HPE GreenLake

Technical white paper

HPE GreenLake: Data Services Cloud Console

Streamlining storage fleet management and operations in the cloud

Contents

Executive summary

As the scale of IT environments continues to expand globally at a rapid pace, operational complexity and costs have likewise grown to unsustainable levels for many organizations. The complications that arise from these scaling environments can immobilize operations, introduce security vulnerabilities, and result in underutilized resources as the challenge of infrastructure tracking and management becomes increasingly difficult.

Data Services Cloud Console (DSCC) provides a unique cloud operating and consumption experience for storage that streamlines data and infrastructure management across your environment while delivering scalability and state-of-the-art security.

With DSCC, the fragmented data management tools, manual processes, and infrastructure silos that impede productivity in conventional on-premises management models are transformed into a scalable operational experience. With the ability to unify an organization's geographically dispersed infrastructure from a single console, DSCC mitigates the traditional burdens of infrastructure management and simplifies the management, protection, and monitoring of your storage assets.

This technical paper is designed for anyone who manages or is interested in managing data and its entire lifecycle in the cloud by using DSCC. The paper walks you through the requirements of accessing DSCC through the HPE GreenLake edge-to-cloud platform and provides a breakdown of the various operations you can perform within DSCC. This document is not intended to provide detailed, prescriptive steps for everything you can do within DSCC. For more information, see the referenced documentation in this paper, the Getting Started section on the Help page in the HPE GreenLake platform, and the Data Services Cloud Console articles within DSCC.

Data Services Cloud Console

Nested within the HPE GreenLake platform, DSCC serves as the primary application platform that offers a unified experience for all storage services and operations. It provides a single portal that enables a consistent front-end user experience for HPE Alletra as well as for all HPE as-a-service storage offerings. DSCC unifies the management of heterogeneous platforms and services that have traditionally required separate, distinct interfaces, resulting in a disjointed user experience.

DSCC is fully integrated into the HPE GreenLake platform. It offers consolidated billing, user management, and device inventory tied in with management for the rest of the HPE portfolio, enabling a cloud-native experience for your entire data center.

Fleet management using Data Services Cloud Console

With the ability to deploy and manage storage from a single console, DSCC delivers a new, unique operational paradigm. Users can set up, configure, and manage their hardware fleet with a variety of tools provided by DSCC. The Setup Service enables simplified installation and deployment of groups of storage arrays with minimal manual intervention. With Data Ops Manager, federated updates enable you to view and apply software updates for your entire install base without having to individually log in and execute the updates on each system. Intent-based provisioning with block storage enables you to provision and protect storage based on your required storage tier while abstracting out individual system information. This capability can help ensure uniform provisioning and utilization across all your arrays, preventing oversubscription of individual systems and wasted resources. HPE GreenLake for Backup and Recovery offers a unified protection experience for VMware® environments and enables simple scheduling and management of local recovery points as well as backups to the cloud.

When combined, these features help construct the unique, unified management experience for all your cloud-enabled storage arrays — your fleet. Fleet management with DSCC enables you to manage very large, complex, and heterogenous data infrastructures with efficiency and agility.

Table 1. Services in DSCC

Service title	Description
Setup Service	Simplifies the installation and deployment of groups of storage arrays with minimal manual intervention and includes blueprints that can be saved as a template in advance and applied to multiple arrays.
Data Ops Manager	Enables you to monitor all of your systems from a single dashboard, create fleet-wide host definitions, and view and apply software updates for your entire install base with federated updates.
Block Storage	Enables simple provisioning and protection of block storage based on the application's required storage tier with intelligent recommendations that can help ensure uniform provisioning and utilization across all arrays.
File Storage	Enables the creation and configuration file shares with unified management, monitoring, and protection of file storage.
Private Cloud Business Edition	Enables the end-to-end management of dHCI environments, including cluster lifecycle management, virtual machine deployment and protection, and integration with public cloud environments.
Backup and Recovery	Enables automated data protection for VMware and AWS resources by using prebuilt and customizable protection policies to create app-consistent, immutable recovery points.
Disaster Recovery	Enables replication of mission-critical applications and data for virtualized environments with near-zero recovery point objectives and the ability to recover your data to any point in time, quickly and efficiently mitigating the effects of ransomware attacks and other outages.
Storage Fabric Management	Enables the configuring, monitoring, and managing your Fibre Channel SAN fabrics, leveraging integration with DSCC-managed storage arrays and providing advanced automation and orchestration capabilities for SAN fabric life cycle management.

Security architecture in Data Services Cloud Console

Managing your storage arrays through DSCC offers industry-standard security assurances, ensuring that the convenience of streamlined cloud management for the arrays never comes at the cost of security vulnerabilities. Communication with DSCC is always initiated by the storage devices over outbound port 443 — DSCC never initiates connection to your storage arrays.

A zero trust security framework is employed with Mutual Transport Layer Security (mTLS) certificate authentication and encrypted communication to establish the secure tunnel and keep traffic between the array and DSCC private and secured. After connectivity has been established between the storage arrays and a DSCC instance, all management commands are initiated through API calls and all traffic over the internet is encrypted. The user data always stays on the array safely in the data center — it is never sent to or stored on DSCC.

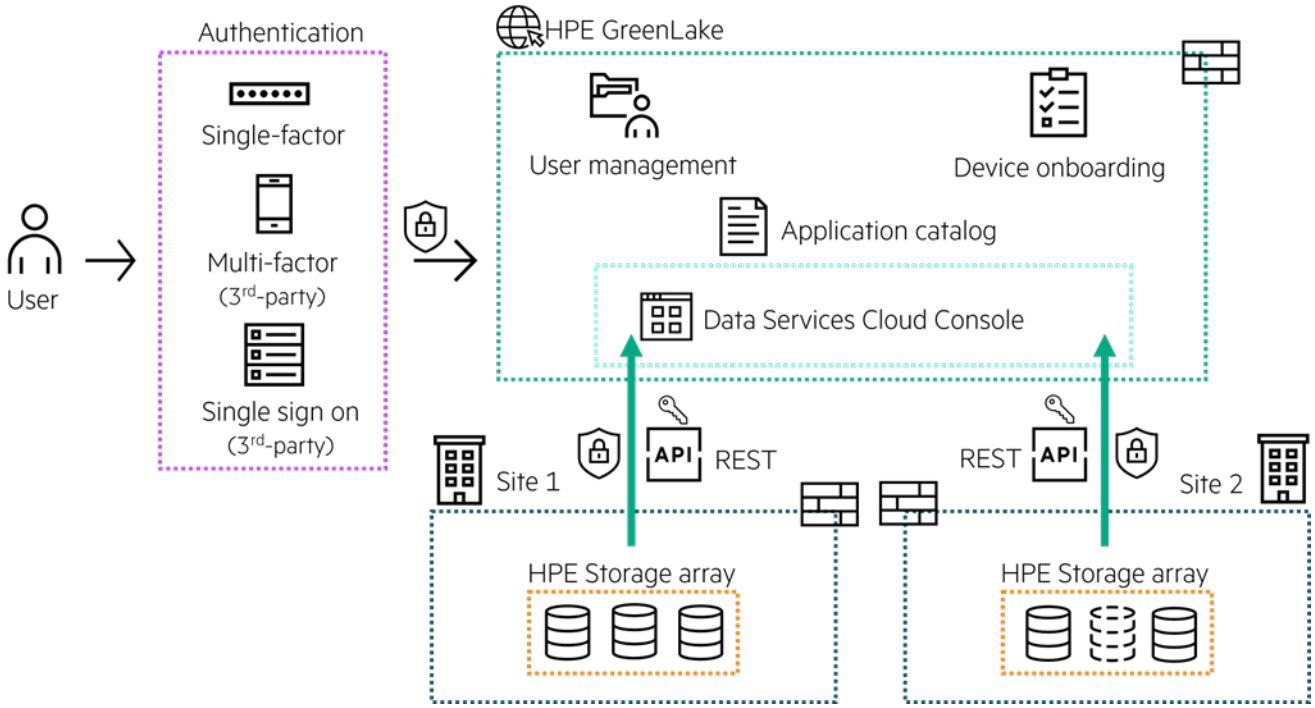

Figure 1. Security and authentication in DSCC

Within the HPE GreenLake platform, multi-factor authentication or single sign-on (SSO) can be configured in addition to single-factor sign-on for user login to meet your organization's security requirements and help prevent rogue user access. Role-based access control gives customizable access to specific functions within DSCC, ensuring that different roles within your organization can be granted appropriate access to devices and services.

Note
For more information about security in DSCC, see the Data Services Cloud Console Security Guide.

Getting started on the HPE GreenLake platform

Before you can begin managing, provisioning, and protecting data in DSCC, you must first have an HPE user account so that you can log into the HPE GreenLake platform at console.greenlake.hpe.com, and you must have a company workspace. Devices and service subscriptions are associated with a company workspace, and your user account ID can be used to manage any company accounts you have been invited to.

Note
For more information about creating an HPE account and an HPE GreenLake company workspace, see the HPE GreenLake Edge to Cloud Platform User Guide.

Adding the DSCC application to your HPE GreenLake workspace

To add the DSCC application to an HPE GreenLake company account, log in to your company account and navigate from the dashboard to the Applications page. Select the DSCC application from the Available Applications section to add it to your company account. DSCC has multiple instances around the world; select the geographical region or regions that are closest to your company's storage arrays.

Configuring user access for DSCC

Your HPE GreenLake user account must have the correct permissions for you to be able to manage your cloud-enabled HPE storage arrays and leverage subscribed services. These are specified in HPE GreenLake as role assignments with the option for resource restriction policies. Without the proper role assigned, you cannot access DSCC and start managing your devices and services.

- **Roles:** A role defines a user as an administrator, an operator, or a read-only observer for a specific application in the HPE GreenLake platform. To manage and provision your storage arrays or manage services in DSCC, you must first assign a role. Multiple roles are available; the administrator role gives you full privileges for managing all devices and services in your environment.
- **Resource Restriction Policy:** You can restrict roles more granularly by creating a resource restriction policy. Resource restriction policies allow you to enable access to a list of specific resources within the application, also known as scopes. The default, pre-created roles automatically have access to all scopes.

Onboarding and assigning devices to DSCC

Before you can manage your installed storage arrays in DSCC, you must first onboard your device to an instance of DSCC in the Device Manager. The onboarding process uses the array serial number and the subscription key. Depending on the device you are onboarding and the version of the operating system, the subscription key can be obtained in several ways:

- The HPE Software Center
- The array onboard UI
- The Cloud Connectivity Wizard

Adding a service subscription

To use services that are not tied directly to a specific storage device within DSCC, such as HPE GreenLake for Backup and Recovery or HPE GreenLake for Storage Fabric Management, you must obtain a valid subscription key for the service and add the key to the Service Subscriptions list in the Subscriptions manager. After the key has been added and validated, you will have access to the service through the DSCC landing page.

Note
For more information about how to install and onboard your devices and add services, visit the HPE Welcome Center or refer to the documentation in the Resources section of this paper.

DSCC end-to-end management experience

Now you can begin managing your storage devices in DSCC. The rest of this document focuses on the tiles within DSCC and how each pertains to a different aspect of the end-to-end storage management experience.

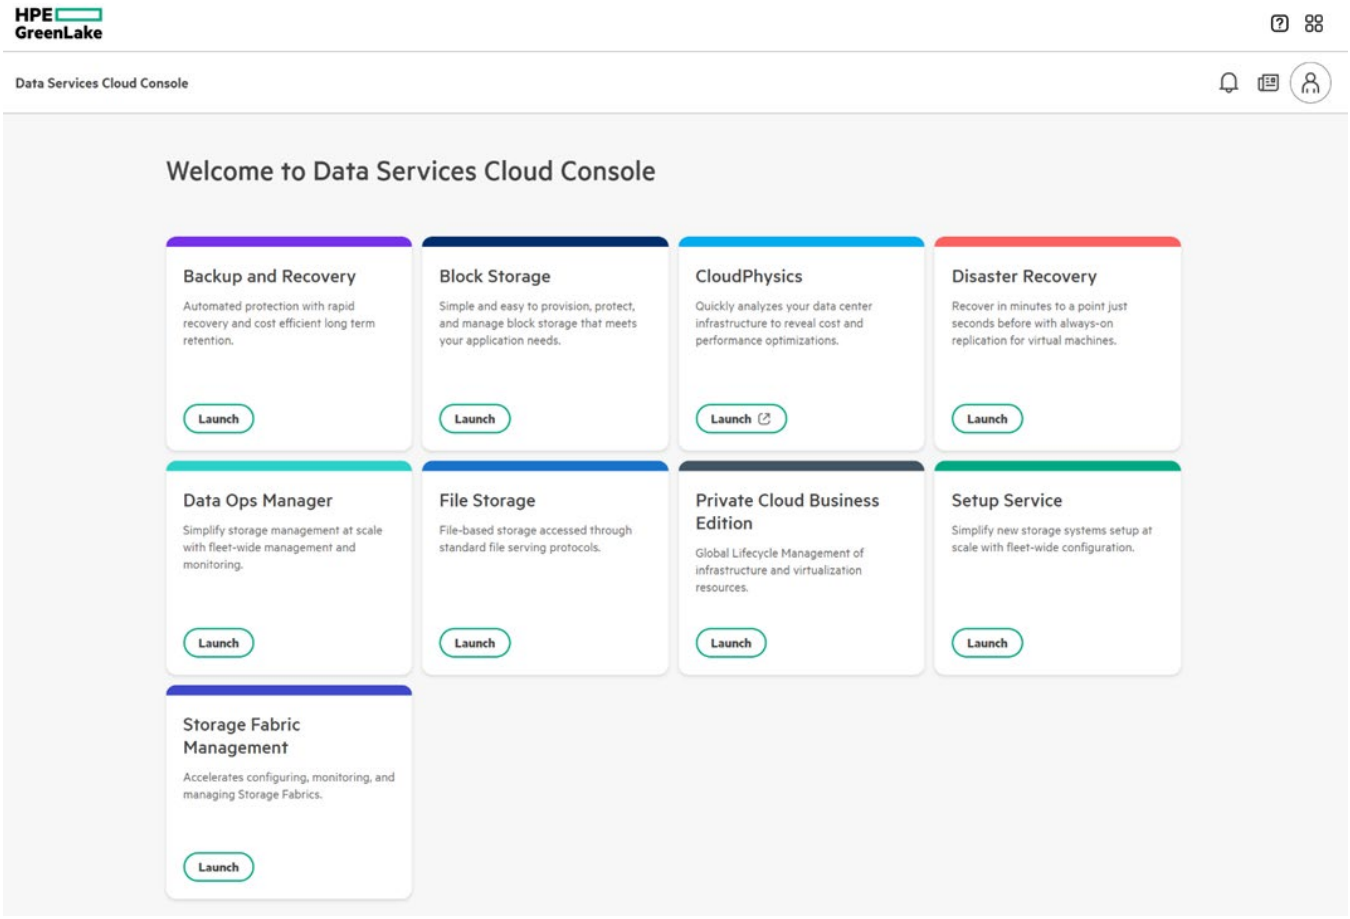

Figure 2. Data Services Cloud Console landing page with the services and applications currently available

Setup Service application

The Setup Service is a cloud-based application that uses wizards to help you set up and initialize newly installed HPE Alletra Storage systems. DSCC also enables you to deploy your new data infrastructure in minutes with automated discovery and activation. You cable up the system, power it on, enter a minimum amount of networking information, and then use the Setup Service within DSCC to do the rest. The Setup Service blueprint feature streamlines the initialization process with the use of custom user-defined templates that can be applied to one or more systems.

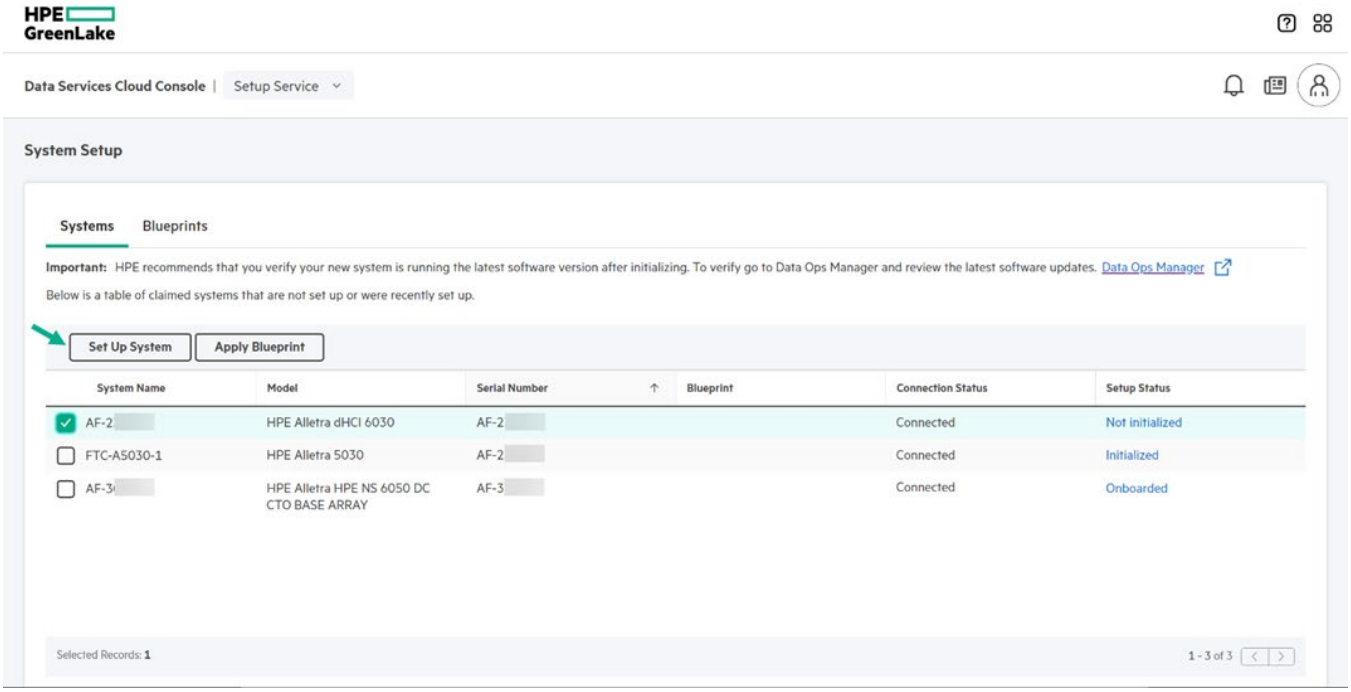

Figure 3. HPE Setup Service in DSCC

Setup Service is particularly useful if you are setting up multiple arrays because all the systems can be initiated and run in parallel from a single console.

Blueprints configured with secrets can facilitate the entirely automated setup of systems after connectivity to DSCC has been established. Secrets enable you to securely create and store credentials in DSCC. If you apply an authorization secret to a blueprint and assign the blueprint to a system that has not yet been initialized, setup according to that blueprint is performed automatically after the system can communicate with DSCC.

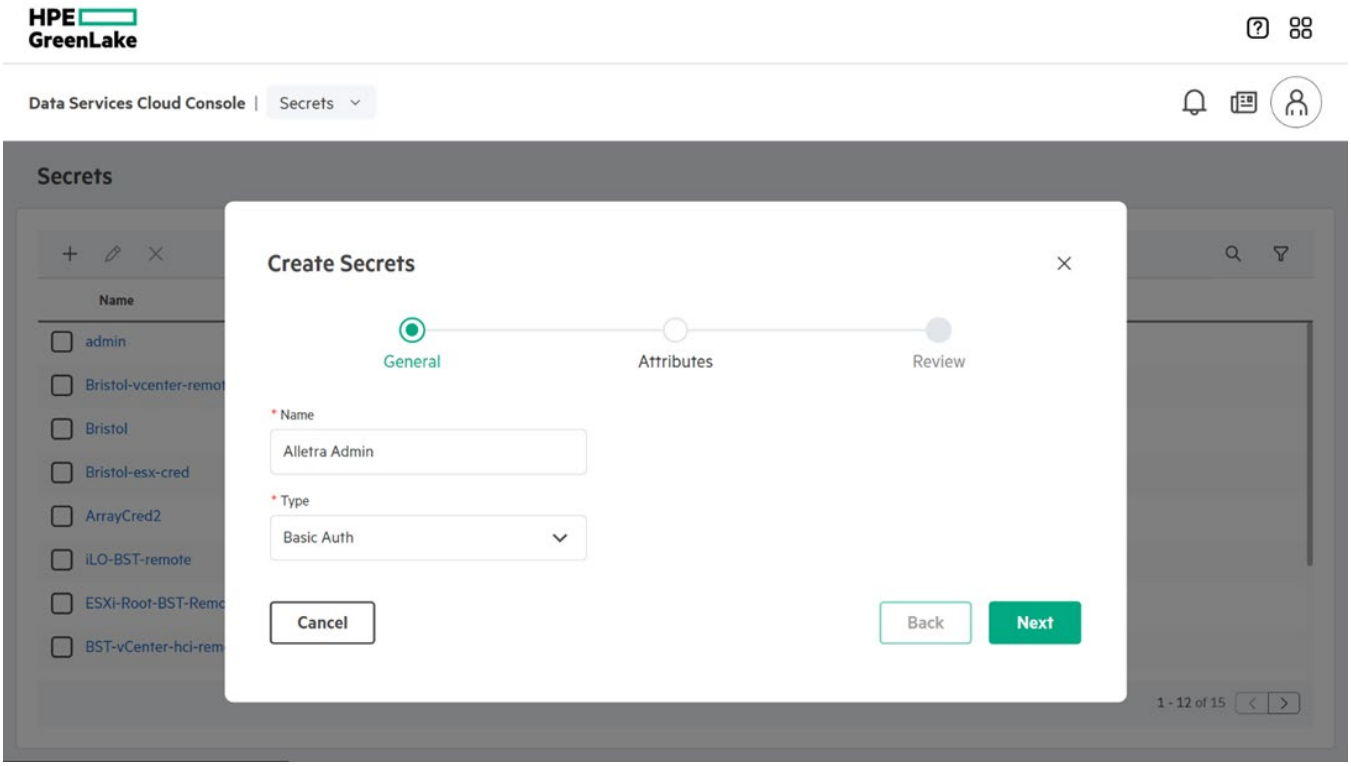

Figure 4. Creating secrets for credentials in DSCC

For more information about how to initialize new devices using the Setup Service, refer to the HPE Welcome Center.

Data Ops Manager

Data Ops Manager serves as the central management hub for HPE arrays and simplifies storage management at scale. The intended users for Data Ops Manager are cloud operators and anyone who will be responsible for infrastructure and fleet management within the organization.

The following hardware platforms are currently supported for management with Data Ops Manager within DSCC:

- HPE Alletra Storage MP
  - HPE GreenLake for Block Storage
  - HPE GreenLake for File Storage
- HPE Alletra 5000
- HPE Alletra 6000
- HPE Alletra 9000
- HPE Primera arrays running HPE Primera OS 4.4 and later
- HPE Nimble Storage AF and HF series arrays (Gen5) running HPE NimbleOS 6.0.0.300 and later

Operations that can performed inside Data Ops Manager include:

- System management, including networking, software updates, and system settings
- Configuration of replication topologies
- System monitoring for issues, capacity trends, and performances statistics
- Creation and management of hosts and host groups

Operations for provisioning and protecting storage — including block volumes, file shares, and virtual machine datastores — are abstracted from Data Ops Manager. These functions are managed exclusively from their specific applications (that is, Block Storage, File Storage, or Private Cloud Business Edition). However, you can view block and file storage performance and capacity from Data Ops Manager as part of the detailed system monitoring.

Dashboard

The Data Ops Manager Dashboard provides at-a-glance summaries of the fleet and enables easy navigation to systems, volumes, and hosts.

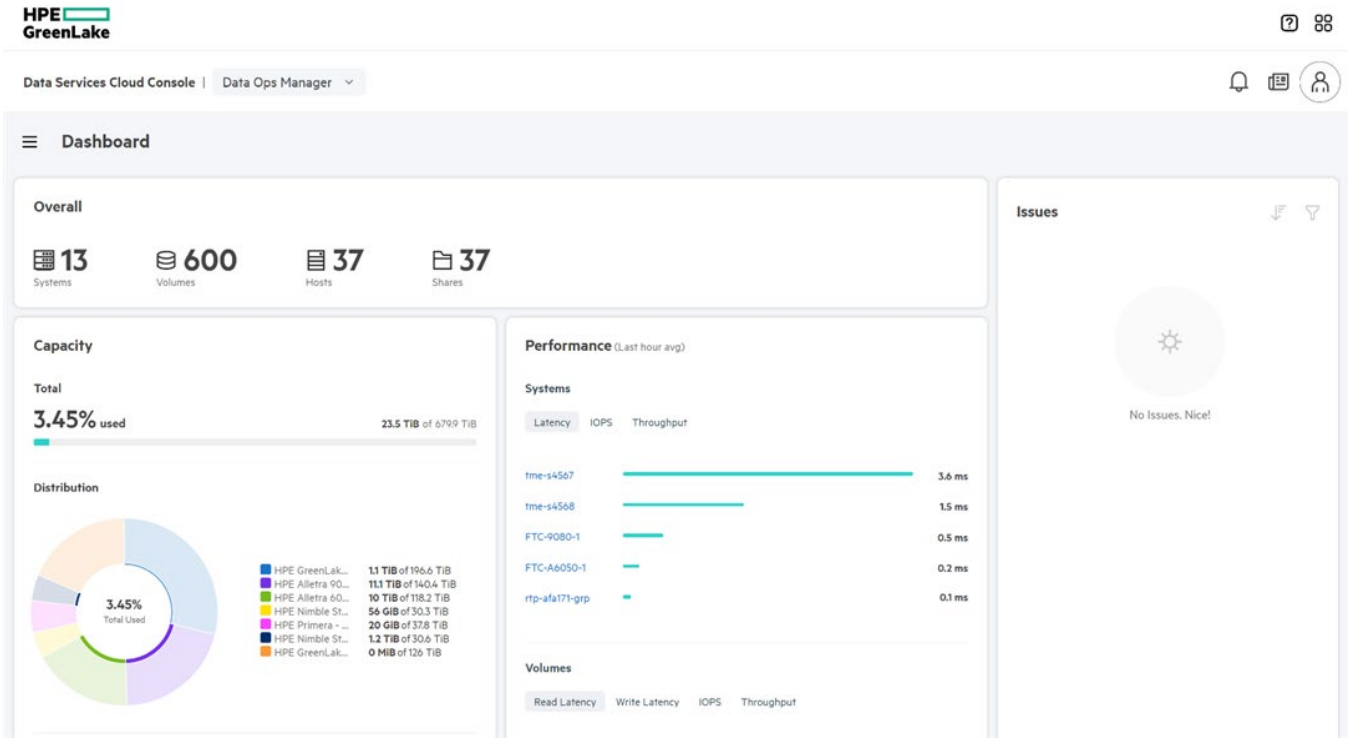

Figure 5. Data Ops Manager dashboard

The dashboard includes the following information:

- The **Overall** panel shows the total number of storage systems being managed, the total number of volumes across all those systems, the number of hosts defined across those systems, and the total number of file shares that have been created.
- The **Capacity** panel shows the physical capacity of the systems in the data center and the provisioned capacity of the volumes across those systems.
  - Within the capacity panels, you can see systems, volumes, and file shares respectively under a range of percentage-full categories: 90%–100% full, 75%–89% full, 50%–74% full, and less than 50% full.
  - The top five systems, volumes, and shares in each percentage full category are displayed.
- The **Performance** panel captures average performance from the last hour for systems and the volumes on those systems.
  - Performance statistics are split out by latency (overall latency for systems, and read/write latencies for volumes), IOPS, and throughput.
  - The top five systems/volumes are displayed for each given performance statistic.
- The **Issues** panel lists reported issues across all systems with the description, severity, category, and date for each issue. Issues shown are specific to those in which a storage administrator is typically interested, such as failed hardware components. The list is sortable by date and filterable by severity, enabling you to easily track unresolved issues over time and filter out less severe issues. You can view the details of each individual issue as well, which provides a more detailed description of the problem and, when applicable or available, possible resolutions.

Each entity on the dashboard (whether system, volume, share, or host) is clickable, enabling easy navigation to more detailed information on a given system or volume.

Systems page

The Systems page within Data Ops Manager provides a complete list of the systems being managed in DSCC. Under the All Systems tab, you can see the name of each system as well as the platform, capacity, and average performance statistics (latency, IOPS, and throughput) from the last hour. You can easily search for systems by name or filter the list according to array type, capacity, or performance. This capability helps you quickly find the systems you are looking for, even if you have hundreds of arrays being managed in DSCC. You can also export the displayed system table in CSV or PDF format.

In the Updates tab on the same page, you can view all the systems that are currently eligible for software updates, the nature and priority of the software updates available, the status of the readiness checks on the system, and the installation progress.

Clicking a system name in the All Systems list or the Updates list takes you to the detailed page for that system, where you can view and perform various functions on a specific system, including:

- View and sort system-specific issues by issue, severity, category, and date/time.
- See current system capacity usage, including overprovisioning ratio and savings ratio, capacity usage on a per-application basis, and capacity trends over the past week, month, or a custom date range.

  A capacity report for the selected range can be downloaded as a CSV or PDF.
- Configure, modify, or remove replication configurations.
- View performance trends for the system over the past eight hours, day, week, month, or a custom date and time range.

  You can customize the performance metrics displayed here based on resource type and metric and export the custom performance report as a CSV or PDF.
- View schematic-level details and the health of specific hardware components such as controller nodes, ports, drive enclosures, and drives.
- View the current software and firmware versions currently on the system as well as the update history and apply available updates.
- Configure and modify the networking settings including management IP, DNS, and proxy servers, and data network interfaces such as Fibre Channel, iSCSI, and replication ports.
- View and modify system settings, alert and monitoring options, and date/time settings.

Performance

Improvements have been made to performance monitoring in recent updates to DSCC and Data Ops Manager. These improvements include the ability to view a custom performance dashboard featuring specific resources versus just the default performance dashboard, which offers a system-wide overview of latency, IOPs, and throughput. The resource types that can be selected depend on the storage system model.

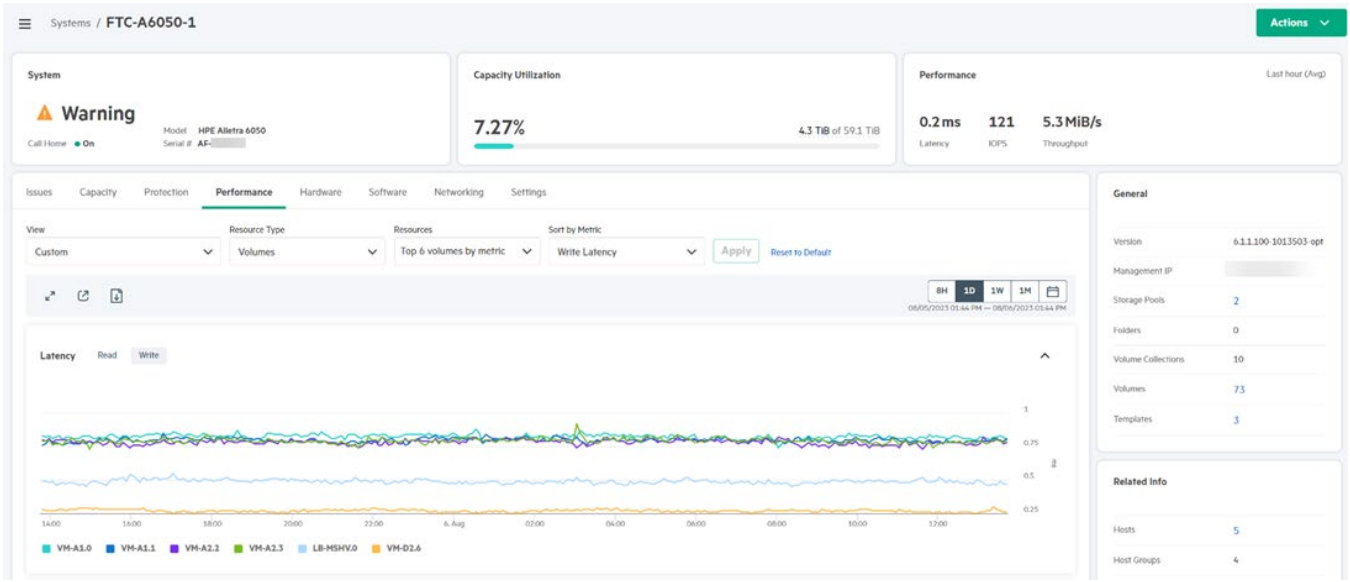

Figure 6. Customized performance dashboard view

For HPE GreenLake for Block Storage systems powered by HPE Alletra Storage MP, advanced performance metrics — including headroom utilization and volume hotspot — have been implemented in the system performance dashboard. Headroom utilization analyzes the system's saturation level and the ability to handle more workload without compromising response times. The top five volumes contributing to the headroom utilization are featured. The higher the headroom utilization, the closer a system is to reaching its limit and potential performance degradation.

The volume hotspots graph shows the volumes that are under more stress as compared to other volumes in the storage system. It measures latency on a per-volume basis to identify volumes with anomalous latencies and displays a heatmap of those anomalies. This graph can help you quickly and efficiently identify performance bottlenecks and resource constrained volumes on your systems.

Figure 7. Headroom utilization and top volume hotspots in the HPE GreenLake for Block Storage performance dashboard

Note
Intent-based provisioning uses headroom utilization when making workload placement recommendations. By following the system recommendations from intent-based provisioning, you can avoid provisioning new workloads on overutilized systems. To learn more, see the HPE GreenLake for Block Storage and Intent-based provisioning sections of this paper.

Software updates

DSCC enables simple software management for a fleet of arrays, bringing customer self-update capabilities to the cloud. In Data Ops Manager you can easily see all the arrays in your fleet that are eligible for updates, perform readiness checks, and install the software updates from a single console.

The information for an available update includes:

- **Version:** The version number of the available update
- **Severity:** The severity category: available, critical, and recommended
- **Type:** The type of software update: OS, Standard Support Release (SSR), Extended Support Release (ESR), drive firmware, and enclosure firmware
- **Readiness checks:** The state of the readiness checks that are run before installing the update: Not checked, ready, warning, and failed

  You can click a failed check to find more information about why the readiness check failed.
- **Time:** The estimated length of time for installing the update
- **Update path:** The sequence of versions required to update from the current version to the target version

The one-step software update feature simplifies the update process by automatically downloading and staging applicable software updates to the arrays and performing the readiness checks. This feature allows software updates to be initiated in one quick and easy step. After a software update has been initiated, the detailed installation progress can be monitored from the system software dashboard.

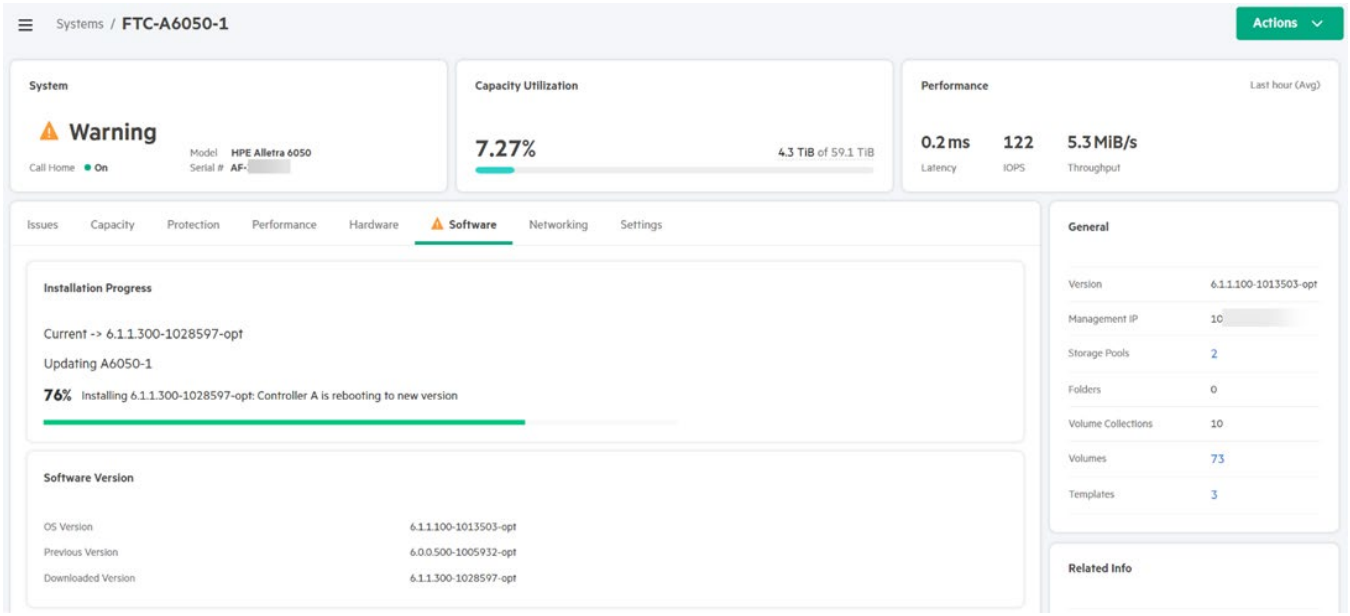

Figure 8. Monitoring an in-progress software update in Data Ops Manager

Protection

In order to protect volumes with remote replication, you must first configure replication between the storage arrays in Data Ops Manager. Replication configurations can be configured between like platforms on supported operating system versions as well as between the following:

- HPE Alletra 9000 and HPE Primera arrays
- HPE Alletra 5000/6000 and HPE Nimble Storage Gen5 arrays

Note

To configure Peer Persistence on a pair of HPE Alletra 5000/6000 or HPE Nimble Storage systems in DSCC, you must first group the systems by using **Actions** > **Add Array to System**. For more information about Peer Persistence with HPE Alletra 5000/6000 and HPE Nimble Storage, refer to the HPE Nimble Storage and HPE Alletra 6000 Peer Persistence deployment considerations paper.

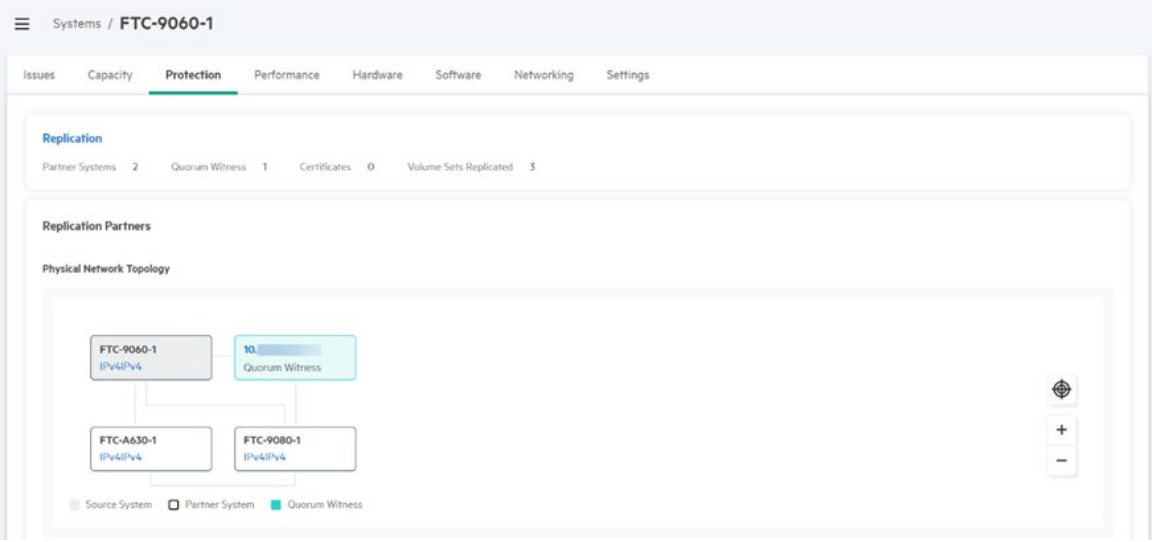

Figure 9. Protection and replication summary page in Data Ops Manager

Use the replication summary on the Systems page to view your replication partners and their volume set information related to replication. The replication summary gives a brief status, and the replication type of the volume sets and can be exported as a CSV or PDF.

Data Access page

The Data Access page in Data Ops Manager enables you to create and manage hosts and host groups that are connected to the storage arrays in the fleet.

- **Host group:** Host groups are logical entities consisting of one or more hosts. In DSCC, storage access is provided through host groups, not individual hosts. This means that hosts must be part of a host group to export volumes to that host through DSCC. Host groups can contain a single host for stand-alone host requirements, and one host can belong to more than one host group, if appropriate. Hosts that are part of a host group should have appropriate application clustering software or a clustered filesystem managing the exported volumes to ensure data consistency and prevent host-caused data integrity problems.
- **Host:** A host represents a physical server in DSCC and is defined by access protocol, operating system, and initiators. Access protocols include FC, iSCSI, and NVMe. Initiators can be selected from a list of available and accessible initiators for the protocol selected or can be created by manually specifying the WWN/IQN/NQN. Hosts must be created as part of a host group in DSCC.

Data Ops Manager simplifies host creation and management by enabling you to create and manage host definitions that apply to every array in the fleet that is connected to the host initiators. A single host group definition can be used to manage the exports of volumes from multiple arrays in the fleet, rather than requiring separate host definitions for each individual system.

Hosts and host groups that have been created on a system outside of DSCC (such as through an array's onboard UI or CLI) display in Data Access with the system name appended in parentheses. You can conveniently merge identical host definitions existing on different systems together so that they display as a single entry in the Data Access page in DSCC.

HPE GreenLake for Block Storage

Offering self-service storage provisioning and protection that meets your workload needs, the HPE GreenLake for Block Storage application focuses on storage lifecycle management for block provisioning, native replication, and snapshot-based data protection and recovery. Block Storage makes it quick and easy to provision, protect, and manage application volume sets.

The intended users for the Block Storage application are cloud consumers: application admins and users who need to provision and present storage to their application hosts but who do not necessarily require access to or have knowledge of the underlying storage arrays in the fleet. By abstracting storage infrastructure management from storage provisioning and protection, Block Storage enables application owners to manage and protect their own block volumes without the assistance of a dedicated storage administrator.

Block Storage separates out the application administration features from Data Ops Manager, leaving only the infrastructure administration in Data Ops Manager.

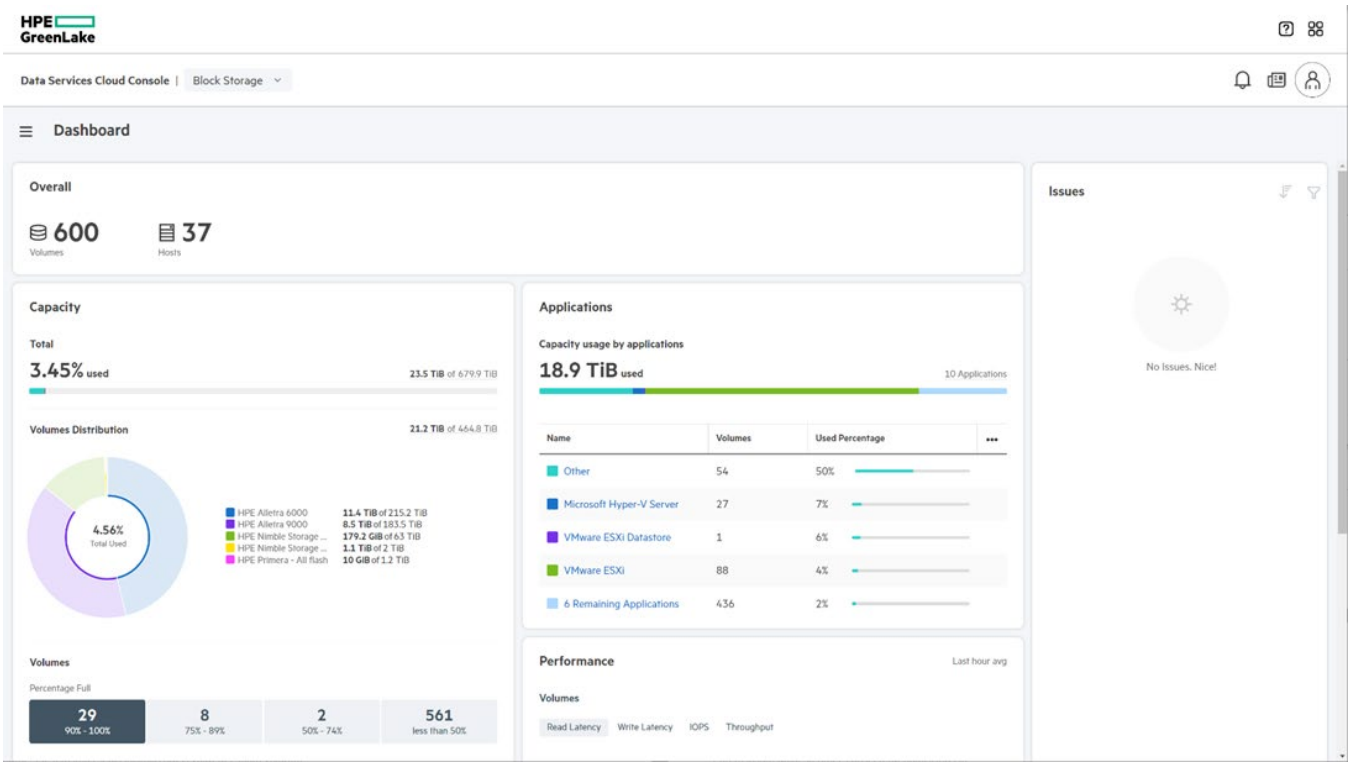

Figure 10. The Block Storage dashboard

You can use Block Storage to:

- Monitor volume and host information, volume issues, and capacity usage per application in an integrated dashboard
- Create, edit, and delete volumes
- Add or remove volumes from existing volume sets
- Provision block storage by storage tier through intent-based provisioning
- Apply predefined or user-defined protection policies to volume sets
- Define custom protection for existing volume sets including local snapshot schedules or remote protection with replication

Intent-based provisioning

Intent-based provisioning enables self-service provisioning of storage tailored to your specific requirements, including:

- **Storage tier:** Select the storage tier that best suits the SLA for the application.
- **Workload type:** Select from a predefined set of common workloads or create a custom type.

  The workload type governs data reduction settings for the volumes.
- **Volume details:** Provide a prefix for the volume names and specify the number and size of the volumes that will be in the set.
- **Host group access:** Select the host group to which the volumes will be exported.
- **Protection policy:** Choose from a predefined list of protection policies, including local snapshots taken at daily or 15-minute intervals, asynchronous replication, synchronous replication, and high-availability Peer Persistence.

  Protection policies for a volume set can be altered or applied later.

Note
The protection policies that are visible in the drop-down list depend on which replication types are available on the systems in that storage tier. For example, if you have only a synchronous replication target configured for a system, you will not see asynchronous replication appear as an option when you create or protect volumes on that system.

HPE GreenLake for File Storage

The HPE GreenLake for File Storage application enables you to create and configure NFS file shares on file clusters that are configured and managed within DSCC.

The entire life cycle of a share and its settings are managed from File Storage. File Storage provides administrators with a unified dashboard with information that relates to the functioning of the provisioned shares. It also enables administrators to manage the life cycle of local users and groups on specific file clusters. All alarms and events on the provisioned shares are shown on the dashboard as well as on the specific share details page.

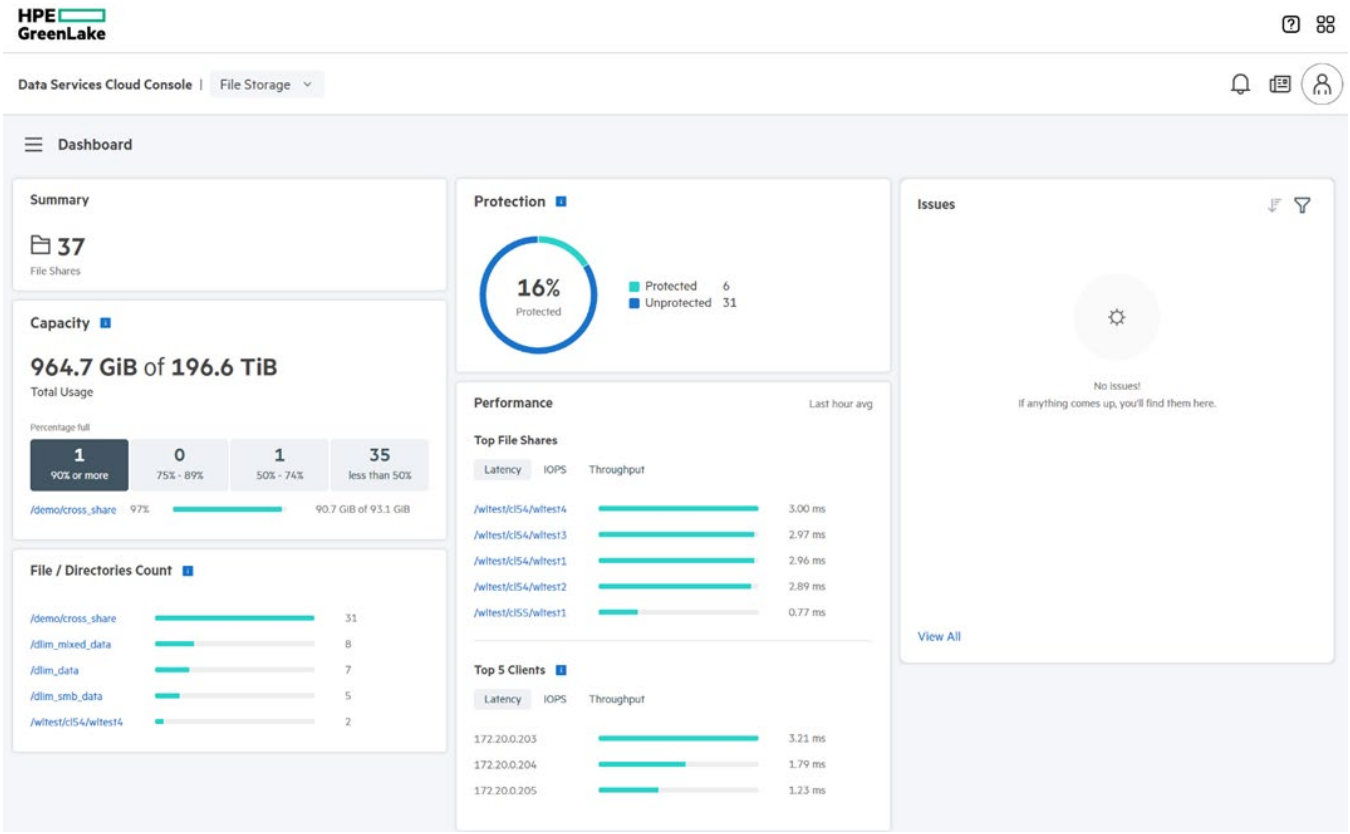

Figure 11. The File Storage dashboard

Features available with HPE GreenLake for File Storage include the ability to:

- Create and manage file shares and file share settings.
- Create and manage local users and groups.
- Create and apply protection policies to file shares.
- Manage how cluster capacity is used by setting quotas for capacity and users/groups.

For more information about HPE GreenLake for File Storage including features and architecture, see the HPE GreenLake for File Storage architecture technical paper.

Private Cloud Business Edition

The HPE GreenLake for Private Cloud Business Edition application extends HPE Alletra dHCI into the hybrid cloud. It provides a single interface for deploying and managing HCI systems across an entire organization. You can also link your public cloud provider accounts so that you can provision and manage those virtual machines from the same HPE interface.

The dashboard provides a global view of all dHCI systems across your organization and includes:

- A summary of public cloud and private cloud virtual machines and private cloud systems
- An overall view of private cloud virtual machine performance, capacity, and utilization
- A hot spot map that shows the top systems consuming the most resources
- A view of the virtual machines that are protected and the ones that are not
- The latest configuration checks report for all private cloud systems

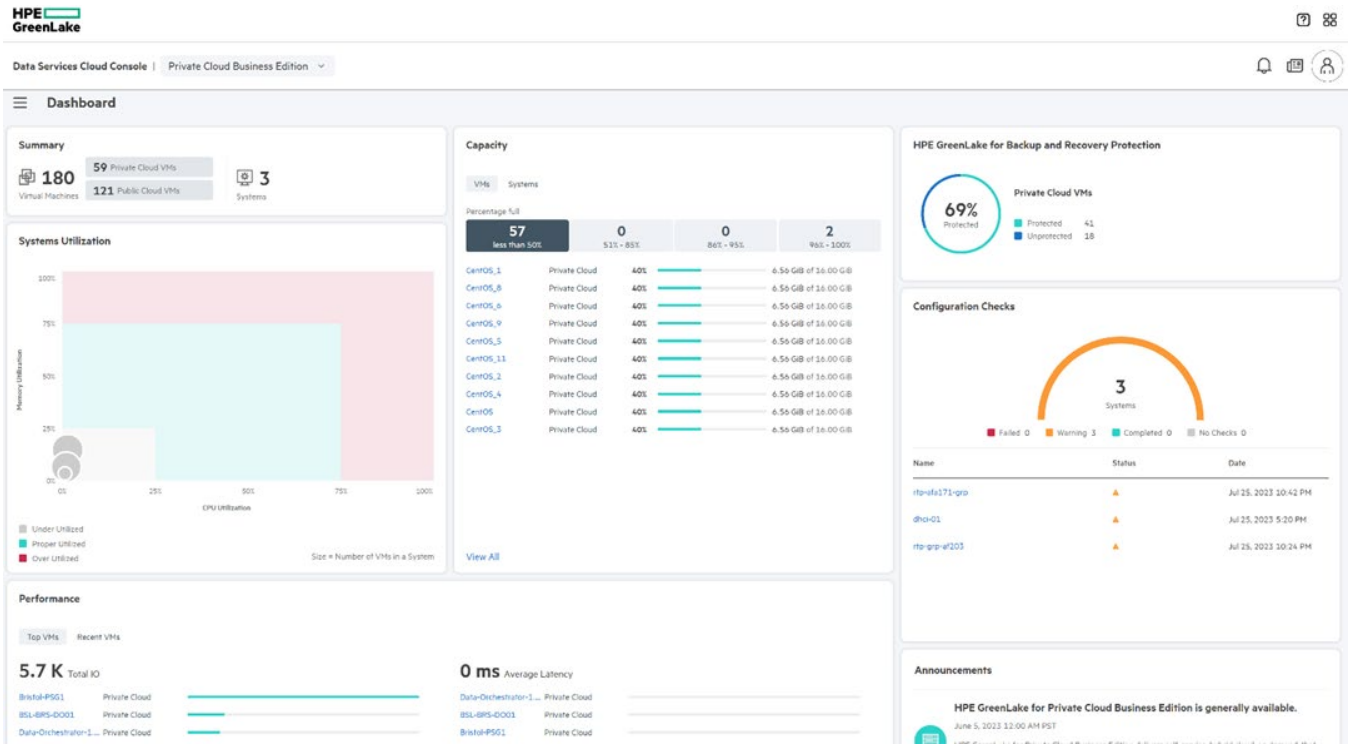

Figure 12. The Private Cloud Business Edition dashboard

HPE GreenLake for Private Cloud Business enables simple and efficient end-to-end management of resources in your hybrid cloud environment, including:

- **Cluster life cycle management:** The systems page enables automated cluster lifecycle management of the full hardware and software stack of dHCI clusters from a single interface. This includes precheck reports and online updates of the server VMware ESXi™ images, the HPE Alletra 5000/6000 or HPE Nimble Storage OS, the server firmware, and host multipath software.
- **Virtual machine provisioning and management:** Enables you to create and provision private cloud virtual machines by using provisioning policies. These provisioning policies include specifications for storage type, deduplication, encryption, and the QoS performance limits. If you have a subscription to HPE GreenLake for Backup and Recovery, you can also add protection policies to your provisioning templates.
- **Cloud account management:** Enables you to link AWS public cloud accounts so that you can provision and manage your Amazon EC2 instances from the same interface as your on-premises environment.

HPE GreenLake for Backup and Recovery

HPE GreenLake for Backup and Recovery is a data protection service that is part of the DSCC application on the HPE GreenLake platform. The service allows you to enable automated data protection for on-premises VMware resources by using prebuilt and customizable protection policies to create VMware consistent immutable recovery points such as local snapshots and backups and cloud backups to managed Amazon Web Services (AWS) S3 stores. You can also back up and restore AWS EBS volumes and EC2 instances.

This service provides unified management through a single cloud console across multiple underlying storage systems in different locations and several backup destinations. It also features a comprehensive dashboard that displays the overall status of your protected resources both on-premises and in the cloud.

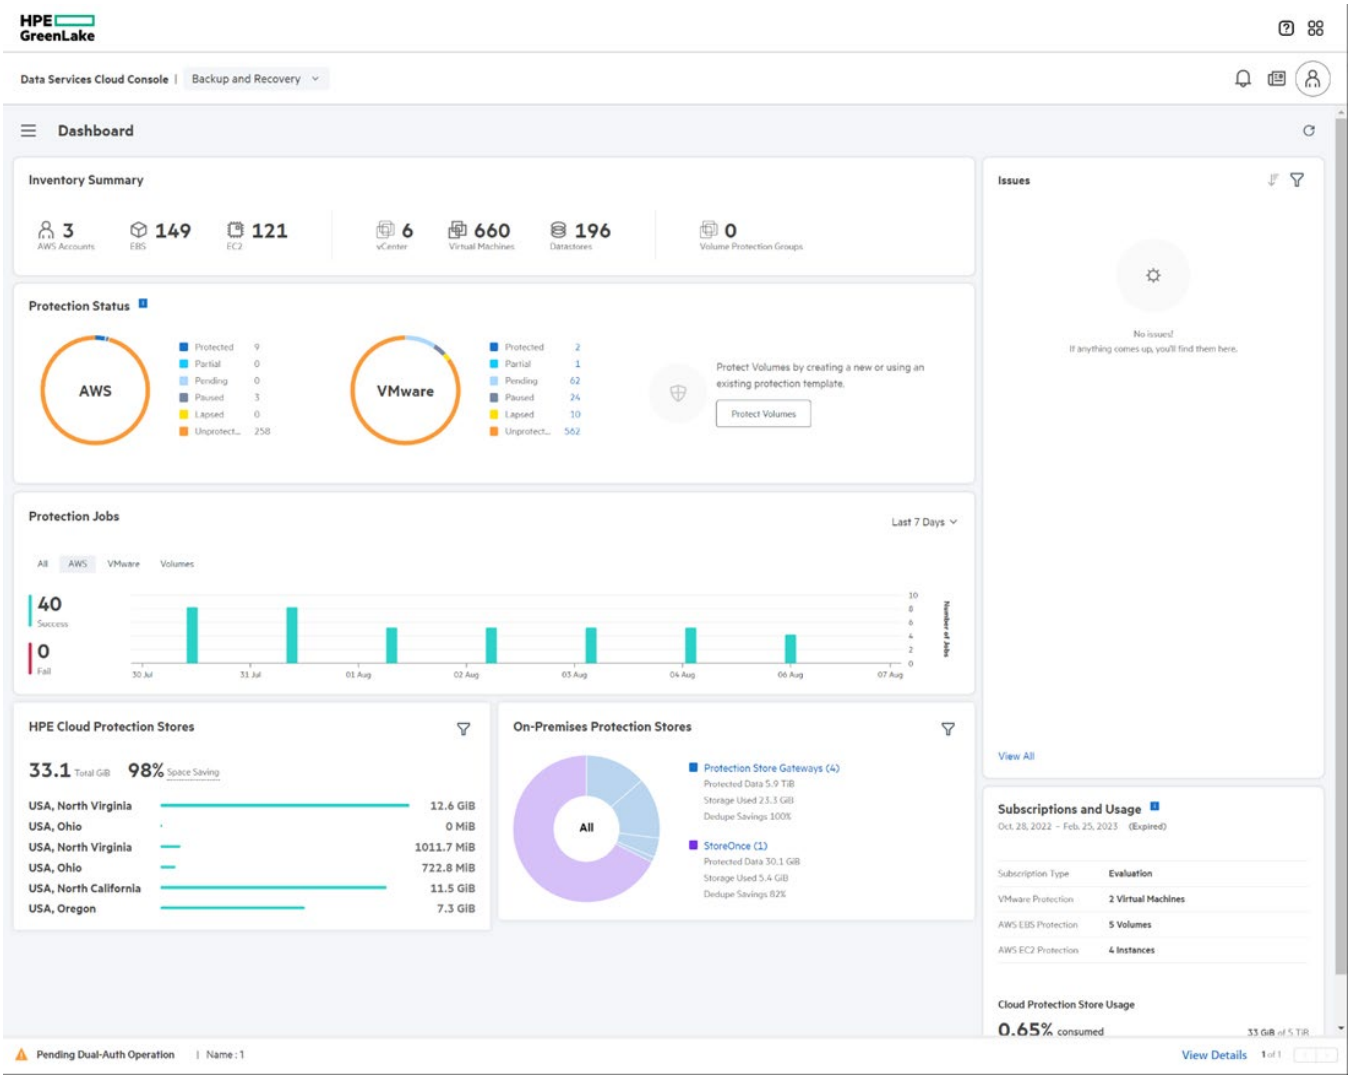

Figure 13. The HPE GreenLake for Backup and Recovery dashboard

Features in HPE GreenLake for Backup and Recovery include:

- Block-based application-consistent snapshots, array-optimized backups, and cloud backups for volumes coming from DSCC-managed storage arrays such as HPE Alletra, HPE Primera, and HPE Nimble Storage
- Storage-agnostic snapshots and backups leveraging VMware Change Block Tracking (CBT), enabling you to protect local datastores and datastores from non-DSCC managed and third-party storage arrays in your VMware environment
- HPE StoreOnce Catalyst backups of protected resources to HPE StoreOnce appliances
- The ability to link your AWS public cloud accounts and protect EBS volumes and EC2 instances
- The ability to protect and restore HPE array volumes that were created in the Block Storage service

- Fully customizable protection policies that enable scheduled snapshots and backups to multiple targets and the ability to make any snapshot or backup immutable.

For more information about the architecture and the latest features that have been made available for HPE GreenLake for Backup and Recovery, refer to the HPE GreenLake for Backup and Recovery technical paper.

HPE GreenLake for Disaster Recovery

HPE GreenLake for Disaster Recovery is an application in the HPE GreenLake Data Services Cloud Console that focuses on enabling the replication of mission-critical applications and data as quickly as possible with minimal data loss in VMware virtualized environments. The service makes it easy to manage, orchestrate, and automate disaster recovery operations across multiple logical groupings of virtual machines.

HPE GreenLake for Disaster Recovery includes the following features:

- A detailed dashboard highlighting a resilience score for the environment, which is calculated based on metrics pertaining to protection groups, including overdue failover tests, percentage of VMs protected, breached RPOs, and breached journal SLAs
- The ability to create customized protection groups, which enable disaster recovery, and continuous journal-based replication to a local or remote VMware vCenter®

  Select the VMs you wish to protect and customize the recovery options, including the recovery host, datastores, failover and failover test network, and recovery folder to create the protection group.
- Live failover and test failover capability, which enable you to select a specific recovery point and fail your protection groups over to the disaster recovery site
- The ability to pull a recovery report, which generates a detailed report for the protection groups managed by your Virtual Manager Appliances and any DR operations that have been executed

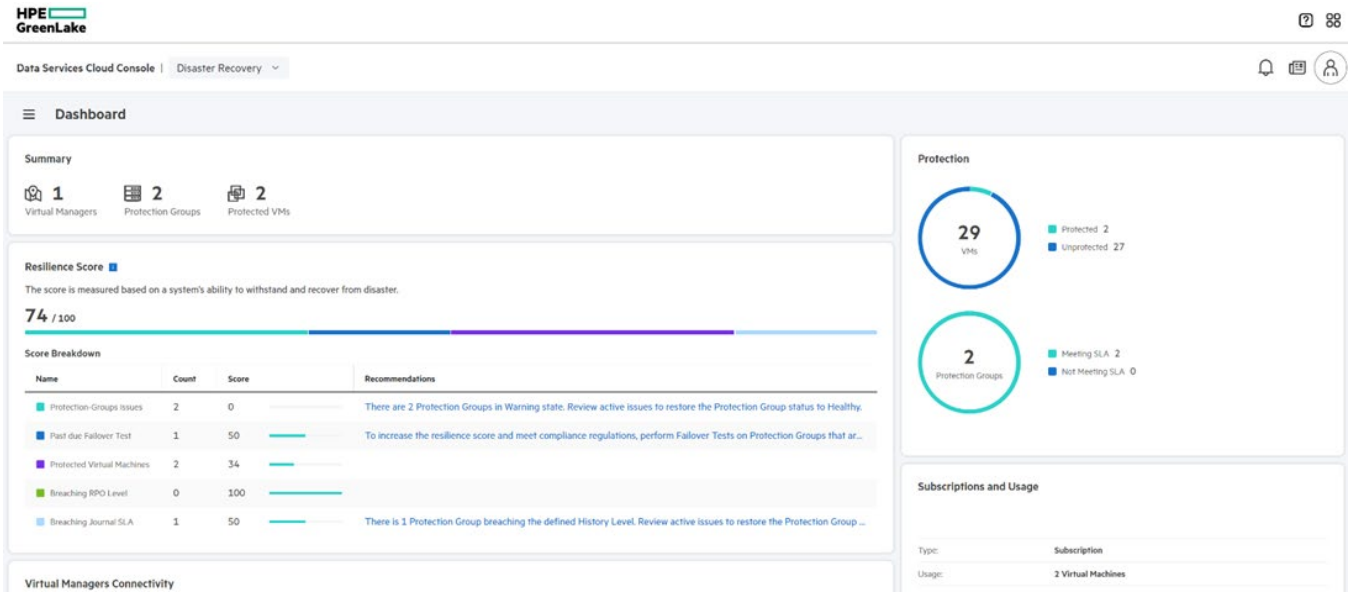

Figure 14. The HPE GreenLake for Disaster Recovery dashboard

With HPE GreenLake for Disaster Recovery, you can quickly recover your data to a historical point before any disruption, enabling you to resume operations following ransomware attacks or other outages.

HPE GreenLake for Storage Fabric Management

HPE GreenLake for Storage Fabric Management provides cloud-based management of Fibre Channel storage fabrics integrated with your DSCC-managed storage arrays. It accelerates configuring, monitoring, and managing your storage networking fabrics while providing advanced automation and orchestration capabilities.

The Storage Fabric Management dashboard gives a health summary of your managed environment, including datacenters, managed device groups, and the switches in those groups.

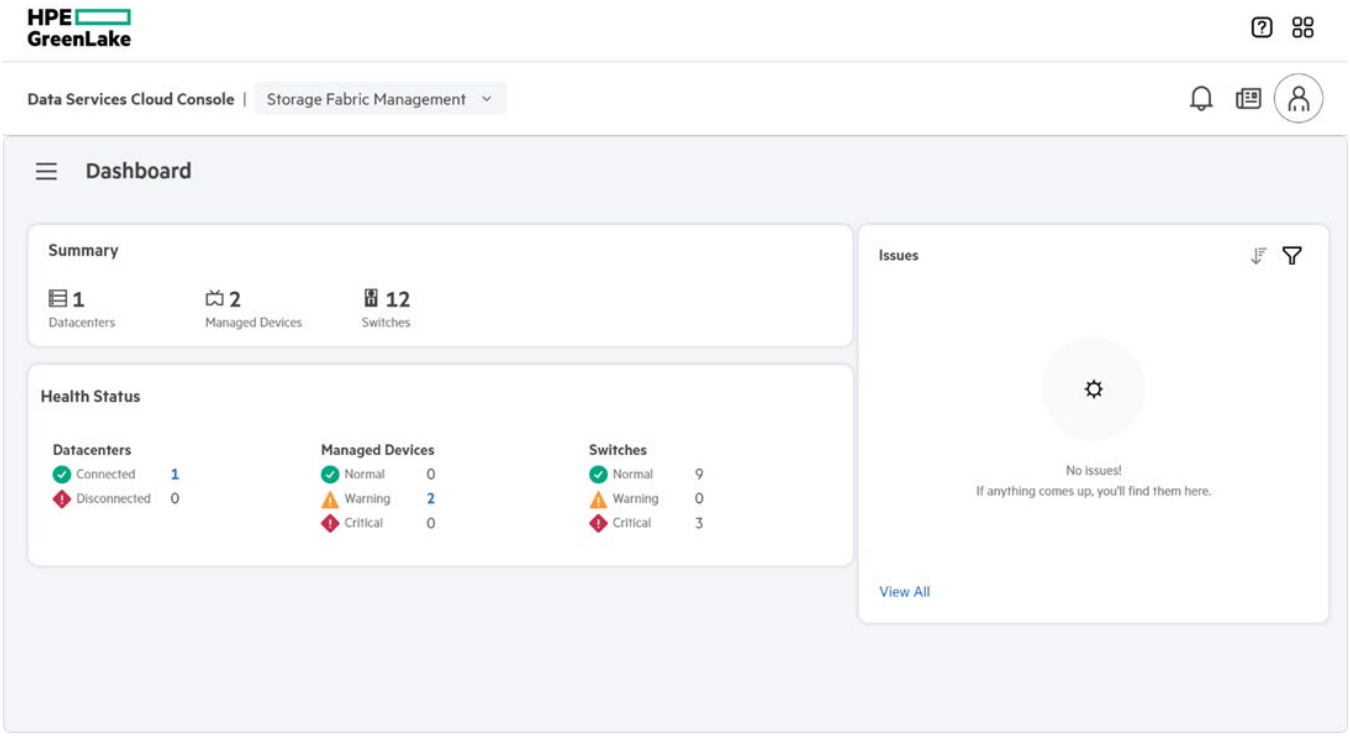

Figure 15. The HPE GreenLake for Storage Fabric Management dashboard

HPE GreenLake for Storage Fabric Management includes the following features:

- Creating FC aliases and zones, and activating zone configurations
- SPOCK and topology validation across all managed devices
- Applying firmware upgrades to managed switches
- Editing physical port properties on the switch, including enabling/disabling the ports as well as license management for ports

Other features

The DSCC tile screen also includes shortcuts to additional features that are not tied directly to a specific application but provide insights into or enable a feature for the entire environment managed within DSCC. These features include Audit logging, dual authorization, secrets, and task monitoring.

Audit logging

Audit logs are an essential security tool for providing records of all events and changes that occur within a system or environment. The DSCC audit log service provides a comprehensive audit trail to assist in monitoring potentially sensitive data or systems for possible security breaches, vulnerability, or misuse of data. It also provides records that serve as evidence in cybersecurity attacks.

See Figure 16 for an example of the audit dashboard, which provides a summary of the events and outcomes over the past hour, eight hours, one day, or one week, depending on the timeframe selected. Under the dashboard, you can review each of the individual events that have been initiated and the outcomes of those events.

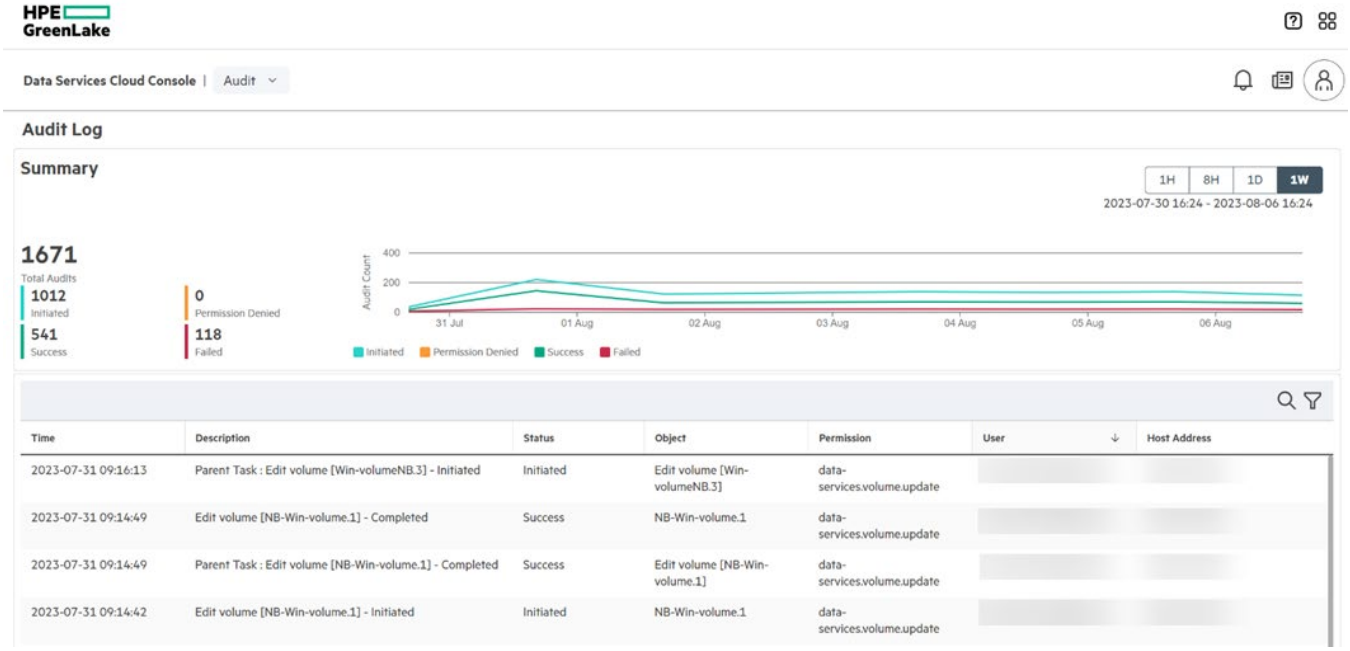

Figure 16. Audit logging in DSCC

Depending on your permissions, you might be able to see events initiated by all users, or you might be restricted to only those events initiated by your user ID.

Dual authorization

Dual authorization is a security process that requires a secondary user to approve requests for certain destructive operations such as deleting resources or disabling features. It can be enabled with HPE GreenLake for Backup and Recovery to require authorization from a secondary user with administrator privileges to delete snapshots and backups.

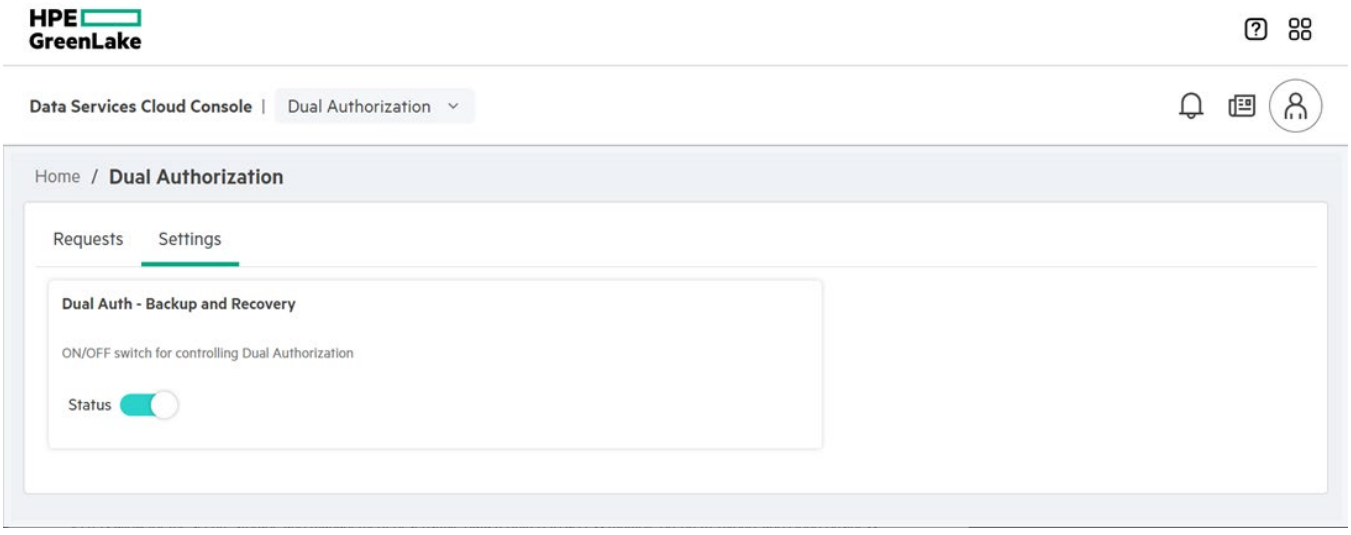

Figure 17. Dual Authorization settings in DSCC

When dual authorization is enabled, users cannot approve their own requests. A secondary authorization is required to disable dual authorization in DSCC.

Secrets

Secrets enable the secure storage and management of sensitive data required to access multiple on-premises, hybrid, and cloud business applications. Secrets can be used by system administrators to securely manage and maintain configuration properties for business applications across the Data Services Cloud Console platform. Secrets can be sets of properties required to access application services, such as credentials or tokens, and may include other data useful for configuring applications and their operations. Secrets data is encrypted, and the encrypted data cannot be returned through the GUI or API.

Secrets can be created for basic authentication, bearer authentication, SSH keypair, certificate, JSON web token, private key, password, and generic types.

Tasks

The tasks page keeps a running log of all tasks initiated and details on the result. Each application has its own task page for monitoring and seeing status details on tasks initiated from that application. The master task list on the DSCC tile page displays tasks for all applications.

Tasks can be sorted and filtered for monitoring on fields, including:

- Status
- Subtasks
- Service or application to which the task pertains
- User who initiated the task
- Date/time range during which a task occurred

Conclusion

Modernizing data storage, access, and protection is imperative for keeping pace with advancing applications, infrastructure, and cyber threats. Architected as a single destination to streamline data management and mobility across edge to cloud, DSCC on HPE GreenLake helps organizations accelerate data-driven innovation and manage the lifecycle of data from anywhere. Self-service access to data infrastructure eliminates the traditional silos and complexity of data management and delivers operational simplicity from a single console. DSCC offers a modernized operational experience that scales with your organization, saving you time and money, and freeing up IT resources so that your teams can focus on innovation.

Resources

HPE Welcome Center

Data Services Cloud Console Getting Started Guide

Data Services Cloud Console Security Guide

HPE GreenLake for Block Storage Architecture

HPE GreenLake for File Storage Architecture

HPE GreenLake for Backup and Recovery

Learn more at

hpe.com/storage

Explore HPE GreenLake

Chat now (sales)

a00124553ENW, Rev. 1

Hewlett Packard Enterprise