: If using SSH to access nmtui, navigating to **Activate a connection** and pressing **Enter** can close the SSH session. To prevent this from happening, enter the following command:
```
sudo systemctl restart NetworkManager
```

Deactivating a Network Interface

To deactivate a network interface:

1. In nmtui, navigate to **Activate a connection** and press **Enter**.
   A window opens, displaying all network adapters attached to the HCP Anywhere Enterprise Portal server.

   Note: An asterisk (*) to the left of a network adapter's name indicates that the network adapter is activated.
2. Navigate to the network adapter you want to deactivate, a network adapter with an asterisk, and

press **Enter**.

3. Navigate to **Back** and press **Enter**.
4. Navigate to **Quit** and press **Enter** to exit nmtui.

The network adapter is deactivated.

Activating a Network Interface

To activate a network interface:

1. In nmtui, navigate to **Activate a connection** and press **Enter**.
   A window opens, displaying all network adapters attached to the HVC Anywhere Enterprise Portal server.
2. Navigate to the deactivated network adapter you want to activate, a network adapter without an asterisk, and press **Enter**.
3. Navigate to **Back** and press **Enter**.
4. Navigate to **Quit** and press **Enter** to exit nmtui.

The network adapter is activated. The asterisk (*) to the left of a network adapter's name indicates that the network adapter is activated.

Configuring Static Routes

To configure a static route for a network interface:

1. In nmtui, navigate to **Edit a connection** and press **Enter**.
   A window opens, displaying all network adapters attached to the HVC Anywhere Enterprise Portal server.
2. Navigate to the network interface for which you want to set a static route and press **Enter**.
   The **Edit connection** screen is displayed.
3. Navigate to **Edit** next to **Routing** and press **Enter**.
4. Navigate to **Add** and press **Enter**.
5. In the fields provided, type the network destination/prefix, the next hop, and the route metric.
   Note: To add another static route, navigate to **Add** and press **Enter**, and then specify the route details.
   To remove an existing route, navigate to **Remove** next to the static route you want to remove and press **Enter**.
6. When done configuring static routes, navigate to **OK** and press **Enter**.
7. Navigate to **OK** and press **Enter**.
8. Navigate to **Back** and press **Enter**.
9. Navigate to **Activate a connection** and press **Enter.**
10. Navigate to **Quit** and press **Enter** to exit nmtui.

Your changes take effect.

Note: If using SSH to access nmtui, navigating to **Activate a connection** and pressing **Enter** can close the SSH session. To prevent this from happening, enter the following command:

```
sudo systemctl restart NetworkManager
```

Configuring a Default Gateway

To set a default gateway for the HCP Anywhere Enterprise Portal server:

1. Log in as root over SSH or through the console to the HCP Anywhere Enterprise Portal.
2. Run the following command:

   `echo "GATEWAY=default_gateway_ip_address" > /etc/sysconfig/network`

   Where *default_gateway_ip_address* is your default gateway IP address.
   For example: `echo "GATEWAY=192.168.90.1" > /etc/sysconfig/network`
3. Restart the network service, by running the following command: `service network restart`

   Your changes take effect.

Additional Installation Instructions for Customers Without Internet Access

The HCP Anywhere Enterprise Portal image requires packages for syslog functionality for offline servers. During a HCP Anywhere Enterprise Portal image installation, these packages are downloaded from the Internet and automatically installed. If you do not have access to the Internet, you can install these packages manually, using the archive file provided with the image installation file.

To install the syslog functionality archive:

1. Install the HCP Anywhere Enterprise Portal image with the latest image file, as described above, prior to running the setup wizard.
2. Get the archive file, **exported_images.tar.gz**, from Hitachi Vantara support and copy it to a local folder on the HCP Anywhere Enterprise Portal machine.
3. Run the following on all servers (including main, Replica, DB, and Preview servers):

   `portal-syslog-client load_images images_archive_path`

   where *images_archive_path* is the path to the local folder where you copied the archive file.
4. Run the setup wizard.

Troubleshooting

If the setup wizard was run prior to loading the new packages, an error message is generated. To resolve it, run the CLI `portal-syslog-client load_images images_archive_path` and restart all servers.

Note: If you do not run the CLI command, the syslog functionality will not work and an error message is generated. For these reasons it is recommended to run the command, however if you don't, apart from the error message, there is no impact on HCP Anywhere Enterprise Portals where syslog functionality is not used.

Chapter 3. Configuring the Primary Server

Configuring the primary server is a one-time operation, the first time you access the HCP Anywhere Enterprise Portal.

Note: All HCP Anywhere Enterprise Portal servers must not run other services. By default, HCP Anywhere Enterprise portal servers are application servers.

If you are installing an additional server, proceed directly with Installing Additional HCP Anywhere Enterprise Portal Servers.

Performing Initial HCP Anywhere Enterprise Portal Setup

This procedure is performed only once, on the primary server.

To perform initial HCP Anywhere Enterprise Portal setup of the primary server:

1. Using a Web browser, browse to the HCP Anywhere Enterprise Portal, via the IP address or DNS.
   The **Setup** wizard opens, displaying the **Select Server operation mode** window.

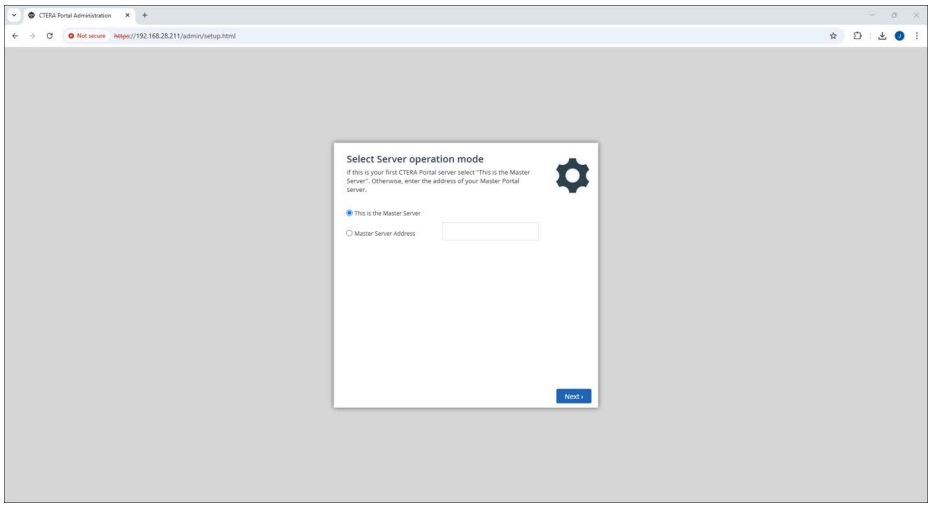

2. Choose **This is the Master Server**.
3. Click **Next**.
   The **License Agreement** window is displayed.
4. Read the agreement and then click **I Accept** and then click **Next**.

The database is initialized and then the **Welcome to HCP Anywhere Enterprise Portal** window is displayed.

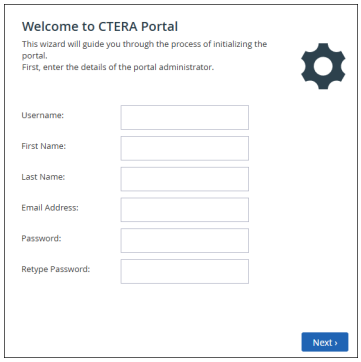

5. Complete the fields as follows:

   Username – The name for your HCP Anywhere Enterprise Portal administrator account.
   First Name – The first name of the administrator.
   Last Name – The last name of the administrator.
   Email Address – The email address of the administrator for notifications.
   Password – The password administrator will use to access the HCP Anywhere Enterprise Portal.
   Retype Password – The password.

6. Click **Next**.

   The **Email Settings** window is displayed.

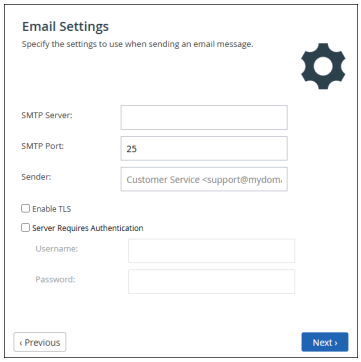

7. Complete the fields as follows:

   SMTP Server – The outgoing mail server address for sending email messages from HCP Anywhere Enterprise Portal to users.
   SMTP Port – The port number for sending email messages from HCP Anywhere Enterprise Portal to users. This port is usually TCP 25.
   Sender – The email address that should appear in the From field of notifications.
   Enable TLS – Select this option to use Transport Layer Security (TLS) encryption for sending email messages from HCP Anywhere Enterprise Portal to users.
   Server requires authentication – Select this option if the SMTP server requires authentication.

   User Name – Type the user name that HCP Anywhere Enterprise Portal should use when authenticating to the SMTP server.
   Password – Type the password that HCP Anywhere Enterprise Portal should use when authenticating to the SMTP server.

Warning: **The username and password you specified for the HCP Anywhere Enterprise Portal administrator are sent to the email address using the information specified here. If the address is incorrect, the email will not arrive and if you have not recorded the administrator username and password details and they are forgotten, you will not be able to access the HCP Anywhere Enterprise Portal.**

8. Click **Next**.

   The **DNS Domain** window is displayed.

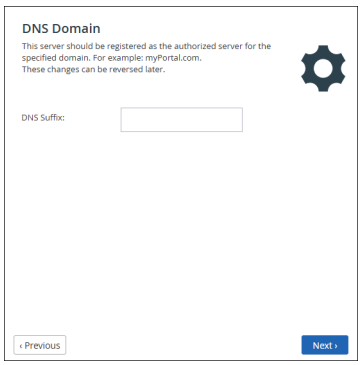

9. In the **DNS Suffix** field, type the DNS suffix to append to each virtual portal's name, in order to create the virtual portal's DNS name.

   For example, if a virtual portal's name is *myportal*, and the DNS suffix is *example.com*, then the virtual portal's DNS name is *myportal.example.com*.

   Note: You can change the DNS domain in the HCP Anywhere Enterprise Portal user interface, in **Settings > Global Settings**.

10. Click **Next**.

    The **Wizard Completed** window is displayed.

11. Click **Finish**.

    The data is saved and a success message is displayed.

12. Click **OK**.

HCP Anywhere Enterprise Portal opens, displaying the **Administrator Login** page.

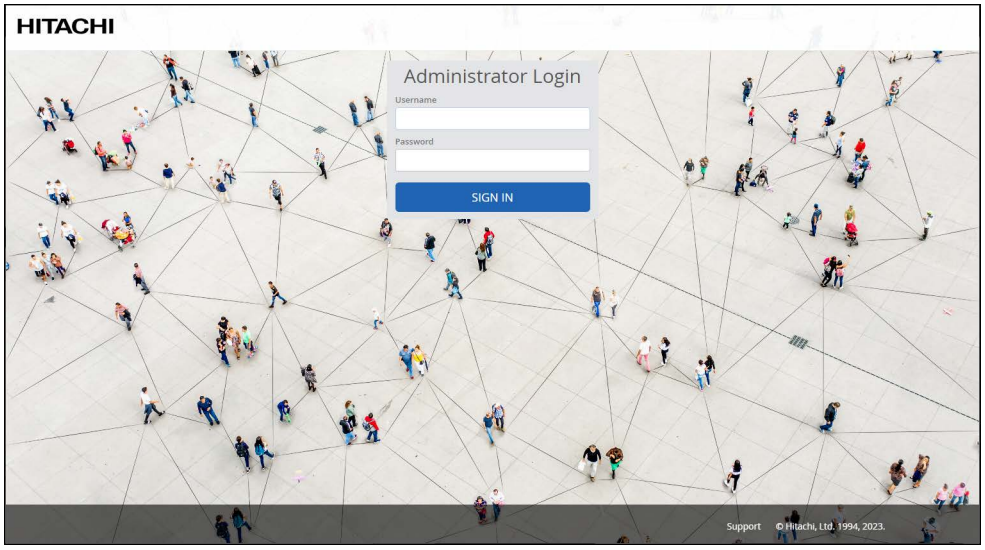

13. Enter the user name and password you specified in step **5** and click **SIGN IN**.

    The portal opens, displaying the **Main > Dashboard** page. By default, HCP Anywhere Enterprise Portal creates a team portal called *portal*. For information about how to rename, view and edit this portal, or create additional team portals, see the *Hitachi Content Platform Anywhere Enterprise Portal Team Administration Guide*.

Warning: **The initial setup includes initializing the PostgreSQL database used by the HCP Anywhere Enterprise Portal. The database must be backed up. Backing up the database is described in Backing Up the HCP Anywhere Enterprise Portal Database.**

Installing the HCP Anywhere Enterprise Portal License

The HCP Anywhere Enterprise Portal includes a trial license for 30 days. This license includes some features that might not be supported in your implementation and excludes other features that you might require. Hitachi Vantara recommends that you immediately install a permanent license that covers your requirements.

Install the permanent license using the following procedure.

To install a license:

1. In the HCP Anywhere Enterprise Portal **Administrator Login** page, sign in as a global administrator.
   The **Main > Dashboard** page for the portal team administrator is displayed.

2. Click the view drop-down and select the **Administration** view.

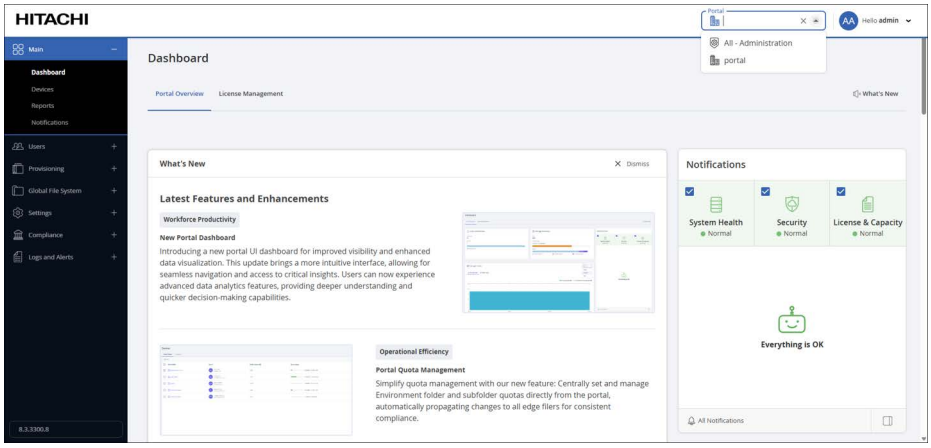

3. In the global administration view, select **Settings > License** in the navigation pane.
   The **License** page is displayed.

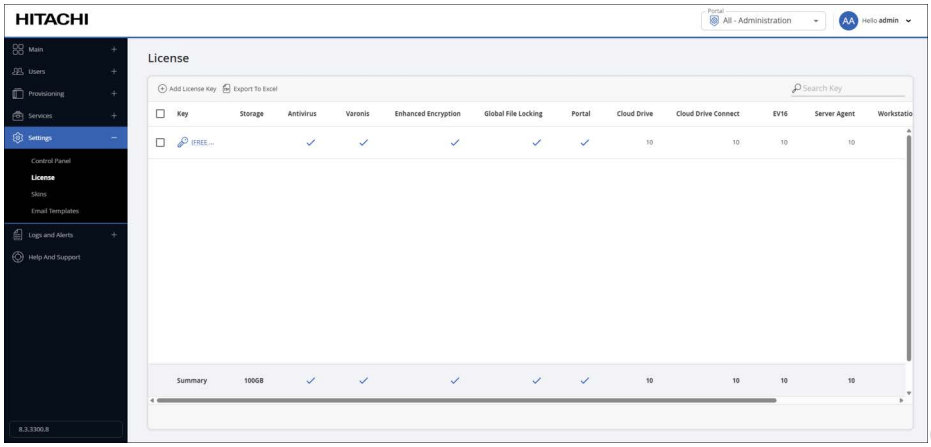

0

4. Click **Add license key**.
   The **Add License Keys** page is displayed.

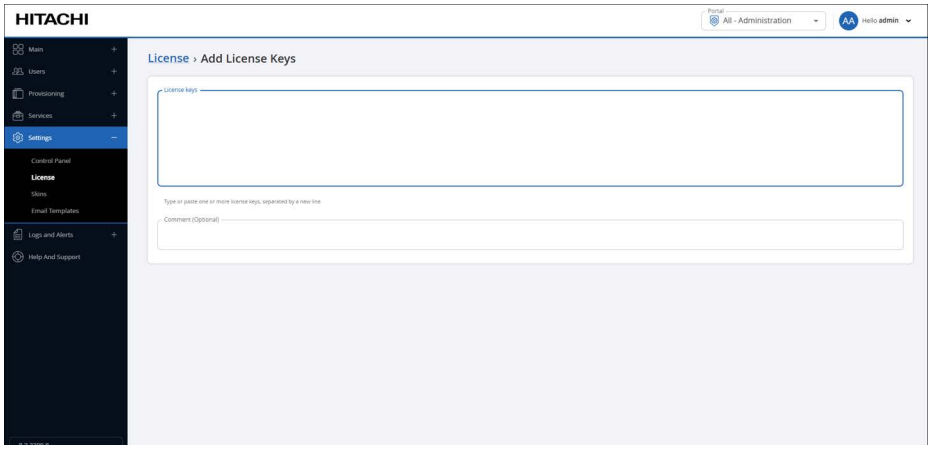

5. Copy the license key you received from Hitachi Vantara, and paste it into the text box.

The system verifies and activates the license key by contacting the Hitachi Vantara Activation service. When the license key is activated, it is associated with this installation of HCP Anywhere Enterprise Portal.

6. Optionally add a comment in the **Comment** field.

   The comment is displayed in the **License** page. You can use this comment to document the purchase order number associated with the license, and the like.
7. Click **Save**.

The system verifies and activates the license key. When the license key is activated, it is associated with this installation of HCP Anywhere Enterprise Portal.

The **License** page displays what is included with HCP Anywhere Enterprise Primary Data Services. Separate licenses are required for some features after the trial license expires, for example for global file locking, Varonis integration, and the Key Management Interoperability Protocol (KMIP).

Installing a TLS Certificate

Perform the following steps to install a certificate on HCP Anywhere Enterprise Portal:

1. Note the HCP Anywhere Enterprise Portal DNS Suffix
2. Obtain a TLS Certificate
3. Generate a Certificate Signing Request
4. Sign the Certificate Request
5. Validate and Prepare Certificates for Upload
6. Install the Signed Certificate on HCP Anywhere Enterprise Portal

Note the HCP Anywhere Enterprise Portal DNS Suffix

You need the HCP Anywhere Enterprise Portal's DNS suffix for use in later steps.

To view the HCP Anywhere Enterprise Portal DNS suffix:

1. In the global administration view, select **Settings > Control Panel** in the navigation pane. The **Control Panel** page is displayed.

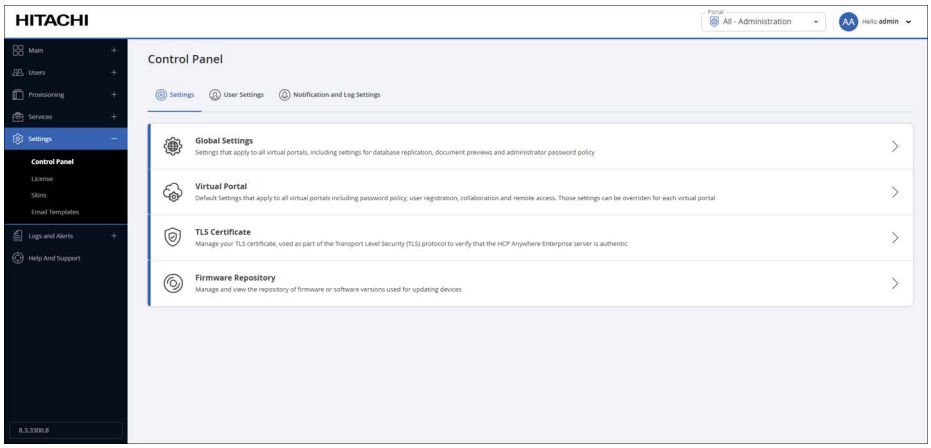

2. Select **Global Settings** under **Settings** in the **Control Panel** page.

The **Global Settings** window is displayed.

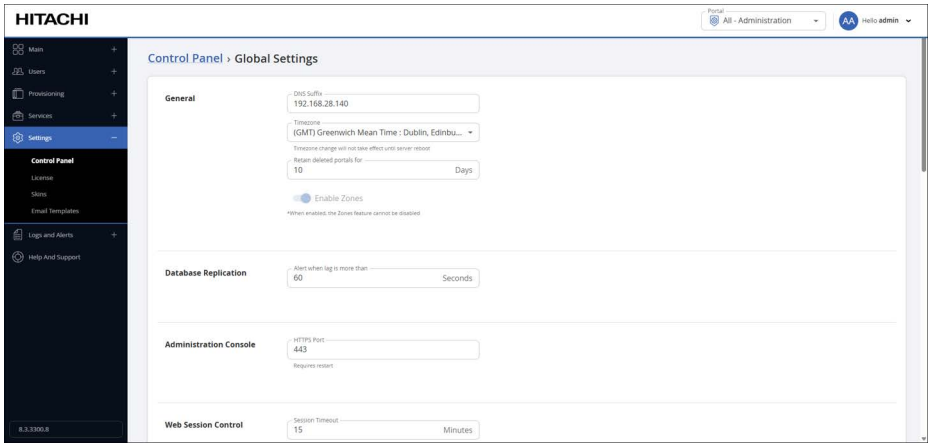

3. Note the HCP Anywhere Enterprise Portal DNS Suffix in the **DNS Suffix** field.

Obtain a TLS Certificate

It is necessary to obtain a valid certificate signed either by a well-known certificate authority, such as GoDaddy, or by your own internal certificate authority.

The TLS certificate can be either of the following:

- **A wildcard certificate**
  A wildcard TLS certificate secures your website URL and an unlimited number of its subdomains. For example, a single wildcard certificate for *.hcp.com can secure both `company01.hcp.com` and `company02.hcp.com`, which may be for virtual portals *company01* and *company02*.
  A wildcard certificate is mandatory if you plan for your service to consist of more than one virtual portal.
- **A domain certificate**
  A domain certificate secures a single domain or subdomain only. For example: `company01.hcp.com`
  This option is relevant if you are planning to provision a single virtual portal only.

Note: To obtain a self-signed certificate for testing and evaluation purposes only, contact Hitachi Vantara Support and specify the HCP Anywhere Enterprise Portal DNS suffix, see **Note the Portal's DNS Suffix**. Hitachi Vantara will generate a self-signed certificate for your DNS suffix and provide you with a ZIP file that you can upload to your HCP Anywhere Enterprise Portal environment.
HCP Anywhere Enterprise Portal also supports certificates with Subject Alternative Names: SAN certificates. This option enables you to secure multiple domain names with a single certificate.

Generate a Certificate Signing Request

You need to generate a certificate signing request, CSR, for your domain. You can generate the CSR from within the HCP Anywhere Enterprise Portal or externally. If you generate the CSR externally, convert the private key file to RSA(PKCS1) format.

Warning: **HCP Anywhere Enterprise Portal generates a built-in certificate that is not suitable for production. This certificate is valid for testing purposes only, as it is not signed by a well-known certificate authority.**

To generate a certificate signing request for your domain:

1. In the global administration view, select **Settings > Control Panel** in the navigation pane. The **Control Panel** page is displayed.

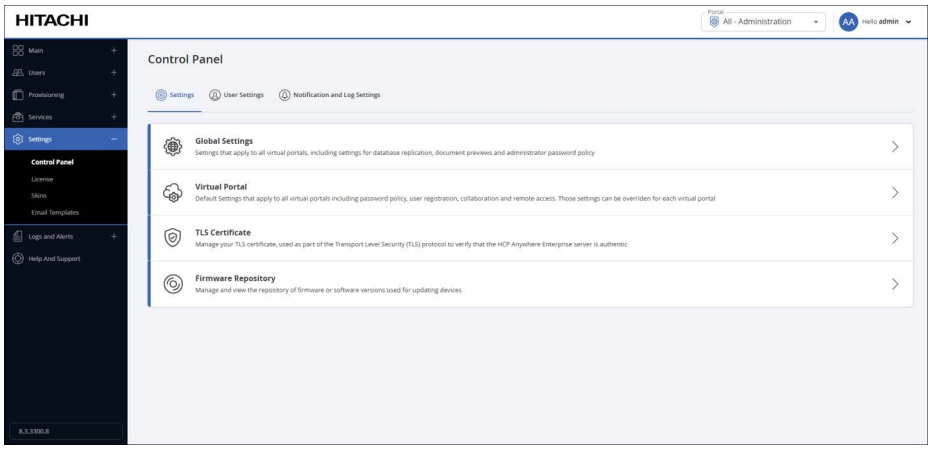

2. Select **TLS Certificate** under **Settings** in the **Control Panel** page.

   The **TLS Certificate** page is displayed.

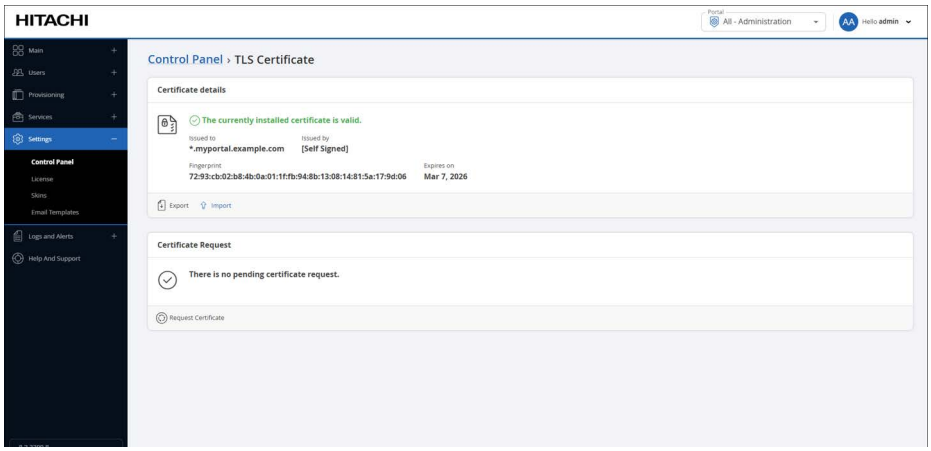

Note: On installing the portal you have a trial license and self-signed certificate.

3. Click **Request Certificate**.

   The **Create a certificate request** page is displayed.

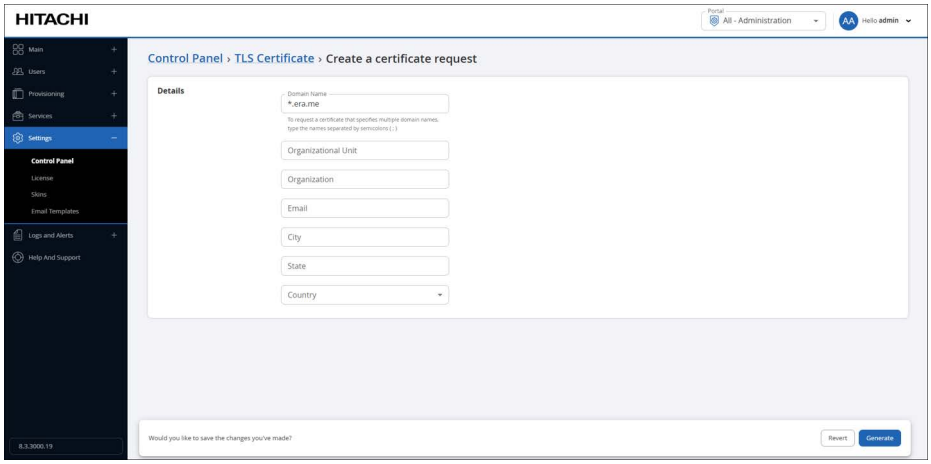

4. In the **Domain Name** field, enter the domain name for which you want to request a certificate.

   The value entered must match the type of certificate you chose to use. For example, if you chose a wildcard certificate, the domain name might be `*.example.com`
   If you chose a domain certificate, the domain name might be `company01.example.com`, where *company01* is the name of your virtual portal.
   If multiple virtual portals are configured, each virtual portal has its own DNS name. In this case, the TLS certificate should be a wildcard certificate with an asterisk before the DNS suffix, for example, `*.example.com`
   If you have only one portal, and do not intended to configure multiple virtual portals, then use a regular TLS certificate and not a wildcard certificate.
   To request a certificate that specifies multiple alternative names, type the multiple names in this field, separated by semicolons (;). The certificate will include the `subjectAltName` certificate extension.
5. Optionally, specify the following:

   Organizational Unit – The name of your organizational unit.
   Organization – The name of your organization.
   Email – Your email address.
   City – Your city.
   State – Your state.
   Country – Your country.
6. Click **Generate**.

   A key pair is generated and stored on the portal.

A message is displayed that you can download the certificate request.

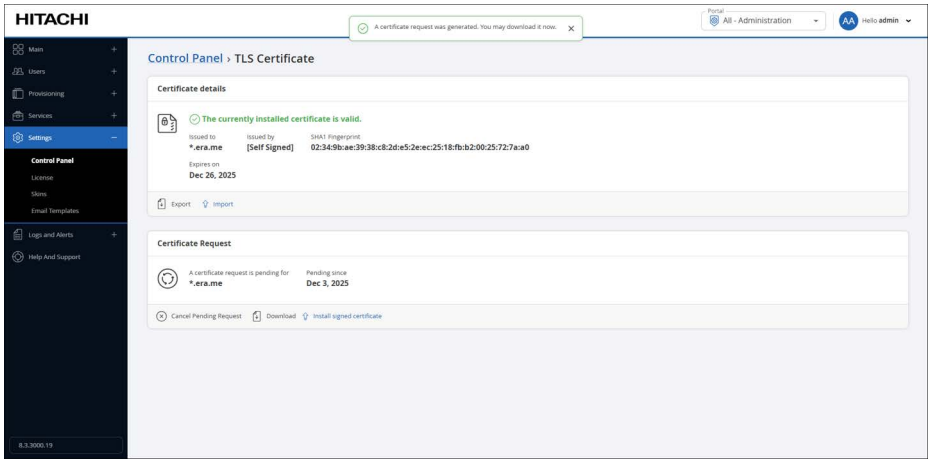

The **Certificate Request** area of the **TLS Certificate** page indicates that the certificate request is pending.

7. Click **Download**.

   The certificate request file *certificate.req* is downloaded to your computer.

   Warning: When you generate the CSR, a private.key file is registered in the HCP Anywhere Enterprise Portal. If you now generate a new CSR, it will override the existing private.key file, and signing the old CSR will result in an error message indicating that the CSR does not match the private.key file. Therefore, do not generate a new CSR before installing the signed certificate.

Sign the Certificate Request

To sign the certificate request:

1. Send the `certificate.req` file you generated to your certificate authority for signing.
   If the request is successful, the certificate authority will send back an identity certificate that is digitally signed with the certificate authority's private key.
   Note: The certificate authority should return a **base-64** encoded identity certificate.
2. Open the identity certificate and verify that the **Issued to** field includes the DNS suffix you provided upon creating the certificate request.
3. Build a certification chain from your identity certificate to your trusted root certificate.

   You need to obtain all of the intermediate certificates, as well as your root certificate authority's self-signed certificate.
   If you are using a well-known certificate authority, the intermediate certificates and the root certificate authority's self-signed certificate can be downloaded from your certificate authority website. If you are using your own internal certificate authority, contact the necessary entity to provide you with the required intermediate and self-signed certificate.
   For example, if the certificate was issued by **Go Daddy Secure Certification Authority** to .era.me, to build the certification chain, you need to obtain a certificate issued to **Go Daddy Secure Certification Authority**.
   To continue the certification chain, you must obtain a certificate issued to the same authority that the previous certificate was issued by. You continue the chain until the certification chain is complete, with the last certificate, which is a self-signed certificate, issued to and by the same entity.

Validate and Prepare Certificates for Upload

To validate and prepare certificates for upload:

1. Verify that none of the certificates in the certificate chain are corrupted or using invalid encoding.
   - Open each certificate in a program such as Notepad or Word, and verify that it contains the following:

   ```
   -------- BEGIN CERTIFICATE -------
   <CERTIFICATE CONTENT>
   -------- END CERTIFICATE -------
   ```
2. Change the identity certificate issued to `*.hcp.me` to `certificate.crt`
3. Change the file extension of the other certificates in the certificate chain to `crt`

   For example, `certificate-name.crt`
4. Archive all of the certificates, the identity certificate, the intermediary certificates, and the root self-signed certificate, in a ZIP file called `certificate.zip`.

Install the Signed Certificate on HCP Anywhere Enterprise Portal

If you have a valid signed certificate, install it and replace the built-in certificate.

1. In the global administration view, select **Settings > Control Panel** in the navigation pane. The **Control Panel** page is displayed.

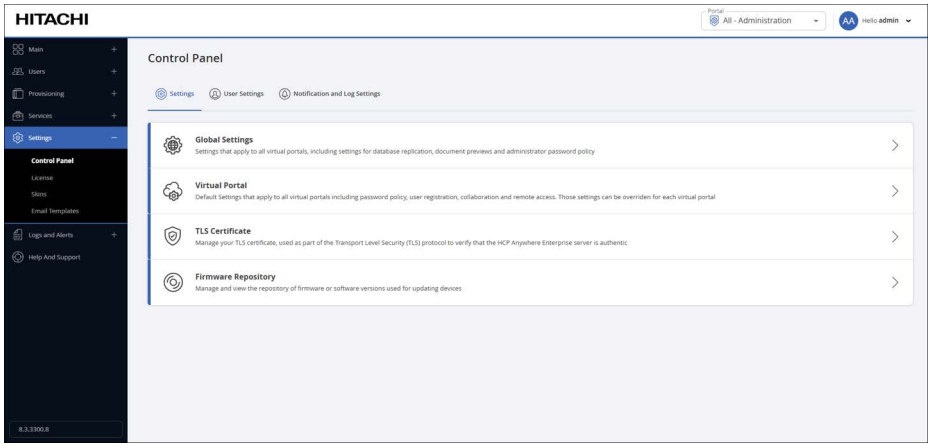

2. Select **TLS Certificate** under **Settings** in the **Control Panel** page.

The **TLS Certificate** page is displayed.

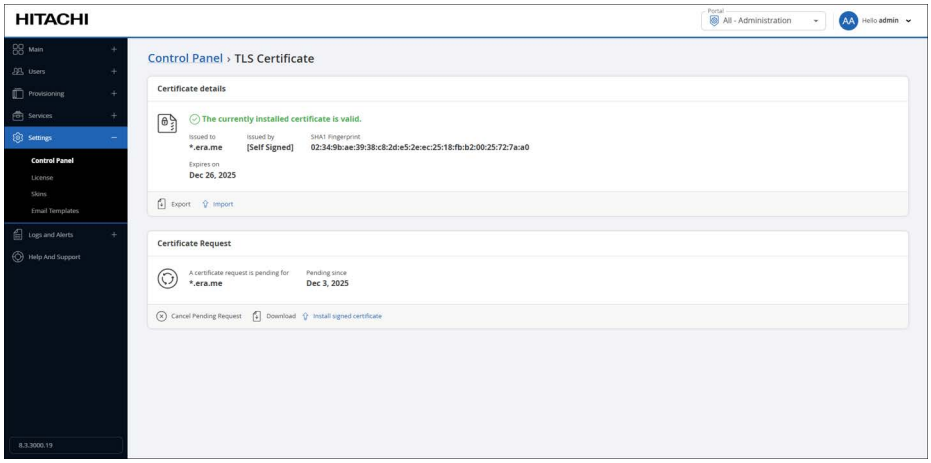

The **Certificate Request** area of the **TLS Certificate** page indicates that the certificate request is pending.

Note: To cancel a pending request, for example, to make changes to the current certificate request, click **Cancel Pending Request** and click **OK** in the confirmation window that is displayed.
The pending certificate request is canceled.

3. Click **Install Signed Certificate** in the **Certificate Request** area.
4. Browse to the certificate, select it and click **Open** and browse to the certificate.zip file you created. All the certificates in the certificate chain must be in the ZIP file in X.509 format, and each file must have a ".crt" extension.

   The certificate is installed on HCP Anywhere Enterprise Portal.

   Note: If you receive the following error: Caused by: java.lang.ClassCastException: class com.safelogic.cryptocomply.asn1.pkcs.PrivateKeyInfo cannot be cast to class org.bouncycastle.openssl.PEMKeyPair (com.safelogic.cryptocomply.asn1.pkcs.PrivateKeyInfo is in module ccj@3.0.0 of loader 'app'; org.bouncycastle.openssl.PEMKeyPair is in unnamed module of loader java.net.URLClassLoader @5ccd43c2) at com.era.security.CA.generateKeyStore(CA.java:427) do the following:
   Using SSH, log in as `root` to the server.
   In the command line, enter the following command `openssl rsa -in private.key -out <private1.key> -traditional`.
   where *private1.key* is the name of the private key.
   Use the private.key in the zip file created in **Validate and Prepare Certificates for Upload**.
5. Click **Finish**.
6. Open HCP Anywhere Enterprise Portal.

If the certificate update was successful, there won't be any security exceptions.

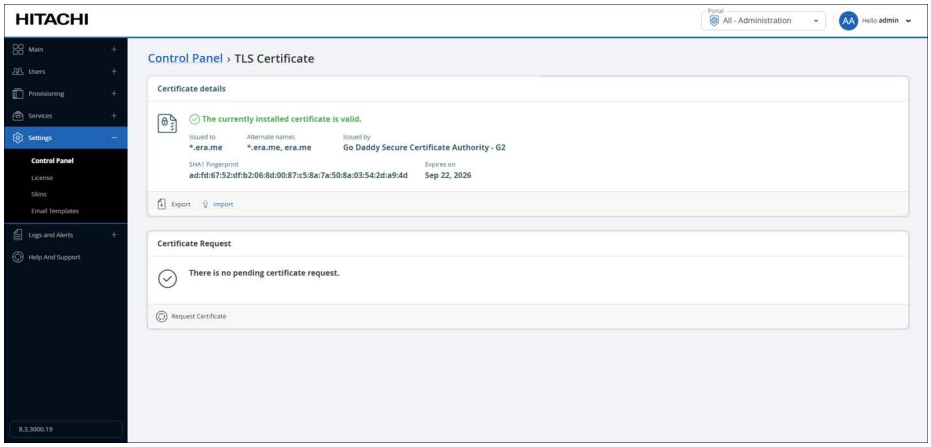

Importing a TLS Certificate

To import a TLS certificate:

1. In the global administration view, select **Settings > Control Panel** in the navigation pane. The **Control Panel** page is displayed.

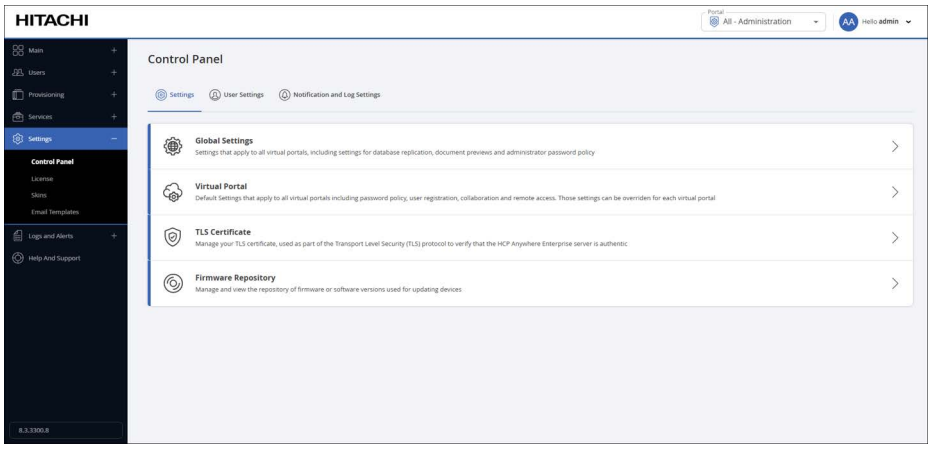

2. Select **TLS Certificate** under **Settings** in the **Control Panel** page.

The **TLS Certificate** page is displayed.

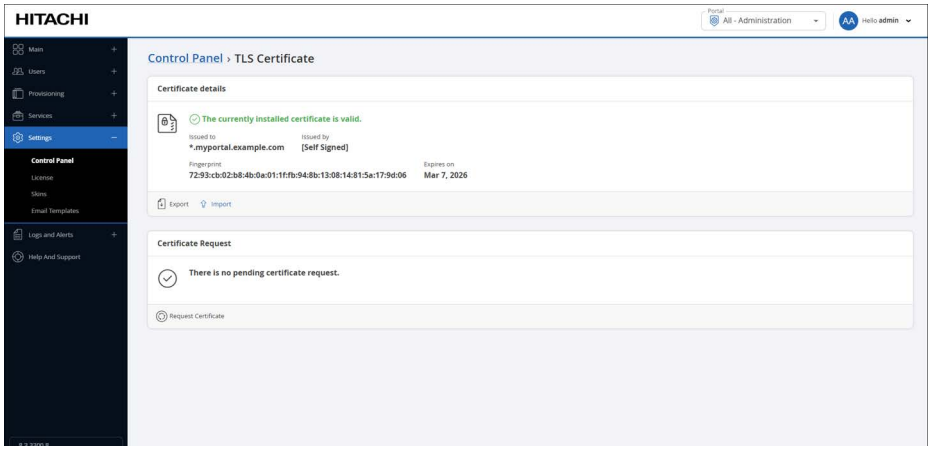

3. Click **Import**.
4. Browse to the ZIP file containing the certificate components, select the file and click **Open**.
5. Log back on to the portal to view the new certificate.

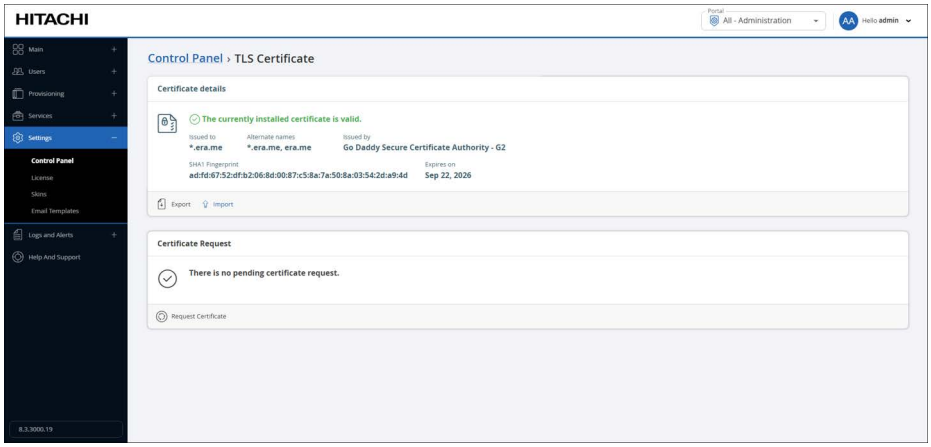

Creating DNS Records

The HCP Anywhere Enterprise Portal includes a built-in DNS server. This server automatically resolves the domain names of all the defined virtual portals, as well as names of devices using the remote access service. In order for this DNS server to work, you must register it using an NS (Name Server) record on your DNS server.

The procedure used for configuring the DNS for remote access depends on the whether you have purchased a dedicated domain (the DNS suffix includes only records for the HCP Anywhere Enterprise Portal) or not (the DNS suffix includes records that are unrelated to the HCP Anywhere Enterprise Portal).

If you have a dedicated domain:

- If you have a dedicated domain for the HCP Anywhere Enterprise Portal – no servers other than the HCP Anywhere Enterprise Portal – then the NS record can be created just once, in that zone.
  For example, for a DNS suffix called *storage.example.com* and two HCP Anywhere Enterprise Portal servers with IPs 123.168.0.3 (primary) and 123.168.0.4 (secondary), you would register:
  ```
  A          srv1.example.com          192.168.10.3
  A          srv2.example.com          192.168.10.4
  ```
  Next, you would create an NS record for each server to the zone *storage.example.com*:
  ```
  NS       storage.example.com       srv1.example.com
  NS       storage.example.com       srv2.example.com
  ```

Configuring a Public NAT Address

Immediately after deploying a HCP Anywhere Enterprise Portal instance, the HCP Anywhere Enterprise Portal server will respond to DNS requests with its private internal IP address. In order to make the HCP Anywhere Enterprise Portal available via the Internet, and to enable the HCP Anywhere Enterprise Portal to respond to DNS queries with the public IP address, you must configure the HCP Anywhere Enterprise Portal's public NAT address.

To configure a Public NAT address:

1. In the global administration view, select **Main > Servers** in the navigation pane.
   The **Servers** page is displayed.

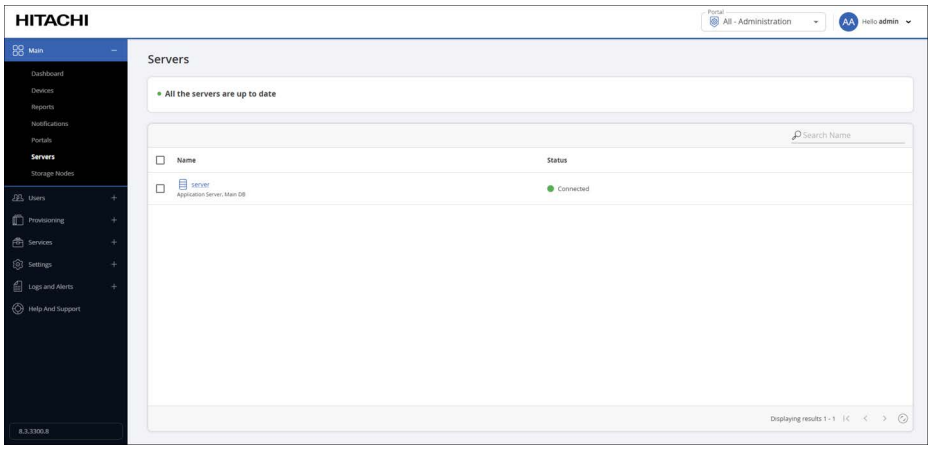

2. Click the server to edit.

   The server window is displayed with the server name as the window title.

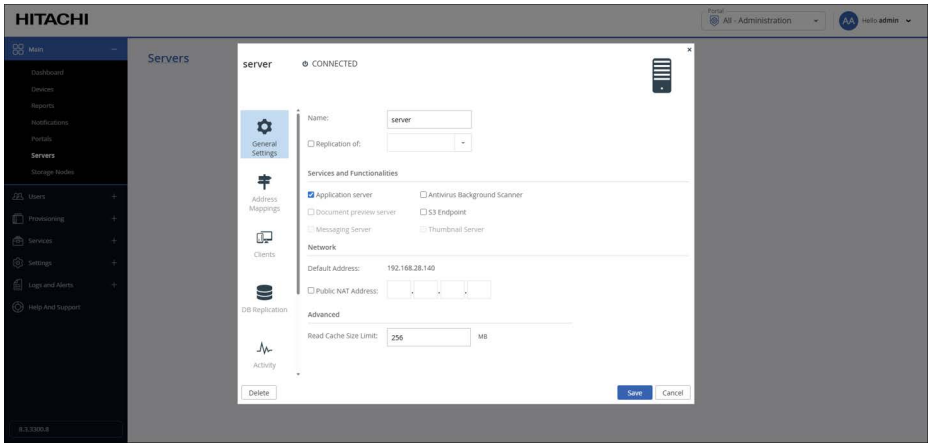

3. Check **Public NAT Address** and enter public IP address.
4. Click **Save**.

Setting Up the Time Zone and Configuring NTP

Note: HCP Anywhere Enterprise Portal is preconfigured with NTP servers. If you want to use different NTP servers, Hitachi Vantara provides the following procedures.

You can configure the time zone for the HCP Anywhere Enterprise Portal. If for whatever reason, the HCP Anywhere Enterprise Portal clock and HCP Anywhere Enterprise Edge Filer clocks are not synced, you can configure the HCP Anywhere Enterprise Portal NTP server, as described below.

To configure the HCP Anywhere Enterprise Portal server's time zone:

1. In the global administration view, select **Settings > Control Panel** in the navigation pane. The **Control Panel** page is displayed.

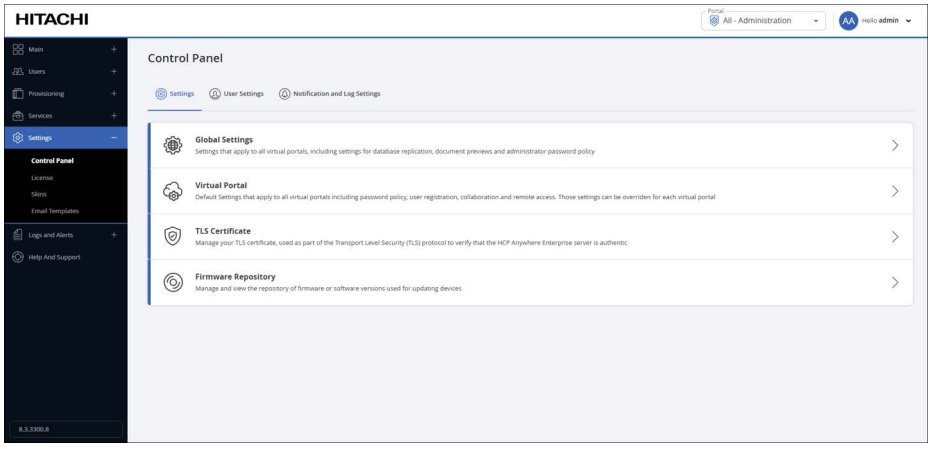

2. Select **Global Settings** under **Settings**.
   The **Global Settings** window is displayed.

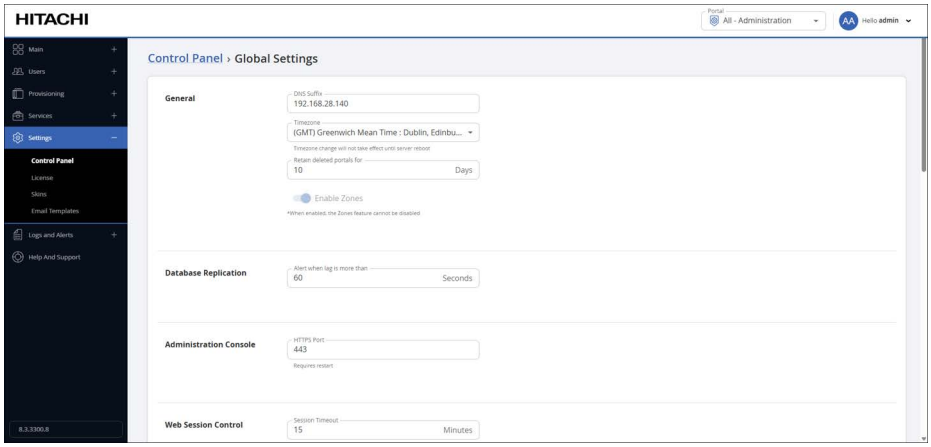

3. Select the correct time zone for the HCP Anywhere Enterprise Portal from the list.
4. Click **Save**.
5. Restart the server as described in Restarting a Server from the User Interface.

To configure NTP in the HCP Anywhere Enterprise Portal server:

1. Log in as `root`, using SSH to the portal server.
2. Edit the *chrony* configuration file `/etc/chrony.conf` and add the following line at the end of the file: `server <NTP_SERVER> iburst prefer` where <*NTP_SERVER*> is the NTP server name or IP address.
3. Run the command to restart the chrony services: `systemctl restart chronyd`.

Restarting a Server from the User Interface

HCP Anywhere Enterprise Portal servers can be restarted from the user interface.

To restart a server:

1. In the global administration view, select **Main > Servers** in the navigation pane. The **Servers** page is displayed.

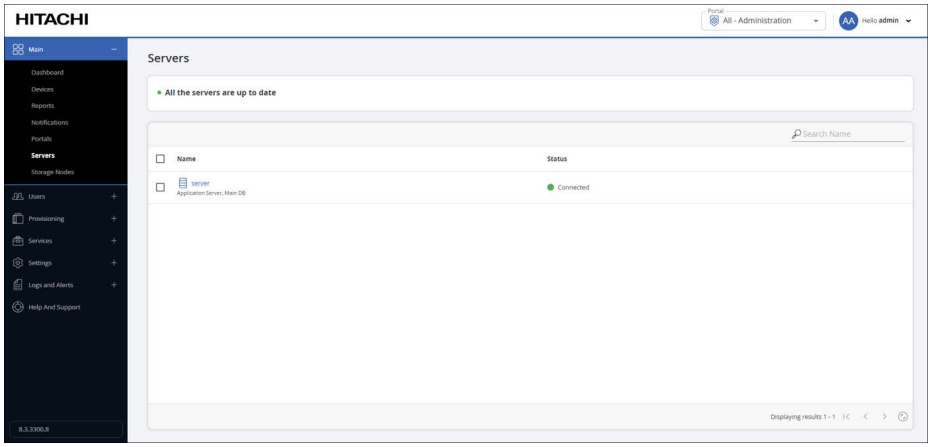

2. Select the server to restart and click **Restart**.

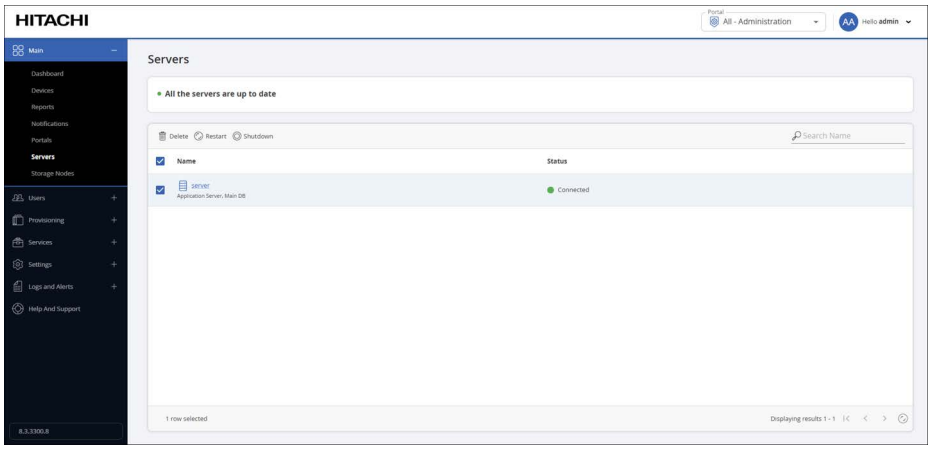

   A confirmation window is displayed.
3. Click **OK** to confirm.

The server is restarted.

Chapter 4. Installing Additional HCP Anywhere Enterprise Portal Servers

All servers except the primary server are additional servers.

Warning: You must complete the setup of each server before setting up another server.

You install additional servers using the same procedure you used to install the primary server, except for the following:

- For messaging servers, you require 4 vCPUs, 32GB RAM, and 250GB data pool (Magnetic).
- For a preview server, you require 4 vCPUs, 16GB RAM, and 60GB data pool (SSD).
- For all other application servers, you require 4 vCPUs, 16GB RAM, and 100GB data pool (Magnetic).

Note: For document previews to work, you must install at least one additional server as a document preview server.

Installing an additional server involves installing a server as described in *Installing HCP Anywhere Enterprise Portal Instances* and then configuring it as an additional server.

In a production environment, the primary server hosts the Hitachi Content Platform Anywhere Enterprise Portal database, referred to as Main DB. In production this server and the server that hosts the replication database, referred to as Replication of main_db_server_name should not be defined as application servers.

To set up an additional server:

Note: If the HTTPS administration port was changed on the primary server, you must enable the relevant port on all servers.

1. Using a Web browser, browse to the new server, via the IP address or DNS.
   The **Setup** wizard opens, displaying the **Select Server operation mode** window.

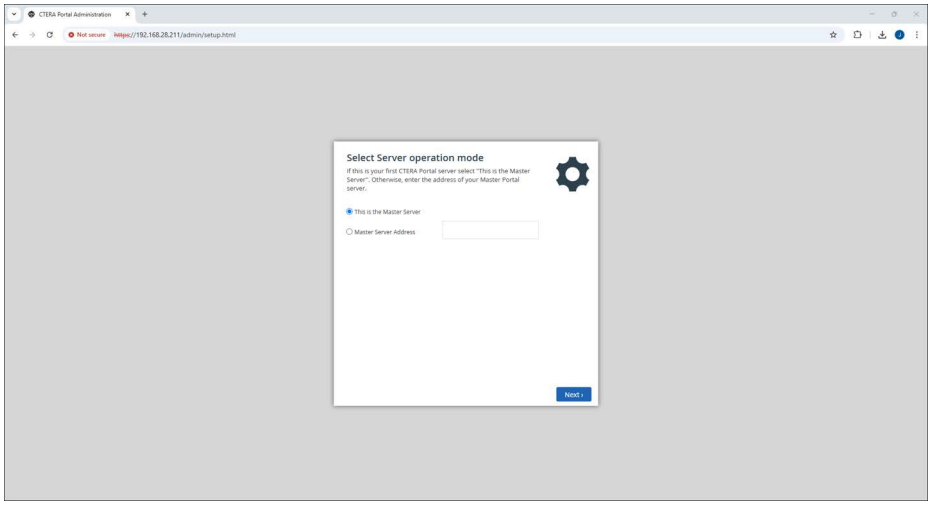

2. Choose **Master Server Address** and enter the address of the primary server.
3. Click **Next**.

The **Master Server Details** window is displayed.

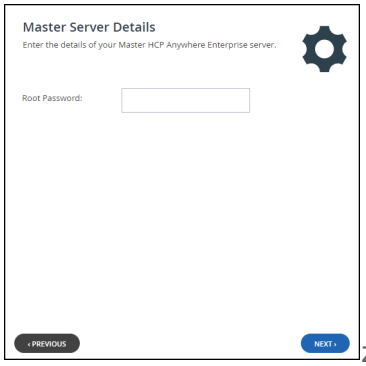

z\

4. Enter the root password for the primary server. The default is `%1Change:Me0`, but this should have been changed on first access to the server.
5. Click **Next**.

   The wizard completes and a 60-second countdown is displayed before the server is up and running.

Note: Completing the procedure takes a few more minutes before the **Administrator Login** page is displayed.

6. Click **OK**.

   HCP Anywhere Enterprise Portal opens, displaying the **Administrator Login** page.

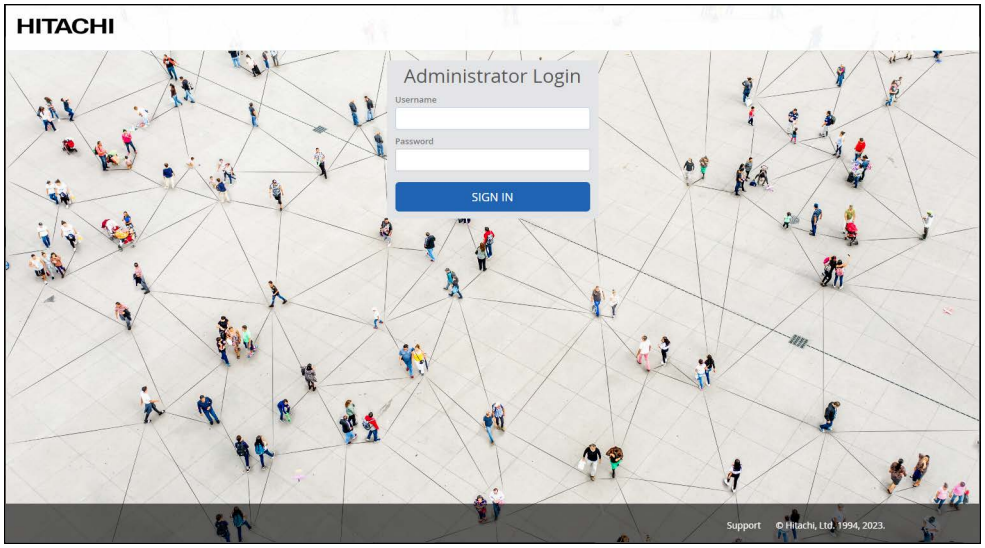

Note: The first access to the **Administrator Login** page can take a few minutes as dockers are loaded to the data pool.
The IP address is the IP of the additional server. The IP address of any server configured as an application server can be used to access the HCP Anywhere Enterprise Portal.

Configuring the Replication Server

The additional server defined with an archive pool is configured as a r**eplication** server of the primary server.

Warning: Configure only one server as a replication server.

To configure the replication server:

1. Sign in to the portal: Enter the user name and password and click **SIGN IN.**
   The HCP Anywhere Enterprise portal opens, displaying the **Main > Dashboard** page.
2. In the global administration view, select **Main > Servers** in the navigation pane.
   The **Servers** page is displayed.
3. Click the server to make the replication server: Designate the server on which you created an archive pool that is not the primary server, as the replication server.
   The server window is displayed with the server name as the window title.
4. Check **Replication of server** and select the primary server from the drop-down box.
5. Uncheck **Application server** for the replication database server, called *Replication of server* in the **Servers** page.
6. Click **Save**.

Configuring Additional Servers

Additional servers are configured to support different services, such as messaging, thumbnails, or document preview servers. In a production environment you have as a minimum a primary database server, a secondary, replication, server and at least two application servers for HA, or three application servers when the HCP Anywhere Enterprise Messaging service is enabled.

To Configure the Additional Server:

1. In the global administration view, select **Main > Servers** in the navigation pane.
   The **Servers** page is displayed.

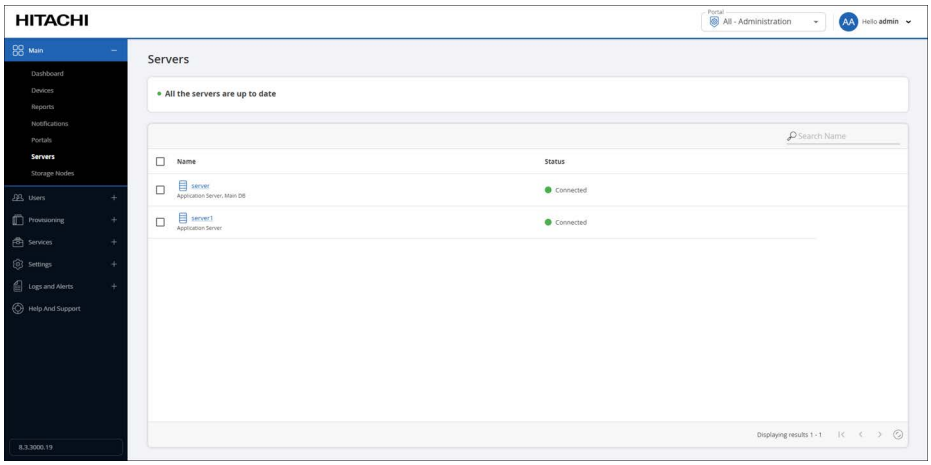

2. If you have already defined another server, that is still configured as an **Application server**, uncheck **Application server** for the primary database server, called *Main DB* in the **Servers** page.

Warning: **You can only sign in to a server that is configured as an Application server, using the default portal DNS or that server IP. Once you have unchecked Application server for the primary database server and secondary, replication, server, you can no longer sign in to the Hitachi Content Platform Anywhere Enterprise Portal using either of these server IPs.**

3. Check the boxes of the services to be provided by this server, for example, as a document preview server.
4. Click **Save**.

Chapter 5. Configuring HCP Anywhere Enterprise Portal Database for Backup and Restore

Backing Up the HCP Anywhere Enterprise Servers and Storage

To back up the HCP Anywhere Enterprise Portal servers and storage you need to use third-party recovery tools.

Note: If the storage fails and you need to use the backup storage, you need to change the definition of the storage node to point to the backup storage bucket in the HCP Anywhere Enterprise Portal.

For HCP Anywhere Enterprise Portal database backup, see Backing Up the HCP Anywhere Enterprise Portal Database.

Using third-party tools can result in a recovery after a disaster not being consistent with the database with the recovery being either older or newer than the database recovery.

If the storage recovery is older than the HCP Anywhere Enterprise Portal database, it can be missing some recent blocks. In this situation, you should rollback the HCP Anywhere Enterprise Portal database to an earlier point-in-time that matches the storage recovery.

If the storage recovery is newer than the HCP Anywhere Enterprise Portal database, since deleted blocks are kept for a minimum of 30 days and these blocks are never modified, as long as the HCP Anywhere Enterprise Portal database is no more than 30 days older, there will be no data loss.

Note: Running HCP Anywhere Enterprise FSCK is usually recommended following a disaster recovery. HCP Anywhere Enterprise FSCK must be run only with approval from Hitachi Vantara support.

Backing Up the HCP Anywhere Enterprise Portal Database

HCP Anywhere Enterprise Portal uses PostgreSQL to store metadata. This database **must** be backed up to ensure continued use of HCP Anywhere Enterprise Portal in order to ensure data and metadata persistence and consistency on the HCP Anywhere Enterprise platform, and to keep Recovery Time Objective (RTO) and Recovery Point Objective (RPO) values to a minimum.

Calculating the Minimum Space Required for the Database Backup

The archive pool must be at least double the data pool.

For example, if the primary database uses 1TB storage and a secondary, replication, database uses another 1TB, then each of these databases require 2TB for backup. The total storage must be at least 6TB: 1TB for the primary database, 1TB for the secondary database and 4TB for the primary and secondary backups (2TB for each backup).

See the following procedures:

- Using PostgreSQL Continuous Archiving – HCP Anywhere Enterprise Portal uses PostgreSQL's built-in continuous archiving mechanism to enable you to roll back to an older version of the portal database.
- Using PostgreSQL Streaming Replication – Streaming replication enables the continuous streaming and replication from a primary database server to a secondary, replication, server.

Configuring PostgreSQL for TLS

By default, PostgreSQL uses unencrypted connections. For more secure connections, you can enable Transport Layer Security (TLS) support on the PostgreSQL server.

Note: All communication between the HCP Anywhere Enterprise Edge Filer, HCP Anywhere Enterprise Drive Connect, and HCP Anywhere Enterprise Drive Share/Protect is encrypted. The communication between postgreSQL running on the primary database server and the secondary, replication server is **not** encrypted by default.

You **cannot** configure PostgreSQL for TLS if the HCP Anywhere Enterprise Portal uses a self-signed certificate or if the certificate has expired.

To enable TLS support on PostgreSQL:

1. Using SSH, log in as `root` to the HCP Anywhere Enterprise Portal primary database server.
2. Edit the **pg_ident.conf** file by adding a mapname and system username for your installation.
3. Restart the HCP Anywhere Enterprise Portal by running the following command: `portal-manage.sh restart`
4. Using SSH, log in as `root` to every server in the portal cluster and edit the **pg_ident.conf** file by making the same changes that were made on the primary database server.
5. On the primary database server, run the following command: `postgres-ssl-mode changingSSLMode true`
6. Restart the HCP Anywhere Enterprise Portal servers.

The HCP Anywhere Enterprise Portal restarts automatically.

Note: If a server was down when the HCP Anywhere Enterprise Portal restarts, when that server restarts, the configuration change will be automatically applied to it and it will then automatically restart.

Adding a New Server to the HCP Anywhere Enterprise Portal Cluster

If you add a new server to the portal cluster, after connecting the new server to the cluster, log in as `root` to the new server and edit the **pg_ident.conf** file as you did to the other servers.

Using PostgreSQL Continuous Archiving

HCP Anywhere Enterprise Portal uses PostgreSQL's built-in continuous archiving mechanism, known as Point-In-Time Recovery (PITR).

After backing up the database for the first time, the *base backup,* Incremental backups of the database are performed using Write Ahead Log (WAL) files, which include all the database transactions and changes. Running a base backup and incremental backups is non-disruptive: the database continues running, without losing any data.

The following diagram illustrates continuous archiving and point-in-time recovery:

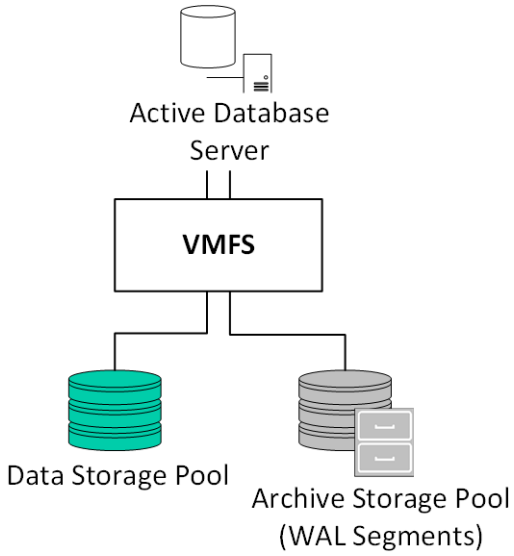

Using WAL logs has the following major benefits:

- Ensures data integrity
- Significantly reduces the number of disk writes

Note: The HCP Anywhere Enterprise Portal database WAL files are located in $pgdatadir/pg_wal which is commonly called the WAL directory.
The WAL file size is 16MB. Once a WAL file reaches 16MB, a new WAL file is created.
The maximum number of WAL files stored in the WAL directory is 128. This means that the WAL directory size can reach 16 (MB) * 128 (files) = 2GB.
When the 128-file threshold is reached, WAL files are recycled.

The database base backup and WAL logs are compressed upon backup:

- The base backup is typically 20%-30% of the size of the original database size. By default, there are always two base backups in the archive pool so the archived base backups typically consume around 40%-60% of the original database size.
- An archived WAL file size is typically 20% the size of the original WAL file.

Configuring PostgreSQL Continuous Archiving

To configure PostgreSQL continuous archiving:

1. Using SSH, log in as `root` to your HCP Anywhere Enterprise Portal primary database server.
2. In the command line, enter the following command to configure the maximum number of days to keep the backups: `portal.sh configure-db-recovery backup-history-days`

   Where *backup-history-days* is the number of days you want to retain a base backup archive before a new one is created. For example, to retain an archive for seven days, run: `portal.sh configure-db-recovery 7`

   An initial base backup of the database is created and the next backup is scheduled based on the backup-history-days parameter. Starting from the second base backup, the first scheduled base backup, there is always two base backups in the archive pool. WAL files are created after the first base backup. When a scheduled base backup is performed, the new base backup replaces both the old base backup that exceeded the backup-history-days, as well as the WAL files created in the period of time between the old base backup and the new base backup.

   Warning: The minimum retention period recommended by Hitachi Vantara is 7 days. If you set the retention period to less than 7 days, you must also have a secondary backup method in order to protect the portal from disasters.

When the command finishes successfully a message is displayed, similar to the following:

```
NOTICE: pg_stop_backup complete, all required WAL segments have been
archived
Done
```

3. Access the HCP Anywhere Enterprise Portal primary database server using Telnet and in the command line enter the following command to enable the portal to start monitoring: `start-replication-monitoring-task`

When the command finishes successfully a message is displayed, similar to the following:

```
NOTICE: pg_stop_backup complete, all required WAL segments have been
archived
Done
```

You can roll back to any older version of the database up until the previous base backup.

Note: You can change the backup-history-days parameter at a later date using the command `portal.sh set_archive_history_days history-days`

Rolling Back PostgreSQL Continuous Archiving to a Previous Point-in-Time

Note: Hitachi Vantara recommends reverting a database only with the help of Hitachi Vantara support.

After continuous archiving has been set up, you can roll back to an older version of the portal database.

Note: If the database has become corrupted and it is not just the data that you want to roll back, you need to revert the primary database to a snapshot, described in Reverting the Primary Database to a Snapshot.

To roll back PostgreSQL continuous archiving to a previous point in time:

1. Using SSH, log in as `root` to the HCP Anywhere Enterprise Portal primary database server.
2. In the command line, enter the following command to view the oldest time possible to roll back to: `portal.sh db-rollback -p`
3. Enter the following command to roll back to a point in time within the available backup range:
   `portal.sh db-rollback -r "point-in-time"`
   Where *point-in-time* is the desired point in time, in the format YYYY-MM-DD hh:mm:ss.
   For example: `portal.sh db-rollback -r "2026-05-06 00:23:21"`
   Note: If the HCP Anywhere Enterprise Messaging service was setup after the rolled back point-in-time, the messaging service status will be gray, as if messaging is not configured, and all existing messaging data will be lost.
4. When rolling back a primary database server, from which a replication server is streaming, described in Using PostgreSQL Streaming Replication, restart the HCP Anywhere Enterprise Portal on the secondary database server, the replicating server, by running:
   `portal-manage.sh restart`

Using PostgreSQL Streaming Replication

Streaming replication enables the continuous streaming and replication of WAL segments from the WAL directory and archived WAL segments from the primary database server to a secondary database server.

The following diagram illustrates streaming replication with continuous archiving and point-in-time recovery:

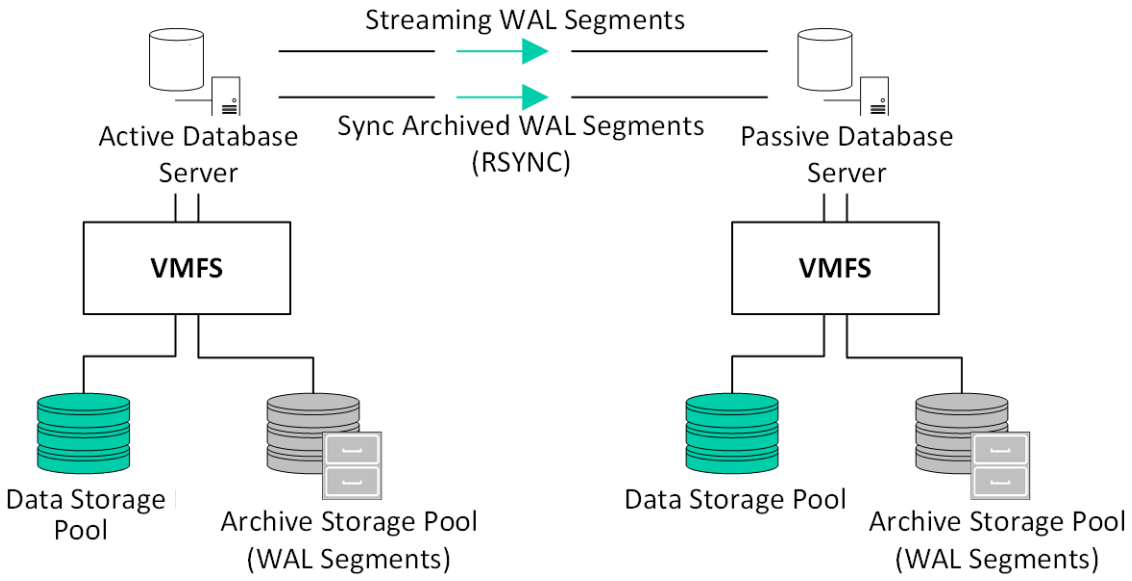

Note: Syncing archived WAL segments is done only on the first setup of the replica.

This process is complementary to configuring continuous archiving. It protects against failures occurring on the primary database server, by streaming the WAL directory and synchronizing the database's archived WAL logs to a secondary database server.

Note: You can also configure the replication so that the archive pool is on the passive, secondary, database and not the active database server with point-in-time recovery on the secondary.

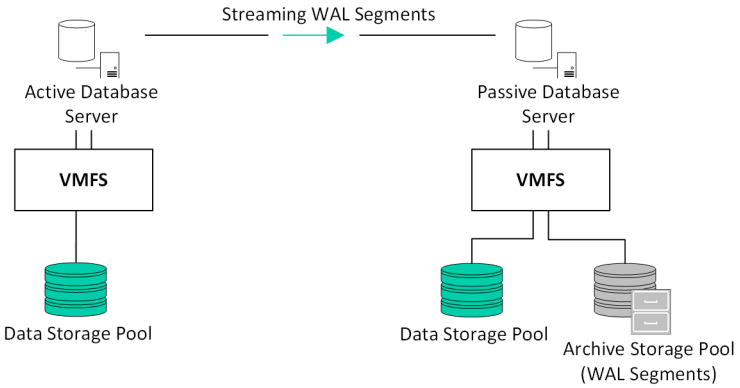

Configuring PostgreSQL Streaming Replication

The secondary server continuously streams the primary server's WAL segments, and synchronizes the primary database server's archived WAL segments every 10 minutes. Streaming replication traffic runs on TCP port 5432, which means you must open these ports for communication between the primary and secondary database servers. The secondary HCP Anywhere Enterprise Portal database server acts as a standby/passive server and cannot be used for load balancing purposes.

Note: The allocated disk space for the PostgreSQL archive pool must be at **least** twice the size of the storage space allocated for the HCP Anywhere Enterprise Portal data pool, as described in Calculating the Minimum Space Required for the Database Backup.
The locations of the streaming and archived WAL segments are `$pgdatadir/pg_wal` and `$DBarchive` respectively.

To configure streaming replication:

- Set up a replication server as described in Configuring Additional Servers.

You configure archiving only on the secondary server when the load on the primary server is very large. In most cases, configure streaming replication, as described above in Configuring PostgreSQL Streaming Replication. Before configuring streaming replication only on the secondary server, you **must** contact Hitachi Vantara support.

To configure archiving only on the secondary server:

1. Set up streaming replication, as described above in Configuring PostgreSQL Streaming Replication.
2. Using SSH, log in as `root` to your HCP Anywhere Enterprise Portal primary database server.
3. Disable archiving on the primary server by editing the `PG_REPLIC_MODE` field in **portal.cfg** file as follows:

   Change `PG_REPLIC_MODE=local_archive_streaming` to `PG_REPLIC_MODE=streaming`.
   You do not need to restart the primary server. However, if users or any tomcat process tries to create a base backup, a warning is issued in the image log. You can ignore this warning.
4. Using SSH, log in as `root` to your secondary server.
5. Enable archiving on the secondary server by editing the **portal.cfg** file on the secondary server, as follows:

Insert the field `SLAVE_ARCHIVE_FORCE=yes`

6. Restart the server by running the following command: `portal.sh restart`

If `SLAVE_ARCHIVE_FORCE=no`, the secondary server archive mode is the same as the primary server mode.

Changing the `SLAVE_ARCHIVE_FORCE` value always requires restarting the server.

Failing Over PostgreSQL Streaming Replication to the Secondary Database Server

The secondary database acts as a passive database, meaning it can only process read requests. In the event that the primary database fails, you have to fail over to the secondary database server, making it active, in order to assure proper continuity of the platform.

Note: You can also switch between the primary and secondary database servers, making the secondary database server the primary database server and the primary database server a secondary server, when the primary database server is still up.

To failover to the secondary server:

1. Using SSH, log in as `root` to the HCP Anywhere Enterprise Portal secondary database server.
2. In the command line, enter the following command: `portal-failover.sh become_master`

The primary database server becomes the secondary server, and the secondary database server becomes the primary server.

Failing Back to the Primary Database Server

When the original primary server comes back online you can failback to it, to return to the original configuration.

To prevent both database servers being identified as the primary database server, triggering inconsistencies in the databases, before you can return the original primary server to being the primary server, you must first define it as a replication, secondary, database.

To failback to the primary database server:

1. When the former primary database server is running again, using SSH, log in as `root` to the original primary database server.
2. Edit the **portal.cfg** file by changing the value of `CTERA_DB_ADDR` to be the new IP address of the original primary database server.
3. In the command line, enter the following command: `portal-failover.sh become_replica`.
4. Wait for the script to finish so that the original primary database is recognized as the secondary database.
5. Log in to the portal as a global administrator and In the global administration view, select **Main > Servers** in the navigation pane and click the replication server name.
6. Click **DB Replication** in the server window that is displayed and under **Database Replication** verify that the **Status** value is set to **OK**.

   If there is a mismatch between the requested WAL files and their location on the server the **Status** value can be set to **Failed** until the mismatch is resolved when the WAL file position reaches the location, which can take a few hours. Hitachi Vantara recommends the following

manual procedure to resolve this issue:

a) Log in to the HCP Anywhere Enterprise Portal as a global administrator and In the global administration view, select **Main > Servers** in the navigation pane.
The **Servers** page is displayed, listing all the servers for the HCP Anywhere Enterprise Portal.

b) Click the replication server name and in the server window that is displayed, under **General Settings**, uncheck **Replication of**.

c) Click **Save**.

d) Click the replication server name again and in the server window that is displayed, under **General Settings**, recheck **Replication of**.

e) Click **Save**.

The replication process will reinitialize which can be monitored by clicking **DB Replication** in the server window and under **Database Replication** verify that the **Status** value is set to **Reinitializing**. After the replication has reinitialized, which can take some time, depending on the HCP Anywhere Enterprise Portal size and the amount of data to be replicated, and the **Status** value is set to **OK**.

7. Only after the **Status** is set to **OK**, in the command line, enter the following command:
`portal-failover.sh become_master`

Monitoring the Database Backup and Streaming Replication

You can monitor every database backup and replication component running on the server, including:

- Streaming replication
- Base backup
- WAL archive (continuous archive)

Monitoring Database Backup and Replication from the Notifications Pane

To monitor database backup and replication from the Notifications page:

- In the global administration view, select **Main > Notifications** in the navigation pane.
The **Notifications** page is displayed.
Any alerts related to database backup or streaming replication are displayed.

Monitoring Database Backup and Replication from the Server Task Manager

If streaming replication is set up, you can monitor it from the Server Task Manager. The HCP Anywhere Enterprise Portal runs a task every few minutes to verify that replication is working as expected. If any issues are detected, the task fails, and the HCP Anywhere Enterprise Portal displays an appropriate notification in the **Notifications** page, and also sends an email alert to the portal administrators.

Note: The task runs the `portal.sh replication_status` command and analyzes the output. For more information see Monitoring Database Backup and Replication from the Server Console.

To monitor database backup and replication from the Server Task Manager:

1. In the global administration view, select **Main > Servers** in the navigation pane. The **Servers** page is displayed.

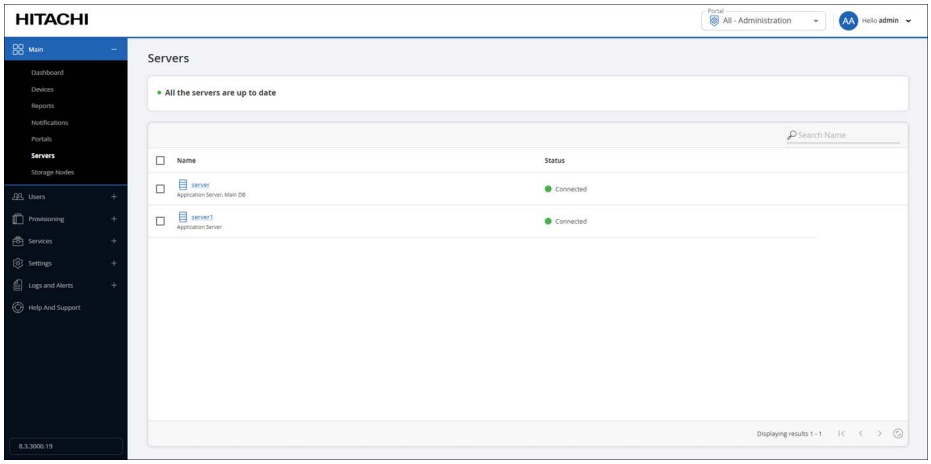

2. Click the primary database server.

   The server window is displayed with the server name as the window title.

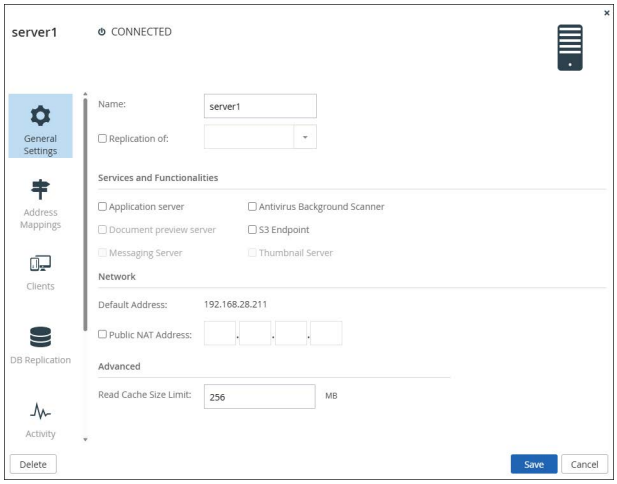

3. Click the **DB Replication** option.

   The HCP Anywhere Enterprise Portal reports the status of its scheduled base backups and transaction log archiving process, as well as additional metrics to help detect when database replication falls behind due to lags in the process. In the event that replication falls behind, HCP Anywhere Enterprise Portal administrators are notified via email. The relevant email templates are *Replication setup failed* and *Replication has errors*.

Backing Up the HCP Anywhere Enterprise Portal Database to an S3 Bucket

In addition to PostgreSQL's built-in continuous archiving mechanism and PostgreSQL streaming replication that enables the continuous streaming and replication from a primary database server to a secondary, replication, server, you can also backup and restore the HCP Anywhere Enterprise Portal database to and from an S3 bucket.

The HCP Anywhere Enterprise Backup to S3 includes backing up critical portal configurations, including networking, certificates, and passwords required to reconnect to the restored database.

Using HCP Anywhere Enterprise Backup to S3 allows administrators to back up the portal settings, base database backup, and archival logs automatically to a designated S3 bucket, ensuring data integrity and availability through consistent backups while providing administrators with a user-friendly interface for backup configuration and status monitoring.

Using HCP Anywhere Enterprise Backup to S3 provides an off-site database backup mechanism for disaster recovery.

Backing up the HCP Anywhere Enterprise Portal database to an S3 bucket and restoring it from the S3 bucket can be done in one of the following ways:

- Automatically, by configuration in the portal user interface, as described in Configuring Backups to an S3 Bucket.
- Using command line scripts for restoring the database.
- Using the HCP Anywhere Enterprise SDK to automate both configuration of the backup and restoring the database.

How It Works

HCP Anywhere Enterprise Backup to S3 uses a network efficient protocol to synchronize the portal database with the specified S3 bucket. It efficiently transfers only the necessary data, new or updated, to ensure the portal database and the S3 destination bucket are identical, saving time and bandwidth compared to a full copy. Updates are incremental and only modified data is transferred from the HCP Anywhere Enterprise Portal database to the S3 bucket. By only transferring changed data, HCP Anywhere Enterprise Backup to S3 reduces bandwidth usage and speeds up the synchronization process.

HCP Anywhere Enterprise Backup to S3 compares the HCP Anywhere Enterprise Portal database to the S3 bucket and only syncs data from the portal database to the S3 bucket that is missing or has been modified. It also deletes data from the S3 bucket that no longer exists in the HCP Anywhere Enterprise Portal database.

For an environment with a large portal database, HCP Anywhere Enterprise Backup to S3 periodically can reduce downtime.

Configuring Backups to an S3 Bucket

Note: An archive pool must have been created on the primary database server. For details refer to the installation documentation for your environment.
Before enabling backing up to an S3 bucket, the base backup of the portal database must be error free.

To configure an S3 bucket for backups from the user interface:

1. Sign in the portal as a global administrator and click Main > Servers.
The **Servers** page is displayed.

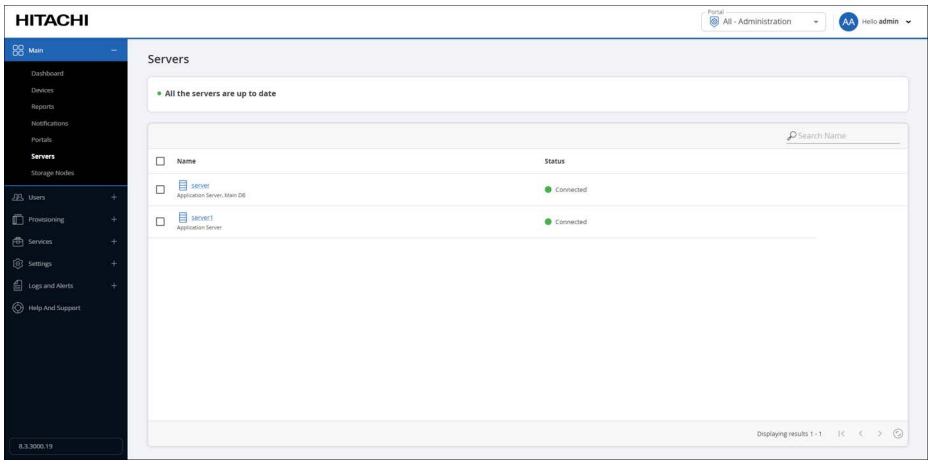

2. Click the primary database server.
The server window is displayed with the server name as the window title.

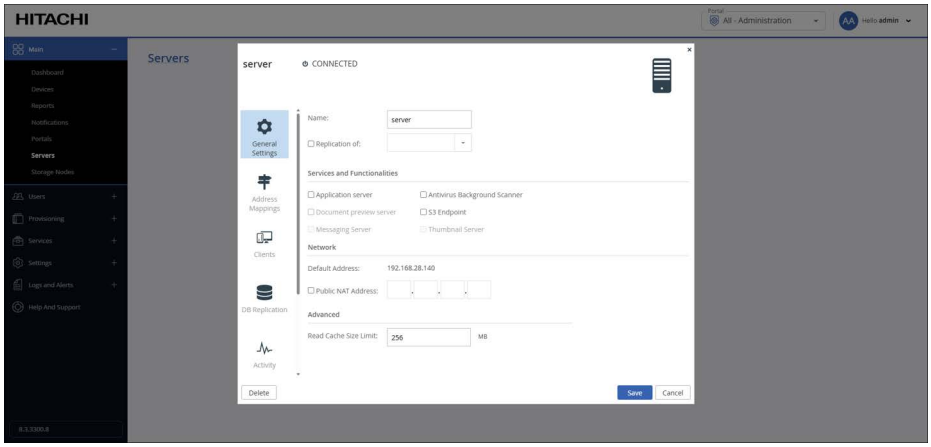

3. Click the **DB Replication** option.
The portal reports the status of its scheduled base backups and transaction log archiving process.
Under **Database backup** you can see that **Backup to Bucket** is not configured.
4. Click the **Backup to Bucket** tab.
5. Check **Enable Backup to Bucket**.
6. Complete the fields as follows:
Backup Interval – The interval between backups from 5 minutes to 43200 minutes (a month).

Under **Bucket details** define the S3 bucket to use for the database backup.

Type – The type of bucket to use. The following S3 buckets are supported:

Amazon S3
Caringo (S3)
Cloudian
Cohesity SmartFiles (S3)
EMC ESC (S3)
Generic (S3)
Google Cloud Storage (S3)
Hitachi Vantara HCP (S3)
IBM Cloud Object Storage (S3)
Isilon (S3)
MinIO (S3)
NetappS3
Nutanix (S3)
Quantum ActiveScale (S3)
QumulosS3
Scalitiy ARTESCA (S3)
Wasabi (S3)
Zadara (S3)

Bucket Name – The name of the bucket. The bucket cannot be a bucket that is used as a storage node nor as the bucket for the HCP Anywhere Enterprise Thumbnails service.

Access Key ID – The bucket access key ID.

Secret Access Key – The bucket secret access key.

Endpoint – The endpoint name of the bucket service. The port for the endpoint can be customized by adding the port after the URL, using a colon (:) separator. The default port is 80.

Use HTTPS – Use HTTPS to connect with the bucket.

7. Click **Save**.

After configuring backups to an S3 bucket the portal database is backed up to the bucket based on the same retention policy on the S3 bucket as is configured on the local archive pool, including the cleanup of old backups.

The **Status** field in the **Backup to Bucket** tab indicated the status of the connection.

Monitoring Backups to an S3 Bucket

The administrator can monitor the status of the backup at any time. The status indicator has the following states:

- Not Configured
- Disabled
- OK
- In Progress
- Failed

Both dashboard alerts and email notifications are provided in case or backup failures, for example when there is a disconnection to the S3 bucket. Errors are also written to the system log to ensure timely administrator intervention.

Restoring From the S3 Bucket

To restore the database you need a clean install of a portal server configured with an archive pool and the same database recovery settings of the original primary database server.

Note: Restoring the database can be done to any location.

You can then use a script on the new server to restore the database from the S3 bucket to the portal server.

Note: You can use the `portal.sh db-rollback` command to restore the database to a specific point in time.

Monitoring Database Backup and Replication from the Server Console

You can view the status of each database backup and replication component.

To monitor database backup and replication from the server console:

Note: This procedure can be performed on both the primary server and on secondary servers.

1. Using SSH, log in as `root` to the HCP Anywhere Enterprise Portal primary database server or secondary database server.
2. Run the following command: `portal.sh replication_status | json_reformat`

   The output is displayed in JSON format. For example:

```
{
  "streaming_replication": {
    "status": "ok",
    "lastSuccess": "256KB"
  },
  "base_backup": {
    "status": "ok",
    "lastSuccess": "none"
  },
  "wal_archive": {
    "status": "failed",
    "lastSuccess": "3 hours"
  }
}
```

   Both the primary and secondary servers have the same output structure. However, the fields have slightly different meanings:

Field	Primary Server	Secondary Server
streaming_replication		
status	Indicates whether a secondary server is streaming from the primary server: **ok** – Streaming replication is running. **failed** – Streaming replication failed. **not configured** – Streaming replication is not configured on this server.	Indicates whether this server is streaming from the primary server: **ok** – Streaming replication is running. **failed** – Streaming replication failed.

Field	Primary Server	Secondary Server
lastSuccess	The difference in size (in KB) between the last sent WAL segment and the secondary server's last checkpoint, a point in the WAL logs at which all data files have been written to the disk. In other words, the amount of data the primary database sent to the secondary database, which has not been processed by the secondary database.	The difference in size (in KB) between the last received WAL segment and the secondary server's last checkpoint, a point in the WAL logs at which all data files have been written to the disk. In other words, the amount of data the primary database sent to the secondary database, but which has not been processed by the secondary database. If the value is `none`, then the secondary database is up to date.

base_backup		
status	Indicates whether the last base backup was completed successfully: **ok** – Last base backup completed successfully. **failed** – Last base backup failed. **not configured** – Base backup is not configured for this server. This means that continuous archiving was not configured on the server.	Indicates whether the last base backup was completed successfully: **ok** – Last base backup completed successfully. **failed** – Last base backup failed.
lastSuccess	The last time a successful backup was run (in days). If the value is `none`, the `status` field is `ok`, and the last base backup was completed successfully.	
wal_archive		

Field	Primary Server	Secondary Server
status	Indicates whether a WAL was successfully archived in the past hour: **ok** – WAL directory was archived successfully in the past hour. **failed** – WAL directory was not archived in the past hour. This does not indicate a problem. You can change the status to `ok`, by running a manual command that forces archiving. For more information, contact Hitachi Vantara support. **not configured** – WAL archiving, that is continuous archiving, is not configured on this server. Archiving occurs after 16MB of data has been written or during a system restart.	Indicates whether a WAL was successfully synchronized in the past hour: **ok** – WAL directory was synchronized successfully in the past hour. **failed** – WAL directory was not synchronized in the past hour.
lastSuccess	The last time a WAL log was archived successfully (in hours). If the `status` field's value is `ok`, then `none` is output.	

3. Run the following command on the secondary server to view the secondary archive synchronization log: `portal-log.sh replication [-f]`

   Include the `-f` flag to display the log in the command window. Otherwise, the log is displayed in vim.

Reverting the Primary Database to a Snapshot

Note: Hitachi Vantara recommends reverting a database only with the help of Hitachi Vantara support.

When reverting a HCP Anywhere Enterprise Portal primary database to a previous snapshot, any files synced to the portal after the time of the snapshot **are missing in the HCP Anywhere Enterprise Portal**. This data is never synced to the portal when the devices are reconnected to the portal unless they physically existed on a device and not just as a stub file.

Note: Even though the data is no longer available, it will still be displayed as stub files on HCP Anywhere Enterprise Edge Filers. You can use these stub files to identify the data that was added to the HCP Anywhere Enterprise Portal after the revert date. These files should be deleted from the HCP Anywhere Enterprise Edge Filer. As long as the stub files exist on the HCP Anywhere Enterprise Edge Filer, the sync status will indicate that there is a sync error.

You have to revert to a snapshot of the virtual machine when the actual database has become corrupted and not the data itself. If it is only the data that needs reverting to a previous point-in-time, rollback the database, as described in Rolling Back PostgreSQL Continuous Archiving to a Previous Point-in-Time. When you rollback the database to a previous point-in-time, after continuous archiving has been set up, any files synced to the portal from the time of the database is rolled back to, to the present are synced when the devices are reconnected to the portal.

To ensure the portal and devices are synced after reverting to a snapshot:

Warning: This procedure must be performed immediately after reverting to the snapshot and before connecting any device to the HCP Anywhere Enterprise Portal.

1. Suspend syncing on all devices connected to the HCP Anywhere Enterprise Portal.
2. Roll back the HCP Anywhere Enterprise Portal primary database, as described in Rolling Back PostgreSQL Continuous Archiving to a Previous Point-in-Time.
3. Using SSH, log in as `root` to every HCP Anywhere Enterprise Portal application server except for the primary database server.
4. In the command line of every HCP Anywhere Enterprise Portal application server except for the primary database server, enter the following command to stop the server:
   `portal-manage.sh stop`
5. Using SSH, log in as `root` to the HCP Anywhere Enterprise Portal primary database server.
6. Access the telnet interface by entering `#telnet 0`
7. At the telnet prompt, input the login command `Login <portal-name> admin <GlobalAdminName>`

   Where *portal-name* is any string, such as *myportal* and *GlobalAdminName* is the actual global administrator username.
8. At the prompt, enter the password for the global administrator.
9. Enter the following command to change the Portal UUID: `change-portal-gvsn`
10. Restart all the application servers you stopped using the following command:
    `portal-manage.sh start`
11. Resume syncing on all devices connected to the HCP Anywhere Enterprise Portal.

Chapter 6. Additional Functionality for HCP Anywhere Enterprise Portal Servers

Enabling Federal Information Processing Standard (FIPS)

The Federal Information Processing Standard (FIPS) Publication 140-3, (FIPS PUB 140-3), is a U.S. government computer security standard used to approve cryptographic modules. Setting the HCP Anywhere Enterprise Portal to be FIPS compliant requires a HCP Anywhere Enterprise Portal restart and impacts performance.

To change the HCP Anywhere Enterprise Portal to be FIPS compliant, contact Hitachi Vantara support.

Enabling/Disabling Remote Support

You can enable remote assistance by the Hitachi Vantara support team. A user called *support* is created, which the Hitachi Vantara Support Team can uses to remotely access and log in to your HCP Anywhere Enterprise Portal.

To enable remote support:

1. Using SSH, log in as `root` to the HCP Anywhere Enterprise Portal primary server.
2. In the command line, enter the following command: `portal-support.sh enable`

To disable remote support:

1. Using SSH, log in as root to your HCP Anywhere Enterprise Portal primary server.
2. In the command line, enter the following command: `portal-support.sh disable`

Load Balancing HCP Anywhere Enterprise Portal Servers

General Load Balancing Best Practices

Probing to test tomcat reachability: Most load balancers have a health check/probing mechanism that checks for ports and services availability. The best scenario is to only use port tests that check if the port is available (checking ports 995 and 443). If a more accurate probing is required, use port 995 probe. With HTTPS use: *portalurl*`/admin/startup`.

Using F5 Load Balancer

Note: If you are not using F5 software for load balancing, the basic principles outlined here can still be applied.

The following describes setting up load balancing based on F5 software. If your version of the F5 software is different to the version described below, contact Hitachi Vantara support for help with your configuration.

Note: Using F5 load balancing to perform SSL offloading requires the following configuration:

- Create an F5 iRule to add **Secure** and **HttpOnly** flags to the JSESSIONID cookie.
- Create an F5 iRule to add **HSTS** flags.
- Disable old insecure encryption algorithms like RC4.
- If F5 is configured to use TLS 1.0, you must change it.

The following best practices are recommended by Hitachi Vantara:

- Configure the **tcp** TCP protocol profile.
  - If **Idle Timeout** is configured, make sure the value is at least 5 minutes, 300 seconds, as HCP Anywhere Enterprise handles its own TCP sessions with keep alives.
  - If **Keep Alive Interval** is configured, make sure the value is greater than the value specified for **Send CTTP keepalive messages every** in the virtual HCP Anywhere Enterprise Portal settings. **Send CTTP keepalive messages every** prevents proxy or load balancer servers from preemptively terminating connection between a HCP Anywhere Enterprise Agent and the HCP Anywhere Enterprise Portal. Hitachi Vantara recommends setting **Send CTTP keepalive messages every** less than half the value specified for **Keep Alive Interval**.
- If **Zero window Timeout** is configured, make sure it is as high as possible. For example, 30000.
- Configure the **source_addr** Persistence profile.
- After setting the profiles, set up the load balancing for the HCP Anywhere Enterprise virtual servers.

Chapter 7. ESXi Specific Management

Increasing the Data or Archive Pool Size

Note: Hitachi Vantara recommends changing the storage with the help of Hitachi Vantara Support.

To increase the storage:

1. For each HCP Anywhere Enterprise Portal instance you have installed, stop HCP Anywhere Enterprise services, by doing the following:
   a) Open an SSH session to the HCP Anywhere Enterprise Portal instance.
   b) Log in as `root` user.
   c) Run the command to stop the HCP Anywhere Enterprise Portal services: `portal-manage.sh stop`
2. In the vSphere Client console, right-click the HCP Anywhere Enterprise Portal virtual machine and click **Edit Settings**.

   The virtual machine settings are displayed.
3. In the **Virtual Hardware** tab, click select **Hard Disk** under **ADD NEW DEVICE**.

   A new hard disk is added to the settings.
4. Specify the disk size you want for the new disk.
5. Expand the **New Hard disk** item and for **Location** browse to the datastore you want to use for either the data pool or archive pool.
6. Select the disk to use.
7. Click **OK**.

   The new VMDK hard disk is added.
8. Log in as `root` over SSH to the HCP Anywhere Enterprise Portal server with the archive pool to increase.
9. Run the following command to identify the name of the device you attached: `fdisk -l`
10. Locate the device name according to its size and usage information.
11. Run the following command to extend the data pool: `portal-storage-util.sh extend_storage` *`deviceName`*

    where *deviceName* is the device name identified when you ran `fdisk -l`.
12. Run the following command to extend the archive pool: `portal-storage-util.sh extend_db_archive_pool` *`deviceName`*

    where *deviceName* is the device name identified when you ran `fdisk -l`.
    For example: `portal-storage-util.sh extend_db_archive_pool sde`
13. For each HCP Anywhere Enterprise Portal instance, starting from the HCP Anywhere Enterprise Portal main database and proceeding to the application servers, start HCP Anywhere Enterprise Portal services.
    a) Log in as root user to each HCP Anywhere Enterprise Portal server over SSH.
    b) Run the following command to start the portal services: `portal-manage.sh start`

Protecting the HCP Anywhere Enterprise Portal Main Database Using vSphere HA

The HCP Anywhere Enterprise Portal platform must be protected from datacenter and virtualization host failure and maintain data integrity and business continuity. This can be achieved through HCP Anywhere Enterprise Portal's native, application level, database archiving and replication procedures but can also be managed at the ESXi layer.

To protect the HCP Anywhere Enterprise Portal platform from datacenter and virtualization host failure at the ESXi layer:

- vSphere HA Hardware and Software Requirements
- Configuring vSphere HA
- Testing vSphere HA Failover

vSphere HA Hardware and Software Requirements

Before starting to configure vSphere HA and protect the HCP Anywhere Enterprise Main Database VM, review VMware's vSphere HA Hardware and Software requirements. These requirements must be addressed prior to any configuration and VM deployments.

Configuring vSphere HA

Perform the following procedure to enable vSphere HA for your datacenter.

To configure vSphere HA:

1. Login in to your vCenter Server using vSphere client.
2. Right-click on the cluster you want to protect using vSphere HA on and click **Settings**.
3. Browse to **vSphere Availability** under **Services**.
4. Click **EDIT** for vSphere HA.

   The **Edit Cluster Settings** window is displayed.
5. Enable **vSphere HA**.
6. In the **Failures and responses** tab make sure that **Enable Host Monitoring** is on and set **Host Failure Response** to **Restart VMs**.
7. In the **Heartbeat Datastores** tab, select the shared datastores that will be used for vSphere HA.

   Note: You must select a minimum of two datastores. For automatic detection of accessible datastores, choose **Automatically select datastores accessible from the hosts**.
8. Click **OK**.

vSphere HA is now enabled on the cluster.

Once vSphere HA is enabled, using the dedicated, shared datastores, virtual machines deployed to hosts participating in your cluster are automatically protected.

Testing vSphere HA Failover

Virtual machines deployed to hosts participating in the vSphere HA-enabled cluster are protected.

To perform a failover test:

1. Deploy the HCP Anywhere Enterprise Portal Main Database VM to one of the hosts is the cluster with vSphere HA set.
2. Upload sample data to the HCP Anywhere Enterprise Global File System through HCP Anywhere Enterprise Portal's web application or via SMB using a HCP Anywhere Enterprise Edge Filer.
3. After successfully uploading files to the HCP Anywhere Enterprise Global File System, power off the vSphere host running the HCP Anywhere Enterprise Portal Main Database.

vSphere will provide an alert indicating that the failover process has started and will begin restarting the HCP Anywhere Enterprise Portal Main Database VM on a second host.

Upon completing the failover, vSphere will provide an alert indicating that the HCP Anywhere Enterprise Portal Main Database virtual machine is back up and running.

The failover and failback processes are performed automatically, requiring no manual intervention.

Chapter 8. Upgrading HCP Anywhere Enterprise Portal

The procedure to upgrade an existing HCP Anywhere Enterprise Portal installation is dependent on whether the upgrade involves upgrading the portal software or both the portal image and the software.

As a general rule, new releases include both image and software upgrades and Hitachi Vantara strongly recommends that you upgrade both. The release notes for each release specifies whether a portal image upgrade is required in addition to a portal software upgrade.

You upgrade both the portal image and the software (Upgrading the HCP Anywhere Enterprise Portal Image and Software) or just the software (Upgrading the HCP Anywhere Enterprise Portal Software).

Hitachi Vantara recommends upgrading software and a new image with the help of Hitachi Vantara Support.

Make sure that the required ports are open.

Upgrading a HCP Anywhere Enterprise Portal to 8.3.3300.x

HCP Anywhere Enterprise Portal platform services require two networks. The second network is internal and not routable and must be defined on every server in the portal cluster.

For ESXi:

1. For every server, right-click the HCP Anywhere Enterprise Portal VM in the navigation page and choose **Edit Settings**.
2. Click **ADD NEW DEVICE > Network Adapter**.

   A new network is added to the list of networks.
3. Click **OK**.
4. Run the following commands after adding the internal network interface to remove the default route (adapter name: eth1):

```
export ETHNAME=`nmcli -t -f NAME,DEVICE connection show | grep
":eth1" | cut -d: -f1`
nmcli connection modify "$ETHNAME" ipv4.ignore-auto-routes yes
nmcli connection modify "$ETHNAME" ipv4.gateway ""
nmcli connection down "$ETHNAME"
nmcli connection up "$ETHNAME"
```

Upgrading the HCP Anywhere Enterprise Portal Image and Software

You have to shutdown all the HCP Anywhere Enterprise Portal servers before upgrading the image.

HCP Anywhere Enterprise Portal requires two networks, one for external client-facing traffic and one for internal communication. You require these two networks on each portal server.

Warning: **If the current network configuration uses eth0 for the portal network and eth1 is free, then the upgrade can proceed. If, however, eth0 or eth1 are not available, you must contact Hitachi Vantara Support, as changes need to be made to portal.cfg before you can continue.**

To upgrade the HCP Anywhere Enterprise Portal image and software:

1. Shutdown all the HCP Anywhere Enterprise Portal servers.
   First stop all application and preview servers. Next stop the main database server and finally stop the replication database server.
   a) Using SSH, log in as `root` to the HCP Anywhere Enterprise Portal server.
   b) Run the following command: `portal-manage.sh stop`
   When the server has stopped, the `Done` message is displayed on the screen.
2. When all servers are in a stop state, add a second network to each server.
3. Upgrade the portal image on each server.
   First upgrade and start the main database server. Next upgrade and restart all application and preview servers. Finally, upgrade and restart the replication database server.
   You restart each server using the command `portal-manage.sh start`.
   a) Using SSH, log in as `root` to the HCP Anywhere Enterprise Portal server.
   b) Copy the upgrade zip file to the machine.
   c) Navigate to the location of the zip file and run the following command to extract the portal image file: `tar xvf file` where *file* is the zip filename. The file contains the rpm packages for the image.
   d) Navigate to the directory where the unzipped file is and run the following command:
   `chmod +x rpm_upgrade.sh && ./rpm_upgrade.sh`

Patch upgrades are relatively quick as only the packages that have changed are updated. An upgrade to the next major version takes more time.

Note: You can also check that the upgrade has succeeded by running `portal-manage.sh status` which displays the image version and service, software, version.
When the release notes specify that the image upgrade also involves a kernel upgrade, an additional restart is also required to complete the kernel upgrade, after restarting each server.
Stop the portal: `portal-manage.sh stop`
Start the portal: `portal-manage.sh start`

Upgrading a HCP Anywhere Enterprise Portal with the HCP Anywhere Enterprise Messaging Service Implemented

After the upgrade, run the following command on every messaging server: curl `http://localhost:8083/connector-plugins`

The expected output:

```
[{"class":"io.confluent.connect.jdbc.JdbcSinkConnec-
tor","type":"sink","version":"10.5.1"},
{"class":"io.confluent.connect.jdbc.JdbcSourceConnec-
tor","type":"source","version":"10.5.1"},
{"class":"io.debezium.connector.postgresql.PostgresConnec-
tor","type":"source","version":"1.9.5.Final"},
{"class":"org.apache.kafka.connect.mirror.MirrorCheckpointConnec-
tor","type":"source","version":"7.2.1-ccs"},
{"class":"org.apache.kafka.connect.mirror.MirrorHeartbeatConnec-
tor","type":"source","version":"7.2.1-ccs"},
{"class":"org.apache.kafka.connect.mirror.MirrorSourceConnec-
tor","type":"source","version":"7.2.1-ccs"}]
```

If there is a different output, run `docker restart kafka_connect_kafka_connect_1`

Upgrading a HCP Anywhere Enterprise Portal with Subfolder Quotas Configured

A migration using the HCP Anywhere Enterprise SDK is required to configure the subfolder quotas in the portal. For details, contact Hitachi Vantara Support.

Upgrading the HCP Anywhere Enterprise Portal Software

You can upgrade the HCP Anywhere Enterprise Portal software in the HCP Anywhere Enterprise Portal user interface or via the command line.

Note: Normally, you upgrade the portal image, which includes upgrading the software.

To upgrade the HCP Anywhere Enterprise Portal software via the user interface:

Warning: You must contact Hitachi Vantara Support before using the UI to upgrade the HCP Anywhere Enterprise Portal software.

To upgrade the HCP Anywhere Enterprise Portal software via CLI:

1. Stop the HCP Anywhere Enterprise Portal servers.
   First stop all application and preview servers. Next stop the main database server and finally stop the replication database server, if available.
   a) Using SSH, log in as `root` to the HCP Anywhere Enterprise Portal server.
   b) Run the following command: `portal-manage.sh stop`
   Once services are stopped, the `Done` message is displayed on the screen.
2. When all servers are in a stop state, upgrade the primary database portal software.

a) Using SSH, log in as `root` to the HCP Anywhere Enterprise Portal server.

b) Upgrade the HCP Anywhere Enterprise Portal software: `portal-manage.sh upgrade upgrade_file`
where *upgrade_file* is the software file provided by Hitachi Vantara.

3. Restart the primary database server: `portal-manage.sh start`

4. After the primary database server is running, using SSH, log in as root to the secondary HCP Anywhere Enterprise Portal servers.

a) Upgrade the HCP Anywhere Enterprise Portal software on each secondary HCP Anywhere Enterprise Portal server: `portal-manage.sh upgrade upgrade_file`
where *upgrade_file* is the software file provided by Hitachi Vantara.

b) Restart the HCP Anywhere Enterprise Portal servers.
First start the replication database server, if available. Next, start the application and preview servers.
Using SSH, log in as `root` to the HCP Anywhere Enterprise Portal server and start the HCP Anywhere Enterprise Portal by running the command: `portal-manage.sh start`

Hitachi Vantara

Corporate Headquarters 2535 Augustine Drive
Santa Clara, CA 95054 USA www.HitachiVantara.com community.HitachiVantara.com

Regional Contact Information
Americas: +1 866 374 5822 or info@hitachivantara.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hitachivantara.com
Asia Pacific: +852 3189 7900 or info.marketing.apac@hitachivantara.com