


HPE GreenLake for Backup and Recovery — Backup-as-a-Service built for the hybrid cloud

A data protection service for on-premises and cloud native workloads



Contents

Executive summary.....	3
Service overview	3
Key benefits of HPE GreenLake for Backup and Recovery.....	4
HPE GreenLake for Backup and Recovery architecture.....	6
Cloud components (AWS and on-premises protection).....	6
On-premises components (on-premises protection only).....	7
HPE GreenLake for Backup and Recovery in use.....	9
Operations	10
Summary.....	17
Appendix A: Storage efficiency test environment.....	18
Resources.....	19



Executive summary

The amount of data is massively exploding in every IT organization and protecting it is as complex, expensive, and challenging as ever. Traditional data protection solutions often grapple with the intricacies of hybrid IT environments. The vast data volumes, diverse data types, increasing security threats, and stringent regulatory mandates surpass their capabilities. They often struggle with scalability and flexibility, failing to provide comprehensive protection, slow recovery times, high operational costs, and complexities in managing backups across multiple platforms. Thus, these solutions are often inadequate for modern businesses that prioritize simplicity, scalability, and efficiency. The ever-evolving digital landscape necessitates a robust, agile, and self-managed approach to data protection.

Delivered through [HPE GreenLake Data Services Cloud Console](#), [HPE GreenLake for Backup and Recovery](#) provides comprehensive data protection for workloads across hybrid cloud — eliminating data silos, multiple administrative touch points, and cumbersome point solutions. The service leverages global Protection Policies for consistent protection of on-premises and cloud native workloads in a simple and efficient manner via a single SaaS console. It brings policy-based automation to protect different enterprise workloads in a few simple steps — within minutes — eliminating the complexities of managing your backup and recovery operations on-premises or in the cloud. There are no additional cloud gateways, agents, backup software, proxies, media servers, or backup targets to manage, and cloud storage is fully managed and scaled automatically by the service. Threats like ransomware and malware are neutralized with built-in encryption, data immutability, dual authorization, and flexibility to store backup copies in an air-gapped manner making them inaccessible to cybercriminals. Organizations lower the cost of protecting data on-premises or in the cloud with consumption-based pricing and ultra-efficient data reduction technologies.

Target audience: Pre-sales consultants, solution architects, IT and storage administrators responsible for the protection of virtual workloads.

Document purpose: This paper describes how HPE GreenLake for Backup and Recovery is differentiated from other solutions by offering data protection with unmatched storage efficiency and ease of use.

Service overview

HPE GreenLake for Backup and Recovery is a backup service designed for the hybrid cloud and delivered through the [HPE GreenLake edge-to-cloud platform](#). It provides a solution for protecting on-premises and cloud native workloads with SaaS simplicity and best-in-class storage efficiency. The data sources and backup targets are illustrated in Figure 1.

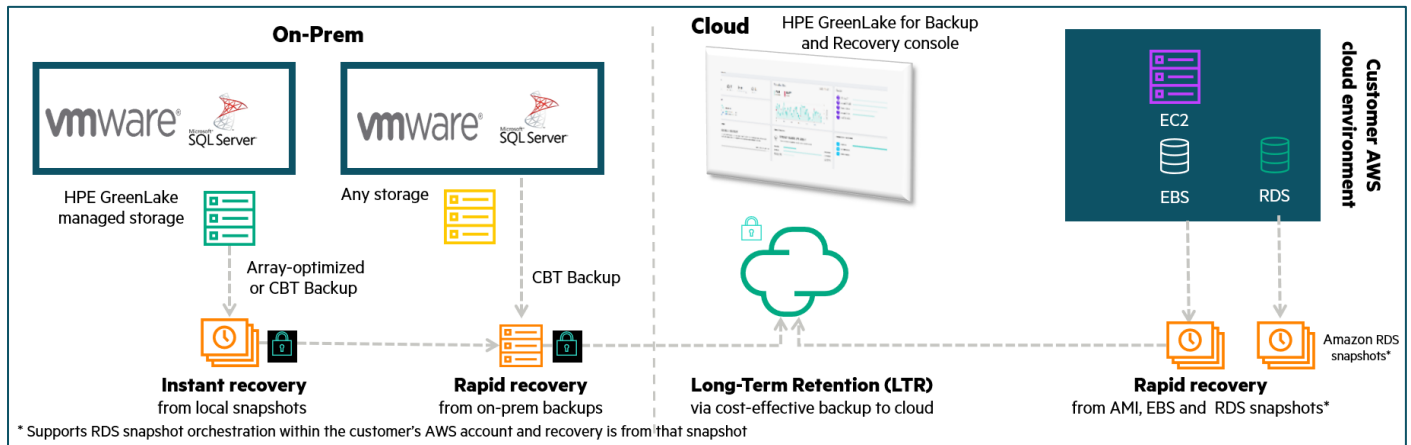


Figure 1. HPE GreenLake for Backup and Recovery provides hybrid cloud data protection for on-premises and cloud native workloads

[Protection Policies](#) provide a simple, flexible, and automated approach for multi-tier protection of business-critical data. The extent of protection that can be configured depends on the resource that is protected.

The on-premises resources (VMware®, MSSQL, HPE Array Volumes) can be protected with:

- Array Snapshots (local/replicated) for instant recovery of data hosted on [HPE GreenLake-managed storage](#)
- On-premises backups (+replicas), in an [On-Premises Protection Store](#) (Protection Store Gateway or HPE StoreOnce System) for rapid recovery of data hosted on HPE GreenLake-managed storage or [Any storage](#)
- Cloud backups in the [HPE Cloud Protection Store](#) for cost-effective, Long-Term Retention (LTR) of data



The cloud native resources (AWS) can be protected with:

- AWS snapshots (local) for rapid recovery
- Cloud backups in the HPE Cloud Protection Store for cost-effective, LTR of data

Key benefits of HPE GreenLake for Backup and Recovery

Effortless protection to meet your SLAs

The HPE GreenLake for Backup and Recovery console makes it easy to manage backup and recovery operations from anywhere. This cloud native console is designed for simplicity of operation. There is no requirement to be a backup and recovery expert or need training to set up and use the service. At its core is the capability for flexible protection that covers on-premises and cloud resources.

HPE GreenLake for Backup and Recovery runs the service in the cloud and includes the integrated management of backup targets, which significantly reduces the hassle of deploying and managing backup infrastructure. This means you can focus on the protection delivered by the service — not the maintenance and operation of the backup infrastructure. Meeting the protection and recovery commitments to your organization and achieving compliance with data governance regulations are assisted with global Protection Policies and Protection Groups that provide auto-protection of existing and newly created resources. APIs are also available for scripted and automated orchestration of protection.

HPE GreenLake for Backup and Recovery is managed through a cloud native console designed to be easy to use. It can be launched from anywhere with internet access via console.greenlake.hpe.com. Figure 2 is an example of the console.

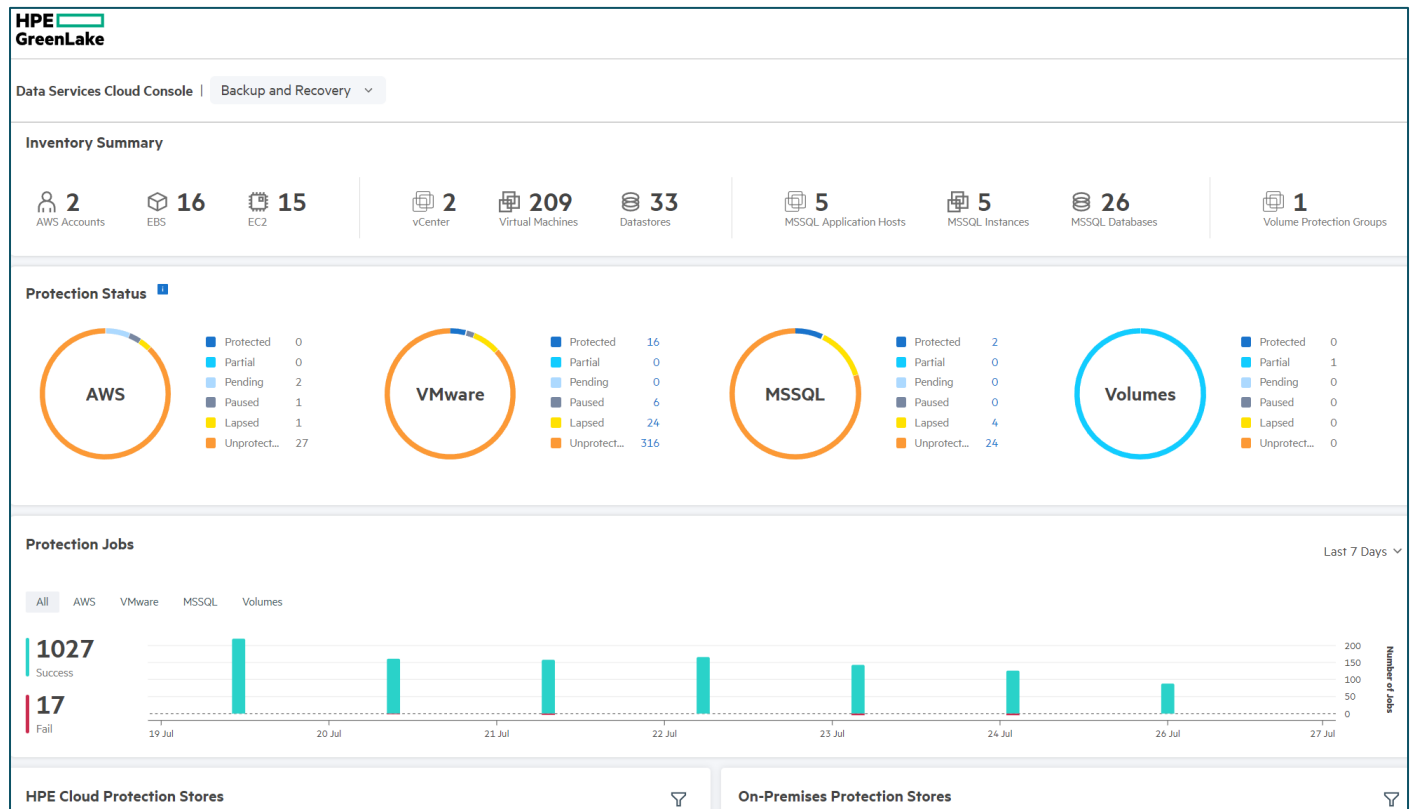


Figure 2. HPE GreenLake for Backup and Recovery dashboard provides a rich view of the overall protection status

Flexible billing, based on usage

HPE GreenLake for Backup and Recovery offers pay-as-you-go¹ usage and billing. Alternatively, the service can be reserved for a committed duration and usage for lower per-unit prices. There are no licenses to manage, just the subscription you choose. Usage-based billing has two billing metrics: the number of protected resources and the capacity of the cloud storage consumed.

¹ May be subject to minimums or reserve capacity may apply.



Elastic scaling provides more capacity when you need it so that resources always match the demand. There is no separate public cloud subscription required for long-term retention and there are no egress charges for recovery.

Unmatched storage efficiency

HPE GreenLake for Backup and Recovery is highly space efficient. It delivers unmatched storage efficiency by using HPE Catalyst technology with small deduplication chunk sizes. Testing has shown that backups consume up to 8x less space than similar solutions so there is less storage to manage, pay for, and move to the cloud for protection. Figure 3 shows the storage efficiency compared to four other solutions that offer similar protection. The relative advantage of HPE GreenLake for Backup and Recovery is clear — enabled by the superior HPE Catalyst deduplication technology.

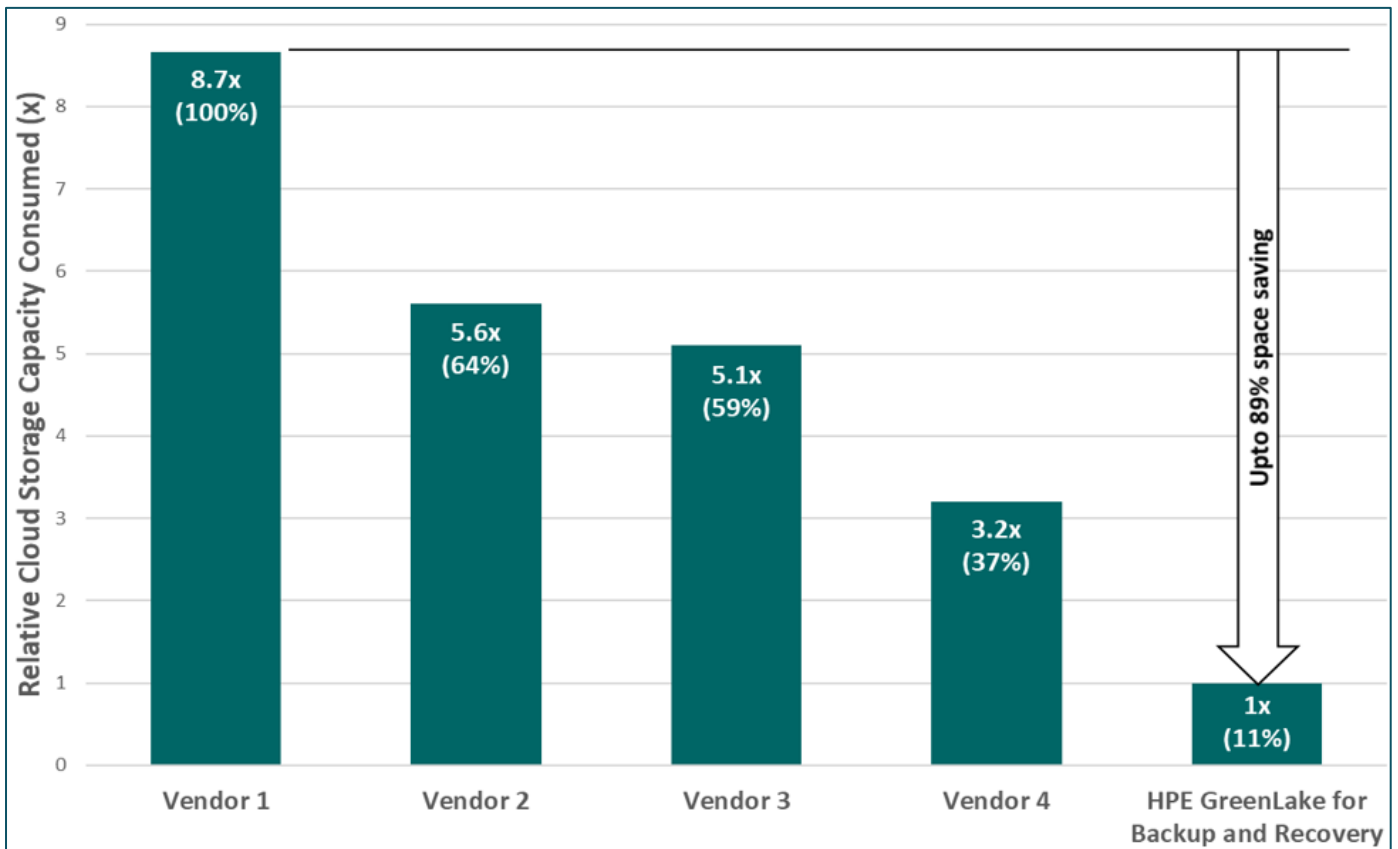


Figure 3. Internal testing showed that HPE GreenLake for Backup and Recovery uses up to 8x less storage capacity than four similar solutions

Note

These test results are based on HPE internal testing conducted in Q2–Q3, 2022 comparing HPE GreenLake for Backup and Recovery with other competing backup solutions. For details, see [Appendix A: Storage efficiency test environment](#).

Secure by design

Ransomware protection is achieved with encrypted and immutable backups. All backups are encrypted for storage and transmission, which make backups unreadable to cyberattacks. Configurable backup data immutability prevents backup data from being modified or deleted by hackers or bad actors. With dual authorization in place, destructive operations require escalation and approval.

Security is built into the HPE GreenLake data services architecture. The Data Services Cloud Console (DSCC) runs in the cloud but is only a management control plane. The data collection is limited strictly to configuration and performance-related data. User data that resides in the protected resources is not exposed to or accessible from the DSCC.

User identity and access for all services are managed centrally by the HPE GreenLake edge-to-cloud platform. Roles and permissions are enforced by using role-based access control (RBAC) to ensure that the correct level of access is given to each user. Roles are used to set user permissions, and scopes are used to limit the resources that each role can access. The account administrator has the option to assign predefined (or custom) roles and scopes provided by the HPE GreenLake edge-to-cloud platform for each service.



The predefined roles include a Backup and Recovery Administrator with full control, or a Backup and Recovery Operator with limited control. All the user activities are audited and monitored for anomalies to help meet security and compliance requirements.

For security details, see [Data Services Cloud Console Security Guide](#).

HPE GreenLake for Backup and Recovery architecture

HPE GreenLake for Backup and Recovery is a cloud native service used for hybrid data protection of resources hosted on-premises or in the cloud. The architecture and data flow are shown in Figure 4. For AWS data protection, no on-premises components are required. For on-premises data protection, the service deploys and manages components in the cloud and on-premises.

For deployment details, see [HPE GreenLake for Backup and Recovery Deployment Considerations Guide](#).

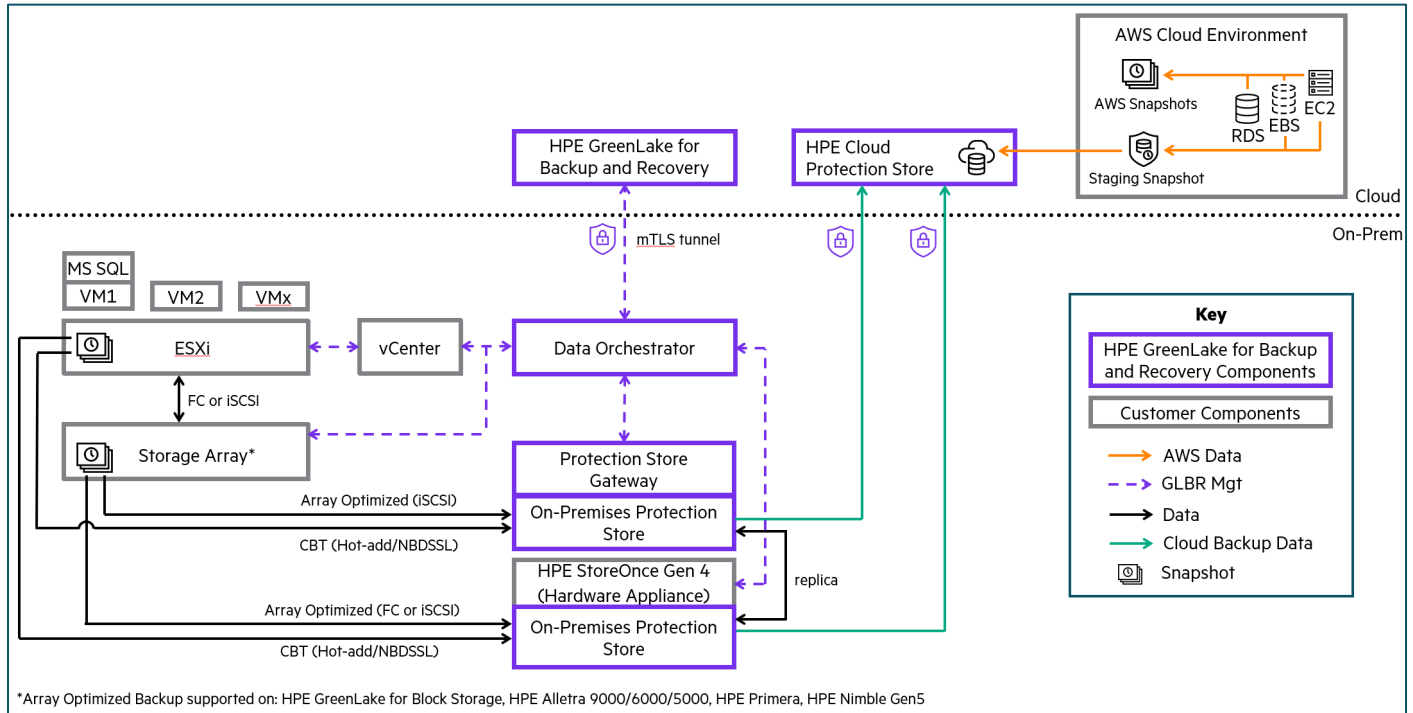


Figure 4. HPE GreenLake for Backup and Recovery provides hybrid data protection with components in the cloud or on-premises

Cloud components (AWS and on-premises protection)

HPE GreenLake for Backup and Recovery is a cloud data service hosted in the [HPE GreenLake edge-to-cloud platform](#). The service removes the management costs of traditional backup and recovery applications and associated storage devices. Multiple sites can be managed from a single cloud console and integration with other HPE cloud services provides a common usage and support experience.

HPE GreenLake for Backup and Recovery components (cloud)

HPE GreenLake for Backup and Recovery

The HPE GreenLake for Backup and Recovery service is flexibly built using a micro-services architecture that enables a high frequency of new features added into the service with zero effort by the user. This is a huge benefit relative to the planning and downtime required to update traditional backup and recovery software.

HPE Cloud Protection Store

The HPE Cloud Protection Store is fully managed by Hewlett Packard Enterprise and provides the destination and capacity for cloud backups. The Cloud Protection Store is automatically deployed by the service in the country/region selected by the user to meet the data sovereignty requirements. Usage is billed based on the capacity used after the cloud backups have been deduplicated and compressed. There are no additional (egress) charges for data recovery.

Other components (cloud)

AWS Cloud Environment

The AWS Cloud Environment is the customer account hosting the EC2, EBS, and RDS resources being protected. AWS snapshots and the latest Staging snapshot, which is the copy sent to the HPE Cloud Protection Store, are stored here.



On-premises components (on-premises protection only)

There are two HPE GreenLake for Backup and Recovery on-premises components deployed as VMware vCenter® VMs: Data Orchestrator (required) and Protection Store Gateway (not required if an HPE StoreOnce System is used for the On-Premises Protection Store). These components are entirely managed and automatically updated by the service. They can be recovered using a guided recovery should they be corrupted or deleted. To further secure HPE GreenLake for Backup and Recovery, only the on-premises Data Orchestrator establishes a connection to the DSCC in the cloud. The Protection Store Gateway never connects to the DSCC, and for cloud protection, it only sends data over an encrypted link to the HPE Cloud Protection Store. These components run a hardened Linux® OS that reduces the attack surface for malware.

For protecting Microsoft SQL Server databases running on on-premises VMware VMs, the service does not require any agent to be installed. The service utilizes the VMware Tools™ installed on the guest operating system to create VSS-aware application-consistent snapshots of the SQL databases.

HPE GreenLake for Backup and Recovery components (on-premises)

Data Orchestrator

The Data Orchestrator manages service operations defined in the cloud. It establishes a mutual Transport Layer Security (mTLS) tunnel to the DSCC Backup and Recovery component in the cloud. This secure tunnel remains active to send northbound management-event data and southbound management requests. The on-premises network proxy or firewall must permit only outbound connections.

On-Premises Protection Store (Protection Store Gateway or HPE StoreOnce System)

Protection Store Gateway

The Protection Store Gateway hosts a service-defined On-Premises Protection Store (any registered VMware vCenter Server® datastore). A Protection Store Gateway processes backup data before it is sent to the On-Premises Protection Store and/or HPE Cloud Protection Store. Communication between the Protection Store Gateway and the console is through the secure Data Orchestrator. The Protection Store Gateway uses the HPE Catalyst engine to deduplicate, compress, and encrypt backup data before it is sent to the On-Premises Protection Store and the HPE Cloud Protection Store. For security, the Protection Store Gateway is deployed as a hardened Linux VM that can be accessed only by the Data Orchestrator. It is locked for any operation by the user or any other application. Multiple Protection Store Gateways can be deployed to add scale and performance. The Protection Store Gateway deployment is flexible to accommodate different sized On-Premises Protection Stores and performance requirements. Sizing guidance is provided during the deployment process within the Backup and Recovery service UI.

Note

[HPE MSA Storage](#) is an ideal choice for provisioning the datastore capacity for the Protection Store Gateway to host the On-Premises Protection Store due to its simplicity, adaptability, and low cost. It offers easy setup, automated tiering for improved performance and cost reduction, and can scale up to over two petabytes of raw capacity. Additionally, its ability to support high-capacity Large Form Factor drives makes it suitable for storing local backup data.

HPE StoreOnce System

An [HPE StoreOnce System](#) is an alternative to host the On-Premises Protection Store. HPE StoreOnce Systems are dedicated, high-efficiency backup appliances. The physical separation of the HPE StoreOnce appliance from the protected VMware environment reduces resource contention and allows scaling on-premises storage without the need to add more VMware storage. Backups to HPE StoreOnce appliances can run over iSCSI or Fibre Channel.

Using an HPE StoreOnce System for on-premises backup storage delivers the same unmatched space saving, encryption, and immutability features as the service-defined On-Premises Protection Stores on a Protection Store Gateway. Any HPE StoreOnce appliance running HPE StoreOnce software 4.3.4 or newer can be connected to HPE GreenLake for Backup and Recovery. Once connected, the service automatically creates an On-Premises Protection Store within the appliance. An HPE StoreOnce appliance can be used by HPE GreenLake for Backup and Recovery and other backup software concurrently.

Other components (on-premises)

VMware ESXi™ and vCenter Servers

VMware ESXi and VMware vCenter Servers host the customer's VM, datastore, and MS SQL Server resources being protected.

Storage (HPE GreenLake managed or any)

HPE GreenLake-managed storage

Includes HPE arrays onboarded into HPE GreenLake, which utilize Array-optimized backups. This backup type uses fast and efficient Array Snapshots and has the least impact on production hosts during the entire backup process. Array-optimized backup is automatically chosen



as the preferred backup mode by HPE GreenLake for Backup and Recovery for the protection of VMware VMFS datastores, vVol containers/VMs, and array volumes provisioned on HPE GreenLake-managed storage.

When datastores of VMware VMs are created on array volumes that are part of a synchronous replication group on HPE GreenLake-managed storage, [Protection Policies](#) can be configured to perform Array-optimized backups from replicated snapshots created by the service on the partner array. Array-optimized backups using replicated snapshots not only offload backup workloads from primary production storage and hosts, but the retained replicated snapshots also act as instant recovery points during primary storage failure.

Any storage

Includes third-party and HPE storage systems not managed by HPE GreenLake, which utilize VMware CBT backups. This backup type currently uses VMware [NBDSSL](#) and [Hot-add](#) transport modes while backing up the VMs. VMware CBT backup supports multiple storage systems and provides a more granular mode of backing up VMware VMs. However, VMware CBT backups have limitations: NBDSSL transport mode provides secure network transfer, but it results in longer backup times and increased resource utilization of VMware ESXi hosts. Hot-add transport mode utilized by the service offers faster speeds when compared to NBDSSL transport mode by directly attaching VM disks to the Protection Store Gateway via the SCSI interface. But Hot-add transport mode requires the Protection Store Gateway VM to have accessibility to the VM datastores and supports only SCSI disks. HPE GreenLake for Backup and Recovery automatically chooses Hot-add as the preferred backup mode whenever the Protection Store Gateway VM is hosted in the same vCenter data center or ESXi cluster as the VMware VM that needs to be protected.

For an additional layer of resiliency, copies (replicas) can be configured for on-premises recovery points. On-premises protection policy schedules can be configured to replicate a recovery point to a separate Protection Store Gateway or HPE StoreOnce System. When using cloud backups, the HPE Cloud Protection Store schedule can be configured to use either the original recovery point or the replica of the recovery point as its source.

For more details, including a current list of supported storage, see [HPE GreenLake for Backup and Recovery QuickSpecs](#).

HPE GreenLake cloud data services (complementary)

[HPE GreenLake for Block Storage](#)

Offers provisioning and managing block storage for enterprise application needs (powered by [HPE Alletra Storage MP](#), which is modular and disaggregated). HPE GreenLake for Backup and Recovery offers data protection for array volumes managed by HPE GreenLake for Block Storage using global Protection Policies.

[HPE GreenLake for Private Cloud Business Edition \(PCBE\)](#)

Offers global lifecycle management of hybrid cloud infrastructure and virtualization resources. A tight integration with HPE GreenLake for Backup and Recovery provides consistent protection Service Level Agreements (SLAs) across the entire global infrastructure, including automatic protection of newly created PCBE resources. This is achieved by adding HPE GreenLake for Backup and Recovery Protection Policies to PCBE VM provisioning policies.

[HPE GreenLake for Disaster Recovery](#)

Offers disaster recovery, including ransomware protection, for mission-critical data using Zerto Continuous Data Protection (near-synchronous replication). From the same cloud console, users can periodically back up hybrid cloud resources using HPE GreenLake for Backup and Recovery and replicate resources using HPE GreenLake for Disaster Recovery.

HPE Catalyst technology (embedded)

Regardless of the backup data type chosen, the backups will be stored space-efficiently after being processed by the HPE Catalyst deduplication and compression technology. This technology is owned and developed by Hewlett Packard Enterprise. The initial application was originally developed for HPE StoreOnce Systems. As a key component of HPE StoreOnce Systems, the HPE Catalyst technology delivers significant backup data space savings for tens of thousands of users. This proven technology is built into HPE GreenLake for Backup and Recovery and delivers unmatched space efficiency.

Central to HPE Catalyst is the deduplication technology that efficiently stores the backup data as chunks. A chunk is stored only once. Duplicate chunks are replaced by a much smaller reference to the existing chunk. The process of chunking and matching is a highly optimized in-line process that has been refined over many releases. While this is conceptually like other deduplication processes, what sets it apart is the efficiency of the process to deliver high-performance backup — and the 4 KB chunk sizes (the smallest of any vendor) that provide unmatched space savings.

All the backups created by the service after the first full backup are synthetic-full backups that remove the hassles of an incremental or incremental-forever backup chain and results in improved space savings in the On-Premises/Cloud Protection Stores. During every backup schedule trigger, only incremental changes are read and deduplicated in 4 KB chunk sizes. The backup process ends with a synthetic full backup creation that is merely a metadata-only rearrangement process performed by the Catalyst server. Thus, the synthetic full backups



are created at incremental backup speed. This mode of backing up the data along with the small chunk sizes used with HPE GreenLake for Backup and Recovery enables fast, efficient on-premises and cloud backups.

The use of HPE Catalyst adds to the security of HPE GreenLake for Backup and Recovery. As well as providing encryption and data immutability, the HPE Catalyst application programming interface (API) is obfuscated from malware, which provides additional protection against ransomware attacks. Malware is incapable of activating within an HPE Catalyst-based protection store because HPE Catalyst does not use standard operating system commands for its operations.

Architecturally, HPE Catalyst is built using an HPE Catalyst client and server. The client executes the chunking and communicates with the server to identify duplicate chunks. The client can be run close to the data and the server run at the most appropriate location for the backup data storage. Only the unique chunks are transmitted to the server to optimize use of network bandwidth for high performance and to control costs. For VMware protection, the HPE Catalyst server runs in both the On-Premises Protection Store and the HPE Cloud Protection Store. For AWS protection, the HPE Catalyst server runs in the HPE Cloud Protection Store.

HPE Catalyst technology is also used when creating backup copies (replicas) to a separate On-Premises Protection Store. This enables a remote backup copy (replica) as an alternative or addition to cloud backup.

HPE GreenLake for Backup and Recovery in use

This section illustrates the ease of use of HPE GreenLake for Backup and Recovery and how simple it is to configure Protection Groups and global Protection Policies to automate protection.

The example use case shown in this section is for a customer named “XYZ” that has empowered their developers and business units outside of IT to be able to deploy workloads on-premises and in AWS on-demand, while also backing-up these new workloads automatically. A global Protection Policy is used for consistent protection of both on-premises and cloud workloads.

Note

HPE GreenLake for Backup and Recovery is a cloud native service that gets updated on a regular basis. There might be occasions when the latest user interface screens and names do not match the screenshots taken at the publication time of this paper.

HPE GreenLake for Backup and Recovery is launched from the Data Services Cloud Console, shown in Figure 5.

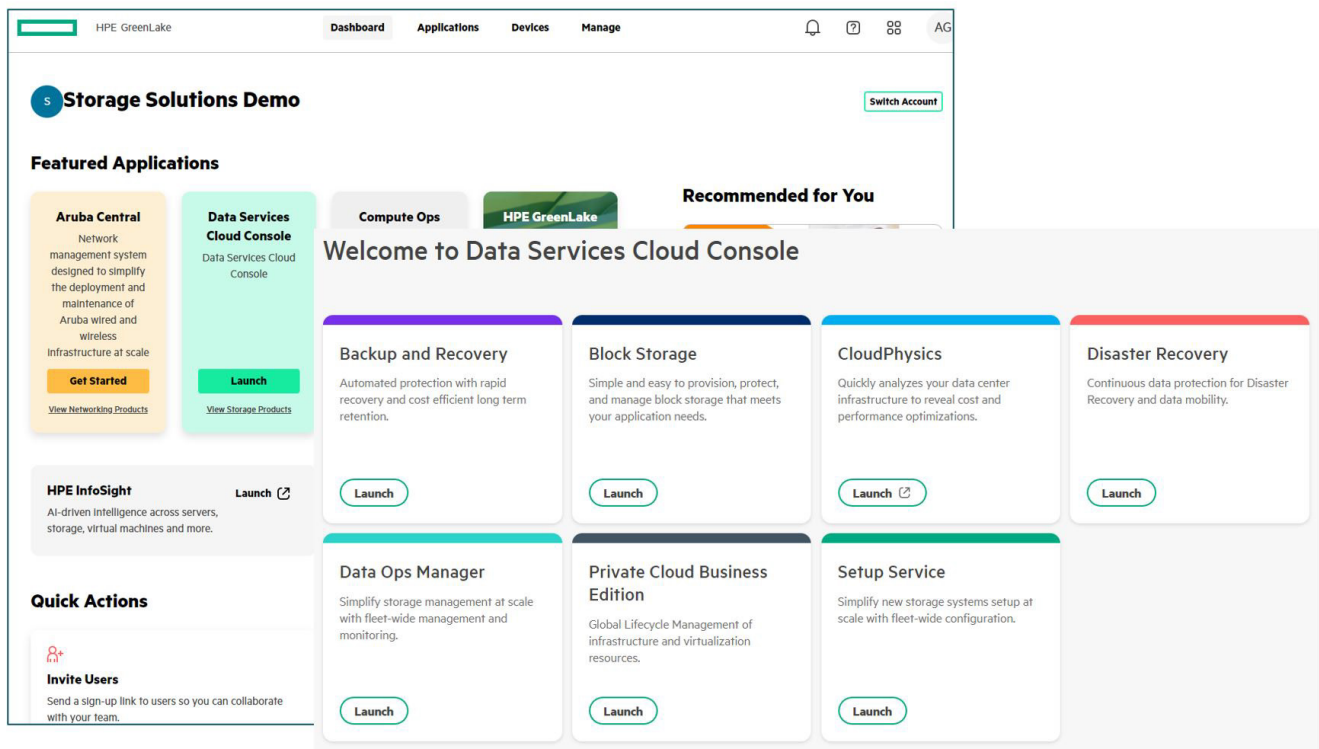


Figure 5. HPE GreenLake console streamlines user, device, and subscription management across all cloud services and is the launching point for data services



Operations

HPE GreenLake for Backup and Recovery makes it easy and reliable to meet recovery SLAs for on-premises and cloud native workloads. The required recovery points and times are used to define global Protection Policies. Automatic protection is achieved by defining Protection Groups for resources (AWS, VMware, MSSQL, HPE Array Volumes) that require similar protection. The resources are then protected by assigning a policy to an individual resource or group.

It is as easy as: define a policy, then group, protect, and recover data.

Consolidated inventory views

The protection status for all registered resources is shown in the HPE Backup and Recovery inventory views. These views of resource types for all locations provide global views that are not available from individual vCenter Server or AWS consoles. Global search and filtering makes it simple to find resources.

Figure 6 is an example of viewing the protection status for customer XYZ's cloud resources (AWS EC2 instances).

Name	Provider State	Protection Status	Protection Policy
<input type="checkbox"/> 00-allang-rh8-ec2-1	Running	Partial	Database Global Policy demo (Protection Group)
<input type="checkbox"/> 00-allang-rh8-ec2-2	Running	Pending	Database Global Policy demo (Protection Group)
<input type="checkbox"/> 00-billo-rh8-ec2-1	Running	Protected	TestDrive All options (Self)
<input type="checkbox"/> 00-billo-win2k19-ec2-2	Running	Unprotected	

Figure 6. Inventory views based on resource type show a consolidated view of protection status for all managed resources

Protection Policies

Protection Policies allow you to set up global rules for defining the extent of protection and lifecycle management of the recovery points, which can be applied to both on-premises and cloud native resources. They include where to store the recovery points (local, on-premises, cloud), how often backups are captured, how long backups are retained, and backup data immutability. For on-premises backups, a backup copy (replica) to a separate location can be configured. The replica schedule is the same as the backup frequency. The replica retention and immutability settings can be configured independently of the backup settings.

Note

The immutable option is only available for Array Snapshots if the array supports immutability. The immutability time is defined by the backup retention period. During this time, an immutable backup cannot be modified or deleted by anyone — not even by users who otherwise would have the appropriate permissions to delete backups.



The protection targets are flexible and can be configured with the following tiers:

Table 1. Protection tier options

Protected resource	Data source	Protection targets	Location coverage
VMware	HPE GreenLake managed storage	<ul style="list-style-type: none"> • Array Snapshot only • Replicated Snapshot only • On-Premises Protection Store (+replica) only • HPE Cloud Protection Store only • Array Snapshot/Replicated Snapshot, On-Premises Protection Store (+replica) • Array Snapshot/Replicated Snapshot, On-Premises Protection Store (+replica), HPE Cloud Protection Store 	<ul style="list-style-type: none"> • Local • Local • On-premises • Cloud • Local, on-premises • Local, on-prem, cloud
VMware	Any storage (not managed by HPE GreenLake)	<ul style="list-style-type: none"> • On-Premises Protection Store (+replica) only • HPE Cloud Protection Store only • On-Premises Protection Store (+replica), HPE Cloud Protection Store 	<ul style="list-style-type: none"> • On-premises • Cloud • On-premises, cloud
MSSQL (virtual)	HPE GreenLake managed storage	<ul style="list-style-type: none"> • Array Snapshot only • On-Premises Protection Store only • HPE Cloud Protection Store only • Array Snapshot, On-Premises Protection Store • Array Snapshot, HPE Cloud Protection Store • Array Snapshot, On-Premises Protection Store, HPE Cloud Protection Store 	<ul style="list-style-type: none"> • Local • On-premises • Cloud • Local, on-premises • Local, cloud • Local, on-premises, cloud
MSSQL (virtual)	Any storage (not managed by HPE GreenLake)	<ul style="list-style-type: none"> • On-Premises Protection Store only • HPE Cloud Protection Store only • On-Premises Protection Store, HPE Cloud Protection Store 	<ul style="list-style-type: none"> • On-premises • Cloud • On-premises, cloud
HPE Array Volumes	HPE GreenLake managed storage	<ul style="list-style-type: none"> • Array Snapshot only • Array Snapshot, On-Premises Protection Store • Array Snapshot, HPE Cloud Protection Store • Array Snapshot, On-Premises Protection Store, HPE Cloud Protection Store 	<ul style="list-style-type: none"> • Local • Local, on-premises • Local, cloud • Local, on-premises, cloud
AWS	AWS account	<ul style="list-style-type: none"> • AWS Snapshot only • HPE Cloud Protection Store Only • AWS Snapshot, HPE Cloud Protection Store 	<ul style="list-style-type: none"> • Local • Cloud • Local, cloud

HPE GreenLake for Backup and Recovery also provides the option to backup the transaction logs of Microsoft SQL Server databases running on VMWare VMs to an Array Snapshot or On-Premises Protection Store. These database log backups are automatically replayed by the service to roll-forward the databases to the desired point in time during the recovery process.

Note

At the time of publication of this paper, protection of AWS RDS instances by HPE GreenLake for Backup and Recovery is limited to AWS Snapshots only. Support for backup to HPE Cloud Protection Store for long-term retention will be delivered as part of regular updates of the service.



Figure 7 is an example of setting up global rules for consistently protecting all customer XYZ's resources. This example shows setting up VMware protection schedules with the **Replicate** option being used to send a copy to another destination. In addition, the recovery point replica is marked **Immutable**, so it cannot be deleted for the duration of the configured retention period.

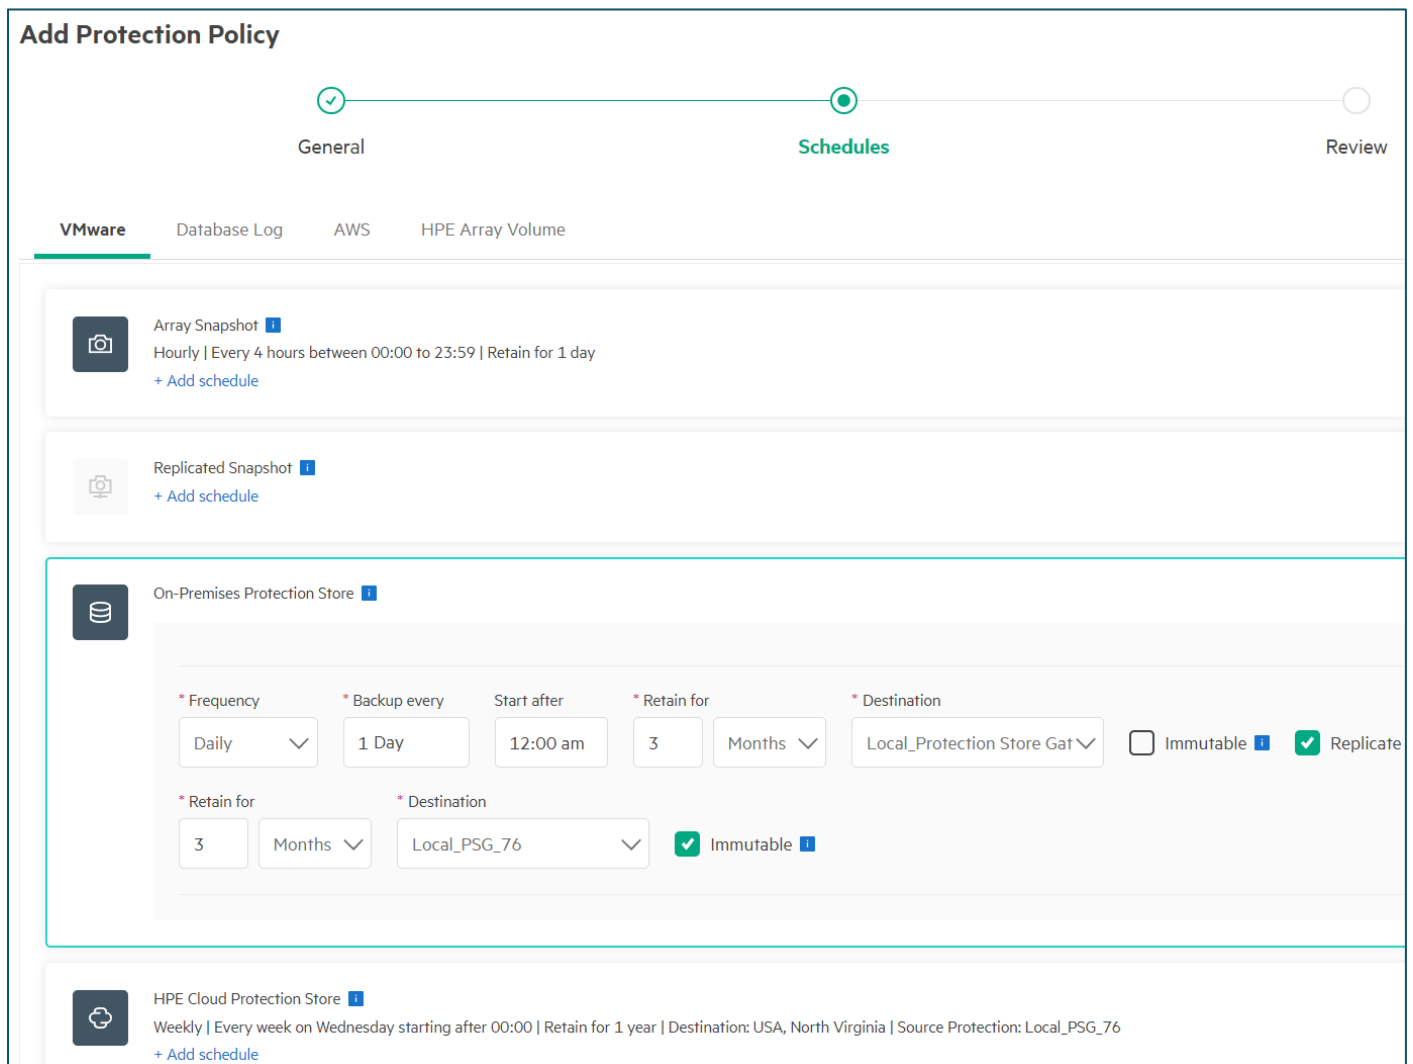


Figure 7. Protection Policies provide the protection service levels the organization needs

Protection Groups

Protection Groups define what is to be protected; these allow you to group a set of resources that require a similar level of protection. Automatic Protection Groups are used to automatically protect resources added to existing underlying groupings (VMware folders/containers/tags/storage replication groups, MSSQL availability groups, storage volume groups, and AWS tags). Custom Protection Groups are used to explicitly select the resources for protection.



Figure 8 is an example of grouping customer XYZ's resources (EC2 instances filtered by AWS tags) for automatic protection.

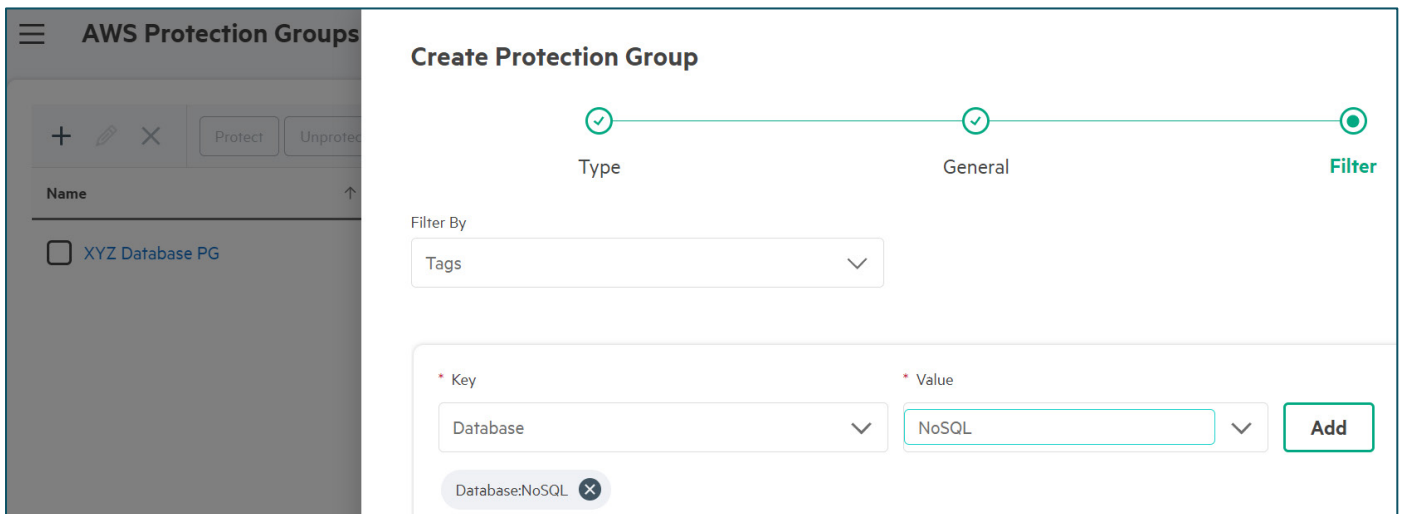


Figure 8. Automatic Protection Groups protect all resources, existing and new, assigned to an underlying grouping

Protecting data

After defining Protection Policies and Protection Groups, you can easily start protection (at scale) by assigning a Protection Policy to a Protection Group or individual resource.


Note

HPE GreenLake for Backup and Recovery blocks the user from protecting Data Orchestrator and Protection Store Gateway VMs to avoid disruption of backup and recovery operations and unnecessary protection of the On-Premises Protection Store. The service provides automated recovery support if these components are damaged or destroyed.


Figure 9 is an example of protecting customer XYZ's resources (AWS Protection Group), using a global Protection Policy that includes all sources and targets. This same policy can be used for consistent protection of all resources.



Protect



Protection Policy



Options

Select a policy to assign to XYZ Database PG

+

Name	Summary																																								
<input checked="" type="checkbox"/> Global Policy Demo	<table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 25%;">Array Snapshot</td> <td style="width: 15%;">VMware</td> <td style="width: 10%;">Hourly</td> <td style="width: 50%;">Every 4 hours between 00:00 to 23:59</td> </tr> <tr> <td>Replicated Snapshot</td> <td>VMware</td> <td>Hourly</td> <td>Every 4 hours between 00:00 to 23:59</td> </tr> <tr> <td>On-Premises Protection ...</td> <td>VMware</td> <td>Daily</td> <td>Every day starting after 00:00</td> </tr> <tr> <td>Array Snapshot</td> <td>HPE Arra...</td> <td>Hourly</td> <td>Every 4 hours between 00:00 to 23:59</td> </tr> <tr> <td>On-Premises Protection ...</td> <td>HPE Arra...</td> <td>Daily</td> <td>Every day starting after 00:00</td> </tr> <tr> <td>AWS Snapshot</td> <td>AWS</td> <td>Daily</td> <td>Every day starting after 00:00</td> </tr> <tr> <td>Database Transaction Log</td> <td>MSSQL</td> <td>Daily</td> <td>Every day starting after 00:00</td> </tr> <tr> <td>HPE Cloud Protection St...</td> <td>VMware</td> <td>Week...</td> <td>Every week on Tuesday starting after ...</td> </tr> <tr> <td>HPE Cloud Protection St...</td> <td>HPE Arra...</td> <td>Week...</td> <td>Every week on Tuesday starting after ...</td> </tr> <tr> <td>HPE Cloud Protection St...</td> <td>AWS</td> <td>Week...</td> <td>Every week on Tuesday starting after ...</td> </tr> </table>	Array Snapshot	VMware	Hourly	Every 4 hours between 00:00 to 23:59	Replicated Snapshot	VMware	Hourly	Every 4 hours between 00:00 to 23:59	On-Premises Protection ...	VMware	Daily	Every day starting after 00:00	Array Snapshot	HPE Arra...	Hourly	Every 4 hours between 00:00 to 23:59	On-Premises Protection ...	HPE Arra...	Daily	Every day starting after 00:00	AWS Snapshot	AWS	Daily	Every day starting after 00:00	Database Transaction Log	MSSQL	Daily	Every day starting after 00:00	HPE Cloud Protection St...	VMware	Week...	Every week on Tuesday starting after ...	HPE Cloud Protection St...	HPE Arra...	Week...	Every week on Tuesday starting after ...	HPE Cloud Protection St...	AWS	Week...	Every week on Tuesday starting after ...
Array Snapshot	VMware	Hourly	Every 4 hours between 00:00 to 23:59																																						
Replicated Snapshot	VMware	Hourly	Every 4 hours between 00:00 to 23:59																																						
On-Premises Protection ...	VMware	Daily	Every day starting after 00:00																																						
Array Snapshot	HPE Arra...	Hourly	Every 4 hours between 00:00 to 23:59																																						
On-Premises Protection ...	HPE Arra...	Daily	Every day starting after 00:00																																						
AWS Snapshot	AWS	Daily	Every day starting after 00:00																																						
Database Transaction Log	MSSQL	Daily	Every day starting after 00:00																																						
HPE Cloud Protection St...	VMware	Week...	Every week on Tuesday starting after ...																																						
HPE Cloud Protection St...	HPE Arra...	Week...	Every week on Tuesday starting after ...																																						
HPE Cloud Protection St...	AWS	Week...	Every week on Tuesday starting after ...																																						

Figure 9. Applying a Protection Policy to a Protection Group or individual resource starts automated protection and provides consistent protection at scale

Note

HPE GreenLake for Backup and Recovery currently supports application-consistent protection for on-premises VMware vCenter workloads, including VMs/datastores and Microsoft SQL Server databases running on the VMs. The service also enables the protection of on-premises enterprise workloads hosted on physical platforms and other custom applications with storage capacity provisioned and managed by [HPE GreenLake for Block Storage](#) using global Protection Policies. However, the protection for these applications is done at the storage volume level with crash-consistent Array Snapshots and array-optimized local/cloud backups. In this case, users can build custom scripts for quiescing or un-quiescing the workloads and leverage HPE GreenLake for Backup and Recovery APIs for application-consistent protection.

Recovering data

After data is protected, you can easily view and recover from recovery points. A recovery point is a snapshot or backup of data at a point in time. The recovery points continue to exist even after a resource is deleted from the inventory. A consolidated view of all recovery points per-protected resource allows you to easily choose where and when to recover from. Data can be recovered to the original or a new resource using any recovery point. On-premises workloads are recoverable from an Array Snapshot, On-Premises Protection Store, or HPE Cloud Protection Store. AWS resources are recoverable from an AWS Snapshot, Staging Snapshot, or HPE Cloud Protection Store.

Figure 10 is an example of choosing a recovery point for customer XYZ's resources (EC2 backup in HPE Cloud Protection Store).



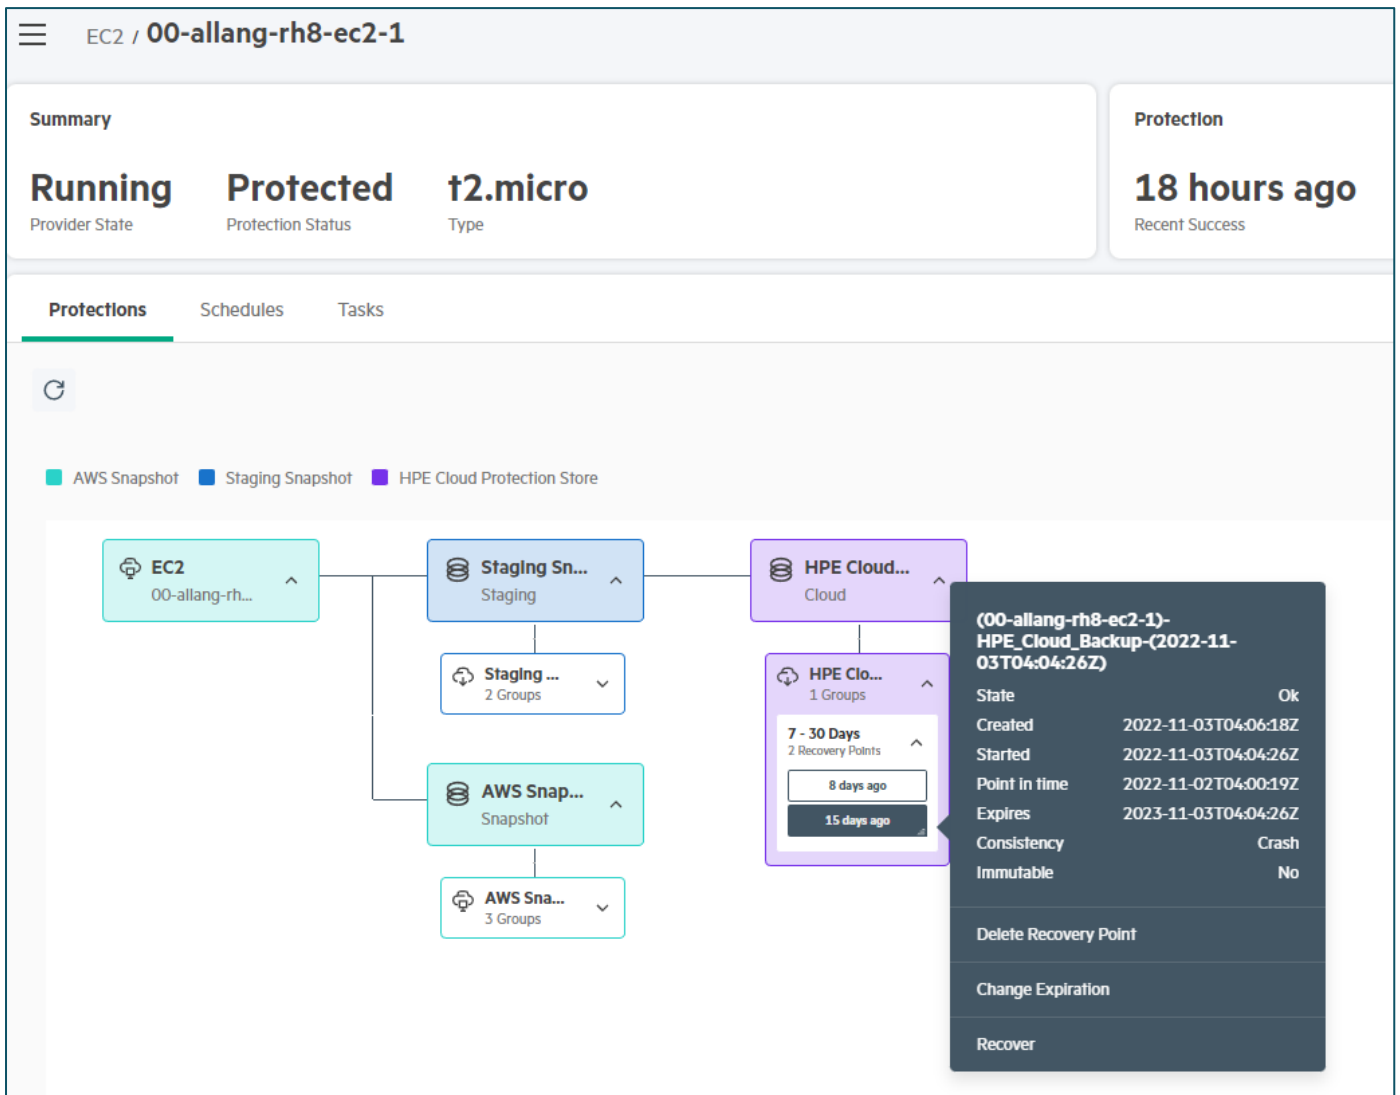


Figure 10. Choosing a recovery point to recover from the desired backup location and point in time

For VMware VMs, granular recovery options are provided to recover the individual virtual disks, files, and folders as shown in Figure 11. For file/folder-level recovery, HPE GreenLake for Backup and Recovery allows the user to explore and search the required file/folder protected with the selected recovery point in the On-Premises Protection Store/HPE Cloud Protection Store and perform recovery to the same VM or a different VM within the same vCenter or an alternate vCenter.

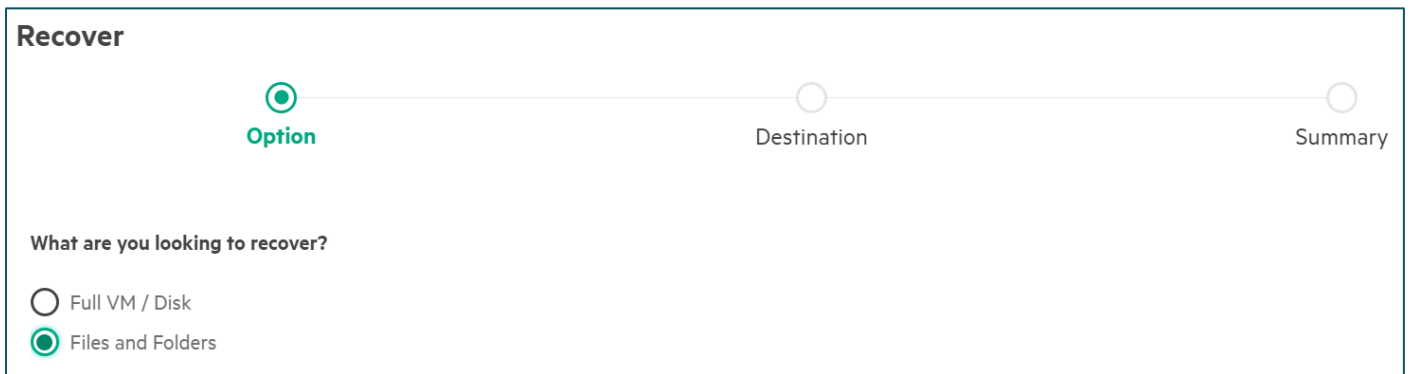


Figure 11. Choosing file/folder recovery automatically builds a searchable catalog of the selected recovery point if it is not pre-indexed by the user



For MSSQL databases hosted on VMware VMs, recovery can be from the desired “Point in Time” or “Point of Failure,” as shown in Figure 12.

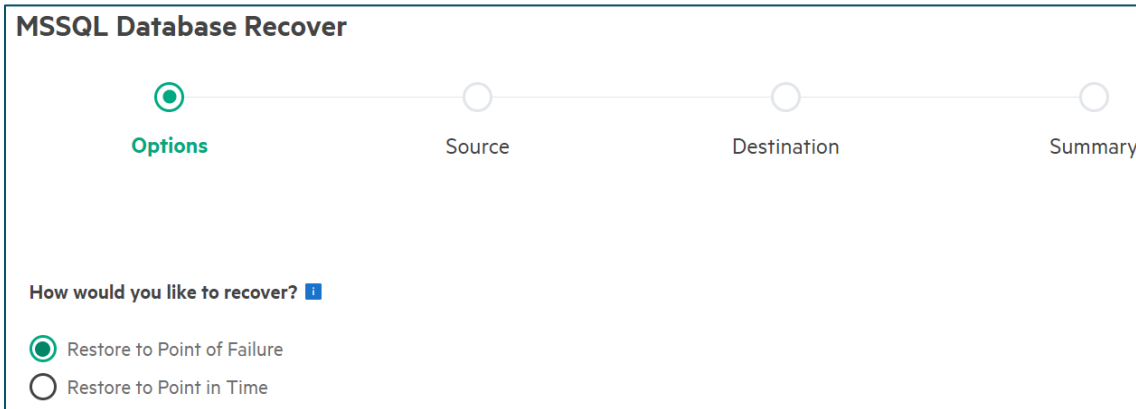


Figure 12. Restore to Point of Failure enables automated recovery of Microsoft SQL databases without any data loss

Recovering on-premises components

The Data Orchestrator and Protection Store Gateway — the required components for on-premises protection — can be recovered in the event of failure, accidental deletion, or hypervisor or site unavailability. The service provides end-to-end automation support for a simple and reliable recovery of these VMs to the same vCenter or an alternate vCenter in a DR site to minimize the downtime and complexities associated with recovering the data and resuming the protection operations after a disaster.

Figure 13 is an example of recovering the Data Orchestrator VM in “Disconnected” State.

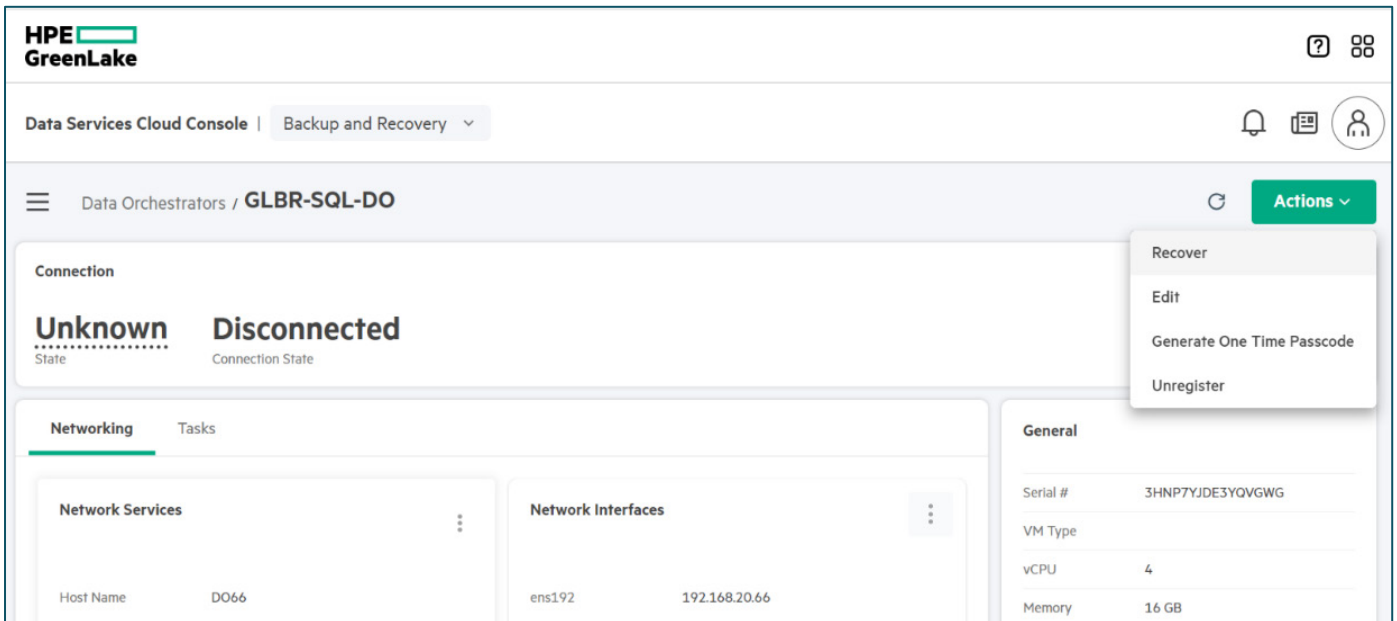


Figure 13. Data Orchestrator recovery process automatically restores the configuration information and the backup catalog

Monitoring and reporting

The protection status of all the resources, the backup space utilization/savings, and the issues that might need attention can be monitored from the HPE GreenLake for Backup and Recovery dashboard. Reports can also be generated for how long Protection Jobs take to complete and the amount of backup data processed/stored. This can help you make data-driven adjustments to schedules and backup target resources to ensure jobs complete within the required SLAs.



Figure 14 is an example of a report of Protection Jobs provided by HPE GreenLake for Backup and Recovery.

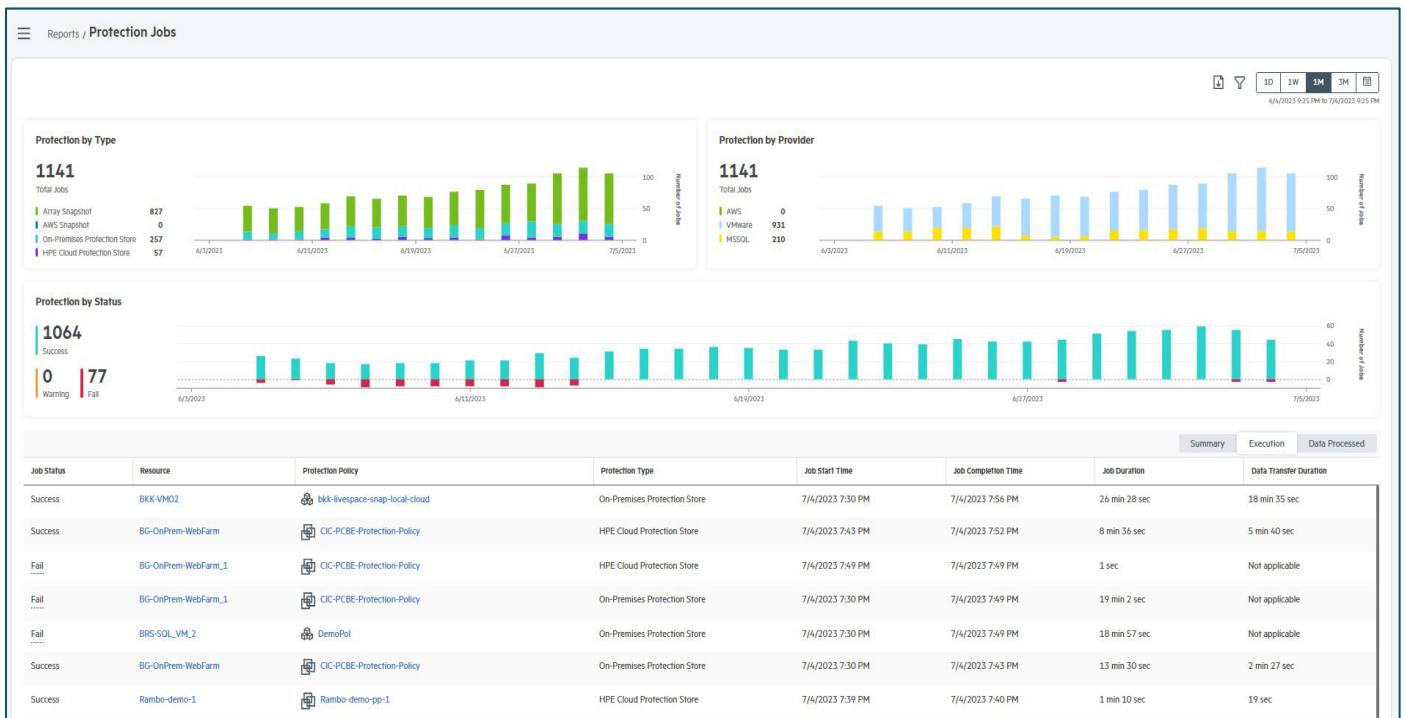


Figure 14. HPE GreenLake for Backup and Recovery "Protection Jobs" report

Summary

HPE GreenLake for Backup and Recovery is the modern way to protect your on-premises and cloud native workloads. This cloud data service, available from the HPE GreenLake edge-to-cloud platform, removes the complexity and costs of deploying traditional backup and recovery software. HPE GreenLake for Backup and Recovery redefines backup and recovery with the simplicity and flexibility of software delivered and managed as a service. Global policy-based protection — configured from a single cloud console — enables you to set up and automate the protection of your VMs in a few simple steps. In the event of a data loss or error, all restore points are easily visualized so the data from the required recovery point can be quickly brought back online. The service is secure and efficient, delivering protection from ransomware, and provides a lower TCO. Security is built into the HPE GreenLake edge-to-cloud platform, plus the backup data is encrypted and optionally set as immutable, along with dual authorization required for backup deletion. The superior storage efficiency provided by HPE Catalyst technology reduces the amount of backup storage capacity needed by up to 8x compared to similar solutions.

The service can co-exist with other backup software for evaluations, and you can try the service — without any commitment for 90 days — at connect.hpe.com/HPE_Backup_and_Recovery_Trial.



Appendix A: Storage efficiency test environment

Testing was conducted at Hewlett Packard Enterprise in Q2–Q3, 2022 by protecting VMs with HPE GreenLake for Backup and Recovery and four other vendors. The testing was designed to represent a typical virtual user environment and protection schedules. The testing used capacity information as reported by the vendor and/or object storage provider.

The testing simulated 50 weekly backups using a typical production workload:

- Five Windows VMs
- 100 GB of different file data per VM
- VMs backed up using HPE GreenLake for Backup and Recovery and each of the four protection vendors
- File data in each VM changed by a script between each backup iteration
- Process repeated 50 times with used cloud capacity measured for each vendor after each backup iteration



Resources

HPE GreenLake for Backup and Recovery QuickSpecs

hpe.com/psnow/doc/a50004269enw?section=Product%20Documentation

HPE GreenLake for Backup and Recovery Deployment Considerations Guide

hpe.com/psnow/doc/a00128576enw

Application page for 90-day HPE GreenLake for Backup and Recovery evaluation

connect.hpe.com/HPE_Backup_and_Recovery_Trial

Data Services Cloud Console Security Guide

hpe.com/psnow/doc/a00113337enw

HPE GreenLake for Backup and Recovery webpage

hpe.com/us/en/storage/data-protection-solutions/backup-recovery.html

Data Services on HPE GreenLake webpage

hpe.com/us/en/storage/data-services-cloud-console.html

HPE GreenLake for hyperconverged webpage

hpe.com/us/en/greenlake/hyperconverged-infrastructure.html

HPE GreenLake for Backup and Recovery protecting dHCI blog

community.hpe.com/t5/around-the-storage-block/unleash-the-power-of-hpe-backup-and-recovery-service-to/ba-p/7177406#.Y5oGlnbML-g

HPE StoreOnce webpage

hpe.com/us/en/storage/storeonce.html

HPE GreenLake Backup and Disaster Recovery with Virtual Machines video

hpe.com/h22228/video-gallery/us/en/v100001959/hpe-greenlake-backup-and-disaster-recovery-with-virtual-machines/video

Learn more at

hpe.com/us/en/storage/data-protection-solutions/backup-recovery.html

Explore **HPE GreenLake** 

 **Chat now (sales)**