Hewlett Packard Enterprise

Installation Guide for HPE Alletra 6000

Notices

Acknowledgments

Support

All documentation and knowledge base articles are available on HPE InfoSight at **https://infosight.hpe.com**. To register for HPE InfoSight, click the *Sign up* link on the HPE Sign In page.

Email: **support@nimblestorage.com**

For all other general support contact information, go to **https://www.hpe.com/us/en/services/nimble-storage.html**.

Revision history

Part number	Publication date	Edition	Summary of changes

Contents

Install the HPE Storage Array

There are several steps to install the HPE storage array, such as providing a suitable operating environment, installing the array into your rack, and making network connections appropriate to your current and anticipated topologies. After the array is configured, you can perform maintenance and upgrade tasks as needed.

Before beginning, download and read the Release Notes, available on HPE InfoSight (**https://infosight.hpe.com**). If you do not have an HPE InfoSight account, click **Create Account** and provide the required information.

Before you can start using the array, you must activate your software subscriptions and enable the array for cloud management, as described in **HPE Cloud Enablement**.

Visit the Welcome Center (**https://infosight.hpe.com/welcomecenter**) for help setting up your array.

Array Management with Data Services Cloud Console

HPE Alletra series arrays require network connectivity to Data Services Cloud Console to perform operations, such as creating a volume, pool, partner, or snapshot. To use Data Services Cloud Console, you must have a cloud account set up with HPE GreenLake. This account consists of both an HPE GreenLake user account and an HPE GreenLake Workspace. You use the cloud account to connect with Data Services Cloud Console.

The following sections provide more details on HPE GreenLake and Data Services Cloud Console.

HPE GreenLake

HPE GreenLake offers a common set of cloud services that enable a consistent, cloud-qualified customer experience. HPE GreenLake is designed to combine the cloud's agility with the governance, compliance, and visibility of the hybrid cloud model.

Key features of HPE GreenLake make it easy for new cloud users to get started while offering powerful capabilities for advanced users as well:

- Global data management for streamlined configuration and deployment of devices. HPE GreenLake supports device management, which enables you to provision and manage multiple devices that have similar configuration requirements with less administrative overhead.

- A secure cloud-based platform.

- A rich API that enables you to implement data management functionality.

Data Services Cloud Console

Data Services Cloud Console is a secure cloud application, deployed on HPE GreenLake, that provides a control plane for simplifying data infrastructure management and delivering data services across edge-to-cloud environments. It contains software-as-a-service (SaaS) offerings that enable global data management, including services such as:

- Data Ops Manager

- Backup and Recovery

- Setup Service

- Block Storage

HPE Cloud Enablement

	Step
1.	**Activate Software Subscriptions**
2.	If you have not already created your HPE GreenLake environment, do the following
	a. **Create an HPE GreenLake User Account**
	b. **Create an HPE GreenLake Workspace**
	c. **Add Data Services Cloud Console to Your HPE GreenLake Workspace**
	d. **Add Roles and Permissions to Your HPE GreenLake User Account**
3.	**Add and Assign the Array to Data Services Cloud Console**

Activate Software Subscriptions

About this task

Software subscriptions must be active before you can start using the array.

Prerequisites

Locate the **Electronic Software Delivery Receipt** email you should have received from HPE. This email contains a link that will enable you to activate your software Subscriptions.

Procedure

1. In the **Electronic Software Delivery Receipt** email, click **Access Your Products**.

2. Log in to **My HPE Software Center** with your HPE Passport account credentials. If you do not have an account, click **Sign In** and then **Sign up** to create one.

3. On the **Activate** page, under **Step 1: Select Products to Activate**, select the products and click **Next**.

4. Under **Step 2: Designate Activatee**, select whether you are activating products for yourself or another user, and then click **Activate**. When activation completes, HPE sends an Activation Receipt.

What to do next

Create an HPE GreenLake user account.

Create an HPE GreenLake User Account

About this task

To start using HPE GreenLake Common Cloud Services, you need to create a user account.

Prerequisites

Locate the HPE Cloud invitation email you should have received from HPE. This email contains a link that will take you directly to the HPE **Sign In** page. If you have a current HPE user account (HPE Support or HPE InfoSight), use the same credentials to sign in, and then go to **Create an HPE GreenLake Workspace**. If you do not have a current user account, do the following to create one.

Procedure

1. In the HPE Cloud invitation email, click **HPE Set Up Account**. Or, go to **https://common.cloud.hpe.com** and click **Sign up** at the bottom of the page.

2. In the **Create an HPE Account** page, specify your account information, and then click **Create Account**.

3. Monitor your email account for a verification email. Click **Verify Email** to verify your identity.

What to do next

Create an HPE GreenLake Company Account.

Create an HPE GreenLake Workspace

About this task

By creating the workspace, you are by default the super admin for the account.

Prerequisites

You must have an existing HPE user account.

Procedure

1. If you are not already, log in to your **HPE GreenLake** user account.

2. Click **Create Workspace**.

3. Provide the following information:

    *   Workspace Name

    *   Workspace Country

    *   Street Address

    *   City, State

    *   ZIP/Postal Code

    *   Phone Number

    *   Email

4. Accept the legal terms.

5. Click **Create Workspace**.

    Your workspace dashboard opens.

What to do next

Add the Data Services Cloud Console to your HPE GreenLake workspace, as described in the following section.

Add Data Services Cloud Console to Your HPE GreenLake Workspace

About this task

To use the Data Services Cloud Console, you need to add it to your HPE GreenLake Workspace.

Procedure

1. In the HPE GreenLake menu bar, click **Applications**.

2. Click **Available Applications**.

3. In the **Data Service Cloud Console** card, click **View Details**.

4. Select the **Region** from which to install the application.

5. Click **Set Up Application**.

6. Select the region in which to deploy the application.

7. Accept the terms and conditions, then click **Deploy**.

   It may take some time to provision the Data Services Cloud Console application to your workspace. Provisioned apps are added to your My Apps page.

   > (!) **IMPORTANT:** Before launching the Data Services Cloud Console, you must add the permissions necessary to access the application, as described in *Add Roles to Your HPE GreenLake User Account*.

Add Roles and Permissions to Your HPE GreenLake User Account

About this task

Add the permissions necessary to access the Data Services Cloud Console.

Procedure

1. In the HPE GreenLake menu bar, click **Manage**.

2. Click **Identity & Access**.

3. In the **Assign Roles** card, click **Assign a Role**.

4. Select the **User**.

5. Select **Data Services Cloud Console** as the **Application**.

6. Select **Administrator** as the **Role**.

7. Leave the **Limit Resource Access** option disabled..

8. Click **Assign Role**.

9. Click **Change Role Assignment** to confirm the assignment.

10. Back on the **Identity & Access** page, click the **Users** card.

11. Find the user, and select **View Details** from the corresponding more (**...**) menu. Verify that the user has the correct permissions.

12. In the HPE GreenLake menu bar, click **Applications**.

13. Click **Data Services Cloud Console**.

14. Click **Launch**.

15. Verify that the following Data Services Cloud Console services are available:

- Setup Service

- Block Storage

- Data Ops Manager

Add and Assign the Array to Data Services Cloud Console

To manage the storage array, you must add and assign it to Data Services Cloud Console.

Before starting, verify that you received an email with the array serial number and subscription key. You will need both to add the array to Data Services Cloud Console. If you do not have this information, skip this procedure until you obtain it. Consider the following resources:

- The array serial number is also located on the rear of the array, on either a pull-out tab, or a label.

- Later in the initial setup process, when you are connecting the array to HPE GreenLake, the Cloud Connectivity Wizard will display the subscription key.

The method for performing this operation depends on whether the array is owned or subscribed.

Add and Assign an Owned Array

Procedure

1. On the HPE GreenLake menu bar, click **Devices**.

2. Click **Add Devices**.

3. Select **Storage Devices** as the **Device Type**, then click **Next**.

4. Select **Purchase or Lease** as the **Ownership Type**.

5. Type the device **Serial Number** and **Subscription Key**, and then click **Enter**. Repeat this step for each additional device. Click **Next** when done.

6. Optionally assign tags to the device. Tags are used to organize/label inventory resources. Click **Next**.

7. Verify your settings, then click **Finish**.

8. In the **Inventory** list, select the device that you just added, then select **Assign to Application**.

9. Select **Data Services Cloud Console** for the **Application**, and then select the **Application Instance** closest to where you have Data Services Cloud Console installed. Click **Finish**, and then **Close**.

10. Back on the **Inventory** page, verify that the array is listed and that it is assigned to the correct Data Services Cloud Console instance.

Add and Assign a Subscribed Array

Procedure

1. Add the array subscription key:

a. On the HPE GreenLake menu bar, click **Manage**.

b. Click **Subscriptions**.

c. Click **Add Device Subscriptions**.

d. Type the array subscription key, and then click **Enter**. Repeat this step for any additional subscriptions.

e. Click **Add Subscriptions**.

2. Add the array serial number:

a. On the menu bar, click **Devices**.

b. Click **Add Devices**.

c. Select **Storage Devices** as the **Device Type**, and then click **Next**.

d. Select **Infrastructure as a Service** as the **Ownership Type**.

e. Type the array **Serial Number**, and then click **Enter**. Repeat this step for each additional array. Click **Next** when done.

f. Optionally assign tags to the storage array. Tags are used to organize/label inventory resources. Click **Next**.

g. Verify your settings, and then click **Finish**.

3. Assign the array to Data Services Cloud Console:

a. At the top of the **Inventory** list, click **Require App Assignments**.

b. Select the array you just added, and then click **Assign to Application**.

c. Select the **Application**, and then select the **Application Instance** (region) closest to where you are installing the array. Click **Next**.

4. Pair the subscription key with the array:

a. At the top of the **Inventory** list, click **Require Subscriptions**.

b. Select the array you just added, and then click **Apply Subscriptions**.

c. Click **Apply Subscriptions**, and then click **Finish**.

d. Select the **Subscription Key**, and then click **Apply Subscriptions**.

e. Click **Finish**.

5. Click **Finish**, and then **Close**.

6. Back on the **Inventory** page, verify that the array is listed and that it is assigned to the correct Data Services Cloud Console instance.

Troubleshoot Connectivity Issues with Data Services Cloud Console

When an array is experiencing connectivity issues, one or more of the following alarms may display. In each case, the array will attempt to connect on its own.

- `Unable to activate array AF-999999 using the activation portal`

- `Array AF-999999 could not connect to the cloud console`

- `Array AF-999999 could not validate the subscription using the cloud console`

If the array is unable to connect:

- Verify that the **firewall rules** allow the array to connect with HPE.

- Check the DNS and proxy settings on the array.

  Network connectivity is required from the Data Orchestrator to Data Services Cloud Console. Network time protocol servers (NTP) must be configured on the Data Orchestrator and must be allowed to be in sync with global time services. DNS name resolution must be configured on the Data Orchestrator and be allowed to resolve global Internet names and domains. The Data Orchestrator will periodically transmit health and statistical telemetry over HTTPS/TLS.

If you are unable to resolve the issue, contact support.

Network Considerations

In preparation for installing your storage array, information in this section will help you:

- Determine your **network topology**

- Plan for the configuration of **IP Addresses**

- Identify options for **subnets**, including (for iSCSI):
  - Discovery IP addressing

  - IP address zones

  - Traffic assignments

- Understand data connections to your iSCSI or Fibre Channel array through **Interface Pairs**

Use the **Installation Checklist** to note all information required before installing your array. For more information, refer to the *GUI Administration Guide* or the *CLI Administration Guide*.

Network Topology

When you initially configure the arrays, you are prompted to choose a network topology. The selection that you make enables the setup wizard to assist in making the optimal networking choices:

- One shared network

- Two dedicated networks

- Advanced network configuration

- Dual Fabric Fibre Channel

(!) **IMPORTANT:** The following diagrams are conceptual examples of the network connections you might use for each topology. These diagrams are not intended to show a complete set of connections.

One Shared Network

In this configuration, the management and data traffic share the same physical network interfaces. If more than one switch is used in this configuration, the switches must be connected to each other by an inter-switch link or trunk port, or they must be configured in a single switch stack if supported by the switch manufacturer. You may want to consider this configuration if you have only one network or if your data traffic is low enough that adding management traffic does not affect the LAN.
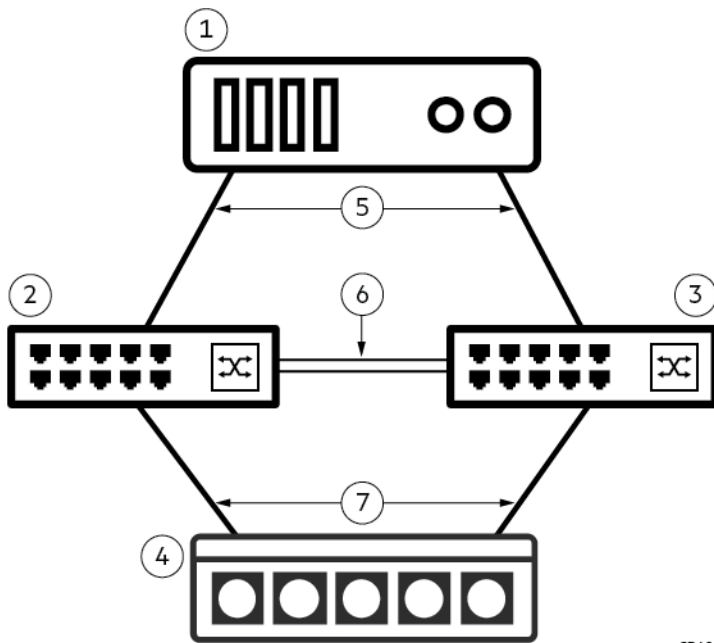
Figure 1: Topology diagram - network connections for one shared network

1	Host/server
2	Ethernet switch 1
3	Ethernet switch 2
4	Array
5	Host data/management connections
6	Inter-switch link
7	Array data/management connections

Two Dedicated Networks

In this configuration, management traffic and data traffic are separated into two subnets. Management traffic uses the ports assigned to the management subnet. All other subnets are designated as data only. Two switches for the management traffic and two switches for the data traffic are recommended for redundancy and high availability. Multiple switches carrying a particular subnet must be connected with an inter-switch link, trunk port, or switch stacking modules. If that subnet is assigned to a particular VLAN, the inter-switch link or trunk port must allow that VLAN.

This configuration option is probably best in cases where security is an issue, where network administrators do not need access to data traffic, where the amount of data traffic is such that it has a dedicated network, or where the network for data traffic is completely separate from the management network.

Choose this option:

* To limit iSCSI traffic from the management network

* To isolate the subnets for network security

* If you want a dedicated broadcast domain reserved for iSCSI traffic
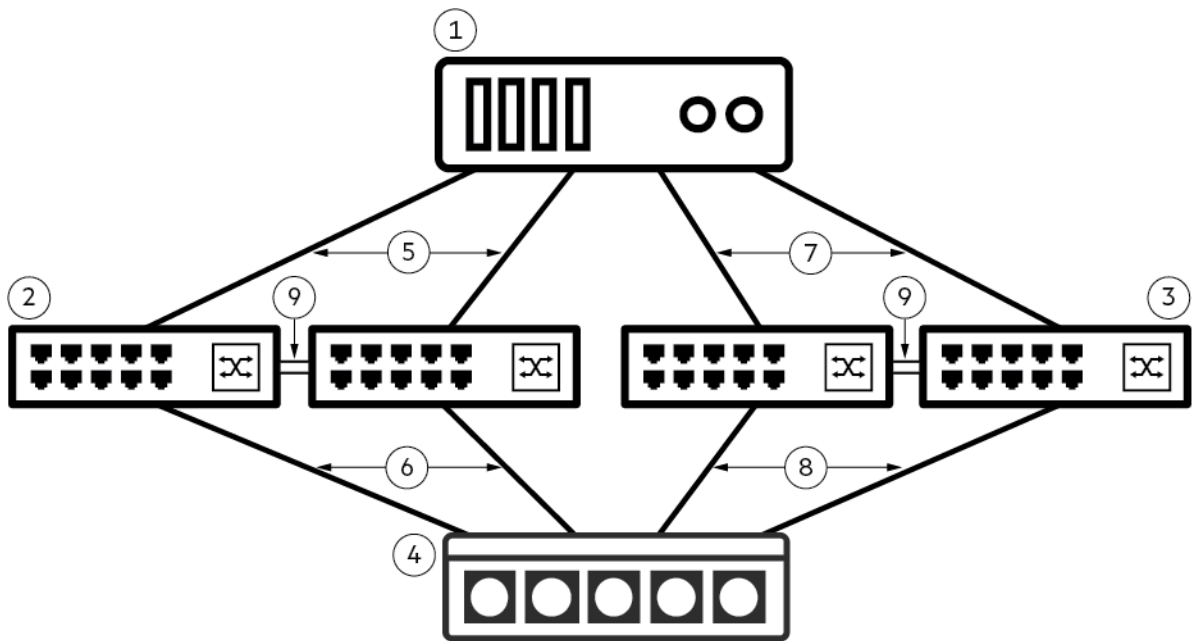
Figure 2: Topology diagram - network connection for two dedicated networks

1	Host/server
2	Ethernet switch (management)
3	Ethernet switch (data)
4	Array
5	Host management connections
6	Array management connections
7	Host data connections
8	Array data connections
9	Inter-switch link

Advanced Network Configuration

From the **Advanced and advanced multipath requirements** selection of the setup wizard, you can define a configuration for a dedicated management port or define any number of ports to handle the desired I/O flow. The advanced network configuration allows the greatest flexibility to suit your network traffic and advanced multipath requirements.

This particular example shows the management and data paths on different networks. Data ports are connected to subnets that carry iSCSI traffic, and you can add more data connections to additional ports as needed. Note that the subnet addresses in the illustration are just examples. This arrangement enables the maximum data throughput by isolating management traffic from the data.
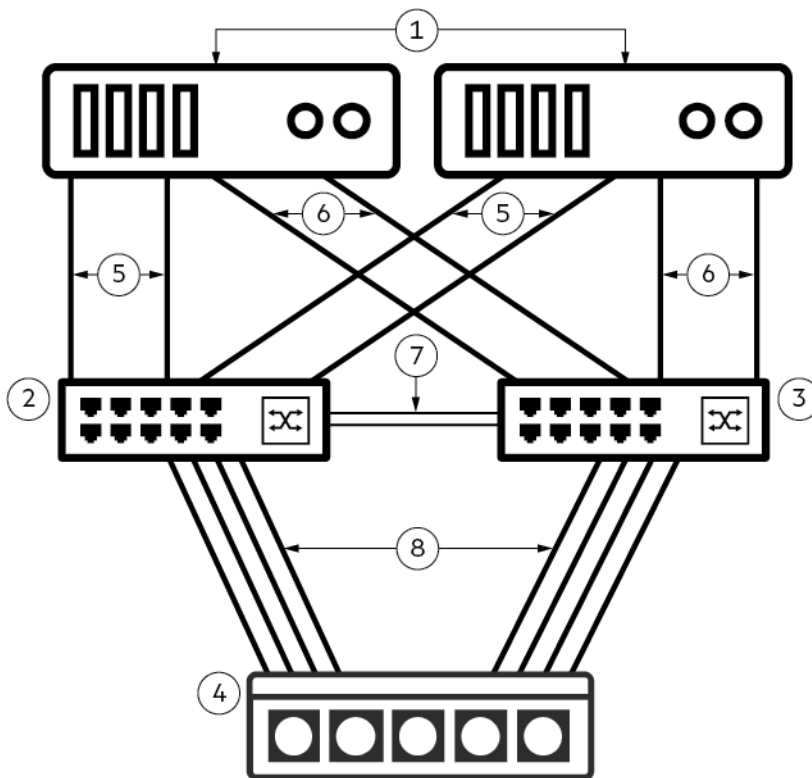
Figure 3: Topology diagram - advanced network configuration, data connections

1	Host/server
2	Ethernet switch 1, subnet 10.10.30.x/24
3	Ethernet switch 2, subnet 10.10.20.x/24
4	Array
5	Host data connections, switch 1
6	Host data connections, switch 2
7	Inter-switch link
8	Array data connections
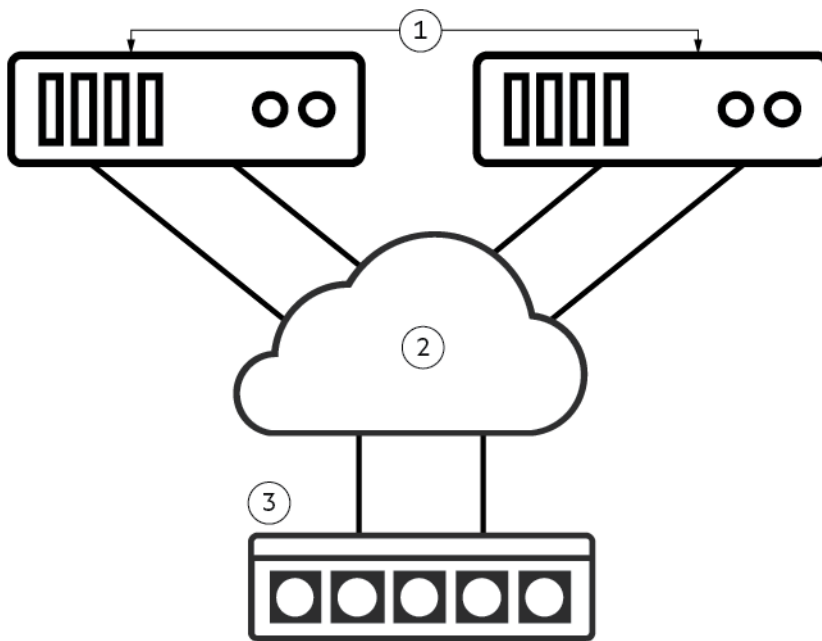
Figure 4: Topology diagram - advanced network configuration, management connections

1	Host/server
2	Management network
3	Array

Dual Fabric Fibre Channel

In this configuration, two Fibre Channel ports on a host HBA are connected to two separate Fibre Channel switches, with each Fibre Channel switch connected to both controllers on the array. The host management port and the controller management ports are connected through an Ethernet switch.

This configuration provides redundant paths for Fibre Channel traffic. If one of the host Fibre Channel ports fails, or one of the Fibre Channel switches fails, the host remains connected to both controllers through the remaining host Fibre Channel port or Fibre Channel switch, and Fibre Channel traffic continues uninterrupted.

Note that only one data connection is shown. You can add more connections and use different ports as needed.
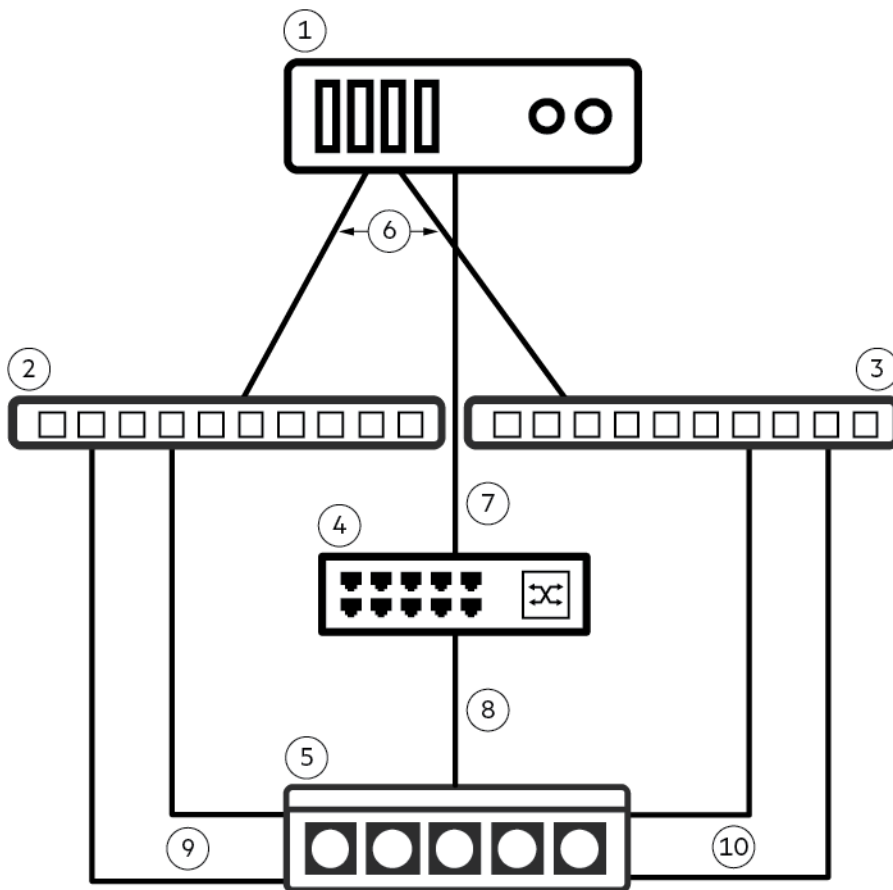
Figure 5: Topology diagram - network connections for dual fabric Fibre Channel

1	Host/server
2	Fibre Channel switch 1
3	Fibre Channel switch 2
4	Ethernet switch
5	Fibre Channel array
6	Host data connections
7	Host management connections
8	Array management connections
9	Data fabric 1 connections
10	Data fabric 2 connections

IP Addresses

An IP address is a 32-bit identifier for devices on a TCP/IP network. IP addresses allow devices on a network, such as servers, switches, and arrays, to communicate with each other. Arrays use IP addresses for the following purposes:

Table 1: Types of IP addresses

IP Address	Purpose
Management	Typically defined on eth1 or on eth1 and eth2 interface, the management IP address provides access to the array OS management interface (GUI, CLI, or API) for the array group. It is also used for volume replication. It resides on the group management subnet and floats across all management only (Mgmt only) and management + data (Mgmt + Data) interfaces.
Secondary	This is a secondary management IP address that is associated with the backup group leader array. In the event of a group leader migration or manual takeover, you can use this IP address to enable the backup group leader to take over the group leader functions. This IP address resides on the group management subnet and floats across all management only (Mgmt only) and management + data (Mgmt + Data) interfaces. While setting up a secondary management IP address is a best practice, it is optional.
Discovery	For iSCSI arrays, each subnet has its own discovery IP address. It enables the iSCSI initiator to discover iSCSI targets for the volumes on the array. You can use this IP address for data as well as management in a single shared network. **NOTE:** Discovery IP addresses are not required for Fibre Channel arrays.
Data	One or more IP addresses can be configured to carry data traffic. One data IP address can be configured for each interface pair (corresponding interfaces on the two controllers). Both controllers use the same IP address but never at the same time because only one controller is active at a time. Other data IP addresses can be configured on different subnets. **NOTE:** In a dedicated network topology, the data IP addresses cannot be the same as the management/iSCSI discovery IP addresses.
Support	Each controller on an array must have a dedicated support IP address, which can be used for troubleshooting and technical support purposes in the event that a controller is not reachable through the management IP address. The support IP addresses must be placed on the group management subnet.

Firewall Rules

If your organization restricts outbound internet traffic you might need to configure rules which allow the Storage SAN (Management IP address and Controller A and B diagnostic IP addresses) to communicate with storage servers on specific ports.

The storage SAN requires the following firewall rules:

Description	Destination	Port	Protocol
AutoSupport and Heartbeat	nsdiag.nimblestorage.com	TCP 443	HTTPS
Software Updates and Downloads	update.nimblestorage.com	TCP 443	HTTPS
Alerts	nsalerts.nimblestorage.com	TCP 443	HTTPS

Table Continued

Description	Destination	Port	Protocol
Streaming Statistics & VMVision Per-VM Alerts	nsstats.nimblestorage.com	TCP 443	HTTPS
HPE Support Secure Tunnel Connection	hogan.nimblestorage.com	TCP 2222	SSH
HPE GreenLake	console.greenlake.hpe.com	TCP 443	HTTPS
Array Initialization	device.cloud.hpe.com	TCP 443	HTTPS
Array Initialization	device.cloud.hpe.com	TCP 443	HTTPS
Data Services Cloud Console	console-**instance name**.data.cloud.hpe.com For example: console-eu1.data.cloud.hpe.com	TCP 443	HTTPS
	tunnel-**instance name**.data.cloud.hpe.com For example: tunnel-eu1.data.cloud.hpe.com	TCP 443	HTTPS
	instance name.data.cloud.hpe.com For example: eu1.data.cloud.hpe.com	TCP 443	HTTPS

NOTE: **instance name** can be **uk1** for the United Kingdom, **eu1** for Europe, **jp1** for Japan, or **us1** for America

NOTE: It is recommended that you use the fully qualified domain name (for example, nsdiag.nimblestorage.com) rather than IP addresses, as IP addresses might change without notice.

Subnets

A subnet is a logical subdivision of a network. It is defined by the first IP address in the network and a netmask that specifies a contiguous range of IP addresses within that network. A subnet can be assigned to one or more network interfaces.

The maximum Transmission Unit (MTU) can be set for a subnet so that it uses either a standard, jumbo, or custom frame size. If you choose to use a custom frame size, you must specify the size in bytes.

Specifying a VLAN ID on a subnet allows more than one subnet to be assigned to an interface using IEEE 802.1Q tagged frames. Subnets without a VLAN ID can only have untagged assignments. Switch port configurations must match the VLAN IDs configured on the subnets for tagged assignments. For more information, refer to the procedure to Configure VLAN Tagging in the *GUI Administration Guide* or *CLI Administration Guide*.

Subnet Traffic Types

Traffic types are used to segregate network traffic into different subnets. A subnet can carry one of the following traffic types.

Table 2: Traffic Types

Traffic Type	Description
Management (Mgmt only)	The subnet carries only management traffic.
Data (Data only)	The subnet carries only data traffic.
Management and Data (Mgmt + Data)	The subnet carries both management and data traffic.

Subnet Traffic Assignments

Traffic assignments determine what type of iSCSI traffic a data subnet carries. You can assign a data subnet on an iSCSI array to carry one of the following kinds of traffic.

NOTE: Traffic assignments are not required for Fibre Channel arrays.

Table 3: Traffic Assignments

Traffic Assignment	Description
iSCSI + Group	The subnet carries both iSCSI data traffic and intra-group communication (traffic between arrays in a group).
iSCSI only	The subnet carries only iSCSI data traffic.
Group only	The subnet carries intra-group communication traffic.

IP Address Zones in Subnets

An IP address zone is a group of host IP addresses and array data IP addresses in a subnet. When using two switches for iSCSI traffic, hosts can achieve better performance by establishing iSCSI connections with data IP addresses inside the same zone, as opposed to establishing iSCSI connections with data IP addresses in a different zone.

NOTE: IP address zones are not required for Fibre Channel arrays.

The IP addresses within a subnet can be divided into IP address zone types.

Table 4: IP Address Zone Types

Zone Type	Description
None	Used for non-iSCSI enabled subnets.
Single	All IP addresses are in one zone. This is the default zoning setting.
	With two network switches, iSCSI connections can be routed over the inter-switch link.
Bisect	One zone includes the IP addresses from the top half of the subnet; for example, 192.168.1.128 to 192.168.1.254. The other zone takes the IP addresses from the bottom half of the subnet; for example, 192.168.1.1 to 192.168.1.127
Even/Odd	The IP addresses are grouped by their last bit. One zone includes the even-numbered IP addresses, such as 192.168.1.2, 192.168.1.4, 192.168.1.6, and so on. The other zone includes the odd-numbered IP addresses, such as 192.168.1.1, 192.168.1.3, 192.168.1.5, and so on.

IP address zones are useful for configurations that use two switches, where you want to establish connections that avoid the Inter-Switch Link. For IP address zones to work, the host and the array must have their data IP addresses configured with half of their IP addresses from one zone connected to one switch and the other half of its IP addresses from the other zone connected to the other switch. For example, assume that:

- There is a single subnet, 192.168.1.0/24.

- There are two zones, defined as Red and Blue.

- Red zone consists of:
    - Host IP 192.168.1.1

    - Array Data IP 192.168.1.3

    - Array Data IP 192.168.1.5

- Blue zone consists of:
    - Host IP 192.168.1.2

    - Array Data IP 192.168.1.4

    - Array Data IP 192.168.1.6

In the IP Address Zone, the host IP addresses in the Red zone only establish connections with the data IP addresses in the Red zone. And the host IP addresses in the Blue zone only establish connections with the data IP addresses in the Blue zone. In this way, iSCSI connections do not use inter-switch link and thereby maximize I/O performance.
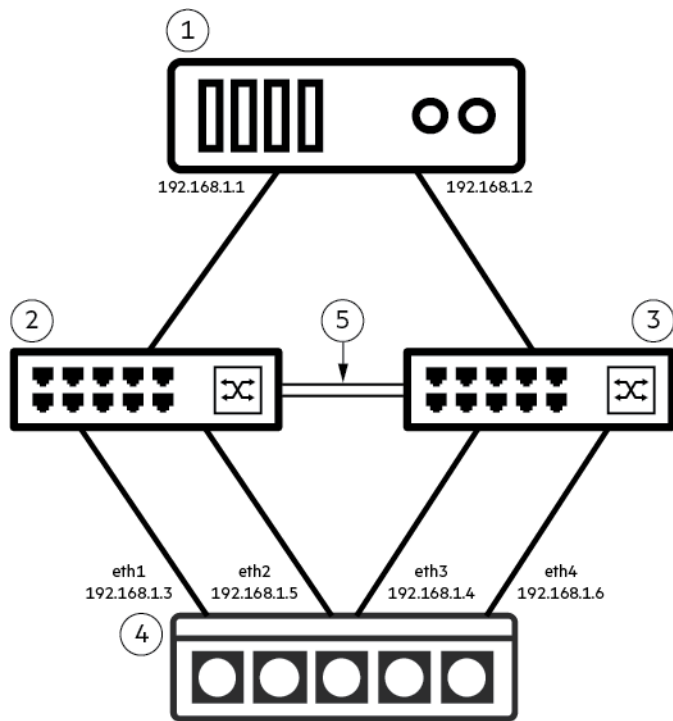
Figure 6: IP Address Zones

1	Host
2	Switch 1
3	Switch 2
4	Array
5	Inter-switch link

Interface Pairs

HPE uses the term interface to identify specific ports on the controller. For iSCSI, an interface is an Ethernet port connection that can be used for data, management, or both; for Fibre Channel, an interface is a Fibre Channel port connection that can be used for data only.

The purpose of interface pairs is to ensure high availability and redundancy in the event of a controller or network switch failure. In the event of a controller failover, proper network connections ensure minimal data service interruption.

To minimize single points of failure, connect as many interface port pairs as possible.

iSCSI Interface Pairs

An iSCSI interface pair refers to a set of two data or management Ethernet connections, one from each controller, to the network switch. The cables for an interface pair connect to the same data or management port on each controller. See **Interface Numbering**.
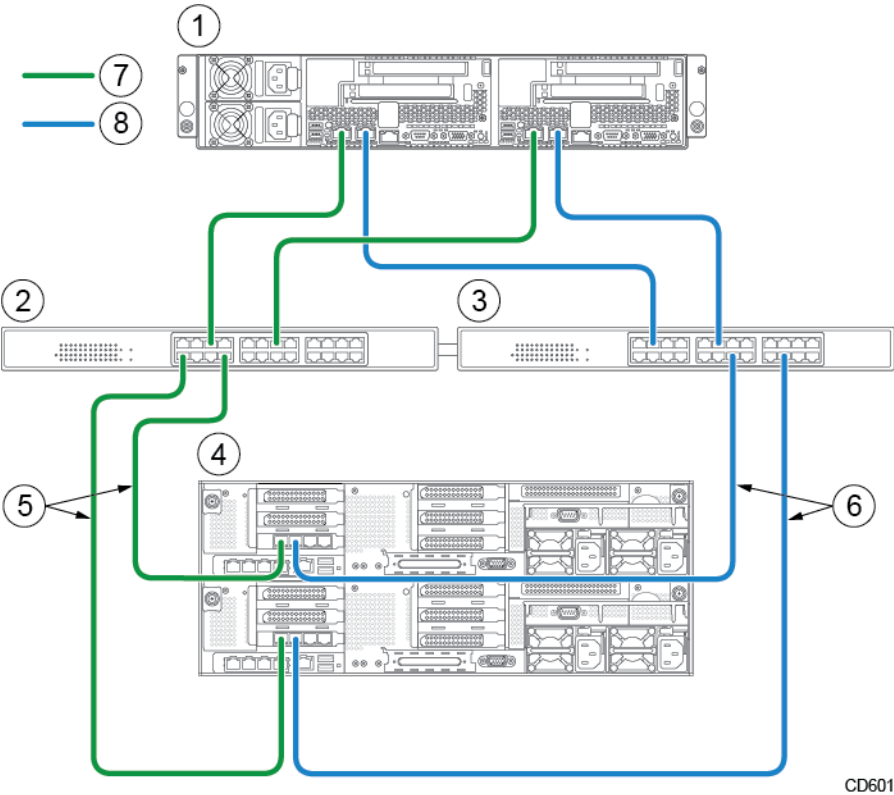
Figure 7: An example of Ethernet connections using interface pairs

1	Host/server
2	Switch 1/subnet 1
3	Switch 2/subnet 2
4	Array
5	Interface pair from data ports eth3a
6	Interface pair from data ports eth3b
7	Data interface pair 1
8	Data interface pair 2

An interface pair is the combination of the same data port on both controllers. In the preceding diagram:

- Data ports eth3a form one interface pair (green)

- Data ports eth3b form another interface pair (blue)

Make sure that each interface pair connects to the SAME network or subnet. In this example, the red interface pair connects to subnet 1 and the blue connects to subnet 2.

Fibre Channel Interface Pairs

A Fibre Channel interface pair refers to a set of two Fibre Channel data connections, one from each controller, to the Fibre Channel switch. The cables for an interface pair connect to the same data port on each controller. See **Interface Numbering**.
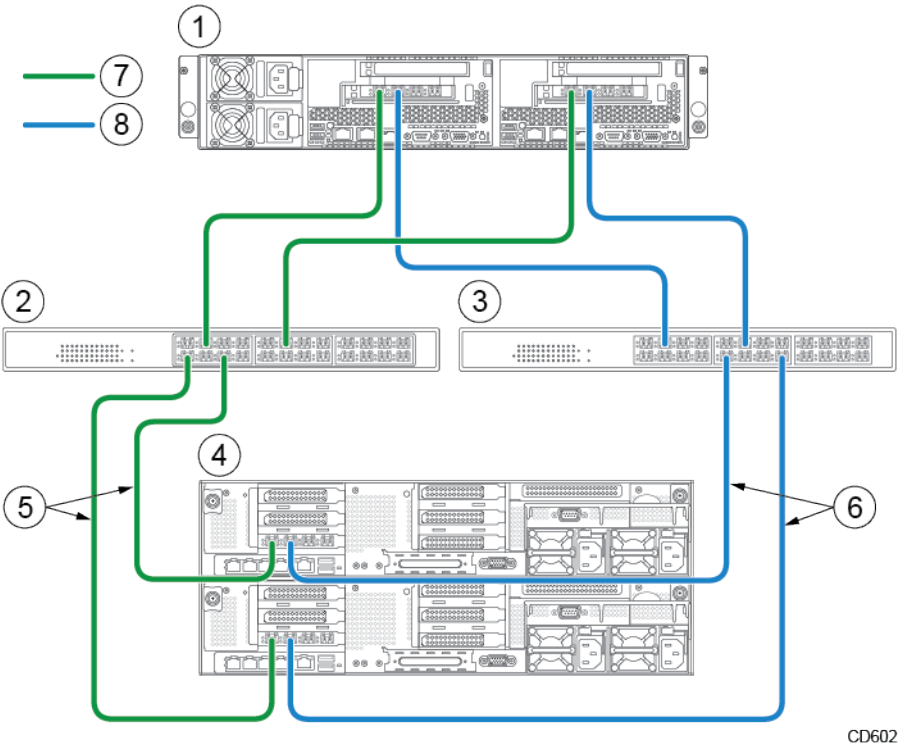
Figure 8: An example of Fibre Channel connections using interface pair

1	Host/server
2	Fibre Channel switch 1
3	Fibre Channel switch 2
4	Array
5	Interface pair from data ports fc3a
6	Interface pair from data ports fc3b
7	Data interface pair 1
8	Data interface pair 2

An interface pair is the combination of the same data port on both controllers. In the preceding diagram:

- Data ports fc3a form one interface pair (green)

- Data ports fc3b form another interface pair (blue)

Interface Numbering

All array models have four onboard 1GBaseT ports. Depending on the array model, adding additional NICs can support up to 12 10GBaseT, SFP, or GbE ports, for a total of up to 14 ports.

Figure 9: Array interface examples (Ethernet and Fibre Channel Ports)
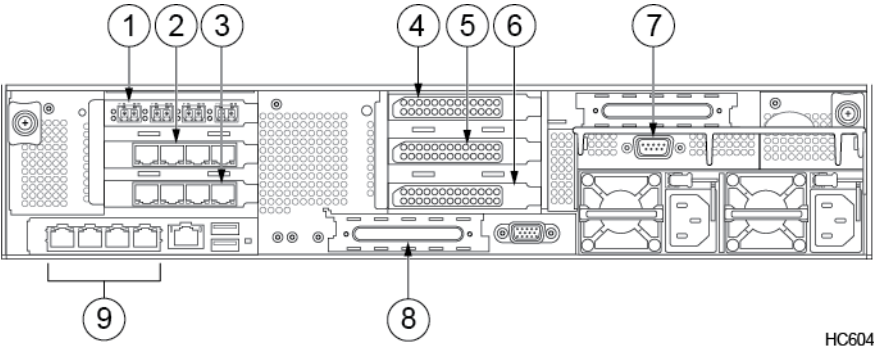
1	Slot 1
2	Slot 2
3	Slot 3
4	Slot 4
5	Slot 5
6	Slot 6
7	Serial port
8	OCP expansion slot
9	Eth0 ports, onboard
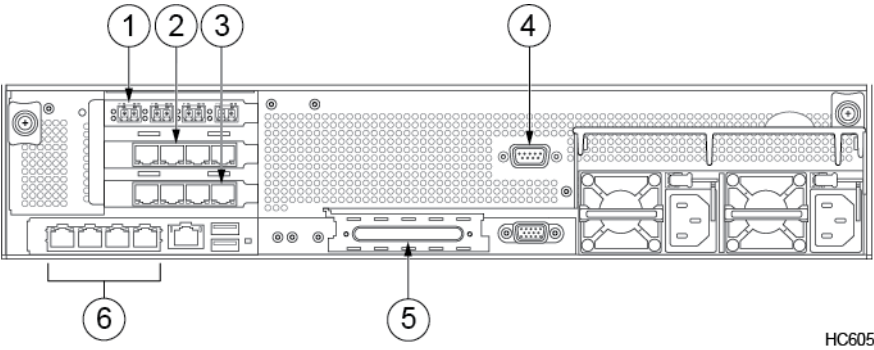
1	Slot 1
2	Slot 2
3	Slot 3
4	Serial port
5	OCP expansion slot
6	Eth0 ports, onboard

For more information on SFPs and connecting 10GbE ports to your array, see .

Installation Checklist

NOTE: During the intial setup process, you will be prompted for the following information. HPE strongly recommends that you gather this information prior to starting the setup process.

All network settings represent IPv4 values.

The latest version of this list can be found at http://infosight.hpe.com/welcomecenter.

Item	Description	Value
Domain		
DNS server	IP address of the DNS server responsible for the domain. At least one server is required. Up to three servers are supported. All configured DNS servers should be reachable, able to resolve the same records, and able to resolve HPE addresses (for example: device.cloud.hpe.com). Network connectivity is required from the Data Orchestrator to Data Services Cloud Console. Network time protocol servers (NTP) must be configured on the Data Orchestrator and must be allowed to be in sync with global time services. DNS name resolution must be configured on the Data Orchestrator and be allowed to resolve global Internet names and domains. The Data Orchestrator will periodically transmit health and statistical telemetry over HTTPS/TLS."	
Domain name	Name of the domain the array will be joining.	
Time		
Time (NTP) server	IP address of the NTP server.	
Region	Geographical region where the array is located.	
Time zone	Time zone in which the array is located.	
Attributes		
HTTP Proxy server		
HTTP Proxy Server Port		
Proxy Server Username		
Proxy Server Password		
Support contact	The person that HPE should contact if the array needs attention.	
HPE InfoSight	Indicate whether to claim the array in HPE InfoSight for predictive support.	

Table Continued

Item	Description	Value
System		
System Name	The name that you want to assign.	
System Country Location	The geographical location where the array is installed (required for setup).	
System IP		
	IPv4: IP address	
	IPv4: Subnet Mask	
	IPv4: Gateway IP address	
	IMPORTANT: All network components must be reachable by the array.	
System Credentials	The administrator name (user account) and password that you want to create.	
Management Networking		
Subnet Name	Name of subnet carrying management traffic.	
Management IP Address	IP address used to access the array.	
Management Netmask	Netmask for the management subnet.	
Default Gateway	IP address of the router or gateway the array will use to access network resources outside of the configured subnet.	
Secondary IP Address		
MTU Size	Standard, Jumbo, or Custom.	
MTU Bytes	When MTU size is set to Custom, this specifies the size in bytes.	
Domain Name	Name of the domain the array will be joining.	
Assigned Subnet Port		
Additional Subnet Port		
Diagnostics A IP Address	Support IP address for controller A.	
Diagnostics B IP Address	Support IP address for controller B.	
Data Networking		
Subnet Name	Name of subnet carrying data traffic.	
Port	Port used for data traffic.	
IP Address	Data port IP address.	
Netmask	Netmask for the data subnet.	
MTU Size	Standard, Jumbo, Custom	
MTU	MTU size in bytes.	
Discovery IP Address	IP address used for iSCSI discovery.	

Rack Mount Safety Precautions

Observe these safety precautions when you install the storage array or expansion shelf in a rack.

Table 5: Rack Mount Safety Precautions

Consideration	Precaution
Operating Ambient Temperature	If the array or expansion shelf is installed in a closed or multi-unit rack assembly, the operating temperature might be greater than the ambient (room) temperature. Provide adequate cooling so the rack environment does not exceed 35°C (95°F) as specified by the rack manufacturer.
Air Flow	Install the equipment in the rack such that each item has ample air flow required for proper cooling. Cold air comes into the front of the array (drive side) and is exhausted out the back of the array (controller side) by the fans. Install the array with the front facing the cold aisle and the back facing the hot aisle in your environment.
Mechanical Loading	Mount the equipment so that the mechanical loads on the rack are even top-to-bottom, front-to-back, and side-to-side.
Circuit Overloading	Be sure there is adequate power service to prevent overloading the circuits. Check the ratings of each item of equipment on each circuit. **NOTE:** The power supplies are 220VAC or 110VAC auto-detecting.
Reliable Grounding	You must provide reliable grounding (earthing) of rack-mounted equipment. In addition, pay attention to supply connections from the branch circuit, such as power strips.

Install the Array

About this task

Use this procedure to install the following array models: HPE Alletra 6010, 6030, 6050, 6070, and 6090.

Before you begin

ⓘ **IMPORTANT:** If you are installing more than one storage array, refer to the *Hardware Guide* available on HPE InfoSight for information about your model.

https://infosight.hpe.com/org/c2165d17-4d55-4954-83e3-60556a7bb6c8/resources/nimble/docs

What you need for each array:

- 19-inch four-post rack. The provided rails fit both square-hole and unthreaded round-hole racks.

- 4U space available in the 19-inch rack

- At least two Ethernet cables for management, one cable per controller

- At least two cables, Ethernet or Fibre Channel, for data

- One or two switches, depending on your array and network topology. Two 10GbE or Fibre Channel switches are recommended for high availability and redundancy.

- Four power outlets, two per controller, on separate circuits

  Current for 800 W power supplies with C14 connectors:
  - Steady state current (approximate): 5.6A @ 240V; 11.7A @120V

  - Startup current (approximate): 4.6A @ 240V; 10.3A @ 120V

  Current for 1600 W power supplies with C14 connectors (240V only):
  - Steady state current (approximate): 8.9A @ 240V

  - Startup current (approximate): 8.2A @ 240V

Before beginning, download and read the Release Notes, available on HPE InfoSight (**https://infosight.hpe.com**). If you do not have an HPE account, click **Sign up** and provide the required information.

Visit the Welcome Center (**http://www.hpe.com/support/storage-welcomecenter**) for help setting up your HPE storage array. Before you can start using the array, you must register the array with Data Services Cloud Console. If you do not have an HPE GreenLake cloud account (**https://common.cloud.hpe.com**), the Welcome Center can guide you.

NOTE:

- Install this product in restricted-access locations only, such as a dedicated equipment room or an electrical closet. Installation and maintenance must be carried out by qualified personnel.

- Position the rack in a location with at least 800 mm rear clearance and 1000 mm front clearance for maintenance access.

- Install the array with the front (drive side) facing the cold aisle and the back (controller side) facing the hot aisle in your environment.
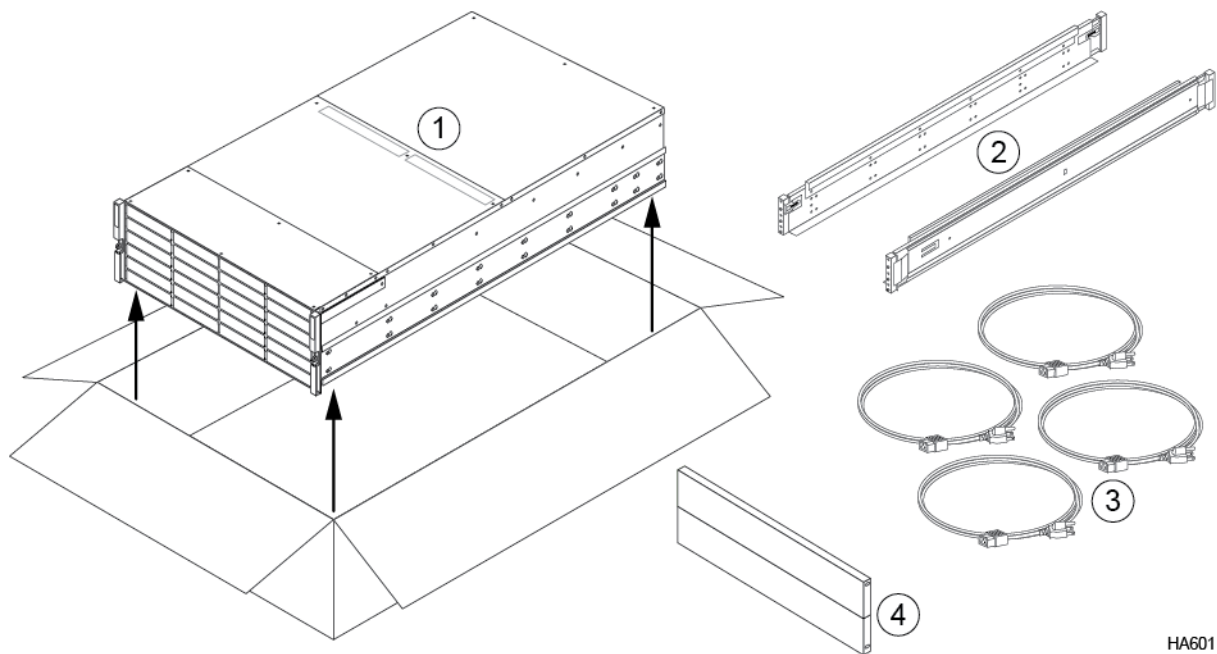
⚠ **CAUTION:** The chassis weighs up to 54.5 kg (120 lb). Always use at least two people or a server lift when lifting the chassis

Prerequisites

Procedure

1. Unpack and inspect the hardware.

   Save the original packaging.

HA601

| 1	Array
2	Outer rails
3	Power cables
4	Bezel

2. Install the outer rail assemblies onto the rack.

   NOTE: To ensure that the array thumbscrews align correctly, align the rail assembly in exactly 4U of rack space. The bottom stud on the rail assembly inserts into the bottom hole of the bottom unit of the 4U rack space allotted for the array.

   a. Align the bottom stud of the rail at the position you want.
      Note the orientations embossed on the outer rail.

   b. Press the keeper latch and position the front studs in the rack. When the studs are engaged, release the keeper latch to secure the front of the rail in the rack.

   c. Align the back of the rail at the same rack position as the front.

   d. Position the back studs in the rack and press the keeper latch. When the studs are engaged, release the keeper latch to secure the back of the rail in the rack.

   e. Repeat steps a through d on the other side.

   The following diagram shows where the rail assemblies should be correctly installed in a rack.
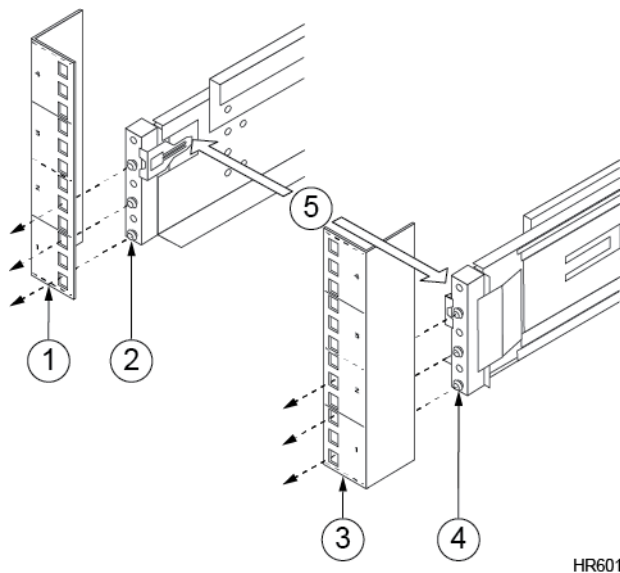
HR601

1	Left front rack post
2	Left rail assembly
3	Right front rack post
4	Right rail assembly
5	Keeper latch

3.

⚠ **CAUTION:** The chassis weighs up to 54.5 kg (120 lb). Always use at least two people or a server lift when lifting the chassis

Slide the chassis into the rack.

a. Insert the chassis into the rack, back side first. Be sure the inner rails insert into the outer rails.

b. Gently slide the chassis into position.

When you hear a click, the inner rails have locked into the rail assembly.
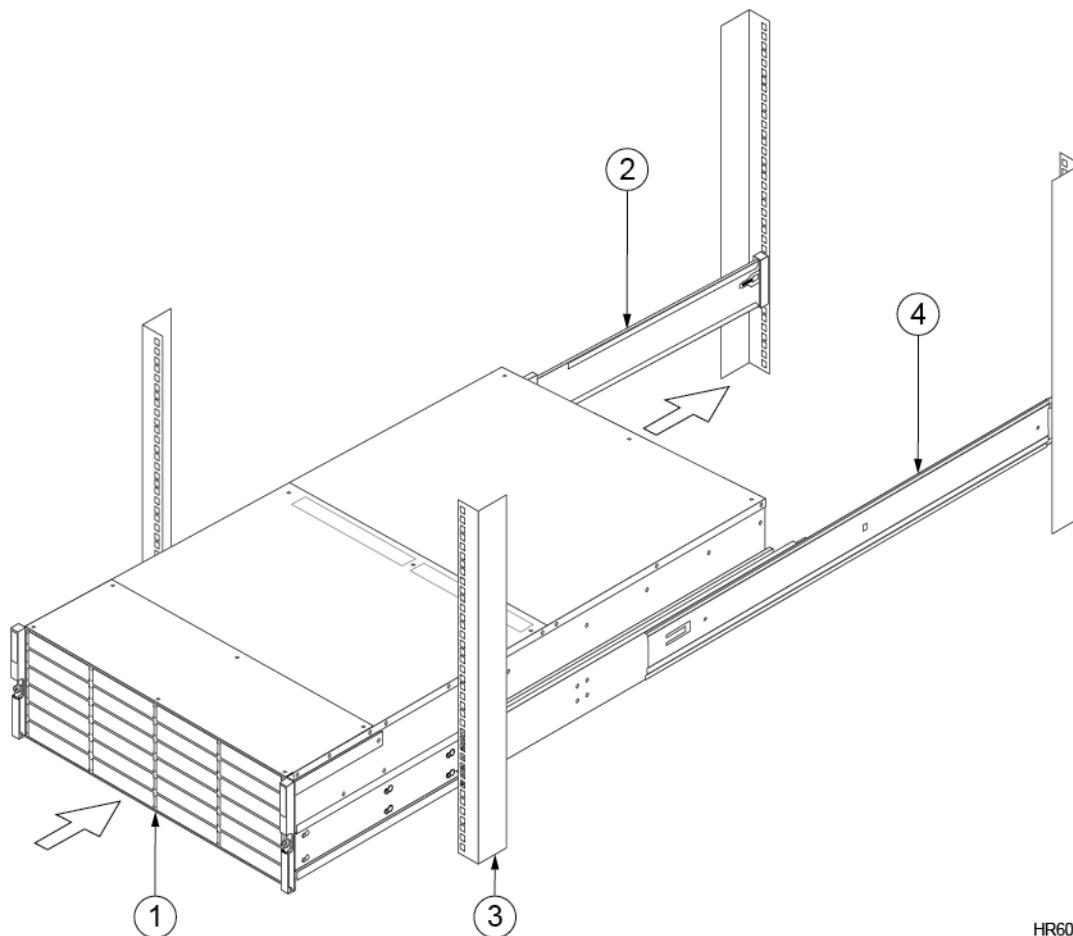
HR602

4. Secure the chassis to the rack.

   The chassis has two handles. Each handle has a thumbscrew that holds the chassis to the rails so it does not slide out of the rack. Tighten the thumbscrew in each handle until secure.
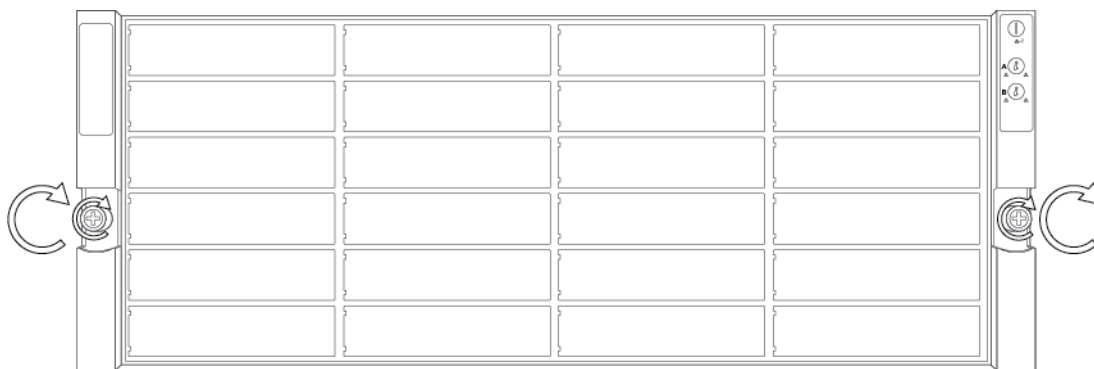
HA607

5. Secure the bezel to the front of the chassis.

   a. Align the four retention pins on the bezel with the four mounting holes on the front of the chassis.
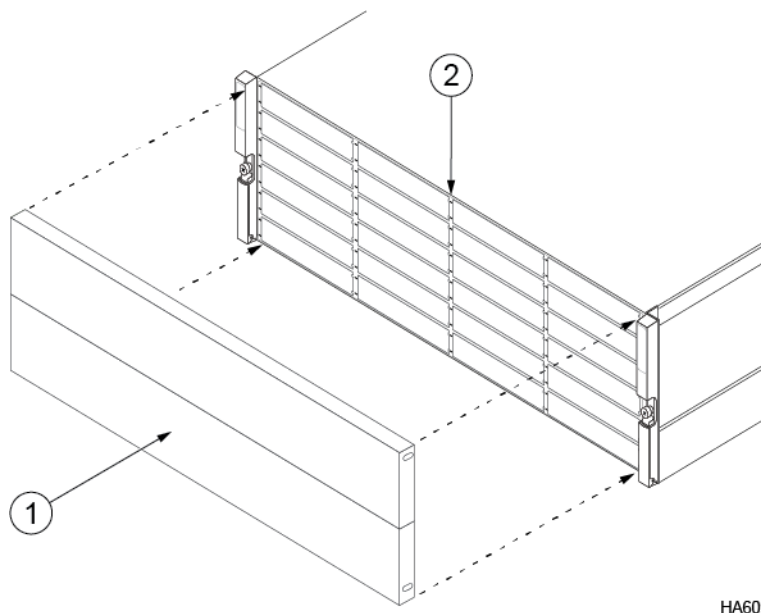
   b. Press gently but firmly until the bezel snaps into place.

HA609

1	Bezel
2	Chassis

6. Cable the array for your desired network topology.

   a. Connect ports for management (commonly the onboard eth0 ports).

   b. Connect ports for data according to your network topology and protocol (eth or fc).

   - For Ethernet, cable the same port on each controller to the same network switch and subnet. The two ports form an interface pair, and iSCSI initiators have a network path to both controllers. If your switches have multiple VLANs configured, make sure that each interface pair is connected to switch ports on the same VLAN.

   - For Fibre Channel, cable the same port on each controller to the same FC switch.

   See the *Hardware Guide* on HPE InfoSight for detailed network configuration examples.

In this example, quad fc ports are installed in Slot 1 and quad eth ports are installed in Slot 2 and Slot 3.

Figure 10: HPE Alletra 6030, 6050, 6070, and 6090
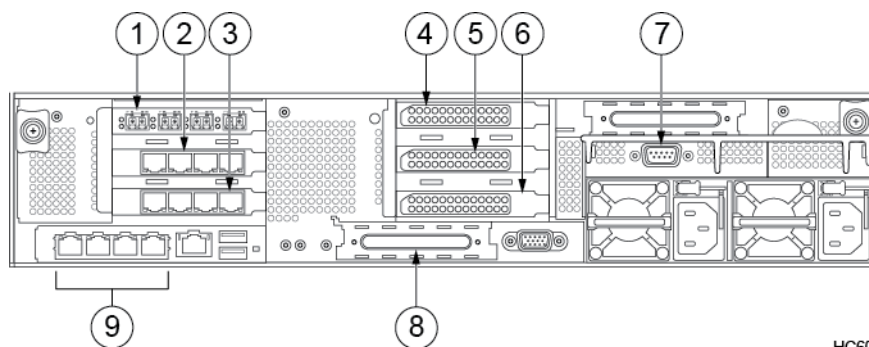
HC604

1	Slot 1
2	Slot 2
3	Slot 3
4	Slot 4
5	Slot 5
6	Slot 6
7	Serial port
8	OCP expansion slot
9	Eth0 ports, onboard

In this example, quad fc ports are installed in Slot 1 and quad eth ports are installed in Slot 2 and Slot 3.

Figure 11: HPE Alletra 6010
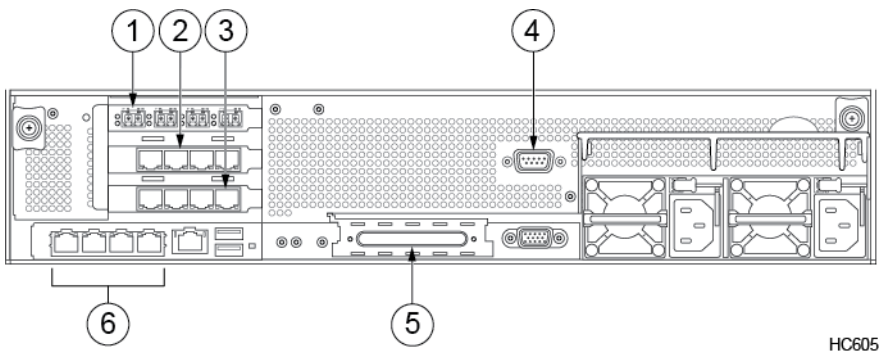
1	Slot 1
2	Slot 2
3	Slot 3
4	Serial port
5	OCP expansion slot
6	Eth0 ports, onboard

Network Connections

Depending on your desired **Network Topology**, the following network configuration examples show how to make the physical network connections to support that topology.

NOTE: Later, be sure to make the appropriate port settings that correspond to these network connections.

Connect the Power to the Array

Prerequisites

The storage array must be installed in a rack or enclosure and the array must be connected to a network.

Procedure

1. Plug power cables into the power supplies.

2. Verify that the connection is secure.

3. Plug the other end of the power cables into the power circuit.

   To minimize the risk of both power supplies for a controller losing power simultaneously, connect each power cable to a different circuit.

   If the array does not come online automatically, press and release the power button ⏻ on the front of the chassis.

Initial Setup Using the HPE Alletra Setup Service

After your array has been installed in the rack and connected to your network, you must complete the initial setup of the array by configuring some basic settings to get the array running on the network and connected to HPE GreenLake.

This chapter describes how to complete the process from a Windows host. If you do not have a Windows host from which to complete the initial setup, you must use the CLI. For more information, see **Initial Setup Using the CLI**.

Download the HPE Storage Toolkit for Windows

About this task

Download the latest version of the HPE Storage Toolkit for Windows to install on your Windows host or Windows VM.

> (!) **IMPORTANT:** Make sure that the version of the toolkit is compatible with the version of the operating system on your array. For compatibility information, see the Validated Configuration Matrix tool, which is available on HPE InfoSight (**https://infosight.hpe.com/resources/nimble/validated-configuration-matrix**).

Prerequisites

What you need on a Windows server:

- Internet connection to HPE InfoSight

- .NET Framework 4.5.2

What you need on a Windows client:

- Internet connection to HPE InfoSight

- Windows OS 7 or later

- .NET Framework 4.5.2

Procedure

1. In your Internet browser, go to **http://www.hpe.com/support/storage-welcomecenter** **https://infosight.hpe.com/welcomecenter**.

2. On the Welcome Center home screen, click the array family.

3. Select the array model and click **Continue**.

4. In the navigation panel, click **Software Configuration** > **Prepare Client**.

5. Follow the instructions to download the toolkit.

6. Save the installation package to your Windows host.

   The installation package has a name similar to `Setup-HPEStorage-x64.x.x.x.x.exe`, where `x64` is the supported microprocessor and `x.x.x.x` is the version number.

7. Download the latest *Toolkit Release Notes*.

   Review the list of Windows Server hotfixes.

> **NOTE:** If you are installing the setup manager alone, then no hotfixes are needed. However, .NET framework requirements still apply. Hotfix requirements are mandatory for any Windows host on which the HPE Storage Connection Manager for Windows is used to connect to array volumes.

Install the HPE Storage Toolkit for Windows on a Client

About this task

You have the option of installing the HPE Storage Toolkit for Windows on a Windows client.

Prerequisites

Download the HPE Storage Toolkit for Windows from HPE InfoSight (**https://infosight.hpe.com**).

See **Download the HPE Storage Toolkit for Windows**.

The client must be running Windows OS 7 or later to install the toolkit.

> (!) **IMPORTANT:** If you plan to install the toolkit in a Remote Desktop session or on a virtual machine, open the Command Prompt and enter:
>
> ```
> mstsc /console
> ```

Procedure

1. Open the `Setup-HPEStorage-x64.x.x.x.x.zip` file and extract the `Setup-HPEStorage-x64.x.x.x.x.exe` installer file.

2. On the Windows client, right-click the installer file and choose **Run as administrator**.

3. In the Welcome to the InstallShield Wizard for the toolkit dialog, click **Next**.

4. In the License Agreement dialog, choose **I accept the terms of the license agreement** and click **Next**.

   > (!) **IMPORTANT:** You must accept the End User License Agreement (EULA) to install the toolkit.

5. In the Logs Directory dialog, do one of the following:

   - Click **Next** to accept the default log location.

   - Click **Change**, browse to the desired directory, click **OK**, and click **Next**.

   In the Custom Setup dialog, **Setup Manager** is selected by default. The setup manager enables first-time setup of the storage array.

6. Specify the Group Management IP address and the username/password credentials for the storage array. The array user should have Power User or above credentials.

   If you do not do this when you set up the toolkit, you can use the PowerShell cmdlet `Set-NWTConfiguration` to specify the array credentials.

   You must register at least one group. The host must be registered with the array group to enable application-synchronized snapshot creation.

7. Click **Install**.

   You can modify this option after clicking **OK** by clicking **Back** to return to the Custom Setup dialog.

If you deselected the Setup Manager, a message is displayed informing you that you need to explicitly select it if iSCSI management is required.

8. Click **OK** to acknowledge the message.

9. In the InstallShield completed dialog, click **Finish**.
   The toolkit software installs onto the Windows host. This process might take a few minutes. The installer log is saved to the `C:\NWT_Install.log` file.

Discover the Array

About this task

The HPE Storage Setup Manager for Windows enables you to create a new group from an uninitialized array. The array is the only member of the group.

NOTE: A newly plugged-in array might not appear immediately on the array details page of the GUI. It might take up to five (5) minutes to update the page. Refresh the browser to see the array on the page immediately after you plug it in.

Prerequisites

The Windows host that you use to initially configure the array must be connected to the same physical subnet as the array management port and have direct (non-routed) access to the array management port.

The Windows host must have the following software installed:

- .NET Framework 4.5.2 or above

- HPE Storage Toolkit for Windows

Before you start, complete the **Installation Checklist**. You need the information from the checklist to perform the initial setup.

⚠ **CAUTION:** If you have multiple unconfigured arrays and you plan to add them to the same group, configure only one array now. Later, add the remaining arrays to the group. For information about how to add an array to a group, refer to the *GUI Administration Guide* or the *CLI Administration Guide*.

NOTE: Subnet and VLAN configuration are separate actions. VLANs are not supported during setup.

Procedure

1. In the Windows Start menu, choose **HPE Storage** > **Setup Manager**.

2. Select one of the uninitialized arrays from the setup manager list and click **Next**.

   NOTE: If the array is not visible in Setup Manager, verify that the array is on the same subnet as the Windows host.

   ---

3. In the Information dialog, click **OK**.

   The HPE Cloud Connectivity Wizard opens to guide you through the process of connecting the storage array to HPE GreenLake.

Connect the Array to HPE GreenLake

About this task

The HPE Cloud Connectivity Wizard guides your through the process of connecting the storage array to HPE GreenLake.

Procedure

1. Click **Get Started**.

2. Read and accept the end user license agree, then click **Continue**.

3. Specify the following network settings, then click **Continue**.

   - IP address

   - Netmask

   - Gateway

   - Domain Name

   - DNS Server IP

4. If using a proxy server to connect to the internet, select **Use proxy** and specify the following, then click **Continue**:

   - HTTP Proxy Server

   - HTTP Proxy Server Port

   - Proxy Server Username

   - Proxy Server Password

5. Specify the following time information, then click **Continue**.

   - Time (NTP) Server

   - Region

   - Country/State/City

   ⓘ **IMPORTANT:** Maintaining the system time within 2 minutes of the correct time is required to initialize and manage an HPE Alletra storage array with Data Services Cloud Console. To maintain the array time within the tolerance, synchronize the array time with an accurate, highly available NTP service that is synchronized with global time services.

   If the array time on an uninitialized array is not within the tolerance, the lack of a connection to Data Services Cloud Console will prevent the array from being initialized. Setting the array time within the tolerance will enable the connection to be established and allow you to begin the initialization.

   If the array time on an initialized array is not within the tolerance, the array will not connect to Data Services Cloud Console. Or, if the array is already connected, the connection will be lost. A lost connection can happen if the array time is set incorrectly, or if the array time drifts beyond the tolerance. Resetting the array time within the tolerance will enable the connection to be reestablished

6. Review the settings, then click **Submit**. The wizard makes the initial connection to HPE GreenLake and returns the subscription key.

7. When the wizard indicates that you have successfully connected the storage array to HPE GreenLake, click **Launch** to open HPE GreenLake.

What to do next

Initialize the storage array.

Initialize the Array Using HPE Alletra Setup Service

About this task

After you have connected the array to Data Services Cloud Console, initialize the array software using the HPE Alletra Setup Service.

Prerequisites

If you completed the installation checklist, refer to it while completing this procedure.

Procedure

1. If not already, log in to **HPE GreenLake**.

2. In the menu at the top of the screen, click **Applications**.

3. On the **My Applications** page, click **Data Services Cloud Console**.

4. On the **Deployment Regions** page, click **Launch** next to the Data Service Cloud Console instance.

5. On the **Welcome** page, click **Launch** in the **Setup Service** tile.

6. Set up the array using Setup Service. For information on using Setup Service, search for the service in the Data Services Cloud Console articles widget.

What to do next

You can now use Data Storage Cloud Console - Block Storage to provision storage and manage array settings.

Initial Setup Using the CLI

After your array has been installed in the rack and connected to your network, you must complete the initial setup of the array by configuring some basic settings to get the array running on the network.

This section describes using the CLI to complete the initial setup. To use the CLI, you must first establish a connection to the active controller on the array.

NOTE: If you have a Windows host, and you cannot establish a serial connection to the active controller on the array, you must use the HPE Storage Setup Manager for Windows to complete the initial setup.

After the initial setup is complete, refer to the *GUI Administration Guide* or the *CLI Administration Guide* for group and volume configuration, plus additional optional configuration tasks.

Set up a Serial Connection

About this task

The CLI requires a serial connection for the initial setup of the array. After setup, the CLI can also use an SSH connection over the network.

Prerequisites

Obtain these items:

- Laptop computer or workstation

- Serial console cables with a DB9 connector at one end and the appropriate connector on the other to mate with the serial interface on the laptop computer or workstation

Procedure

1. Connect the serial port on your laptop computer or workstation to the serial port on the controller.

HC606
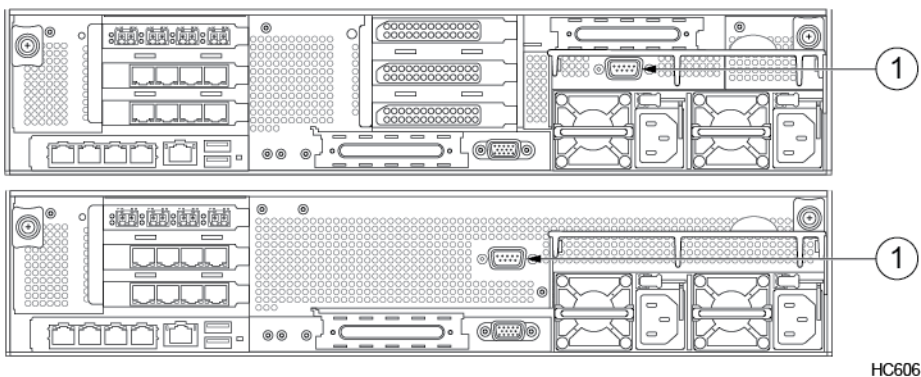
1	Serial port

2. Run the serial console software with the following settings:

Data bits:	8
Parity:	None

Table Continued

Flow controller:	None
Stop bits:	1
Speed:	115.2Kbps

3. Log in as user `admin` with the password `admin`.

> **NOTE:** You can only log into the active controller. If you cannot log in, move the laptop computer or workstation connection to the other controller and try to log in again.

After you have logged in, the CLI is ready to accept commands.

What to do next

Type `?` to see a list of commands.

Type the command followed by `--help` to see the usage information for that command.

Complete the Initial Setup Using the CLI

Prerequisites

Before you start, fill out the **Installation Checklist**. You need the information from the checklist to perform the initial CLI setup.

> **NOTE:** Subnet and VLAN configuration is a separate action.

Procedure

1. After you are logged into the active controller, type:
   `setup`

   The Terms and Conditions are displayed.

2. Click `Enter` to read the Terms and Conditions, then type `yes` to accept them.

3. Type the following information as requested. Example values are shown here:

   - Array name: `greyhound`

   - Group name: `storehouse`

   - Management IP address: `10.12.128.79`

   - Management netmask: `255.255.0.0`

   - Default gateway IP address: `10.12.255.254`

   - Domain name: `admin.storage.com`

   - New password:

4. Specify whether to include IPs in the group certificates: `yes` or `no`.

- Type `yes` to include the array's IP address in the certificate. When using this option, you will have to create a new certificate if the array's IP address changes.

- Type `no` to exclude the array's IP address from the certificate..

5. Specify whether you want to continue the setup in the GUI.

   a. Type `yes` to open the GUI. Then complete .

   b. Type `no` to continue with the CLI. Then complete the remaining steps in this procedure.

6. Type the following information as requested. Example values are shown here:

   - Comma-separated list of DNS server(s): `10.64.0.13,10`

   - NTP server name: `time.ntpserver.com`

   - Time zone: `America/Los_Angeles`

   - Subnet label for NIC eth0a: `management`

   - Subnet for NIC eth0a (use slash notation, for example, 10.11.0.0/16): `10.12.0.0/16`

   - Comma-separated list of subnet types for the subnet on NIC eth0a [management,data]: `management`

   - MTU for the subnet on NIC eth0a [standard | jumbo | <integer>]: `standard`

   - Subnet label for NIC eth0b: `management`

   - Subnet label for NIC tg1a: `Data1`

   - Subnet for NIC tg1a (use slash notation, for example, 10.11.0.0/16): `172.0.1.0/22`

   - Comma-separated list of subnet types for the subnet on NIC tg1a [management,data]: `data`

   - MTU for the subnet on NIC tg1a [standard|jumbo|<integer>]: `standard`

   - Subnet label for NIC tg1b: `Data2`

   - Subnet for NIC tg1b (use slash notation, for example, 10.11.0.0/16): `172.0.2.0/23`

   - Comma-separated list of subnet types for the subnet on NIC tg1b [management,data]: `data`

   - MTU for the subnet on NIC tg1b [standard|jumbo|<integer>]: `standard`

   - Subnet label for NIC tg1c: `Data1`

   - Subnet label for NIC tg1d: `Data2`

   - Data IP address for NIC tg1a: `172.0.1.83`

   - Data IP address for NIC tg1b: `172.0.2.83`

   - Data IP address for NIC tg1c: `172.0.1.84`

   - Data IP address for NIC tg1d: `172.0.2.84`

   - Discovery IP address for subnet Data1(172.0.1.0/255.255.255.0): `172.0.1.82`

   - Discovery IP address for subnet Data2(172.0.2.0/255.255.255.0): `172.0.2.82`

   - Support IP address for controller A: `10.207.9.83`

- Support IP address for controller B: `10.207.9.84`

- Secondary management IP address: `10.12.128.81`

7. Specify whether to allow HPE to use Google Analytics in the GUI. HPE uses Google Analytics to gather data related to GUI usage. The data gathered is used to evaluate and improve the product. Enter yes or no: `yes`

8. After the setup is complete, disconnect the Serial-USB-VGA adapter (dongle) and serial cable.

What to do next

Open an SSH session to the management IP address of the array and log in with the user name `admin` and the password `admin`.

If you need to change the admin password, type `useradmin --passwd`, press **Enter**, and follow the prompts.

Next, see **After Installation and Initial Setup**.

After Installation and Initial Setup

Software Updates, Host Connection, Groups, Volumes (iSCSI and FC)

Continue your configuration by following instructions in the *GUI Administration Guide* or *CLI Administration Guide* to create performance policies, set up connections, and complete setup in the HPE Storage Setup Manager for Windows. HPE arrays support both iSCSI and Fibre Channel protocols, and provides information for your specific host connection method to create initiator groups and volumes. Additional configuration options include creating and managing storage pools, volume collections, snapshots, replication strategies, security policies, and performance monitoring.

Learn about these post-installation configuration tasks in the *GUI Administration Guide* or the *CLI Administration Guide*.

Windows Environments

The HPE Storage Toolkit for Windows, mentioned briefly in this guide, includes features that go beyond setting up the array. The VSS requester and the VSS hardware provider enable you to take application-consistent snapshots on an array. The HPE Storage Connection Manager for Windows sets up the optimum number of data connections, and finds the best data connection to use under MPIO. The Windows Toolkit includes a DSM kernel driver that claims and aggregates data paths for the array.

See the *Windows Integration Guide* for more information.

VMware Environments

HPE storage integration with VMware consists of features preinstalled in the array OS and features installed separately. The vStorage APIs for Array Integration (VAAI) enable the WRITE SAME, UNMAP, ATS, Copy Offload, and XCOPY APIs. The HPE Storage vCenter plugin creates and manages datastores on the array. The OS also includes a VASA Provider for vVol support. VMware Synchronized Snapshots enable application consistent snapshots within VMware environments. SRA for SRM lets you set up disaster recovery plans. Connection manager installs on the ESXi host. Connection Manager consists of two components:

- HPE Storage Connection Service – CIM provider. Optimized connection management.

- HPE Storage Path Selection Plugin – Optimized path management.

See the *VMware Integration Guide* for more information.

Linux Environments

The array fits well into a generic Linux environment. For help setting up MPIO in a Linux environment, see the *Linux Integration Guide.*

Using the array OS

The array OS has a multitude of features only touched upon in this document. Over time, you might need to update your network settings, performance policies, or replication partners, as well as create new volumes and groups and make administrative changes to your array.

Anytime you have a question, in the GUI, click **Help** at the top right of the OS screen.

In the CLI type **?** to see a list of commands.

Also refer to the *GUI Administration Guide*, the *CLI Administration Guide*, or the *Command Reference*.

Updates and Upgrades

When the time comes to update the OS on your array, refer to the *GUI Administration Guide* or the *CLI Administration Guide* for information on how to update the OS using the GUI or CLI.

The *Command Reference* also covers updates.

Depending on which model array you have, there are upgrade paths for the controllers and cache drives. When you are ready to upgrade, contact support first.

See the quick start guide that ships with the upgrade component. The *Hardware Guide* also covers upgrades.

Maintenance and Repair

The array is a proven and reliable product. But on occasion, you may need to replace a controller, power supply, drive, or other component.

Refer to the *Hardware Guide* for more information.

How Do I Obtain these Guides?

The guides discussed here are available for free download from the HPE InfoSight portal at **https://infosight.hpe.com/**.

Questions?

Support offers troubleshooting and other guidance by phone and by email.

All documentation and knowledge base articles are available on the HPE InfoSight portal at **https://infosight.hpe.com/**. To register on HPE InfoSight, click the *Enroll Now* link on the main page.

Email: **support@nimblestorage.com**

For all other general support contact information, go to **https://hpe.com/us/en/services/nimble-storage.html**.