HPE GreenLake

Technical white paper

HPE Alletra 6000 Storage simplified setup with Data Services Cloud Console

Initializing the HPE Alletra 6000 storage array and onboarding to Data Services Cloud Console

Contents

Executive summary

Data Services Cloud Console (DSCC) provides a cloud operation and consumption experience for your storage and streamlines data management across your environment. It allows you to leverage unique features such as intent-based provisioning and manage heterogeneous systems with a single interface. And it is all managed globally through a cloud portal accessible from anywhere, on any device.

DSCC also allows you to deploy your new data infrastructure in minutes with automated discovery and activation. You cable up the system, power it on, enter a minimal amount of networking information, and then use the Setup Service within DSCC to do the rest. The DSCC Setup Service is a cloud-based application that uses wizards to help you set up and initialize newly installed HPE Alletra 6000 storage systems. The blueprint feature streamlines the initialization process with the use of custom, user-defined templates that can be applied to one or more systems. For example, if you have multiple systems that share the same settings, such as DNS and NTP servers or support contact, you can create a blueprint and then apply it to each new system you initialize. The Setup Wizard will be pre-populated with the appropriate data. The Setup Service is particularly useful is you are setting up multiple arrays because you can initialize all the systems simultaneously from a single console and begin managing those arrays as part of your storage fleet immediately after setup is complete – all from DSCC.

Document purpose

This technical paper is intended for anyone who is deploying a new HPE Alletra 6000 storage system in their data center. It walks through all the steps required to get an HPE Alletra 6000 up and running after the hardware has been installed, cabled, and powered on. Steps covered include onboarding the array into the HPE GreenLake Edge-to-Cloud platform, setting up the networking required to connect the system to DSCC via the HPE Cloud Connectivity Wizard, and completing the software initialization using the DSCC Setup Service.

Overview

Before the HPE Alletra 6000 OS 6.0.0.400, the initialization of an HPE Alletra 6000 array was accomplished entirely on-premises through the onboard UI or CLI on the HPE Alletra system. With OS 6.0.0.400 and later, this process has been greatly streamlined with the help of the HPE Cloud Connectivity Wizard and the Setup Service within DSCC.

With the HPE Cloud Connectivity Wizard, you enter the minimal amount of networking information that will allow your device to make the initial secure connection to HPE GreenLake and DSCC. After that, the rest of the software initialization is completed with the Setup Service within DSCC where you can take advantage of blueprints—custom templates that reduce the amount of information that must be entered manually for each setup—and immediately begin managing your new system with the rest of your fleet in Data Ops Manager.

Initialization and DSCC onboarding steps

The initialization and onboarding process includes the follow steps:

1. Onboard the array in the Device Manager page of the HPE GreenLake Edge-to-Cloud platform using the serial number and subscription key that ship with your software entitlement.
2. Use the HPE Storage Setup Manager, which is part of the HPE Storage Toolkit, to connect to the HPE Alletra 6000 to run the Cloud Connectivity Wizard.
3. The Cloud Connectivity Wizard runs in a browser and allows you to enter enough networking information, including management IP and subnet, NTP server, and proxy server to make an initial secure connection to DSCC through an encrypted (mTLS) link. When this has completed, a link is provided to launch HPE GreenLake.
4. The last step in the process is to enter the information for system initialization in the DSCC Setup Service Wizard. Here you have the option to either manually step through each part of the setup or apply a previously created blueprint. Applying a blueprint will auto-populate one or more configured fields that may apply to multiple systems deployed in your environment.

When the process is complete, your device is successfully initialized, and you can begin managing your device in DSCC and the onboard UI.

Logging on to the HPE GreenLake Edge-to-Cloud Platform

The first task in onboarding a new HPE Alletra 6000 storage system requires logging in to the HPE GreenLake Edge-to-Cloud platform. If this is your first time logging in to the platform, you will need to create a user account and a company account.

This document includes a high-level overview of the process of creating a user account and a company account. The process is covered in detail in the HPE GreenLake Edge-To-Cloud Platform User Guide, which is available in the Help area of the HPE GreenLake landing page.

Creating an HPE GreenLake user account

The first time you log on to HPE GreenLake, you are prompted to sign up and create a user account. If you already have an HPE InfoSight or HPE Passport account, you can use those credentials to sign in. To begin, click Sign up at the bottom of the Sign In screen as shown in Figure 1. Then follow the steps in the HPE GreenLake Edge-To-Cloud Platform User Guide.

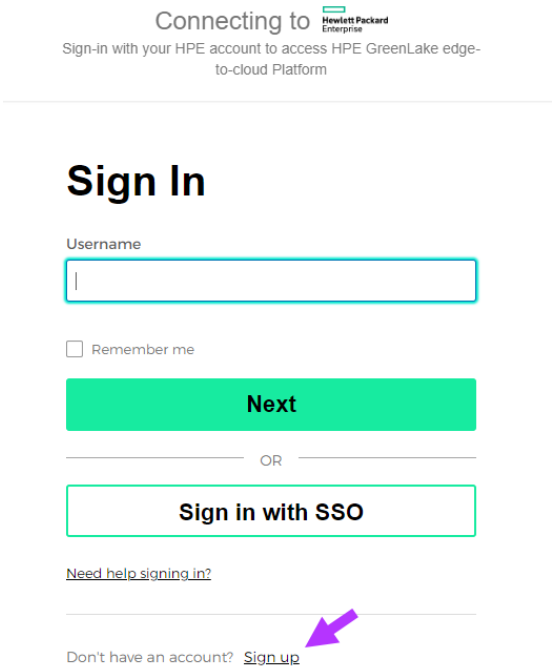

Figure 1. Click Sign Up to begin the process of creating an HPE GreenLake user account.

Creating an HPE GreenLake company account

The HPE GreenLake company account is used to manage your HPE Alletra storage systems. You will have the opportunity later to invite other users from your company to the account so that they can view or manage your HPE Alletra storage, depending on the permissions that they are given.

After creating your email and user account, you will be prompted to create a new account for your company when you first log in to HPE GreenLake. Follow the steps in the HPE GreenLake Edge-To-Cloud Platform User Guide to enter the required information and create the account.

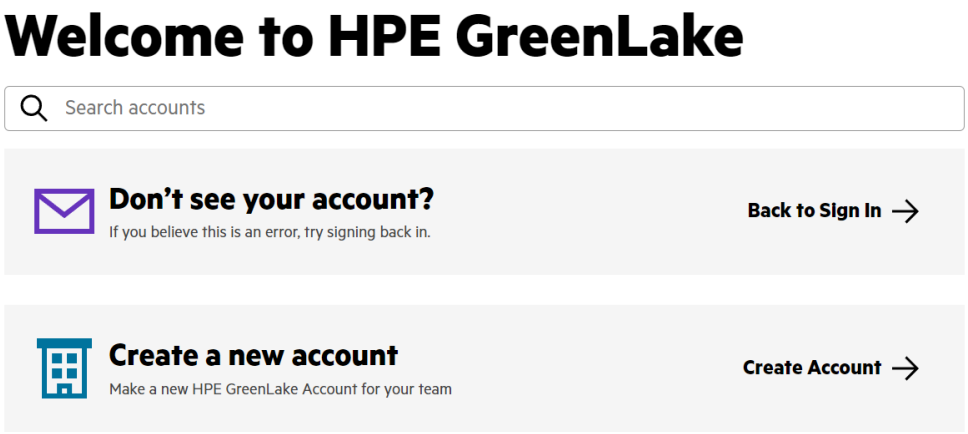

Figure 2. Creating a new HPE GreenLake company account

Adding the Data Services Cloud Console app to your HPE GreenLake company account

Data Services Cloud Console is the application within HPE GreenLake that provides the unified data operations and management for your HPE Alletra storage array. You must select the Data Services Cloud Console app from the App Catalog and add it to your company account. Data Services Cloud Console has multiple instances around the world, so you must also select the geographical region that is closest to your company's storage arrays. If you have storage arrays installed in multiple geographies, you must deploy the application for each region.

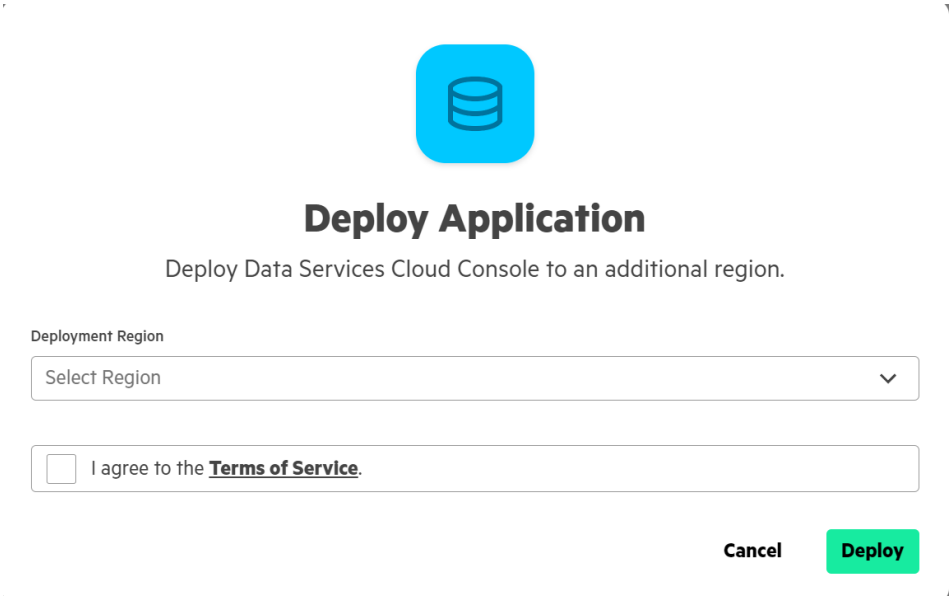

Figure 3. Deploy the DSCC application to your company account.

Adding Data Services Cloud Console permissions to your user account

Your HPE GreenLake user account must have the correct permissions for you to be able to manage your HPE Alletra array. These permissions are specified in HPE GreenLake as role and scope assignments. Without the proper roles and scopes, you cannot access the DSCC Setup Service, which is needed for the next step in the onboarding process. You must also have the correct privileges to be able to manage your storage array. If you created your company account, you will automatically be an account administrator for the Common Cloud Service. In this step, you will assign yourself the administrator role for Data Services Cloud Console and add the scopes that will allow you to manage the array.

Configuring user access

The DSCC administrator role must be added to the user that will be used to onboard the HPE Alletra array. See the HPE GreenLake Edge-to-Cloud Platform User Guide for detailed instructions.

The user who created the company account is automatically assigned administrator privileges for the Common Cloud Service for that organization. The role of administrator for DSCC is not added automatically; it must be added manually to the user account in order to perform the simplified setup of new systems. Figure 4 shows a list of predefined roles available for DSCC.

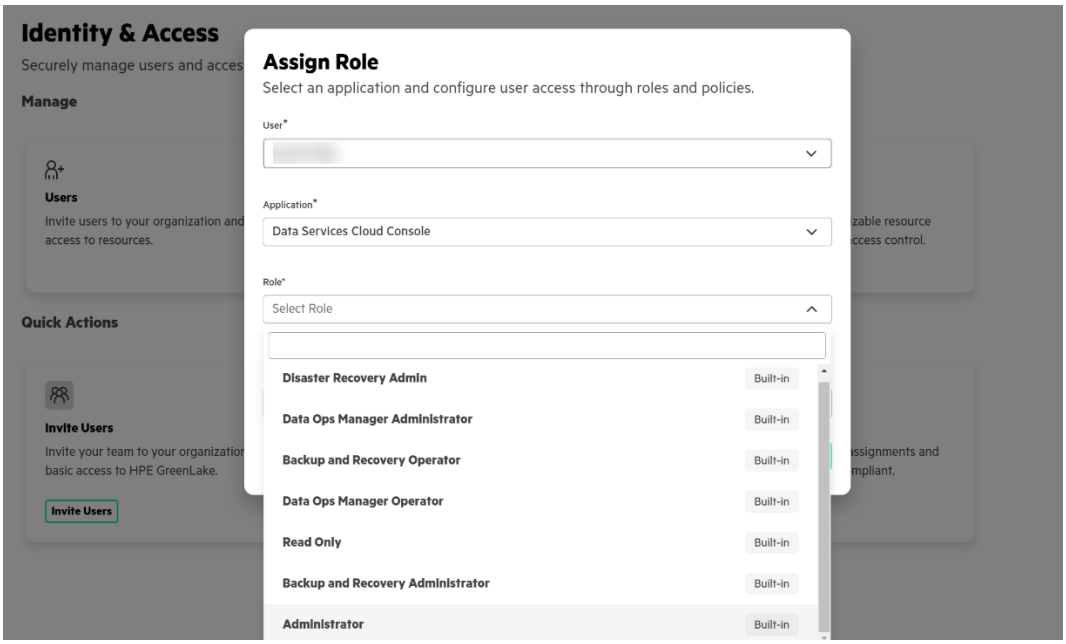

Figure 4. Assigning roles in DSCC for a company account.

Device onboarding to Data Services Cloud Console

Before your array can connect to DSCC via the HPE Cloud Connectivity wizard, you must first onboard the device to an instance of DSCC within your company account. The onboarding process to add the device to your account and assign it to DSCC requires the array serial number and, for purchased or leased storage systems, the subscription key that you receive when you activate your software entitlement.

Onboarding a device to Data Services Cloud Console

To onboard a device to Data Services Cloud Console, complete the following steps:

1. On your company's main DSCC dashboard, from the Onboard Devices box, click Manage Devices.
2. At the top right-hand side of the screen, click Add Devices to add a new device.
3. From the drop-down list, select Storage Devices and click Continue.
4. Select the applicable ownership model for your device and enter the following information:
   a. For Infrastructure as a Service, enter the serial number of the device.
   b. For Purchase or Lease, enter the serial number and the subscription key from the order entitlement activation for the device.

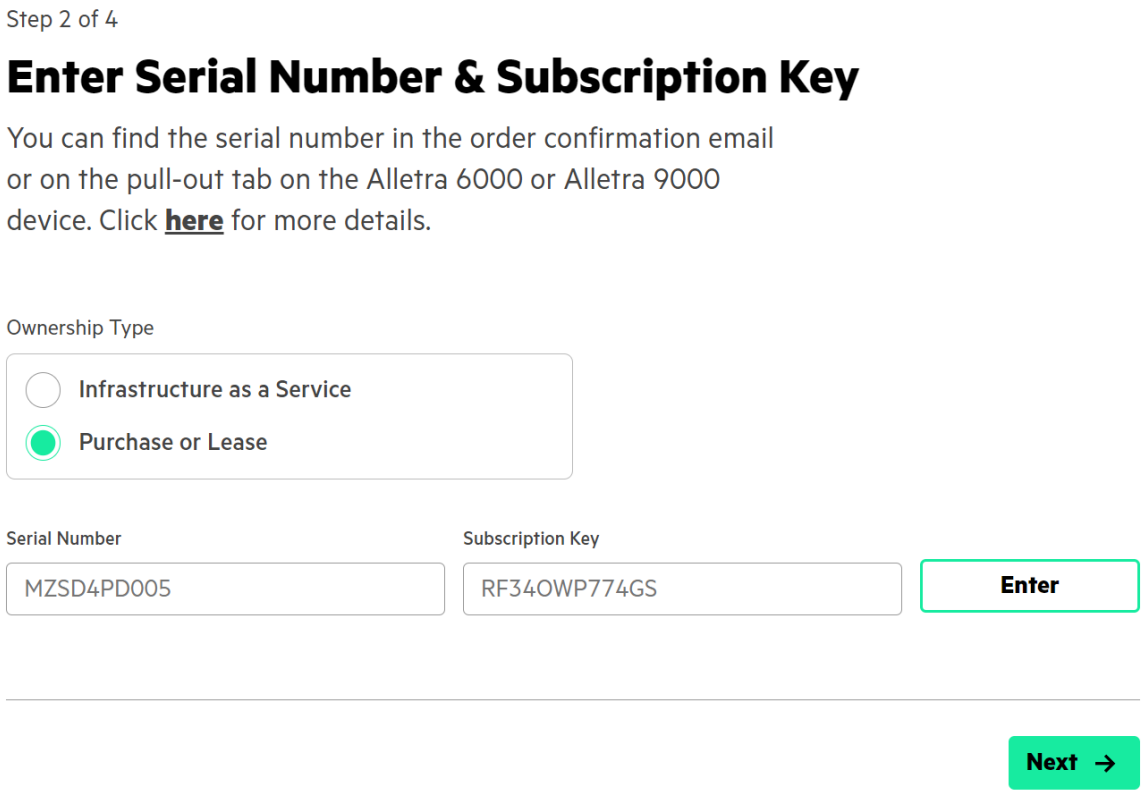

Figure 5. Enter the serial number and subscription key of your new device to onboard it to your HPE GreenLake company account.

5. If you have multiple devices to add, enter the serial numbers and subscription keys if applicable for your additional devices. After you have finished adding your devices, click Next to continue.
6. Optionally, you can create and assign tags to your new device. A tag is a form of metadata that can be used to categorize resources within your HPE GreenLake company account based on purpose, owner, environment, or other criteria. Click Next to continue.
7. Review the devices being added and click Finish.

   It only takes a few minutes for your devices to be added to HPE GreenLake. When they are added, they appear in the list of devices. If you company account has only a single instance of DSCC associated with it, your device is added to that instance automatically. If that is the case, you can omit Steps 7 and 8 and proceed directly to the "Discovering the HPE Alletra 6000 System" section of this paper.
8. If you have multiple instances of Data Services Cloud Console, you must assign your devices to the DSCC application. Place a check mark to the left of each device you want to assign and click Assign to Application.

   You may assign multiple devices to the same geographic instance of DSCC simultaneously. If you have multiple devices in different regions, you will need to run the Assign to Application process for each instance.
9. From the drop-down list on the left, select Data Services Cloud Console as the application, and from the drop-down list on the right, select the geographic location that corresponds to this device or devices.

   The DSCC instances that are visible in the drop-down list are the ones you previously assigned to your company account.
10. Click Finish when you are done.

Note

If you do not see DSCC in the drop-down list under the Application heading, you must add it to your company account. Follow the process in the section titled "Adding the Data Services Cloud Console app to your HPE GreenLake company account" in this paper.

Discovering the HPE Alletra 6000 System

After the storage is cabled to the network and powered on, the next task is to discover and connect to the storage array. There are multiple ways to initially set up an array. Hewlett Packard Enterprise recommends using the HPE Storage Setup Manager for Windows, which is a component in the HPE Storage Toolkit for Windows.

HPE Storage Setup Manager

Installed as a feature of the HPE Storage Toolkit for Windows, the HPE Storage Setup Manager makes it easy to discover and set up uninitialized arrays on the same system. To discover uninitialized arrays on the network, the HPE Storage Toolkit must be installed on a Windows host that either:

- Is connected to the same physical subnet as the array
- Has direct (non-routed) access to the array

You can download the latest version of the HPE Storage Toolkit from HPE InfoSight under Resources → HPE Alletra 6000, Alletra 5000, Nimble Storage → Software Downloads.

When you are installing the HPE Storage Toolkit, you can choose among several features that can be installed. The only feature required for discovering the array and launching the HPE Cloud Connectivity Wizard is the HPE Storage Setup Manager as shown in Figure 6.

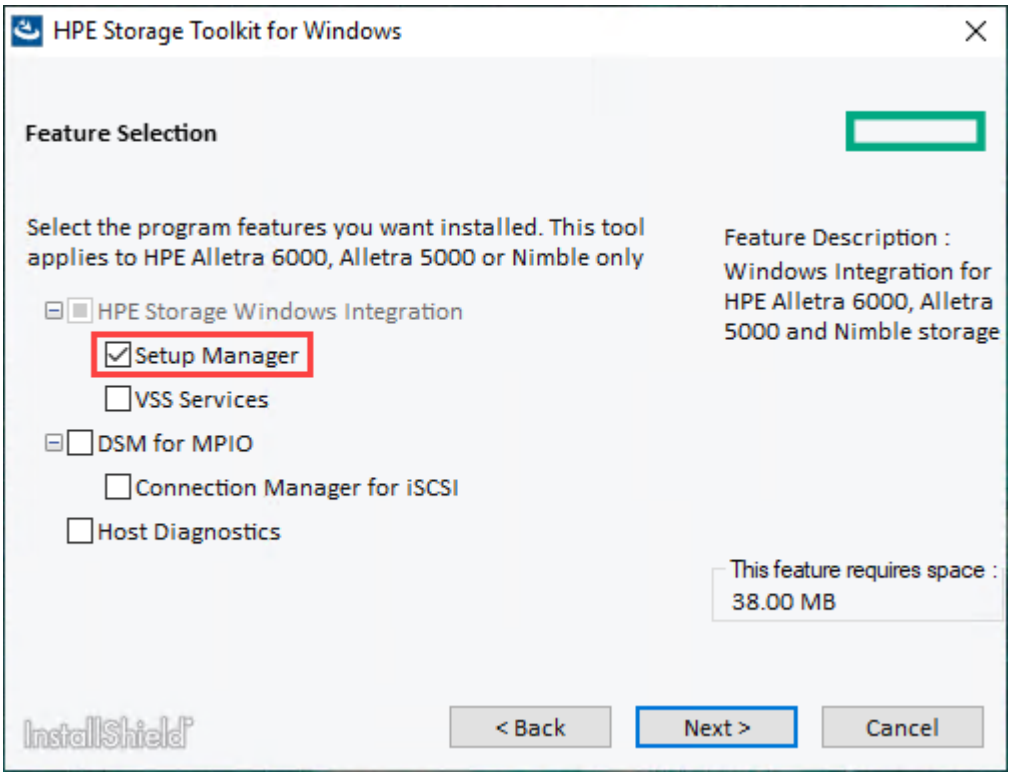

Figure 6. Select Setup Manager when installing the HPE Storage Toolkit for Windows.

After the HPE Storage Setup Manager has been installed, launch the program. A window will open with the list of uninitialized arrays that were discovered on the network. Select an array from the list and click Next to launch the HPE Cloud Connectivity Wizard.

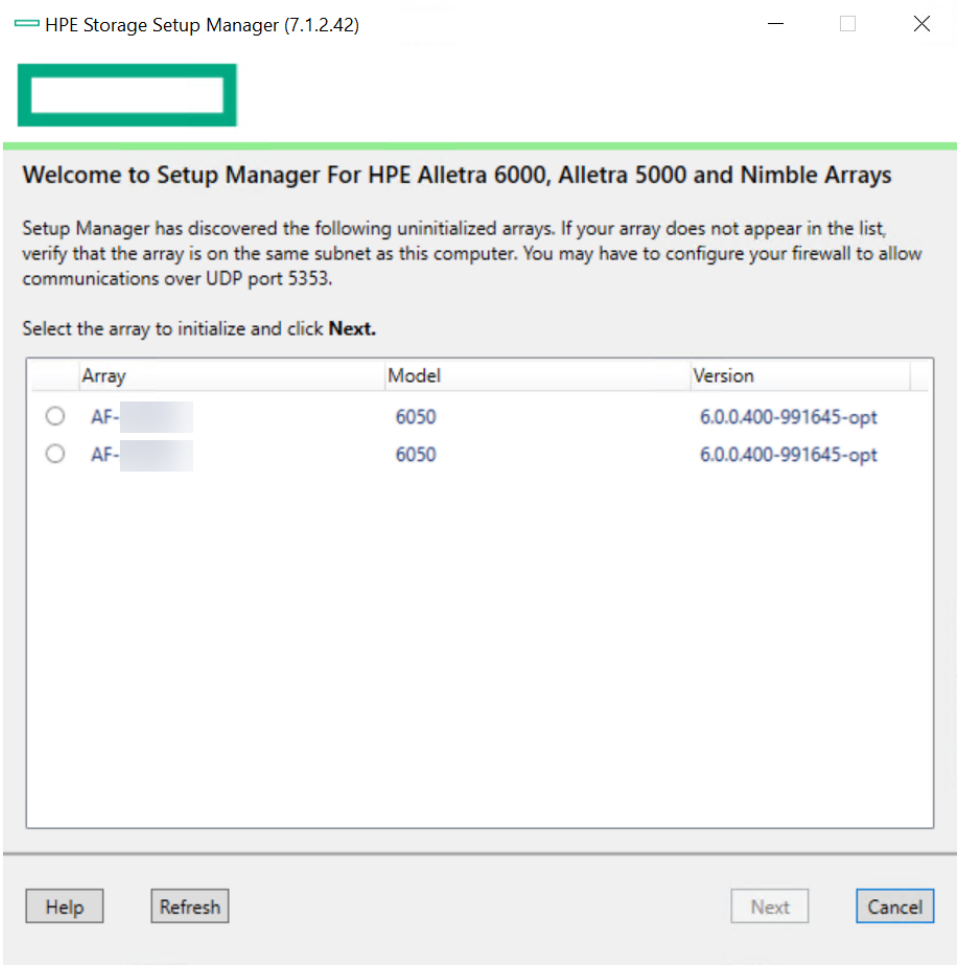

Figure 7. The list will populate with all uninitialized systems discovered on the network.

Note
If you do not have a Windows host from which to complete the initial discovery, you must use the CLI to perform the initial setup of the array. For more information, see the CLI Administration Guide for the operating system version on your array.

Running the HPE Cloud Connectivity Wizard

The HPE Cloud Connectivity Wizard is launched in a browser after you select the discovered array in the HPE Storage Setup Manager. The wizard runs at the site on the array and can be thought of as a pre-setup process that requires you to enter only the networking information that is needed to allow the array to connect to DSCC through the secure tunnel.

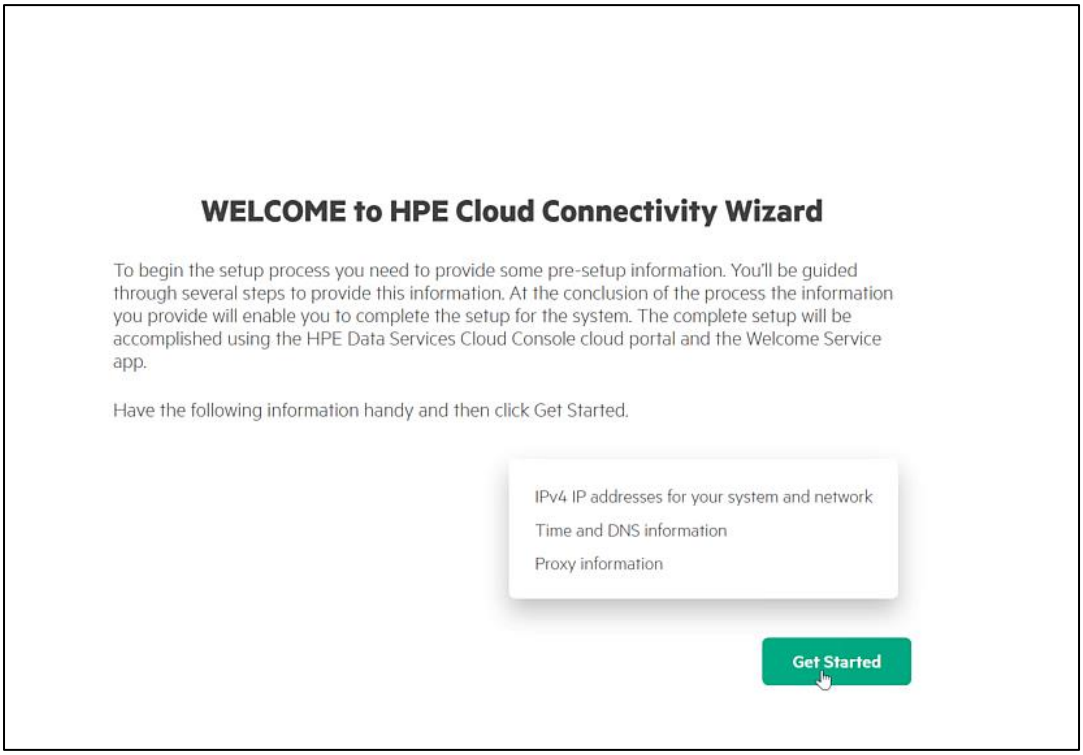

Figure 8. The HPE Cloud Connectivity Wizard will launch in your browser after you select a system to initialize.

The following information must be entered into the HPE Cloud Connectivity Wizard:

- Networking—The IP address, netmask, and gateway that will be used for array management.
- DNS server—The IP addresses of at least one and up to three DNS servers.

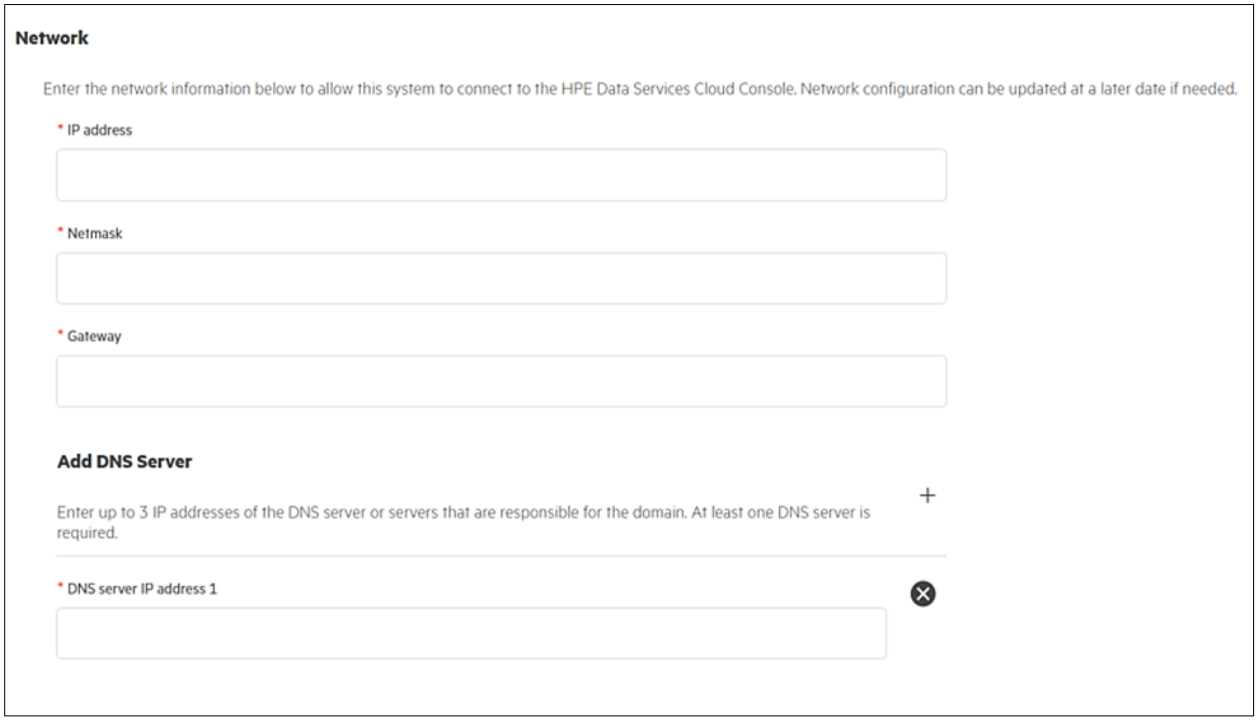

Figure 9. Entering the network information into the HPE Cloud Connectivity Wizard

- Proxy—The proxy server, if required, to connect the device to the internet. Check the box that says Use proxy and enter the address of the proxy server and proxy port. Optionally, you can also enter a proxy server username and password.

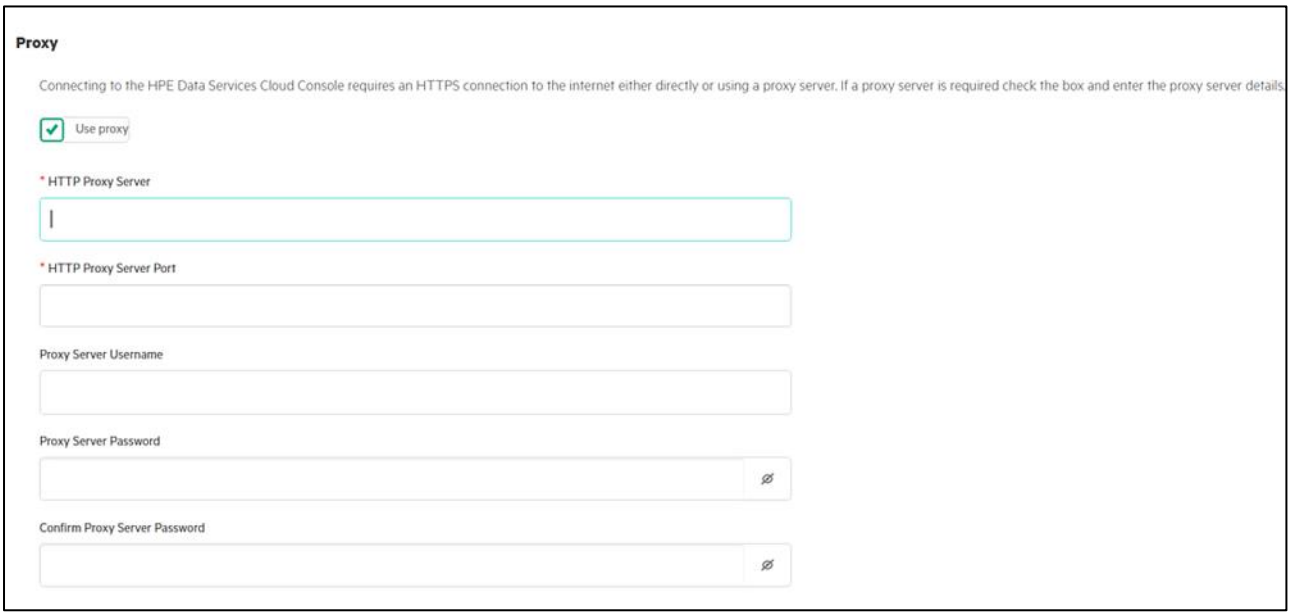

Figure 10. Entering the proxy information into the HPE Cloud Connectivity Wizard

- Time—The address of an NTP server or manually specified date and time. If the time is entered manually, it is very important to verify that it is the correct time. If the time entered is not accurate, you may experience issues trying to connect the storage to DSCC. Hewlett Packard Enterprise highly recommends that you use an NTP server instead of manually entering the time.

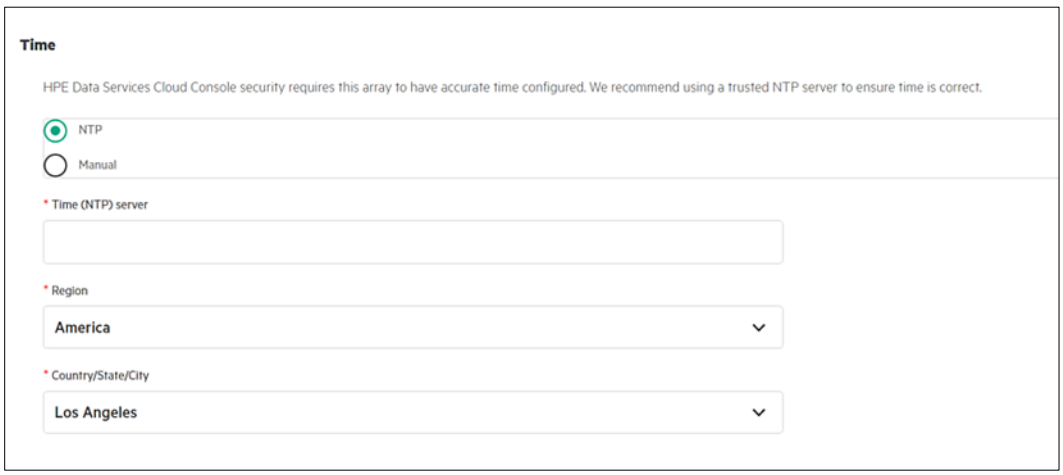

Figure 11. Entering the NTP server or manual date and time and the time zone into the HPE Cloud Connectivity Wizard

After all the information has been entered and confirmed, the array will take a few minutes to apply the settings and make an initial connection to the cloud. Each HPE Alletra system ships from the factory with a device certificate installed that allows the array to connect through an mTLS encrypted tunnel to HPE GreenLake and DSCC. If the wizard returns an error because an entry was incorrectly made in the wizard, the Modify button allows you to go back and make any necessary changes. After those changes have been made, the process continues where it left off.

After the HPE Cloud Connectivity Wizard has completed successfully and the system is connected to DSCC, it displays 100% finished as shown in Figure 12.

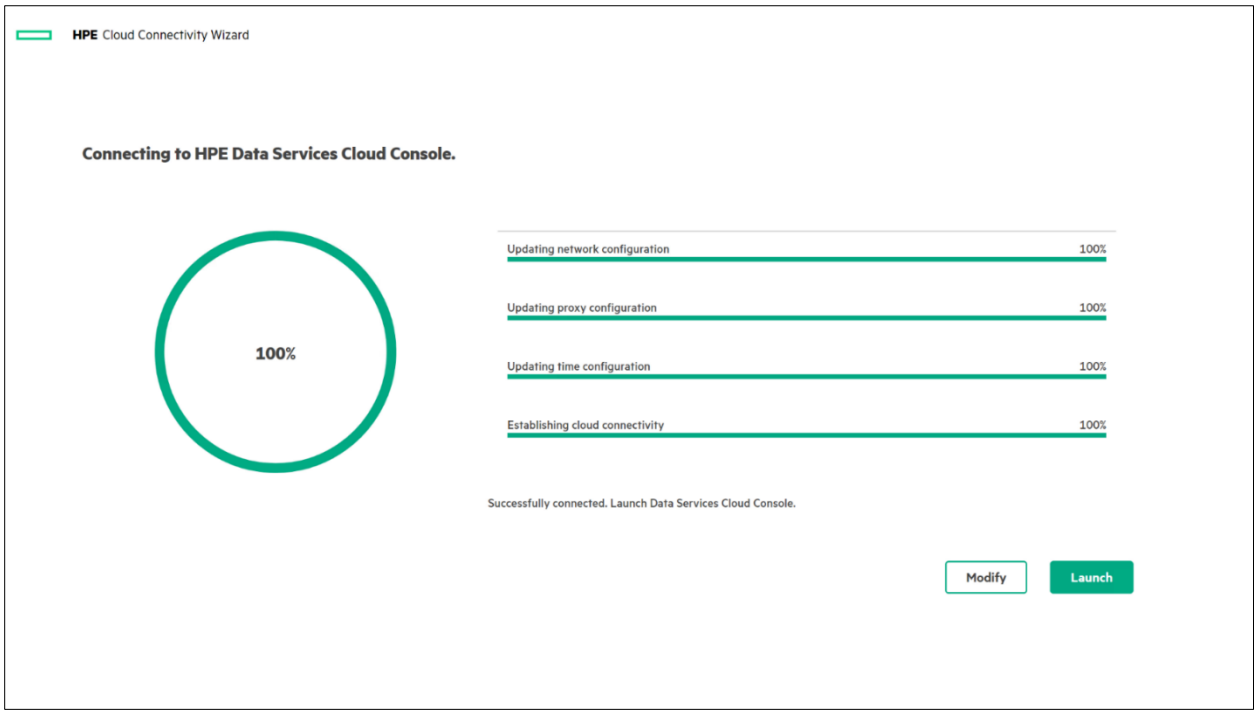

Figure 12. HPE Cloud Connectivity Wizard has completed successfully.

Click Launch Data Services Cloud Console to launch the HPE GreenLake Edge-to-Cloud Platform and complete the array initialized with the Setup Service in DSCC.

HPE Alletra 6000 software initialization

The software initialization of the HPE Alletra 6000 system is initiated from the Setup Service within DSCC. You can initialize a single array or start the initialization of multiple arrays and have them all run in parallel. You also have the option to use blueprints that you created previously to facilitate the process. This can be a significant time-saver, especially if you are onboarding multiple arrays at the same time.

Running the Data Services Cloud Console System Setup Wizard

To run the System Setup Wizard:

1. From your company's dashboard, launch the Data Services Cloud Console application.
2. Launch Setup Service.

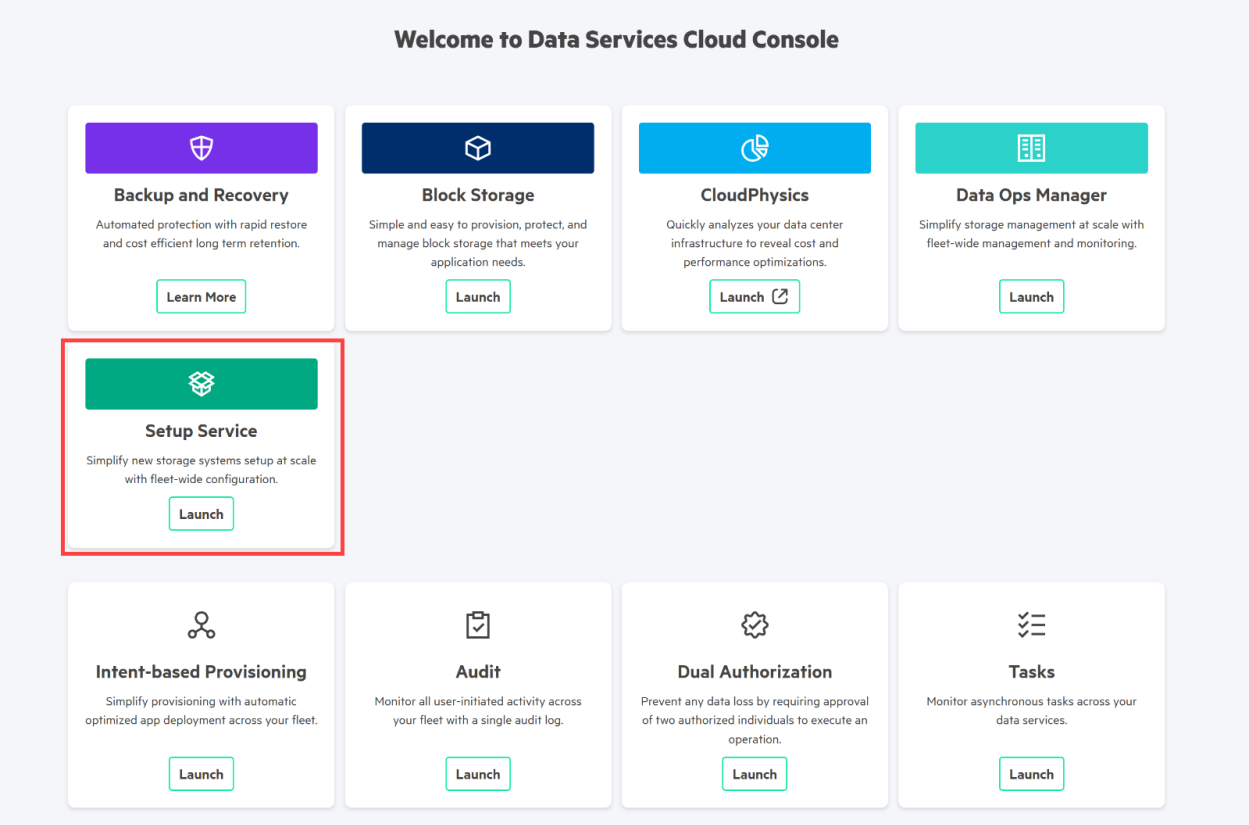

Figure 13. Launching Setup Service from the DSCC landing page

A list will be displayed of the HPE Alletra systems that are either not set up or were recently set up. Any system that has not yet been set up displays in the Setup Status column as "Not Initialized."

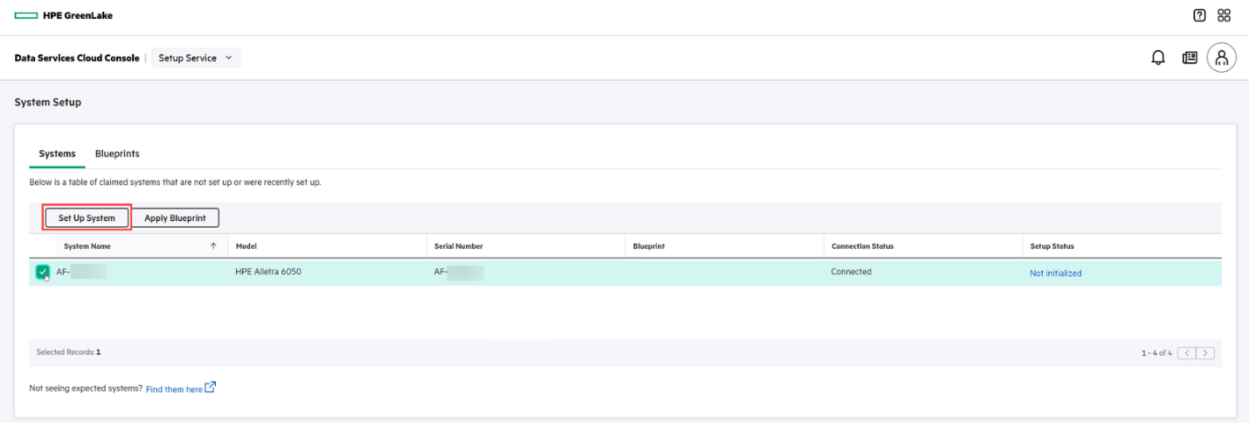

Figure 14. Selecting a system to set up with the HPE Alletra System Setup Wizard

3. Select the system you want to initialize and click Set Up System.
4. Click Get Started to begin entering data into the HPE Alletra Setup Wizard.

   Any data that was previously entered into the HPE Cloud Connectivity Wizard, such as the DNS or time information, is populated into the Setup Wizard; there is no need to enter it again. However, if you want to change any of the settings that were entered into the Cloud Connectivity Wizard, you may enter the information over any pre-populated settings, and the changes will be applied.

   Appendix A: DSCC Setup Service Wizard Settings lists the information that must be entered into the HPE Alletra Setup Wizard.

5. Review your settings in the Review and Finalize step within the HPE Alletra Setup Wizard. If everything looks right, click Submit to start the system initialization.

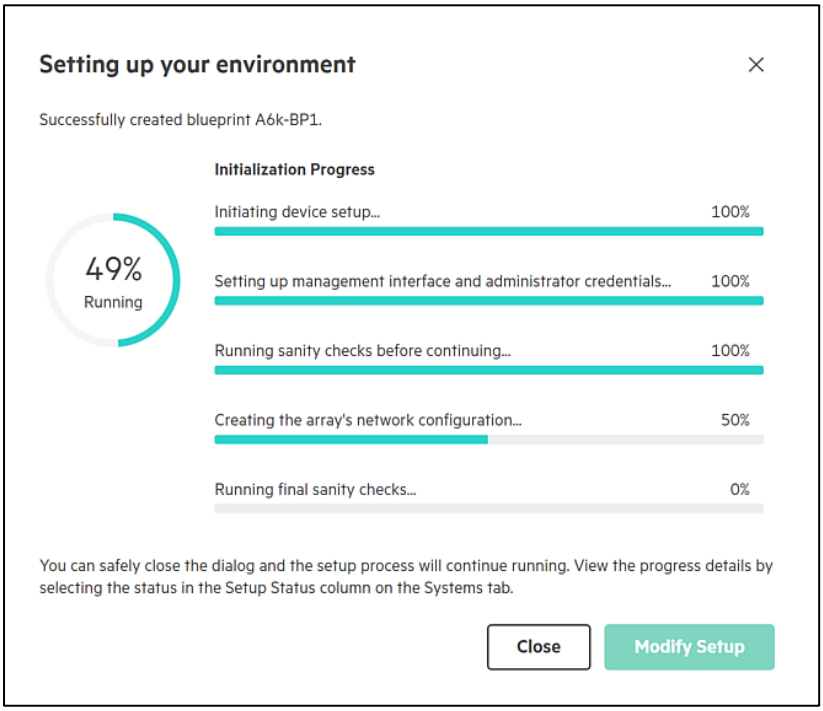

Figure 15. A dialog box allows you to monitor the progress of your system initialization.

If the process experiences an error, you can click Modify Setup to return to the wizard and change any settings that might have been entered incorrectly. The initialization process will continue from where it left off.

You can choose to keep the dialog box open, or you can close it—the process will continue running. This enables you to initialize multiple systems and allow their setups to continue in parallel. If you close the dialog box, you can return to it at any time by clicking Initializing under the Setup status column for the system on the System Setup page.

6. When the initialization of your HPE Alletra system is complete, the dialog box status displays "100% Completed." In the list of systems in the Systems tab, the Setup Status will change to "Initialized."

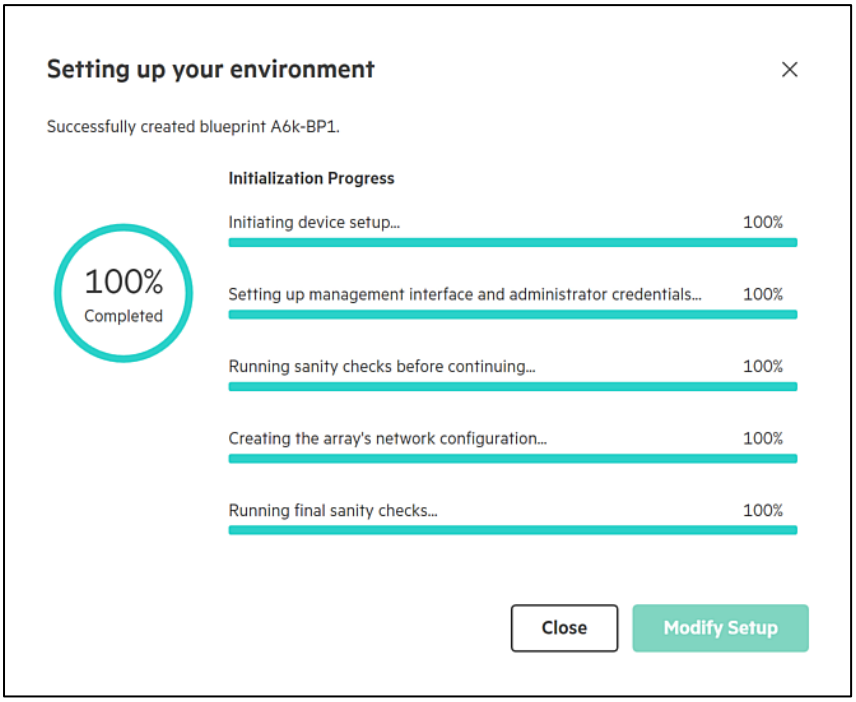

Figure 16. System initialization has completed successfully.

7. Now that the initialization of your system is complete, you can proceed to the Data Ops Manager app in Data Services Cloud Console to begin managing your storage and to the Block Storage app to begin provisioning.

DSCC Setup Service blueprints

Using a DSCC blueprint as a template enables you to enter settings one time and then apply them to each system that you need to initialize. Blueprints also allow you to prepare for the initialization of your storage arrays ahead of the date when you plan to initialize them. You can create the blueprints at any time with the settings that will apply to your new arrays, and when it is time to install and initialize the array, you can apply the blueprint you created earlier. When you use a blueprint to set up a system, you have the option to retain or re-enter information in any of the prepopulated fields.

There are two fields that cannot be pre-populated for an HPE Alletra 6000 in a blueprint. These fields must be filled out manually for each system:

- The administrator credentials used for the on-premises management tools such as the onboard UI and the CLI
- Proxy credentials for configured proxy servers (if applicable)

Blueprints can be used for heterogeneous HPE Alletra storage platforms, which means you can create common blueprints that can be applied to both HPE Alletra 6000 and HPE Alletra 9000 systems. This simplifies the setup of environments with both business-critical and mission-critical tier HPE Alletra arrays.

Fields that are applicable to one platform but not the other, such as storage networking, are ignored when the blueprint is applied to a different model system. For information specific to HPE Alletra 9000 setup, see the HPE Alletra 9000 Storage Simplified Setup white paper.

Creating a blueprint

To create a new blueprint, from the System Setup dashboard in the DSCC Setup Service, launch the HPE Alletra Blueprint Wizard.

1. From your company's dashboard, launch Data Services Cloud Console.
2. In the Setup Service box, click Launch.
3. Select the Blueprints tab.
4. Click the + sign to launch the wizard.
5. Click Get Started to begin entering data into the wizard.

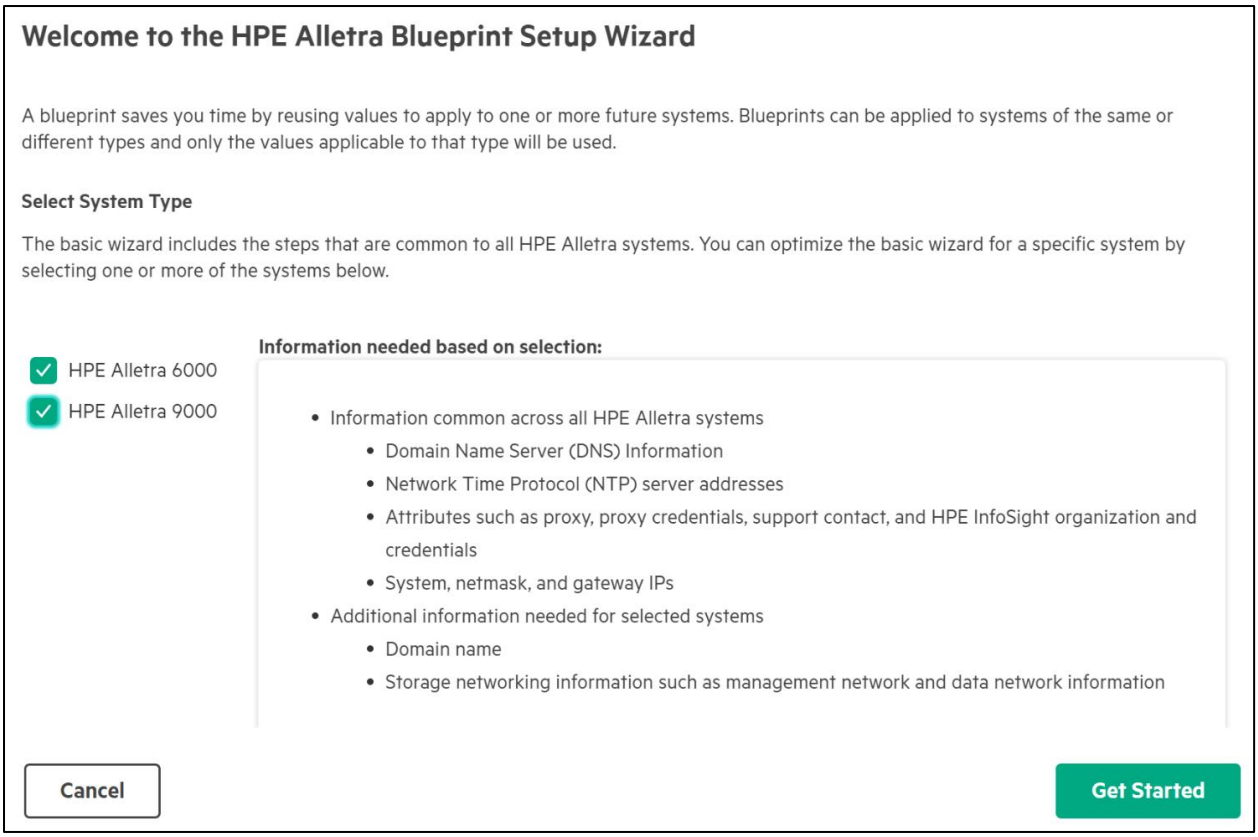

Figure 17. The HPE Alletra Blueprint Setup Wizard allows you to create a blueprint for the HPE Alletra 6000, HPE Alletra 9000, or both.

6. After you have finished entering information into the wizard, give the blueprint a name, review the settings, and click Submit.
7. Your blueprint is now saved and visible in the list of blueprints. You can return to the Blueprints tab at any time and edit the information saved in a blueprint by selecting the blueprint you want to modify and clicking the pencil icon.

Blueprints can also be created as part of the System Setup Wizard during the Review and Finalize step, allowing you to save the fields that were entered while initializing a storage system in order to apply them to future systems.

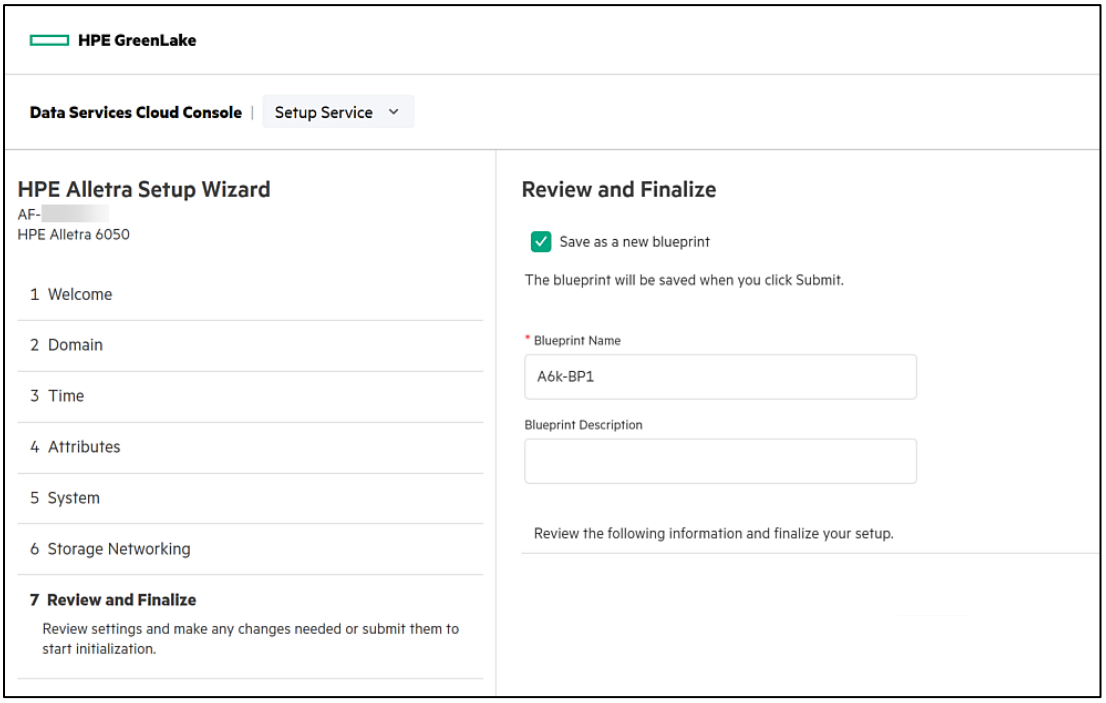

Figure 18. Creating a blueprint from the HPE Alletra Setup Wizard

Setting up a system using an HPE Alletra blueprint

After a blueprint has been created, you can use it to set up uninitialized systems.

1. In the Systems tab in Setup Service, select the system you want to initialize and click Apply Blueprint.

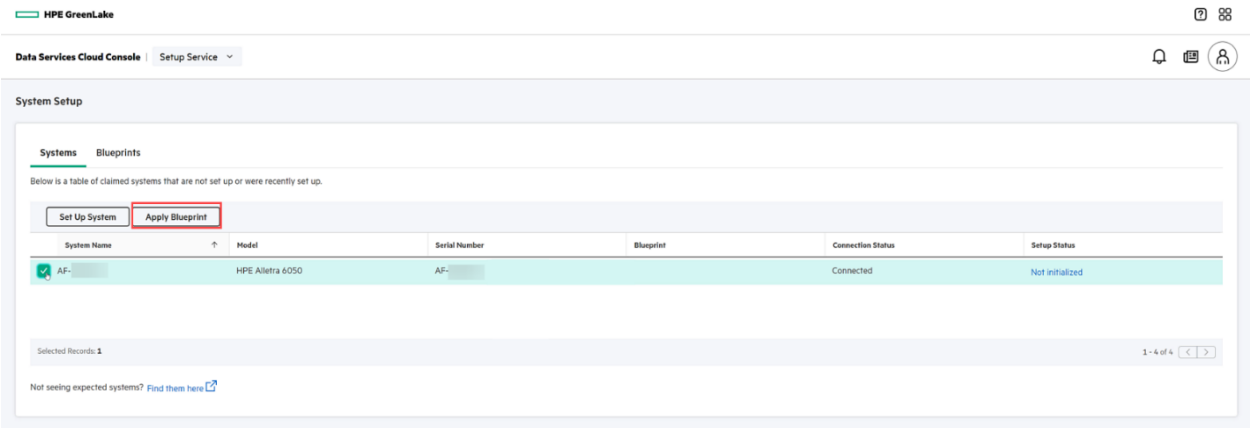

Figure 19. Applying a blueprint to an uninitialized HPE Alletra storage system

2. From the list of available blueprints displayed, select the one you want to apply to this HPE Alletra system and click Apply.

   The HPE Alletra System Setup Wizard launches. This process follows the same flow as described in the "Running the Data Services Cloud Console System Setup Wizard" section in this paper; the only difference is that the fields you have entered in the blueprint are pre-populated. You have the option of keeping any of the information in the Setup Wizard or re-entering the information to change the settings to apply to this storage array.

3. Review your settings in the Review and Finalize step within the HPE Alletra Setup Wizard. If everything looks right, click Submit to start the system initialization.

   A dialog box appears on the screen to enable you to monitor the steps in the initialization process. If the process experiences an error, you can click Modify Setup to return to the wizard and change any settings that might have been entered incorrectly. The initialization process will continue from where it left off.

Summary

Initialization of an HPE Alletra 6000 storage system can now be performed in the HPE GreenLake Edge-to-Cloud platform with the HPE Cloud Connectivity Wizard and the DSCC Setup Service. The HPE Cloud Connectivity Wizard is used to provide enough networking information to allow the HPE Alletra system to make an initial connection to DSCC, after which the DSCC Setup Service can be used to initialize the storage arrays.

With Setup Service, you can initialize multiple HPE Alletra storage systems in parallel and take advantage of blueprints to accelerate the deployment of your fleet. With the setup, management, and provisioning of the cloud native HPE Alletra storage arrays available under the single, unified DSCC interface, managing the end-to-end lifecycle of your data has never been easier.

Appendix A: DSCC Setup Service Wizard Settings

The HPE Alletra System Setup Wizard and the HPE Alletra Blueprint Wizard share a common set of inputs. The only exception is that there are a few fields that cannot be pre-populated in a blueprint but must be filled out manually for each system. Fields that cannot be entered into the blueprint are noted in Table 1. Table 1 summarizes the information that must be entered into the HPE Alletra System Setup and HPE Alletra Blueprint Wizards for the HPE Alletra 6000.

Table 1. Summary of fields in the HPE Alletra System Setup and HPE Alletra Blueprint Wizards that are applicable to the HPE Alletra 6000

Field	Required	Can be entered into blueprint	Notes
DNS servers	Yes	Yes	One DNS server is required, up to a maximum of three servers.
Domain	Yes	Yes	The network domain that the system will be a part of.
NTP servers	Yes	Yes	One NTP server is required, up to a maximum of three servers. Acceptable formats are hostname or IPv4.
System time zone	Yes	Yes	Pertains to the geographic location where the storage array resides.
HTTP proxy server	No	Yes	Proxy server hostname or IPv4 address and port. Proxy server credentials can be added in the System Setup wizard, but not the Blueprint wizard.
Support contact	Yes	Yes	Name, country, preferred language, company, phone number, and contact email address.
SMTP Server	No	Yes	The hostname or IPv4 address of an SMTP server. If an SMTP server is configured, you will also need to specify a from email address to identify which system is sending alerts.
System Information	Yes	Yes	System name and system country location. The system name will also be the group name.
Network information	Yes	Yes	Primary management IPv4 address, subnet mask, and gateway.
System credentials	Yes	No	Admin user account for on-premises management.
Management networking	Yes	Yes	Includes the secondary management IPv4 address, MTU size, optional additional management interface, and the diagnostic IPv4 addresses for the controller nodes.
Data networking	For iSCSI only systems	Yes	For iSCSI only systems, at least one iSCSI data subnet is required. iSCSI data subnets are optional for Fibre Channel systems.

For more detail on the various fields that are part of Setup Service, see the Data Services Cloud Console Articles embedded in DSCC. You can search for Setup Service to get a full description of all fields that are part of the HPE Alletra System Setup Wizard and the HPE Alletra System Blueprint Wizard.

Resources, contacts, or additional links

HPE Welcome Center

HPE GreenLake Edge-to-Cloud Platform User Guide

HPE Alletra 6000 CLI Administration Guide

HPE Storage Toolkit Windows Integration Guide

HPE Alletra 9000 Storage Simplified Setup

Learn more at

hpe.com/storage

a00126367ENW

Hewlett Packard Enterprise