Data Protector for Cloud Workloads 7.1

January 2025

Table of Contents

What's new in this release

What's new in this release

- New: CEPH RBD and Nutanix VG support for a secondary backup destination
- New: Support for installation on SLES 15
- New: Restore individual disks to a datastore for Proxmox platform
- New: Enable-cross restore between OpenStack and Virtuozzo platforms
- Improvements: update for UI views

Overview

In this section, we'll briefly discuss the architecture and main features of Data Protector for Cloud Workloads as well as some typical use case scenarios.

The Main Features section briefly summarizes the key functionalities of the Data Protector for Cloud Workloads solution.

In the Architecture section, you will learn what the main components of the Data Protector for Cloud Workloads solution are, as well as find out how to place them in your deployment.

In the Support Matrix you can check versions of supported virtualization platforms, backup, and cloud providers.

Platform requirements present what are hardware and software requirements needed to run Data Protector for Cloud Workloads components.

High availaibility section provides guidance to plan Data Protector for Cloud Workloads solutions resistant to failures.

Sizing Guide is the place where we present key information that the user needs to collect before the installation process.

Main Features

Backup

- Support for a wide range of of platforms:
    - Virtual Machines:
        - Red Hat Virtualization
        - oVirt
        - Oracle Linux Virtualization Manager
        - Nutanix Acropolis Hypervisor (AHV)
        - Citrix Hypervisor (XenServer) with CBT support
        - XCP-ng with CBT support
        - Proxmox VE
        - Oracle VM
        - OpenStack (incremental backups for Ceph RBD-based environments)
        - libvirt hypervisors (KVM, PowerKVM, KVM for IBM z, Xen)
        - SSC//Platform
        - Huawei FusionCompute
    - Containers:
        - Kubernetes (deployment-level protection for Persistent Volumes)
        - Red Hat OpenShift (deployment-level protection for Persistent Volumes)
    - Cloud:
        - Amazon EC2
        - GCP GCE
        - Microsoft 365
    - Storage:
        - Ceph RBD (with snapshot difference support)
        - Nutanix Files (with Changed-File Tracking)
        - Nutanix Volume Groups (with Changed-Region Tracking)
    - Applications:

        - generic backup mechanism for custom backup process
        - ready templates for commonly used applications
- Snapshot Management (Copy Data Management)
- Advanced backup features:
    - Snapshot consistent technology (quiesced/application-consistent snapshots or FS freeze)
    - Pre/post snapshot remote command execution on VM to enable operations such as DB quiesce
    - CBT/CFT for faster incremental backups
    - VM disk exclusion option
    - automatic backup repetition
- Backup SLAs:
    - VM automatic policy assignment based on regular expressions and tags
    - backup job prioritization
    - multiple policy rules (with different scheduling and backup destinations) for the same protected object
- Multi-node architecture:
    - better scalability
    - automatic task load balancing
    - suitable for geographically dispersed environments
- Built-in Data Protector for Cloud Workloads DB backup

Recovery

- File-level restore using mountable backups
    - directly via web browser
    - transfer to the remote host over SSH and WinRM
- Mounted backups - RAW disks shareable over iSCSI (for direct block-access to your backup data)
- Recovery plans for automated DR
    - on-demand restore of multiple VMs when needed

    - scheduled basis for testing
- Customizable networking and disk layout during restore
- oVirt/RHV/OLVM instant restore
- Individual disk recovery

Backup storage

- Integration with different backup destinations:
    - File-based:
        - Synthetic backup provider using XFS or NFS 4.2
        - Any mounted file system (local or remote, especially GlusterFS for replication, CephFS, NFS, SMB, and many more)
        - Dell EMC Data Domain (BoostFS integration)
    - Object Storage:
        - Amazon S3 (with Amazon Glacier as a 2nd tier archive storage),
        - S3-compliant storage (IBM Cloud, Oracle Cloud, Scality RING)
        - Google Cloud Storage
        - Microsoft Azure Blob Storage
        - OpenStack Swift
    - Enterprise-grade backup providers:
        - MicroFocus Data Protector
- Built-in data deduplication with Virtual Data Optimizer (VDO)
- Backup Copy - secondary backup destination to store data in more than one location
- Pre/post backup destination access command execution to execute custom operations on external storage providers such as replication

Security

- RBAC for administrative accounts
- Audit-log for administrative actions

- Customizable logging configuration for external SIEM support
- Data-at-rest encryption for file system backup destination
- Ransomware protection
    - Immutable Backup (XFS-based backup destination) that protects backup data from being encrypted by ransomware

UI and Integration

- Central, easy to use and modern management with HTML5-based web UI
- Advanced reporting directly in the UI and with e-mails
- Event notifications using e-mail, Slack or custom API call
- CLI for advanced administrators
- Open API for 3rd party software integration (REST API)
- Easy deployment in OpenShift environments using ready operator
- LDAP authentication
- OpenStack Horizon plugin
- oVirt/RHV/OLVM console integration
- Multi-language support:
    - English
    - Chinese
    - Spanish
    - German
    - French
    - Japanese
    - Polish

Architecture

High-level Architecture

Use Data Protector for Cloud Workloads to back up data from your virtualization platforms, M365 Cloud and storage providers. You can back up data to and recover data from a local filesystem or an NFS/CIFS share, object storage (cloud providers), or Micro Focus Data Protector.

Detailed Architecture

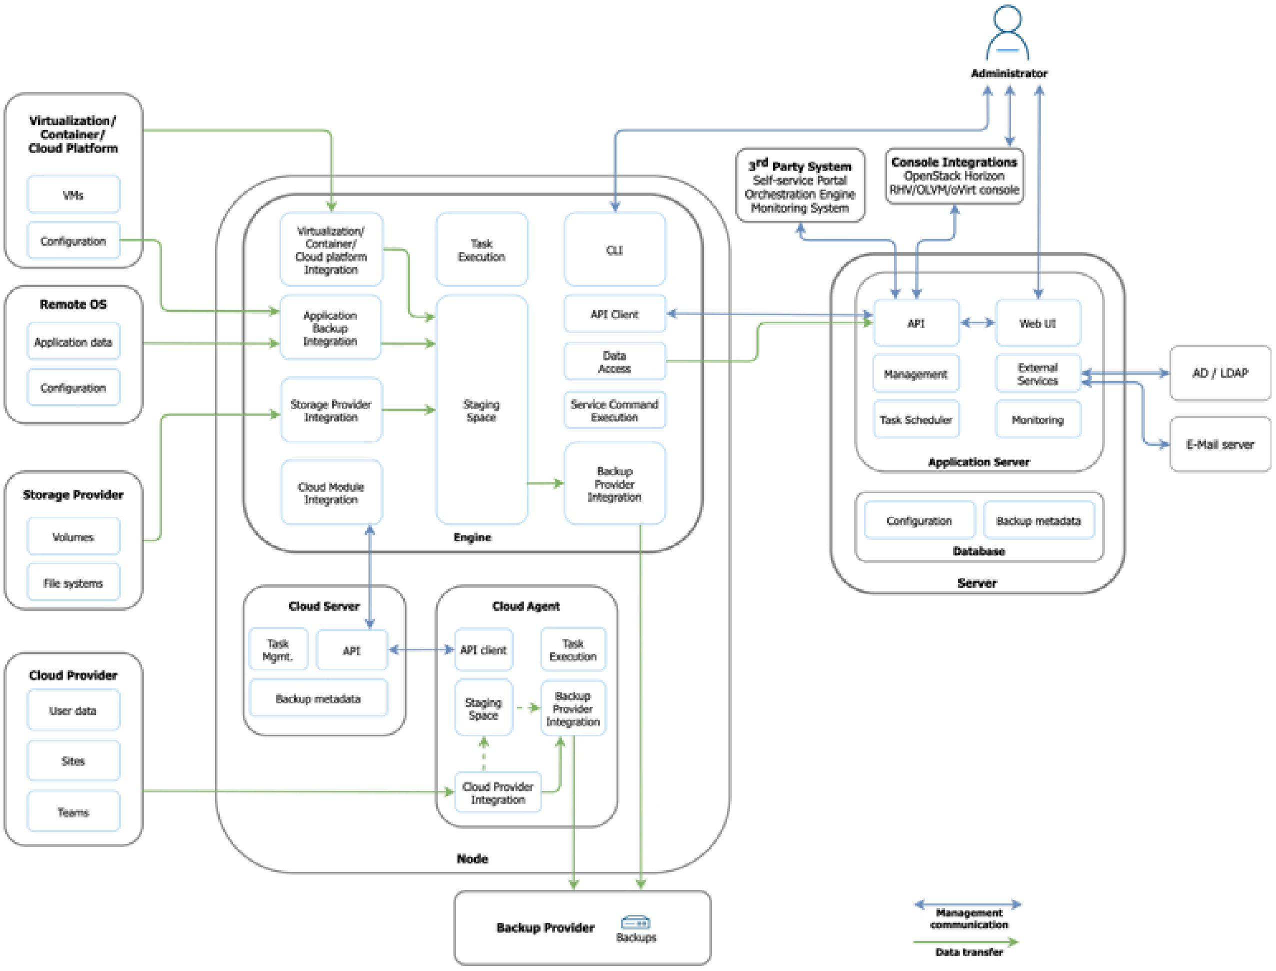

Data Protector for Cloud Workloads consists of 2 main components:

- **Data Protector for Cloud Workloads Server** - the central point of Data Protector for Cloud Workloads management, provides administrative Web UI, APIs and is a central repository of metadata
- **Data Protector for Cloud Workloads Node** - data mover that performs backups, restores, and mounts:
    - multiple nodes can be deployed for scalability or other reasons,
    - all nodes are managed by the server and need to be registered to the server.

Component placement

- **Data Protector for Cloud Workloads Server and Node can be installed on the same host.**
- The Server can be installed on a physical machine or VM - externally deployed nodes require network connectivity to the Server and PowerProtect DD target(s).
- Nodes may be deployed as physical or virtual systems unless the selected backup strategy requires the Node to be installed as a VM on a Hypervisor Cluster (especially when the "disk attachment" export mode is mentioned).
- Both components are installed on a CentOS 8 Stream or RHEL 8 minimal.

For detailed deployment scenarios refer to the following sections:

Virtual Environments

Microsoft 365

Applications

Storage Providers

Typical workflows

There are several standard workflows in Data Protector for Cloud Workloads and they result in a set of tasks:

- **Backup**
  - **Export** - a task that creates backup or snapshot and exports data to the staging space
  - **Store** - a task that moves data to the backup destination
- **Restore to filesystem**
  - **Restore** - a task that gets data from a backup provider and puts data in the staging space
- **Restore to a virtualization platform**
  - **Restore** - a task that gets data from a backup provider and puts data in the staging space (if it is a full backup that is being restored residing on the file system backup provider - this task just informs where files are waiting for import task)
  - **Import** - a task that imports data to the virtualization platform and recreates VM
- **Restore for a mount (file-level restore)**
  - **Restore** - a task that gets data from a backup provider and puts data in the staging space (if it is a full backup that is being restored residing on the file system backup provider - this task just informs where files are waiting for mount task)
  - **Mount** - mounts backup on the Data Protector for Cloud Workloads Node and either allows user to browse files or exposes backup over iSCSI, so that remote iSCSI initiator can access it)
- **Snapshot**
  - **Snapshot** - a task that creates a local persisted snapshot of a VM in the hypervisor environment according to a policy that was assigned to the VM - snapshots that are no longer needed (according to the policy) will be removed

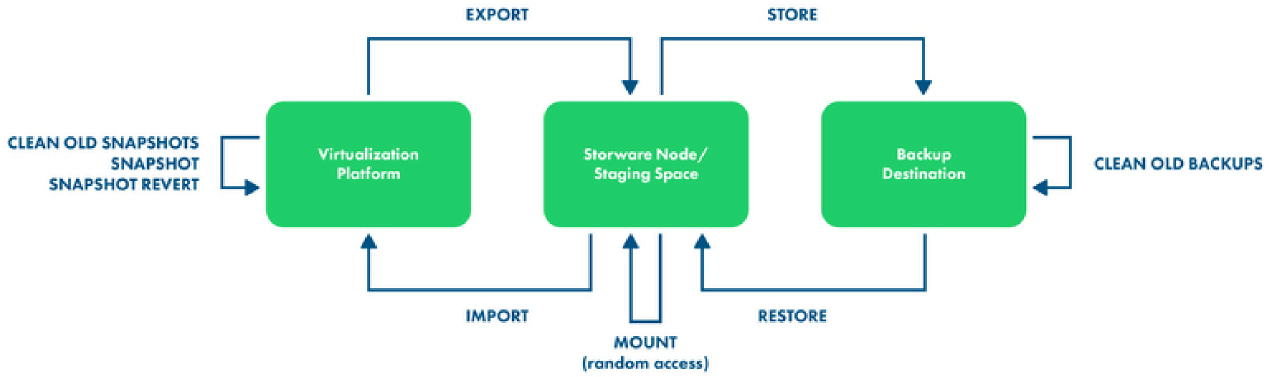
EXPORT
STORE
CLEAN OLD SNAPSHOTS
SNAPSHOT
SNAPSHOT REVERT
Virtualization Platform
Storware Node/ Staging Space
Backup Destination
CLEAN OLD BACKUPS
IMPORT
RESTORE
MOUNT (random access)

Typical Scenarios

Backup & Recovery

The core functionality of Data Protector for Cloud Workloads is an **agentless backup** for multiple virtualization, container, cloud platforms, storage providers and applications.

With snapshot-based backups, you don't have to install an agent inside VMs or customize your hypervisors.

Backups performed by Data Protector for Cloud Workloads usually are crash-consistent, but you can enable **application consistency** or enhance the backup process with your own **custom pre/post snapshot remote command execution**.

Snapshots are exported from your virtualization platform and can be stored in the backup provider of your choice. You can use enterprise-grade backup providers, object storage, or just a file system as your target.

This means that Data Protector for Cloud Workloads can act as a **stand-alone solution** or as a **proxy to your existing storage or enterprise backup provider.**

You also can periodically restore your VMs to verify if your backups are consistent.

With mounted backups, you can also **restore individual files** from your backups via a Web UI or directly from Data Protector for Cloud Workloads Node.

Disaster Recovery

Real disasters can sometimes happen - with Data Protector for Cloud Workloads you can configure your backups to be performed in one datacenter and - if necessary - restore them in a second datacenter.

Data Protector for Cloud Workloads can use replicated file systems or other built-in backup provider mechanisms to allow you to keep a copy in the secondary data center.

During DR, you can use Recovery Plans to restore multiple VMs to a predefined location.

Snapshot Management

Backups are usually quite an intensive operation. Snapshots have to be exported and stored, which usually means that you can't perform them too often. With Data Protector for Cloud Workloads, you can use Snapshot Management policies to periodically create additional snapshots on your VMs without the need to export them.

When you need to restore a VM to the most recent saved state, you can quickly revert to a snapshot that Data Protector for Cloud Workloads has created for you.

Application Backup & Recovery

There are many cases where VM-level backup may not be enough. Applications such as databases usually have their own mechanisms that guarantee consistent backups. As we are aware, in many situations you need to have the option to customize the backup process - therefore Data Protector for Cloud Workloads provides a **generic mechanism** for multiple scenarios.

You can prepare a custom script or invoke any backup command that produces backup artifacts (or just initiates the external backup process) on a remote host and stores backups to your backup provider.

With Application backup, you can extend your protection capabilities to:

- any remote applications with their own mechanisms
- hypervisor configuration
- files on remote hosts (physical, virtual, or containers)

    - this includes shares, mounted object-storage buckets, LVM block devices, or virtually anything which can be presented as a file
- initiating external backup processes such as RMAN

Support Matrix

Virtualization Platforms

Huawei FusionCompute

Supported backup strategies: CBT

Supported versions	8.0, 8,1, 8.2, 8.3, 8.5, 8.6
The last snapshot is kept on the hypervisor for incremental backups	Yes
Access to hypervisor OS required	No
Proxy VM required	No

Full backup	Supported
Incremental backup	Supported
Restore	Supported
File-level restore	Supported
VM disk exclusion	Supported
Quiesced snapshots	Not supported
Snapshots management	Supported
Pre/post command execution	Supported
Access to VM disk backup over iSCSI	Supported
VM name-based policy assignment	Supported
VM tag-based policy assignment	Not supported
Power-on VM after restore	Not supported

KVM

Supported backup strategies: SSH transfer

Supported versions	QEMU 2.1 and higher
The last snapshot is kept on the hypervisor for incremental backups	Yes
Access to hypervisor OS required	Yes
Proxy VM required	No

Full backup	Supported
Incremental backup	Supported *
Restore	Supported
File-level restore	Supported
VM disk exclusion	Supported
Quiesced snapshots	Supported
Snapshots management	Not supported
Pre/post command execution	Supported
Access to VM disk backup over iSCSI	Supported
VM name-based policy assignment	Supported
VM tag-based policy assignment	Not supported
Power-on VM after restore	Supported

** Not supported for LVM disks*

ⓘ Supported VM storage formats: QCOW2/RAW, Ceph RBD volume, LVM volume, LVM-thin volume

Nutanix AHV

Supported backup strategies: Disk attachment

Supported versions	5.5, 5.6, 5.8, 5.9, 5.10, 5.11, 5.15, 5.16, 5.17, 5,18, 5.19, 5.20, 6.0, 6.1, 6.5, 6.6, 6.7, 6.8
The last snapshot is kept on the hypervisor for incremental backups	Yes
Access to hypervisor OS required	No
Proxy VM required	Yes

Full backup	Supported
Incremental backup	Supported
Restore	Supported
File-level restore	Supported
VM disk exclusion	Supported
Quiesced snapshots	Supported
Snapshots management	Supported
Pre/post command execution	Supported
Access to VM disk backup over iSCSI	Supported
VM name-based policy assignment	Supported
VM tag-based policy assignment	Supported *
Power-on VM after restore	Supported

** When using Prism Central*

! Backup of virtual machines with vTPM enabled is not supported

Proxmox VE

Supported backup strategies: Export storage repositories, SSH transfer

	Export storage repository	SSH transfer
Supported versions	5.2, 5.3, 5.4, 6.0, 6.1, 6.2, 6.3, 6.3, 6.4, 7.0, 7.1, 7.2, 7.3, 7.4, 8.0, 8.1, 8.2	5.2, 5.3, 5.4, 6.0, 6.1, 6.2, 6.3, 6.3, 6.4, 7.0, 7.1, 7.2, 7.3, 7.4, 8.0, 8.1, 8.2
The last snapshot is kept on the hypervisor for incremental backups	Yes	Yes
Access to hypervisor OS required	Yes	Yes
Proxy VM required	Yes	No

	Export storage repository	SSH transfer
Full backup	Supported	Supported
Incremental backup	Not supported	Supported
Restore	Supported	Supported
File-level restore	Supported	Supported
VM disk exclusion	Not supported	Supported
Quiesced snapshots	Supported	Supported
Snapshots management	Supported	Supported
Pre/post command execution	Supported	Supported
Access to VM disk backup over iSCSI	Not supported	Supported
VM name-based policy assignment	Supported	Supported
VM tag-based policy assignment	Not supported	Not supported
Power-on VM after restore	Supported	Supported

OpenNebula

Supported backup strategies: CBT

Supported versions	6.6, 6.7, 6.8, 6.9
The last snapshot is kept on the hypervisor for incremental backups	No
Access to hypervisor OS required	No
Proxy VM required	Yes

Full backup	Supported
Incremental backup	Supported
Restore	Supported
File-level restore	Supported
VM disk exclusion	Supported
Quiesced snapshots	Supported
Snapshots management	Supported
Pre/post command execution	Supported
Access to VM disk backup over iSCSI	Supported
VM name-based policy assignment	Supported
VM tag-based policy assignment	Supported
Power-on VM after restore	Not Supported (always on)

ⓘ During the data export phase, the hypervisor may experience higher CPU usage

OpenStack

Supported backup strategies: Disk attachment, Image transfer, CBT

	Disk attachment	CBT	SSH Transfer
Supported versions	Queens, Rocky, Stein, Train, Ussuri, Victoria, Wallaby, Xena, Yoga, Zed, Antelope, Bobcat, Caracal	Queens, Rocky, Stein, Train, Ussuri, Victoria, Wallaby, Xena, Yoga, Zed, Antelope, Bobcat, Caracal	Queens, Rocky, Stein, Train, Ussuri, Victoria, Wallaby, Xena, Yoga, Zed, Antelope, Bobcat, Caracal
The last snapshot is kept on the hypervisor for incremental backups	Yes	No	Yes
Access to hypervisor OS required	No	No	Yes
Proxy VM required	Yes	Yes	No

	Disk attachment	CBT	SSH Transfer
Full backup	Supported *	Supported	Supported
Incremental backup	Not supported	Supported	Supported
Restore	Supported	Supported	Supported
File-level restore	Supported	Supported	Supported
VM disk exclusion	Supported	Supported	Supported
Quiesced snapshots	Supported **	Supported **	Supported **
Snapshots management	Supported ***	Supported ***	Not supported

Pre/post command execution	Supported	Supported	Supported
Access to VM disk backup over iSCSI	Supported	Supported	Supported ****
VM name-based policy assignment	Supported	Supported	Supported
VM tag-based policy assignment	Supported	Supported	Supported
Power-on VM after	Not supported	Not supported	Not supported

** Ceph RBD volumes only*
*** Hypervisor dependent*
**** Without snapshot revert*
***** RAW/LVM disks only*

Oracle Linux Virtualization Manager

Supported backup strategies: Disk attachment, Image transfer, CBT

	Disk attachment	Image transfer	CBT
Supported versions	4.3, 4.4, 4.5	4.3, 4.4, 4.5	4.4, 4.5
The last snapshot is kept on the hypervisor for incremental backups	No	Yes	No
Access to hypervisor OS required	No	No	No
Proxy VM required	Yes	No	No

	Disk attachment	Image Transfer	CBT

Full backup	Supported	Supported	Supported
Incremental backup	Not supported	Supported	Supported
Restore	Supported	Supported	Supported
File-level restore	Supported	Supported	Supported
VM disk exclusion	Supported	Supported	Supported
Quiesced snapshots	Supported	Supported	Supported
Snapshots management	Supported	Supported	Supported
Pre/post command execution	Supported	Supported	Supported
Access to VM disk backup over iSCSI	Supported	Supported *	Supported
VM name-based policy assignment	Supported	Supported	Supported
VM tag-based policy assignment	Supported	Supported	Supported
Power-on VM after restore	Supported	Supported	Supported

** Only for RAW disk types*

Direct LUN disks are not supported

Oracle VM

Supported backup strategies: Export storage repository

Supported versions	3.4
The last snapshot is kept on the hypervisor for incremental backups	n/a
Access to hypervisor OS required	No
Proxy VM required	No

Full backup	Supported
Incremental backup	Not supported
Restore	Supported
File-level restore	Supported
VM disk exclusion	Supported
Quiesced snapshots	Not supported
Snapshots management	Not supported
Pre/post command execution	Not supported
Access to VM disk backup over iSCSI	Supported
VM name-based policy assignment	Not supported
VM tag-based policy assignment	Supported
Power-on VM after restore	Not Supported

oVirt

Supported backup strategies: Disk attachment, Image transfer, CBT

	Disk attachment	Image transfer	CBT
Supported versions	4.0, 4.1, 4.2, 4.3, 4.4, 4.5	4.3, 4.4, 4.5	4.4, 4.5
The last snapshot is kept on the hypervisor for incremental backups	No	Yes	No
Access to hypervisor OS required	No	No	No
Proxy VM required	Yes	No	No

	Disk attachment	Image Transfer	CBT
Full backup	Supported	Supported	Supported
Incremental backup	Not supported	Supported	Supported
Restore	Supported	Supported	Supported
File-level restore	Supported	Supported	Supported
VM disk exclusion	Supported	Supported	Supported
Quiesced snapshots	Supported	Supported	Supported
Snapshots management	Supported	Supported	Supported
Pre/post command execution	Supported	Supported	Supported
Access to VM disk backup over iSCSI	Supported	Supported *	Supported
VM name-based policy assignment	Supported	Supported	Supported
VM tag-based policy assignment	Supported	Supported	Supported
Power-on VM after restore	Supported	Supported	Supported

** Only for RAW disk types*

Red Hat Virtualization

Supported backup strategies: Disk attachment, Image transfer, CBT

	Disk attachment	Image transfer	CBT
Supported versions	4.0, 4.1, 4.2, 4.3	4.3, 4.4	4.4
The last snapshot is kept on the hypervisor for incremental backups	No	Yes	No

Access to hypervisor OS required	No	No	No
Proxy VM required	Yes	No	No

	Disk attachment	Image Transfer	CBT
Full backup	Supported	Supported	Supported
Incremental backup	Not supported	Supported	Supported
Restore	Supported	Supported	Supported
File-level restore	Supported	Supported	Supported
VM disk exclusion	Supported	Supported	Supported
Quiesced snapshots	Supported	Supported	Supported
Snapshots management	Supported	Supported	Supported
Pre/post command execution	Supported	Supported	Supported
Access to VM disk backup over iSCSI	Supported	Supported *	Supported
VM name-based policy assignment	Supported	Supported	Supported
VM tag-based policy assignment	Supported	Supported	Supported
Power-on VM after restore	Supported	Supported	Supported

** Only for RAW disk types*

SC//Platform

Supported backup strategies: Export storage domain, disk attachment

	Export storage domain	Disk attachment

Supported versions	8.9	8.9
The last snapshot is kept on the hypervisor for incremental backups	No	Yes
Access to hypervisor OS required	No	No

	Export storage repository	SSH transfer
Full backup	Supported	Supported
Incremental backup	Not supported	Supported
Restore	Supported	Supported
File-level restore	Not supported	Supported
VM disk exclusion	Supported	Supported
Quiesced snapshots	Not supported	Not supported
Snapshots management	Supported	Supported
Pre/post command execution	Supported	Supported
Access to VM disk backup over iSCSI	Supported	Supported
VM name-based policy assignment	Supported	Supported
VM tag-based policy assignment	Supported	Supported
Power-on VM after restore	Supported	Supported

Virtuozzo

Supported backup strategies: Disk attachment

Supported versions	4.7, 5.0, 5.2, 5.3, 5.4, 6.0, 6.1, 6.2
The last snapshot is kept on the hypervisor for incremental backups	Yes

Access to hypervisor OS required	No
Proxy VM required	Yes

Full backup	Supported
Incremental backup	Supported
Restore	Supported
File-level restore	Supported
VM disk exclusion	Supported
Quiesced snapshots	Supported *
Snapshots management	Supported **
Pre/post command execution	Supported
Access to VM disk backup over iSCSI	Supported
VM name-based policy assignment	Supported
VM tag-based policy assignment	Supported
Power-on VM after restore	Not Supported (always on)

** Hypervisor dependent*

*** Without snapshot revert*

XCP-ng

Supported backup strategies: Single image (XVA), CBT

	Single image (XVA)	CBT
Supported versions	7.4, 7.5, 7.6, 8.0, 8.1, 8.2, 8.3	7.4, 7.5, 7.6, 8.0, 8.1, 8.2, 8.3
The last snapshot is kept on the hypervisor for incremental backups	Yes	Yes

Access to hypervisor OS required	No	Yes

	Single image (XVA)	CBT
Full backup	Supported	Supported
Incremental backup	Supported *	Supported
Restore	Supported	Supported
File-level restore	Not supported	Supported
VM disk exclusion	Not supported	Supported
Quiesced snapshots	Supported	Supported
Snapshots management	Supported	Supported
Pre/post command execution	Supported	Supported
Access to VM disk backup over iSCSI	Not supported	Supported
VM name-based policy assignment	Supported	Supported
VM tag-based policy assignment	Supported	Supported
Power-on VM after restore	Supported	Supported

** Not supported when using a synthetic backup destination*

XenServer (Citrix Hypervisor)

Supported backup strategies: Single image (XVA), CBT

	Single image (XVA)	CBT
Supported versions	6.5, 7.0, 7.1, 7.2, 7.3, 7.4, 7.5, 7.6, 8.0, 8.1, 8.2	6.5, 7.0, 7.1, 7.2, 7.3, 7.4, 7.5, 7.6, 8.0, 8.1, 8.2

The last snapshot is kept on the hypervisor for incremental backups	Yes	Yes
Access to hypervisor OS		

	Single image (XVA)	CBT
Full backup	Supported	Supported
Incremental backup	Supported	Supported *
Restore	Supported	Supported
File-level restore	Not supported	Supported
VM disk exclusion	Not supported	Supported
Quiesced snapshots	Supported	Supported
Snapshots management	Supported	Supported
Pre/post command execution	Supported	Supported
Access to VM disk backup over iSCSI	Not supported	Supported
VM name-based policy assignment	Supported	Supported
VM tag-based policy assignment	Supported	Supported
Power-on VM after restore	Supported	Supported

** Requires XenServer 7.3 or higher*

Containers

Kubernetes

Supported backup strategies: Helper pod, Ceph RBD

	Helper pod	Ceph RBD

Minimal version	1.10	1.10
The last snapshot is kept on the system for incremental backups	Yes	Yes
Access to OS required	No	No
Proxy VM required	No	No

	Helper pod	Ceph RBD
Full backup	Supported	Supported
Incremental backup	Not supportd	Supported *
Restore	Supported	Supported
File-level restore	Not supported	Supported *
Volume exclusion	Supported	Supported
Quiesced snapshots	Supported **	Supported **
Snapshots management	Not supported	Not supported
Pre/post command execution	Supported ***	Supported ***
Access to VM disk backup over iSCSI	Not supported	Supported *
Name-based policy assignment	Supported	Supported
Tag-based policy assignment	Supported	Supported
Power-on after restore	Supported	Supported
StatefulSet	Supported	Supported

** When using Ceph RBD as Persistent Volume*

*** Deployment pause*

**** Only 'post'*

Proxmox VE

Supported backup strategies: Export storage repository

Supported versions	5.2, 5.3, 5.4, 6.0, 6.1, 6.2, 6.3, 6.3, 6.4, 7.0, 7.1, 7.2, 7.3, 7.4, 8.0, 8.1, 8.2
The last snapshot is kept on the hypervisor for incremental backups	Yes
Access to hypervisor OS required	No
Proxy VM required	Yes

Full backup	Supported
Incremental backup	Supported
Restore	Supported
File-level restore	Supported
VM disk exclusion	Supported
Quiesced snapshots	Supported
Snapshots management	Supported
Pre/post command execution	Supported
Access to VM disk backup over iSCSI	Supported
VM name-based policy assignment	Supported
VM tag-based policy assignment	Supported *
Power-on after restore	Supported

** When using Prism Central*

Red Hat OpenShift

Supported backup strategies: Helper pod, Ceph RBD

	Helper pod	Ceph RBD
Minimal version	4.10	4.10
The last snapshot is kept on the system for incremental backups	Yes	Yes
Access to OS required	No	No
Proxy VM required	No	No

	Helper pod	Ceph RBD
Full backup	Supported	Supported
Incremental backup	Not supportd	Supported *
Restore	Supported	Supported
File-level restore	Not supported	Supported *
Volume exclusion	Supported	Supported
Quiesced snapshots	Supported **	Supported **
Snapshots management	Not supported	Not supported
Pre/post command execution	Supported ***	Supported ***
Access to VM disk backup over iSCSI	Not supported	Supported *
Name-based policy assignment	Supported	Supported
Tag-based policy assignment	Supported	Supported
Power-on after restore	Supported	Supported
StatfuSet	Supported	Supported

** When using Ceph RBD as Persistent Volume*

*** Deployment pause*

**** Only 'post'*

Cloud

Amazon EC2

Supported backup strategies: Disk attachment, CBT

	Disk attachement	CBT
The last snapshot is kept on the hypervisor for incremental backups	No	No
Access to hypervisor OS required	No	No
Proxy VM required	Yes	Yes

	Disk attachement	CBT
Full backup	Supported	Supported
Incremental backup	Not supported	Supported
Restore	Supported	Supported
File-level restore	Supported	Supported
VM disk exclusion	Supported	Supported
Quiesced snapshots	Not supported	Not supported
Snapshots management	Supported	Supported
Pre/post command execution	Supported	Supported
Access to VM disk backup over iSCSI	Supported	Supported
VM name-based policy assignment	Supported	Supported

VM tag-based policy assignment	Supported	Supported

Microsoft Azure

Supported backup strategies: Disk attachment, CBT

	Disk attachement	CBT
The last snapshot is kept on the hypervisor for incremental backups	No	No
Access to hypervisor OS required	No	No
Proxy VM required	Yes	Yes

	Disk attachement	CBT
Full backup	Supported	Supported
Incremental backup	Not supported	Supported
Restore	Supported	Supported
File-level restore	Supported	Supported
VM disk exclusion	Supported	Supported
Quiesced snapshots	Not supported	Not supported
Snapshots management	Not supported	Not supported
Pre/post command execution	Supported	Supported
Access to VM disk backup over iSCSI	Supported	Supported
VM name-based policy assignment	Supported	Supported
VM tag-based policy assignment	Supported	Supported
Power-on VM after restore	Not supported (always on)	Not supported (always on)

Microsoft 365

Mailbox messages

Full backup	Supported
Incremental backup	Supported
Restore	Supported
Single item restore	Supported
Restore to another path	Supported
Restore to another account	Supported
Local restore (raw data)	Supported
Restore to PST	Supported

Mailbox messages archive

Full backup	Supported
Incremental backup	Supported
Restore	Supported
Single item restore	Supported
Restore to another path	Supported
Restore to another account	Supported
Local restore (raw data)	Supported
Restore to PST	Supported

Contacts

Full backup	Supported
Incremental backup	Supported
Restore	Supported
Single item restore	Supported
Restore to another path	Supported
Restore to another account	Supported
Local restore (raw data)	Supported
Restore to PST	Not supported

Calendars

Full backup	Supported
Incremental backup	Supported
Restore	Supported
Single item restore	Supported
Restore to another calendar	Supported
Restore to another account	Supported
Local restore (raw data)	Supported
Restore to PST	Not supported

OneDrive for Business

Full backup	Supported
Incremental backup	Supported
Restore	Supported
Single item restore	Supported

Restore to another path	Supported
Restore to another account	Supported
Local restore (raw data)	Supported

Sharepoint sites

Full backup	Supported
Incremental backup	Supported
Restore	Supported
Single item restore	Supported
Restore to another path	Supported
Restore to another site	Not supported
Local restore (raw data)	Supported

Sharepoint pages

Full backup	Supported
Incremental backup	Supported
Restore	Supported
Single item restore	Supported
Restore to another path	Supported
Restore to another site	Not supported
Local restore (raw data)	Supported

Sharepoint list items

Full backup	Supported

Incremental backup	Supported
Restore	Supported
Single item restore	Supported
Restore to another path	Supported
Restore to another site	Not supported

Sharepoint document libraries

Full backup	Supported
Incremental backup	Supported
Restore	Supported
Single item restore	Supported
Restore to another path	Supported
Restore to another site	Not supported
Local restore (raw data)	Supported

Teams channel

Full backup	Supported
Incremental backup	Supported
Restore	Supported
Single item restore	Not supported
Restore to another team	Not supported
Local restore (messeges history)	Supported

Teams 1on1 chat

Full backup	Supported
Incremental backup	Supported
Restore	Supported
Single item restore	Not supported
Restore to another team	Not supported
Local restore (messeges history)	Supported

Teams files

Full backup	Supported
Incremental backup	Supported
Restore	Supported
Single item restore	Supported
Restore to another team	Not supported
Local restore (raw data)	Supported

Storware Backup and Recovery uses Microsoft Teams Export API to export Teams data.

Utilizing the Microsoft Teams Export API generates additional costs for Microsoft tenant.

Storage Providers

Ceph RBD

Source type: RBD Volume (RBD Export/RBD-NBD).

Requires Red Hat Ceph Storage version 4.0 or newer or Ceph v14.2.0 Nautilus or newer

Full backup	Supported
Incremental backup	Supported (RBD snap-diff)
Restore	Supported
Single item restore	Supported
Access to files backup over iSCSI	Supported
Name-based policy assignment	Supported

Nutanix Files (AFS)

Source type: NFS and Samba shares

Full backup	Supported
Incremental backup	Supported (CFT API)
Restore	Supported
Single item restore	Supported
Access to files backup over iSCSI	Supported
Name-based policy assignment	Supported

Nutanix Volume Groups

Source type: Disk attachment

Full backup	Supported
Incremental backup	Supported (CBT API)
Restore	Supported

Single item restore	Supported
Snapshot management	Supported
Access to files backup over iSCSI	Supported
Name-based policy assignment	Supported

Ceph RBD

Source type: RBD Volume (RBD Export/RBD-NBD).

Requires Red Hat Ceph Storage version 4.0 or newer or Ceph v14.2.0 Nautilus or newer

Full backup	Supported
Incremental backup	Supported (RBD snap-diff)
Restore	Supported
Single item restore	Supported
Access to files backup over iSCSI	Supported
Name-based policy assignment	Supported

Backup destinations

Filesystem

Generic filesystem

Supported version	n/a
Syntetic backup	Not supported

Random Access	Supported
Deduplication	Supported *
Encryption	Supported

** When using VDO*

XFS filesystem

Supported version	Linux 4.15 and newer, xfsprogs 4.17 and newer
Syntetic backup	Supported
Random Access	Supported
Deduplication	Supported *
Encryption	Not supported
Pre/post command execution	Supported

** When using VDO*

Object storages

Amazon S3/S3-comatible

Syntetic backup	Not supported
Random Access	Not supported
Deduplication	n/a
Encryption	Supported
Pre/post command execution	Supported

Impossible Cloud

Syntetic backup	Not supported
Random Access	Not supported
Deduplication	n/a
Encryption	Supported
Pre/post command execution	Supported

Google Cloud Storage

Syntetic backup	Not supported
Random Access	Not supported
Deduplication	n/a
Encryption	Supported
Pre/post command execution	Supported

Microsoft Azure Blob Storage

Syntetic backup	Not supported
Random Access	Not supported
Deduplication	n/a
Encryption	Supported
Pre/post command execution	Supported

OpenStack Swift

Syntetic backup	Not supported

Random Access	Not supported
Deduplication	n/a
Encryption	Supported
Pre/post command execution	Supported

Enterprise backup providers

OpenText Data Protector

Supported version	25.1
Syntetic backup	Not supported
Random Access	Not supported
Deduplication	Supported
Encryption	Provider dependent
Pre/post command execution	Supported

Integration plugins

Red Hat Virtualization UI plugin	oVirt we admin 4.3 and newer
oVirt UI Plugin	oVirt we admin 4.3 and newer
Oracle Linux Virtualization Manager UI Plugin	oVirt we admin 4.3 and newer
OpenStack UI Plugin	Horizon 17.0.0 and never

Platform Requirements

System requirements

Operating System

- CentOS Linux Stream 8
- CentOS Linux Stream 9
- Red Hat Enterprise Linux 8.x
- Red Hat Enterprise Linux 9.x
- Oracle Linux 8.x
- Oracle Linux 9.x
- AlmaLinux 8.x
- AlmaLinux 9.x
- Rocky Linux 8.x
- Rocky Linux 9.x
- SUSE Linux Enterprise Server (SLES) 15

> - Using Red Hat Enterprise Linux requires an active subscription.
> - Minimal installation is required.

MariaDB

Storware Backup & Recovery server requires a MariaDB database server.

- Minimum supported MariaDB version: 10.6
- Latest supported MariaDB version: 10.11

We recommend installing MariaDB from the official repository.

ⓘ If you need to install MariaDB packages without accessing an external repository during installation you also can download RPMs and install them manually as described here

Hardware Requirements

Minimum requirements for all-in-one installation (server and node on the same host):

- 64-bit 8 cores processor
- 10 GB RAM
- 20GB free disk space for the operating system and installation
- Free disk space for data staging
  - You can estimate the free space requirement using the following equation: `(Size of the biggest virtual machine) * (number of parallel backup threads)`

Minimum requirements for server (standalone installation):

- 64-bit 4 cores processor
- 4 GB RAM
- 20GB free disk space for the operating system and Storware Backup and Recovery installation

Minimum requirements for node (standalone installation):

- 64-bit 4 cores processor
- 6 GB RAM
- 20GB free disk space for the operating system and installation
- Free disk space for data staging

- You can estimate the free space requirement using the following equation: `(Size of the biggest virtual machine) * (number of parallel backup threads)`

Network requirements

Communication between node and server

Source	Destination	Ports	Description
Node	Server	443/tcp or 8181/tcp	Node ↔ Server communication over HTTPS (port 443 or 8181)
Server	Node	111/tcp, 111/UDP, 2049/tcp, 2049/UDP, ports specified in `/etc/sysconfig/nfs` - variables `MOUNTD_PORT` (TCP and UDP), `STATD_PORT` (TCP and UDP), `LOCKD_TCPPORT` (TCP), `LOCKD_UDPPORT` (UDP)	NFS access to browse mountable backups and logs from administrative portal (using IP that is detected as the source IP - shown in the Node list in the portal)

Network consideration

- Depending on where the node is located you need to verify if data will not pass via low-bandwidth links.
- Access to the internet network from the node may be required in the following scenarios:
  - Installation, when using the external repositories

    - Backup and restore of Amazon EC2, Google Cloud Platform, Azure Cloud and M365

- Node requires connectivity with backup destinations
- Node needs connectivity with the Hypervisor or Hypervisor Manager.
- If a netcat transfer is used for Red Hat Virtuallization/oVirt/Oracle Linux VM/Proxmox VE/KVM stand-alone environments - **16000-16999** ports must be reachable from the hypervisors to the node which is responsible for those hypervisors.

Nutanix AHV

Disk attachment

Connection URL: `https://PRISM_HOST:9440/api/nutanix/v3` (Prism Central or Prism Elements)

> **Note:** when connecting via Prism Central, the same credentials will be used to access all Prism Elements

Source	Destination	Ports	Description
Node	Prism Elements (and optionally Prism Central if used)	9440/tcp	API access to the Nutanix manager

Network Ports

OpenStack

Disk attachment

Connection URL: `https://KEYSTONE_HOST:5000/v3`

Source	Destination	Ports	Description
Node	Keystone, Nova, Glance, Cinder	ports that were defined in endpoints for OpenStack services	API access to the OpenStack management services - using endpoint type that has been specified in hypervisor manager details
Node	Ceph monitors	3300/tcp, 6789/tcp	if Ceph RBD is used as the backend storage - used to collect changed-blocks lists from

SSH transfer

Connection URL: `https://KEYSTONE_HOST:5000/v3`

> **Note:** You also must provide SSH credentials to all hypervisors that have been detected during inventory sync

Source	Destination	Ports	Description
Node	Hypervisor	22/tcp	SSH access
Hypervisor	Node	netcat port range defined in node configuration - by default 16000-16999/tcp	optional netcat access for data transfer
Node	Ceph monitors	3300/tcp, 6789/tcp, 10809/tcp	if Ceph RBD is used as the backend storage - used for data transfer over NBD

OpenNebula

Disk attachment

Connection URL: `https://MANAGER_HOST`

Source	Destination	Ports	Description
Node	Manager Host	XML-RPC API port - 2633/tcp by default	API access to the OpenNebula management services

oVirt/RHV/OLVM

Export storage domain

Connection URL: `https://RHV_MGR_HOST/ovirt-engine/api`

Source	Destination	Ports	Description
Node	oVirt/RHV/OLVM manager	443/tcp	oVirt/RHV/OLVM API access

		If Node is hosting staging space: 111/tcp, 111/UDP, 2049/tcp, 2049/UDP, ports specified in `/etc/sysconfig/nf`	

Disk attachment

Connection URL: `https://MANAGER_HOST/ovirt-engine/api`

Source	Destination	Ports	Description
Node	oVirt/RHV/OLVM manager	443/tcp	oVirt/RHV/OLVM API access

Disk Image Transfer

Connection URL: `https://MANAGER_HOST/ovirt-engine/api`

Source	Destination	Ports	Description
Node	oVirt/RHV/OLVM manager	443/tcp	oVirt/RHV/OLVM API access
Node	oVirt/RHV/OLVM hypervisor	54322/tcp	oVirt/RHV/OLVM ImageIO services - for data transfer (primary source)
Node	oVirt/RHV/OLVM manager	54323/tcp	oVirt/RHV/OLVM ImageIO services - for data transfer (fallback to ImageIO Proxy)

SSH Transfer

Connection URL: `https://MANAGER_HOST/ovirt-engine/api`

> **Note:** You also must provide SSH credentials to all hypervisors that have been detected during inventory sync

Source	Destination	Ports	Description
Node	oVirt/RHV/OLVM manager	443/tcp	oVirt/RHV/OLVM API access
Node	oVirt/RHV/OLVM hypervisor	22/tcp	SSH access for data transfer
oVirt/RHV/OLVM hypervisor	Node	netcat port range defined in node configuration - by default 16000-16999/tcp	optional netcat access for data transfer

Change-Block Tracking

Connection URL: `https://MANAGER_HOST/ovirt-engine/api`

Source	Destination	Ports	Description
Node	oVirt/RHV/OLVM manager	443/tcp	oVirt/RHV/OLVM API access
Node	oVirt/RHV/OLVM hypervisor	54322/tcp	oVirt/RHV/OLVM ImageIO services - for data transfer (primary source)
Node	oVirt/RHV/OLVM manager	54323/tcp	oVirt/RHV/OLVM ImageIO services - for data transfer (fallback to ImageIO Proxy)

Oracle VM

Export storage domain

Connection URL: `https://MANAGER_HOST:7002`

Source	Destination	Ports	Description
Node	OVM manager	7002/tcp	OVM API access
Hypervisor	Node	If Node is hosting staging space: 111/tcp, 111/UDP, 2049/tcp, 2049/UDP, ports specified in `/etc/sysconfig/nfs` - variables `MOUNTD_PORT` (TCP and UDP), `STATD_PORT` (TCP and UDP), `LOCKD_TCPPORT` (TCP), `LOCKD_UDPPORT` (UDP), otherwise check the documentation of your NFS storage provider	if staging space (export storage repository) is hosted on the Node - NFS access
Node and hypervisor	shared NFS storage	check the documentation of your NFS storage provider	if staging space (export storage repository) is hosted on the shared storage - NFS access

Citrix XenServer/xcp-ng

> **Note:** all hosts in the pool must be defined

Single image (XVA-based)

Source	Destination	Ports	Description

Node	Hypervisor	443/tcp	API access (for data transfer management IP is used, unless `transfer NIC` parameter is configured in hypervisor details)

Changed-Block Tracking

Source	Destination	Ports	Description
Node	Hypervisor	443/tcp	API access (for data transfer management IP is used, unless `transfer NIC` parameter is configured in hypervisor details)
Node	Hypervisor	10809/tcp	NBD access (data transfer IP is returned by hypervisor)

KVM/Xen stand-alone

SSH transfer

Source	Destination	Ports	Description
Node	Hypervisor	22/tcp	SSH access
Hypervisor	Node	netcat port range defined in node configuration - by default 16000-16999/tcp	optional netcat access for data transfer

Node	Ceph monitors	3300/tcp, 6789/tcp, 10809/tcp	if Ceph RBD is used as the backend storage - used for

Proxmox VE

Export storage repository

Source	Destination	Ports	Description
Node	Hypervisor	22/tcp	SSH access
Hypervisor	Node	If Node is hosting staging space: 111/tcp, 111/UDP, 2049/tcp, 2049/UDP, ports specified in `/etc/sysconfig/nfs` - variables `MOUNTD_PORT` (TCP and UDP), `STATD_PORT` (TCP and UDP), `LOCKD_TCPPORT` (TCP), `LOCKD_UDPPORT` (UDP), otherwise check the documentation of your NFS storage provider	if staging space (export storage domain) is hosted on the Node - NFS access
Node and hypervisor	shared NFS storage	check the documentation of your NFS storage provider	if staging space (export storage domain) is hosted on the shared storage - NFS access

SSH transfer

Source	Destination	Ports	Description
Node	Hypervisor	22/tcp	SSH access
Hypervisor	Node	netcat port range defined in node configuration - by default 16000-16999/tcp	optional netcat access for data transfer

Microsoft 365

Source	Destination	Ports	Description
Node	Microsoft 365	443/tcp	Microsoft 365 API access

You can find more detailed description about Office 365 URLs and IP address ranges on this page.

To successfully synchronize M365 user account, it must fulfill following requirements:

- has an email,
- is not filtered by location, country or office location (user filter in UI),
- field `user type` is set to `Member`,
- has a license or is a shared mailbox.

Security Requirements

User Permissions

User `vprotect` must be a member of group "disk".

Sudo privileges are required for the following commands:

Data Protector for Cloud Workloads Node:

- `/usr/bin/targetcli`
- `/usr/sbin/exportfs`
- `/usr/sbin/kpartx`
- `/usr/sbin/dmsetup`
- `/usr/bin/qemu-nbd`
- `/usr/bin/guestmount`
- `/usr/bin/fusermount`
- `/bin/mount`
- `/bin/umount`
- `/usr/sbin/parted`
- `/usr/sbin/nbd-client`
- `/usr/bin/tee`
- `/opt/vprotect/scripts/vs/privileged.sh`
- `/usr/bin/yum`
- `/usr/sbin/mkfs.xfs`
- `/usr/sbin/fstrim`
- `/usr/sbin/xfs_growfs`
- `/usr/bin/docker`
- `/usr/bin/rbd`
- `/usr/bin/chown`
- `/usr/sbin/nvme`
- `/bin/cp`
- `/sbin/depmod`
- `/usr/sbin/modprobe`
- `/bin/bash`
- `/usr/local/sbin/nbd-client`
- `/bin/make`

Data Protector for Cloud Workloads Server:

- `/opt/vprotect/scripts/application/vp_license.sh`

- `/bin/umount`
- `/bin/mount`

SELinux

PERMISSIVE - currently it interferes with the mountable backups (file-level restore) mechanism. Optionally can be changed to ENFORCING if the file-level restore is not required.

Sizing Guide

The best strategy is to **plan** your backup environment/procedure **before implementing** it. In this chapter, we have collected generic hints and guides which you might find useful while thinking about your Data Protector for Cloud Workloads implementation.

1. Collect information about the `TotalSizeOfData` to be protected in your environment
   - this is the size of your VMs/Storage that will be transferred within the backup window
     - for general sizing, assume all backups to be full
   - if your staging space is separate from the backup destination, also check what are the biggest VMs/Storage that may end up in your staging area
2. Assume `BackupWindow` length - backups are usually executed overnight, so 10h-12h is common practice
3. Run a **test transfer** on a test file to estimate the maximum achievable bandwidth per thread (`SingleThreadTransfer`) from the hypervisor (or manager) to the node
   - we recommend 10 simultaneous transfers with the result divided by 10 threads to see if other limitations of the environment do not impact the total transfer rate (one such common limitation is disk read performance on the virtualization platform)
   - all the methods usually use snapshots to do backup - check if snapshot removal in your environment does not take a significant amount of time, as it is a highly resource-intensive operation that impacts overall backup time - especially when running multiple export jobs in parallel
4. Estimate **the number of the nodes**
   - required bandwidth per node:

     `RequiredBandwidth = TotalSizeOfData / BackupWindow`
   - the total number of export tasks (note that other aspects such as snapshot handling, file system scanning during export, and infrastructure bottlenecks when using multiple threads will usually impact the overall speed):

     `TotalNumberOfExportTasks = RequiredBandwidth / (70% * SingleThreadTransferSpeed)`

- the number of nodes:

  `NumberOfTheNodes = TotalNumberOfExportTasks / 10` **(10 is the recommended maximum number of export tasks per Node)**
- note that the `TotalSizeOfData` does not mean that it is only a full backup, as you can mix full and incremental backups
- granularity is a single hypervisor or storage provider, so at the maximum, you cannot have more nodes than hypervisor storage providers in your environment
- if you have multiple clusters and you want to use the disk-attachment method, this automatically implies a minimum of 1 node per cluster

5. Estimate the total **store rate** in the backup destination
   - if multiple nodes are required, add up the total amount of data from all nodes
   - if your backup destination is accessible over LAN
     - do a test transfer from the node to the backup destination to verify if the performance on the backup destination is able to receive such a load
6. `NumberOfExportTasksPerNode`
   - we recommend using the same node configuration for multiple nodes, so the same limit value will be applied to all nodes sharing the configuration
   - this implies that we recommend assuming this value as follows (rounded down):

     `NumberOfExportTasksPerNode = TotalNumberOfExportTasks / NumberOfTheNodes`
7. `NumberOfStoreTasksPerNode` usually depends on destination backup performance
   - we recommend a value equal to the `NumberOfExportTasksPerNode` or higher
   - reduce this value only if your backup provider starts to have significant I/O latency eventually leading to a slower write rate than with the lower number of threads - this will typically result in higher staging space occupation as backups will be kept for a longer period of time in the temporary space
8. **Node resource requirements:**
   - **CPU**: Assume 0.5 CPU per task, minimum 2 cores - rounded up to get the full core count supported by the hypervisor or physical server (it may be

   required to round up 2.5 cores to 4 vCPUs if the hypervisor on which the node is deployed doesn't allow to 3 vCPUs to be assigned)

    - if SSH transfer (without netcat) or client-side deduplication is used (like VDO) assume 1 CPU per task
  - **Memory**: 256 MB RAM per task, with a minimum of 2GB
  - **Staging space:** if not shared with the backup destination - the biggest VM/Storage multiplied by the number of tasks
  - when counting tasks for each node assume: `NumberOfExportTasksPerNode + NumberOfStoreTasksPerNode`

9. **Server resource requirements:**
   - **CPU**: Assume 0.5 CPU per task, minimum 2 cores - rounded up to full core count supported by the hypervisor or physical server (it may be required to round up 2.5 cores to 4 vCPUs if the hypervisor on which the node is deployed doesn't allow 3 vCPUs to be assigned)
   - **Memory**: 256 MB RAM per task, with a minimum of 6GB
   - when counting tasks assume:

     `TotalNumberOfExportTasks + TotalNumberOfStoreTasks`
10. Finally, if the resulting node count is too big:
    - divide your VMs into multiple backup policies with separate schedules so that some full backups of your VMs will be done on Monday, some on Tuesday, while the rest will run incremental backups at the same time - this will reduce the value of `TotalSizeOfData` in the previous equations
    - check if the backup window cannot be extended
      - exports usually impact infrastructure more, while store tasks can also safely be done during the day
      - Data Protector for Cloud Workloads will start backups only within the backup window, but once the tasks are started, they may continue even after your backup window ends

Notes on sizing using different setups

Disk-attachment methods

- read data from locally attached drives (which may use LAN or SAN behind the scenes depending on your virtualization platform setup) and write it to the staging space (local or remote)
- run a read test from one device to the staging storage to estimate the processing rate
- this method requires usually 1 node per cluster, so treat each cluster separately
- this method also requires time for attachment/detachment of drives

Export storage repository methods

- export data from a specific host to the staging space of Data Protector for Cloud Workloads via NFS
- run a test export on any VM to estimate the export speed in your environment
- export methods also usually have limitations on the hypervisor side, so OVM can process only 1 export job simultaneously using a specific set of storage repositories (on which the VM disks reside, and to which the VM is being exported), which may impact overall performance - consult your hypervisor documentation to check for export process limitations
- it is common to share a backup destination with staging space from an external backup provider via NFS (not from each node) so that exports are done directly to the backup destination storage

Direct export from hypervisors or underlying storage

- if you can enable the netcat in SSH Transfer methods, it should result in 2-3 times faster transfer rates compared to standard SSH
- export tasks run against stand-alone hypervisors will be automatically balanced, while those managed by hypervisor managers will be subject to global and per source limits
  - this means that if you configure a maximum number of export tasks per source to 5 and the global number to 10, you will have no more than 5 export tasks running against a single manager regardless of the number of hosts
- when using Ceph RBD in KVM stand-alone or the OpenStack SSH Transfer method, the actual transfer is done directly from Ceph monitors, and this is the

network path that needs to be checked when estimating bandwidth - use `rbd export` or mount a test volume over RBD-NBD to test it

Backup destination

- if you plan to use common storage for the staging space and backup destination, your reads from the source will be limited by the write speed of your backup destination
- make sure you have the appropriate bandwidth between the nodes and the backup destination
- verify if the backup destination is able to process IOPS coming from multiple sources - it is common to assume the export rate as the minimum required store rate

Deployment

Deployment

1. Start with **overview**:
   - Architecture
   - Support Matrix
   - Platform Requirements
2. Check **where node should be installed** for your environment:
   - Virtual Environments
   - Microsoft 365
   - Applications
   - Storage Providers
3. **Install** using one of the following options:
   - Quick Installation using all-in-one script (recommended)
   - Installation using Ansible playbook
   - Installation with RPMs

   Regardless of the installation option you choose:

   - The node requires **staging space** - assume a number of concurrent export and store tasks and multiply it by the biggest VM size (**for example:** 6 export tasks + 4 store tasks * 100 GB should require around 1 TB)
   - The Staging space configuration will guide you to prepare storage on the spare drive
   - Data Protector for Cloud Workloads is installed in the `/opt/vprotect` folder and staging space is assumed to be in `/vprotect_data` - these are the defaults and should not be changed.
4. Run the **configuration wizard** (Initial configuration), where you can (or do manually the following steps):
   - upload the license
   - configure connection to the source you would like to protect

    - configure backup destination - we recommend to use Synthetic File System
    - configure backup SLA (policies and schedules)
    - configure backup of your internal DB (for DR purposes)
5. Once you have configured source, backup destination and backup SLA - **initiate backup and restore** operations:

- Virtual Environments
- Microsoft 365
- Applications
- Storage Providers

Quick Installation using all-in-one script

Using this method of installation you will deploy the server and node on the same host. The installation script will perform the following actions:

- Install the server
- Install the node
- Generate an SSL certificate based on the hostname

The installation script is deploying components using the Ansible playbooks

Installation steps

1. Log in to the machine using SSH
2. The installation requires root privileges
3. Download Data Protector for Cloud Workloads package.
4. Extract this package on the host where you're installing it:

   ```
   tar xvf DP-for-Cloud-Workloads-XXX.tgz
   ```

5. Move extracted repository to temp directory:

   ```
   mv elX/* /tmp/DP-for-Cloud-Workloads-repo
   ```

6. As a root run:

   ```
   ./DP-for-Cloud-Workloads-local-install.sh
   ```

7. Move extracted repository to temp directory:

   ```
   mv elX/* /tmp/DP-for-Cloud-Workloads-repo
   ```

8. If your netork is using proxy server ppdate the proxy details in /opt/vprotect/vprotect.env for cloud backup/restore and restart the server and node services.

Using script to install only server or only node component

This script allows to install just server or node (which can be registered to the existing server).

Server only installation

Before running the installation command - export the following variable:

```
export SBR_INSTALL_NODE=n
```

Node only installation

Before running the installation command - export the following variables:

```
export SBR_INSTALL_SERVER=n
export SBR_SERVER_FQDN=your.server.host.com
export SBR_NODE_NAME=your-node-name
```

where `your.server.host` should be a FQDN of your server component, and `your-node-name` should be a unique name for a node being installed. Optionally, you can export `SBR_ADMIN_USER` if you want to register your nodes using non-admin accounts.

Disabling password prompts

If you want to use this script without being prompted for admin user or database password you can export `DP_ADMIN_PASS` and `DP_DB_PASS` variables.

Post-installation

Now you should be able to log in to Data Protector for Cloud Workloads Server using https://<DP4CW_server_IP> with local node registered and running.

Remember to prepare your staging space as described in the Staging space configuration.

Now proceed with the Initial configuration instructions to configure access to the hypervisors and backup destinations.

By default, Data Protector for Cloud Workloads has one admin account - `admin` with the password `vPr0tect` (with a zero)

Update

This package replaces previous installation. Database model and any dependencies may be updated during update. All configuration stored in the database or migrated to the new model automatically if necessary.

Server Upgrade

1. Create database backup - run as root over SSH on the Data Protector for Cloud Workloads Server

   ```
   /opt/vprotect/scripts/backup_db.sh /path/to/backup/file.tgz
   ```

2. Extract this package on the hosts with Data Protector for Cloud Workloads Server or Node:

   ```
   tar xvf DP-for-Cloud-Workloads-XXX.tgz
   ```

3. Update Data Protector for Cloud Workloads Server using RPMs in elX folder

   ```
   yum update elX/DP-for-Cloud-Workloads-server-XXX.elX.x86_64.rpm
   ```

4. Update each Data Protector for Cloud Workloads Nodes using RPMs in elX folder

```
yum update elX/DP-for-Cloud-Workloads-node-XXX.elX.x86_64.rpm
```

5. Update each Data Protector for Cloud Workloads Cloud Server:

   ```
   yum update elX/DP-for-Cloud-Workloads-cloudserver-XXX.elX.x86_64.rpm
   ```

6. Update each Data Protector for Cloud Workloads Cloud Agent:

   ```
   yum update elX/DP-for-Cloud-Workloads-cloudagent-XXX.elX.x86_64.rpm
   ```

7. Log in to the Data Protector for Cloud Workloads Server using https://<DP4CW_server_IP> with nodes updated and running.

Notice, that you may need to refresh your browser cache after update - for Chrome use CTRL+SHIFT+R (Windows/Linux) / CMD+SHIFT+R (MacOS)

Downgrade

1. Downgrade Server with yum:

   ```
   yum downgrade vprotect-server
   ```

2. On the Data Protector for Cloud Workloads Server host stop the Server service, restore database using your DB password and start server again (these can be found in `/opt/vprotect/payara.properties` file)

   ```
   systemctl stop vprotect-server
   mysql -u vprotect -pDBPASSWORD -e "drop database vprotect"
   mysql -u vprotect -pDBPASSWORD -e "create database vprotect"
   gunzip < PATH_TO_GZIPPED_BACKUP | mysql -u vprotect -pDBPASSWORD vprot
   systemctl start vprotect-server
   ```

3. On the Data Protector for Cloud Workloads Nodes hosts - downgrade nodes with

   ```
   yum downgrade vprotect-node
   ```

4. Make sure all nodes are running and optionally start service on each Data Protector for Cloud Workloads host

   ```
   systemctl start vprotect-node
   ```

Deinstallation

1. Remove packages with yum:

   ```
   yum remove vprotect-server
   yum remove vprotect-node
   ```

2. To remove configuration files:

   ```
   rm -rf /opt/vprotect
   ```

3. To remove all remaining MariaDB data:

   ```
   yum remove mariadb mariadb-server
   rm -rf /var/lib/mysql
   rm /etc/my.cnf
   Optional step:
   rm ~/.my.cnf
   rm -f /var/log/mariadb
   rm -f /var/log/mariadb/mariadb.log.rpmsave
   rm -rf /usr/lib64/mysql
   rm -rf /usr/share/mysql
   ```

4. To remove users:

   ```
   userdel vprotect
   userdel mysql
   ```

5. To remove certificate generated during installation:

   ```
   keytool -delete -keystore /usr/lib/jvm/jre/lib/security/cacerts -alias
   ```

   Replace **CERT_ALIAS** with the hostname of your OS. Default password for keystore is "changeit".

Installation using Ansible playbook

You can install the complete Data Protector for Cloud Workloads solution using the following 2 roles, available on Ansible Galaxy:

- Data Protector for Cloud Workloads Server: https://galaxy.ansible.com/xe0nic/ansible_vprotect_server
- Data Protector for Cloud Workloads Node: https://galaxy.ansible.com/xe0nic/ansible_vprotect_node

This approach installs a server and one or more nodes on remote hosts and generates an SSL certificate based on the server hostname. The end result should be the same as an RPM-based installation without the staging setup. Configuration (such as backup destination definition or hypervisor connectivity) still needs to be done after installation. You can also add more nodes in the future if necessary.

Prerequisites

> You can find list of all supported operating systems in this chapter

You need to prepare CentOS or RHEL minimal for Data Protector for Cloud Workloads (both roles can be installed on the same or different hosts). The Ansible control host should have Ansible installed so that it uses Python 3.x.0

This example assumes that you have `root` access to this host and you have configured your Ansible to connect with SSH public keys to your host. For example:

generate key:
`ssh-keygen -f ~/.ssh/id_rsa -P ""`

and copy it to your CentOS/RHEL box:
`ssh-copy-id -i ~/.ssh/id_rsa.pub root@YOUR_HOST`

The nodes will communicate with the Data Protector for Cloud Workloads Server via port 8181, so they need to be able to access it using the server's FQDN (this needs to be resolvable).

Installation

> Before installing Data Protector for Cloud Workloads we highly recommend doing a system update and reboot.

This example assumes that you want to install both the Data Protector for Cloud Workloads Server and Node **using a single playbook** and **on the same host.** However, keep in mind that you may also install them separately by providing different target hosts and using separate playbooks like in the examples in the readme roles (links above).

Run these on the system from which you run Ansible playbooks:

- Download the installation package from the Micro Focus download page
- Upload the installation package to all hosts (server and nodes)
- On each host, extract the archive:

  ```
  tar xvf your-package.tgz
  ```

- The installation package contains a package repository in the `el8` folder which will be added automatically by the installation script. Move the extracted repository to the temp directory.

  ```
  mv elX/* /tmp/DP-for-Cloud-Workloads-repo
  ```

- Install Ansible roles:

  ```
  ansible-galaxy install xe0nic.ansible_vprotect_server
  ansible-galaxy install xe0nic.ansible_vprotect_node
  ```

- Install additional collections

  ```
  ansible-galaxy install -r  ~/.ansible/roles/xe0nic.ansible_vprotect_se
  ansible-galaxy install -r  ~/.ansible/roles/xe0nic.ansible_vprotect_nc
  ```

- Create a playbook directory and change it to a working directory, i.e: `mkdir dp4cw && cd dp4cw`
- Create an inventory file - `hosts` and refer to the location to where you extracted the repository
  - in this example we have specified one node and server, but you can define more nodes (each one must be in a separate line and have a unique node name)
  - the server can be on a different host
  - we recommend having at least one node installed together with the server to run DB backups

  ```
  [all:vars]
  ansible_user = root
  vprotect_repo = file:///tmp/DP-for-Cloud-Workloads-repo
  admin_pass=password
  db_pass=password

  [server]
  192.168.1.2

  [nodes]
  192.168.1.2 node_name=node1
  ```

  where:

  - `admin_pass` - password for admin user
  - `db_pass` - password for mysql root user
  - `node_name` - name under which node will be registered

> If you don't provide password for admin user and mysql root user, it will be set to **vPr0tect**

- Create a playbook file - `site.yml`:

```
---

- hosts: server
  roles:
  - xe0nic.ansible_vprotect_server

- hosts: nodes
  roles:
  - xe0nic.ansible_vprotect_node
```

- Run the playbook: `ansible-playbook -i hosts site.yml`
- After installation, you should be able to log in to your Data Protector for Cloud Workloads Server: `https://<DP4CW_server_IP>` and your nodes should be registered and running. By default, Data Protector for Cloud Workloads has one admin account - `admin` with the password `vPr0tect` (with a zero).
- After the initial log in you can configure single sign-on using LDAP or Keycloak.
- Remember to prepare your staging space as described in the Staging space configuration.
- Now proceed with the Initial configuration instructions to configure access to the hypervisors and backup destinations.

Variables

These two roles use just a few variables. Both plays use the `server_fqdn` variable. If not defined, the server play sets the variable `server_fqdn` to the hostname reported by the OS on which it is installed. The server play will generate an SSL certificate for this FQDN, and node play will automatically use this value if defined. You can also provide this variable manually (either in the `hosts` file or with the extra vars switch in the `ansible-playbook` command, `-e "server_fqdn=vprotect.server.local"`

Node play needs a `node_name` for the registration process. If not provided, it will just use the hostname reported by the OS, however, keep in mind that it needs to be **unique** for each node. We recommend that you set them in the host inventory file.

Optionally, you may want to set a `db_password` for the root DB access which is set during server installation. Note, that the Server service uses its own account with an auto-generated password.

By default, Data Protector for Cloud Workloadst uses MariaDB 10.4 for CentOS - you can control the source, distribution and version of your MariaDB with the following variables (with their respective default values):

```
mariadb_version: "10.4"
mariadb_distro: "centos7-amd64"
mariadb_repo_url: "http://yum.mariadb.org/{{ mariadb_version }}/{{ mariad
mariadb_repo_gpg_key: "https://yum.mariadb.org/RPM-GPG-KEY-MariaDB"
```

Installation with RPMs

Procedure

Create a repository file

The repository file must be created on each host where the product components will be deployed.

1. Download the Data Protector for Cloud Workloads packages
2. Extract your package (replace the name with the downloaded package name):

   ```
   tar xvf your-package.tgz
   ```

3. Move the extracted repository to the temp directory.

   ```
   mv el9 /tmp/DP-for-Cloud-Workloads-repo
   ```

4. Create a repository file `/etc/yum.repos.d/vProtect.repo` with the following content:

For Red Hat Enterprise Linux 8 and compatible

```
# Data Protector for Cloud Workloads - Enterprise backup solution for vir
[DP-for-Cloud-Workloads]
name = Data Protector for Cloud Workloads
baseurl = file:///tmp/DP-for-Cloud-Workloads-repo
gpgcheck = 0
```

For Red Hat Enterprise Linux 9 and compatible

```
# Data Protector for Cloud Workloads - Enterprise backup solution for vir
[DP-for-Cloud-Workloads]
name = Data Protector for Cloud Workloads
baseurl = file:///tmp/DP-for-Cloud-Workloads-repo
gpgcheck = 0
```

For SUSE Linux Enterprise Server 14 and compatible

```
# Data Protector for Cloud Workloads - Enterprise backup solution for vir
[DP-for-Cloud-Workloads]
name = Data Protector for Cloud Workloads
baseurl = file:///tmp/DP-for-Cloud-Workloads-repo
gpgcheck = 0
```

Create a repository file for MariaDB

Installing MariaDB is required only on the host where the server is deployed.

1. Generate repository file at MariaDB download site
2. Copy and paste the generated repo file into `/etc/yum.repos.d/MariaDB.repo`

Red Hat Enterprise Linux or and compatible

Server installation

1. Install package "sudo":

   ```
   dnf install sudo
   ```

2. Install the server using the following command:

   ```
   dnf install vprotect-server
   ```

Node installation

1. Install the node using the following command

   ```
   dnf install vprotect-node
   ```

SUSE Linux Enterprise Server and compatible

Server installation

1. Add Desktop Application Tools module:

   ```
   SUSEConnect -p sle-module-desktop-applications/15.4/x86_64
   ```

2. Add Development tools module:

   ```
   SUSEConnect -p sle-module-development-tools/15.4/x86_64
   ```

3. Install package "sudo":

   ```
   zypper install sudo
   ```

4. Install the Storware Backup and Recovery server using the following command:

   ```
   zypper install vprotect-server
   ```

Node installation

1. Install the node using the following command

   ```
   zypper install vprotect-node
   ```

Server configuration

1. Configure access to the database. Run the following command:

   ```
   vprotect-server-configure
   ```

2. Start the server service:

   ```
   systemctl start vprotect-server
   ```

Open a firewall port

By default, the server service listens on port 8181. Open the port using the following commands:

```
firewall-cmd --add-port=8181/tcp --permanent
firewall-cmd --complete-reload
```

(optional) Forward the default HTTPS port 443 to port 8181:

```
/opt/vprotect/scripts/./ssl_port_forwarding_firewall-cmd.sh
```

Node staging space

- Prepare your staging space (on the Data Protector for Cloud Workloads Node host):
    - If you just started with Data Protector for Cloud Workloads, and do not know what is staging space, follow the steps described in the Staging space configuration
    - **if your path is different than** `/vprotect_data` it is recommended to create a symlink `/vprotect_data` pointing to your staging space mount point, e.g.:

      ```
      ln -s /mnt/staging /vprotect_data
      ```

Node registration

1. Each installed node needs to be registered in the server:

   ```
   vc node inst register --name=<node name> --login=<admin user> --passwo
   ```

   where:

   - <node name> - the name under which the node will appear in the system
   - <admin user> - the login of the administrative user
   - <server address>:<port> - address and port of the installed server

   Example:

   ```
   vc node inst register --name=node1 --login=admin --password=vPr0tect -
   ```

2. Start the node service:

```
systemctl start vprotect-node
```

3. Run the script to configure the operating system. Script changes the QEMU user/group to vprotect, disables SELinux, adds product to the disk group and sudoers policy to allows run privileged commands:

   ```
   vprotect-node-configure
   ```

4. Reboot the host to apply the operating system changes:

   ```
   reboot
   ```

Post-installation

Now you should be able to log in to Data Protector for Cloud Workloads Server using https://<DP4CW_server_IP> with local node registered and running.

Remember to prepare your staging space as described in the Staging space configuration.

Now proceed with the Initial configuration instructions to configure access to the hypervisors and backup destinations.

By default, Data Protector for Cloud Workloads has one admin account - `admin` with the password `vPr0tect` (with a zero)

Update

This packages replaces previous installation. Database model and any dependencies may be updated during update. All configuration stored in the database or migrated to the new model automatically if necessary.

Server Upgrade

1. Create database backup - run as root over SSH on the Data Protector for Cloud Workloads Server

   ```
   /opt/vprotect/scripts/backup_db.sh /path/to/backup/file.tgz
   ```

2. Extract this package on the hosts with Data Protector for Cloud Workloads Server or Node:

   ```
   tar xvf DP-for-Cloud-Workloads-XXX.tgz
   ```

3. Update Data Protector for Cloud Workloads Server using RPMs in elX folder

   ```
   yum update elX/DP-for-Cloud-Workloads-server-XXX.elX.x86_64.rpm
   ```

4. Update each Data Protector for Cloud Workloads Nodes using RPMs in elX folder

   ```
   yum update elX/DP-for-Cloud-Workloads-node-XXX.elX.x86_64.rpm
   ```

5. Update each Data Protector for Cloud Workloads Cloud Server:

   ```
   yum update elX/DP-for-Cloud-Workloads-cloudserver-XXX.elX.x86_64.rpm
   ```

6. Update each Data Protector for Cloud Workloads Cloud Agent:

   ```
   yum update elX/DP-for-Cloud-Workloads-cloudagent-XXX.elX.x86_64.rpm
   ```

7. Log in to the Data Protector for Cloud Workloads Server using https://<DP4CW_server_IP> with nodes updated and running.

Notice, that you may need to refresh your browser cache after update - for Chrome use CTRL+SHIFT+R (Windows/Linux) / CMD+SHIFT+R (MacOS)

Downgrade

Downgrade Server with yum:

```
yum downgrade vprotect-server
```

On the Data Protector for Cloud Workloads Server host stop the Server service, restore database using your DB password and start server again (these can be found in `/opt/vprotect/payara.properties` file)

```
systemctl stop vprotect-server
mysql -u vprotect -pDBPASSWORD -e "drop database vprotect"
mysql -u vprotect -pDBPASSWORD -e "create database vprotect"
gunzip < PATH_TO_GZIPPED_BACKUP | mysql -u vprotect -pDBPASSWORD vprotect
systemctl start vprotect-serverOn the Data Protector for Cloud Workloads
```

```
yum downgrade vprotect-node
```

Make sure all nodes are running and optionally start service on each Data Protector for Cloud Workloads host

```
systemctl start vprotect-node
```

Uninstall

Remove packages with yum:

```
yum remove vprotect-server
yum remove vprotect-node
```

To remove configuration files:

```
rm -rf /opt/vprotect
```

To remove all remaining MariaDB data:

```
yum remove mariadb mariadb-server
rm -rf /var/lib/mysql
rm /etc/my.cnf
Optional step:
rm ~/.my.cnf
rm -f /var/log/mariadb
rm -f /var/log/mariadb/mariadb.log.rpmsave
rm -rf /usr/lib64/mysql
rm -rf /usr/share/mysql
```

To remove users:

```
userdel vprotect
userdel mysql
```

To remove certificate generated during installation:

```
keytool -delete -keystore /usr/lib/jvm/jre/lib/security/cacerts -alias CE
```

Replace **CERT_ALIAS** with the hostname of your OS. Default password for keystore is "changeit".

Backup Destinations

A backup destination is a storage location where Data Protector for Cloud Workloads keeps VMs, Containers, Cloud, and applications backup copies. To configure a backup destination, you can use the following storage types:

- File System
- Deduplication Appliances
- Object Storage
- Enterprise Backup Providers

The backup destination is defined by the backup provider configuration and retention settings. Each policy can be backed up to the selected backup destination. Backup destinations must be assigned to the nodes in the node configuration.

> **Note:** removal of any backup destination leaves data in the backup provider without an option to re-attach it in the future.

Pre/post access command execution

- Prepare your scripts
  - the pre-script is invoked before every access to the Backup Destination - common usage - create and mount the remote volume
  - the post-script is executed after Node finishes store, restore, and clean-up operations
- The following environment variables are set before each execution - you can use them later in your scripts:
  - `VP_VM_GUID` - GUID of the VM in Data Protector for Cloud Workloads
  - `VP_VM_UUID` - UUID of the VM used by the hypervisor or hypervisor manager
  - `VP_VM_NAME` - the name of the VM

    - `VP_VM_TMP_DIR` - path to the folder containing files in the staging
    - `VP_BD_GUID` - GUID of the Backup Destination being accessed
    - `VP_BD_NAME` - the name of the Backup Destination being accessed
    - `VP_CONTAINER_NAME` - standard container name generated by Data Protector for Cloud Workloads that can be used for names of the volumes (format `<VM-NAME>__<PART-OF-UUID>`, for example `Centos 7__8d3ef6f1`, may contain special characters)
    - `VP_EXPORT_PATH` - an export path from Node Configuration, can be used as the mount root for backup destination volumes
    - `VP_TASK_TYPE` - the name of the task type, e.g.: STORE / RESTORE / DELETE_VM / OLD_BACKUPS_REMOVAL - to distinguish operation type when scripts are being invoked
- Upload your scripts to the **node**, where the `vprotect` user is able to access them
- Optionally, you may need to add a new file in the `/etc/sudoers.d/` directory to enable the `vprotect` user to execute privileged script (like chown operations in some file system locations): `%vprotect ALL=(root) NOPASSWD: /opt/vprotect/scripts/myscripts/privileged.sh`
- Open the "BACKUP DESTINATIONS" section from the left menu:

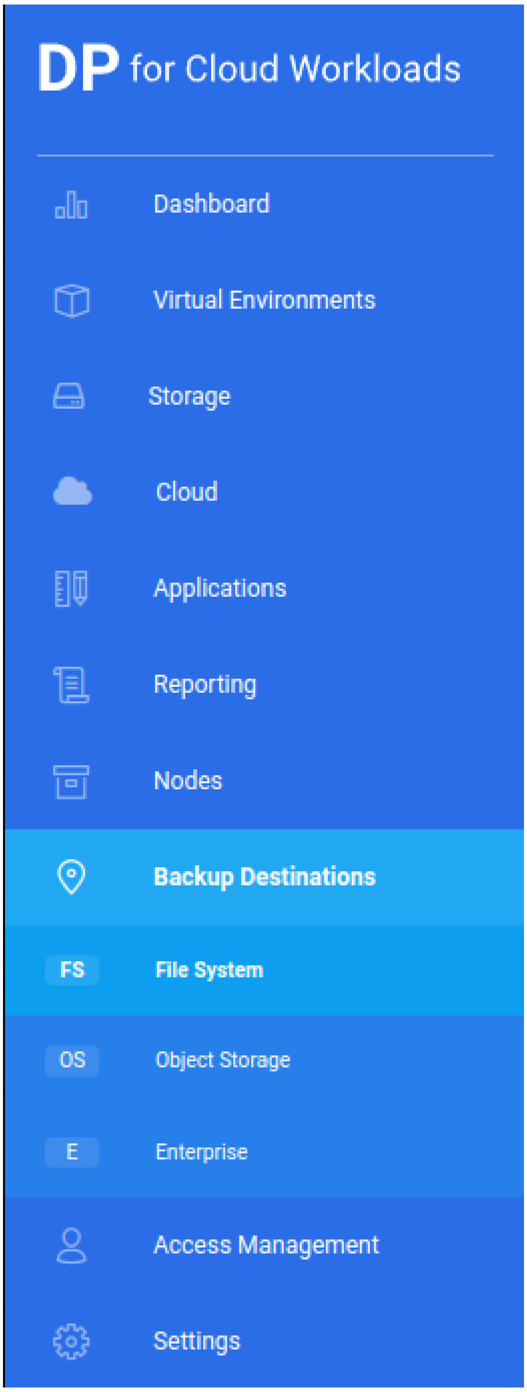

- Open your Backup Destination (click on the name)
- Provide pre/post access command arguments (the first argument is the command executed locally on the **node**):

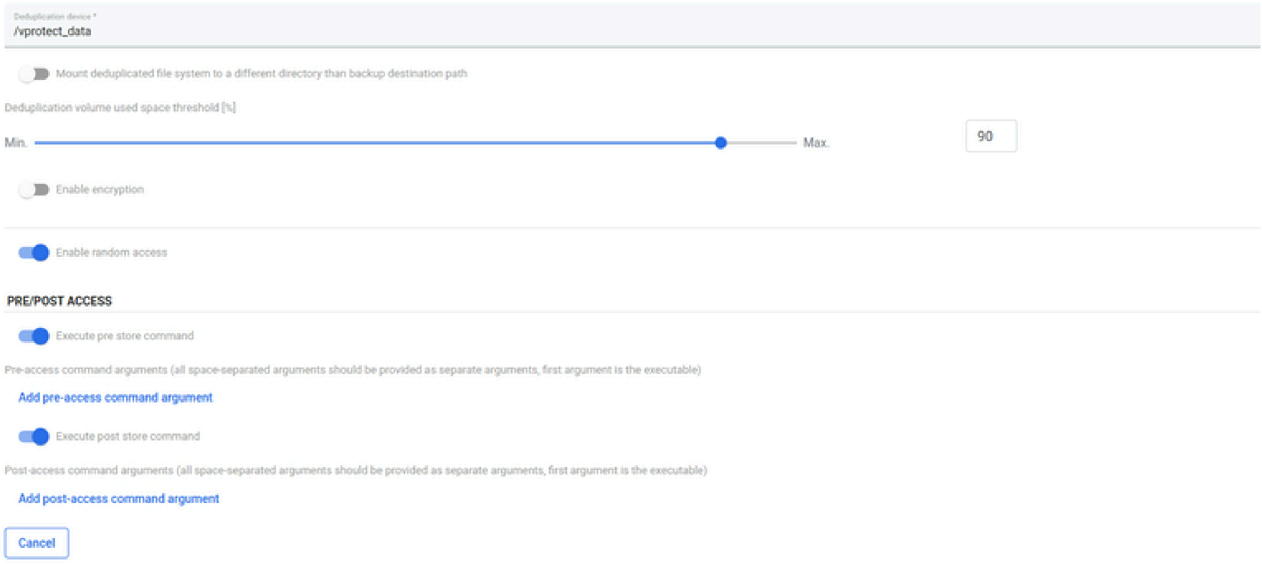

Deduplication device *
/vprotect_data

○ Mount deduplicated file system to a different directory than backup destination path

Deduplication volume used space threshold [%]

Min. ——●— Max. 90

○ Enable encryption

● Enable random access

PRE/POST ACCESS

● Execute pre store command

Pre-access command arguments (all space-separated arguments should be provided as separate arguments, first argument is the executable)

● Execute post store command

Post-access command arguments (all space-separated arguments should be provided as separate arguments, first argument is the executable)

Add pre-access command argument
Add post-access command argument
Cancel

Encryption

Target	Supported	Source	Key stored	Generated	Algorithm
Filesystem	Yes	Data Protector for Cloud Workloads Node	Generated based on metadata in the database. Separated keys are generated per object.	automatically	AES
Filesystem (synthetic/XFS)	Yes	Data Protector for Cloud Workloads Node	Generated based on metadata in the database. Separated keys are generated per object.	automatically	AES
			Generated based on metadata in		

MS Azure Blob Storage	Yes	Data Protector for Cloud Workloads Node	the database. Separated keys are generated per object.	automatically	AES
Amazon S3	Yes	Server Side (Backup Destination own mechanism, not managed by Data Protector for Cloud Workloads)	Generated based on metadata in the database. Separated keys are generated per object.	automatically	n/a
S3 compatible	Yes	Server Side (Backup Destination own mechanism, not managed by Data Protector for Cloud Workloads)	Generated based on metadata in the database. Separated keys are generated per object.	automatically	n/a
Cloudian S3	Yes	Server Side (Backup Destination own mechanism, not managed by Data Protector for Cloud Workloads)	Generated based on metadata in the database. Separated keys are generated per object.	automatically	n/a
		Server Side (Backup Destination own	Generated based on metadata in		

Alibaba Cloud OSS	Yes	mechanism, not managed by Data Protector for Cloud Workloads)	the database. Separated keys are generated per object.	automatically	n/a
Nutanix Objects	Yes	Server Side (Backup Destination own mechanism, not managed by Data Protector for Cloud Workloads)	Generated based on metadata in the database. Separated keys are generated per object.	automaticallya	n/a
OpenStack Swift	Yes	Server Side (Backup Destination own mechanism, not managed by Data Protector for Cloud Workloads)	Generated based on metadata in the database. Separated keys are generated per object.	automatically	n/a
Scality Ring	Yes	Server Side (Backup Destination own mechanism, not managed by Data Protector for Cloud Workloads)	Generated based on metadata in the database. Separated keys are generated per object.	automatically	n/a

File System

File System

This section presents the key steps necessary for configuring a file system as your backup destination. You can use a:

- local File system or remote (NFS, SMB, etc.) or attach a block device with enabiling there Virtual Data Optimizer (VDO)
- Synthetic File System

File system

In this section, we'll show you how to set up a file system (it can be a local or remote file system, but this example assumes that you have a dedicated disk that you're going to use as a backup destination with a local XFS file system)

> **Note:**
>
> - Any remote FS like **NFS, SMB, etc.** - needs to be mounted by the user, and the vprotect user/group must own the directories within the backup destination. Data Protector for Cloud Workloads expects an already mounted file system and mount point in the backup destination.
> - You should add this file system to your /etc/fstab file on the node so that it gets mounted automatically if the OS is rebooted.
> - Consider using the same file system for the staging and backup destination (this boosts storage tasks, as no data needs to be copied again) - in such a scenario, the only difference would be that the presented /backupdestination mount point becomes a subdirectory of the staging space (usually /vprotect_data/backups).

Preparation

1. Log in to Data Protector for Cloud Workloads Node and create the mount directory as in the example `/backupdestination`

   ```
   mkdir /backupdestination
   ```

2. List all existing disks and find your drive:

   ```
   [root@vProtect01 ~]# fdisk -l | grep dev
   Disk /dev/sda: 32.2 GB, 32212254720 bytes, 62914560 sectors
   /dev/sda1   *        2048     1026047      512000   83  Linux
   /dev/sda2         1026048    62914559    30944256   8e  Linux LVM
   Disk /dev/sdc: 500 GB, 17179869184 bytes, 33554432 sectors
   Disk /dev/sdb: 21.5 GB, 21474836480 bytes, 41943040 sectors
   Disk /dev/mapper/centos-root: 28.5 GB, 28462546944 bytes, 55590912 sec
   Disk /dev/mapper/centos-swap: 3221 MB, 3221225472 bytes, 6291456 secto
   ```

3. Prepare a filesystem on it:

   ```
   mkfs.xfs -K /dev/sdc
   ```

4. Add permission for the Data Protector for Cloud Workloads user to access the directory `/backupdestination`

   - we assume here that you use a separate file system than your staging space
   - as an alternative, you also can point Data Protector for Cloud Workloads to use a subdirectory on the same file system as your staging space, for example `/vprotect_data/backups` (which you probably don't have to initialize at this point, as you may have already prepared it in the Staging space configuration, and you can just jump to the Web UI part in the next steps).

   ```
   chown vprotect:vprotect -R /backupdestination
   ```

5. Add this line to the `/etc/fstab` file to automatically mount new the filesystem after reboot:

   ```
   /dev/sdc    /backupdestination    xfs    defaults 0 0
   ```

   or if you want to store backups on NFS share then it will look like this (where 10.50.1.28 is your host):

   ```
   10.50.1.28:/example_nfs_share /backupdestination nfs defaults  0 0
   ```

6. Check if the fstab entry is OK and mount the filesystem:

   ```
   mount /backupdestination
   ```

7. Log in to the Data Protector for Cloud Workloads web UI.
8. Go to **Backup Destinations.**
9. Click on **Create Backup Destination**, choose a **File system.**
10. Type the name for the new backup destination, set the retention, and select at least one node configuration.
11. Usually, you have to decide if your backup destination is a separate entity from the staging space.

    - If the **staging space is different than your backup storage destination:**
      - In **Storage paths** type `/backupdestination` - this path will be used to mount the prepared file system (XFS) on top of the VDO volume.

- If the **staging space needs to be the same as your backup storage destination:**
  - In **Storage paths** type `/vprotect_data/backups`, where you point to a subdirectory (for example `backups` on your staging space path `/vprotect_data).`

12. Save the configuration.

Virtual Data Optimizer (VDO)

In this section, you can find information on how to enable deduplication using basically any block storage available. We assume that you have prepared your storage provider and have exposed the block device to the system where Data Protector for Cloud Workloads Node is installed.

Preparation

> Disable Secure Boot option for the VM to allow VDO work properly. Run below command to check status of Secure Boot option:
>
> ```
> mokutil --sb-state
> ```

1. Log in to Data Protector for Cloud Workloads Node and create a mount directory as in the example `/backupdestination`

   ```
   mkdir /backupdestination
   ```

2. List all existing disks, and find your drive. Let's assume `/dev/sdc` is your empty block device that you want to use:

   ```
   [root@vProtect01 ~]# fdisk -l | grep dev
   Disk /dev/sda: 32.2 GB, 32212254720 bytes, 62914560 sectors
   /dev/sda1   *        2048     1026047      512000   83  Linux
   /dev/sda2         1026048    62914559    30944256   8e  Linux LVM
   Disk /dev/sdc: 500 GB, 17179869184 bytes, 33554432 sectors
   Disk /dev/sdb: 21.5 GB, 21474836480 bytes, 41943040 sectors
   Disk /dev/mapper/centos-root: 28.5 GB, 28462546944 bytes, 55590912 sec
   Disk /dev/mapper/centos-swap: 3221 MB, 3221225472 bytes, 6291456 secto
   ```

3. Log in to the Data Protector for Cloud Workloads web UI.
4. Go to **Backup Destinations.**
5. Click on **Create Backup Destination**, choose a **File system.**
6. Type a name for the new backup destination, set the retention, and select at least one node configuration.

7. Based on whether the staging space is same as backup destination or not, do one of the following:
    - If the **staging space is different than your backup destination** storage:
        - In **Storage paths** type `/backupdestination` - this path will be used to mount the prepared file system (XFS) on top of the VDO volume.
        - Check **Enable deduplication.**
        - Provide your block device (for example `/dev/sdc`) as your Deduplication device.
    - If the **staging space needs to be the same as your backup destination** storage:
        - In **Storage paths** type `/vprotect_data/backups` - this path assumes that `/vprotect_data` is your staging space path and `backups` is a subdirectory of the staging space.
        - Check **Enable deduplication.**
        - Provide your block device (for example `/dev/sdc`) as your **Deduplication device.**
        - Enable **Mount deduplicated file system to a different directory than backup destination path** and provide the mount point - your staging space path, for example `/vprotect_data` - this will force Data Protector for Cloud Workloads to mount XFS on top of VDO in the staging space directory rather than in the backup subdirectory.

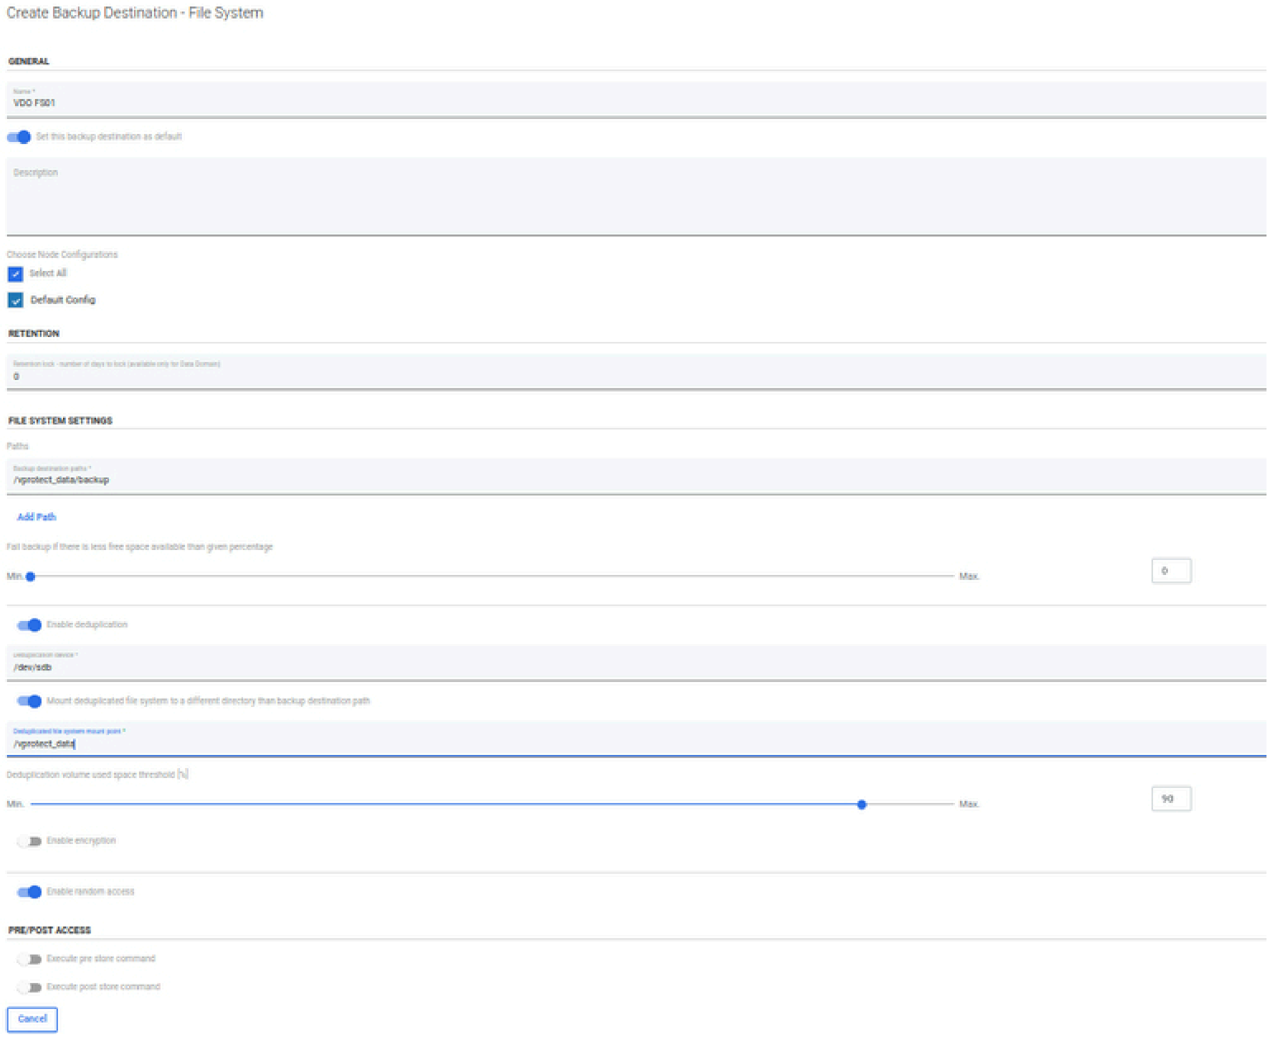

Create Backup Destination - File System

GENERAL

Name *
VDO FS01

● Set this backup destination as default

Description

Choose Node Configurations

☑ Select All

☑ Default Config

RETENTION

Retention lock - number of days to lock (available only for Data Domain)
0

FILE SYSTEM SETTINGS

Paths

Backup destination paths *
/vprotect_data/backup

Fail backup if there is less free space available than given percentage

Min. Max. 0

● Enable deduplication

Deduplication device *
/dev/sdb

● Mount deduplicated file system to a different directory than backup destination path

Deduplicated file system mount point *
/vprotect_data

Deduplication volume used space threshold [%]

Min. Max. 90

○ Enable encryption

● Enable random access

PRE/POST ACCESS

○ Execute pre store command

○ Execute post store command

> **Note**:
>
> Only one file system backup destination with deduplication using VDO pointing to a specific directory can be used. If you want to add another backup destination using the same VDO device, but just a different subdirectory, create it without deduplication enabled.

Importing existing VDO volumes to LVM

The python-based VDO management software has been deprecated and removed from RHEL 9/CentOS 9 Stream. It has been replaced by the LVM-VDO integration. If you are using VDO on RHEL 8/CentOS 8 Stream and plan to upgrade to version 9, you need to convert VDO volume.

In this example we have VDO volume called VDOexample created and managed by Data Protector for Cloud Workloads.

```
[root@dp4cw-node ~]# lsblk
NAME         MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda            8:0    0   40G  0 disk
|-sda1         8:1    0  600M  0 part /boot/efi
|-sda2         8:2    0    1G  0 part /boot
`-sda3         8:3    0 38.4G  0 part
  |-cs-root  253:0    0 34.4G  0 lvm  /
  `-cs-swap  253:1    0    4G  0 lvm  [SWAP]
sdb            8:16   0  100G  0 disk
`-VDOexample 253:2    0  300G  0 vdo  /backups
sr0           11:0    1 1024M  0 rom
```

1. On Data Protector for Cloud Workloads Node, stop vprotect-node service.

```
[root@dp4cw-node ~]# systemctl stop vprotect-node
```

2. Unmount VDO volume from backup destination path.

```
[root@dp4cw-node ~]# umount /backups
```

3. Convert VDO volume. Change `/dev/sdb` to the device on which you have created VDO.

```
[root@dp4cw-node ~]# lvm_import_vdo /dev/sdb
Convert VDO device "/dev/sdb" to VDO LV "vdovg/vdolvol"? [y|N]: Yes
Stopping VDO VDOexample
Converting VDO VDOexample
    Opening /dev/sdb exclusively
    Loading the VDO superblock and volume geometry
    Checking the VDO state
    Converting the UDS index
    Converting the VDO
    Conversion completed for '/dev/sdb': VDO is now offset by 2097152
Physical volume "/dev/sdb" successfully created.
Volume group "vdovg" successfully created
WARNING: Logical volume vdovg/vdolvol_vpool not zeroed.
Logical volume "vdolvol_vpool" created.
WARNING: Converting logical volume vdovg/vdolvol_vpool to VDO pool vol
WARNING: Using invalid VDO pool data MAY DESTROY YOUR DATA!
Logical volume "vdolvol" created.
Converted vdovg/vdolvol_vpool to VDO pool volume and created virtual v
```

4. Rename volume group and logical volume names. They must be the same as the original VDO volume name.

```
[root@dp4cw-node ~]# vgrename vdovg VDOexample
Volume group "vdovg" successfully renamed to "VDOexample"
[root@dp4cw-node ~]# lvrename /dev/VDOexample/vdolvol /dev/VDOexample/
Renamed "vdolvol" to "VDOexample" in volume group "VDOexample"
```

5. On Data Protector for Cloud Workloads Server machine, create a vprotect database backup and copy it to safe place. Wait for all tasks to finish before stopping the vprotect-server service.

```
[root@dp4cw-server ~]# stop systemctl vprotect-server
[root@dp4cw-server ~]# /opt/vprotect/scripts/backup_db.sh
[root@dp4cw-server ~]# cp /tmp/vprotect_db.sql.gz /root
```

6. Login to mysql and execute below SQL query.

```
[root@dp4cw-node ~]# mysql -uroot -p vprotect

update filesystembackupdestination
inner join backupdestination on filesystembackupdestination.guid = bac
set filesystembackupdestination.dedupvolume = CONCAT('/dev/', REGEXP_R
where filesystembackupdestination.dedupvolume is not null;

MariaDB [vprotect]> quit
```

7. Start vprotect-server service.

```
[root@dp4cw-server ~]# systemctl start vprotect-server
```

8. Proceed with the system upgrade of the Data Protector for Cloud Workloads Node machine. After the reboot, you should have new LVM-VDO mounted on your backupdestination directory.

```
[root@dp4cw-node ~]# lsblk
NAME                               MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                                  8:0    0   40G  0 disk
├─sda1                               8:1    0  600M  0 part /boot/efi
├─sda2                               8:2    0    1G  0 part /boot
└─sda3                               8:3    0 38.4G  0 part
├─cs-root                        253:0    0 34.4G  0 lvm  /
└─cs-swap                        253:1    0    4G  0 lvm  [SWAP]
sdb                                  8:16   0  100G  0 disk
└─VDOexample-vdolvol_vpool_vdata   253:2    0  100G  0 lvm
└─VDOexample-vdolvol_vpool-vpool 253:3    0  300G  0 lvm
    └─VDOexample-VDOexample        253:4    0  300G  0 lvm  /backups
```

Synthetic File System

A synthetic file system allows us to store and use incremental backups as if they were full backup files, but they take up a fraction of full file size.

To start using Synthetic File System read Prerequisites for **Synthetic filesystem XFS/NFS 4.2**

XFS

Prerequisites

> **Note:**
>
> - The only prerequisite to using synthetic XFS as a backup destination is that the selected storage path is on the XFS
> - For a basic setup of file systems on the Node check File system

Creating a Synthetic Filesystem Backup Destination

1. Select File System from Backup Destinations,

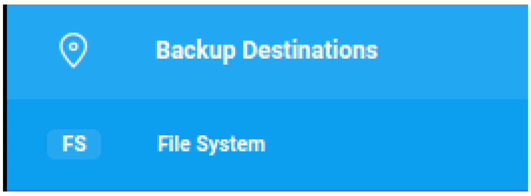

2. Select Create Backup Destination → File System (Synthetic),

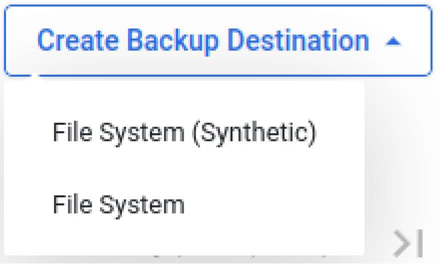

3. The rest of the configuration is similar to a regular File system.

> **When setting the path, make sure it's actually on the XFS!**

Deduplication Appliances

This section presents the key steps necessary for configuring integration with deduplication appliances as your backup destination. You can use NFS or SMB to attach Dell EMC Data Domain.

Dell EMC Data Domain

Create a new Backup Destination (Dell EMC Data Domain)

- Go into the backup destination menu and click on Create a backup destination.
- Provide a name and description for the new backup destination.
- Specify the retention days for full and incremental backups.
- Specify the retention versions for full and incremental backups.
- Choose and assign the node configuration to which you want to attach the new backup destination.
- Add to one or more storage paths.
  - `example - /vprotect_data/backupdestination`
- Save the configuration.

DD Boost FS Plugin

- To boost the backup process we recommend using a **single** Storage Unit and mtree, and subfolders on the BoostFS for multiple backup destinations (with possibly different retention settings).
  - No additional data copy is needed in the store phase if staging is using the same file system as the backup destination.
  - The Setup assumes a single Storage Unit and a single mtree for all backup destinations.
  - The staging space should always be a top directory, and all backup destinations should be defined as separate subfolders of this file system.
  - Data Protector for Cloud Workloads handles retention, and each backup destination may have different retention configured.
  - A single Storage Unit will also affect replication as it has to cover all backup destinations, and may replicate temporary data from the staging space or mounted backups.

- **Sharing** the same BoostFS across multiple nodes allows the administrator to create backups on one node (one host/environment) and restore using a different node (to a different host/environment).
  - UID/GID ownership and permissions must allow Data Protector for Cloud Workloads to read/write contents of the BoostFS share.
  - To meet these requirements, the user and group named vprotect that was created during the installation process must have the same UID and GID on each Data Protector for Cloud Workloads Node machine. You can create this before installing Data Protector for Cloud Workloads packages or change it after installation.
- Data Domain User requirements:
  - user must have **backup-operator** management role
  - user must be assigned to DD Boost Storage Unit

Prepare your PowerProtect DD as a backup destination:

- Login to PowerProtect DD and create an NFS Storage Unit called `DP-for-Cloud-Workloads`.

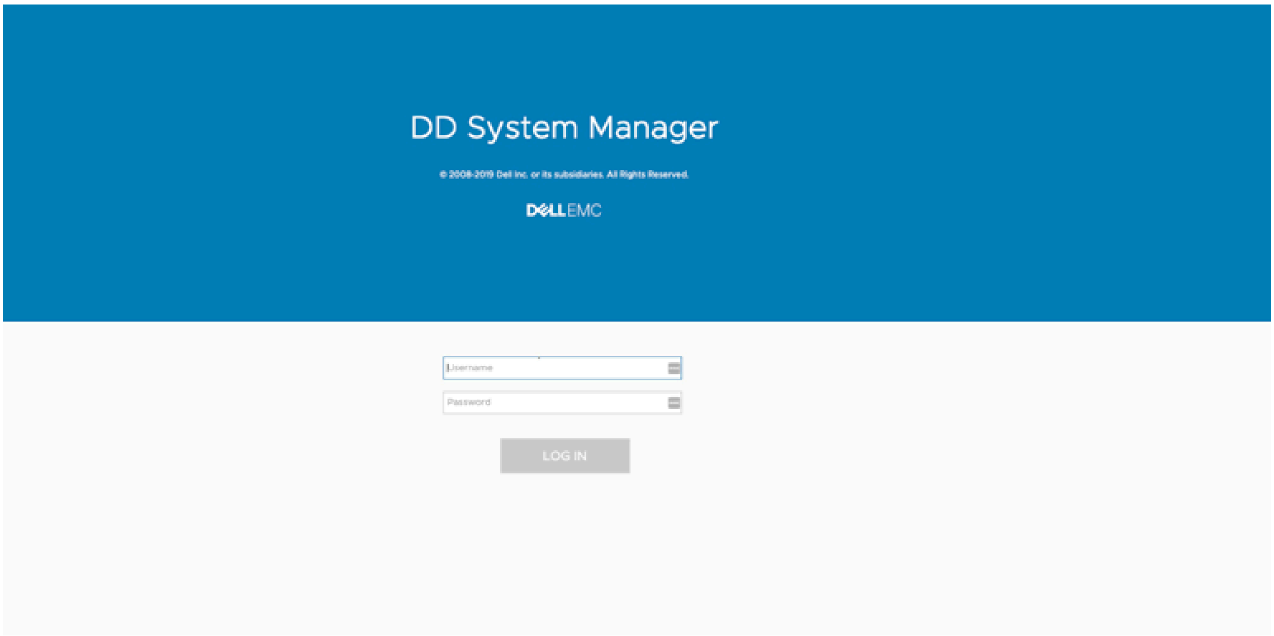
DD System Manager
© 2008-2019 Dell Inc. or its subsidiaries. All Rights Reserved.
DELLEMC
Username
Password
LOG IN

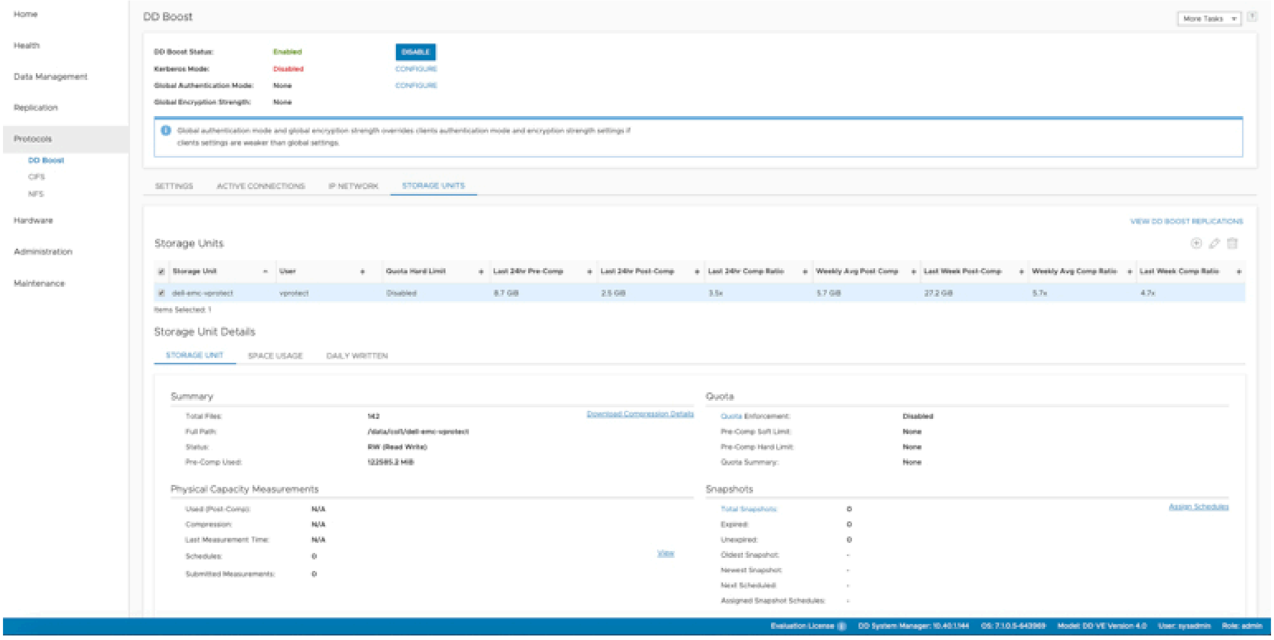

DD Boost

DD Boost Status: Enabled
Kerberos Mode: Disabled
Global Authentication Mode: None
Global Encryption Strength: None

Global authentication mode and global encryption strength overrides clients authentication mode and encryption strength settings if clients settings are weaker than global settings.

Storage Units

☑	Storage Unit	User	Quota Hard Limit	Last 24hr Pre-Comp	Last 24hr Post-Comp	Last 24hr Comp Ratio	Weekly Avg Post Comp	Last Week Post-Comp	Weekly Avg Comp Ratio	Last Week Comp Ratio
☑	dell-emc-vprotect	vprotect	Disabled	8.7 GB	2.5 GB	3.5x	5.7 GB	27.2 GB	5.7x	4.7x

Items Selected: 1

Storage Unit Details

Summary

Total Files: 142
Full Path: /data/col1/dell-emc-vprotect
Status: RW (Read Write)
Pre-Comp Used: 122585.2 MiB

Physical Capacity Measurements

Used (Post-Comp): N/A
Compression: N/A
Last Measurement Time: N/A
Schedules: 0
Submitted Measurements: 0

Quota

Quota Enforcement: Disabled
Pre-Comp Soft Limit: None
Pre-Comp Hard Limit: None
Quota Summary: None

Snapshots

Total Snapshots: 0
Expired: 0
Unexpired: 0
Oldest Snapshot: -
Newest Snapshot: -
Next Scheduled: -
Assigned Snapshot Schedules: -

Evaluation License DD System Manager: 10.40.1.144 OS: 7.1.0.5-643969 Model: DD VE Version 4.0 User: sysadmin Role: admin

- Download BoostFS RPM from the Dell EMC site.
- Install BoostFS:

```
rpm -ivh DDBoostFS-7.0.0.0-633922.rhel.x86_64.rpm
```

- Save the password for BoostFS.

```
# Syntax
/opt/emc/boostfs/bin/boostfs lockbox set -d [DataDomain_IP_OR_DNS_NAME
# Example
/opt/emc/boostfs/bin/boostfs lockbox set -d 10.1.10.100 -u vprotect -s
```

- Add the /etc/fstab entry:

```
# Syntax
[DataDomain_IP_OR_DNS_NAME]:/[Storage_Area_Name] /[Mount_Point] boostf
# Example
10.1.10.100:/vprotectbackup /vprotect_data boostfs defaults,_netdev,bf
```

- Mount the fstab entry:

```
mount -a
```

- For a manual, one-time mount you can run this command:

```
# Syntax
/opt/emc/boostfs/bin/boostfs mount -o allow-others=true -d [DataDomain
# Example
/opt/emc/boostfs/bin/boostfs mount -o allow-others=true -d 10.1.10.100
```

- Confirm with `df -h` that your `/vprotect_data` is mounted
  Note: Remember to specify the backup destination path as a subdirectory of /vprotect_data if you would like to use the same storage unit as a staging space and backup destination - for example: /vprotect_data/my-backups.

  ```
  mkdir /vprotect_data/my-backups
  ```

- Set ownership to the vprotect user on the directory /vprotect_data.

  ```
  chown vprotect:vprotect -R /vprotect_data
  ```

- Set ownership to the vprotect user and data domain group on the directory /vprotect_data/my-backups.

  ```
  chown vprotect:gid /vprotect_data/my-backups
  ```

  where 'gid' is the GID of data domain user specified in Synthetic DD Boost backup destination configuration.

- Set read and write privileges for both user and group to the directory /vprotect_data/my-backups.

  ```
  chmod 0775 /vprotect_data/my-backups
  ```

Object Storage

Object Storage

A backup destination is a storage location where Data Protector for Cloud Workloads keeps VMs, Containers, Cloud, and application backup copies. Data Protector for Cloud Workloads supports different types of object storage.

- Alibaba Cloud OSS
- AWS S3 or S3-compatible
- Ceph Rados Gateway
- Cloudian S3
- Wasabi
- Google Cloud Storage
- IBM Cloud Object Storage
- Microsoft Azure Blob Storage
- Nutanix Objects
- OpenStack SWIFT
- Oracle Cloud Infrastructure Object Storage
- Scality RING

Alibaba Cloud OSS

Overview

Alibaba Cloud is an S3-compatible backup provider. Configuration as a backup destination is similar to AWS S3.

Example

After logging in, go to the Object Storage Service and create a new bucket.

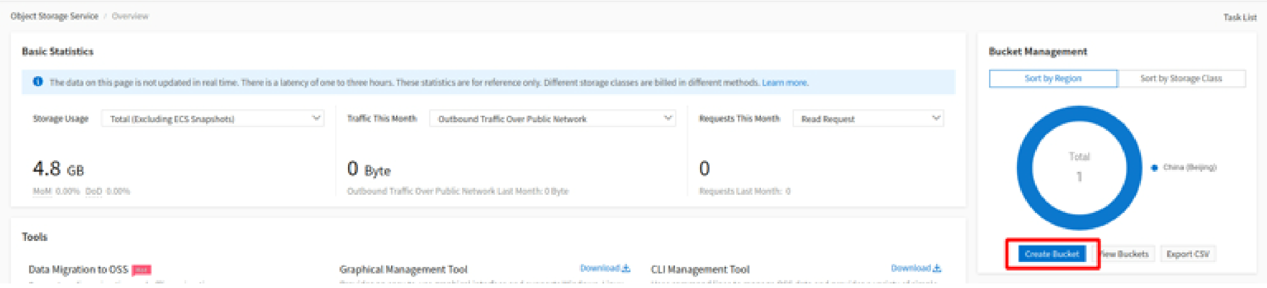

The data on this page is not updated in real time. There is a latency of one to three hours. These statistics are for reference only. Different storage classes are billed in different methods. Learn more.

Storage Usage: Total (Excluding ECS Snapshots)
4.8 GB
MoM 0.00% DoD 0.00%

Traffic This Month: Outbound Traffic Over Public Network
0 Byte
Outbound Traffic Over Public Network Last Month: 0 Byte

Requests This Month: Read Request
0
Requests Last Month: 0

Provide necessary details for your bucket and enable versioning.

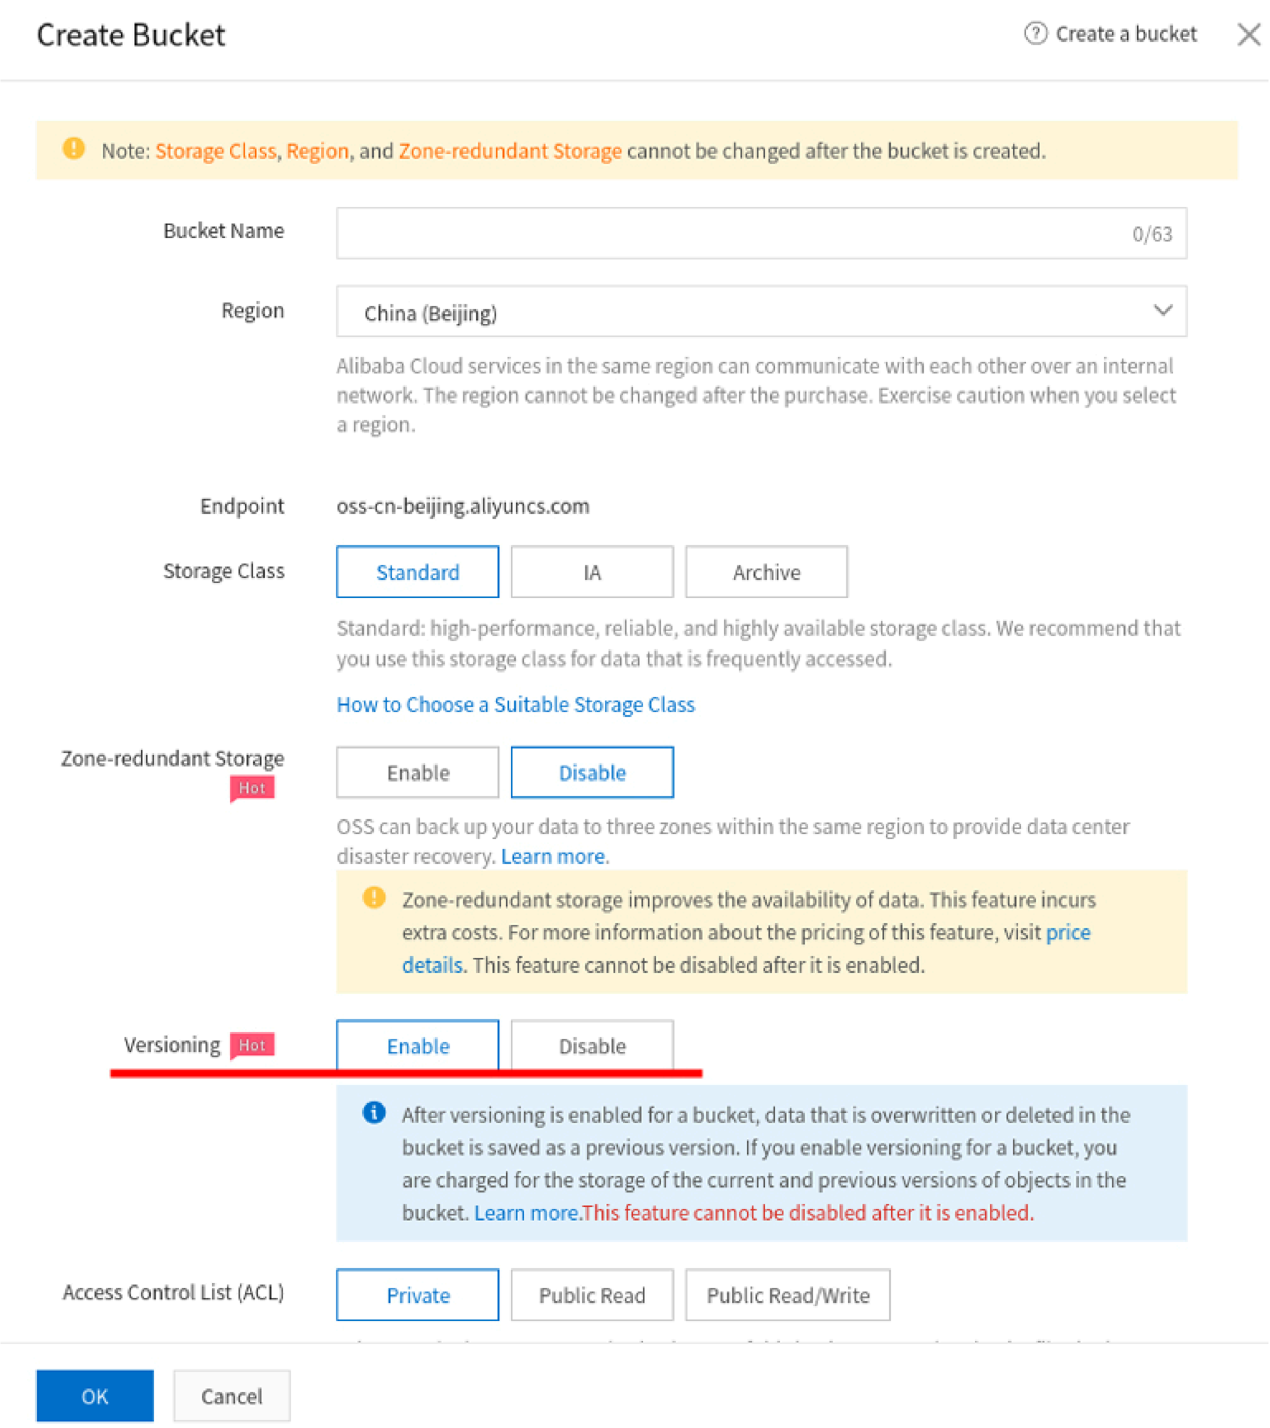

Create Bucket

Create a bucket

Note: Storage Class, Region, and Zone-redundant Storage cannot be changed after the bucket is created.

Bucket Name: 0/63

Region: China (Beijing)

Alibaba Cloud services in the same region can communicate with each other over an internal network. The region cannot be changed after the purchase. Exercise caution when you select a region.

Endpoint: oss-cn-beijing.aliyuncs.com

Storage Class: [✓ Standard] IA Archive

Standard: high-performance, reliable, and highly available storage class. We recommend that you use this storage class for data that is frequently accessed.

How to Choose a Suitable Storage Class

Zone-redundant Storage Hot: Enable [✓ Disable]

OSS can back up your data to three zones within the same region to provide data center disaster recovery. Learn more.

Zone-redundant storage improves the availability of data. This feature incurs extra costs. For more information about the pricing of this feature, visit price details. This feature cannot be disabled after it is enabled.

Versioning Hot: [✓ Enable] Disable

After versioning is enabled for a bucket, data that is overwritten or deleted in the bucket is saved as a previous version. If you enable versioning for a bucket, you are charged for the storage of the current and previous versions of objects in the bucket. Learn more.This feature cannot be disabled after it is enabled.

Access Control List (ACL): [✓ Private] Public Read Public Read/Write

Next, go to Manage AccessKey Management and create new AccessKey

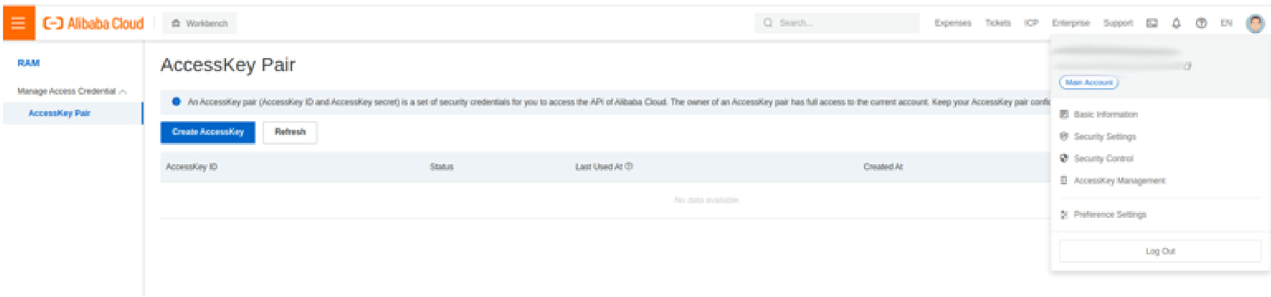

AccessKey Pair

An AccessKey pair (AccessKey ID and AccessKey secret) is a set of security credentials for you to access the API of Alibaba Cloud. The owner of an AccessKey pair has full access to the current account. Keep your AccessKey pair confi

AccessKey ID	Status	Last Used At	Created At
No data available.			

Now go to the Backup destination tab on the Data Protector for Cloud Workloads dashboard and change the sub-tab to object storage. Provide the bucket name and key credentials, and then configure the remaining options according to your requirements:

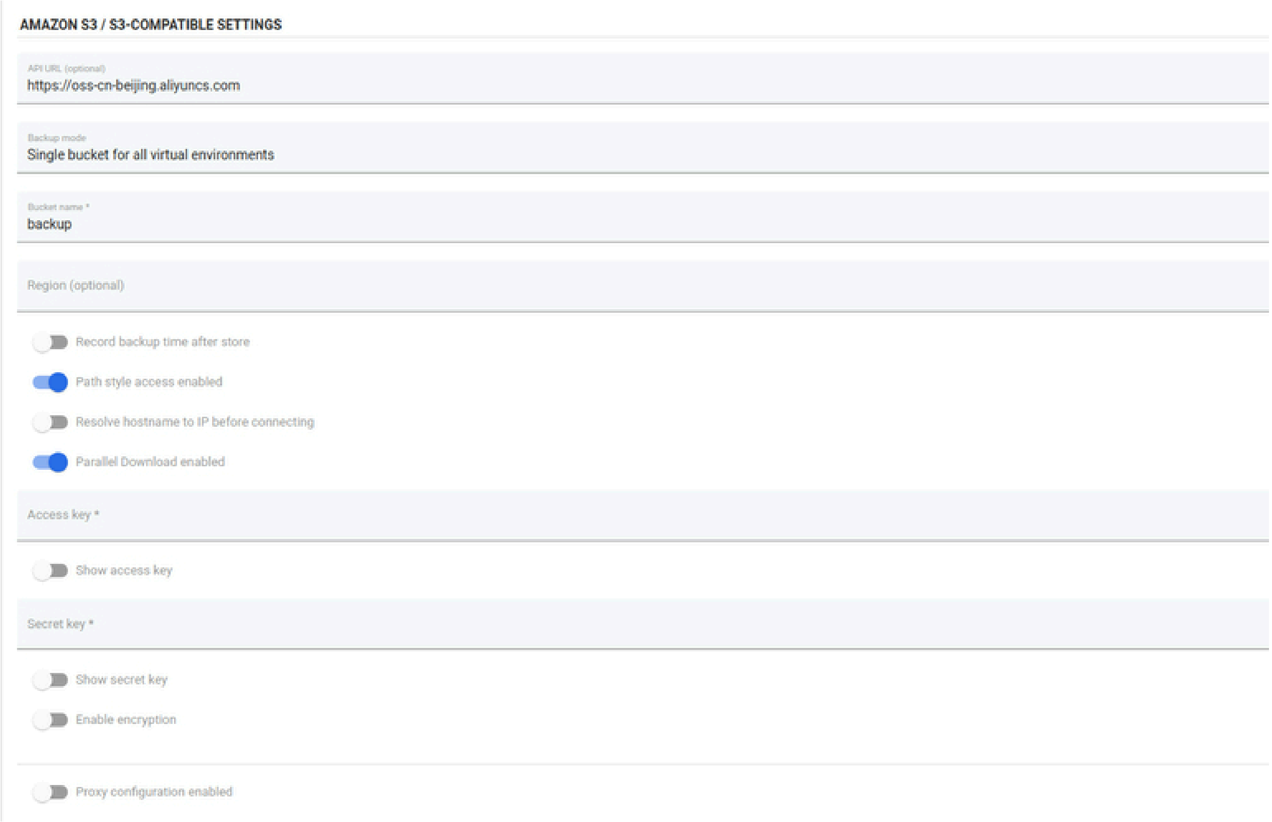

AMAZON S3 / S3-COMPATIBLE SETTINGS

API URL (optional): https://oss-cn-beijing.aliyuncs.com

Backup mode: Single bucket for all virtual environments

Bucket name *: backup

Region (optional)

☐ Record backup time after store

☑ Path style access enabled

☐ Resolve hostname to IP before connecting

☑ Parallel Download enabled

Access key *

☐ Show access key

Secret key *

☐ Show secret key

☐ Enable encryption

☐ Proxy configuration enabled

AWS S3 or S3-compatible

Overview

Data Protector for Cloud Workloads can store backups in AWS S3 or S3-compatible backup providers. In most cases, you just need to prepare a bucket (with versioning enabled if possible) and generate an access/secret key for Data Protector for Cloud Workloads. Data Protector for Cloud Workloads can be installed in AWS (if EC2 backup is used), but in most cases, S3 is used just as a cloud backup provider for on-prem environments.

Typical use cases are:

- When AWS is used - choose a single bucket with **versioning enabled** - all backup objects will have names in `/container_name/path/to/backup` format, where `container_name` typically is the VM name with an identifier.
- When a 3rd party is used - you need to verify:
    - Which strategy is supported by the vendor - for example Scality requires a single bucket without versioning.
    - When timestamp recording of the object should occur - for example Scality does it after data is stored (unlike AWS).

Data Protector for Cloud Workloads is also able to **encrypt** backups before sending backups (client-side encryption: SSE-C). Once enabled, new data is stored as encrypted with keys generated and kept by Data Protector for Cloud Workloads. For performance improvements, we also recommend using AWS Direct Connect to access S3. Otherwise, backups would be sent over the Internet, which could result in poor performance.

> **Note:** S3 has a **limit of 5TB** per object. This means that depending on the virtualization platform and backup format used by export/import mode you may have a limit of 5TB per VM (if it is Proxmox VMA or Citrix XVA image-based backup) or per VM disk (in most cases). Bigger files are currently not supported.

Permissions

Depending on the selected mode, you may have different permission sets. For a single bucket, you need to use the access keys of a user that has the ability to control objects within the bucket over the specific bucket - here is an example of IAM policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1568968204280",
      "Action": [
        "s3:DeleteObject",
        "s3:DeleteObjectTagging",
        "s3:DeleteObjectVersion",
        "s3:DeleteObjectVersionTagging",
        "s3:GetBucketTagging",
        "s3:GetBucketVersioning",
        "s3:GetObject",
        "s3:GetObjectRetention",
        "s3:GetObjectTagging",
        "s3:GetObjectVersion",
        "s3:GetObjectVersionTagging",
        "s3:ListBucket",
        "s3:ListBucketVersions",
        "s3:PutObject",
        "s3:PutObjectTagging",
        "s3:PutObjectVersionTagging",
        "s3:RestoreObject"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:s3:::BACKUP_DESTINATION_BUCKET/*"
    }
  ]
}
```

You can also use a predefined role and create a user from the AWS console: https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html#Using_CreateAccessKey

> **Note:** It is recommended to periodically rotate your access/secret keys. More information can be found here: https://aws.amazon.com/blogs/security/how-to-rotate-access-keys-for-iam-users/. After changing the key in AWS, remember to update it in Data Protector for Cloud Workloads as well.

Bucket replication

Even though S3 is a highly available service, you may want to be prepared in case of a region failure. We recommend following this guide https://docs.aws.amazon.com/AmazonS3/latest/dev/replication.html to set up bucket replication so that your data is replicated to another region in a worst-case scenario. Remember to point Data Protector for Cloud Workloads to the replicated bucket in case of a disaster.

Glacier/Deep Archive support

Data Protector for Cloud Workloads is able to move older backups to a Glacier/Deep Archive storage tier. In the S3 backup provider settings, you need to enable the `Move old versions to other storage class` toggle and provide extended retention settings.

Keep in mind that Data Protector for Cloud Workloads will try to restore it to S3 with an expiration set to 2 days. You'll notice that although the task is running, no progress is taking place as it is waiting for the object to be restored from Glacier to S3. This **may take several hours** as Glacier doesn't provide instant access to archival data. Once this part is completed, Data Protector for Cloud Workloads will proceed with regular restore from a temporary S3 object.

Costs

When storing backups in S3, additional charges will occur for stored backups. Retention setting in Data Protector for Cloud Workloads can limit the storage costs of stored backups.

Please visit https://aws.amazon.com/s3/pricing/ to check current AWS S3 pricing.

Example

Now we will show you how to quickly create S3 storage and integrate it with Data Protector for Cloud Workloads as a backup destination.
After logging in, expand the services tab and choose S3 under the Storage section:

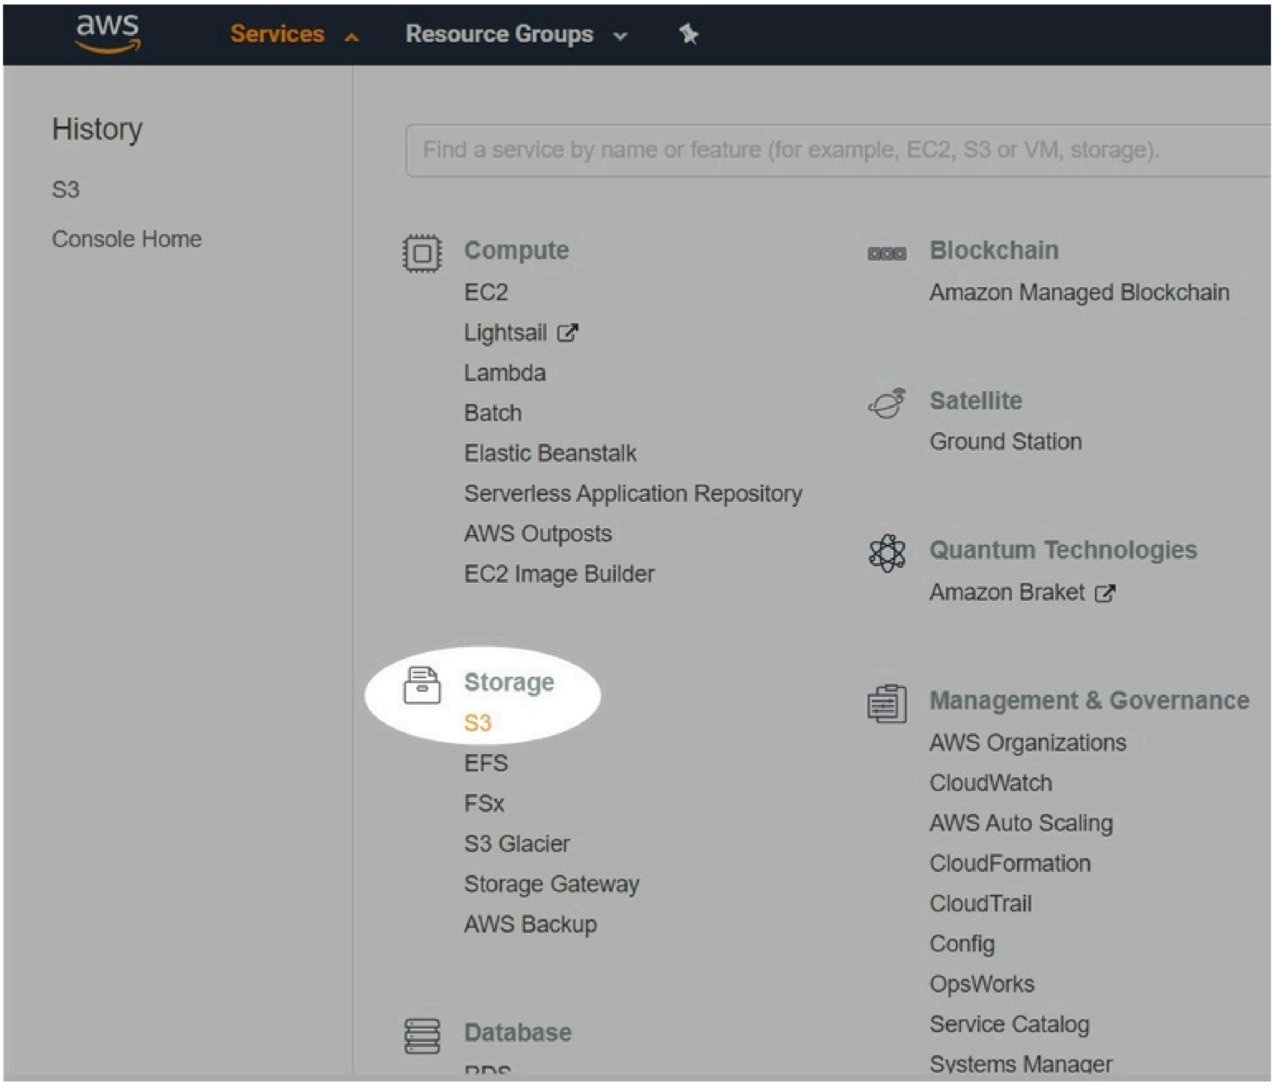

Now create a new bucket for your backups:

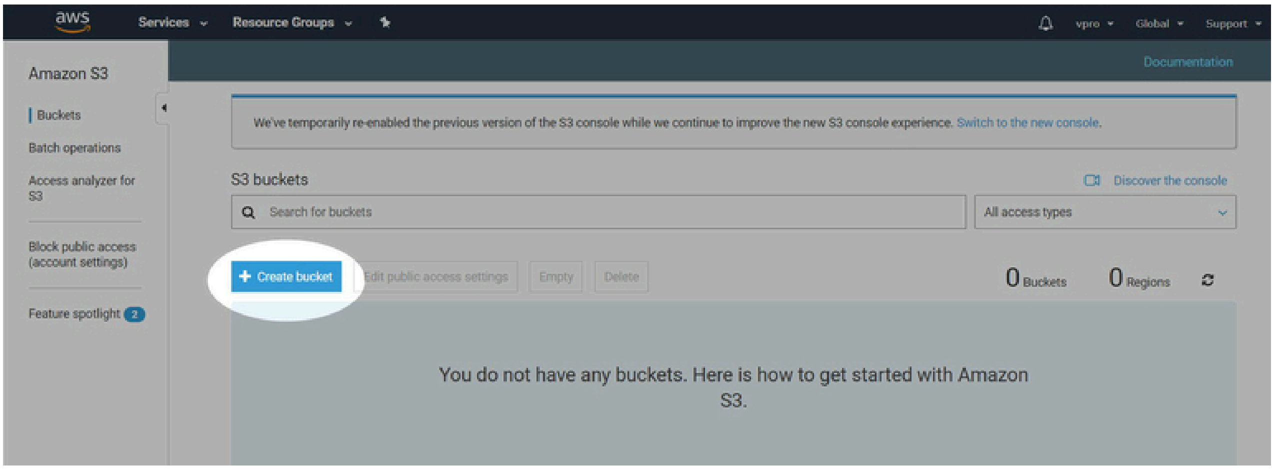

We've temporarily re-enabled the previous version of the S3 console while we continue to improve the new S3 console experience. Switch to the new console.

You do not have any buckets. Here is how to get started with Amazon S3.

In "Configure options" activate versioning:
(In all other tabs, you can leave the default settings)

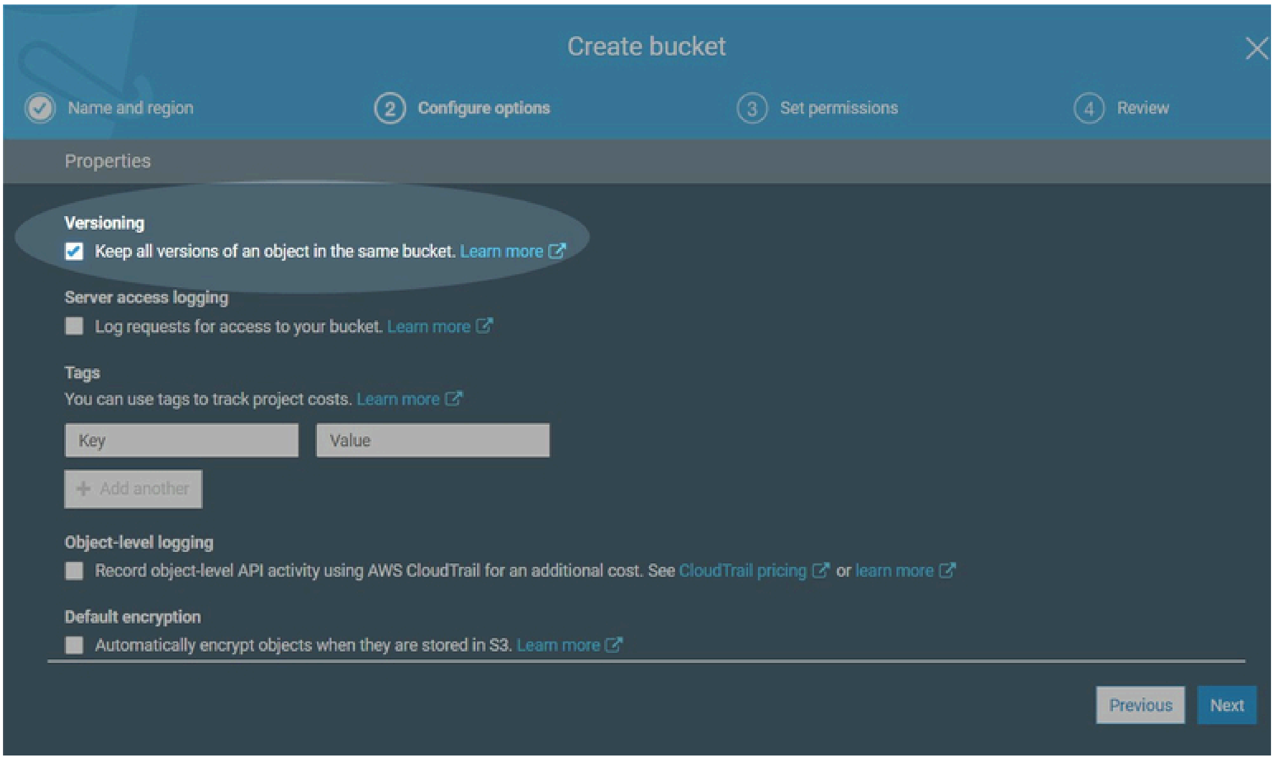

Versioning

☑ Keep all versions of an object in the same bucket. Learn more

Server access logging

□ Log requests for access to your bucket. Learn more

Tags

You can use tags to track project costs. Learn more

Object-level logging

□ Record object-level API activity using AWS CloudTrail for an additional cost. See CloudTrail pricing or learn more

Default encryption

□ Automatically encrypt objects when they are stored in S3. Learn more

After creating a bucket, we need to create a new user with appropriate permissions:

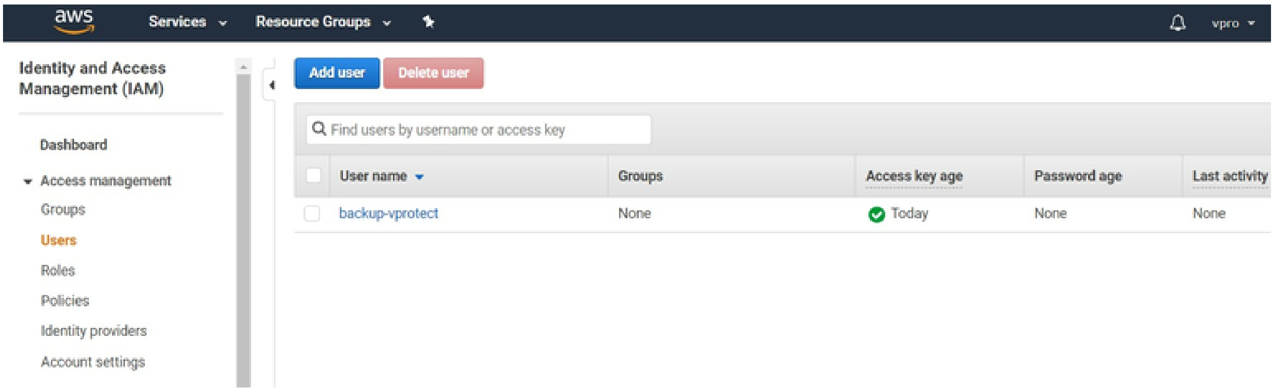

	User name	Groups	Access key age	Password age	Last activity
□	backup-vprotect	None	✓ Today	None	None

Remember to choose the "Programmatic access" account type:

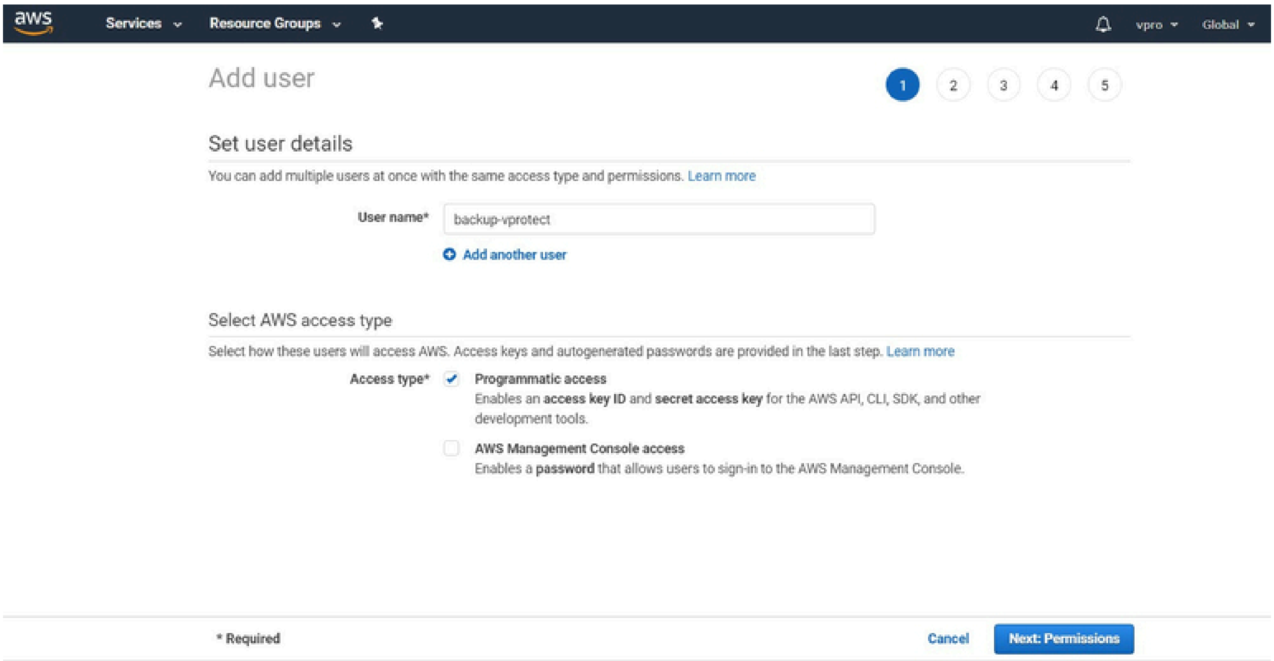

Add user

Set user details

You can add multiple users at once with the same access type and permissions. Learn more

User name* backup-vprotect

Select AWS access type

Select how these users will access AWS. Access keys and autogenerated passwords are provided in the last step. Learn more

Access type* ☑ Programmatic access
Enables an **access key ID** and **secret access key** for the AWS API, CLI, SDK, and other development tools.

□ AWS Management Console access
Enables a **password** that allows users to sign-in to the AWS Management Console.

* Required

From the predefined roles, you can choose "AmazonS3FullAccess" or you can create a new one as described in the Permissions section:

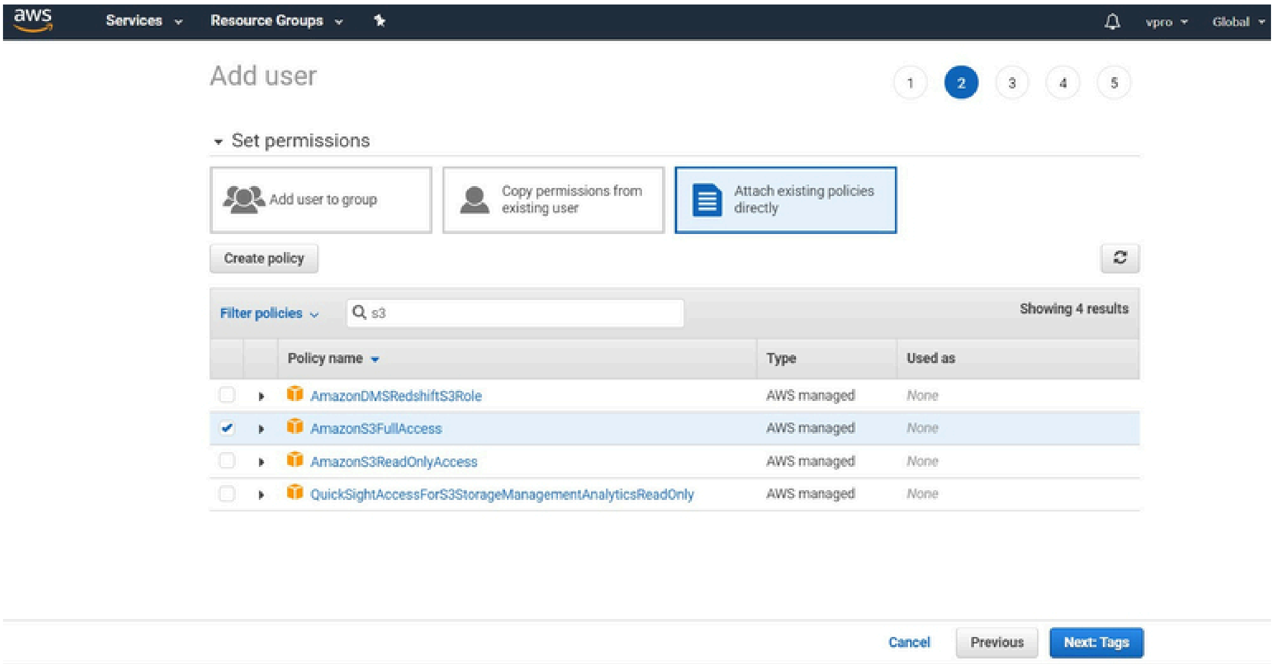

Add user

Set permissions

○ Add user to group ○ Copy permissions from existing user ● Attach existing policies directly

Q s3

		Policy name	Type	Used as
□	▸	AmazonDMSRedshiftS3Role	AWS managed	None
☑	▸	AmazonS3FullAccess	AWS managed	None
□	▸	AmazonS3ReadOnlyAccess	AWS managed	None
□	▸	QuickSightAccessForS3StorageManagementAnalyticsReadOnly	AWS managed	None

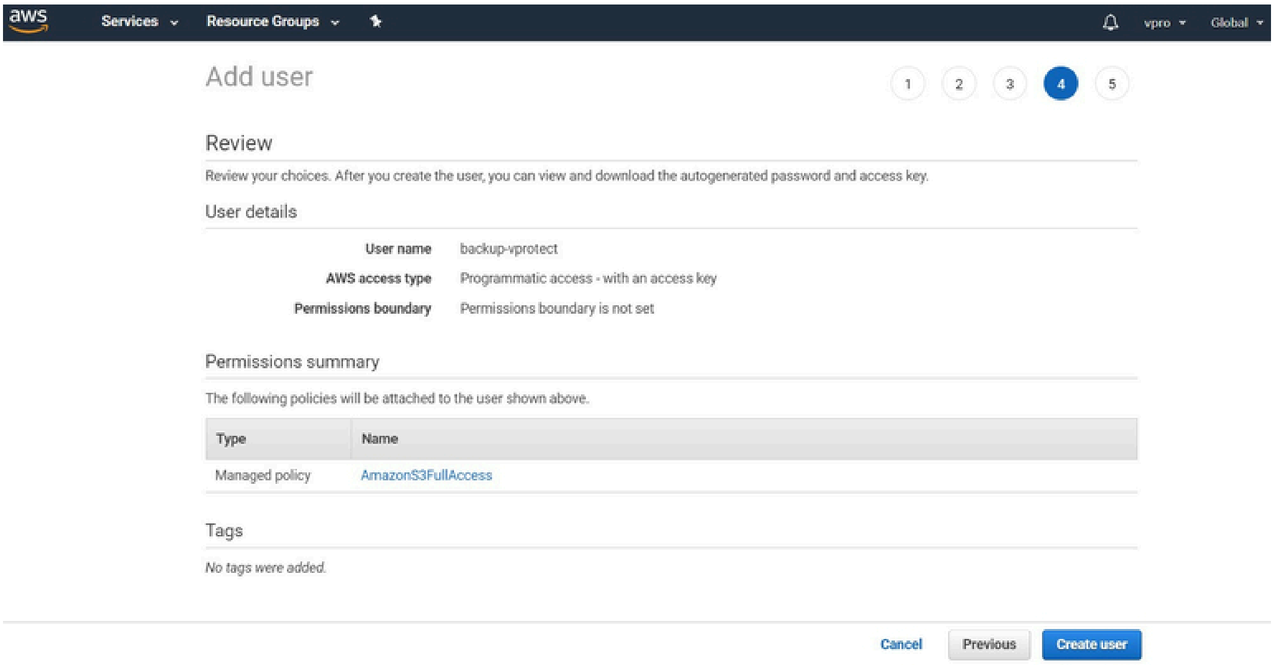

Add user

Review

Review your choices. After you create the user, you can view and download the autogenerated password and access key.

User details

User name	backup-vprotect
AWS access type	Programmatic access - with an access key
Permissions boundary	Permissions boundary is not set

Permissions summary

The following policies will be attached to the user shown above.

Type	Name
Managed policy	AmazonS3FullAccess

Tags

No tags were added.

> Remember to download the .csv or copy the key credentials manually:

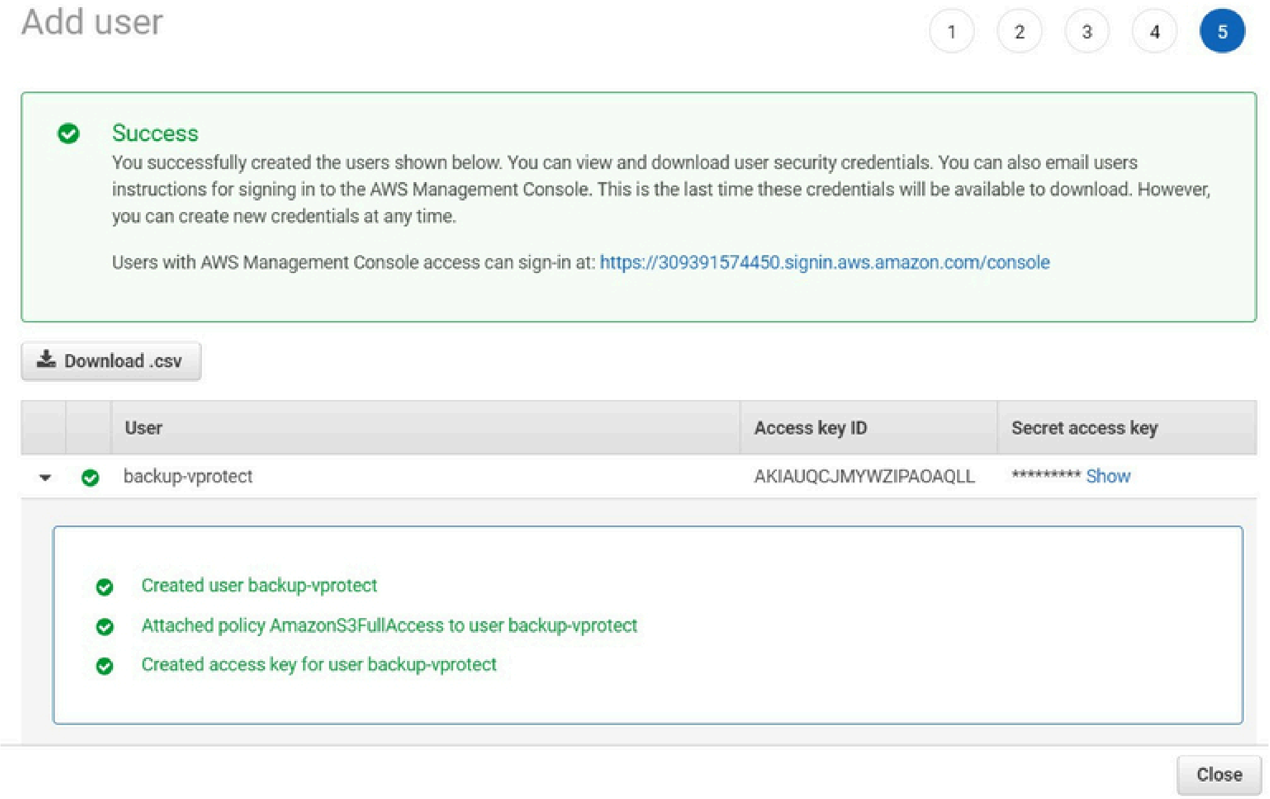

Add user

Success

You successfully created the users shown below. You can view and download user security credentials. You can also email users instructions for signing in to the AWS Management Console. This is the last time these credentials will be available to download. However, you can create new credentials at any time.

Users with AWS Management Console access can sign-in at: https://309391574450.signin.aws.amazon.com/console

		User	Access key ID	Secret access key
▾	✓	backup-vprotect	AKIAUQCJMYWZIPAOAQLL	********* Show

- ✓ Created user backup-vprotect
- ✓ Attached policy AmazonS3FullAccess to user backup-vprotect
- ✓ Created access key for user backup-vprotect

Now go to the Backup destination tab on the Data Protector for Cloud Workloads dashboard and change the sub-tab to object storage. Provide the bucket name and key credentials, and then configure the remaining options according to your requirements:

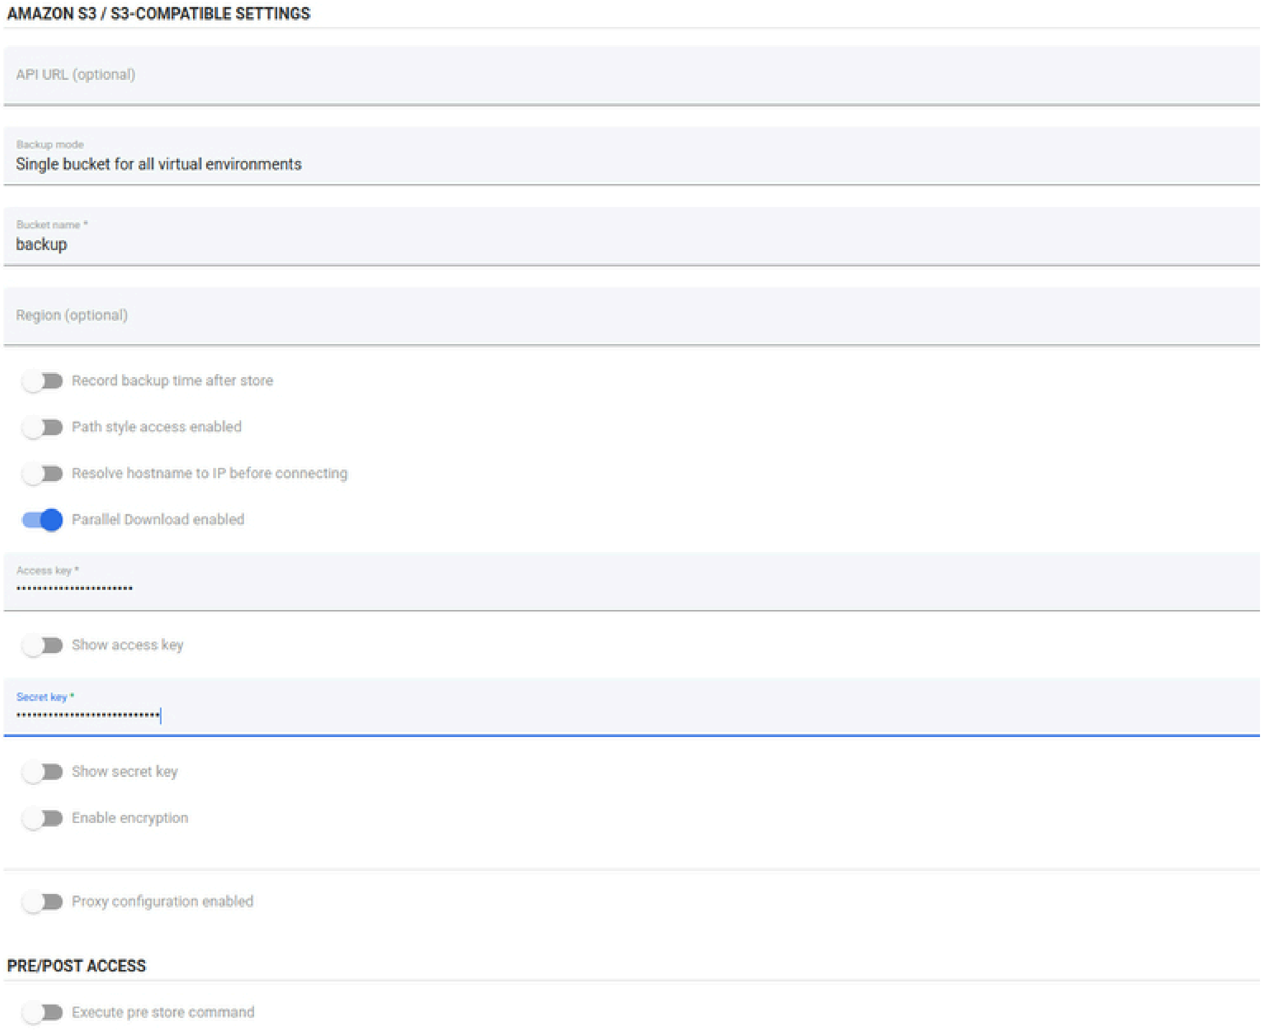

AMAZON S3 / S3-COMPATIBLE SETTINGS

API URL (optional)

Backup mode
Single bucket for all virtual environments

Bucket name *
backup

Region (optional)

○ Record backup time after store

○ Path style access enabled

○ Resolve hostname to IP before connecting

● Parallel Download enabled

Access key *
••••••••••••••••••••••••

○ Show access key

Secret key *
••••••••••••••••••••••••••••••

○ Show secret key

○ Enable encryption

○ Proxy configuration enabled

PRE/POST ACCESS

○ Execute pre store command

Ceph Rados Gateway

Overview

Ceph Object Gateway supports a RESTful API that is compatible with the basic data access model of the Amazon S3 API. Ceph Object Gateway is an object storage interface built on top of librados to provide applications with a RESTful gateway to Ceph Storage Clusters. Ceph Object Storage supports two interfaces:

- **S3-compatible**: Provides object storage functionality with an interface that is compatible with a large subset of the Amazon S3 RESTful API.
- **Swift-compatible**: Provides object storage functionality with an interface that is compatible with a large subset of the OpenStack Swift API.

Example

Log in to the ceph dashboard. Open Object gateway and then go to "Buckets".

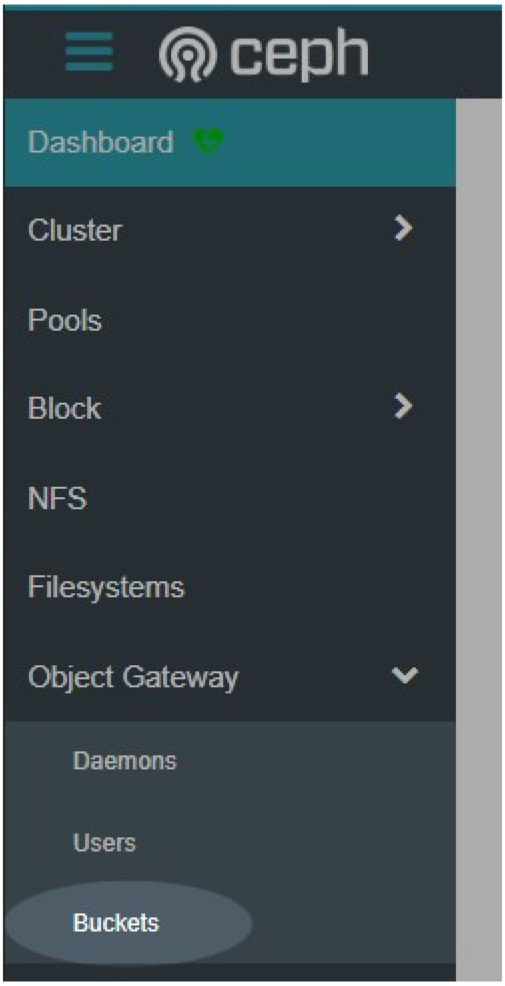

Then click on the "Create" button.

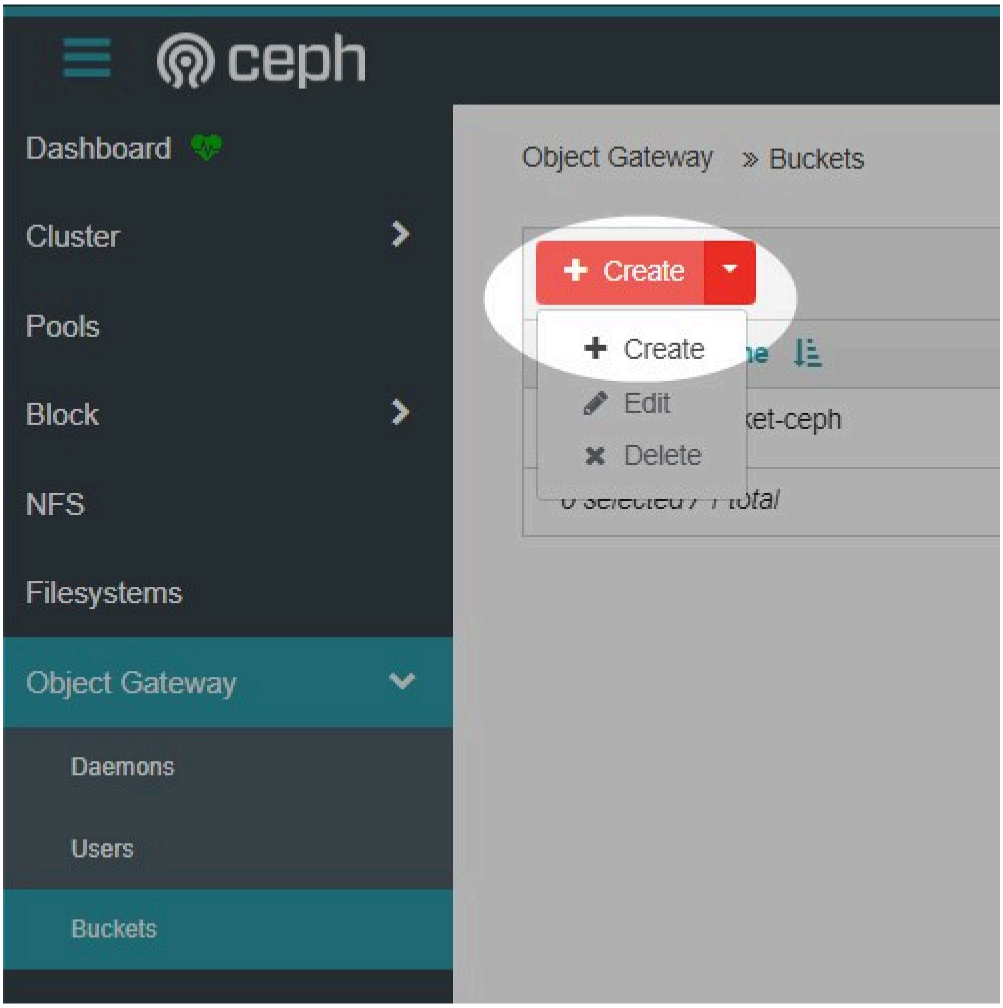

Fill in the required fields.

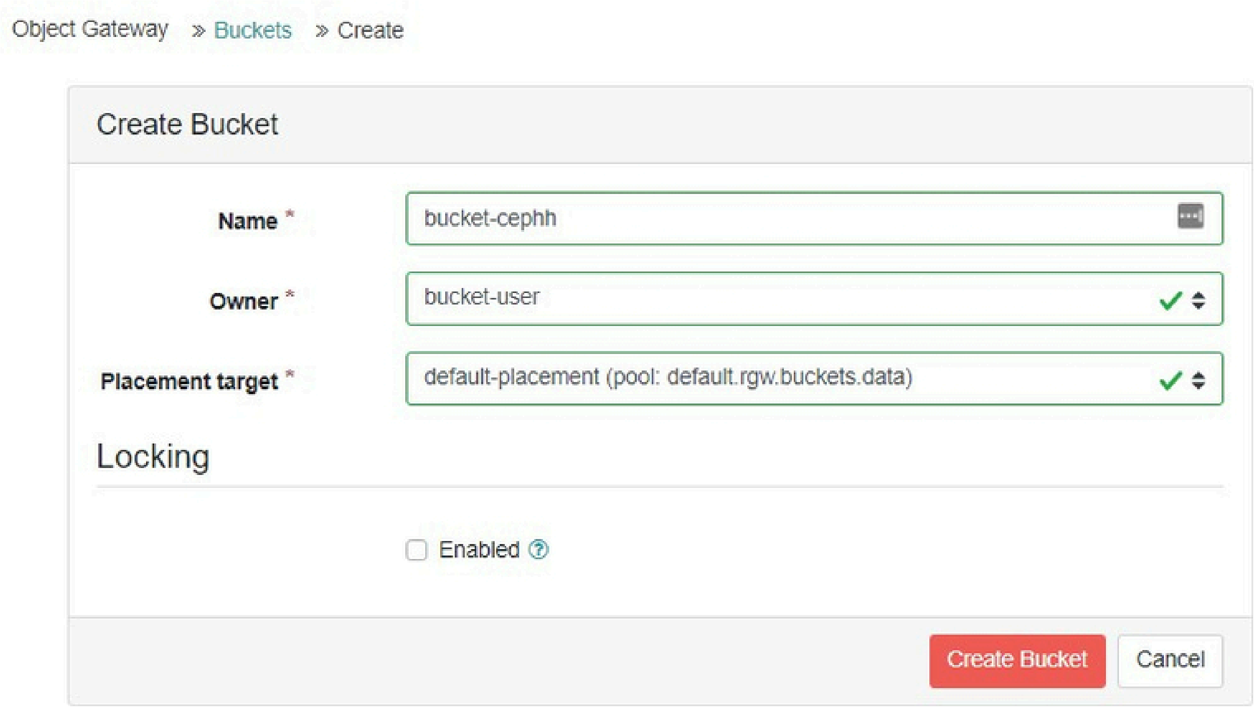

Create Bucket

Name *	bucket-cephh
Owner *	bucket-user
Placement target *	default-placement (pool: default.rgw.buckets.data)

Locking

□ Enabled

Now create a dedicated access account for the backup destination. Open the Users tab under the object gateway menu.

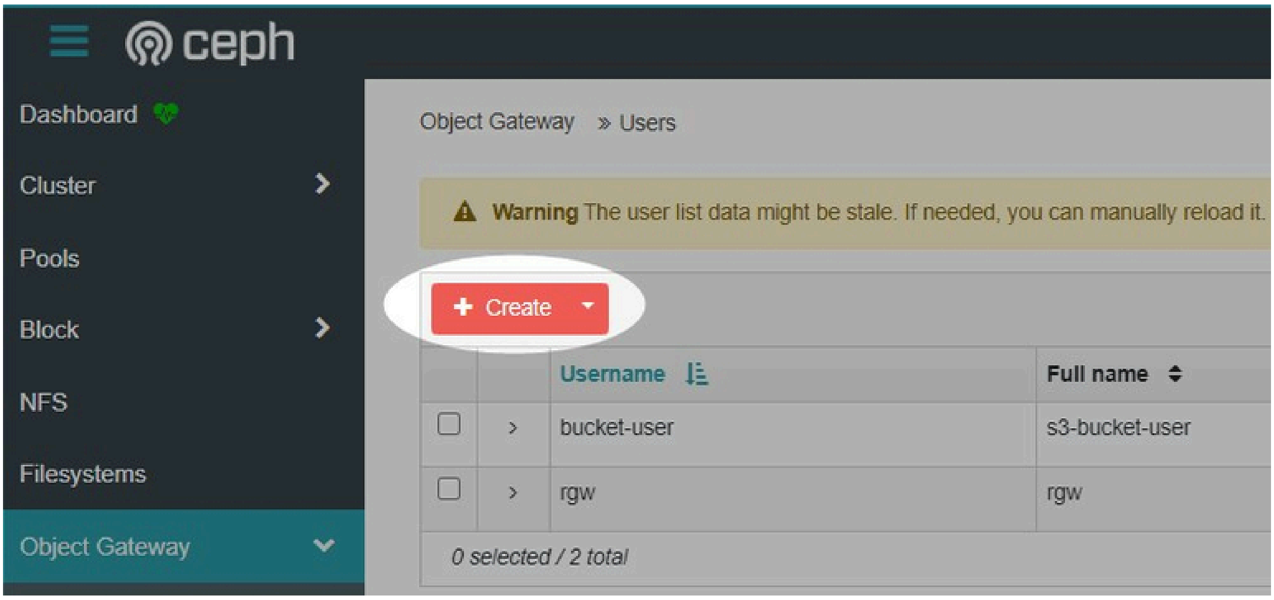

Object Gateway » Users

Warning The user list data might be stale. If needed, you can manually reload it.

		Username	Full name
□	›	bucket-user	s3-bucket-user
□	›	rgw	rgw

0 selected / 2 total

Fill in the username field, you can leave the other settings as default.

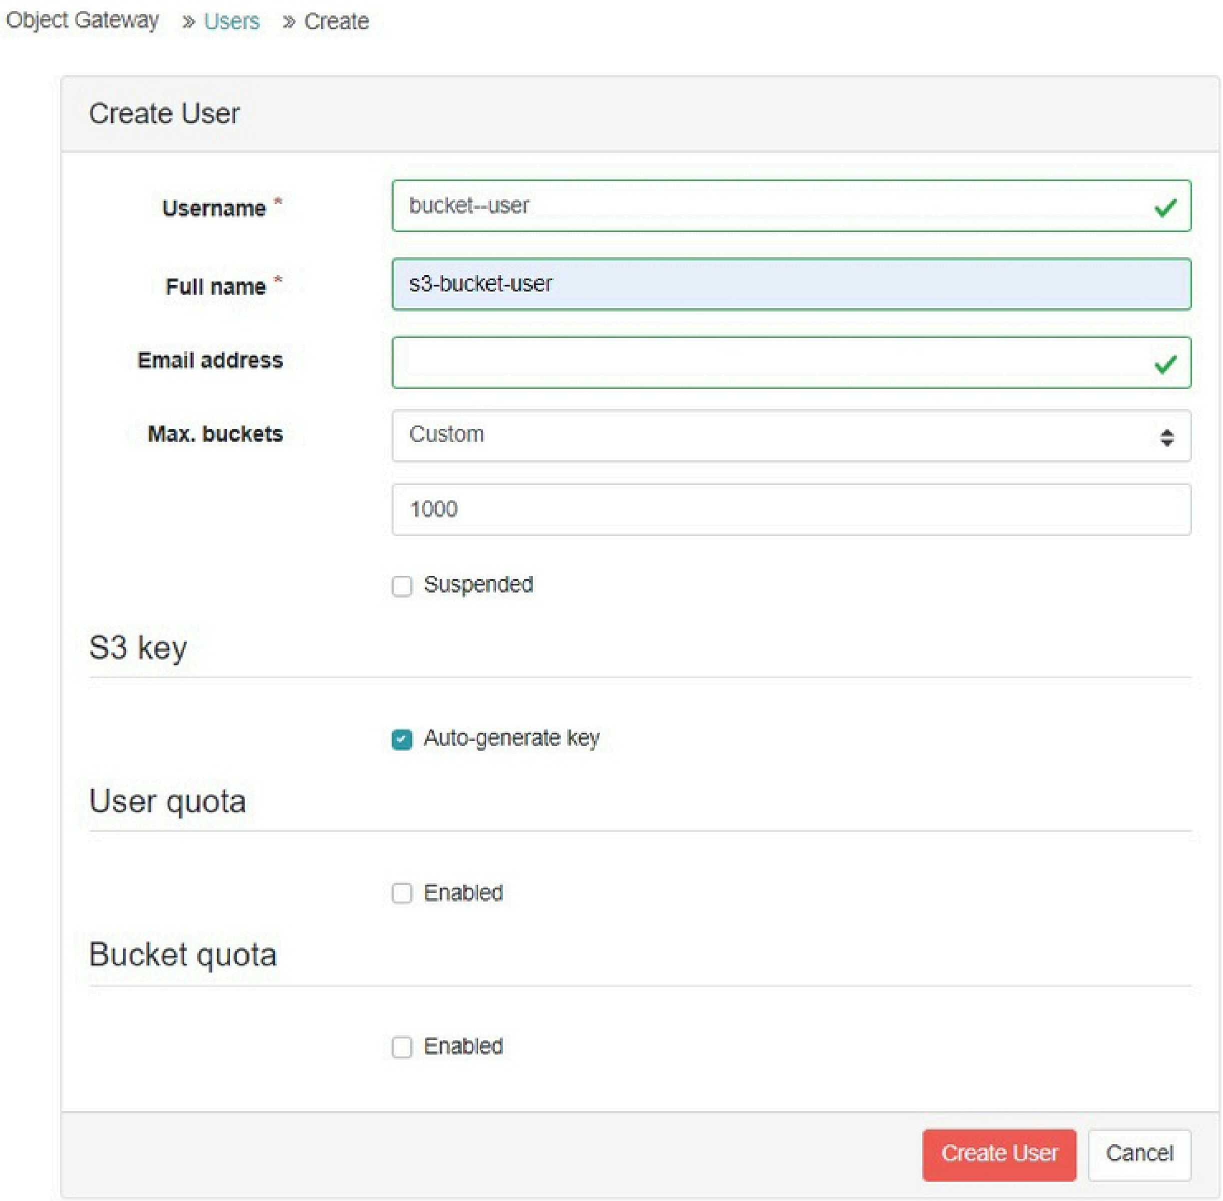

Object Gateway » Users » Create

Create User

Username * bucket--user

Full name * s3-bucket-user

Email address

Max. buckets Custom

1000

□ Suspended

S3 key

☑ Auto-generate key

User quota

□ Enabled

Bucket quota

□ Enabled

To see the account key and secret key, expand the user details and open the keys tab, click on the key, and then on the show button.

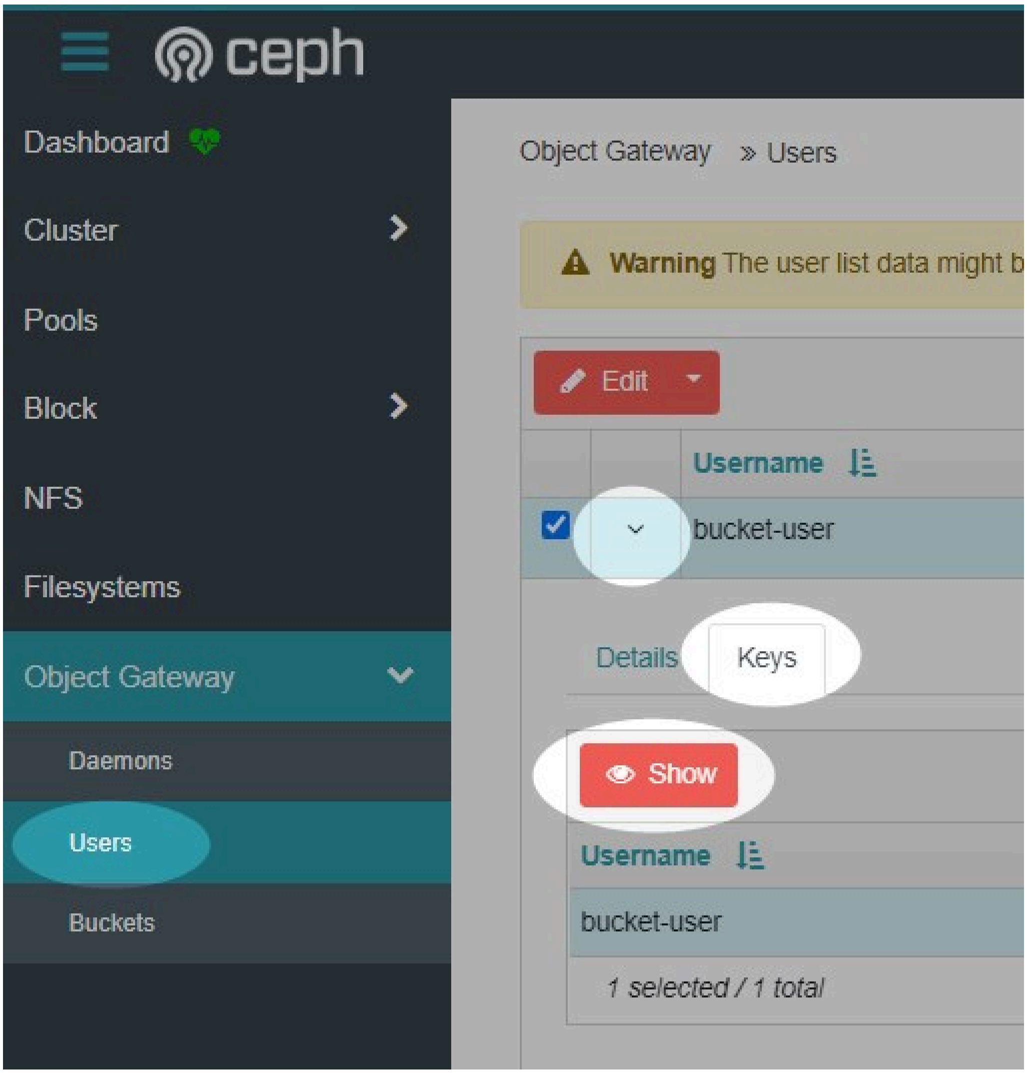

Warning The user list data might b

Username
bucket-user

Username
bucket-user

1 selected / 1 total

The access key and secret key will be needed to create a backup destination in Data Protector for Cloud Workloads.

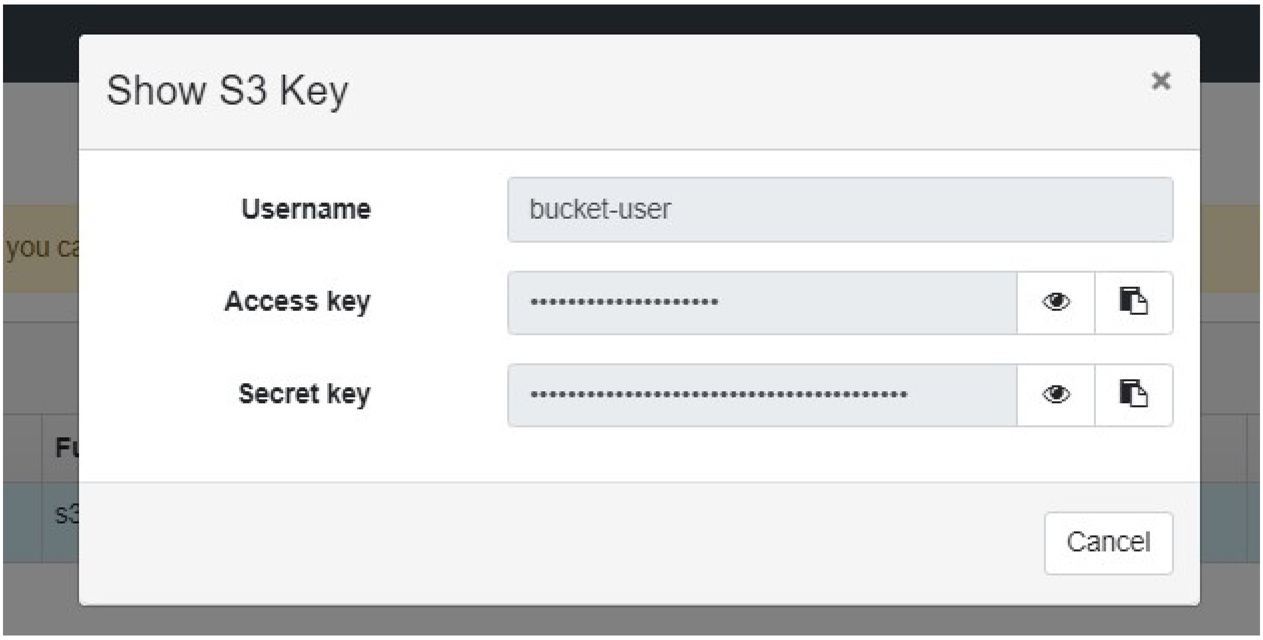

Show S3 Key

Username	bucket-user
Access key	••••••••••••••••••••
Secret key	••

Now we can go to the Data Protector for Cloud Workloads Dashboard. Open the "Backup Destination" tab from the left side menu and choose "Amazon S3 / S3-compatible" as the new type of backup destination.

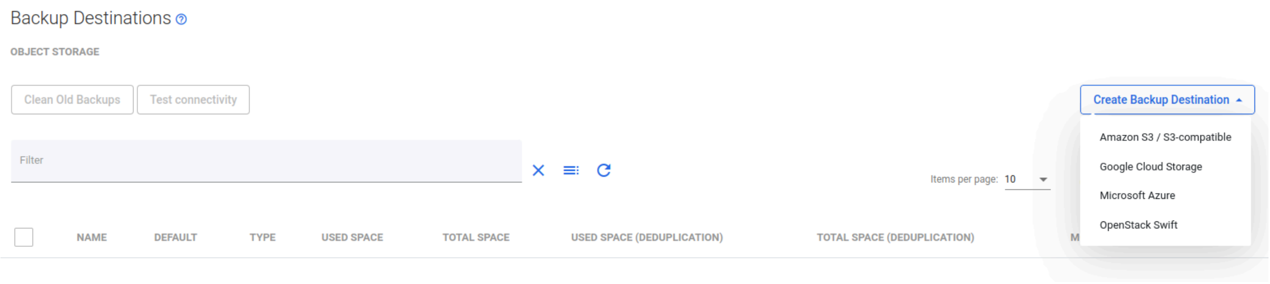

By default, Ceph provides S3 via port 8000. Also, remember to enable the "record backup time after store" option.

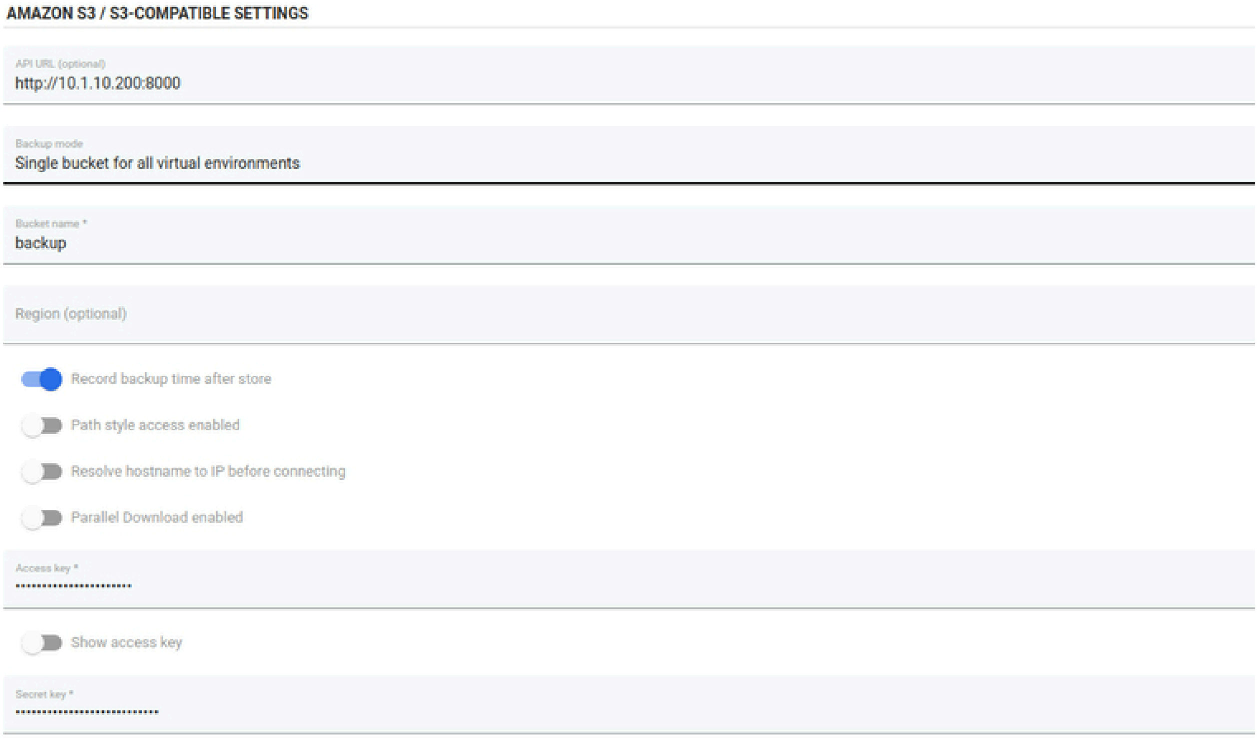

AMAZON S3 / S3-COMPATIBLE SETTINGS

API URL (optional)
http://10.1.10.200:8000

Backup mode
Single bucket for all virtual environments

Bucket name *
backup

Region (optional)

● Record backup time after store

○ Path style access enabled

○ Resolve hostname to IP before connecting

○ Parallel Download enabled

Access key *
••••••••••••••••••••••••

○ Show access key

Secret key *
••••••••••••••••••••••••••••••

Cloudian S3

Overview

Cloudian is an S3-compatible backup provider. Configuration as the backup destination is similar to AWS S3.

Example

After logging in, create a new bucket for your backups

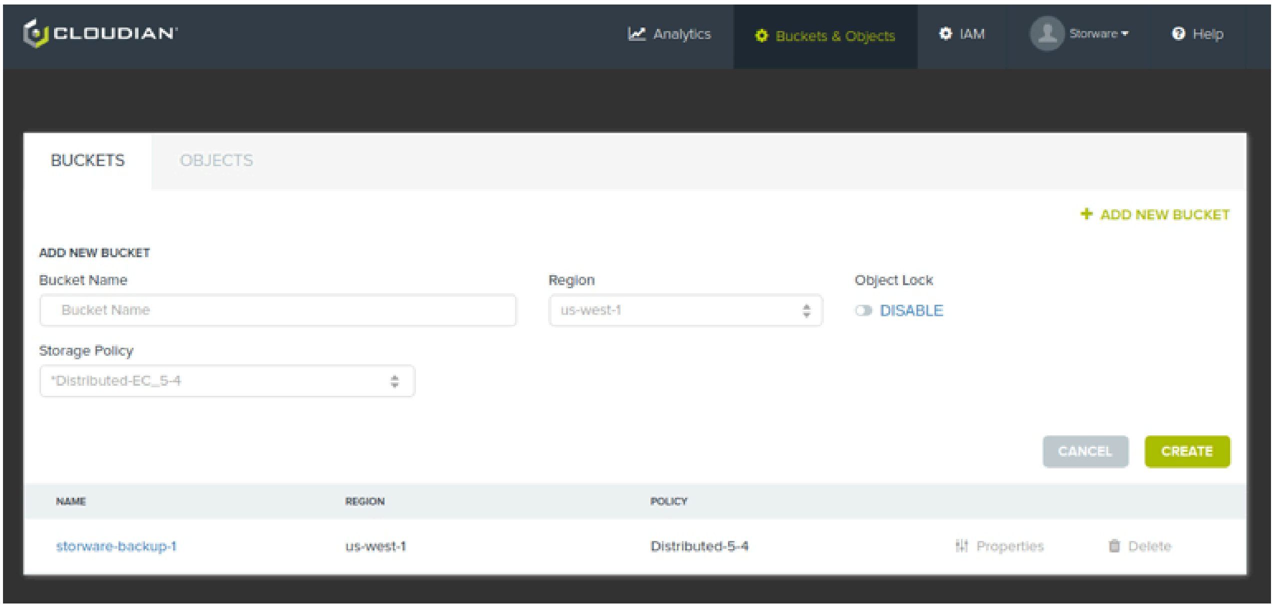

ADD NEW BUCKET

Bucket Name: Bucket Name

Region: us-west-1

Object Lock: DISABLE

Storage Policy: *Distributed-EC_5-4

NAME	REGION	POLICY		
storware-backup-1	us-west-1	Distributed-5-4	Properties	Delete

Next, go to security credentials and generate a new access key.

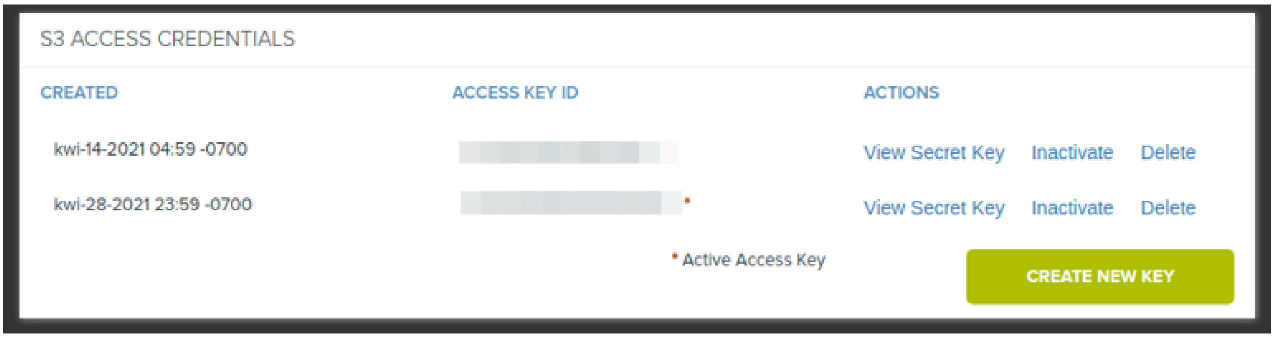

S3 ACCESS CREDENTIALS

CREATED	ACCESS KEY ID	ACTIONS		
kwi-14-2021 04:59 -0700	[redacted]	View Secret Key	Inactivate	Delete
kwi-28-2021 23:59 -0700	[redacted] *	View Secret Key	Inactivate	Delete

* Active Access Key

Now go to the Backup destination tab on the Data Protector for Cloud Workloads dashboard and change the sub-tab to object storage. Provide the bucket name and key credentials, and then configure the remaining options according to your requirements. Also, enable `Path style access enabled` option:

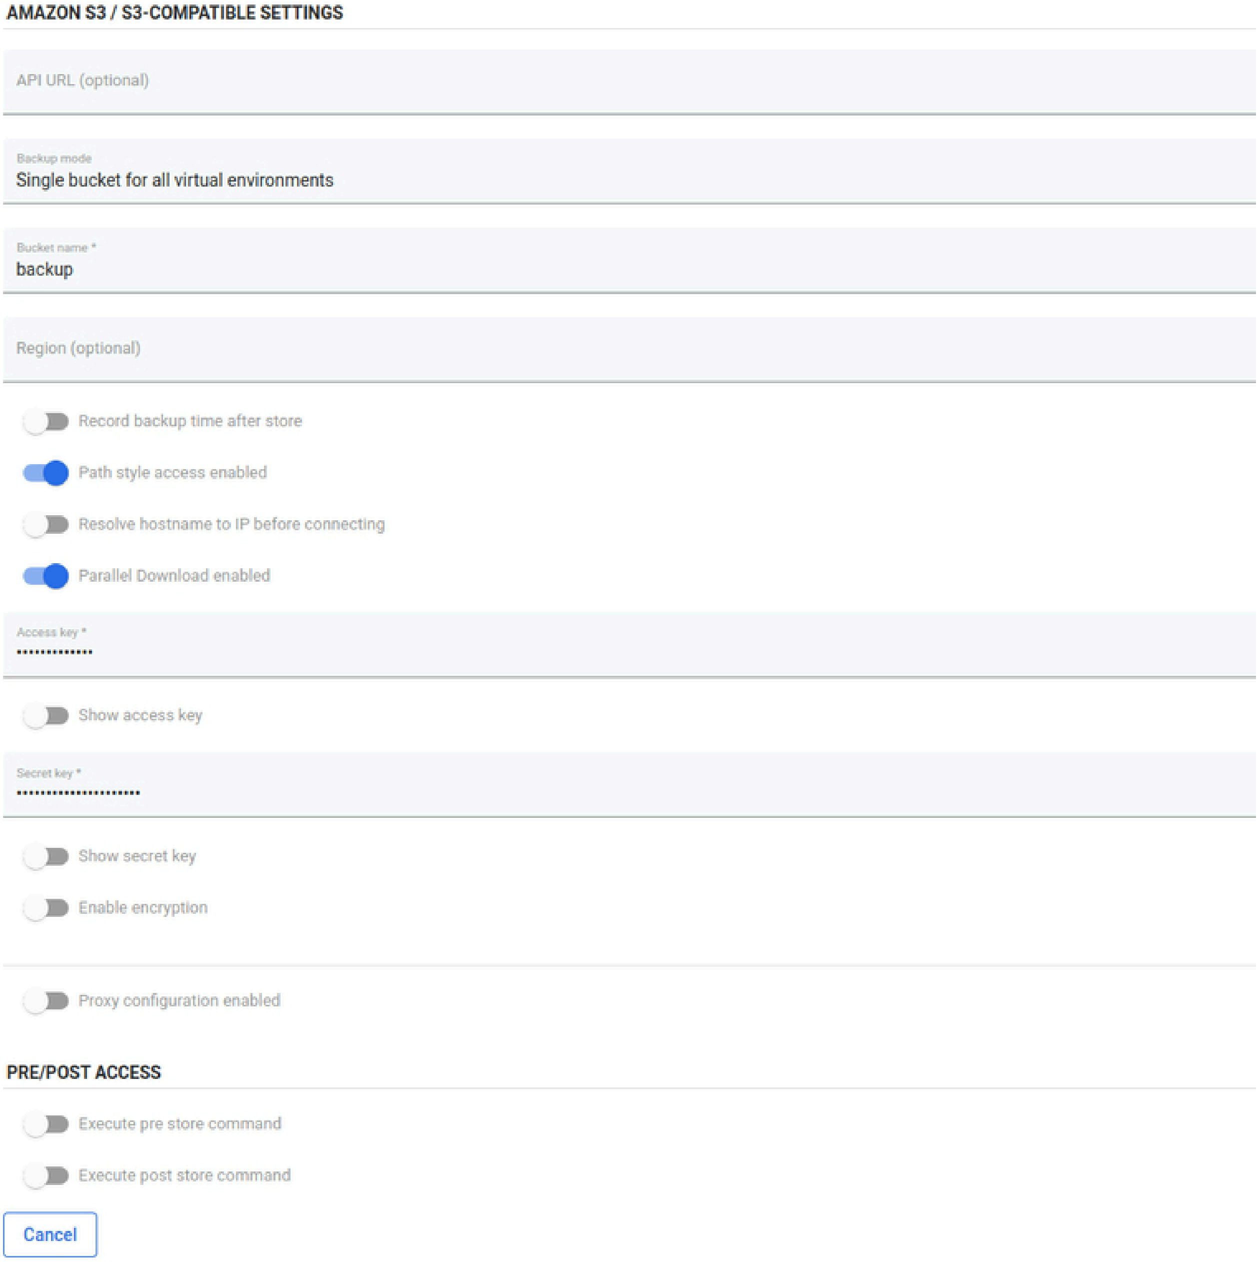

AMAZON S3 / S3-COMPATIBLE SETTINGS

API URL (optional)

Backup mode
Single bucket for all virtual environments

Bucket name *
backup

Region (optional)

○ Record backup time after store
● Path style access enabled
○ Resolve hostname to IP before connecting
● Parallel Download enabled

Access key *
•••••••••••••

○ Show access key

Secret key *
•••••••••••••••••••••

○ Show secret key
○ Enable encryption

○ Proxy configuration enabled

PRE/POST ACCESS

○ Execute pre store command
○ Execute post store command

Google Cloud Storage

Google Cloud Storage allows data to be stored and accessed on Google Cloud Platform infrastructure. It combines the performance and scalability of Google's cloud with advanced security and sharing capabilities.

How to use GCS as a backup destination for Data Protector for Cloud Workloads:

1. Create a project: Click here for more info about **Creating and Managing Projects**.
2. Create a bucket: Click here for more info about **Creating Storage Buckets**.

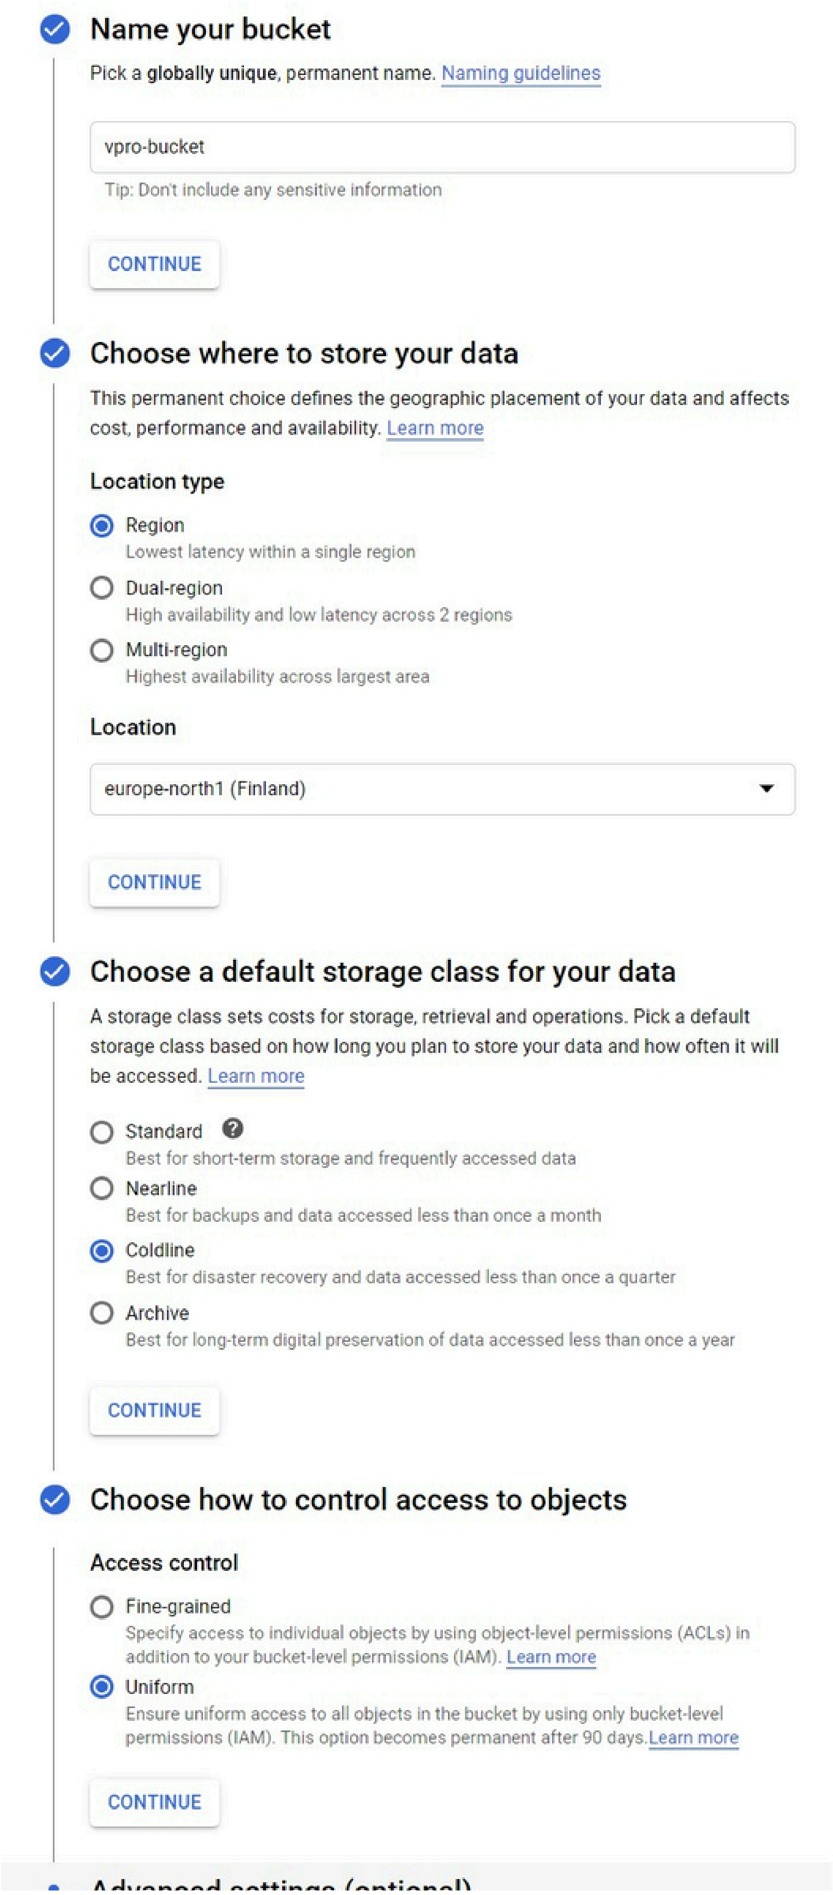

Name your bucket

Pick a **globally unique**, permanent name. Naming guidelines

vpro-bucket

Tip: Don't include any sensitive information

Choose where to store your data

This permanent choice defines the geographic placement of your data and affects cost, performance and availability. Learn more

Location type

● Region
Lowest latency within a single region

○ Dual-region
High availability and low latency across 2 regions

○ Multi-region
Highest availability across largest area

Location

europe-north1 (Finland)

Choose a default storage class for your data

A storage class sets costs for storage, retrieval and operations. Pick a default storage class based on how long you plan to store your data and how often it will be accessed. Learn more

○ Standard
Best for short-term storage and frequently accessed data

○ Nearline
Best for backups and data accessed less than once a month

● Coldline
Best for disaster recovery and data accessed less than once a quarter

○ Archive
Best for long-term digital preservation of data accessed less than once a year

Choose how to control access to objects

Access control

○ Fine-grained
Specify access to individual objects by using object-level permissions (ACLs) in addition to your bucket-level permissions (IAM). Learn more

● Uniform
Ensure uniform access to all objects in the bucket by using only bucket-level permissions (IAM). This option becomes permanent after 90 days.Learn more

Advanced settings (optional)

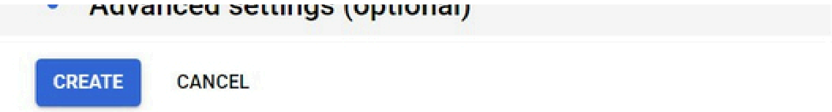

Advanced settings (optional)

CREATE CANCEL

3. Enable versioning in your bucket: Click [here](#) for more info about **Enabling Object Versioning**.
4. Generate a service account key: Click [here](#) for more info about **Creating service account keys**. The service account key should have the **Role** set to **Storage Admin and Service Usage Consumer**.

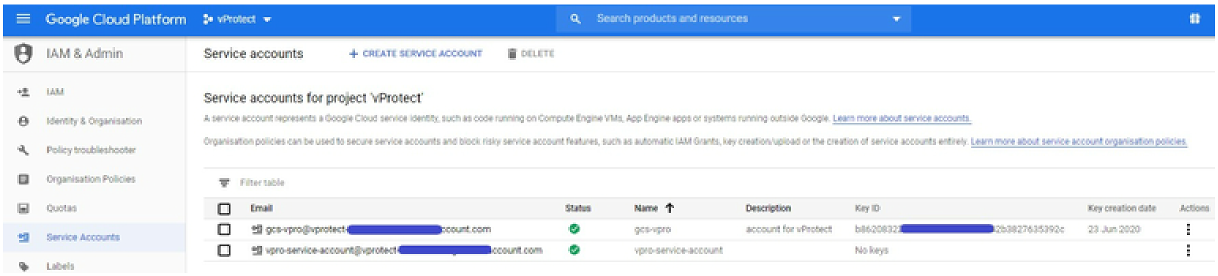

Service accounts

Service accounts for project 'vProtect'

A service account represents a Google Cloud service identity, such as code running on Compute Engine VMs, App Engine apps or systems running outside Google. Learn more about service accounts.

Organisation policies can be used to secure service accounts and block risky service account features, such as automatic IAM Grants, key creation/upload or the creation of service accounts entirely. Learn more about service account organisation policies.

□	Email	Status	Name ↑	Description	Key ID	Key creation date	Actions
□	gcs-vpro@vprotect-[redacted]ccount.com	✓	gcs-vpro	account for vProtect	b86208322[redacted]2b3827635392c	23 Jun 2020	⋮
□	vpro-service-account@vprotect-[redacted]ccount.com	✓	vpro-service-account		No keys		⋮

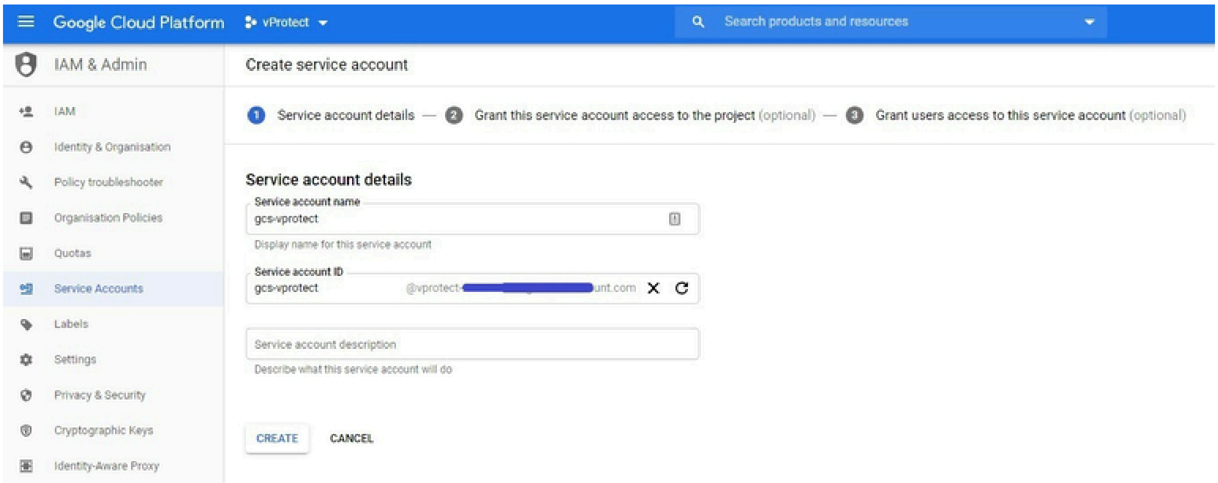

Create service account

1 Service account details — 2 Grant this service account access to the project (optional) — 3 Grant users access to this service account (optional)

Service account details

Service account name: gcs-vprotect

Display name for this service account

Service account ID: gcs-vprotect @vprotect-[redacted]unt.com

Service account description

Describe what this service account will do

CREATE CANCEL

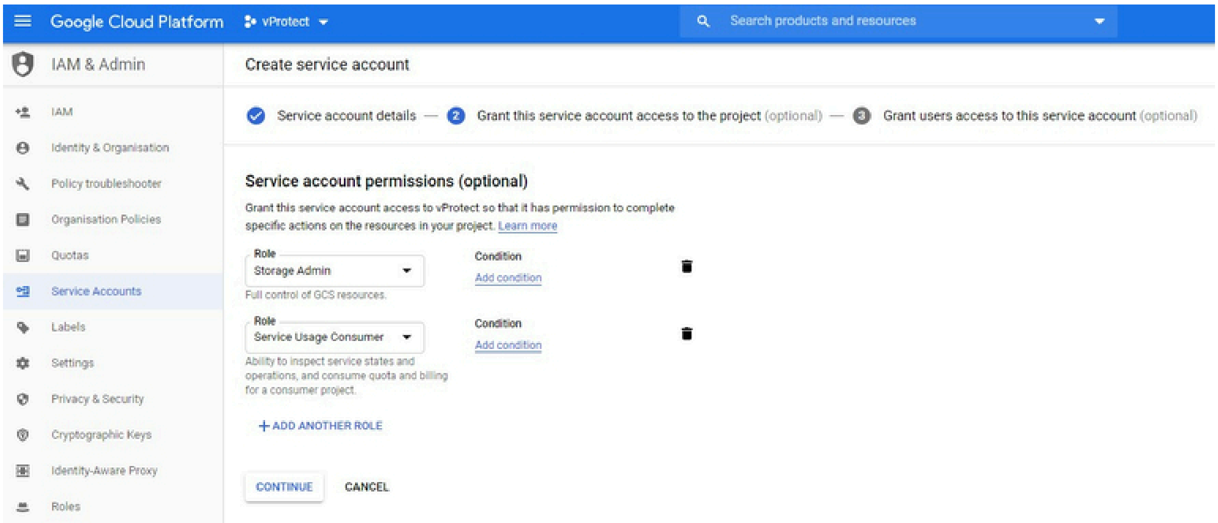

Create service account

✓ Service account details — 2 Grant this service account access to the project (optional) — 3 Grant users access to this service account (optional)

Service account permissions (optional)

Grant this service account access to vProtect so that it has permission to complete specific actions on the resources in your project. Learn more

Role: Storage Admin — Condition: Add condition

Full control of GCS resources.

Role: Service Usage Consumer — Condition: Add condition

Ability to inspect service states and operations, and consume quota and billing for a consumer project.

+ ADD ANOTHER ROLE

CONTINUE CANCEL

- You can leave the third tab - Grant users access to this service account (optional).\
- To generate an account key, click on the "three-dot" button next to your service account and then click on "create key". You should then see the

window below - click on create to download the JSON file. You'll need its content in the last step.

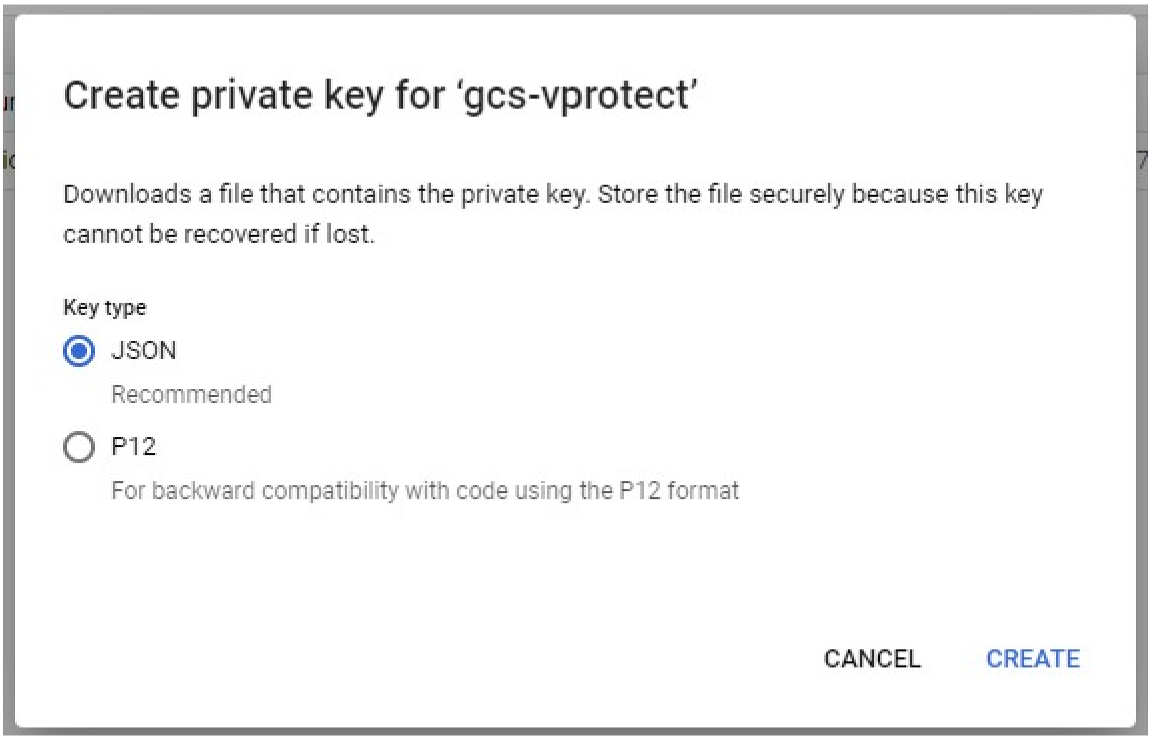

Create private key for 'gcs-vprotect'

Downloads a file that contains the private key. Store the file securely because this key cannot be recovered if lost.

Key type

● JSON

Recommended

○ P12

For backward compatibility with code using the P12 format

5. After the key is created, open your Data Protector for Cloud Workloads Web UI (you can also use **CLI**), click on **BACKUP DESTINATIONS**, then on the **Create Backup Destination** button, and then select **Google Cloud Storage** from the drop-down list. In addition to the standard properties, you need to specify:
6. The **Bucket name** was specified during bucket creation.
7. The **Service account key** - paste the content of the service account key .json file created before.

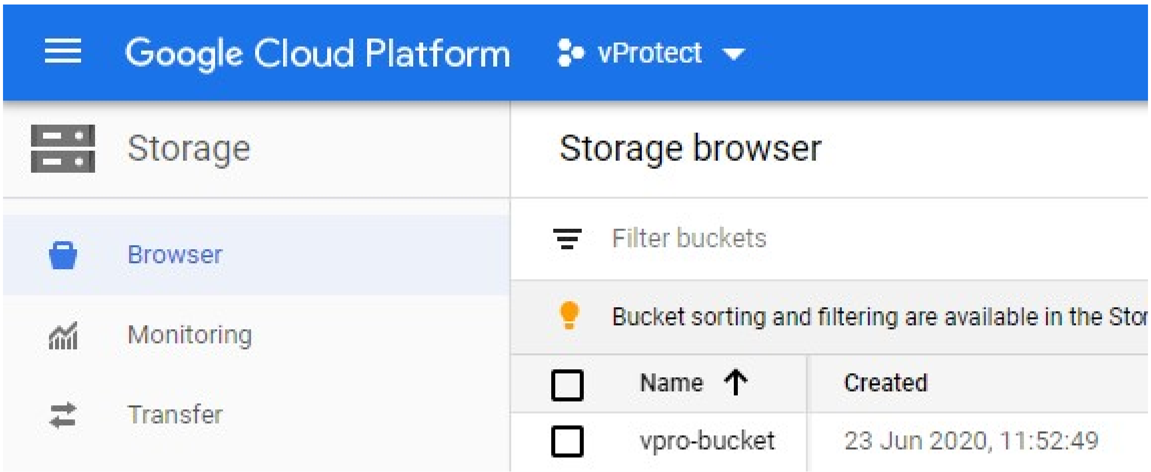

Bucket sorting and filtering are available in the Sto

□	Name ↑	Created
□	vpro-bucket	23 Jun 2020, 11:52:49

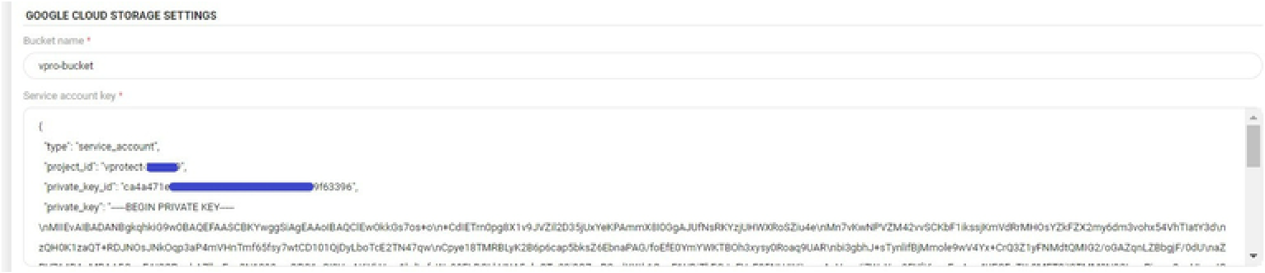

GOOGLE CLOUD STORAGE SETTINGS

Bucket name *

vpro-bucket

Service account key *

```
{
  "type": "service_account",
  "project_id": "vprotect-[redacted]",
  "private_key_id": "ca4a471e[redacted]9f63396",
  "private_key": "-----BEGIN PRIVATE KEY-----
\nMIIEvAIBADANBgkqhkiG9w0BAQEFAASCBKYwggSiAgEAAoIBAQCIEwOkk0s7os+o\n+CdIETm0pg8X1v9JVZII2D35jUxYeKPAmmX8I0GgAJUfNsRKYzjUHWXRoSZiu4e\nMn7vKwNPVZM42vvSCKbF1ikssjKmVdRrMHOsYZkFZX2my6dm3vohx54VhTlatY3d\n
zQH0K1zaQT+RDJNOsJNkOqp3aP4mVHnTmf65fsy7wtCD101QjDyLboTcE2TN47qw\nCpye18TMRBLyK2B6p6cap5bksZ6EbnaPAG/foEfE0YmYWKTBOh3xysy0Roaq9UAR\nbi3gbhJ+sTynlifBjMmole9wV4Yx+CrQ3Z1yFNMdtQMIG2/oGAZqnLZ8bgjF/0dU\naZ
```

Now, you can store Data Protector for Cloud Workloads backups on Google Cloud Storage.

IBM Cloud Object Storage

Overview

IBM Cloud Object Storage is a push-button deployed cloud storage service and is available in IBM Cloud global data centers. It offers leading data protection, high durability, and fast access to your data. You can use it to store and protect data with easy-to-use management features to organize your data and to configure finely-tuned access controls.

Example

Log in to your IBM Cloud account. On the main dashboard, you will see the "Create a resource" button - Click on it.

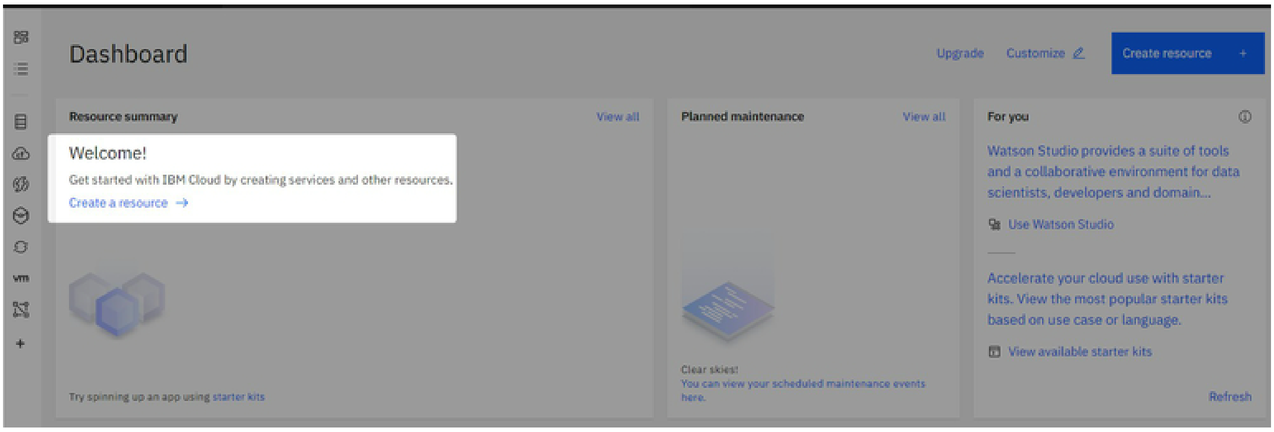

Welcome!

Get started with IBM Cloud by creating services and other resources.

Try spinning up an app using starter kits

Clear skies!
You can view your scheduled maintenance events here.

Watson Studio provides a suite of tools and a collaborative environment for data scientists, developers and domain...

Accelerate your cloud use with starter kits. View the most popular starter kits based on use case or language.

On the next screen, search for a resource named "Object Storage" and click on it.

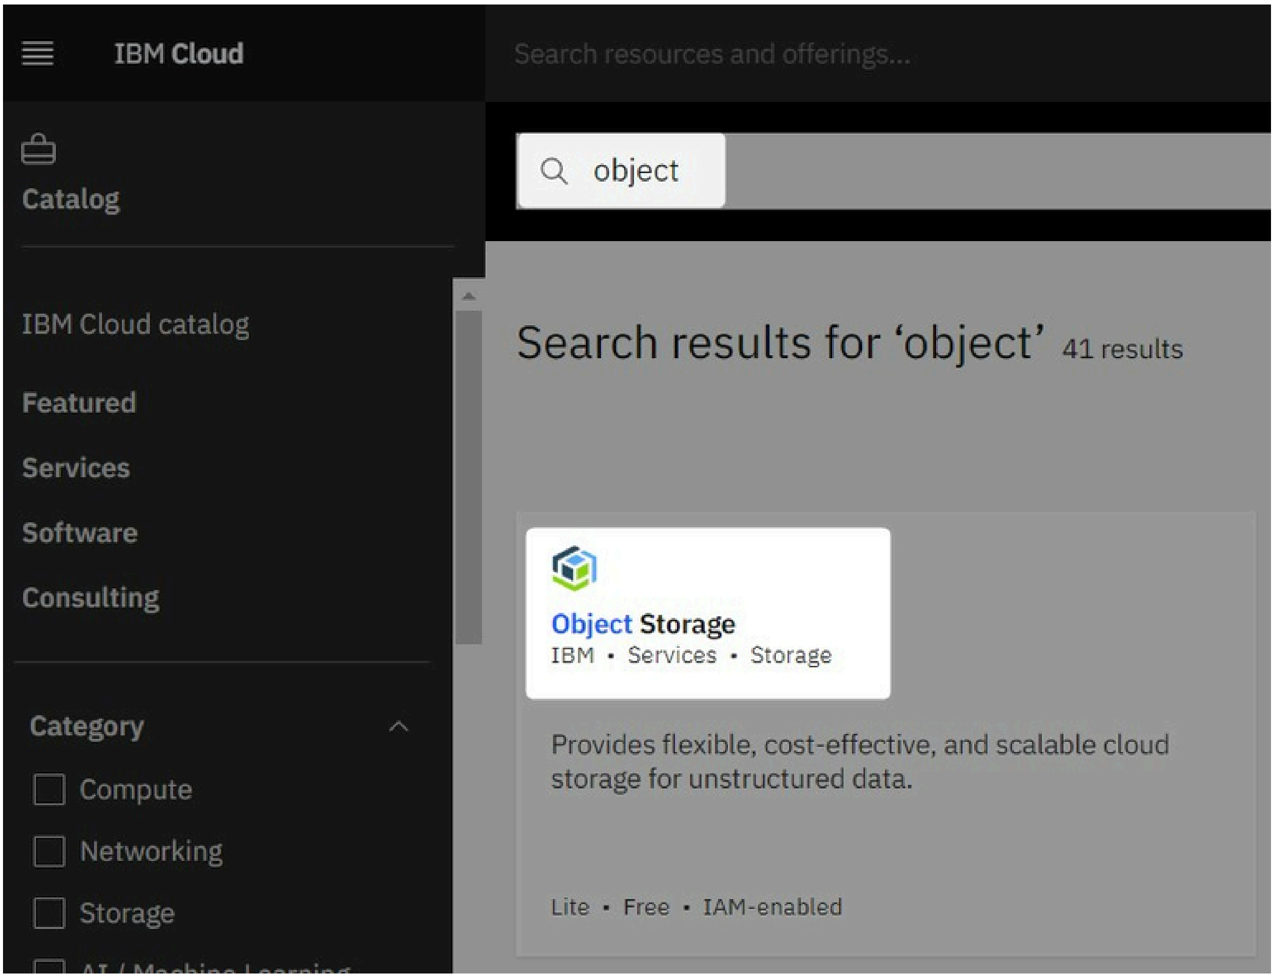

Search results for 'object' 41 results

Object Storage
IBM • Services • Storage

Provides flexible, cost-effective, and scalable cloud storage for unstructured data.

Lite • Free • IAM-enabled

On the next screen, you can choose piercing plans, etc. Select the options according to your requirements.

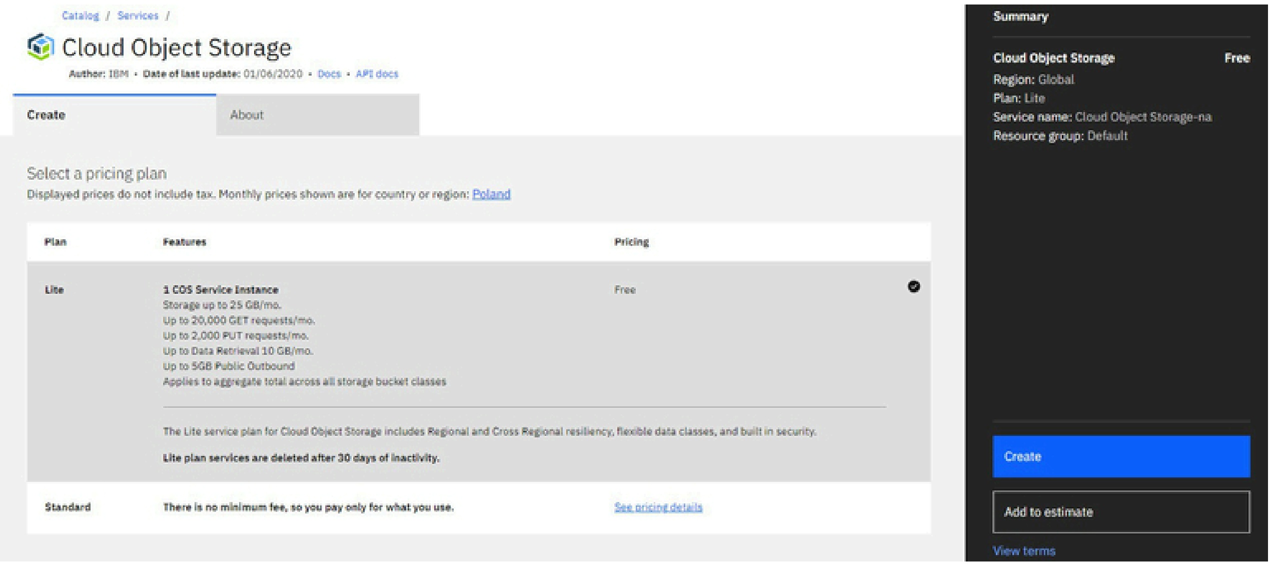

Cloud Object Storage

Author: IBM • Date of last update: 01/06/2020 • Docs • API docs

Select a pricing plan

Displayed prices do not include tax. Monthly prices shown are for country or region: Poland

Plan	Features	Pricing	
Lite	**1 COS Service Instance** Storage up to 25 GB/mo. Up to 20,000 GET requests/mo. Up to 2,000 PUT requests/mo. Up to Data Retrieval 10 GB/mo. Up to 5GB Public Outbound Applies to aggregate total across all storage bucket classes The Lite service plan for Cloud Object Storage includes Regional and Cross Regional resiliency, flexible data classes, and built in security. **Lite plan services are deleted after 30 days of inactivity.**	Free	●
Standard	There is no minimum fee, so you pay only for what you use.	See pricing details	

Summary

Cloud Object Storage **Free**
Region: Global
Plan: Lite
Service name: Cloud Object Storage-na
Resource group: Default

After creating a storage resource we need to create a bucket.

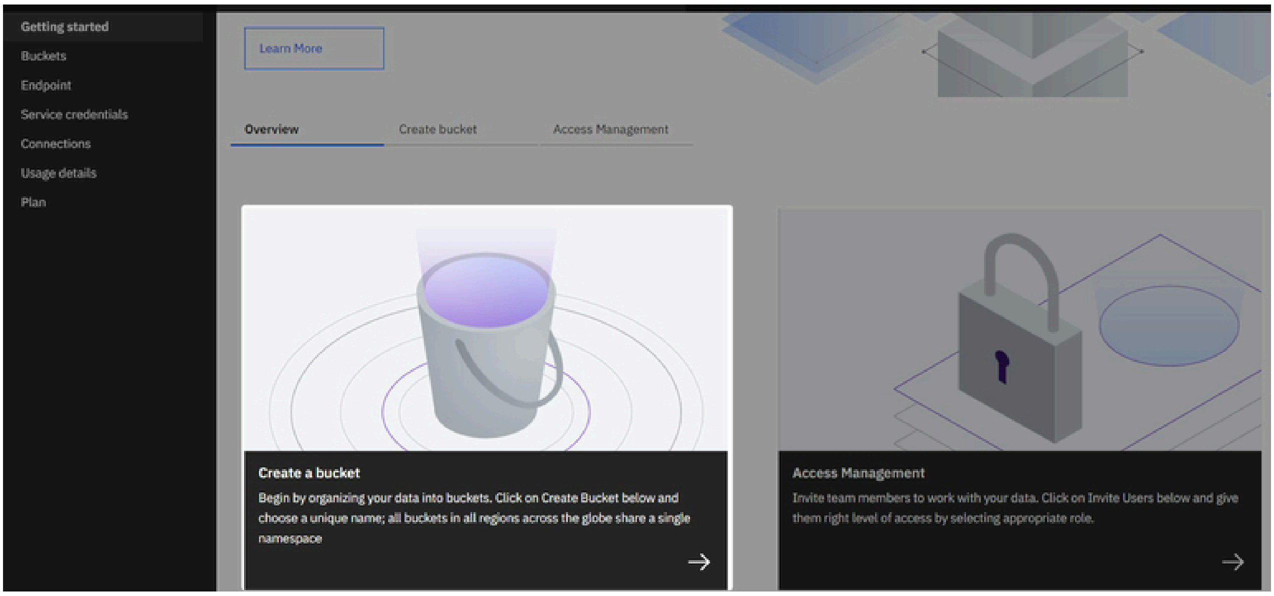

Create a bucket

Begin by organizing your data into buckets. Click on Create Bucket below and choose a unique name; all buckets in all regions across the globe share a single namespace

Access Management

Invite team members to work with your data. Click on Invite Users below and give them right level of access by selecting appropriate role.

You can choose predefined templates or select the option to create a bucket with your own settings. In this example, we will choose "Custom bucket".

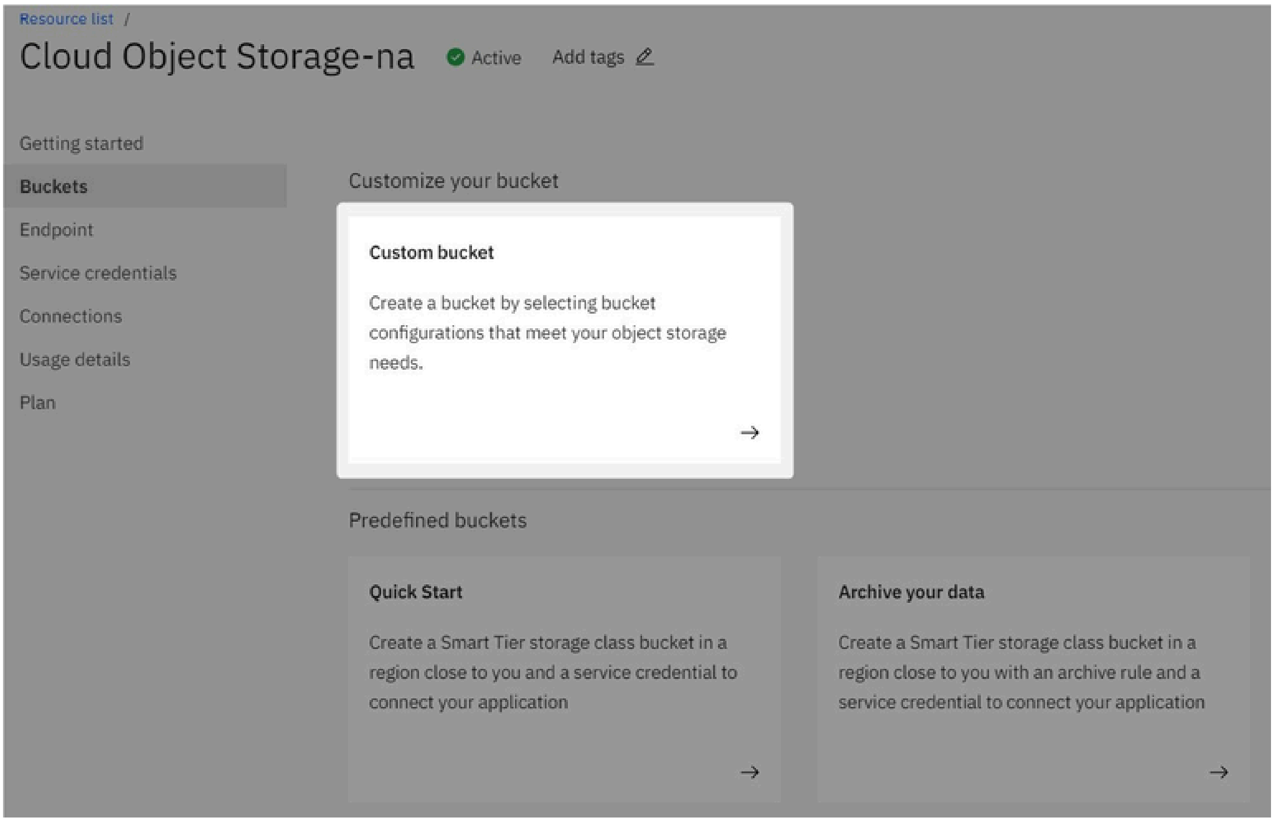

Custom bucket

Create a bucket by selecting bucket configurations that meet your object storage needs.

Quick Start

Create a Smart Tier storage class bucket in a region close to you and a service credential to connect your application

Archive your data

Create a Smart Tier storage class bucket in a region close to you with an archive rule and a service credential to connect your application

Data Protector for Cloud Workloads has no special requirements for the bucket, all options can be configured according to customer needs.

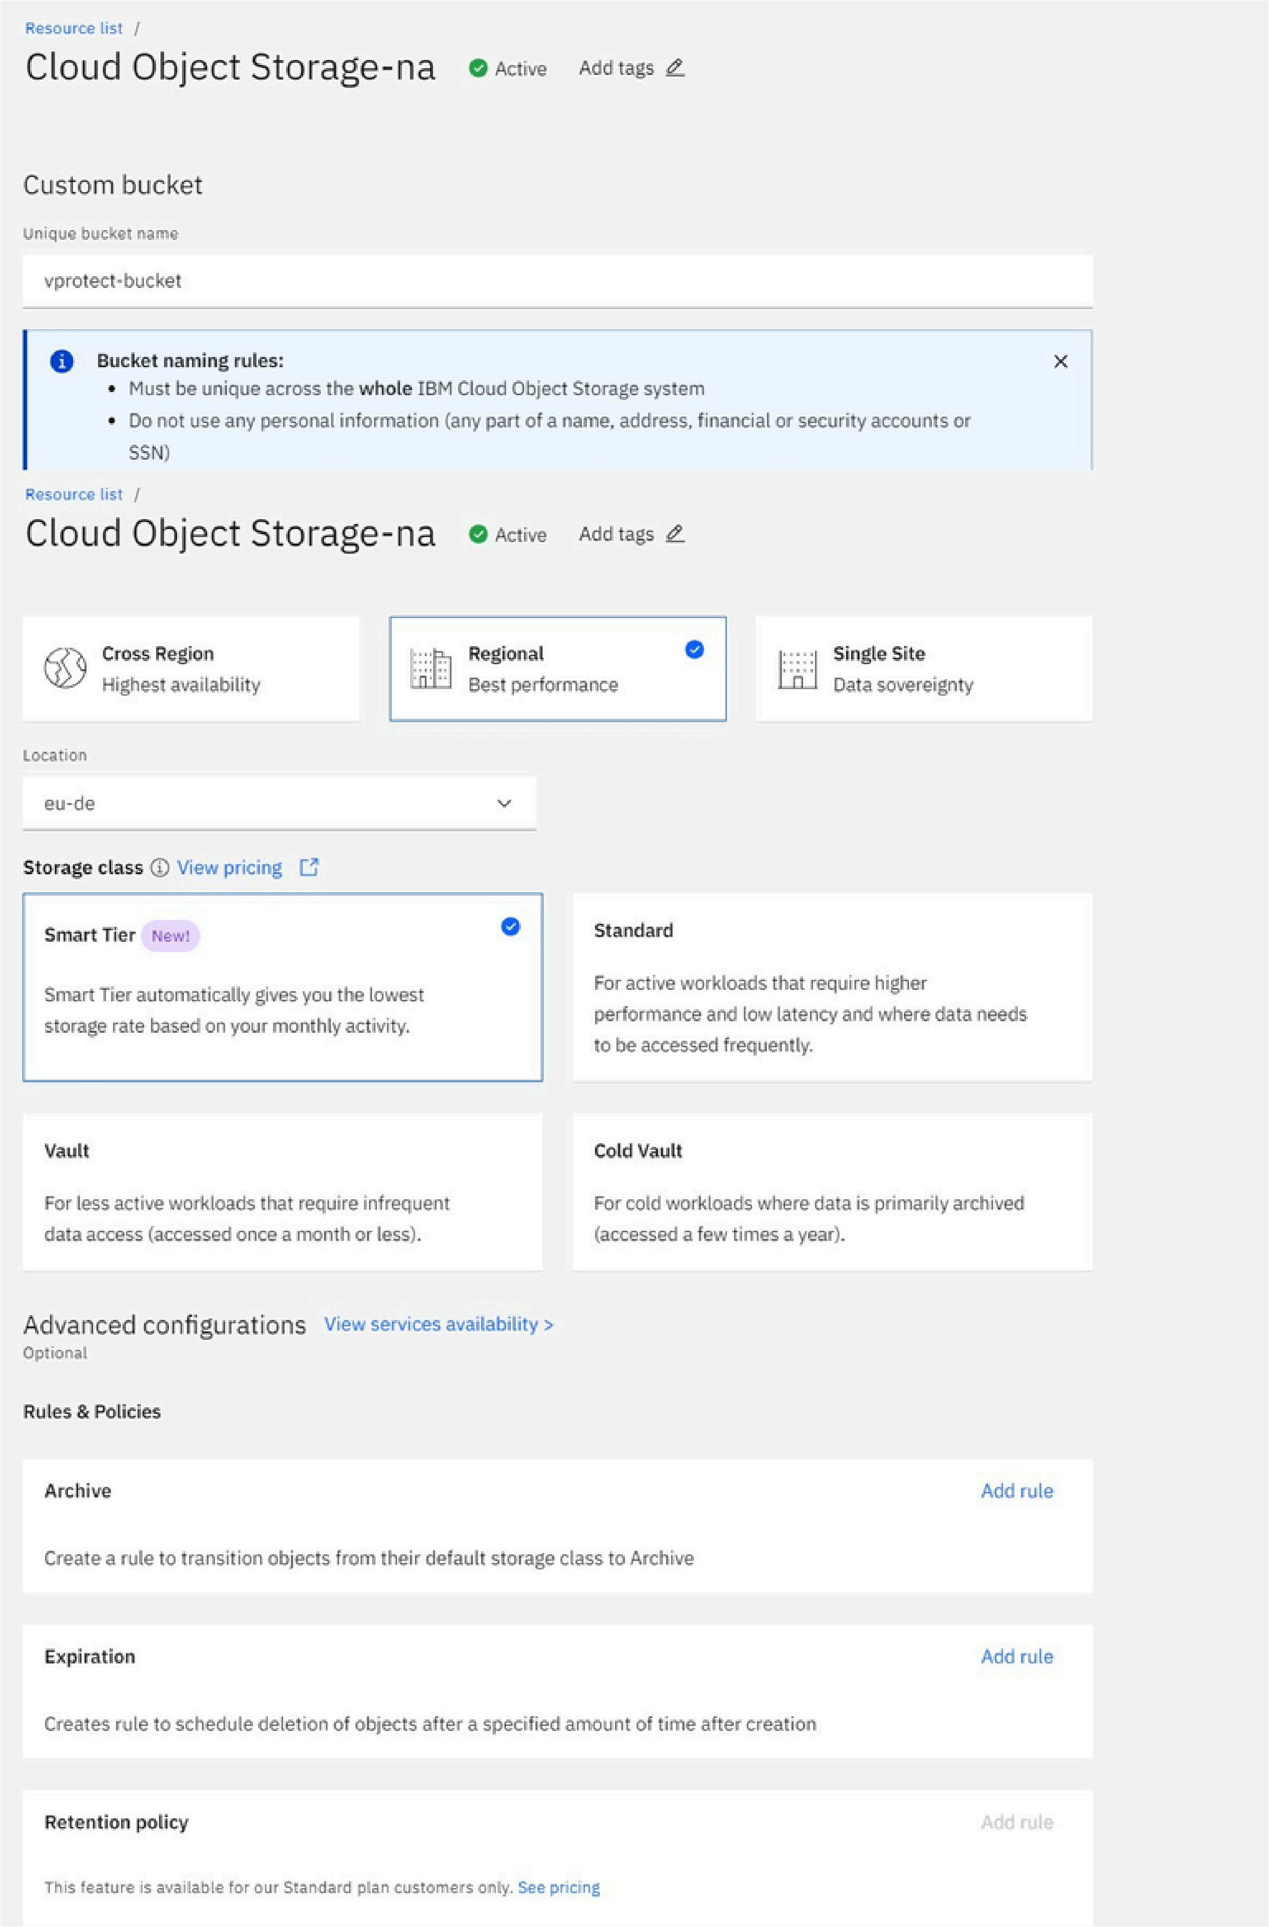

Resource list /

Cloud Object Storage-na ● Active Add tags

Custom bucket

Unique bucket name

vprotect-bucket

Bucket naming rules:

- Must be unique across the **whole** IBM Cloud Object Storage system
- Do not use any personal information (any part of a name, address, financial or security accounts or SSN)

Resource list /

Cloud Object Storage-na ● Active Add tags

Cross Region
Highest availability

Regional ✓
Best performance

Single Site
Data sovereignty

Location

eu-de

Storage class View pricing

Smart Tier New! ✓

Smart Tier automatically gives you the lowest storage rate based on your monthly activity.

Standard

For active workloads that require higher performance and low latency and where data needs to be accessed frequently.

Vault

For less active workloads that require infrequent data access (accessed once a month or less).

Cold Vault

For cold workloads where data is primarily archived (accessed a few times a year).

Advanced configurations View services availability >

Optional

Rules & Policies

Archive Add rule

Create a rule to transition objects from their default storage class to Archive

Expiration Add rule

Creates rule to schedule deletion of objects after a specified amount of time after creation

Retention policy Add rule

This feature is available for our Standard plan customers only. See pricing

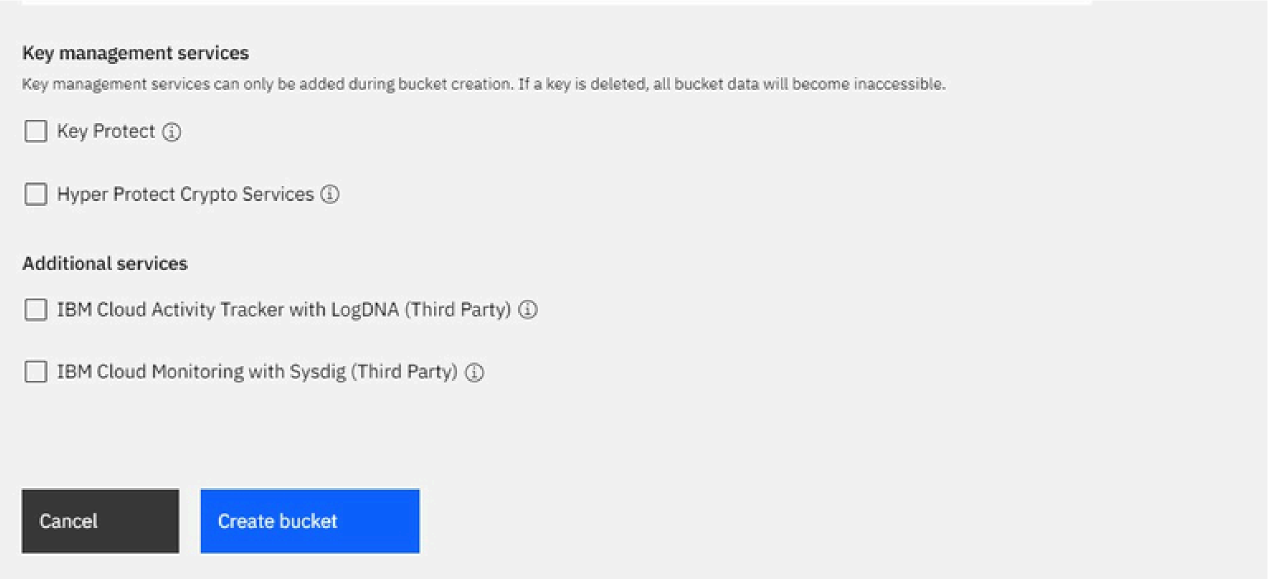

Key management services

Key management services can only be added during bucket creation. If a key is deleted, all bucket data will become inaccessible.

□ Key Protect ⓘ

□ Hyper Protect Crypto Services ⓘ

Additional services

□ IBM Cloud Activity Tracker with LogDNA (Third Party) ⓘ

□ IBM Cloud Monitoring with Sysdig (Third Party) ⓘ

After creating the bucket, you'll see the objects page. From the menu on the left select the configuration tab. You will see a summary of the resource you have created. To create a backup destination you will need the ***"public" address from the endpoints section*** from here.

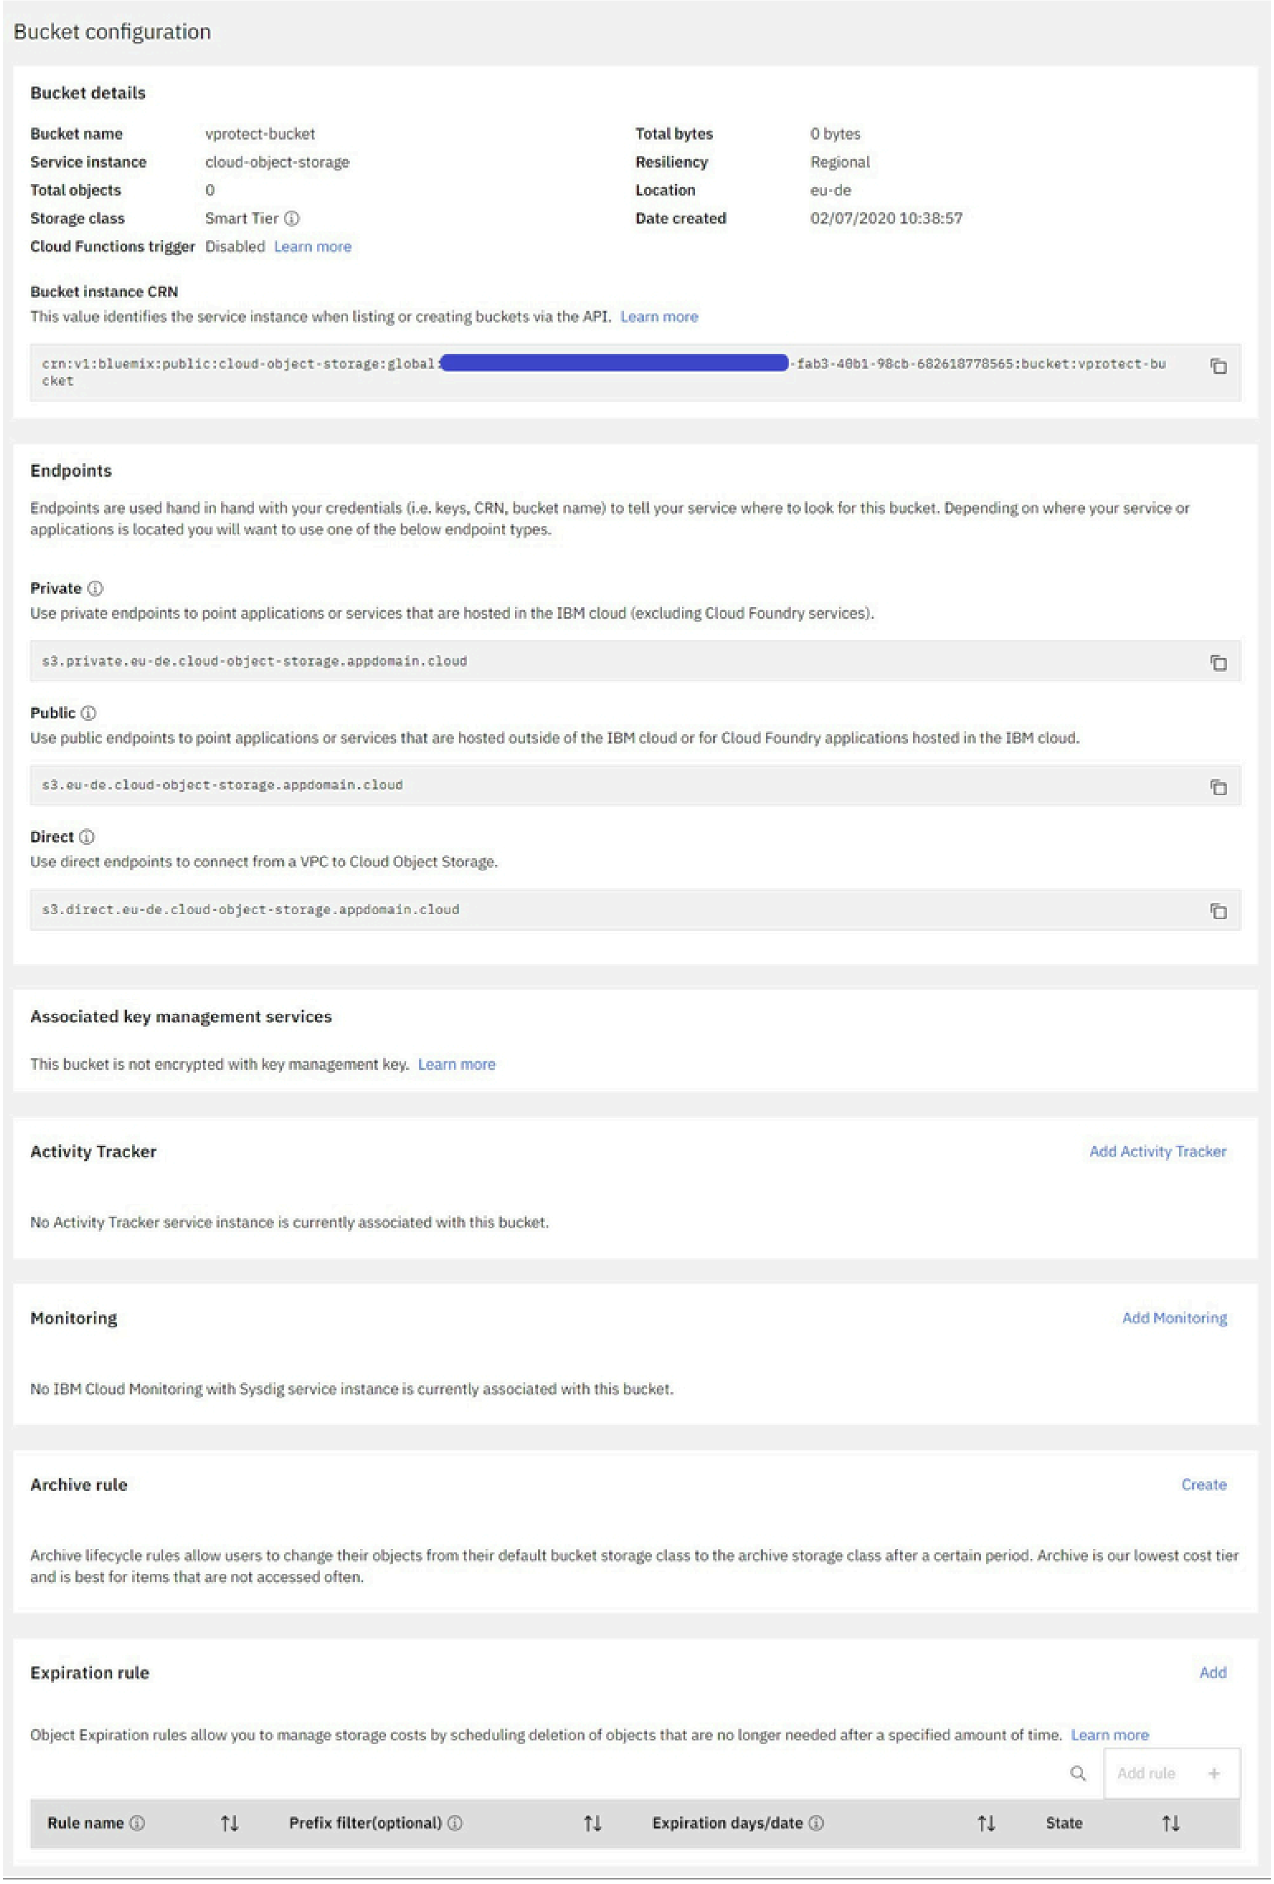

Bucket configuration

Bucket details

Bucket name	vprotect-bucket	**Total bytes**	0 bytes
Service instance	cloud-object-storage	**Resiliency**	Regional
Total objects	0	**Location**	eu-de
Storage class	Smart Tier ⓘ	**Date created**	02/07/2020 10:38:57
Cloud Functions trigger	Disabled Learn more		

Bucket instance CRN

This value identifies the service instance when listing or creating buckets via the API. Learn more

```
crn:v1:bluemix:public:cloud-object-storage:global:[redacted]-fab3-40b1-98cb-682618778565:bucket:vprotect-bu
cket
```

Endpoints

Endpoints are used hand in hand with your credentials (i.e. keys, CRN, bucket name) to tell your service where to look for this bucket. Depending on where your service or applications is located you will want to use one of the below endpoint types.

Private ⓘ

Use private endpoints to point applications or services that are hosted in the IBM cloud (excluding Cloud Foundry services).

```
s3.private.eu-de.cloud-object-storage.appdomain.cloud
```

Public ⓘ

Use public endpoints to point applications or services that are hosted outside of the IBM cloud or for Cloud Foundry applications hosted in the IBM cloud.

```
s3.eu-de.cloud-object-storage.appdomain.cloud
```

Direct ⓘ

Use direct endpoints to connect from a VPC to Cloud Object Storage.

```
s3.direct.eu-de.cloud-object-storage.appdomain.cloud
```

Associated key management services

This bucket is not encrypted with key management key. Learn more

Activity Tracker

No Activity Tracker service instance is currently associated with this bucket.

Monitoring

No IBM Cloud Monitoring with Sysdig service instance is currently associated with this bucket.

Archive rule

Archive lifecycle rules allow users to change their objects from their default bucket storage class to the archive storage class after a certain period. Archive is our lowest cost tier and is best for items that are not accessed often.

Expiration rule

Object Expiration rules allow you to manage storage costs by scheduling deletion of objects that are no longer needed after a specified amount of time. Learn more

Rule name ⓘ ⇅	Prefix filter(optional) ⓘ ⇅	Expiration days/date ⓘ ⇅	State ⇅

We are almost done here, now we need to create API access and a secret key. Go to "Service credentials" on the left side menu then create new credentials using the blue button on the right.

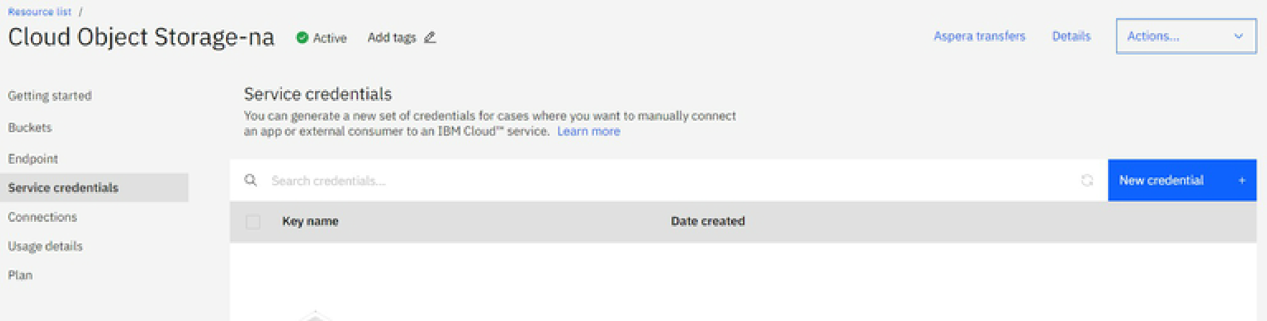

Service credentials

You can generate a new set of credentials for cases where you want to manually connect an app or external consumer to an IBM Cloud™ service. Learn more

There are two important options on this screen. You must select the appropriate role (for Data Protector for Cloud Workloads it is the "Writer" role) and select the option "Include HMAC credential".

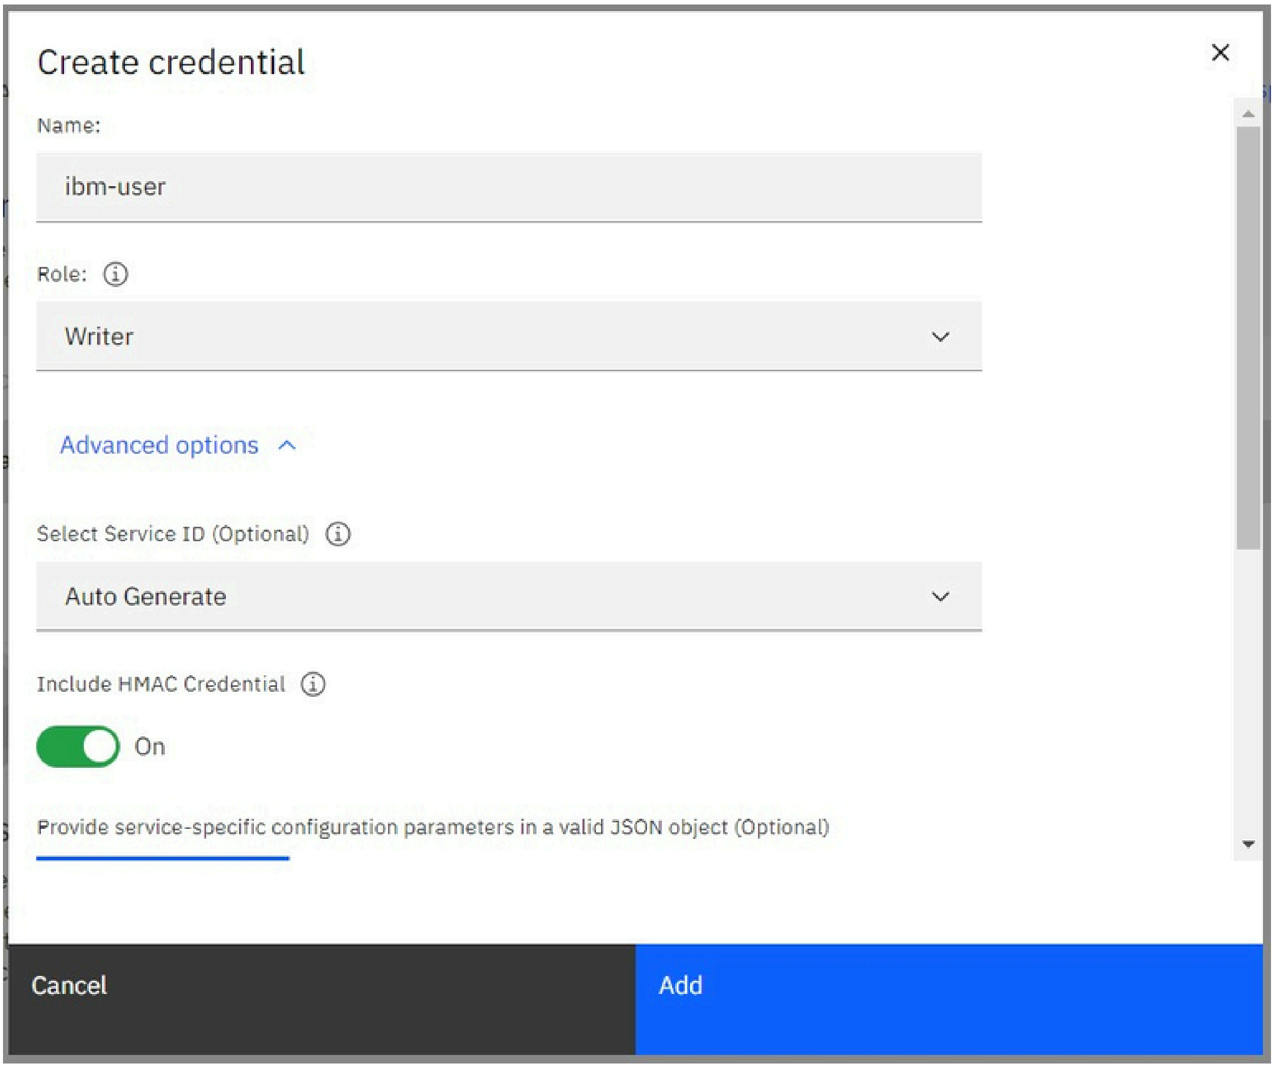

Create credential

Name: ibm-user

Role: Writer

Select Service ID (Optional): Auto Generate

Include HMAC Credential: On

Provide service-specific configuration parameters in a valid JSON object (Optional)

Now expand the detailed information about the created credentials by clicking on the arrow next to the name. What we need is "access_key_id" and "secret_access_key".

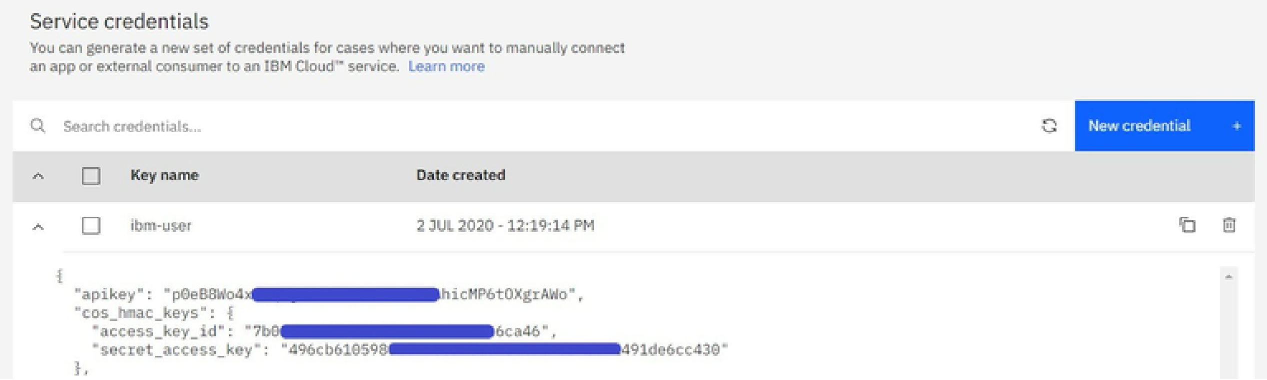

Service credentials

You can generate a new set of credentials for cases where you want to manually connect an app or external consumer to an IBM Cloud™ service. Learn more

Key name	Date created
ibm-user	2 JUL 2020 - 12:19:14 PM

```
{
  "apikey": "p0eB8Wo4x[redacted]hicMP6tOXgrAWo",
  "cos_hmac_keys": {
    "access_key_id": "7b0[redacted]6ca46",
    "secret_access_key": "496cb610598[redacted]491de6cc430"
  },
```

Now we can log in to the Data Protector for Cloud Workloads Dashboard and create a backup destination. Go to the backup destination tab on the left side menu and then choose "Amazon S3 / S3-compatible".

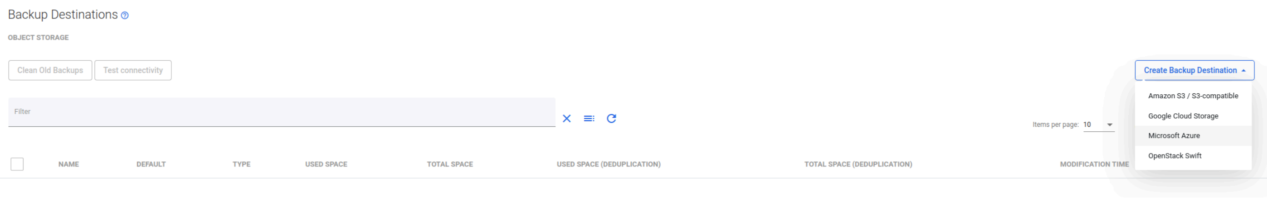

Backup Destinations

OBJECT STORAGE

NAME	DEFAULT	TYPE	USED SPACE	TOTAL SPACE	USED SPACE (DEDUPLICATION)	TOTAL SPACE (DEDUPLICATION)	MODIFICATION TIME

As IBM cloud storage is compatible with Amazon-S3, many settings will be very similar. However, remember to enter the API URL (remember about "https://" at the beginning), select the "Record backup time after store" option, and enter the region.

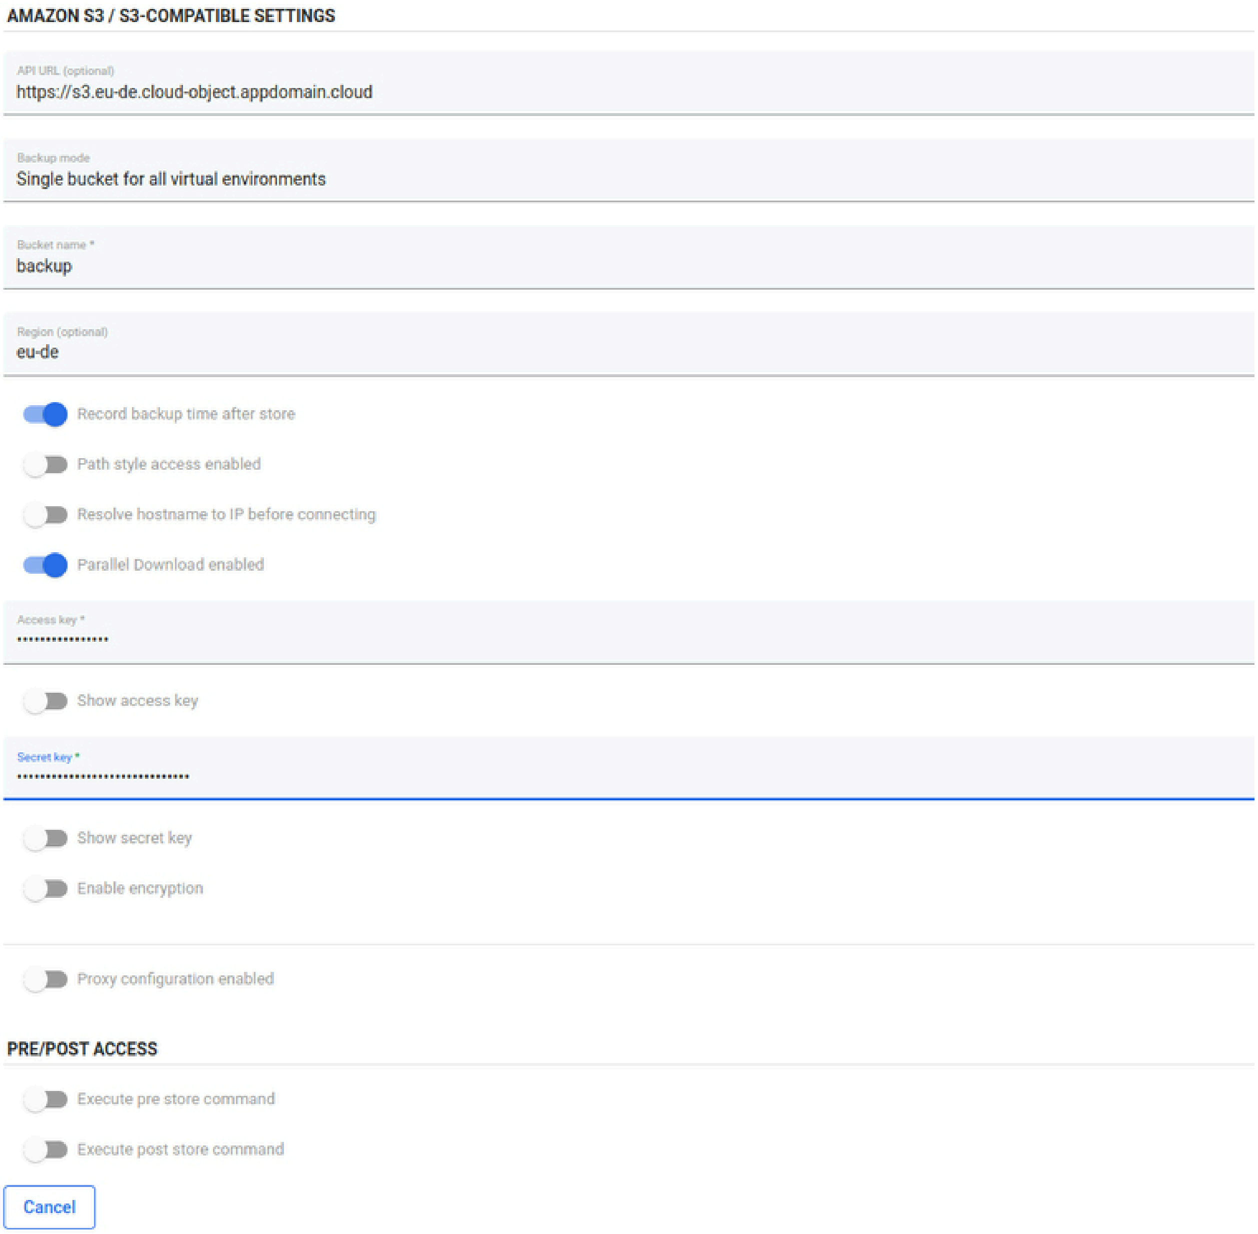

AMAZON S3 / S3-COMPATIBLE SETTINGS

API URL (optional)
https://s3.eu-de.cloud-object.appdomain.cloud

Backup mode
Single bucket for all virtual environments

Bucket name *
backup

Region (optional)
eu-de

● Record backup time after store

○ Path style access enabled

○ Resolve hostname to IP before connecting

● Parallel Download enabled

Access key *
••••••••••••••••

○ Show access key

Secret key *
••••••••••••••••••••••••••••••

○ Show secret key

○ Enable encryption

○ Proxy configuration enabled

PRE/POST ACCESS

○ Execute pre store command

○ Execute post store command

Microsoft Azure Blob Storage

Data Protector for Cloud Workloads supports integration with MS Azure Blob Storage. An Azure storage account contains all of your Azure Storage data objects: blobs, files, queues, tables, and disks. The storage account provides a unique namespace for your Azure Storage data that is accessible from anywhere in the world over HTTP or HTTPS. If you don't already know Azure Blob storage, read this great documentation https://docs.microsoft.com/en-gb/azure/storage/blobs/.

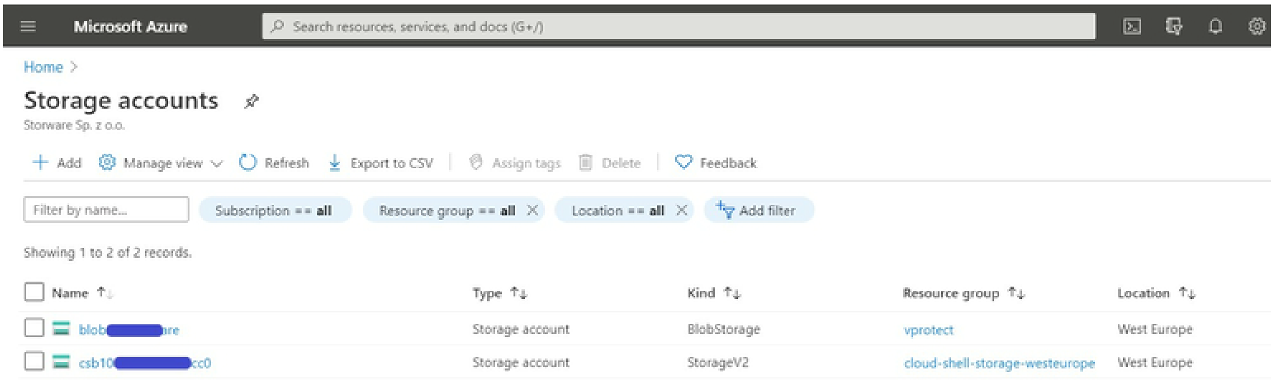

Storage accounts

Storware Sp. z o.o.

Showing 1 to 2 of 2 records.

Name	Type	Kind	Resource group	Location
blob[redacted]re	Storage account	BlobStorage	vprotect	West Europe
csb10[redacted]cc0	Storage account	StorageV2	cloud-shell-storage-westeurope	West Europe

To configure Azure as a backup destination for Data Protector for Cloud Workloads, we just need:

- The storage account name
- One of the account keys

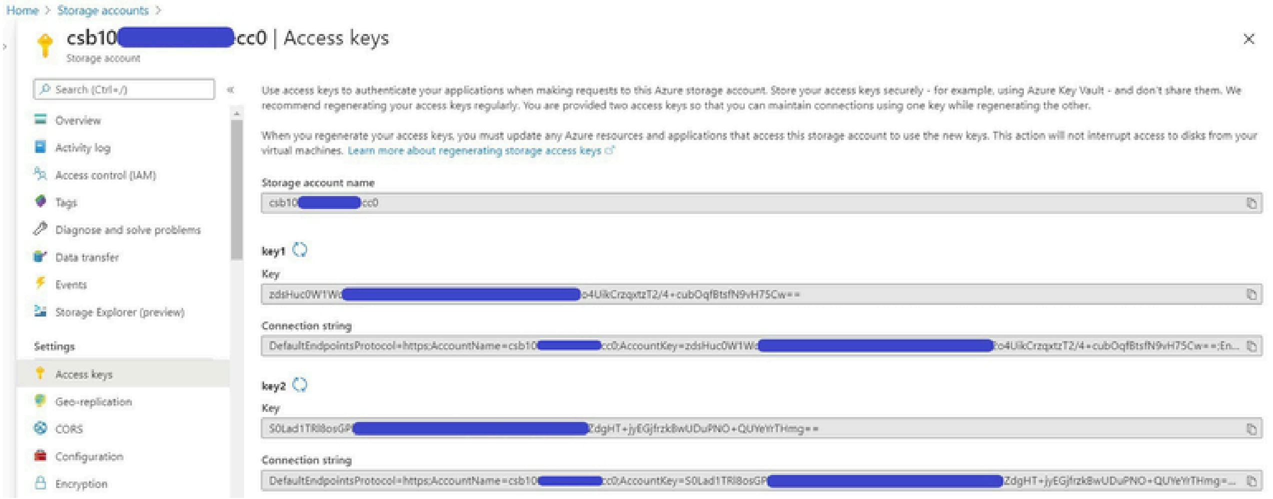

csb10[redacted]cc0 | Access keys

Storage account

Use access keys to authenticate your applications when making requests to this Azure storage account. Store your access keys securely - for example, using Azure Key Vault - and don't share them. We recommend regenerating your access keys regularly. You are provided two access keys so that you can maintain connections using one key while regenerating the other.

When you regenerate your access keys, you must update any Azure resources and applications that access this storage account to use the new keys. This action will not interrupt access to disks from your virtual machines. Learn more about regenerating storage access keys

Storage account name: csb10[redacted]cc0

key1

Key: zdsHuc0W1Wd[redacted]o4UikCrzqxtzT2/4+cubOqfBtsfN9vH75Cw==

Connection string: DefaultEndpointsProtocol=https;AccountName=csb10[redacted]cc0;AccountKey=zdsHuc0W1Wd[redacted]o4UikCrzqxtzT2/4+cubOqfBtsfN9vH75Cw==;En...

key2

Key: S0Lad1TRl8osGP[redacted]ZdgHT+jyEGjfrzkBwUDuPNO+QUYeYrTHmg==

Connection string: DefaultEndpointsProtocol=https;AccountName=csb10[redacted]cc0;AccountKey=S0Lad1TRl8osGP[redacted]ZdgHT+jyEGjfrzkBwUDuPNO+QUYeYrTHmg=...

Now you can go to the backup destinations tab in Data Protector for Cloud Workloads and create a new Microsoft Azure backup destination.

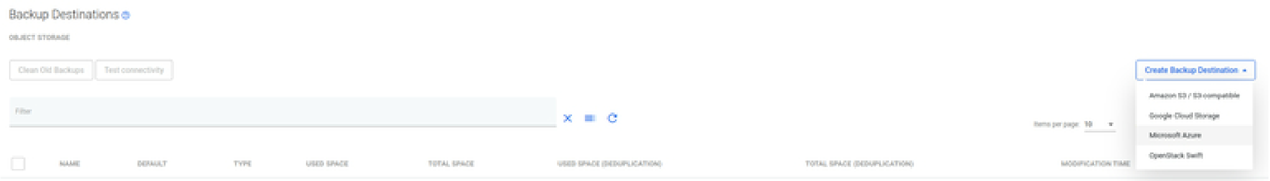

You just need to provide an account name, bucket name and key.

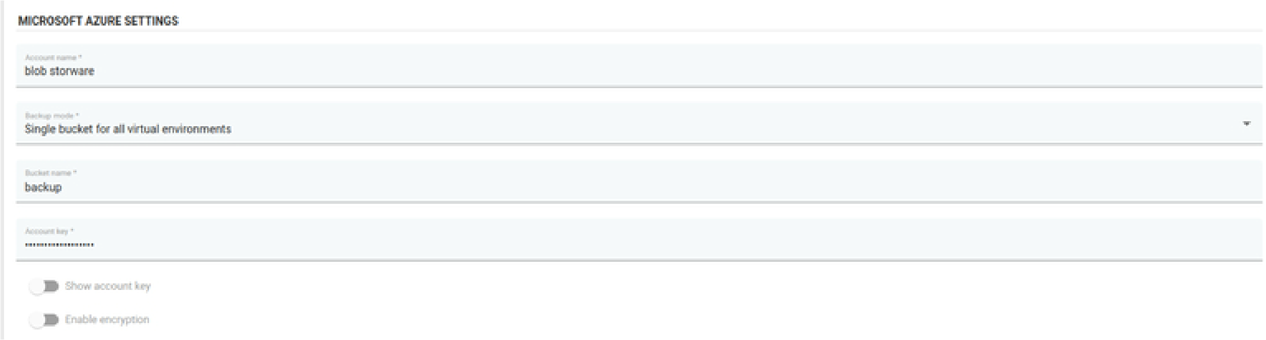

MICROSOFT AZURE SETTINGS

Account name *
blob storware

Backup mode *
Single bucket for all virtual environments

Bucket name *
backup

Account key *
••••••••••••••••••••

Show account key

Enable encryption

And that's all. As you see, in a few minutes you can integrate Data Protector for Cloud Workloads with Azure Blob storage to securely store your backups

Nutanix Objects

Nutanix Objects is an S3-compatible backup provider. Configuration as the backup destination is similar to AWS S3.

Example

In the Data Protector for Cloud Workloads system, go to the **Backup Destinations -> Object Storage** tab, then press the **Create Backup Destination** button and select the **Amazon S3 / S3-compatible** option.

In this step, complete the name, retention, add: API URL, Access key, and Secret key, indicate the name of the bucket to be used.

Then go to the **AMAZON S3/S3-COMPATIBLE SETTINGS** the segment in which you should **deselect** the **Parallel Download enabled** option for Nutanix Objects.

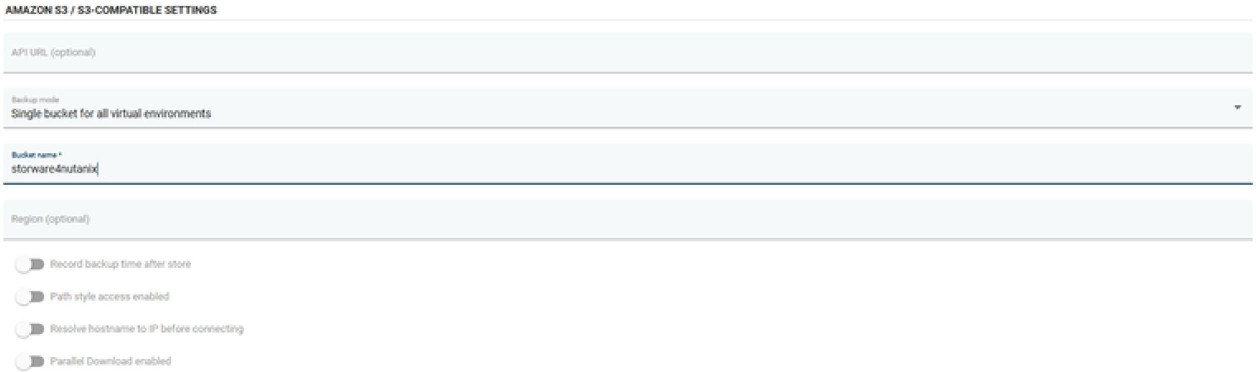

AMAZON S3 / S3-COMPATIBLE SETTINGS

API URL (optional)

Backup mode
Single bucket for all virtual environments

Bucket name *
storware4nutanix

Region (optional)

Record backup time after store

Path style access enabled

Resolve hostname to IP before connecting

Parallel Download enabled

> When using Nutanix Objects version 3.5, the region "us-east-1" may be required.

After entering the settings, press the **Save** button to be able to use Nutanix Objects as Backup Destination.

OpenStack SWIFT

Data Protector for Cloud Workloads supports integration with OpenStack SWIFT.

Example

In the Data Protector for Cloud Workloads system, go to the **Backup Destinations** -> **Object Storage** tab, then press the **Create Backup Destination** button and select the **OpenStack Swift** option.

Enter the name of the new backup destination, assign it to **Node Configuration** and set up the retention.

Next, provide settings specific to **OpenStack Swift**:

- Authentication URL - URL pointing to authentication service, it should be similar to the following

  ```
  https://SWIFT_HOST:5000/v3/auth/tokens
  ```

- User name - domain formatted username used by Data Protector for Cloud Workloads to log into OpenStack Swift
- Authentication method - BASIC / TEMPAUTH / KEYSTONE / KEYSTONE_V3
  - in the case of KEYSTONE_V3 authentication method, you also need to enter **Authentication method scope**, **Domain** and **Project**
- Name of Swift service intended to be used
- Number of thread used (Swift connector supports multithreading)
- Endpoint interface type - type of interface used by connector (PUBLIC / INTERNAL / ADMIN)

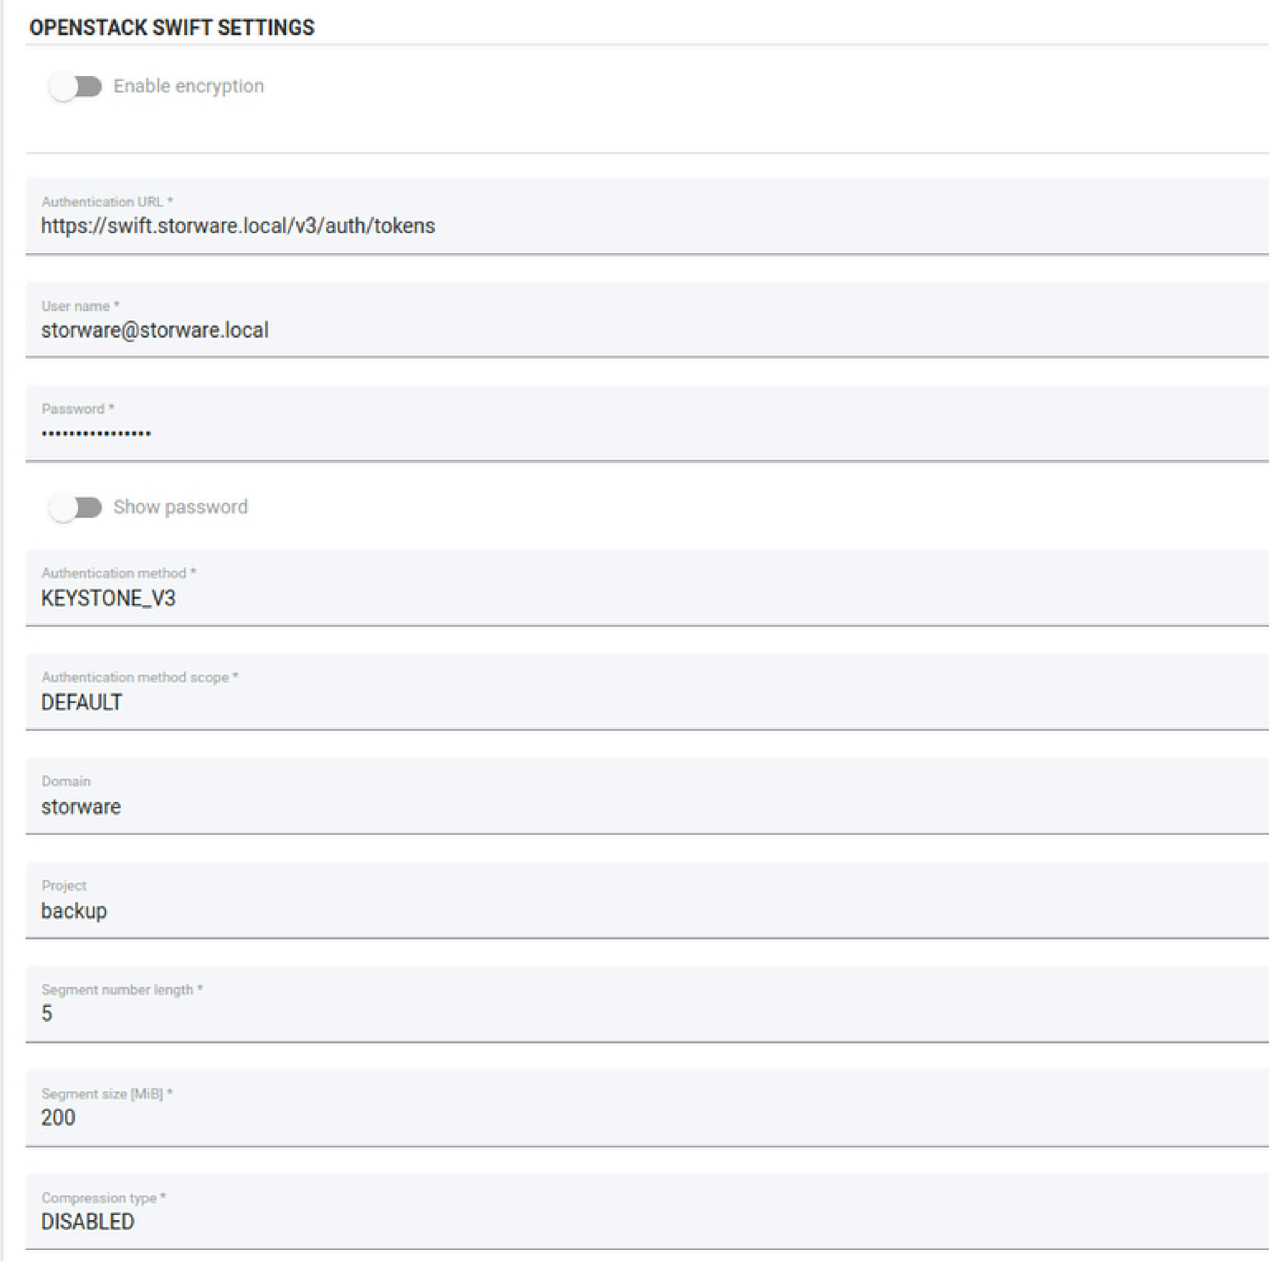

OPENSTACK SWIFT SETTINGS

Enable encryption

Authentication URL *
https://swift.storware.local/v3/auth/tokens

User name *
storware@storware.local

Password *
••••••••••••••••

Show password

Authentication method *
KEYSTONE_V3

Authentication method scope *
DEFAULT

Domain
storware

Project
backup

Segment number length *
5

Segment size [MiB] *
200

Compression type *
DISABLED

Oracle Cloud Infrastructure Object Storage

Overview

The Oracle Cloud Infrastructure Object Storage service is an internet-scale, high-performance storage platform that offers reliable and cost-efficient data durability. The Object Storage service can store an unlimited amount of unstructured data of any content type, including analytic data and rich content.

Example

Log in to the Oracle cloud dashboard, expand the left side menu and go to the Object Storage tab.

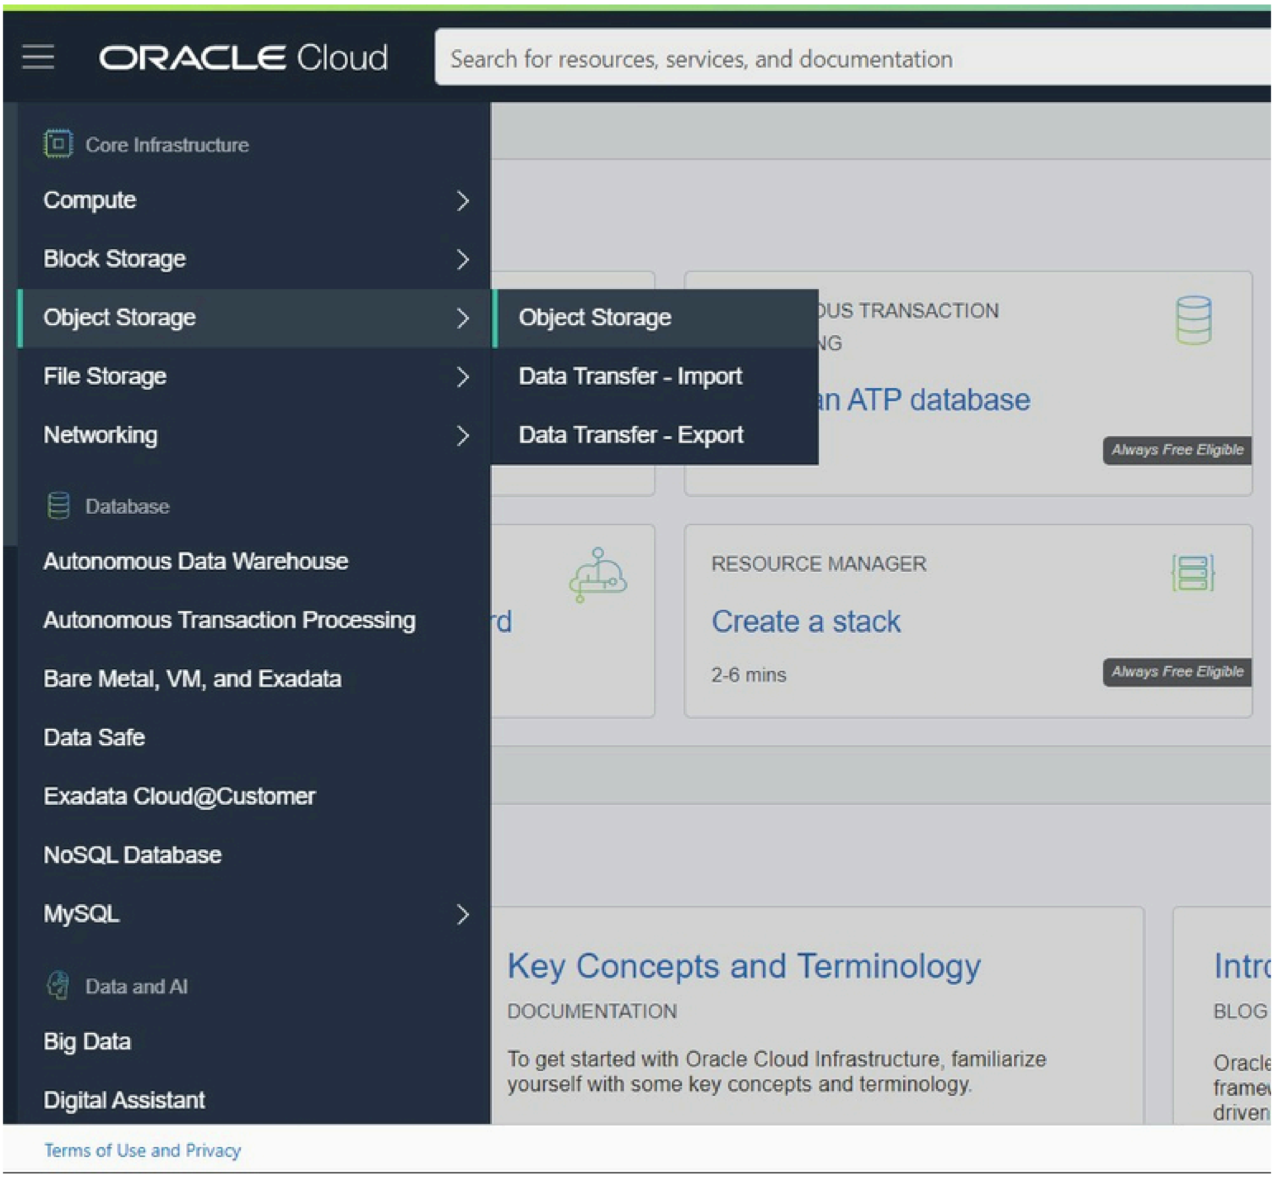

To get started with Oracle Cloud Infrastructure, familiarize yourself with some key concepts and terminology.

Oracle frame driven

Now let's create a new bucket.

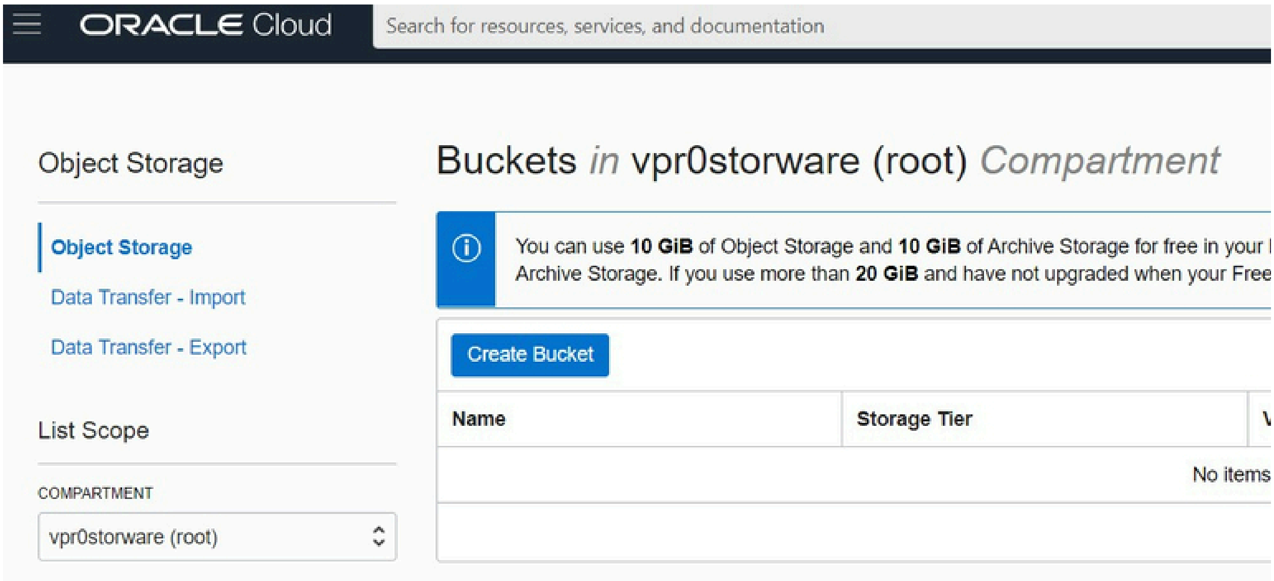

Buckets *in* vpr0storware (root) *Compartment*

You can use **10 GiB** of Object Storage and **10 GiB** of Archive Storage for free in your Archive Storage. If you use more than **20 GiB** and have not upgraded when your Free

Name	Storage Tier	V
		No items

We do not require specific bucket settings for Data Protector for Cloud Workloads. The bucket name will be needed when we want to create a backup destination in Data Protector for Cloud Workloads.

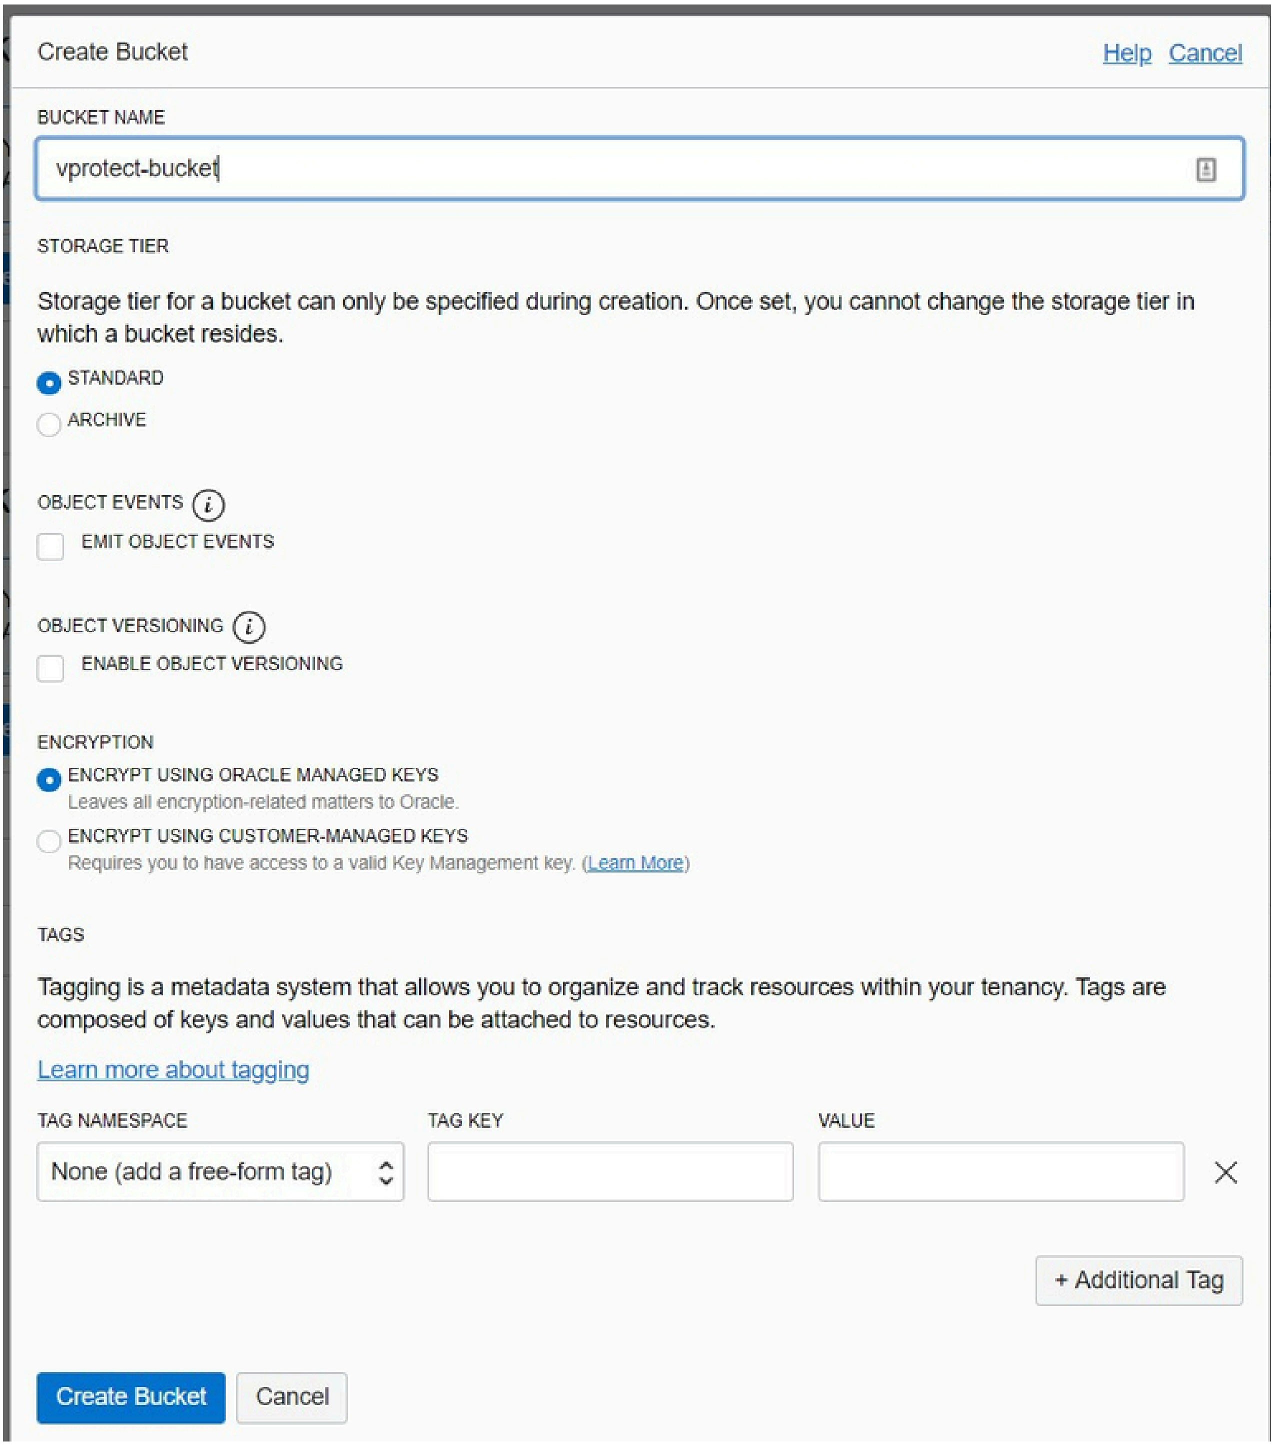

Create Bucket

BUCKET NAME

vprotect-bucket

STORAGE TIER

Storage tier for a bucket can only be specified during creation. Once set, you cannot change the storage tier in which a bucket resides.

● STANDARD

○ ARCHIVE

OBJECT EVENTS

□ EMIT OBJECT EVENTS

OBJECT VERSIONING

□ ENABLE OBJECT VERSIONING

ENCRYPTION

● ENCRYPT USING ORACLE MANAGED KEYS
Leaves all encryption-related matters to Oracle.

○ ENCRYPT USING CUSTOMER-MANAGED KEYS
Requires you to have access to a valid Key Management key. (Learn More)

TAGS

Tagging is a metadata system that allows you to organize and track resources within your tenancy. Tags are composed of keys and values that can be attached to resources.

Learn more about tagging

TAG NAMESPACE	TAG KEY	VALUE
None (add a free-form tag)		

After creating the bucket, you'll see a list of buckets. Click on the name to view the details of the object. Remember the "namespace", we also need it when creating a backup destination.

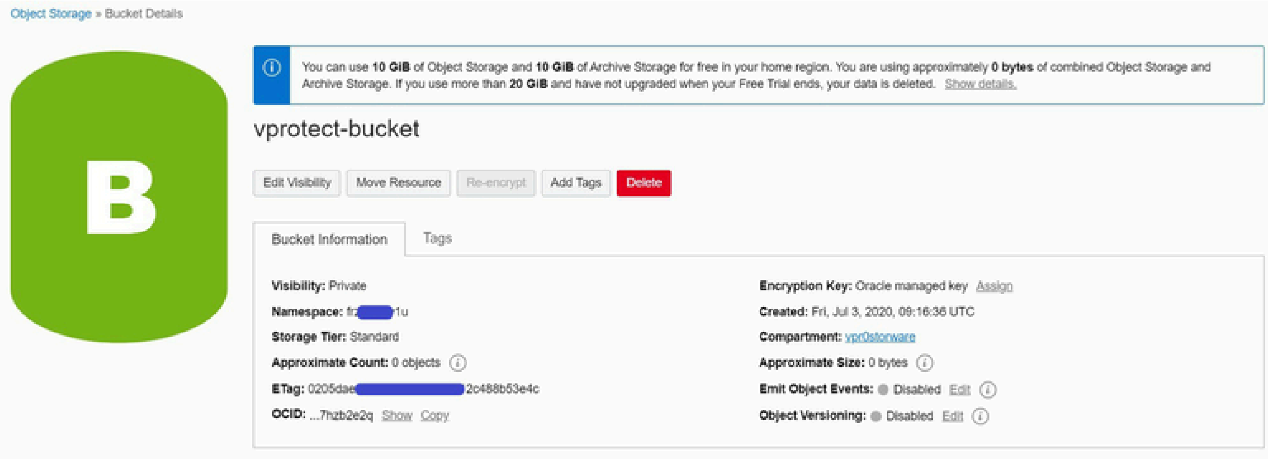

You can use **10 GiB** of Object Storage and **10 GiB** of Archive Storage for free in your home region. You are using approximately **0 bytes** of combined Object Storage and Archive Storage. If you use more than **20 GiB** and have not upgraded when your Free Trial ends, your data is deleted. Show details.

vprotect-bucket

Visibility: Private
Namespace: fr[redacted]1u
Storage Tier: Standard
Approximate Count: 0 objects
ETag: 0205dae[redacted]2c488b53e4c
OCID: ...7hzb2e2q Show Copy

Encryption Key: Oracle managed key Assign
Created: Fri, Jul 3, 2020, 09:16:36 UTC
Compartment: vpr0storware
Approximate Size: 0 bytes
Emit Object Events: Disabled Edit
Object Versioning: Disabled Edit

Now we need to create a user that we will use to authenticate our backup destination. Go to the Users tab under the Identity tab in the menu on the left.

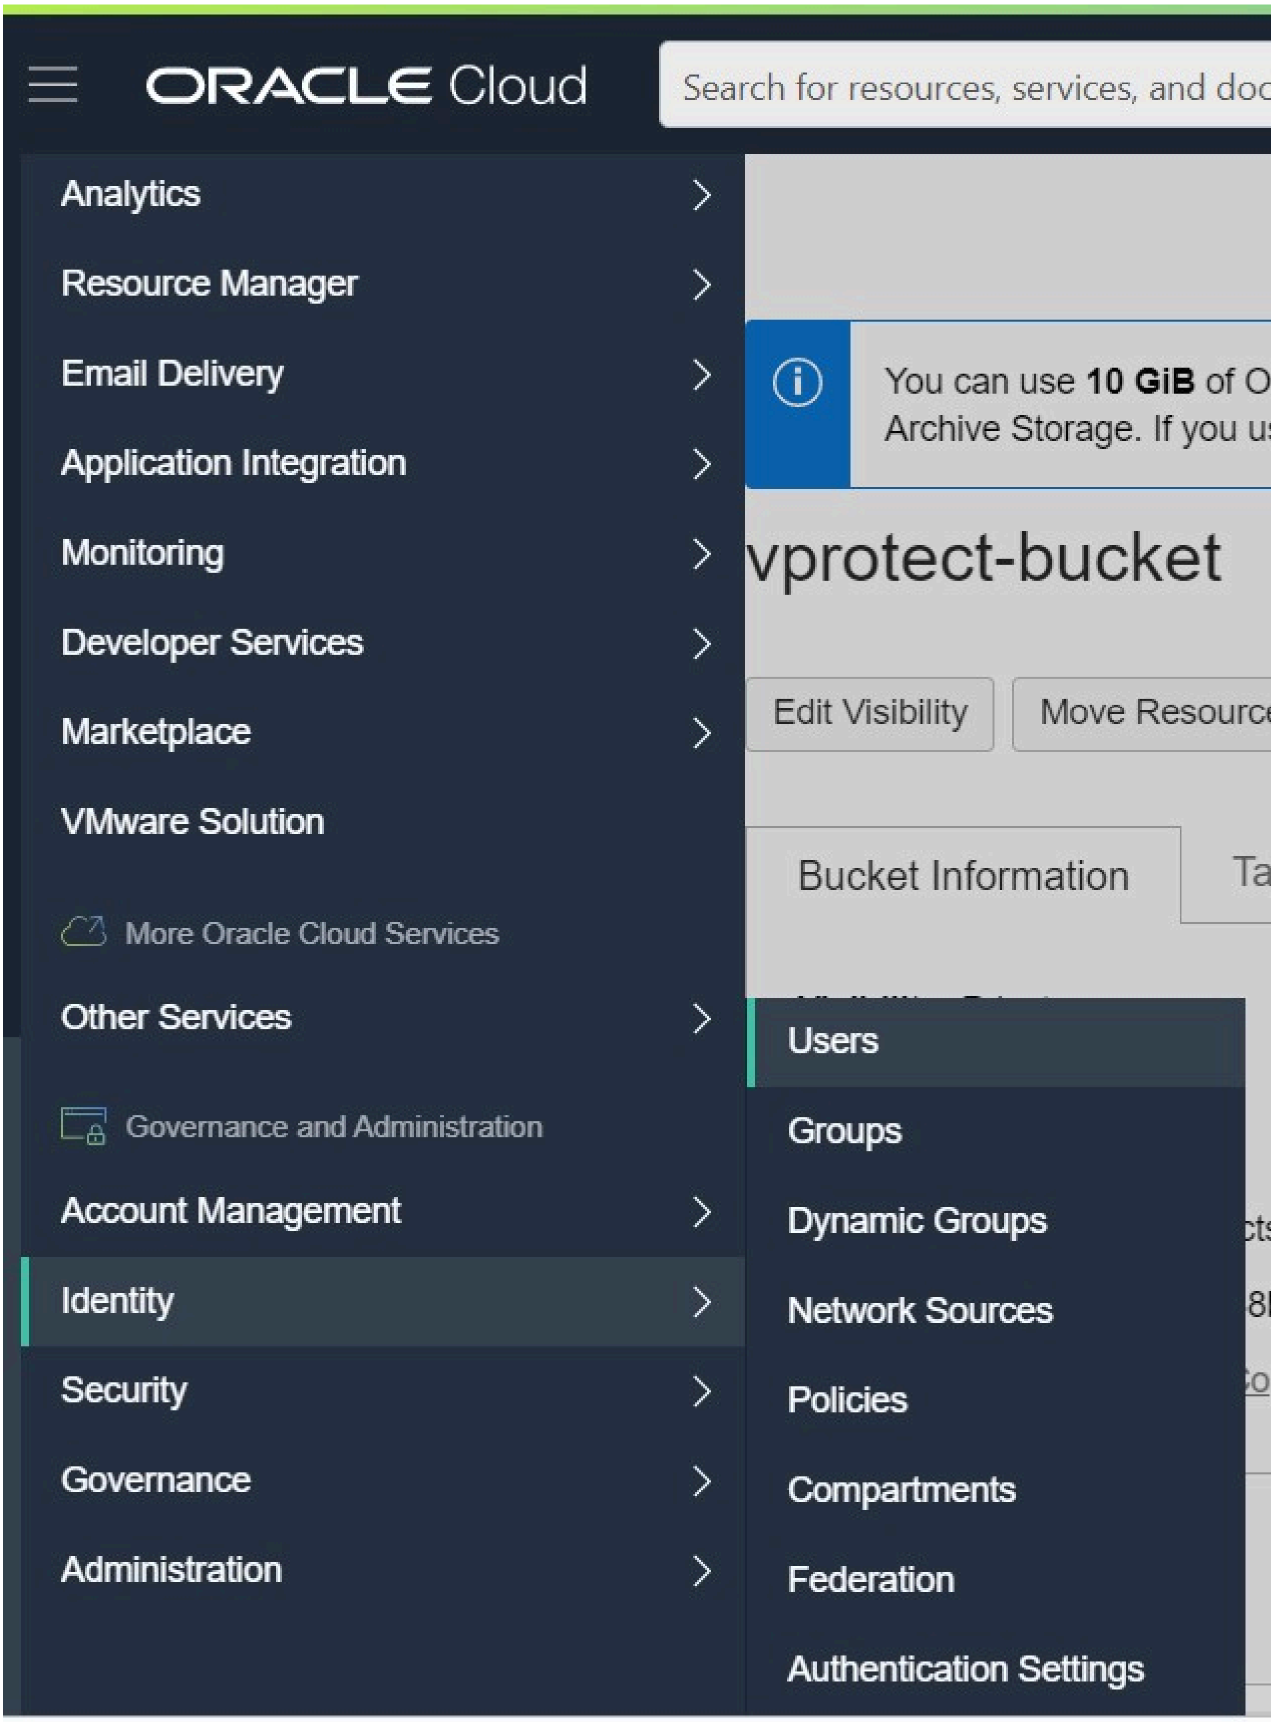

Now create a new user.

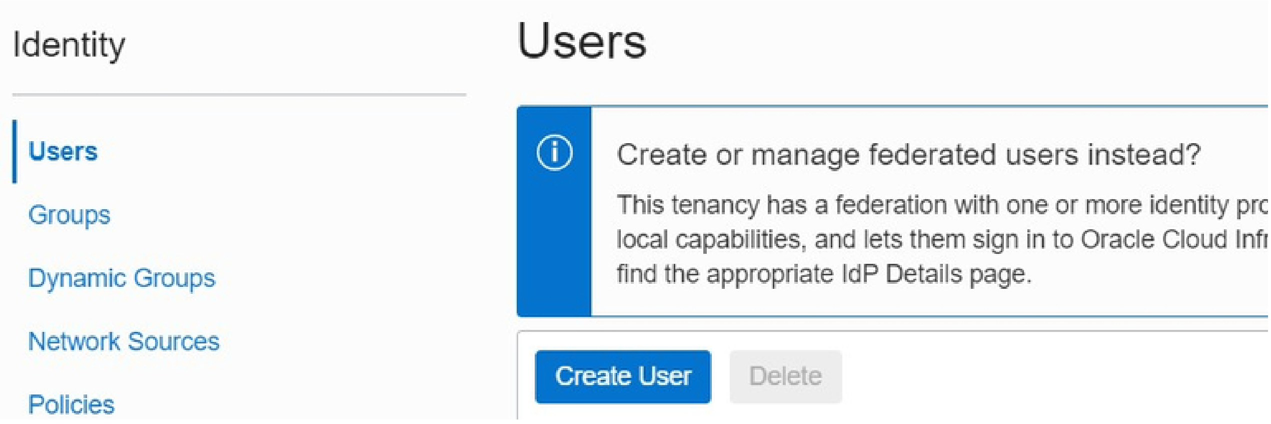

Users

Create or manage federated users instead?

This tenancy has a federation with one or more identity pro
local capabilities, and lets them sign in to Oracle Cloud Infr
find the appropriate IdP Details page.

Fill in the required fields.

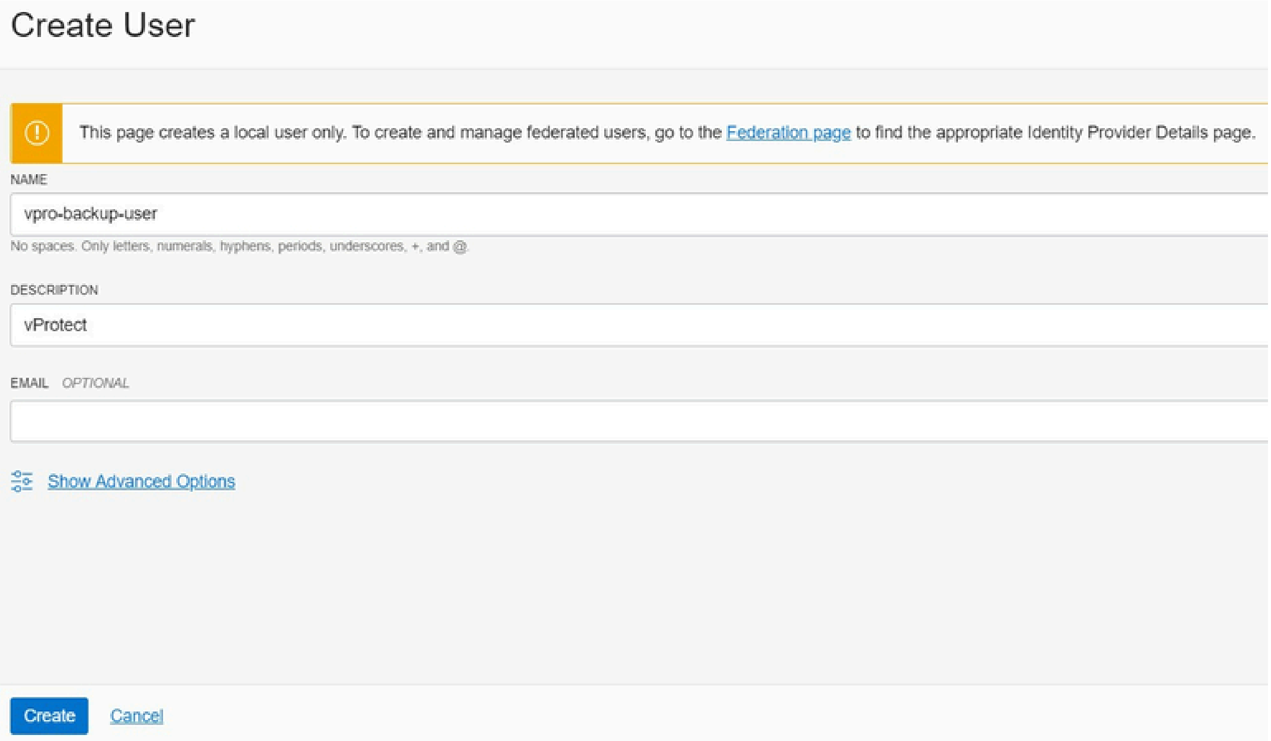

Create User

This page creates a local user only. To create and manage federated users, go to the Federation page to find the appropriate Identity Provider Details page.

NAME: vpro-backup-user

No spaces. Only letters, numerals, hyphens, periods, underscores, +, and @.

DESCRIPTION: vProtect

EMAIL OPTIONAL:

Then go to the Groups page, which you can also find under the identity tab in the left side menu.

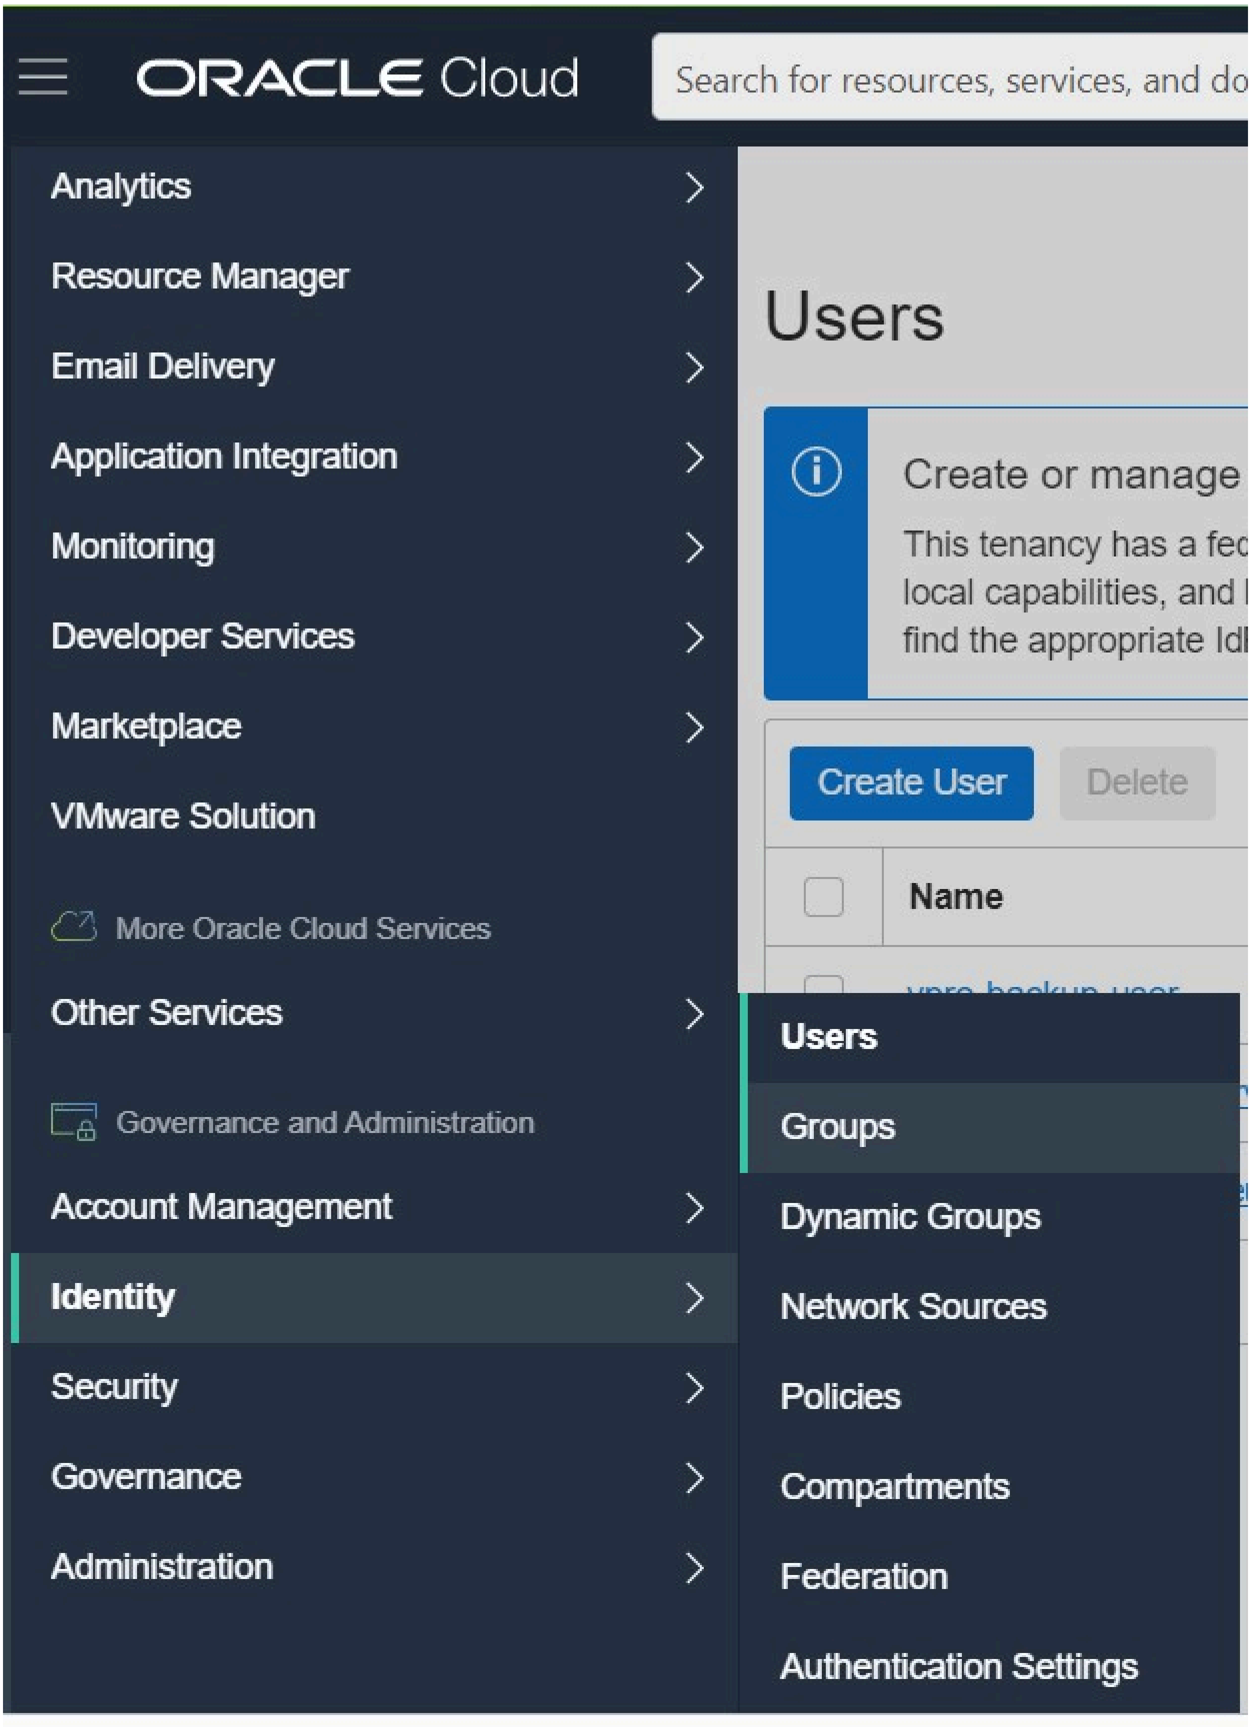

Now click on the existing group "Administrators".

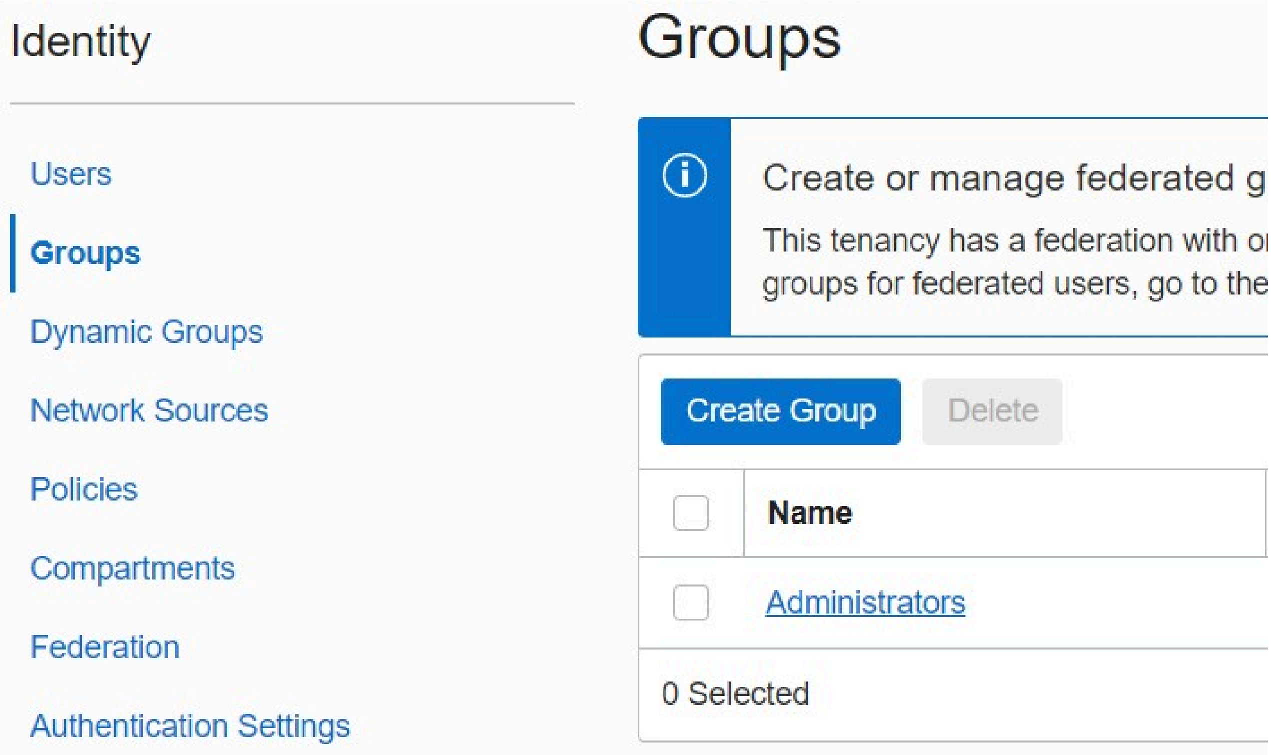

Create or manage federated g

This tenancy has a federation with o
groups for federated users, go to the

□	Name
□	Administrators

0 Selected

Now click on "Add User to Group" and choose the user you created previously.

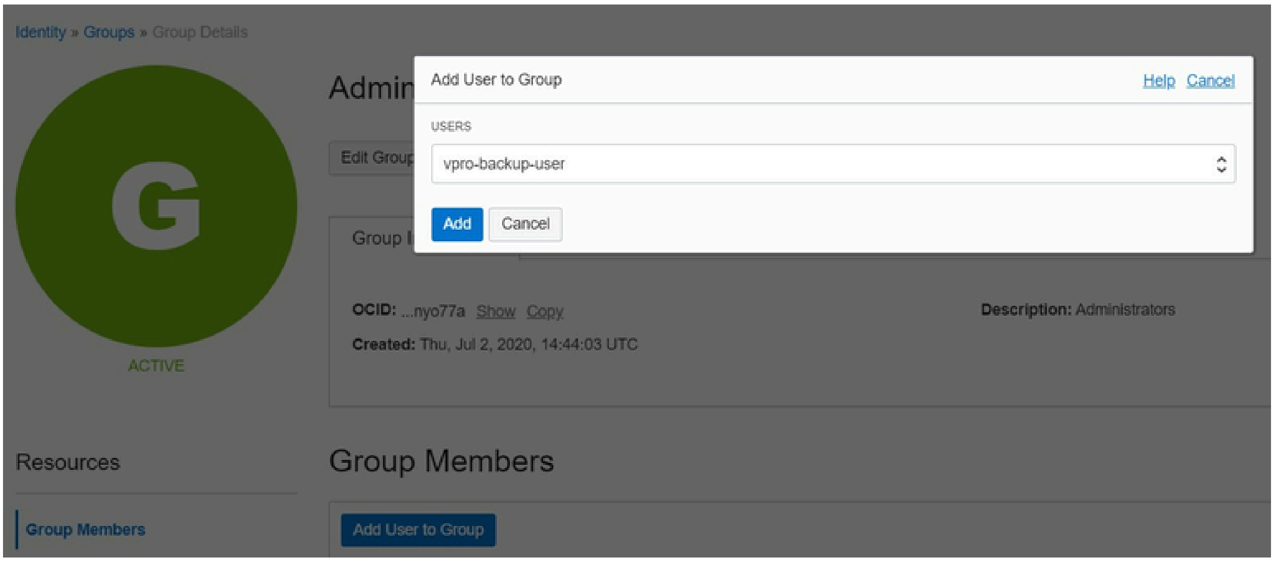

USERS: vpro-backup-user

OCID: ...nyo77a Show Copy

Created: Thu, Jul 2, 2020, 14:44:03 UTC

Description: Administrators

Go back to the Users page and go to the details page of our user.

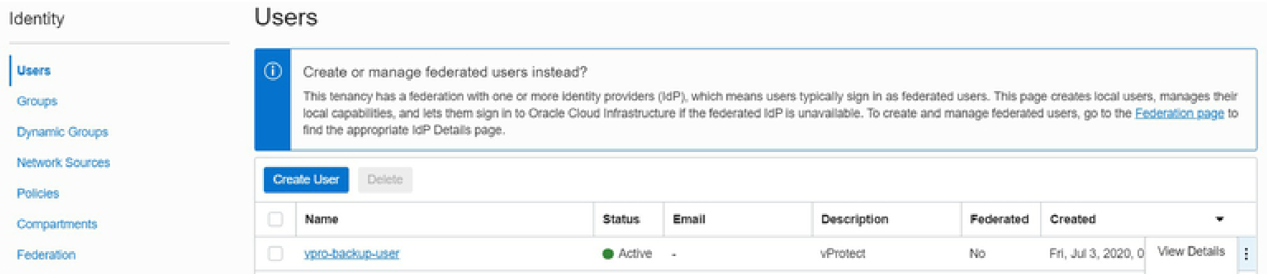

Create or manage federated users instead?

This tenancy has a federation with one or more identity providers (IdP), which means users typically sign in as federated users. This page creates local users, manages their local capabilities, and lets them sign in to Oracle Cloud Infrastructure if the federated IdP is unavailable. To create and manage federated users, go to the Federation page to find the appropriate IdP Details page.

□	Name	Status	Email	Description	Federated	Created
□	vpro-backup-user	● Active	-	vProtect	No	Fri, Jul 3, 2020, 0

Scroll down and open the "Customer Secret Keys" tab. Click on "Generate Secret Key".

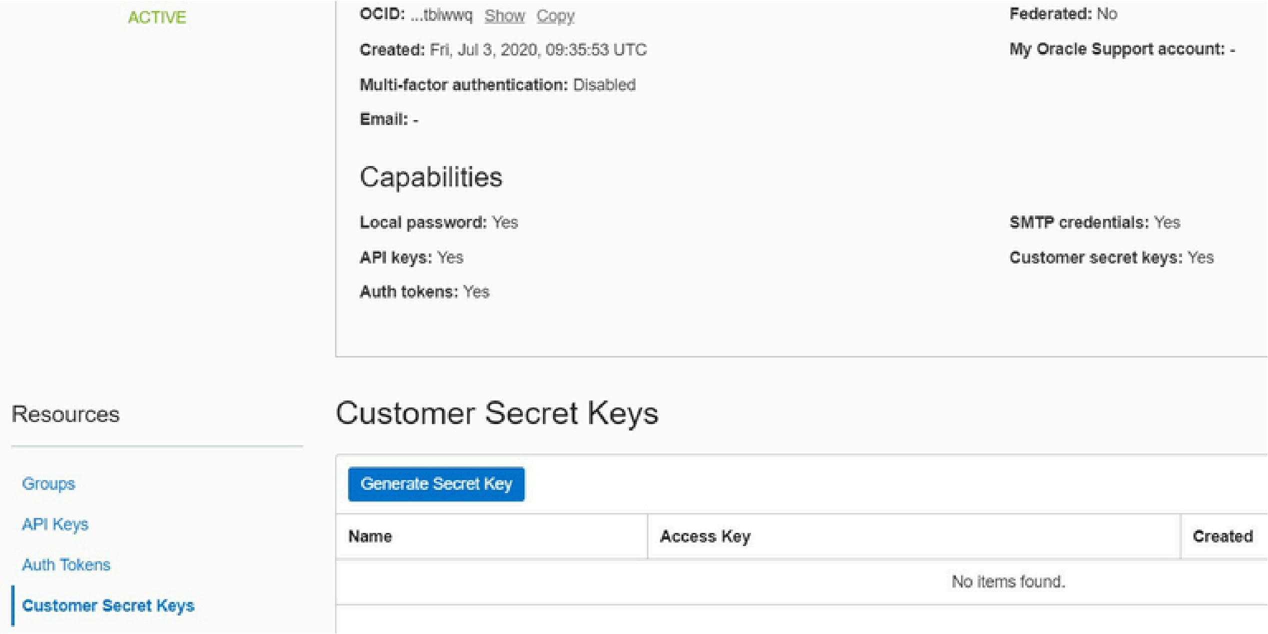

ACTIVE

OCID: ...tblwwq Show Copy

Created: Fri, Jul 3, 2020, 09:35:53 UTC

Multi-factor authentication: Disabled

Email: -

Federated: No

My Oracle Support account: -

Capabilities

Local password: Yes

API keys: Yes

Auth tokens: Yes

SMTP credentials: Yes

Customer secret keys: Yes

Customer Secret Keys

Name	Access Key	Created
No items found.		

Enter any name.

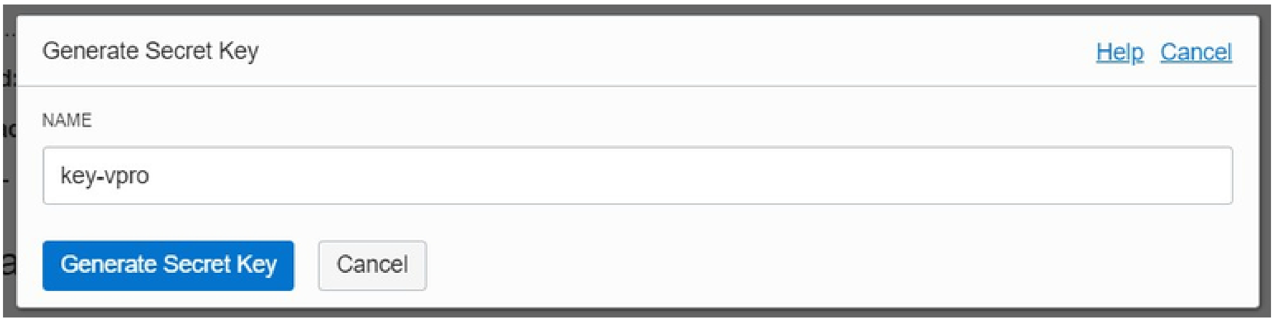

Generate Secret Key

NAME

key-vpro

As you see in the note below, copy and save the secret key because you can only do this now.

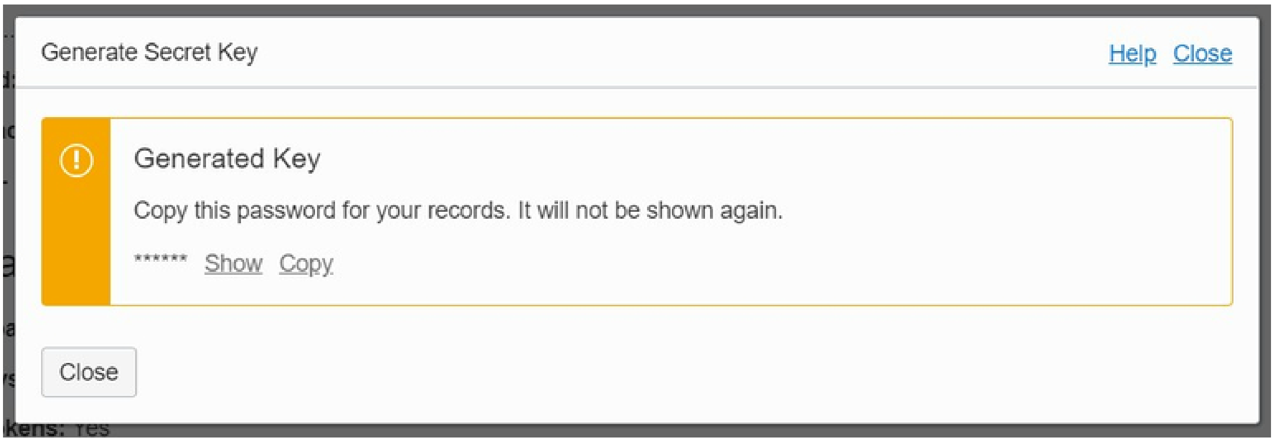

Generate Secret Key

Generated Key

Copy this password for your records. It will not be shown again.

****** Show Copy

After generating the secret key, you can view the access key, just move the mouse over it.

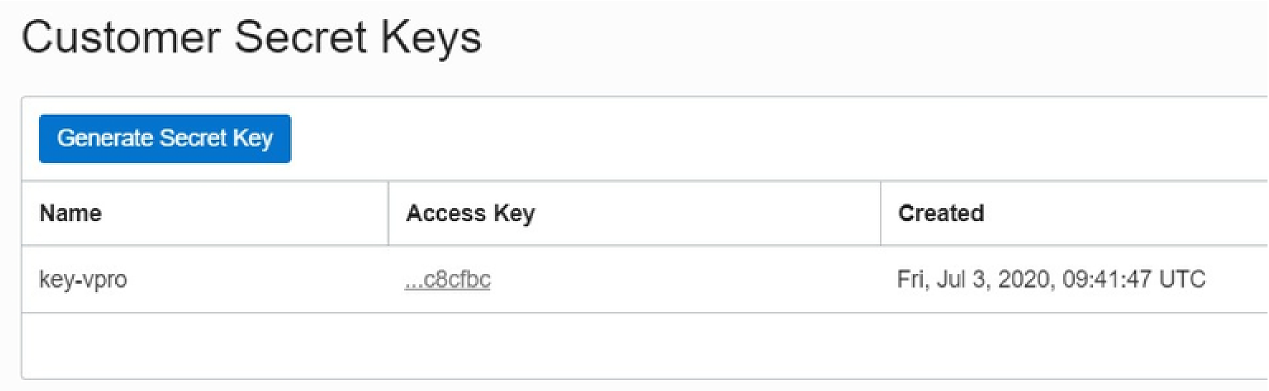

Customer Secret Keys

Generate Secret Key

Name	Access Key	Created
key-vpro	...c8cfbc	Fri, Jul 3, 2020, 09:41:47 UTC

Now we can go to the Data Protector for Cloud Workloads Dashboard. Open the "Backup Destination" tab from the left side menu, then the sub-tab "Object Storage" and choose "Amazon S3 / S3-compatible" as the new type of backup destination.

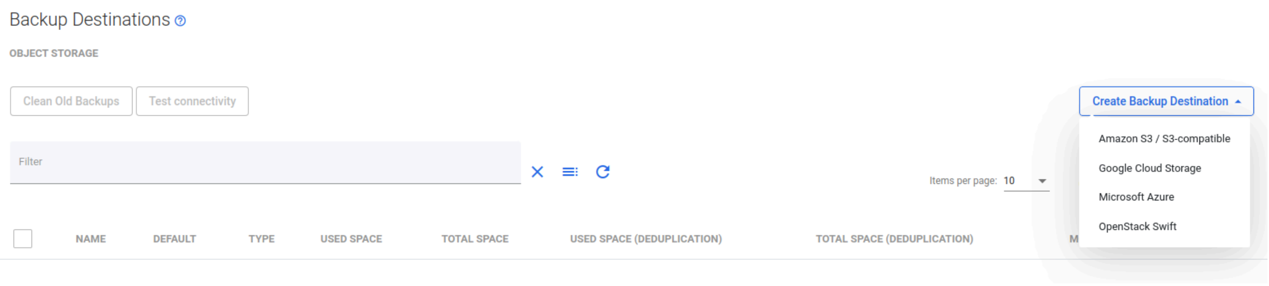

Backup Destinations

OBJECT STORAGE

Clean Old Backups | Test connectivity

Filter

Items per page: 10

□	NAME	DEFAULT	TYPE	USED SPACE	TOTAL SPACE	USED SPACE (DEDUPLICATION)	TOTAL SPACE (DEDUPLICATION)	M

First, let's focus on the "S3-Compatible" section.
To generate an API URL, you will need this site: https://docs.cloud.oracle.com/en-us/iaas/api/#/en/s3objectstorage/20160918/
As We mentioned earlier, you will need an object storage namespace (choose the API URL from the list according to your region).
Then provide your bucket name and region, and finally switch on "Record time after backup" and "Path style access enabled".
Configure the rest of the settings as desired.

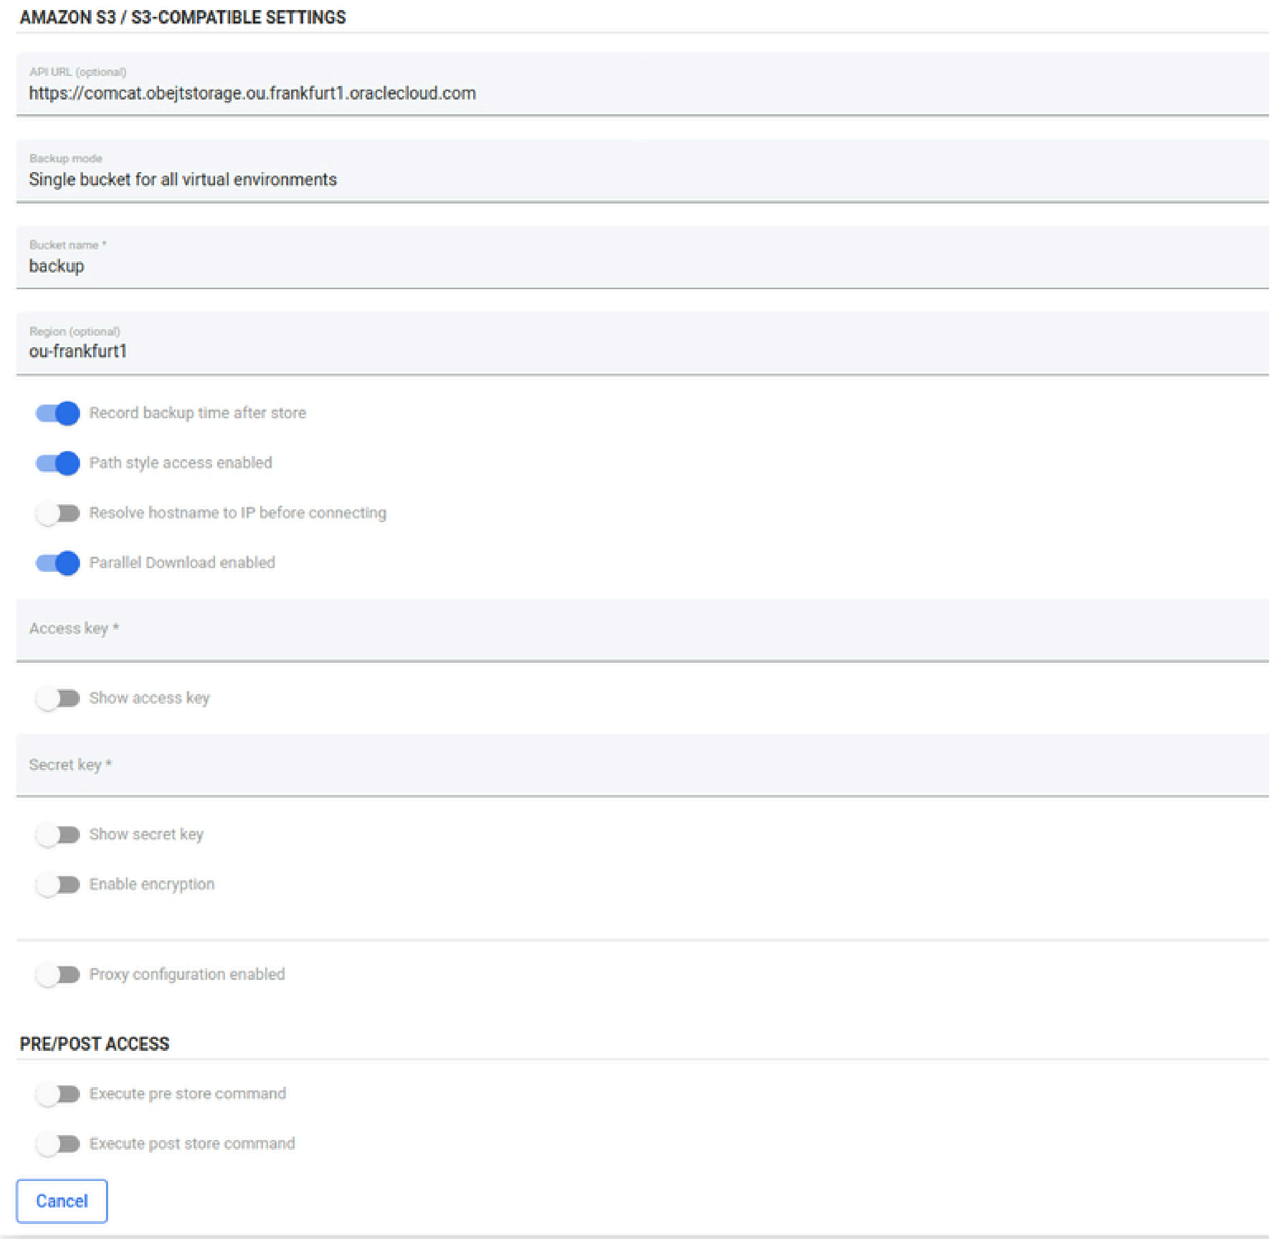

AMAZON S3 / S3-COMPATIBLE SETTINGS

API URL (optional)
https://comcat.obejtstorage.ou.frankfurt1.oraclecloud.com

Backup mode
Single bucket for all virtual environments

Bucket name *
backup

Region (optional)
ou-frankfurt1

● Record backup time after store

● Path style access enabled

○ Resolve hostname to IP before connecting

● Parallel Download enabled

Access key *

○ Show access key

Secret key *

○ Show secret key

○ Enable encryption

○ Proxy configuration enabled

PRE/POST ACCESS

○ Execute pre store command

○ Execute post store command

Scality RING

Overview

Scality Ring offers an object storage solution with a native and comprehensive S3 interface. Scality S3 Connector is the first AWS S3-compatible object storage for enterprise S3 applications with secure multi-tenancy and high performance.
AWS has achieved incredible traction with services such as S3 for a wide variety of cloud application and service provider businesses. However, for many service providers and enterprise corporations who require an on-premises deployment model in order to maintain control over sensitive data, for performance optimization, or for reasons of security or compliance – Scality's new S3 Connector for the RING provides an optimal solution. The S3 Connector offers a solution that is application-compatible with AWS S3 at both the data API level and also with the rapidly evolving AWS multi-tenancy model termed IAM (Identity and Access Management).

Example

In this example, we will show you how to use the Scality S3 connector to create the backup destination for Data Protector for Cloud Workloads.
It assumes that the S3 connector is installed and configured

We will start by creating a user. Launch the S3 connector user interface.

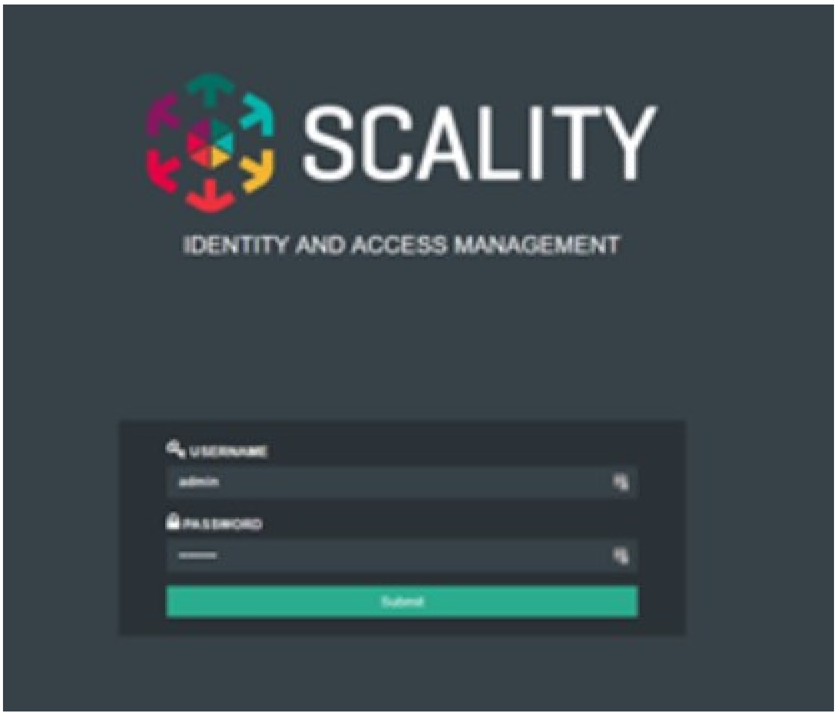

USERNAME: admin

PASSWORD: ••••••

Log in as an account user using the password set in *Setting an account Password* from the S3 console GUI.
Select the user to open the user management window.
Click Add user to open the add user window.
Enter the user name and make sure to check the box for "FullAccessGroup".

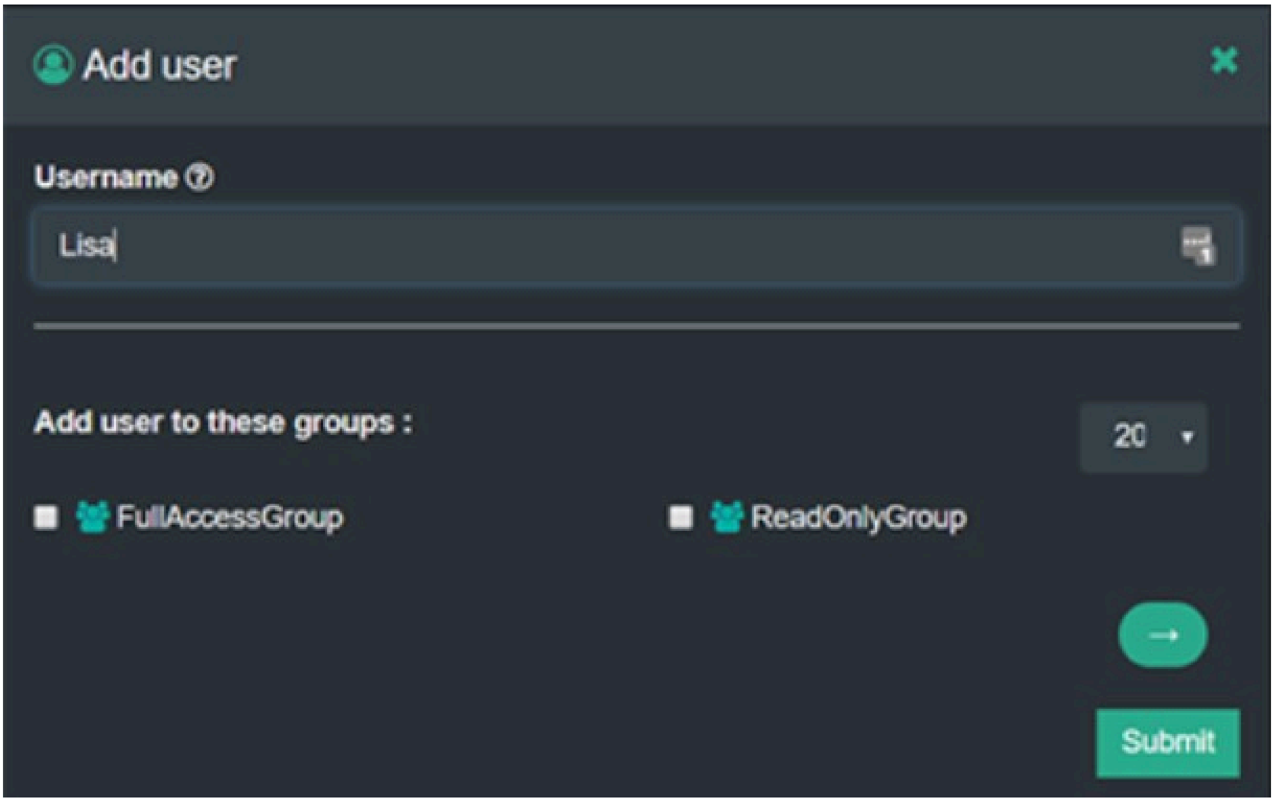

Username: Lisa

Add user to these groups :

□ FullAccessGroup □ ReadOnlyGroup

The user management panel displays the user name and the Amazon Resource Name (ARN).
Now we will generate the access and secret keys for the user.
Click on the key icon in the Actions column of the user row.

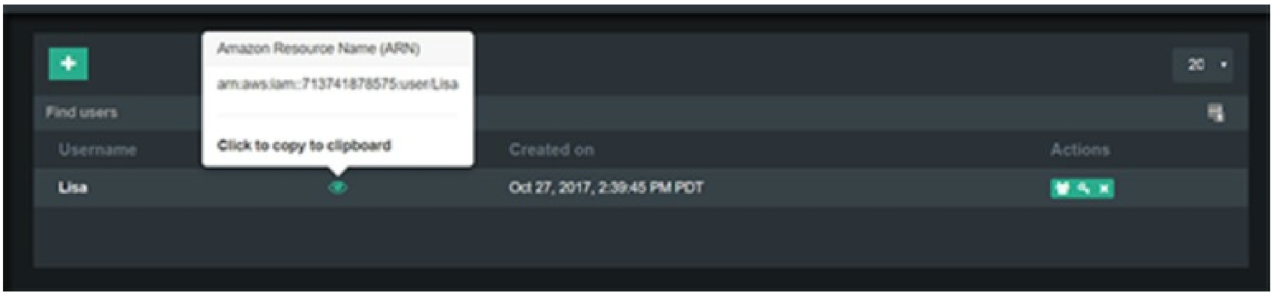

Amazon Resource Name (ARN)

arn:aws:iam::713741878575:user/Lisa

Click to copy to clipboard

Username	Created on	Actions
Lisa	Oct 27, 2017, 2:39:45 PM PDT	

Click on Generate a new key.

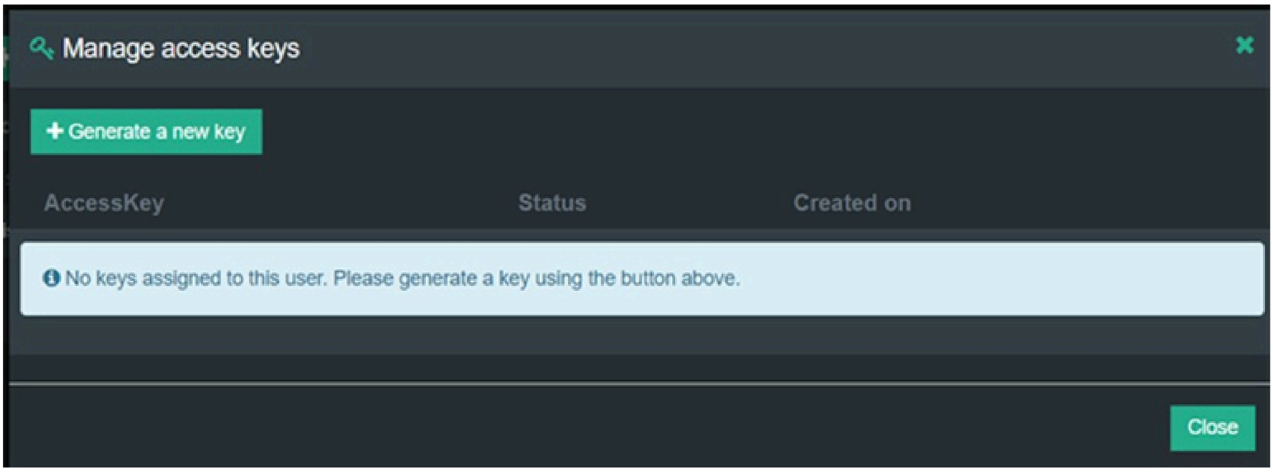

AccessKey	Status	Created on

No keys assigned to this user. Please generate a key using the button above.

Click on Proceed to generate the user's AccessKey and SecretAccessKey.

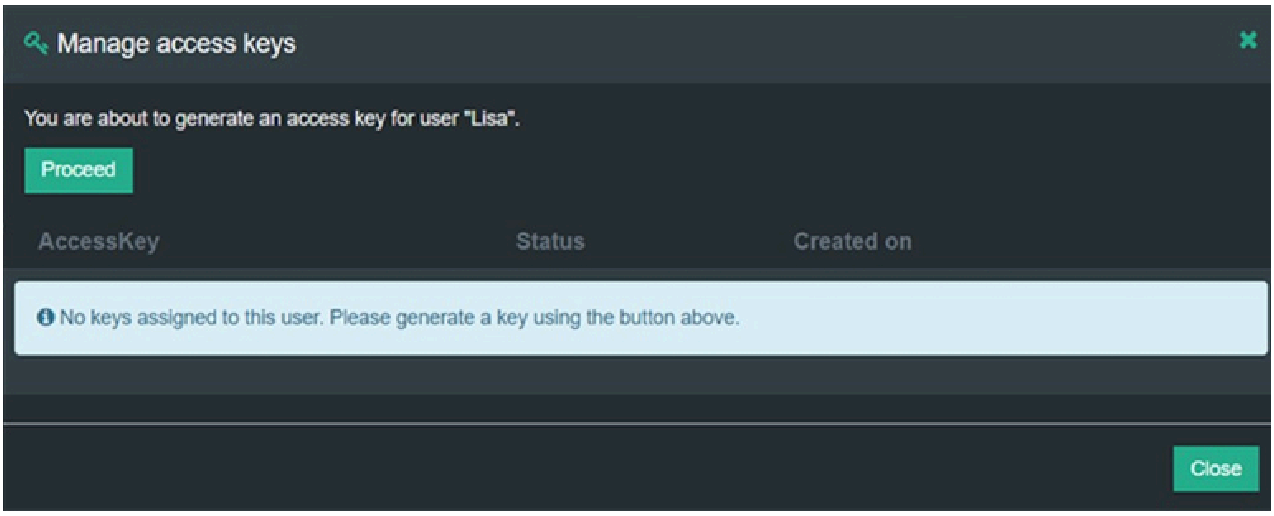

You are about to generate an access key for user "Lisa".

AccessKey	Status	Created on

No keys assigned to this user. Please generate a key using the button above.

Copy and save the SecretAccessKey to a secure location. It is not shown again and cannot be recovered later.

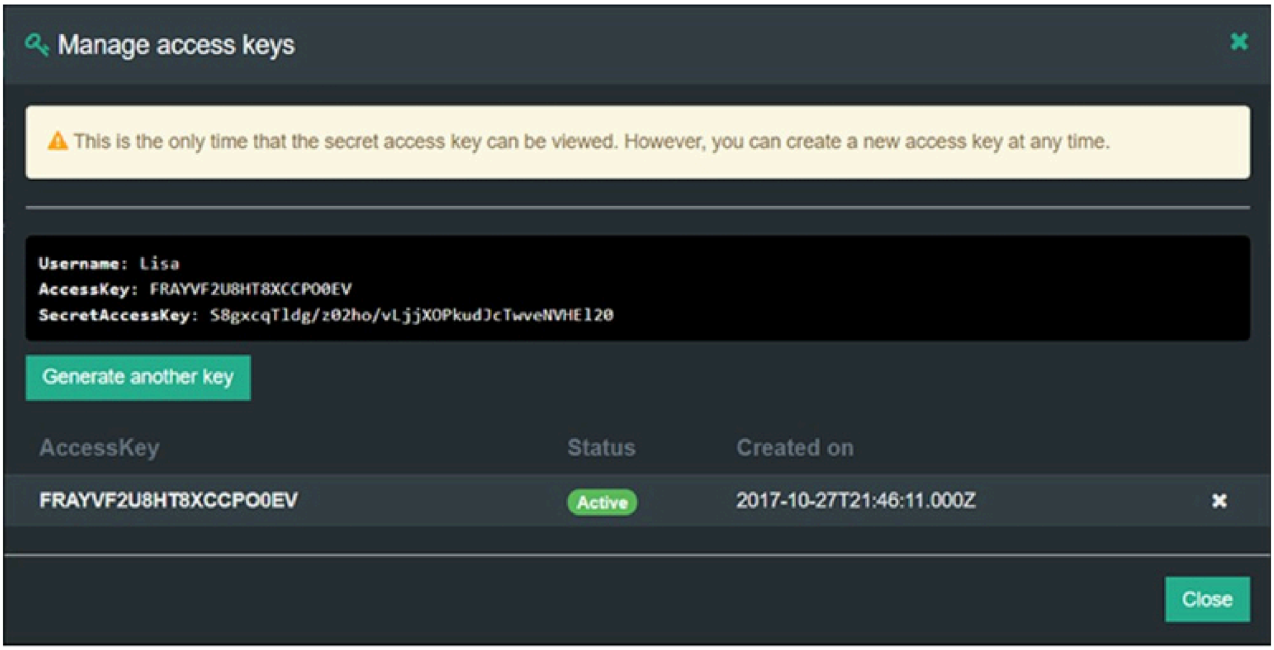

Manage access keys

⚠ This is the only time that the secret access key can be viewed. However, you can create a new access key at any time.

```
Username: Lisa
AccessKey: FRAYVF2U8HT8XCCPO0EV
SecretAccessKey: S8gxcqTldg/z02ho/vLjjXOPkudJcTwveNVHE120
```

AccessKey	Status	Created on	
FRAYVF2U8HT8XCCPO0EV	Active	2017-10-27T21:46:11.000Z	✖

Now we can go to bucket creation. Go to the S3 Browser interface.

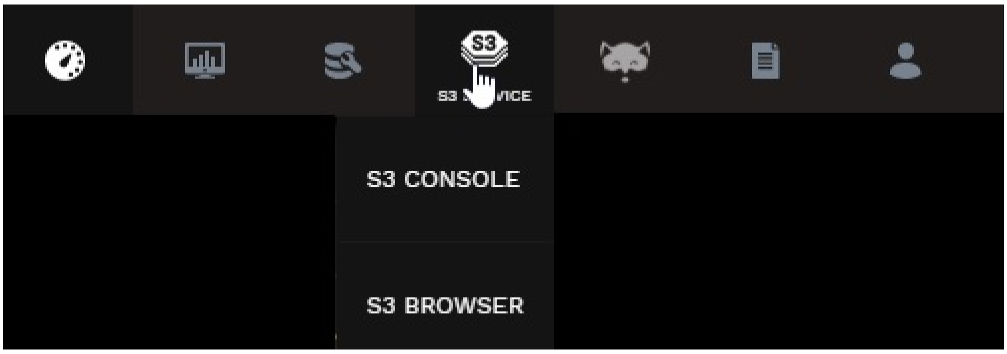

The S3 Browser opens the main window, from which one can see the entire roster of buckets.

Click the Create Bucket button in the top left of the main window.

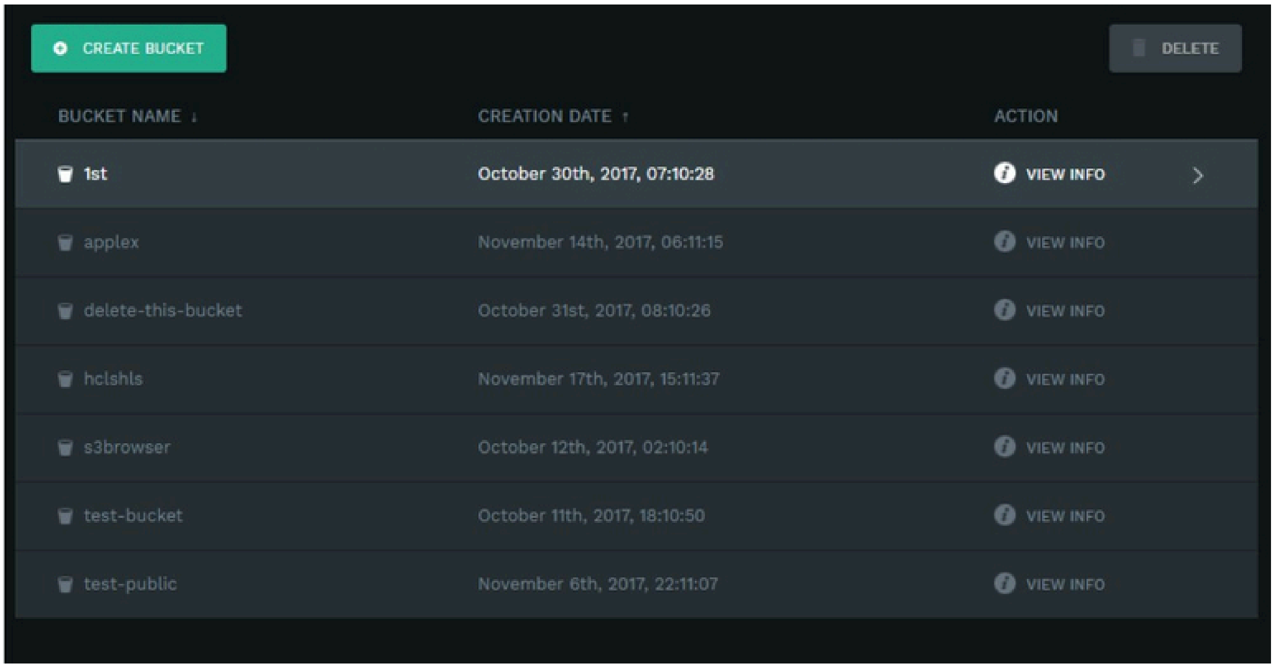

BUCKET NAME ↓	CREATION DATE ↑	ACTION
1st	October 30th, 2017, 07:10:28	VIEW INFO
applex	November 14th, 2017, 06:11:15	VIEW INFO
delete-this-bucket	October 31st, 2017, 08:10:26	VIEW INFO
hclshls	November 17th, 2017, 15:11:37	VIEW INFO
s3browser	October 12th, 2017, 02:10:14	VIEW INFO
test-bucket	October 11th, 2017, 18:10:50	VIEW INFO
test-public	November 6th, 2017, 22:11:07	VIEW INFO

Enter a name for the new bucket and click on Create button.

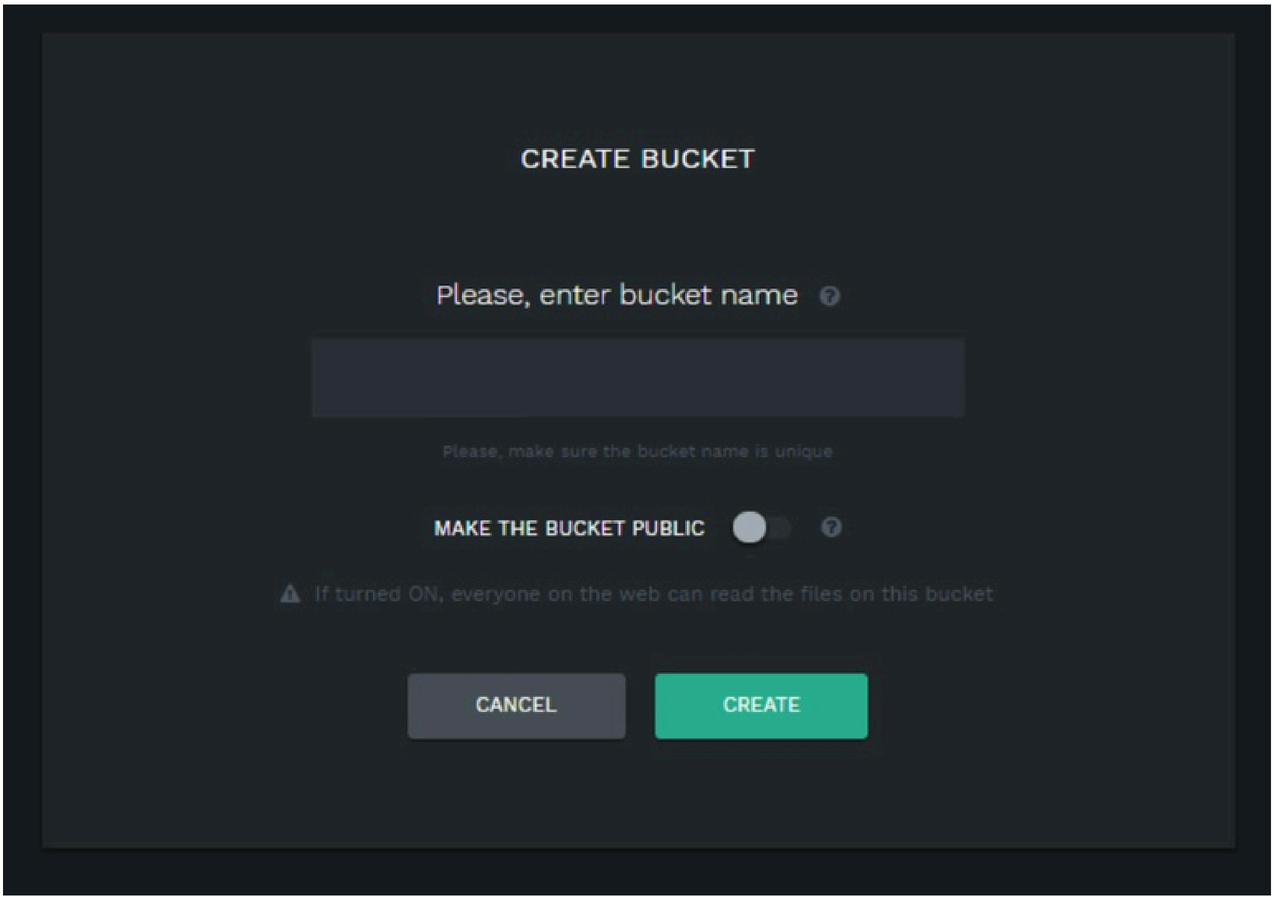
CREATE BUCKET

Please, enter bucket name

Please, make sure the bucket name is unique

MAKE THE BUCKET PUBLIC

If turned ON, everyone on the web can read the files on this bucket

That's all on the Scality side. Now we can go to Data Protector for Cloud Workloads.

Open the "Backup Destination" tab from the left side menu and choose "Amazon S3 / S3-compatible" as the new type of backup destination.

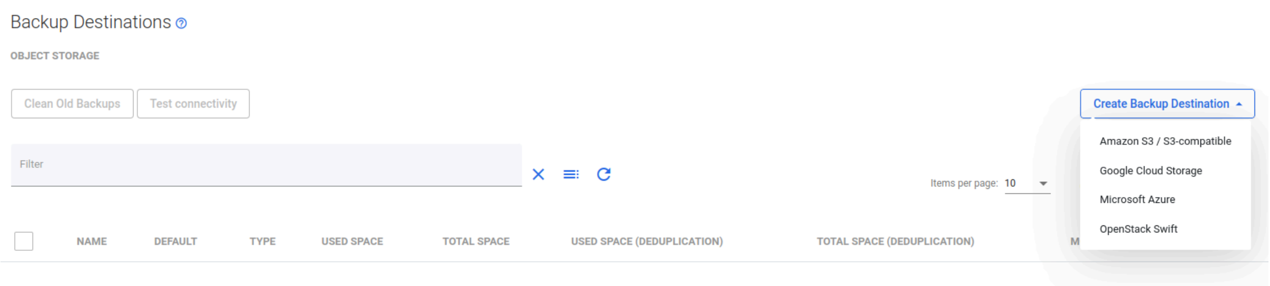

Like in other S3-compatible backup destinations, you have to fill in the fields below and provide the access and secret key.

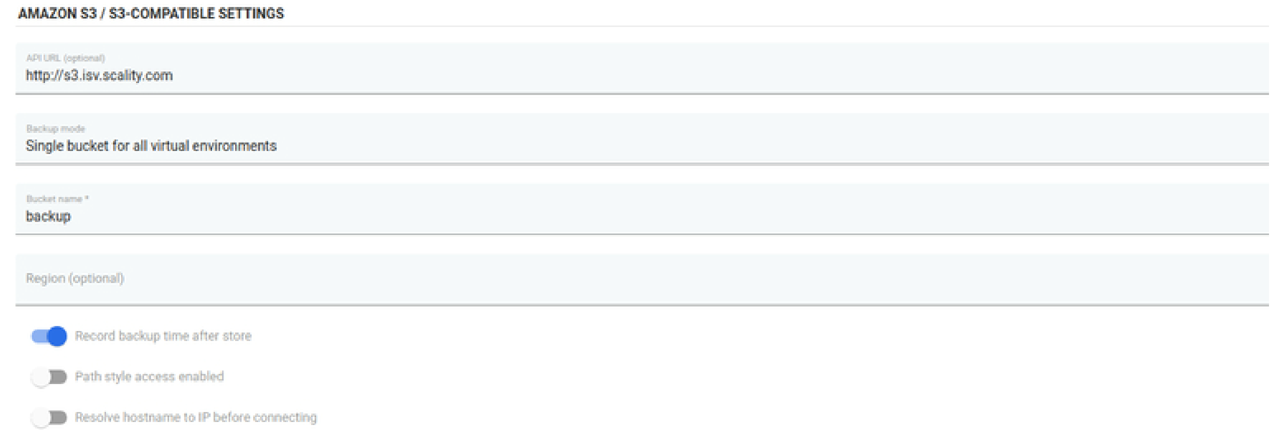

AMAZON S3 / S3-COMPATIBLE SETTINGS

API URL (optional)
http://s3.isv.scality.com

Backup mode
Single bucket for all virtual environments

Bucket name *
backup

Region (optional)

● Record backup time after store

○ Path style access enabled

○ Resolve hostname to IP before connecting

That's it, you can now safely store your backups.

Enterprise Backup Providers

Enterprise Backup Providers

Micro Focus Data Protector

To integrate Data Protector for Cloud Workloads with Micro Focus Data Protector list device names by running following command on Data Protector for Cloud Workloads Node:

```
[root@protectorvp ~]# /opt/omni/bin/omnidownload -list_devices

Device Name                      Host
=========================================================================
dyskd_gw1                        win-srv-proxy
vp_gw2                           protector11.lab.local
vp_protectorlab_gw1         protectorlab.lab.local
vp_protectorvp_gw1               protectorvp.lab.local
=========================================================================
```

Next, go to Data Protector for Cloud Workloads and go to **Backup Destinations** → **Enterprise**. Click **Create Backup destination** and choose **Micro Focus Data Protector**. Type the name for new backup destination and provide **Device name** which you get from first step.

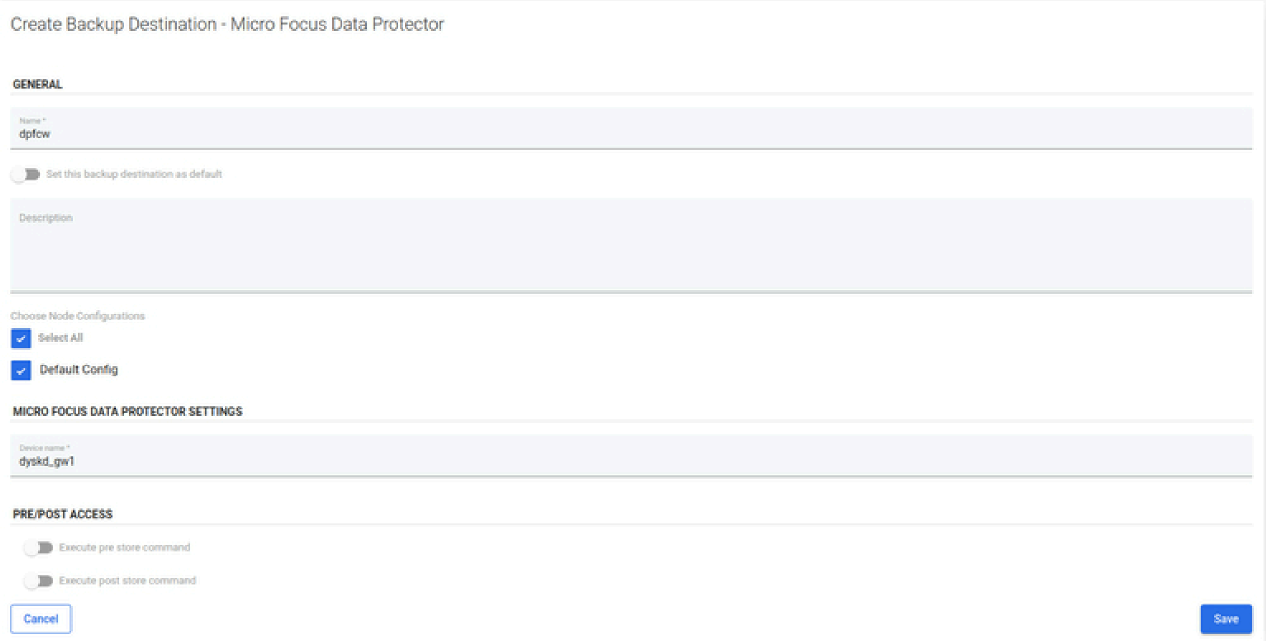

Create Backup Destination - Micro Focus Data Protector

GENERAL

Name *
dpfcw

Set this backup destination as default

Description

Choose Node Configurations

☑ Select All

☑ Default Config

MICRO FOCUS DATA PROTECTOR SETTINGS

Device name *
dyskd_gw1

PRE/POST ACCESS

Execute pre store command

Execute post store command

Initial Configuration

Node

1. Set up the backup destinations (examples):
   - File System
   - Virtual Data Optimizer (VDO)
2. For backup strategies involving **disk attachment** mode, follow these steps: LVM setup on Data Protector for Cloud Workloads Node for disk attachment backup mode.

Server

1. Upload your license key:
   - if you don't have it, you can contact the OpenText
   - log in to the web UI and go to the **Settings** → **License** and upload your `license.key` file.
2. It is **highly recommended** to set up Data Protector for Cloud Workloads DB backup - the database is key to restoring your Data Protector for Cloud Workloads environment and later all of the backups that you need.
3. Admin account setup:
   - for audit purposes, it is recommended to add individual admin accounts using the Access Management section

> **Note:** make sure to set the correct **time zone** for each user - the default admin account has **UTC** by default.

Configuration Wizard

- The configuration wizard can be accessed from the main dashboard by clicking on the "configuration wizard" button on the right.

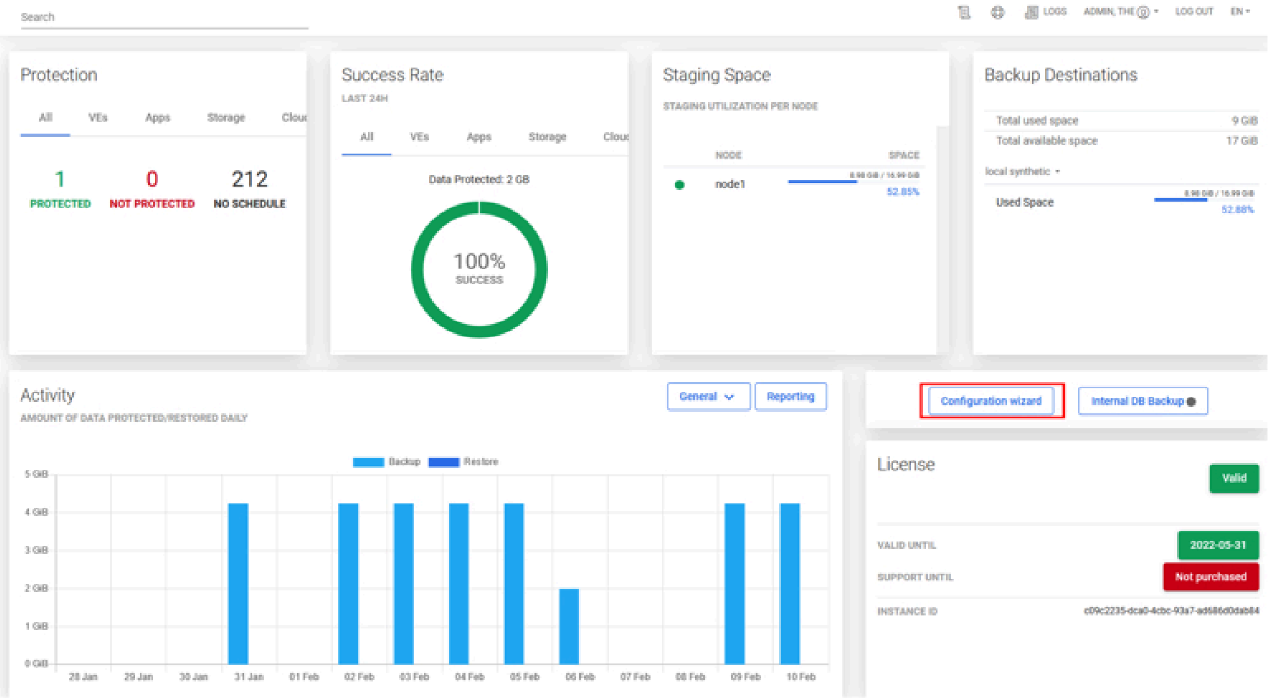

Welcome page - nodes

- On the welcome page, you should see the Data Protector for Cloud Workloads Nodes summary. You need at least one fully running node to continue. If you meet this requirement, click on the Next button.

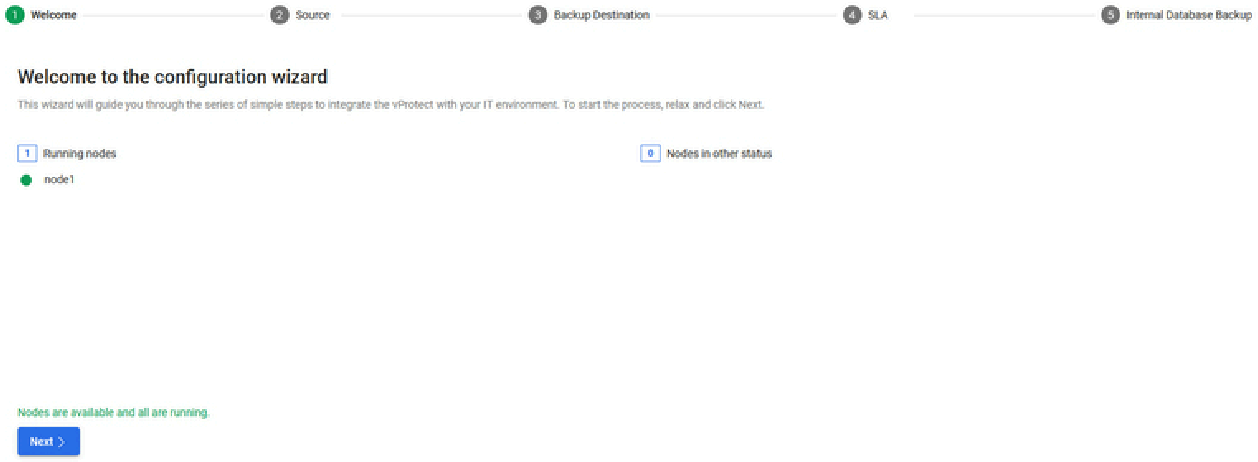

Welcome to the configuration wizard

This wizard will guide you through the series of simple steps to integrate the vProtect with your IT environment. To start the process, relax and click Next.

1 Running nodes

● node1

0 Nodes in other status

Nodes are available and all are running.

Add a hypervisor

- In the Hypervisor section, you will start by selecting the hypervisor manager or hypervisor that you want to add. You can repeat this step if you have many types of virtualization providers.

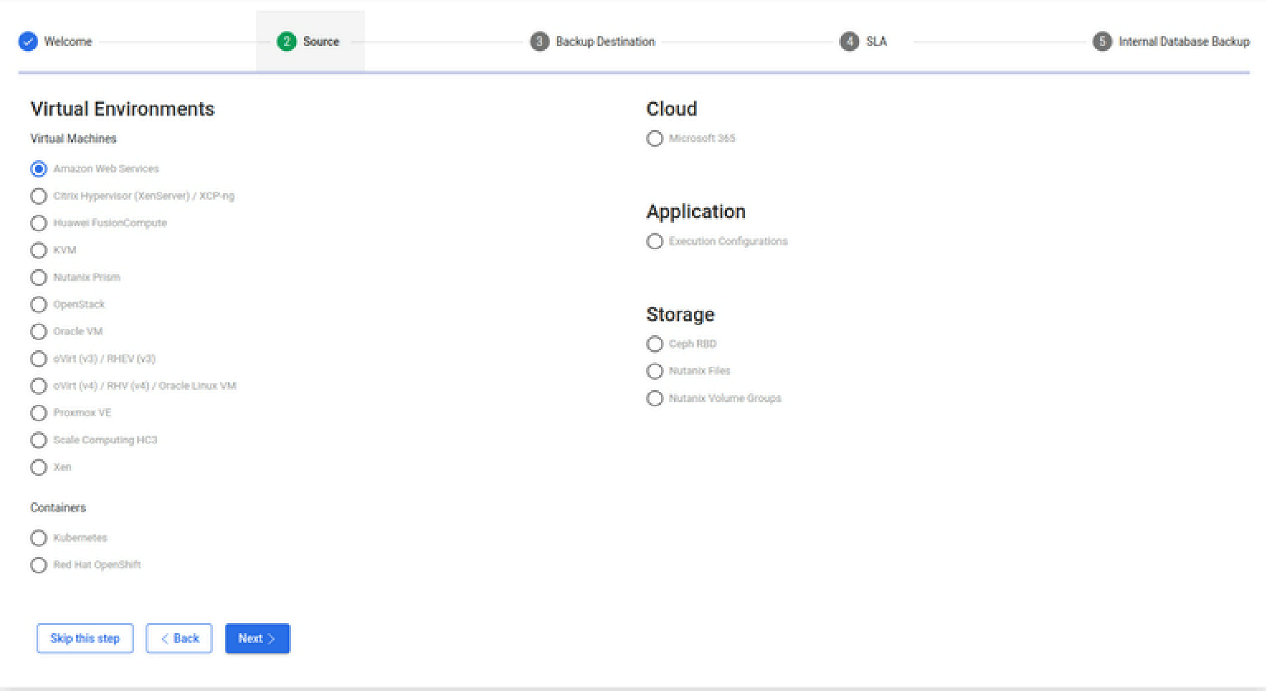

✓ Welcome — 2 Source — 3 Backup Destination — 4 SLA — 5 Internal Database Backup

Virtual Environments

Virtual Machines

● Amazon Web Services
○ Citrix Hypervisor (XenServer) / XCP-ng
○ Huawei FusionCompute
○ KVM
○ Nutanix Prism
○ OpenStack
○ Oracle VM
○ oVirt (v3) / RHEV (v3)
○ oVirt (v4) / RHV (v4) / Oracle Linux VM
○ Proxmox VE
○ Scale Computing HC3
○ Xen

Containers

○ Kubernetes
○ Red Hat OpenShift

Cloud

○ Microsoft 365

Application

○ Execution Configurations

Storage

○ Ceph RBD
○ Nutanix Files
○ Nutanix Volume Groups

- For the Citrix hypervisor (as an example) you have to enter the following parameters

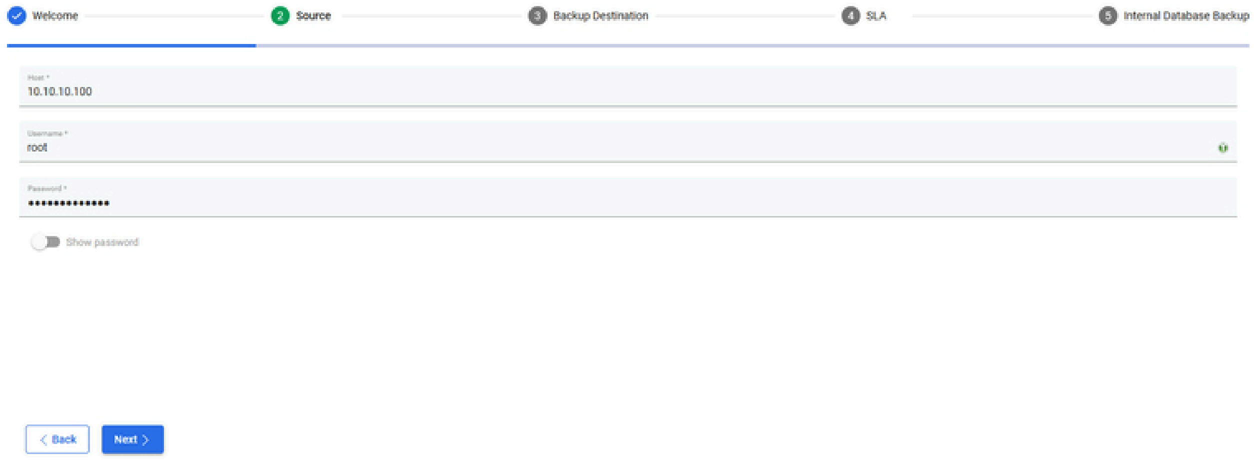

✓ Welcome — 2 Source — 3 Backup Destination — 4 SLA — 5 Internal Database Backup

Host *
10.10.10.100

Username *
root

Password *
••••••••••••••

Show password

- Choose node configuration

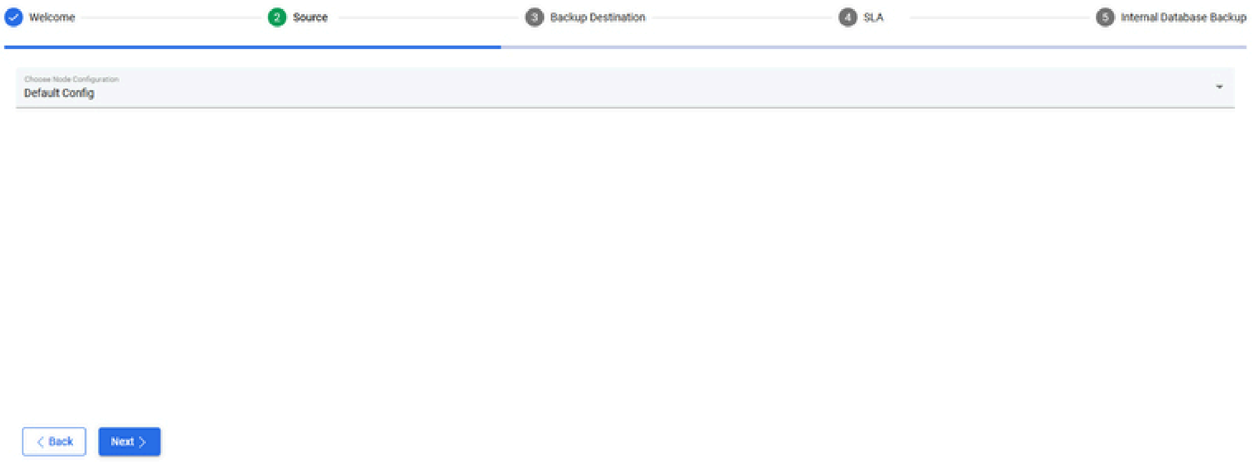

Choose Node Configuration
Default Config

Select a backup strategy for your hypervisor

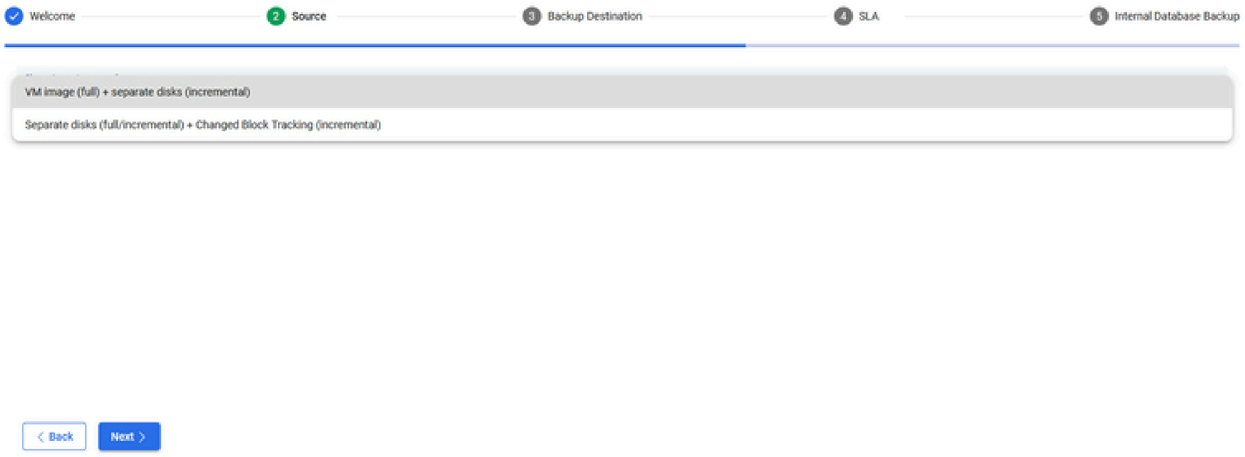

VM image (full) + separate disks (incremental)

Separate disks (full/incremental) + Changed Block Tracking (incremental)

- Optionally, you can add an additional NIC for transfer purposes (provide IP address)

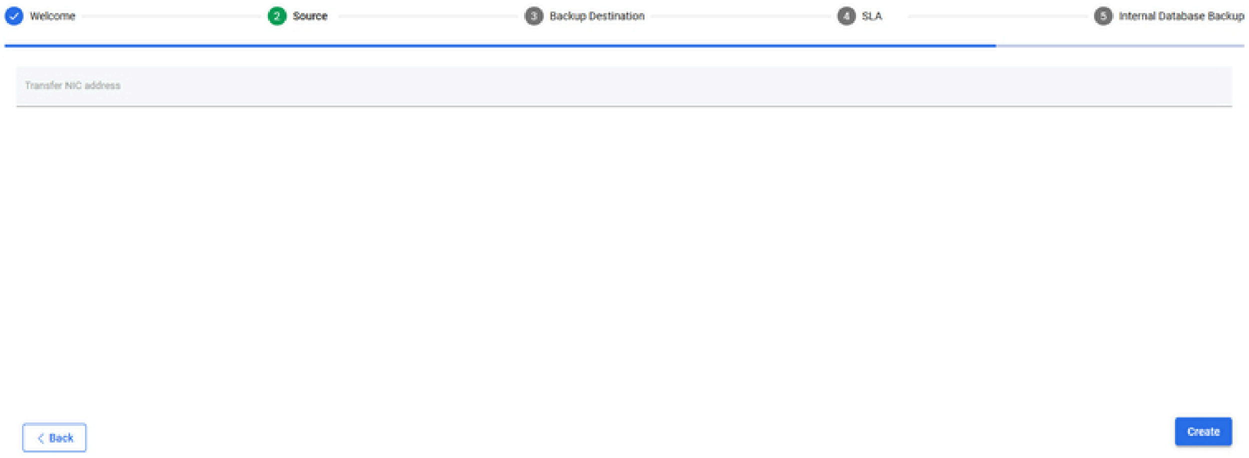

Transfer NIC address

- At the end, you will see a popup window that allows you to run inventory synchronization. After that, you should see all the virtual machines from that hypervisor.

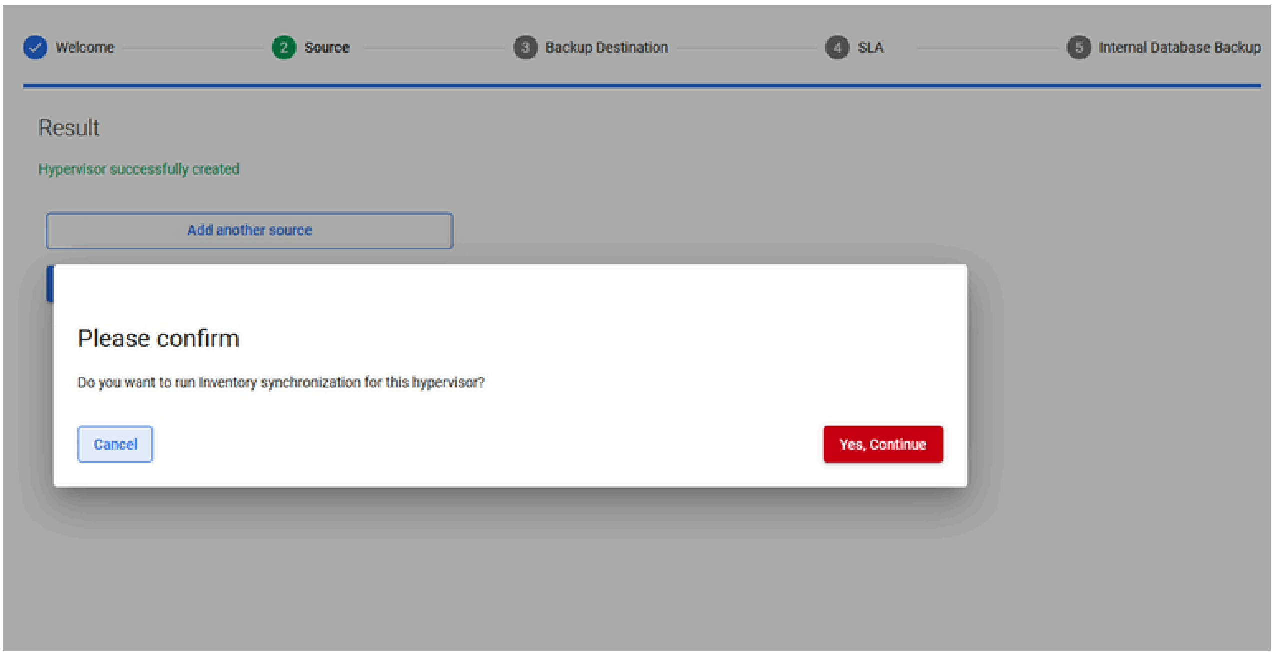

Result

Hypervisor successfully created

Please confirm

Do you want to run Inventory synchronization for this hypervisor?

Add backup destination

- In the next section, you can add a backup destination. In this case, you can also repeat the whole process so that you can add multiple providers using the wizard.
- Choose a backup destination (we used File System as an example)

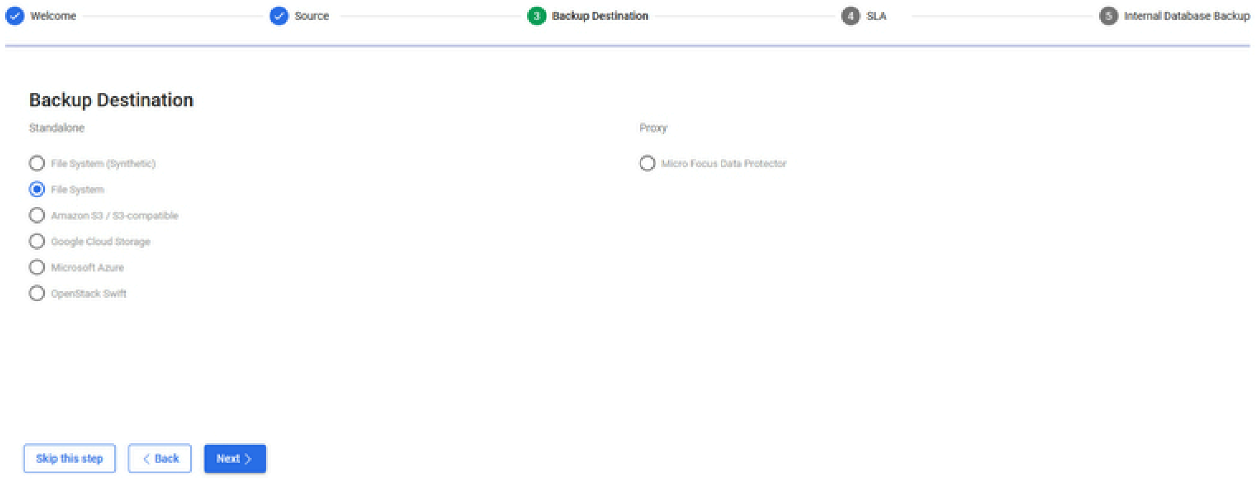

Backup Destination

Standalone

○ File System (Synthetic)
● File System
○ Amazon S3 / S3-compatible
○ Google Cloud Storage
○ Microsoft Azure
○ OpenStack Swift

Proxy

○ Micro Focus Data Protector

- First, enter a name for your backup destination

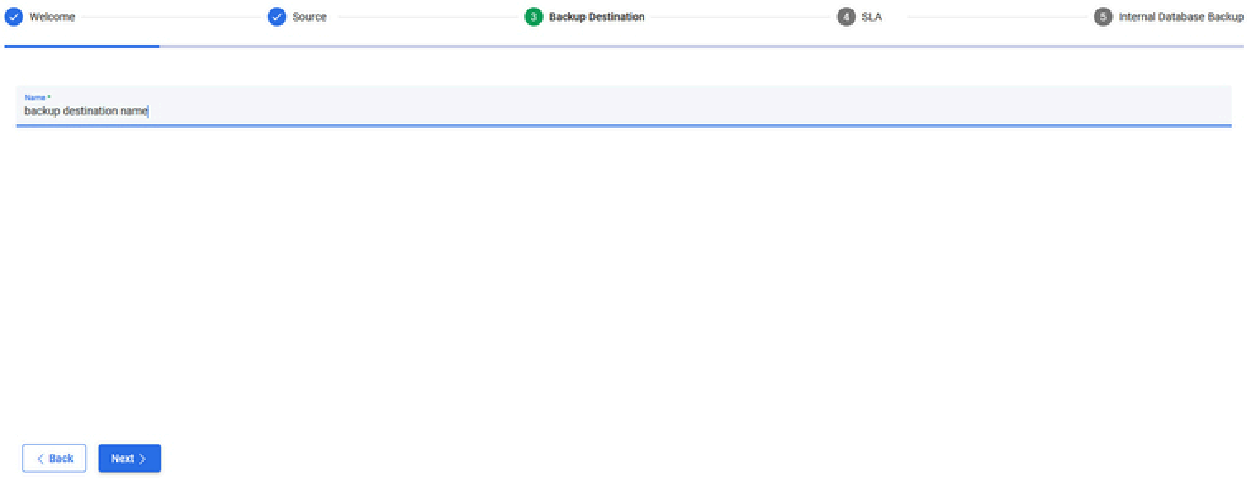

Name *
backup destination name

- Choose, if you want to use deduplication based on [Virtual Data Optimizer (VDO)]

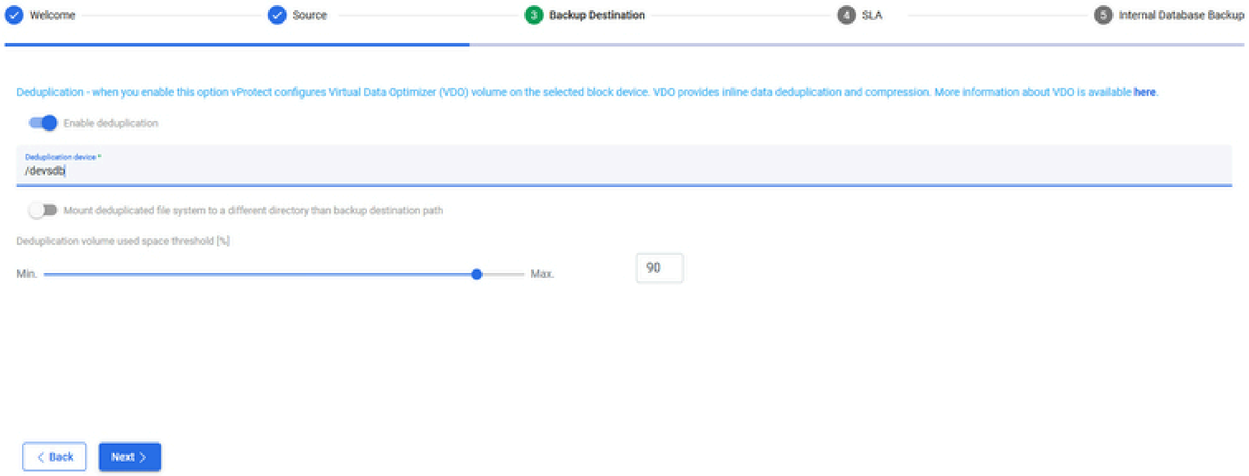

Deduplication - when you enable this option vProtect configures Virtual Data Optimizer (VDO) volume on the selected block device. VDO provides inline data deduplication and compression. More information about VDO is available here.

Enable deduplication

Deduplication device *
/devsdb

Mount deduplicated file system to a different directory than backup destination path

Deduplication volume used space threshold [%]

Min. Max. 90

- Set up a storage path, where your data should be stored

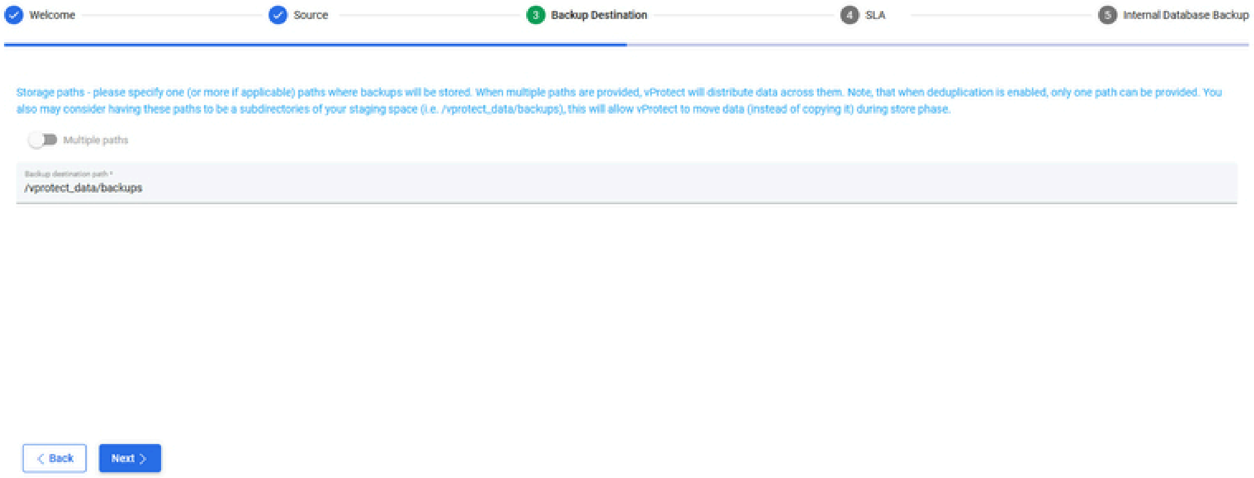

Storage paths - please specify one (or more if applicable) paths where backups will be stored. When multiple paths are provided, vProtect will distribute data across them. Note, that when deduplication is enabled, only one path can be provided. You also may consider having these paths to be a subdirectories of your staging space (i.e. /vprotect_data/backups), this will allow vProtect to move data (instead of copying it) during store phase.

Multiple paths

Backup destination path *
/vprotect_data/backups

- Optionally you can enable encryption (AES-256 algorithm) - if you enable it, remember that you will not benefit from deduplication.

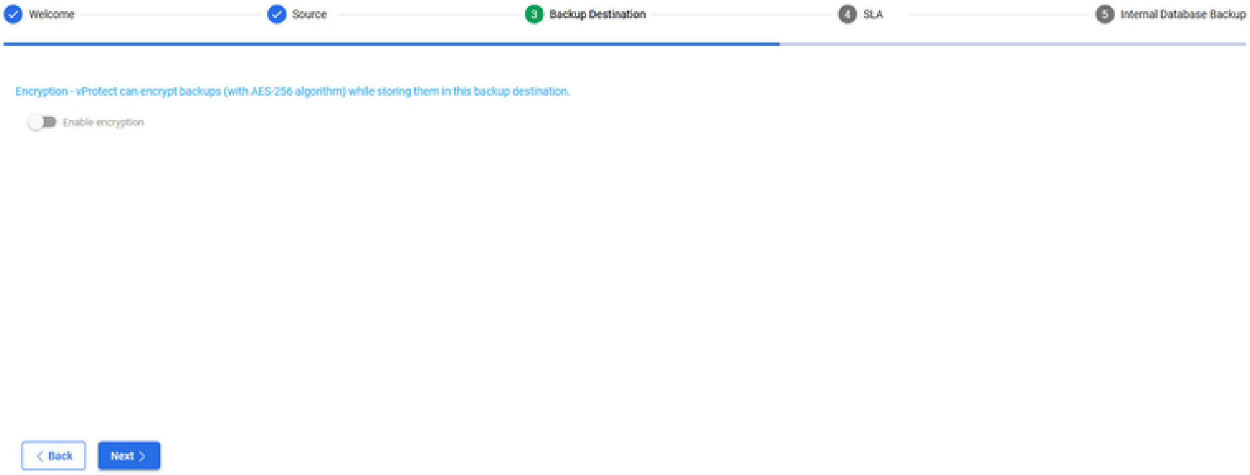

Encryption - vProtect can encrypt backups (with AES-256 algorithm) while storing them in this backup destination.

○ Enable encryption

- configuration for pre/post execution command. If you use a File System with VDO, skip this step.

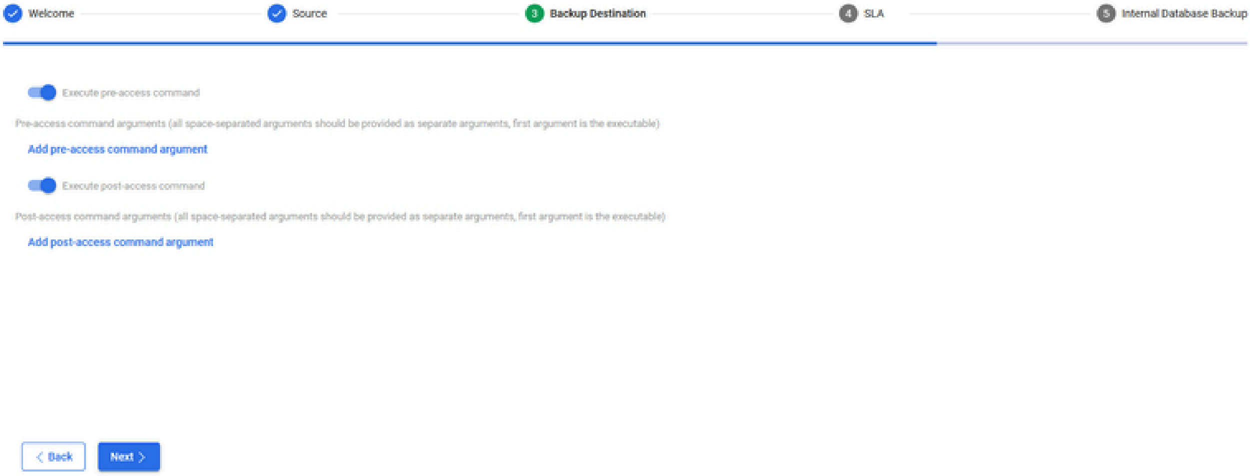
Welcome
Source
3 Backup Destination
4 SLA
5 Internal Database Backup
Add pre-access command argument
Add post-access command argument
< Back
Next >

● Execute pre-access command

Pre-access command arguments (all space-separated arguments should be provided as separate arguments, first argument is the executable)

● Execute post-access command

Post-access command arguments (all space-separated arguments should be provided as separate arguments, first argument is the executable)

- Decide if you want to set up this backup destination as the default one.

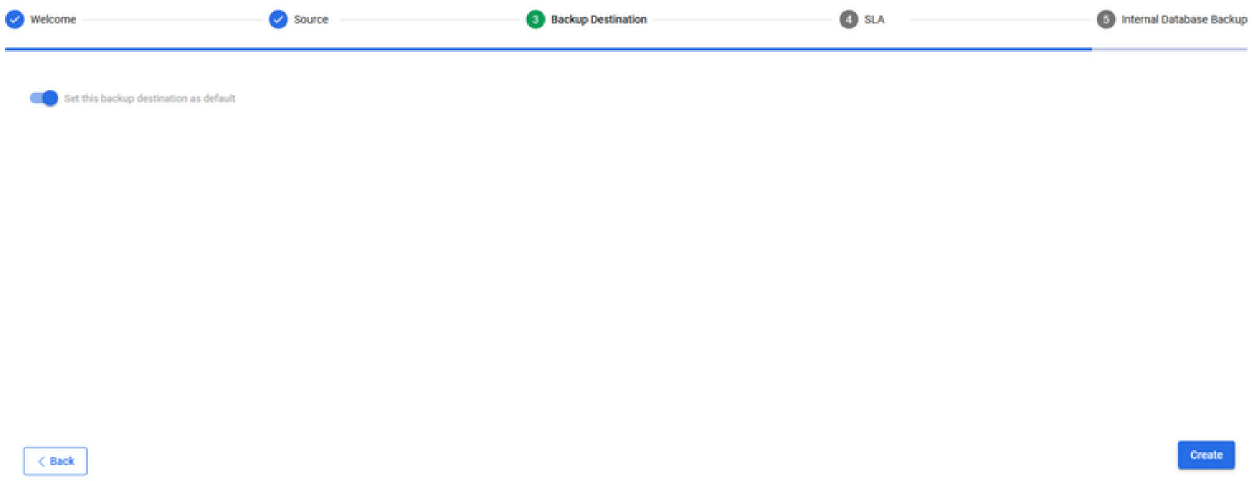
Welcome
Source
3 Backup Destination
4 SLA
5 Internal Database Backup
Set this backup destination as default
Back
Create

- Finish this step by going to the next section or adding another backup destination

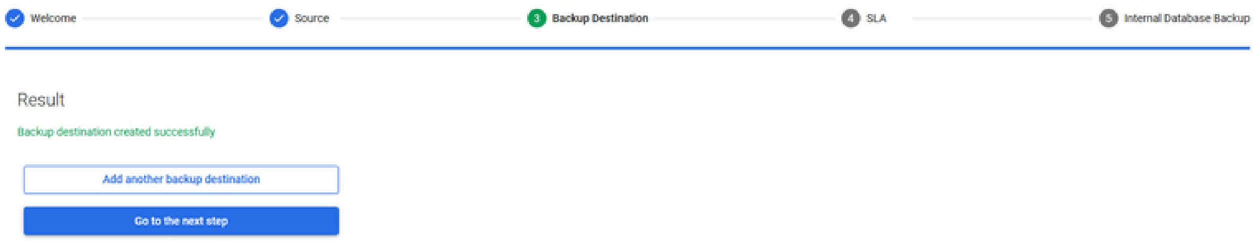

Result

Backup destination created successfully

Add SLA

In this example, we will add SLA for Virtual Environment backup.

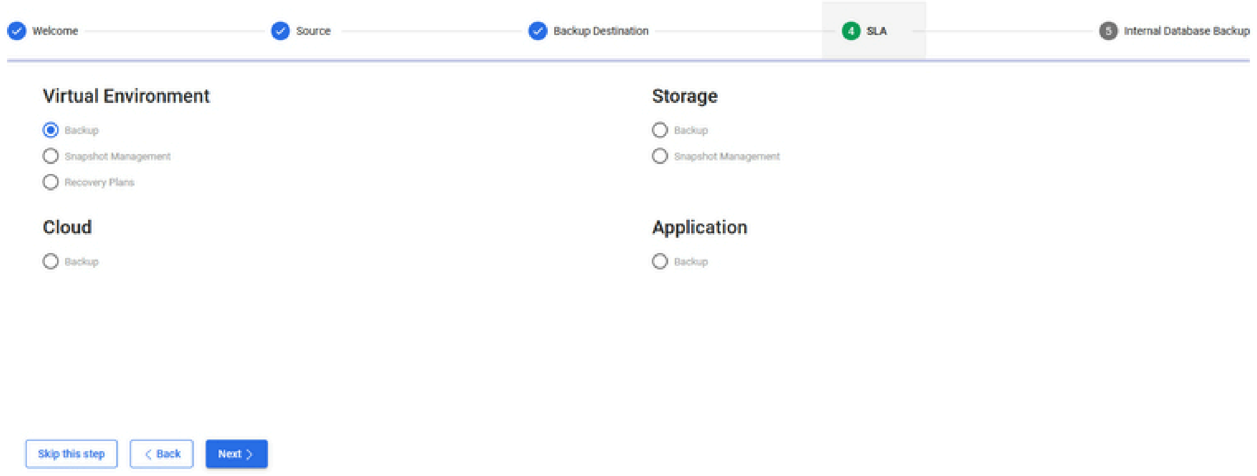

Virtual Environment

● Backup
○ Snapshot Management
○ Recovery Plans

Storage

○ Backup
○ Snapshot Management

Cloud

○ Backup

Application

○ Backup

Add policy

- Choose a name for the policy, auto-remove non-present virtual environments (if Data Protector for Cloud Workloads should remove VM from a policy that no longer exists) tick the checkbox, and set the priority

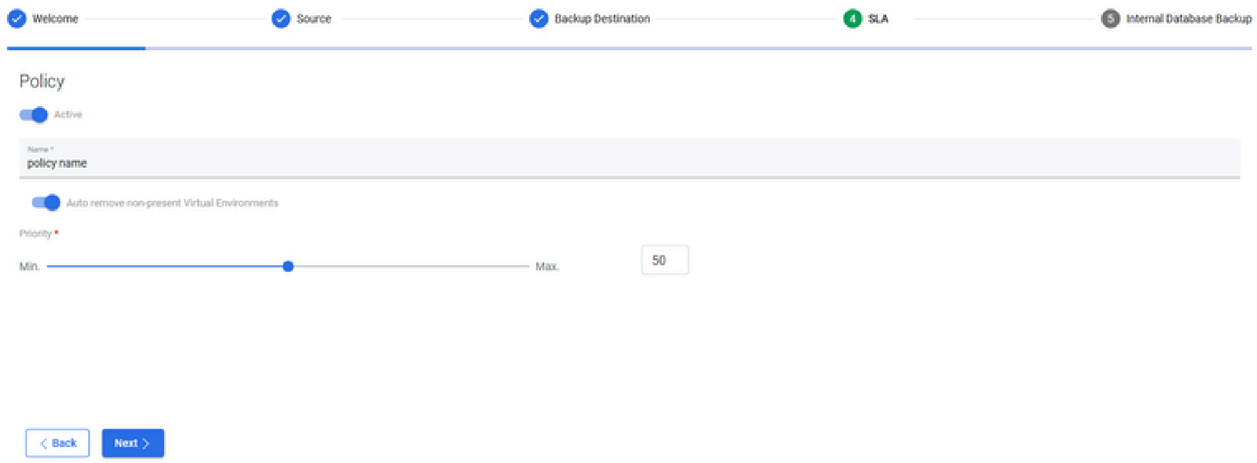

Policy

Active

Name *
policy name

Auto remove non-present Virtual Environments

Priority *

Min. Max. 50

- Choose if you want to use auto assign mode based on tags and regular expressions (matched against the VM name, `.*` matches all characters 0 or more times)

> **Note:** check the Administration section for details of Backup SLAs to each protected platform

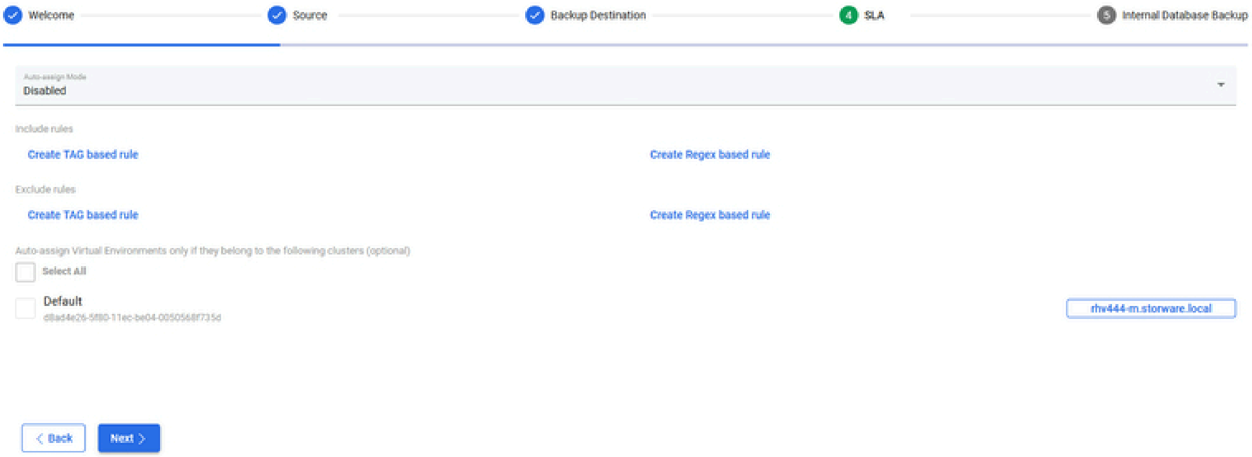

Auto-assign Mode: Disabled

Include rules

Exclude rules

Auto-assign Virtual Environments only if they belong to the following clusters (optional)

□ Select All

□ Default d8ad4e26-5f80-11ec-be04-0050568f735d — rhv444-m.storware.local

- Manually add the VMs if you do not want to use the auto-assignment mode

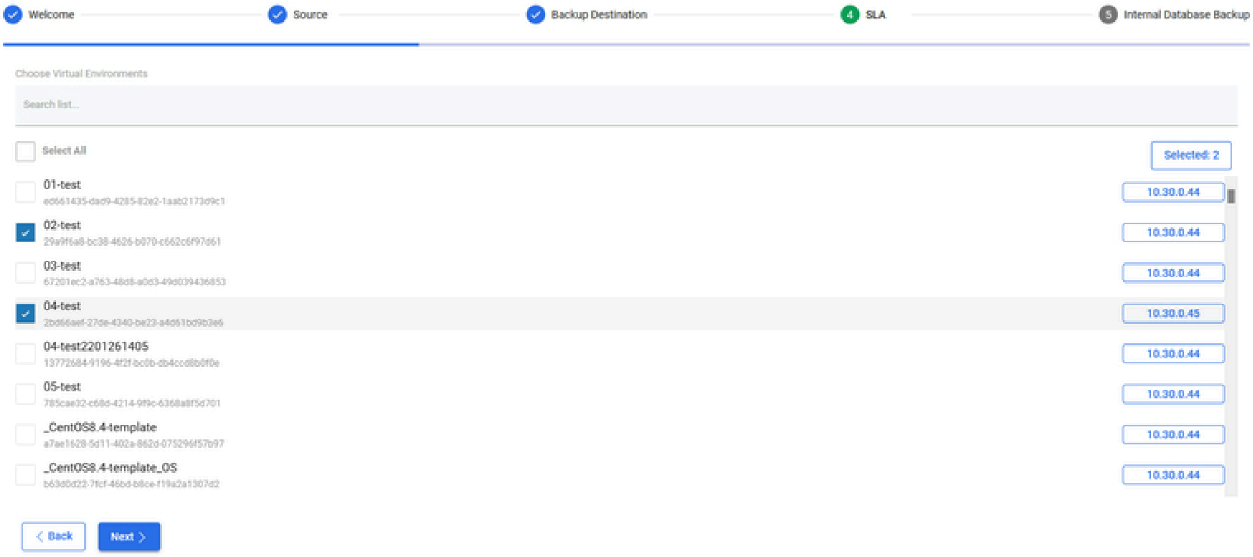

Choose Virtual Environments

Search list...

□ Select All — Selected: 2

□	01-test ed661435-dad9-4285-82e2-1aab2173d9c1	10.30.0.44
☑	02-test 29a9f6a8-bc38-4626-b070-c662c6f97d61	10.30.0.44
□	03-test 67201ec2-a763-48d8-a0d3-49d039436853	10.30.0.44
☑	04-test 2bd66aef-27de-4340-be23-a4d61bd9b3e6	10.30.0.45
□	04-test2201261405 13772684-9196-4f2f-bc0b-db4ccd8b0f0e	10.30.0.44
□	05-test 785cae32-c68d-4214-9f9c-6368a8f5d701	10.30.0.44
□	_CentOS8.4-template a7ae1628-5d11-402a-862d-075296f57b97	10.30.0.44
□	_CentOS8.4-template_OS b63d0d22-7fcf-46bd-b8ce-f19a2a1307d2	10.30.0.44

- Choose a backup destination target for this policy

> **Note:** You can now customize retention. Each backup destination has its own retention settings. Whichever condition is met first (either number of versions has been reached or the backup is older than the given limit), it is removed from the backup destination.

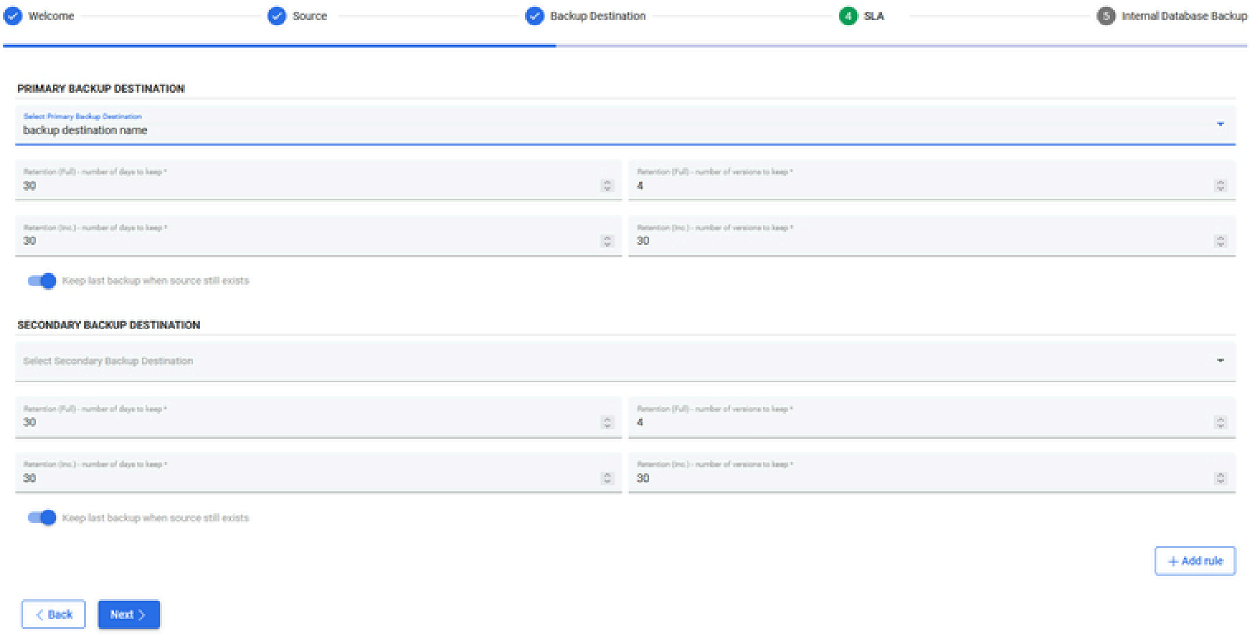

PRIMARY BACKUP DESTINATION

Select Primary Backup Destination: backup destination name

Retention (Full) - number of days to keep *: 30

Retention (Full) - number of versions to keep *: 4

Retention (Inc.) - number of days to keep *: 30

Retention (Inc.) - number of versions to keep *: 30

Keep last backup when source still exists

SECONDARY BACKUP DESTINATION

Select Secondary Backup Destination

Retention (Full) - number of days to keep *: 30

Retention (Full) - number of versions to keep *: 4

Retention (Inc.) - number of days to keep *: 30

Retention (Inc.) - number of versions to keep *: 30

Keep last backup when source still exists

Configure the following thresholds:

- Fail rest of the backup tasks if more than X % of EXPORT tasks already failed
- Fail rest of the backup tasks if more than X % of STORE tasks already failed

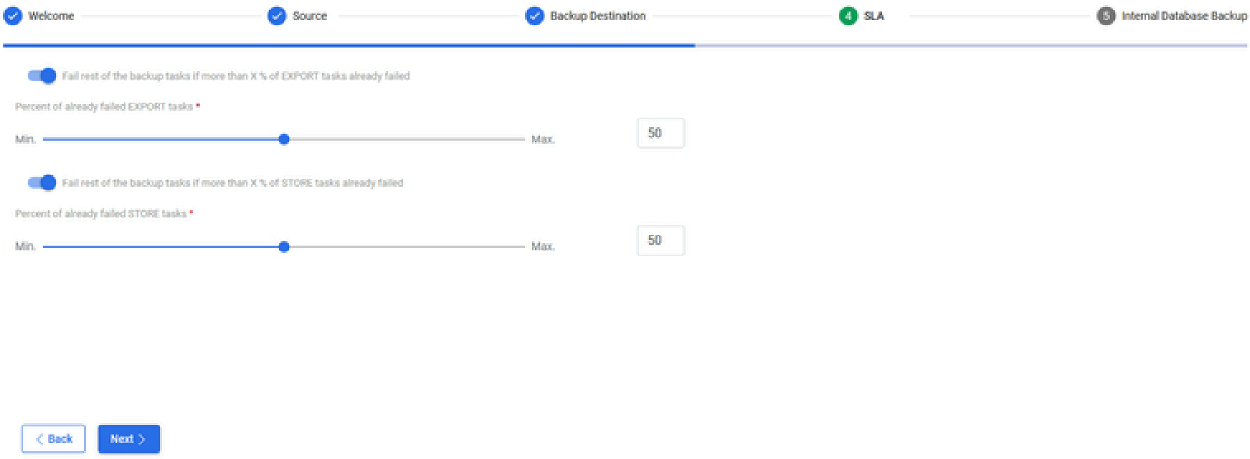

Fail rest of the backup tasks if more than X % of EXPORT tasks already failed

Percent of already failed EXPORT tasks *: 50

Fail rest of the backup tasks if more than X % of STORE tasks already failed

Percent of already failed STORE tasks *: 50

Add schedule

- Choose a name for the schedule and define the type:
  - Full
  - Incremental

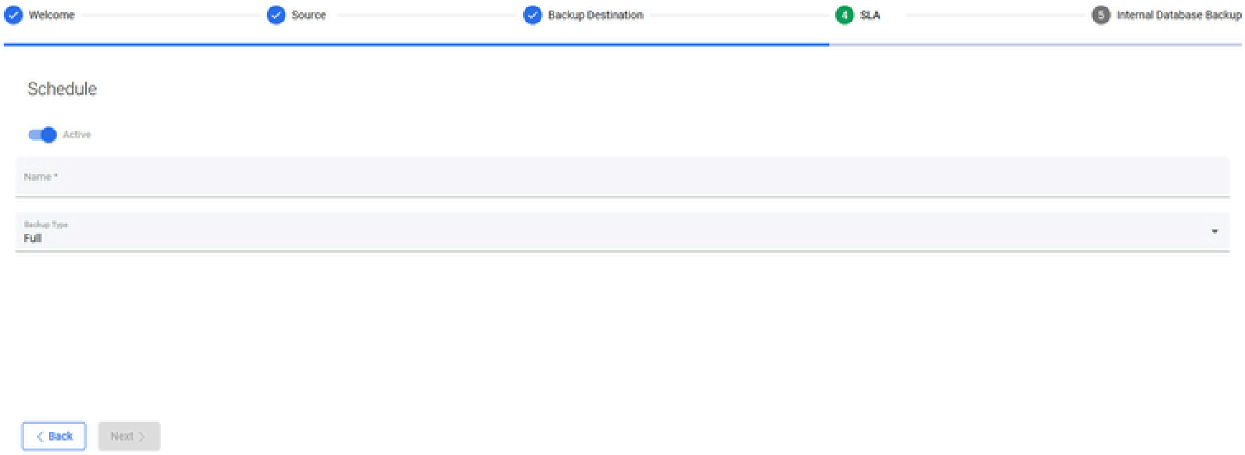

Schedule

Active

Name *

Backup Type
Full

- Define the execution type:
  - time
  - interval
- Define the start window length
- Choose the time of day for backup

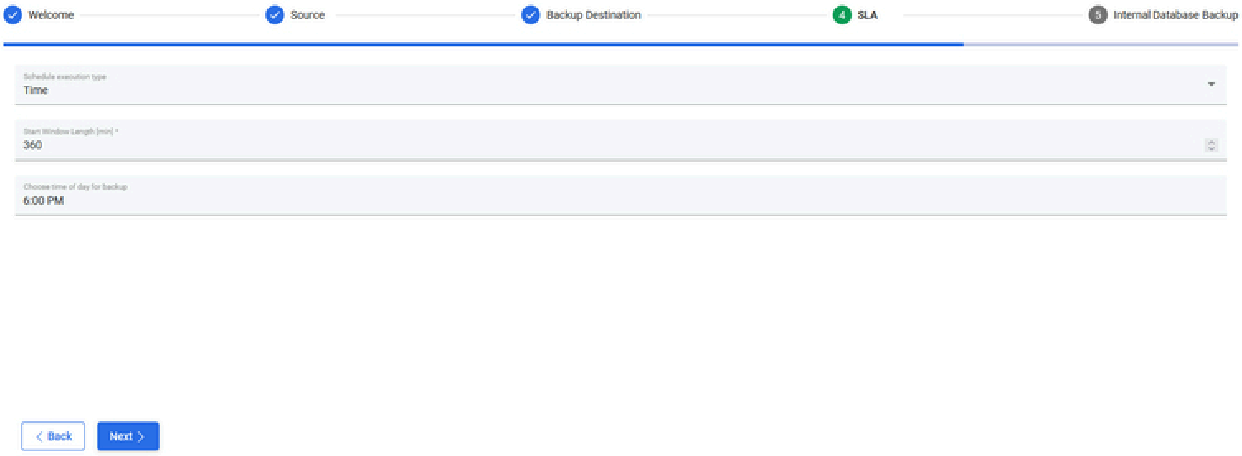

Schedule execution type
Time

Start Window Length [min] *
360

Choose time of day for backup
6:00 PM

- Choose
  - days (required).
  - day of week occurrence (optional)
  - selected months (optional)

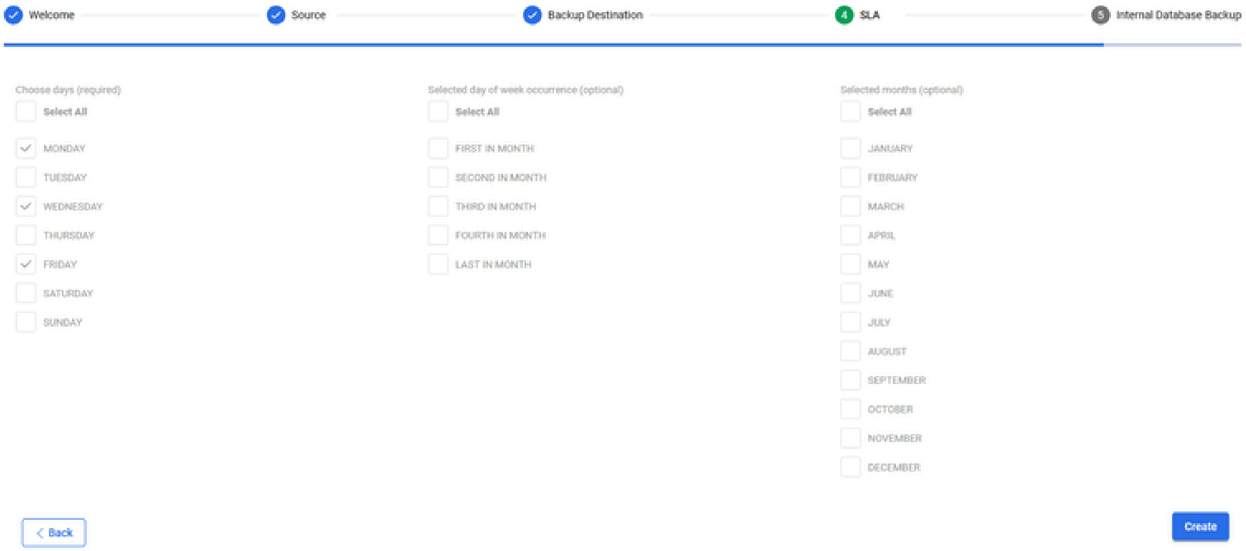

Choose days (required)

□ Select All

☑ MONDAY
□ TUESDAY
☑ WEDNESDAY
□ THURSDAY
☑ FRIDAY
□ SATURDAY
□ SUNDAY

Selected day of week occurrence (optional)

□ Select All

□ FIRST IN MONTH
□ SECOND IN MONTH
□ THIRD IN MONTH
□ FOURTH IN MONTH
□ LAST IN MONTH

Selected months (optional)

□ Select All

□ JANUARY
□ FEBRUARY
□ MARCH
□ APRIL
□ MAY
□ JUNE
□ JULY
□ AUGUST
□ SEPTEMBER
□ OCTOBER
□ NOVEMBER
□ DECEMBER

- Finish this step by going to the next section or adding another SLA.

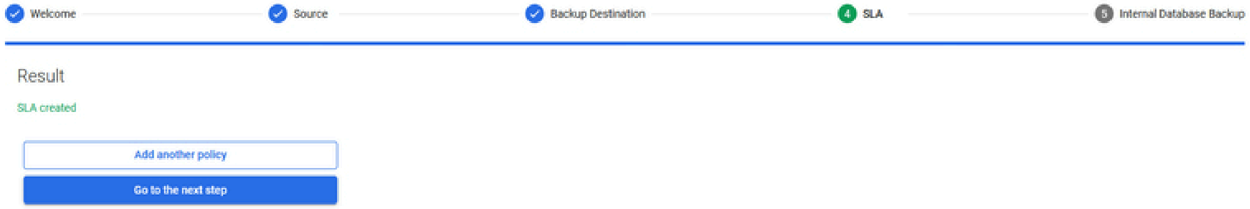

Result

SLA created

Add internal DB backup

- Choose which node config should be used to perform a Data Protector for Cloud Workloads DB backup

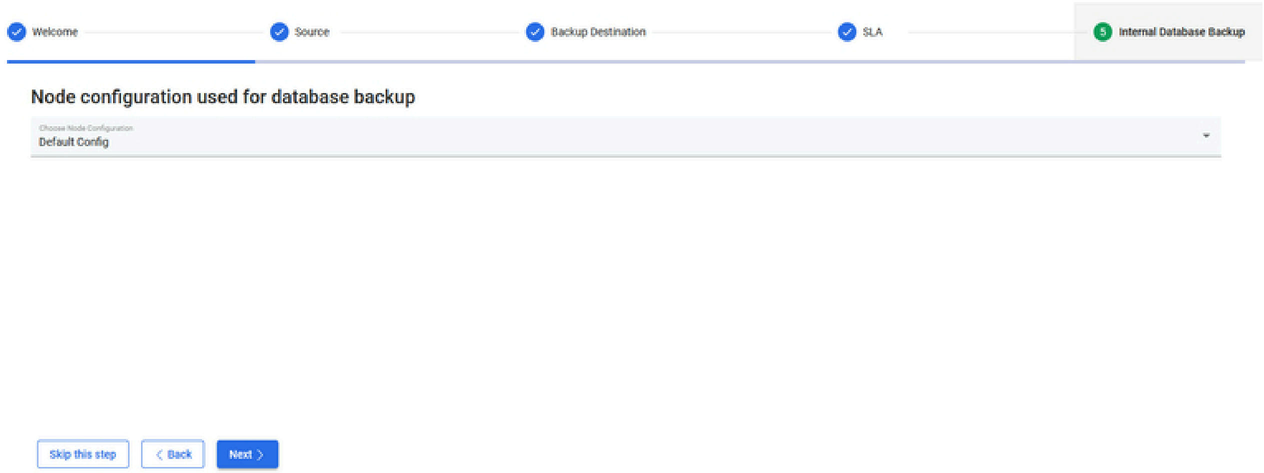

Node configuration used for database backup

Choose Node Configuration: Default Config

- Choose the backup destination for the DB backup

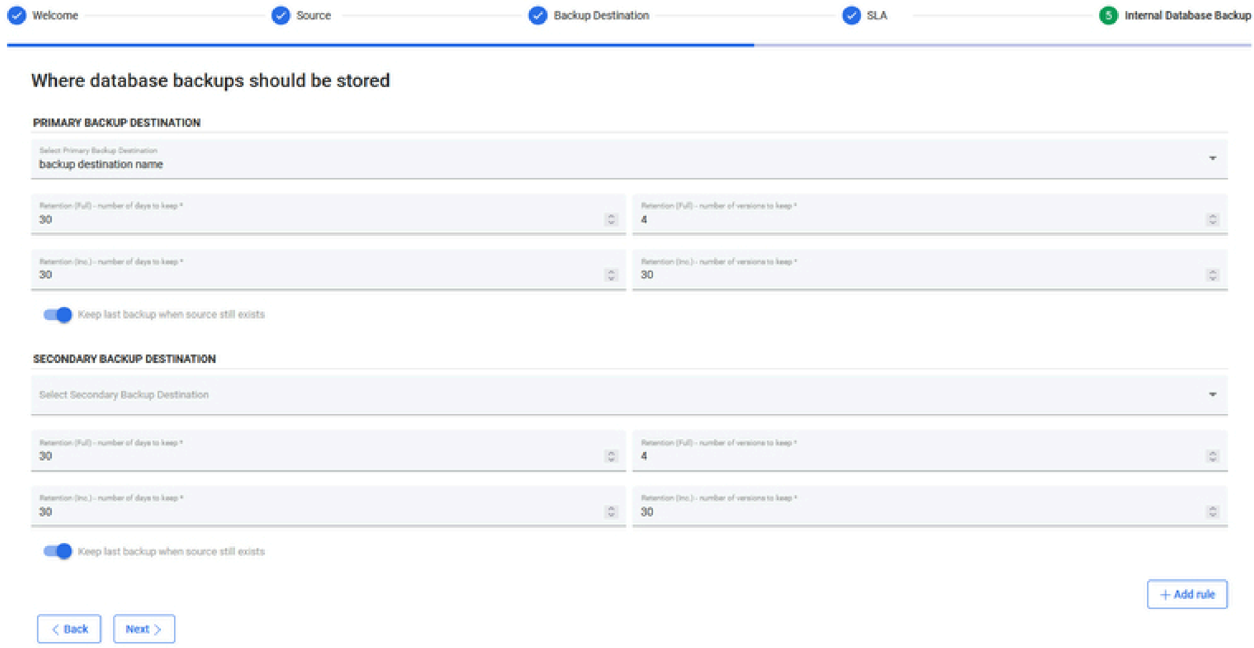

Where database backups should be stored

PRIMARY BACKUP DESTINATION

Select Primary Backup Destination: backup destination name

Retention (Full) - number of days to keep *: 30

Retention (Full) - number of versions to keep *: 4

Retention (Inc.) - number of days to keep *: 30

Retention (Inc.) - number of versions to keep *: 30

Keep last backup when source still exists

SECONDARY BACKUP DESTINATION

Select Secondary Backup Destination

Retention (Full) - number of days to keep *: 30

Retention (Full) - number of versions to keep *: 4

Retention (Inc.) - number of days to keep *: 30

Retention (Inc.) - number of versions to keep *: 30

Keep last backup when source still exists

Choose when the DB backup should be run (daily basis)

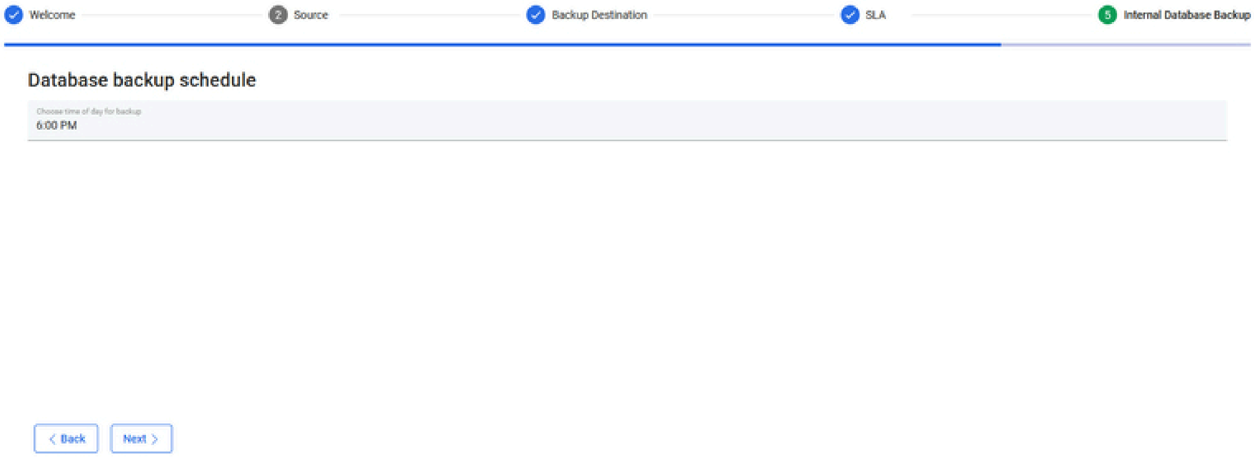

Database backup schedule

Choose time of day for backup
6:00 PM

- Finalize the configuration and/or run the backup manually (on demand)

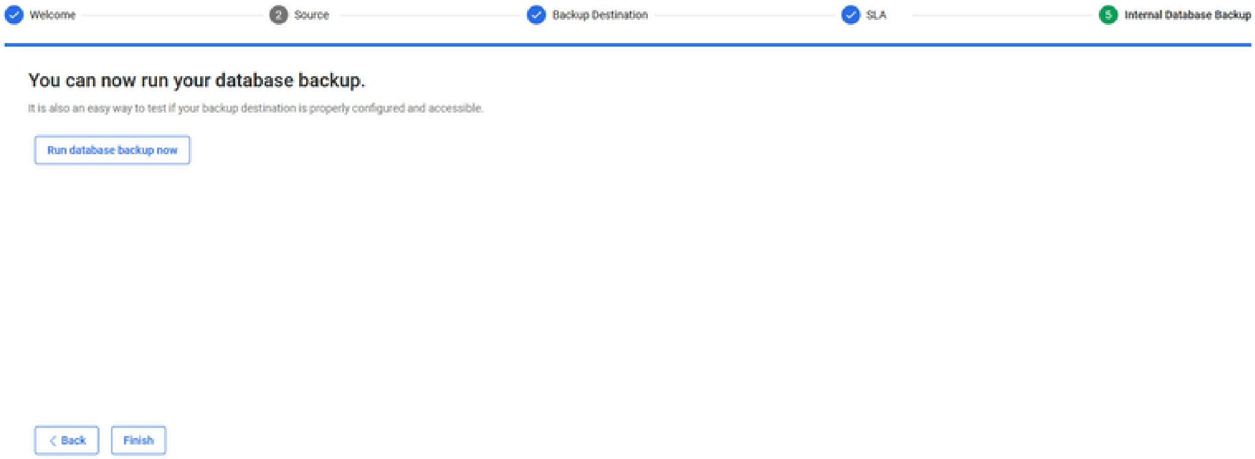

You can now run your database backup.

It is also an easy way to test if your backup destination is properly configured and accessible.

- you are ready to go!

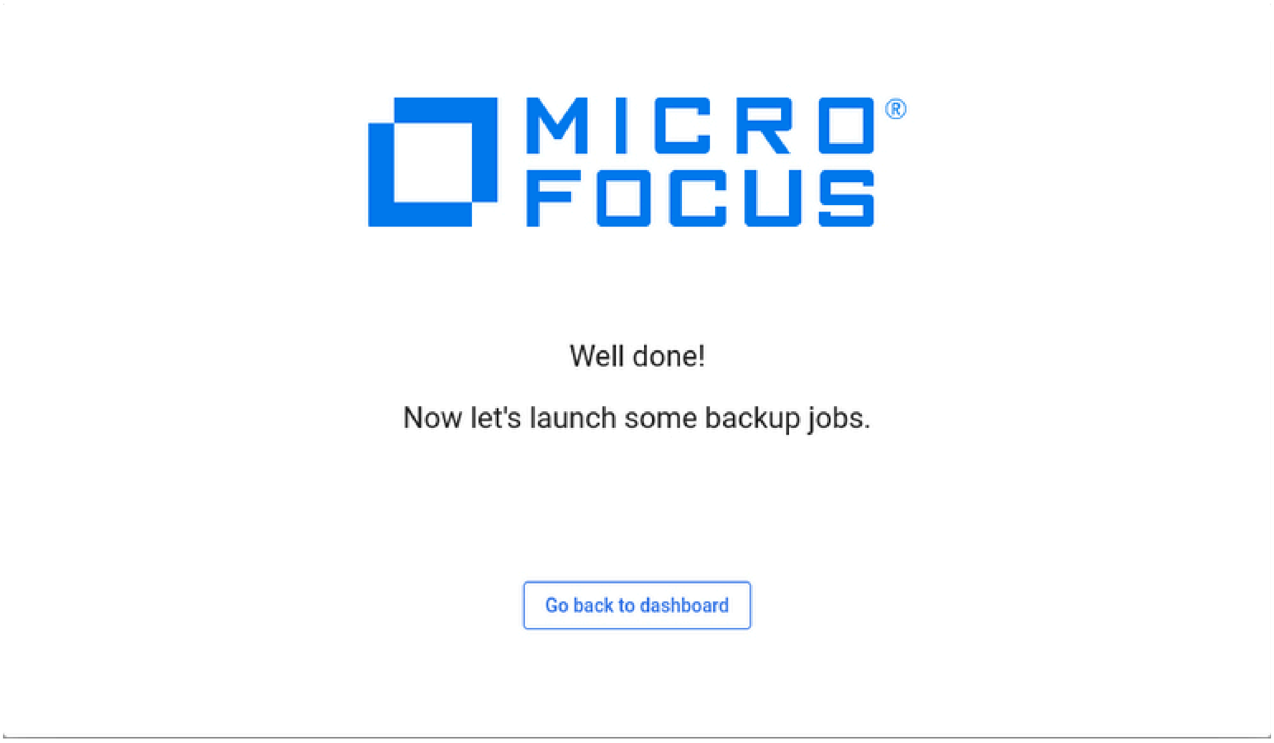

Well done!

Now let's launch some backup jobs.

High Availability

In this scenario, we are going to set up two Data Protector for Cloud Workloads Servers in High Availability, Active/Passive mode. This is possible by using techniques such as a pacemaker, corosync, and DRBD. At least a basic understanding of these is highly desirable. This how-to is intended for RPM-based systems such as Red Hat / CentOS. If you run Data Protector for Cloud Workloads on a different OS, you may need to refer to your distribution docs.

Our environment is built of the following elements:

1. server1 - first Data Protector for Cloud Workloads Server + Data Protector for Cloud Workloads Node, IP: 10.40.1.50
2. server2 - second Data Protector for Cloud Workloads Server + Data Protector for Cloud Workloads Node, IP: 10.40.1.52
3. Cluster IP: 10.40.1.100 - We will use this IP to connect to our **active** Data Protector for Cloud Workloads service. This IP will float between our servers and will point to an active instance.
4. DRBD (optionally with VDO) for data replication and deduplication between nodes.
5. MariaDB master ↔ master replication

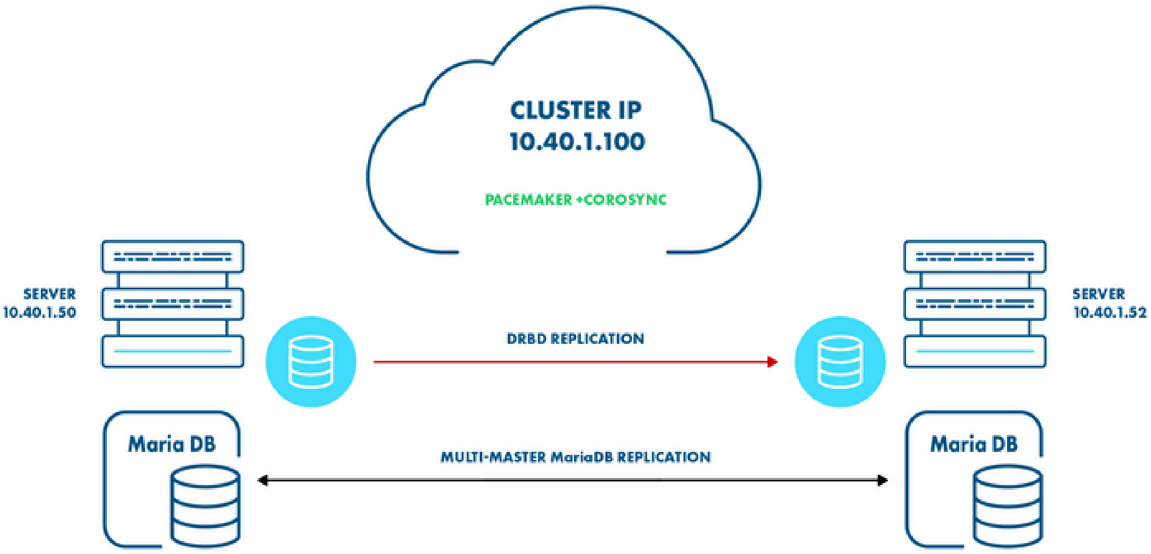

HA cluster setup

Preparing the environment

- Stop and disable the Data Protector for Cloud Workloads Server, node and database as the cluster will manage these resources.

```
systemctl disable vprotect-server vprotect-node mariadb
```

- **Use yum to check if you have any updates pending**

```
# yum update
```

- It is a good idea to check ***/etc/hosts,*** especially if you installed Data Protector for Cloud Workloads using the ***All in one quick installation*** method, as you might find an entry such as:

  ```
  127.0.0.1 <your_hostname_here>
  ```

  Delete it as this prevents the cluster from functioning properly (your nodes will not "see" each other).

Now we can proceed with installation of the required packages.

- **On both servers run**

```
# yum install -y pacemaker pcs psmisc policycoreutils-python
```

- **Add a firewall rule to allow HA traffic** - TCP ports 2224, 3121, and 21064, and UDP port 5405 (both servers)

```
# firewall-cmd --permanent --add-service=high-availability
success
# firewall-cmd --reload
success
```

While testing, depending on your environment, you may encounter problems related to network traffic, permissions, etc. While it might be a good idea to temporarily disable the firewall and SELinux, we do not recommend disabling that mechanism in the production environment as it creates significant security issues. **If you choose to disable the firewall, bear in mind that Data Protector for Cloud Workloads will no longer be available on ports 80/443. Instead, connect to ports 8080/8181 respectively.**

```
# setenforce 0
# sed -i.bak "s/SELINUX=enforcing/SELINUX=permissive/g" /etc/selinux/conf
# systemctl mask firewalld.service
# systemctl stop firewalld.service
# iptables --flush
```

- **Enable and start PCS daemon**

```
# systemctl enable pcsd.service
# systemctl start pcsd.service
```

Cluster configuration

Earlier installation of a pcs package automatically creates a user ***hacluster*** with no password authentication. While this may be good for running locally, we will require a password for this account to perform the rest of the configuration, so let's

- **configure the same password on both nodes**

```
# passwd hacluster
Changing password for user hacluster.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
```

Corosync configuration

- On node 1, issue a command to authenticate as a **hacluster** user:

```
[root@vprotect1 ~]# pcs cluster auth vprotect1 vprotect2
Username: hacluster
Password:
vprotect1: Authorized
vprotect2: Authorized
```

- **Generate and synchronize the corosync configuration**

```
[root@vprotect1 ~]# pcs cluster setup --name mycluster vprotect1 vprotect
```

Take a look at your output, which should look similar to below:

```
Destroying cluster on nodes: vprotect1, vprotect2...
vprotect1: Stopping Cluster (pacemaker)...
vprotect2: Stopping Cluster (pacemaker)...
vprotect1: Successfully destroyed cluster
vprotect2: Successfully destroyed cluster

Sending 'pacemaker_remote authkey' to 'vprotect1', 'vprotect2'
vprotect1: successful distribution of the file 'pacemaker_remote authkey'
vprotect2: successful distribution of the file 'pacemaker_remote authkey'
Sending cluster config files to the nodes...
vprotect1: Succeeded
vprotect2: Succeeded

Synchronizing pcsd certificates on nodes vprotect1, vprotect2...
vprotect1: Success
vprotect2: Success
Restarting pcsd on the nodes in order to reload the certificates...
vprotect1: Success
vprotect2: Success
```

- **Enable and start your new cluster**

```
[root@vprotect1 ~]# pcs cluster start --all && pcs cluster enable --all
vprotect1: Starting Cluster (corosync)...
vprotect2: Starting Cluster (corosync)...
vprotect1: Starting Cluster (pacemaker)...
vprotect2: Starting Cluster (pacemaker)...
vprotect1: Cluster Enabled
vprotect2: Cluster Enabled
```

OK! We have our cluster enabled. We have not created any resources (such as a floating IP) yet, but before we proceed we still have a few settings to modify.

Because we are using only two nodes, we need to

- **disable default quorum policy**

(this command should not return any output)

```
[root@vprotect1 ~]# pcs property set no-quorum-policy=ignore
```

We should also

- **define default failure settings**

```
[root@vprotect1 ~]# pcs resource defaults failure-timeout=30s
[root@vprotect1 ~]# pcs resource defaults migration-threshold=3
```

These two settings combined will define how many failures can occur for a node to be marked as ineligible for hosting a resource and after what time this restriction will be lifted. We define the defaults here, but it may be a good idea to also set these values at the resource level, depending on your experience.

As long we are not using any fencing device in our environment (and here we are not) we need to:

- **disable stonith**

```
[root@vprotect1 ~]# pcs property set stonith-enabled=false && crm_verify
```

The second part of this command verifies running-config. These commands normally do not return any output.

Resource creation

Finally, we have our cluster configured, so it's time to proceed to

- **resource creation**

First, we will create a resource that represents our ***floating IP*** 10.40.1.100. Adjust your IP and cidr_netmask, and you're good to go.

IMPORTANT: From this moment on we need to use this IP when connecting to our Data Protector for Cloud Workloads Server.

```
[root@vprotect1 ~]# pcs resource create "Failover_IP" ocf:heartbeat:IPadd
```

Immediately, we should see our IP is up and running on one of the nodes (most likely on the one we issued this command for).

```
[root@vprotect1 ~]# ip a
[..]
2: ens160:  mtu 1500 qdisc mq state UP group default qlen 1000
    link/ether 00:50:56:a6:9f:c6 brd ff:ff:ff:ff:ff:ff
    inet 10.40.1.50/22 brd 10.40.3.255 scope global ens160
       valid_lft forever preferred_lft forever
    inet 10.40.1.100/22 brd 10.40.3.255 scope global secondary ens160
       valid_lft forever preferred_lft forever
    inet6 fe80::250:56ff:fea6:9fc6/64 scope link
       valid_lft forever preferred_lft forever
```

As you can see, our floating IP 10.40.1.100 has been successfully assigned as the second IP of interface ens160. This is what we wanted!

We should also check if the Data Protector for Cloud Workloads web interface is up and running. We can do this by opening the web browser and typing in https://<DP4CW_server_IP>. At this point we should see:

DATA PROTECTOR
FOR CLOUD WORKLOADS
User name
Password
LOGIN

The next step is to

- **define a resource responsible for monitoring network connectivity**

```
[root@vprotect1 ~]# pcs resource create ping ocf:pacemaker:ping dampen=5s
[root@vprotect1 ~]# pcs constraint location Failover_IP rule score=-INFIN
```

Note that you need to use **your gateway IP** in the ***host_list*** parameter

Finally, we have to define a set of cluster resources responsible for other services crucial for Data Protector for Cloud Workloads as Data Protector for Cloud Workloads Node and the Data Protector for Cloud Workloads Server itself. We will logically link these services with our floating IP. Whenever the floating IP disappears from our server, these services will be stopped. We also have to define the proper order for services to start and stop, as for example starting the Data Protector for Cloud Workloads-server without a running database makes little sense.

- **Resource creation**

```
[root@vprotect1 ~]#  pcs resource create "vProtect-node" systemd:vprotect
[root@vprotect1 ~]# pcs resource create "vProtect-server" service:vprotec
```

It is OK for these commands not to return any output.

- **Resource colocation**

```
[root@vprotect1 ~]# pcs constraint colocation add Failover_IP with vProte
```

To finish with, we can set which server is more preferred for running our services

- **Set node preference**

```
[root@vprotect1 ~]# pcs constraint location Failover_IP prefers vprotect1
[root@vprotect1 ~]# pcs constraint location vProtect-group prefers vprote
```

We have made it to the end. At this point, our pacemaker HA cluster is functional.

However, there are still two things we need to consider, that is:

1. Creating DB replication
2. Setting up DRBD for /vprotect_data (optionally with VDO)

Setting up VDO+DRBD

In this section, we will prepare our deduplicated and replicated filesystem mounted in /vprotect_data.

Using a deduplicated FS is optional but highly recommended. If you don't intend to use it, skip the part regarding VDO configuration.

Note: If you are altering existing Data Protector for Cloud Workloads configuration it is very important to preserve the /vprotect_data contents and transfer them to the new filesystem. You may also need to re-create your backup_destination if you previously had one in this directory. Setting up VDO and DRBD will cause all data to be wiped from the configured volume.

Installation is split into the steps below that you need to follow to get the job done.

- **Stop the Data Protector for Cloud Workloads Server and node**

```
# systemctl stop vprotect-server vprotect-node
```

No output means everything went OK.

- **On both nodes install the equired repositories and packages**

```
# rpm --import https://www.elrepo.org/RPM-GPG-KEY-elrepo.org
# rpm -Uvh https://www.elrepo.org/elrepo-release-7.0-4.el7.elrepo.noarch.
Retrieving https://www.elrepo.org/elrepo-release-7.0-4.el7.elrepo.noarch.
Preparing...                          ################################# [
Updating / installing...
   1:elrepo-release-7.0-4.el7.elrepo  ################################# [
```

The next command can produce quite a few lines, so I've truncated the output, however the idea is simple: install drbd packages:

```
[root@vprotect1 ~]# yum install -y kmod-drbd84 drbd84-utils

Installed:
drbd84-utils.x86_64 0:9.6.0-1.el7.elrepo
```

If you have not disabled SELinux and the firewall, remember to

- **configure them on both nodes**

  ```
  # semanage permissive -a drbd_t
  # firewall-cmd --add-port=7788/tcp --permanent
  success
  # firewall-cmd --complete-reload
  success
  ```

  Don't forget to repeat these steps on the second node

Now that we have the necessary software installed, we must prepare an identical size block device on both nodes. A block device can be a hard drive, a hard drive partition, software RAID, LVM Volume, etc. In this scenario, we are going to use a hard drive connected as ***/dev/sdb***.

To add a DRBD resource we create the file ***/etc/drbd.d/vprotect.res*** with the content below. Be sure to change the "address" so that t reflects your network configuration.

Also, the node names (server1 and server2) must match your ***uname -n*** output.

```
resource replicate {
protocol C;
    on vprotect1 {
                device /dev/drbd0;
                disk /dev/sdb;
                address 10.40.1.50:7788;
                meta-disk internal;
        }
    on vprotect2 {
                device /dev/drbd0;
                disk /dev/sdb;
                address 10.40.1.52:7788;
                meta-disk internal;
        }
```

We now have config in place and can create and bring our resource online.

- **On both nodes, run**

  ```
  # drbdadm create-md replicate
  initializing activity log
  initializing bitmap (4800 KB) to all zero
  Writing meta data...
  New drbd meta data block successfully created.
  ```

  then bring the volume online

  ```
  # drbdadm up replicate
  ```

  You can verify if the device is up & running by issuing

```
# lsblk
NAME                    MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda                       8:0    0   16G  0 disk
├─sda1                    8:1    0    1G  0 part /boot
└─sda2                    8:2    0   15G  0 part
├─vg_vprotect-lv_root 253:0    0 13.4G  0 lvm  /
└─vg_vprotect-lv_swap 253:1    0  1.6G  0 lvm  [SWAP]
sdb                       8:16   0  150G  0 disk
└─drbd0                 147:0    0  150G  1 disk
```

However, if we check

```
[root@vprotect1 ~]# drbdsetup status replicate
replicate role:Secondary
disk:Inconsistent
peer role:Secondary
replication:Established peer-disk:Inconsistent
```

we will notice we need to start synchronization before we can use our volume.

- **On the first server, run**

```
[root@vprotect1 ~]# drbdadm primary --force replicate
[root@vprotect1 ~]# drbdsetup status replicate
replicate role:Primary
disk:UpToDate
peer role:Secondary
replication:SyncSource peer-disk:Inconsistent done:0.22
```

This way we have successfully started the process of replication between servers with vprotect1 as the ynchronization source.

If you don't want to create a VDO device, then create and mount your filesystem:

```
[root@vprotect1 ~]# mkfs.xfs -K /dev/drbd0
[root@vprotect1 ~]# mount /dev/mapper/drbd0 /vprotect_data/ && chown -
```

- **Create VDO volume** (optional)

By issuing the command below we will create a VDO volume called ***vdo_data*** and put in at the top our DRBD volume. Afterwards, we format it with XFS and mount it in /vprotect_data.

```
[root@vprotect1 ~]# vdo create --name=vdo_data --device=/dev/drbd0 --v
Creating VDO vdo_data
Starting VDO vdo_data
Starting compression on VDO vdo_data
VDO instance 0 volume is ready at /dev/mapper/vdo_data

[root@vprotect1 ~]# mkfs.xfs -K /dev/mapper/vdo_data
meta-data=/dev/mapper/vdo_data   isize=512    agcount=4, agsize=262144
    =                       sectsz=4096  attr=2, projid32bit=1
    =                       crc=1        finobt=0, sparse=0
data     =                       bsize=4096   blocks=104857600, imaxpc
    =                       sunit=0      swidth=0 blks
naming   =version 2              bsize=4096   ascii-ci=0 ftype=1
log      =internal log           bsize=4096   blocks=51200, version=2
    =                       sectsz=4096  sunit=1 blks, lazy-count=1
realtime =none                   extsz=4096   blocks=0, rtextents=0

[root@vprotect1 ~]# mount /dev/mapper/vdo_data /vprotect_data/ && chow
```

- **Copy the VDO config to the second node**

  ```
  [root@vprotect1 ~]# scp /etc/vdoconf.yml root@vprotect2:/etc/vdoconf.y
  ```

- **Disable VDO automatic startup**

  As this resource will be managed by the cluster, we need to disable auto startup of this service ***on both nodes.***

  ```
  # systemctl disable vdo
  ```

Final cluster settings

At this point, we have three components set up. To fully utilize our HAcluster and eliminate the need for manual intervention we should add the resources and settings below to our cluster.

Issue these commands on one node only as it will propagate to the cluster settings.

```
[root@vprotect1 ~]#  pcs cluster cib drbd_cfg
[root@vprotect1 ~]#  pcs -f drbd_cfg resource create replicate ocf:linbit
         drbd_resource=replicate op monitor interval=10s --group fs_group

[root@vprotect1 ~]#  pcs -f drbd_cfg resource master replicateClone repli
         master-max=1 master-node-max=1 clone-max=2 clone-node-max=1 \
         notify=true --group fs_group

[root@vprotect1 ~]#  pcs -f drbd_cfg resource create vdo_resource ocf:hea
[root@vprotect1 ~]#  pcs -f drbd_cfg resource create fs_resource ocf:hear
[root@vprotect1 ~]#  pcs cluster cib-push drbd_cfg --config

[root@vprotect1 ~]#  pcs constraint colocation add vdo_resource with repl
[root@vprotect1 ~]#  pcs constraint order start vdo_resource then fs_reso
[root@vprotect1 ~]#  pcs constraint order start replicateClone then vdo_r
[root@vprotect1 ~]#  pcs constraint colocation add vProtect-group with fs
[root@vprotect1 ~]#  pcs constraint colocation add vdo_resource with repl
[root@vprotect1 ~]#  pcs constraint order promote replicateClone then sta
```

Here we have created a temporary file ***drbd_cfg*** and inside this file we have added our drbd_resource called ***replicate***, plus a Master/Slave set for this resource.

Afterwards, we have the definition of the vdo_resource and fs_resource in one fs_group followed by an update of the cluster configuration.

As a second step, we have put in place several resource colocations and constraints which allow us to control the order and existence of newly created resources.

We need still to

- Make sure that our node is pointed to a localhost address. Check the ***Nodes*** UI section.

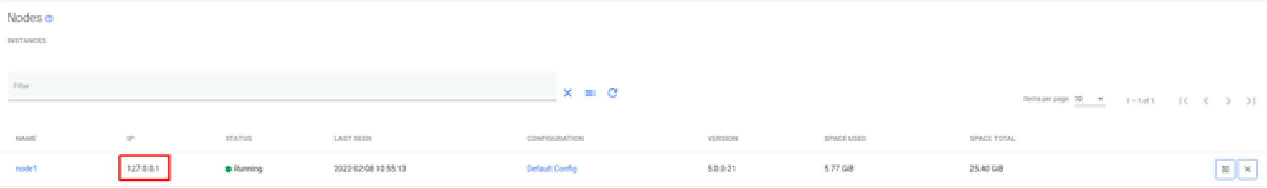

NAME	IP	STATUS	LAST SEEN	CONFIGURATION	VERSION	SPACE USED	SPACE TOTAL
node1	127.0.0.1	● Running	2022-02-08 10:55:13	Default Config	5.0.0-21	5.77 GiB	25.40 GiB

If the node's IP is different than 127.0.0.1, delete the node and re-register it using

```
[root@vprotect1 ~]# vprotect node -e <Node_Name> admin http://127.0.0.1:8
```

- copy our license and node information from the first node to the second node:

```
[root@vprotect1 ~]# scp -pr /opt/vprotect/.session.properties
[root@vprotect1 ~]# scp -pr /opt/vprotect/license.key
```

MariaDB replication

In this section, we will cover how to setup master↔master MariaDB replication.

- On both nodes, if you have the firewall enabled, allow communication via port **3306**

```
# firewall-cmd --add-port=3306/tcp --permanent
# firewall-cmd --complete-reload
```

Steps to run on the first server1 node: 10.40.1.50

This server will be the source of DB replication.

- **Stop the Data Protector for Cloud Workloads Server, node and database**

```
[root@vprotect1 ~]# systemctl stop vprotect-server vprotect-node mariadb
```

- **Edit the config file**, enable binary logging and start MariaDB again. Depending on your distribution, the config file location may vary, most likely it is /etc/my.cnf or /etc/my.cnf.d/server.cnf

  In the ***[mysqld]*** section, add the lines:

```
[root@vprotect1 ~]# vi /etc/my.cnf.d/server.cnf
log-bin
server_id=1
replicate-do-db=vprotect
[root@vprotect1 ~]# systemctl start mariadb
```

- Now **log in into your MariaDB**, create a user used for replication and assign appropriate rights to it.

  For the purpose of this task, we will set the username to 'replicator' and the password to 'R3pLic4ti0N'

```
[root@vprotect1 ~]# mysql -u root -p
Enter password:
[..]
MariaDB [(none)]> create user 'replicator'@'%' identified by 'R3pLic4ti0N
Query OK, 0 rows affected (0.026 sec)

MariaDB [(none)]> grant replication slave on *.* to 'replicator'@'%';
Query OK, 0 rows affected (0.001 sec)

MariaDB [(none)]> FLUSH PRIVILEGES;
Query OK, 0 rows affected (0.001 sec)
```

Don't log out just yet, we need to check the master status and

- **write down the log file name and position**, as it is required for proper slave configuration.

```
MariaDB [(none)]> show master status;
+----------------------+----------+--------------+------------------+
| File                 | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+----------------------+----------+--------------+------------------+
| vprotect1-bin.000007 |    46109 |              |                  |
+----------------------+----------+--------------+------------------+
```

- Dump the Data Protector for Cloud Workloads database and copy it onto the second server (vprotect2).

```
[root@vprotect1 ~]# mysqldump -u root -p vprotect > /tmp/vprotect.sql
[root@vprotect1 ~]# scp /tmp/vprotect_rep.sql root@vprotect2:/tmp/
```

Steps to run on the 2nd server, server2: 10.40.1.52

For the reader's convenience, I have only highlighted the differences in configuration between server1 and server2, and omitted the output of some commands if they are the same as on the previous node.

- **Stop the Data Protector for Cloud Workloads Server, Node and database**
- Edit the MariaDB config file. **Assign a different server id**, for example: 2. Then start MariaDB.

```
[root@vprotect2 ~]# vi /etc/my.cnf.d/server.cnf
log-bin
server_id=2
replicate-do-db=vprotect
[root@vprotect2 ~]# systemctl start mariadb
```

- **Load the database dump** copied from server1.

```
[root@vprotect2 ~]# mysql -u root -p vprotect < /tmp/vprotect.sql
```

At this point, we have two identical databases on our two servers.

- **Log in to the MariaDB instance, create a replication user with a password**. Use the same user as on server1. Grant the necessary permissions.
- Set the master host. You ***must*** use the user_master_log_file and master_log_pos written down earlier. Change the IP of the master host to match your network configuration.

```
MariaDB [(none)]> STOP SLAVE;
MariaDB [(none)]> CHANGE MASTER TO MASTER_HOST = '10.40.10.50', MASTER_US
Query OK, 0 rows affected (0.004 sec)
```

- Start the slave, check the master status and **write down the file name and position.**

```
MariaDB [(none)]> start slave;
Query OK, 0 rows affected (0.001 sec)

MariaDB [(none)]> SHOW MASTER STATUS;
+----------------------+----------+--------------+------------------+
| File                 | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+----------------------+----------+--------------+------------------+
| vprotect2-bin.000002 |   501051 |              |                  |
+----------------------+----------+--------------+------------------+
1 row in set (0.000 sec)
```

Go back to the first server (server1)

- On **storreaw1**, stop the slave then change the master host using the parameters noted down in the previous step. Also, change the master host IP to match your network configuration.

```
MariaDB [(none)]> stop slave;
MariaDB [(none)]> MariaDB [(none)]>  change master to master_host='10.40.
Query OK, 0 rows affected (0.004 sec)
MariaDB [(none)]> start slave;
Query OK, 0 rows affected (0.001 sec)
```

At this point, you have successfully configured MariaDB master↔master replication.

Testing the setup

Automatic

The fastest way to test our setup is to invoke

```
# pcs node standby vprotect1
```

to put server1 into standby mode, which prevents it from hosting any cluster resources.

After a while, you should see your resources up and running on server2.

Note that if you perform normal OS shutdown (not a forced one), the pacemaker will wait for a long time for a node to come back online, which in fact will prevent completion of shutdown. As a result, resources ***will not*** switch correctly to the other node.

Manual

If you want to dive a little bit deeper, we have prepared instructions on how to manually move a filesystem resource from the first node to the second.

1. Stop vprotect services.

```
systemctl stop vprotect-server && systemctl stop vprotect-node
```

2. Unmount the FS used by DRBD/VDO on the primary server (here server1).

```
[root@vprotect1 ~]# drbdadm role replicate
Primary/Secondary
[root@vprotect1 ~]# umount /vprotect_data/
```

3. If you are using a VDO device, stop it.

```
[root@vprotect1 ~]# vdo stop -n vdo_data
Stopping VDO vdo_data
```

4. Demote the primary replication server (still server1) to secondary server.

```
[root@vprotect1 ~]# drbdadm secondary replicate
```

On the second server

1. Promote the second server (here server2) to the primary DRBD role.

```
[root@vprotect2 ~]# drbdadm    primary replicate
```

2. Start the VDO.

```
[root@vprotect2 ~]# vdo start -n vdo_data
Starting VDO vdo_data
Starting compression on VDO vdo_data
VDO instance 2 volume is ready at /dev/mapper/vdo_data
```

3. Mount the filesystem on the second server.

```
[root@vprotect2 ~]# mount /dev/mapper/vdo_data /vprotect_data/
```

Now you have your replicated volume mounted on the second node.

Common tasks

Common tasks

This section presents several supplementary tasks that may be needed in Data Protector for Cloud Workloads deployment. This includes tasks such as HTTPS setup, SSH public key authentication with your hypervisors, VMs or libvirt/qemu package installation.

Staging space configuration

General

Data Protector for Cloud Workloads Node needs staging space available in `/vprotect_data` by default. It is common to use PowerProtect DD for both the staging and backup destination. This will result in instant "store" processing, without the need to copy data from the staging space to the backup destinations. It is common to just attach an empty drive and mount it.

When using separate storage (usually local disks) for the staging space, consider its requirements. Staging space size depends on the number and size of simultaneous backups - as a rule of thumb make it approximately equal to the number of expected simultaneous backup threads multiplied by the size of your biggest VM.

In any case - make sure the staging space is always mounted in the `/vprotect_data` folder, and that the vprotect user is able to have full permissions to this file system.

Example - Local filesystem

You also can use a plain file system for staging space (and optionally for backup destination). Here are steps assuming you have a local (physical or virtual) disk.

- List all existing disks, and find your dedicated disk (let's say - `/dev/sdc`):

```
[root@vProtect01 ~]# fdisk -l | grep dev
Disk /dev/sda: 32.2 GB, 32212254720 bytes, 62914560 sectors
/dev/sda1   *        2048     1026047      512000   83  Linux
/dev/sda2         1026048    62914559    30944256   8e  Linux LVM
Disk /dev/sdc: 500 GB, 17179869184 bytes, 33554432 sectors
Disk /dev/sdb: 21.5 GB, 21474836480 bytes, 41943040 sectors
Disk /dev/mapper/centos-root: 28.5 GB, 28462546944 bytes, 55590912 sector
Disk /dev/mapper/centos-swap: 3221 MB, 3221225472 bytes, 6291456 sectors
```

- If you have a new clean disk prepare a filesystem on it:

```
mkfs.xfs -K /dev/sdc
```

- Test mount your existing filesystem in the created directory:

```
mount /dev/sdc /vprotect_data
```

- Set ownership to `vprotect` user on directory `/vprotect_data`:

```
chown vprotect:vprotect -R /vprotect_data
```

- Add a line to `/etc/fstab`file, to automatically mount new filesystem after reboot:

```
/dev/sdc    /vprotect_data    xfs    defaults 0 0
```

- Mount

```
mount -a
```

- Confirm with `df` that your `/vprotect_data` is mounted
- Restart your `vprotect-node`service:

```
systemctl restart vprotect-node
```

Enabling HTTPS connectivity for nodes

The default certificate presented by the application server uses `localhost.localdomain`. This works only for local node installations (server and node on a single host).

> **Note:**
>
> - You can use the default certificate - remember that you may need to use the `./node_add_ssl_cert.sh` script after future updates to refresh the certificate on the node
> - For the default certificate - jump to the Node configuration and use the localhost.localdomain instead of the `dp4cw.local`
> - When registering the node locally over HTTPS note that the URL you should use is `localhost.localdomain` - **NOT** `localhost`
> - When registering a node via HTTPS, please note that the server must have an FQDN that is different from the IP address (hostname like `10.10.10.10` can be processed incorrectly).

This section presents the steps necessary for generating an SSL certificate, for setup Data Protector for Cloud Workloads to use it and how to register a remote node.

Data Protector for Cloud Workloads Server (when using own certificate)

This section describes certificate generation and import on the Data Protector for Cloud Workloads Server side. It uses a self-signed certificate. If you would like to use CSR and your own CA instead - check for additional steps described in the next section.

1. SSH to Data Protector for Cloud Workloads Server host

2. Generate the key and certificate (remember to provide a valid DP4CW Server DNS hostname - in our example it was dp4cw.local):

```
[root@dp4cw.local ~]# openssl req -x509 -newkey rsa:4096 -keyout dp4cw
Generating a 4096 bit RSA private key
......................................................................
......................................................................
writing new private key to 'dp4cw.key'
Enter PEM pass phrase:
Verifying - Enter PEM pass phrase:
-----
You are about to be asked to enter information that will be incorporat
into your certificate request.
What you are about to enter is what is called a Distinguished Name or
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [XX]:PL
State or Province Name (full name) []:
Locality Name (eg, city) [Default City]:Warsaw
Organization Name (eg, company) [Default Company Ltd]: your Company
Organizational Unit Name (eg, section) []:
Common Name (eg, your name or your server's hostname) []:dp4cw.local
Email Address []:
```

3. Create the PKCS12 bundle from the certificate and the key:

```
[root@localhost ~]# openssl pkcs12 -export -in dp4cw.crt -inkey dp4cw.
Enter pass phrase for dp4cw.key:
Enter Export Password:
Verifying - Enter Export Password:
```

4. Create a keystore for the Data Protector for Cloud Workloads Server with the PKCS12 bundle:

```
[root@localhost ~]# keytool -importkeystore -destkeystore /opt/vprotec
Enter destination keystore password:
Re-enter new password:
Enter source keystore password:
```

5. Change ownership on the keystore to the `vprotect` user:

```
chown vprotect:vprotect /opt/vprotect/keystore.jks
```

6. Edit `/opt/vprotect/payara.properties`, change the path to the keystore and password (use password generated in step 3 of that instruction):

```
javax.net.ssl.keyStore=/opt/vprotect/keystore.jks
javax.net.ssl.keyStorePassword=[keystorepassword]
```

7. Restart the Server:

   ```
   systemctl stop vprotect-server
   systemctl start vprotect-server
   ```

Data Protector for Cloud Workloads Node (any SSL certificate)

1. SSH to Data Protector for Cloud Workloads Node host
2. Make sure that your nodes resolve the hostname (FQDN) of the Data Protector for Cloud Workloads Server. You also can add an entry in the `/etc/hosts` like this (example IP: 1.2.3.4):

   ```
   1.2.3.4 dp4cw.local
   ```

3. Check with your browser that `https://DP4CW_HOST:8181` presents the certificate that you have just generated. You also can execute the openssl client from the node to print it (check the hostname that you have provided in the certificate):

   ```
   openssl s_client -connect dp4cw.local:8181 < /dev/null
   ```

4. Import the server certificate using the script under the /opt/vprotect/scripts folder:

   ```
   cd /opt/vprotect/scripts
   ./node_add_ssl_cert.sh [SERVER_HOST] [PORT] [KEYSTORE_PASS]
   ```

   - [SERVER_HOST] - FQDN name of Data Protector for Cloud Workloads Server
   - [PORT] - port for SSL communication on Data Protector for Cloud Workloads Server (you need to open it on server `# firewall-cmd --permanent --add-port=[PORT]/tcp && firewall-cmd --reload`)
   - [KEYSTORE_PASS] - password which you defined in step 3 of that instruction

   > **Note:**

> If you have node on the same host as server, You could use default variables of script (and you can use script without arguments). Default variables are:
>
> - SERVER_HOST = `127.0.0.1`
> - PORT = `8181`
> - KEYSTORE_PASS = `changeit`
>
> It applies if you would not generated any certificate.

5. Register the node with the NODE_NAME of your choice, the ADMIN_USER user name which you would like to use and the URL to Data Protector for Cloud Workloads API, and provide the password when prompted:

   ```
   vprotect node -r NODE_NAME ADMIN_USER http(s)://DP4CW_SERVER:PORT/api
   ```

   > **Examples:**
   >
   > - Remote server with a generated certificate:
   >
   > ```
   > vprotect node -r node1 admin https://dp4cw.local:8181/api`
   > ```
   >
   > - Local installation with default certificate:
   >
   > ```
   > vprotect node -r node1 admin https://localhost.localdomain:8181/api
   > ```

Notes on using your own certificate with CSR and your own CA

When using CSR to get a trusted certificate, you need to replace step 2 in Data Protector for Cloud Workloads Server (when using own certificate) with several steps including CSR generation, and download the CRT signed by your CA. The steps are as follows:

1. Generate the CSR - answer the same set of questions as above:openssl req -new -newkey rsa:2048 -nodes -keyout dp4cw.key -out dp4cw.csr.
2. Send your CSR and have it signed by your CA.

3. Download your CRT file and save it as dp4cw.crt (note that you should have your working directory set to `/opt/vprotect`).
4. Download your CA certificate chain (for example for a singleca.crt) and import it with the CA_ALIAS of your choice as follows:

   ```
   keytool -import -trustcacerts -keystore /usr/lib/jvm/jre/lib/security/
   ```

5. Now continue from PKCS12 bundle generation (step 3 in the section above).

LVM setup on Data Protector for Cloud Workloads Node for disk attachment backup mode

> **Note:** This is required for backup of virtual environments when using disk attachment mode, such as Nutanix backups.

Data Protector for Cloud Workloads Node attaches VM disks that potentially are clones of its own (for example if Node deployed from the template) - you need to configure LVM on the Node so that it doesn't scan for LVM volumes where disks are being attached.

1. Set the following variables in `/etc/lvm/lvm.conf` in `devices` section - so that only system volumes are being detected by LVM daemon (in this example sda disk with 2 partitions - sda1 and sda2):

   ```
   devices {
            filter = [ "a|^/dev/sda|", "a|^/dev/sda1|", "a|^/dev/sda2|",
            global_filter = [ "a|^/dev/sda|", "a|^/dev/sda1|", "a|^/dev/s
   }
   ```

2. Check with `vgscan -vvv` that your OS volumes are still being detected:

   ```
   Allocated VG vg_vprotect at 0x55914f19fac0.
   Importing logical volume vg_vprotect/lv_root.
   Importing logical volume vg_vprotect/lv_swap.
   ```

3. Reboot:

   ```
   reboot
   ```

Full versions of libvirt/qemu packages installation

Make sure that your `libvirt` supports the `virsh blockcommit` operation. CentOS distribution requires you to install the full `libvirt` and `qemu-img` from the `oVirt` repository. This can be done like this:

1. Install oVirt repo:

```
yum install http://resources.ovirt.org/pub/yum-repo/ovirt-release42.rpm -
```

2. Update the packages

```
yum update -y
```

which should replace `qemu` related packages with full versions from the oVirt repo.

SSH public key authentication

General

Instead of using password authentication - anywhere where you're able to provide SSH credentials (hypervisors, VMs applications, etc) you also have the public key alternative.**.
By default, Data Protector for Cloud Workloads uses the `/opt/vprotect/.ssh/id_rsa` path, however, you also can override it with your own path*.
***(this needs to be owned by `vprotect` user and make sure it has the `0400` permission set.**
*****You don't have to pass a passphrase, you can leave this parameter blank.***

> **Note:** Data Protector for Cloud Workloads does not support keys other than "RSA"

Example

1. Generate a key or use yours and store it as `/opt/vprotect/.ssh/id_rsa` (make sure that the `vprotect` user and group own the file)

- example key generation:

```
[root@vProtect3 vprotect]# sudo -u vprotect ssh-keygen -t rsa -m PEM
Generating public/private rsa key pair.
Enter file in which to save the key (/opt/vprotect/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /opt/vprotect/.ssh/id_rsa.
Your public key has been saved in /opt/vprotect/.ssh/id_rsa.pub.
The key fingerprint is:
SHA256:86HSLKYwl7maDR7U1oIH1Y6VDtRFNJgHgfdjikg3VnQ vprotect@vProtect3
The key's randomart image is:
+---[RSA 2048]----+
|   .o=+XE        |
|   .o X...       |
|  .  O o         |
|  .+=.o +        |
| .o+=o.oS..      |
| ..o.+.o + .     |
|  = + + + .      |
| . O + o         |
|  +.+            |
+----[SHA256]-----+
```

2. Use `ssh-copy-id` to upload your public key (as `vprotect` user) to the KVM host:

```
sudo -u vprotect ssh-copy-id -i /opt/vprotect/.ssh/id_rsa.pub root@HYPERV
```

3. Check if you're able to log in to the hypervisor using the local `vprotect` user without being asked for the password:

```
[root@vProtect3]# sudo -u vprotect ssh -i /opt/vprotect/.ssh/id_rsa root@
Last failed login: Mon Jan 29 17:53:01 CET 2018 from 10.50.1.107 on ssh:n
There was 1 failed login attempt since the last successful login.
Last login: Mon Jan 29 17:52:39 2018 from 10.50.1.107
[root@dKVM ~]# logout
```

4. Now you should be able to index VMs regardless of the password set for the hypervisor (the key should be used instead)

5. Provide path to key (default: /opt/vprotect/.ssh/id_rsa) in Data Protector for Cloud Workloads dashboard

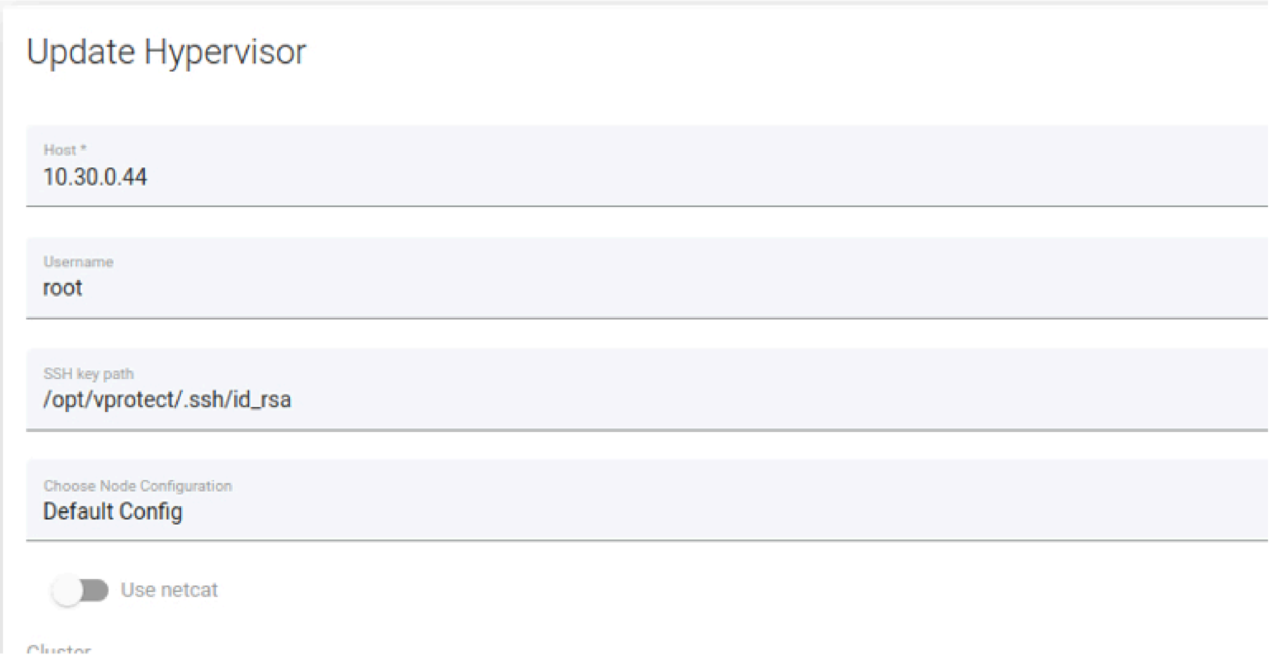

Update Hypervisor

Host *
10.30.0.44

Username
root

SSH key path
/opt/vprotect/.ssh/id_rsa

Choose Node Configuration
Default Config

○ Use netcat

Cluster

Enabling HTTP(S) Proxy for Data Protector for Cloud Workloads

You can configure the system to communicate through an HTTP(S) proxy. You can configure the `HTTP_PROXY` and `HTTPS_PROXY` environment variables using the `vprotect.env` file.

1. Edit the vprotect.env file that is located in `/opt/vprotect/vprotect.env`. Uncomment the following lines and specify the correct proxy address:

   ```
   http_proxy="proxy.address:8080"
   https_proxy="proxy.adress:8080"
   no_proxy="localhost,127.0.0.1"
   ```

   > Make sure to change proxy.address to the address of your proxy, which can be either IP address or FQDN.

2. Restart the Data Protector for Cloud Workloads Node and Server to apply the changes.

   ```
   systemctl restart vprotect-node vprotect-server
   ```

Repeat above steps for each host where the Server and/or Node is installed.

Protecting Virtual Environments

Protecting Virtual Environments

Data Protector for Cloud Workloads supports multiple on-premise virtualization platforms. In this section, you will find what backup methods are supported and the specific steps that are needed for each of them to be integrated with Data Protector for Cloud Workloads.

- Virtual Machines
- Cloud
- Containers
- Backup & Restore

Virtual Machines

Protecting virtual environments

In this chapter, You will know how to add and protect your Virtual Machines such as:

- Nutanix Acropolis Hypervisor (AHV)
- Red Hat Openshift Virtualization
- Red Hat Virtualization
- oVirt
- Oracle Linux Virtualization Manager
- Oracle VM
- Proxmox VE
- KVM/Xen
- OpenStack
- OpenNebula
- Virtuozzo
- Citrix Hypervisor (XenServer)
- XCP-ng
- Huawei FusionCompute
- SC//Platform

Nutanix Acropolis Hypervisor (AHV)

Nutanix Acropolis Hypervisor (AHV)

General

Data Protector for Cloud Workloads supports the Nutanix AHV platform by using a VM called “Proxy VM”. The node invokes commands on your hypervisor manager to snapshot and attach drives of a specific VM to itself (Proxy VM). The proxy VM is able to read the data from the attached disk snapshots and forward them to the backup provider.

This strategy allows you to exclude drives from a backup that you do not need. Remember that you need to install at least 1 Proxy VM per cluster so that the drives the node tries to attach are reachable.

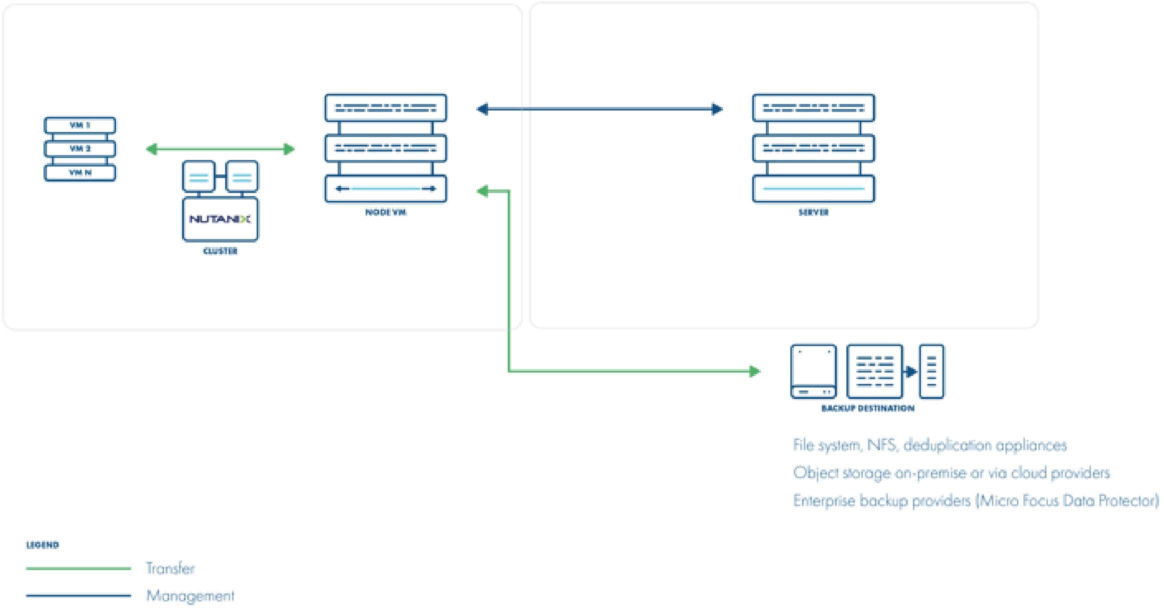
VM 1
VM 2
VM N
NUTANIX
CLUSTER
NODE VM
SERVER
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot using hypervisor's API

- optionally application-consistent snapshot can be done if enabled and guest tools installed inside - the type of snapshot is selected based on is QuiesceBeforeSnapshot setting and passed as part of the snap request. The created snapshot might end up being of a different type (depending on the presence of tools
- optional application consistency using pre/post snapshot command execution
- metadata exported from API
- snapshot disks are mounted one by one to the Proxy VM
- data read directly on the Proxy VM
- incremental backups using CBT API - only changed blocks are read from the attached disk
- the last snapshot kept on the hypervisor for the next incremental backup (if at least one schedule assigned to the VM has the backup type set to incremental)
- restore creates empty disks on the Proxy VM, imports merged data then recreates VM and reattaches volumes to the target VM

A general explanation of "The dynamically attached disks slot offset" parameter for Data Protector for Cloud Workloads Node proxy VM

Our best practice is to use a proxy machine with one disk device for the purposes of the operating system if you are using the "Disk attachment" backup strategy. Due to the simplification of the configuration of the environment, we also do not achieve any benefits for this element of the environment.

Our experience shows that after adding a new node to the environment, is good to perform a test backup and check the logs from which disk device Data Protector for Cloud Workloads Node want to start the backup. Depending on the proxy virtual machine configuration, Data Protector for Cloud Workloads will select the appropriate disk or you need to manually set the offset parameter. Rather, we do not encounter this type of situation when a virtual machine has only one disk device.

Recommendations on how to set up the environment for Data Protector for Cloud Workloads

- As the backup strategy for the Nutanix environment depends on attaching and detaching disk devices to Proxy VM, we recommend simplifying the hardware configuration of this machine. If your backup destination allows having staging space on the same storage as the backup destination, one disk device should be sufficient for the proxy virtual machine's operating system purposes.
- If it is not possible to have only one disk device for Proxy VM, read the Example section. We explained what you need to do to make sure your Data Protector for Cloud Workloads backups are good.
- If your backup destination requires that Proxy VM need to have staging space on a local disk device, then Staging space must be on a volume coming from container storage. Otherwise, Data Protector for Cloud Workloads may select the wrong device during backup.
- Our recommendation is also to configure LVM filters on Proxy VM. You need to add all OS disks and partitions, follow these steps: LVM setup on Data Protector for Cloud Workloads Node for disk attachment backup mode

Things to Know About "How to Add Nutanix Hypervisor Manager to Data Protector for Cloud Workloads"

- When adding Nutanix hypervisor managers use a URL similar to the following:

```
https://PRISM_HOST:9440/api/nutanix/v3
```

- Nutanix environments require the Data Protector for Cloud Workloads Node to be installed in one of the VMs residing on the Nutanix cluster. Data Protector for Cloud Workloads should automatically detect the VM with the Data Protector for Cloud Workloads Node during the inventory synchronization operation.
- Data Protector for Cloud Workloads requires that there be a user with "cluster-admin" privileges on Prism, to process the backup/restore job.
- You can specify either a Prism Element or a Prism Central as hypervisor manager. If Prism Central is specified credentials for Prism Central and each Prism Element must be the same.
- Hypervisor tags are supported only with Prism Central
- Volume groups attached to the VMs are not affected by snapshot, hence neither backup nor snapshot revert on such volumes is going to include them.

You can deploy more nodes in each cluster and map individual hypervisors to them:

- This should statically load balance jobs based on a hypervisor
- Each node will handle VMs that reside on the particular hypervisor (which because of data locality may be faster than backup of VMs from other hosts
- VMs that don't have hypervisor assigned are handled by the node from the hypervisor manager
- Each node needs to run inventory synchronization to record its Proxy VM UUID on which it is installed

Example

How to start back up for Nutanix AHV Hypervisor

- Create Proxy VM into Nutanix cluster (with one of the supported OS: Platform Requirements)

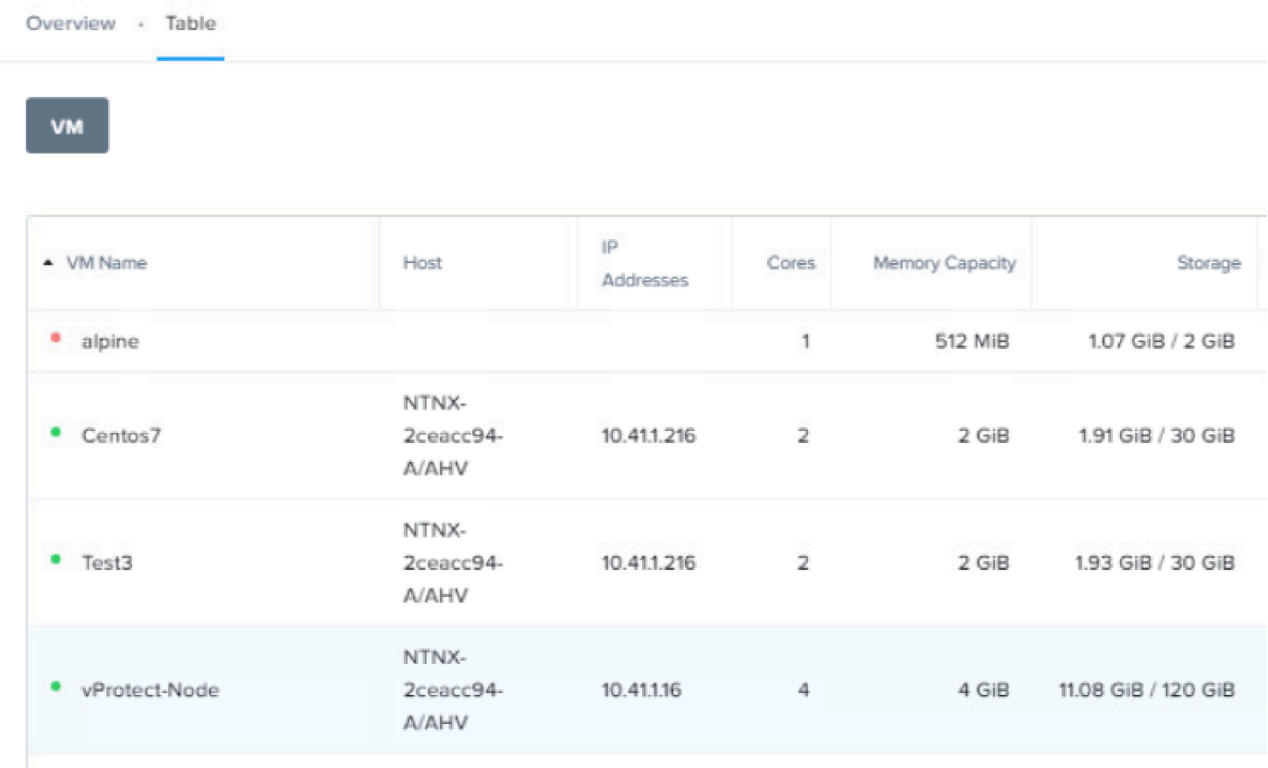

▲ VM Name	Host	IP Addresses	Cores	Memory Capacity	Storage
● alpine			1	512 MiB	1.07 GiB / 2 GiB
● Centos7	NTNX-2ceacc94-A/AHV	10.41.1.216	2	2 GiB	1.91 GiB / 30 GiB
● Test3	NTNX-2ceacc94-A/AHV	10.41.1.216	2	2 GiB	1.93 GiB / 30 GiB
● vProtect-Node	NTNX-2ceacc94-A/AHV	10.41.1.16	4	4 GiB	11.08 GiB / 120 GiB

- Install Data Protector for Cloud Workloads Node (How to install Data Protector for Cloud Workloads Node installation by RPM)

- Login to Data Protector for Cloud Workloads Dashboard and add hypervisor manager *Remember that if you add prism central all credentials must be the same (for prism elements and prism central)*

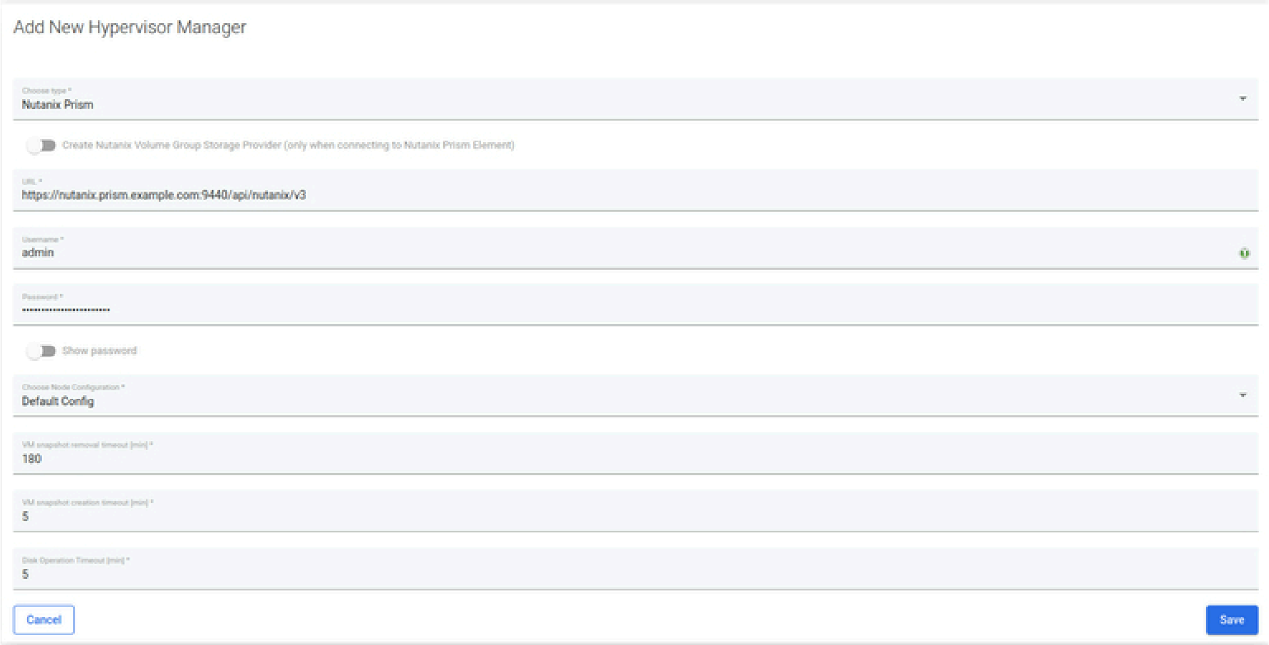

Add New Hypervisor Manager

Choose type *
Nutanix Prism

Create Nutanix Volume Group Storage Provider (only when connecting to Nutanix Prism Element)

URL *
https://nutanix.prism.example.com:9440/api/nutanix/v3

Username *
admin

Password *
••••••••••••••••••••••••

Show password

Choose Node Configuration *
Default Config

VM snapshot removal timeout [min] *
180

VM snapshot creation timeout [min] *
5

Disk Operation Timeout [min] *
5

- Run inventory synchronization task, after that you should see all Nutanix hosts under the hypervisor tab

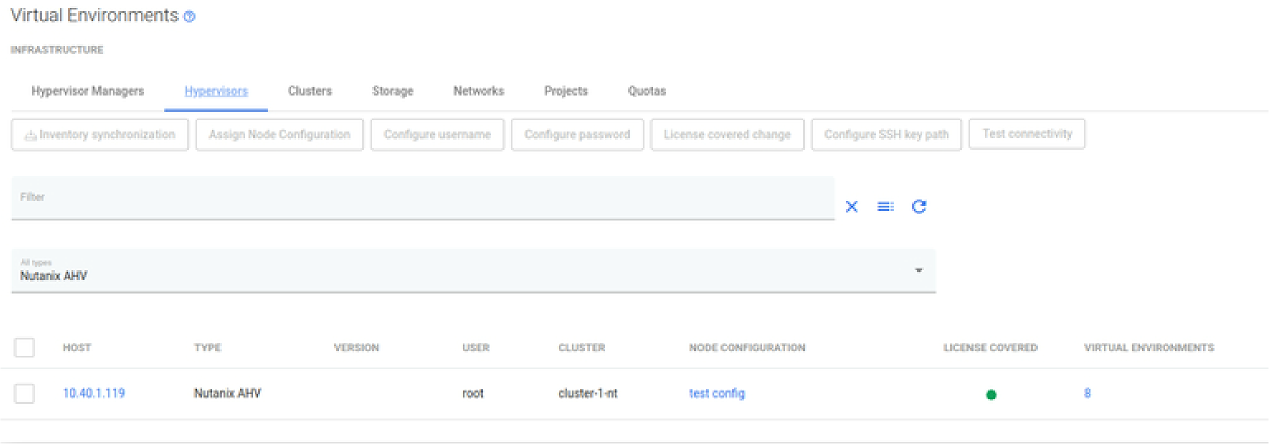

Virtual Environments

INFRASTRUCTURE

All types
Nutanix AHV

□	HOST	TYPE	VERSION	USER	CLUSTER	NODE CONFIGURATION	LICENSE COVERED	VIRTUAL ENVIRONMENTS
□	10.40.1.119	Nutanix AHV		root	cluster-1-nt	test config	●	8

- As we describe above, we can back up Nutanix VMs thanks to the disk attachment backup strategy. As this is one of the most demanding methods, at this point we recommend that you perform a few easy tests to make sure that the backup you are going to perform is correct.
- Connect via SSH to the Proxy VM. Enter "lsblk" to check the disk devices that belong to the machine. In this example, we have two disk devices:
    - 1. /dev/sda - with three partitions /dev/sda1, /dev/sda2, /dev/sda3

- 2. /dev/sdb - with one partition /dev/sdb1

This information will be needed for the next steps: configuring the lvm filter and checking if we need to correct the value of the parameter "dynamically disk attachment offset".

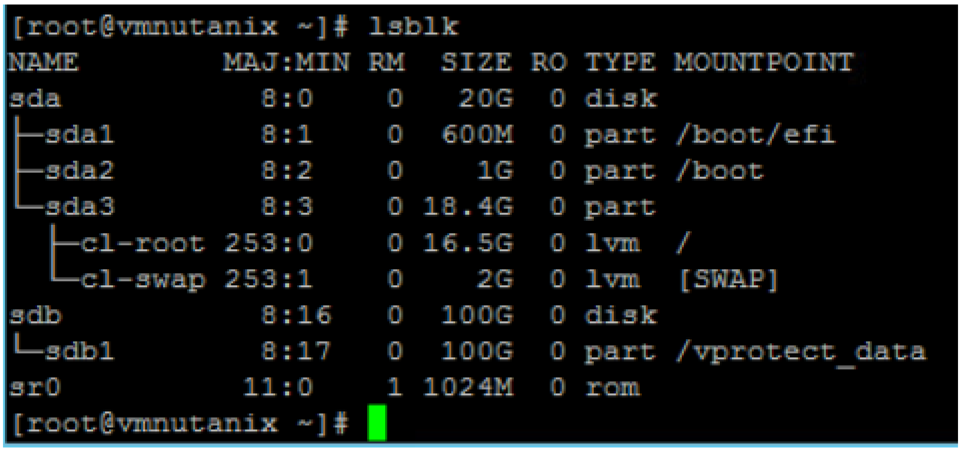

```
[root@vmnutanix ~]# lsblk
NAME          MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda             8:0    0    20G  0 disk
├─sda1          8:1    0   600M  0 part /boot/efi
├─sda2          8:2    0     1G  0 part /boot
└─sda3          8:3    0  18.4G  0 part
  ├─cl-root   253:0    0  16.5G  0 lvm  /
  └─cl-swap   253:1    0     2G  0 lvm  [SWAP]
sdb             8:16   0   100G  0 disk
└─sdb1          8:17   0   100G  0 part /vprotect_data
sr0            11:0    1  1024M  0 rom
[root@vmnutanix ~]#
```

- We'll start by configuring the lvm filter.
  - Global article about LVM: LVM setup manual
  - Remember to reboot VM after changes
  - Remember that the structure of this file is important and you need to put the filter lines back in their original place. Open in a text editor `/etc/lvm/lvm.conf` uncomment and replace the line: `filter = [ "a|.*|" ]` to `filter = [ "a|^/dev/sda|", "a|^/dev/sda1|", "a|^/dev/sda2|", "a|^/dev/sda3|", "a|^/dev/sdb|", "a|^/dev/sdb1|", "r|.|" ]` **and** `global_filter = [ "a|.*|" ]` **to** `global_filter = [ "a|^/dev/sda|", "a|^/dev/sda1|", "a|^/dev/sda2|", "a|^/dev/sda3|", "a|^/dev/sdb|", "a|^/dev/sdb1|", "r|.|" ]`
- Now we can move on to the "dynamically disk attachment offset" tests. *You need to do this only if Proxy VM has more than one disk device for OS purposes*
  - Switch Data Protector for Cloud Workloads Node logs (Proxy VM) to Debug mode: How to Enable Debug mode
  - Run a test backup - try to choose a small VM to not wait too long
  - After the backup is complete, download the log file from our dashboard

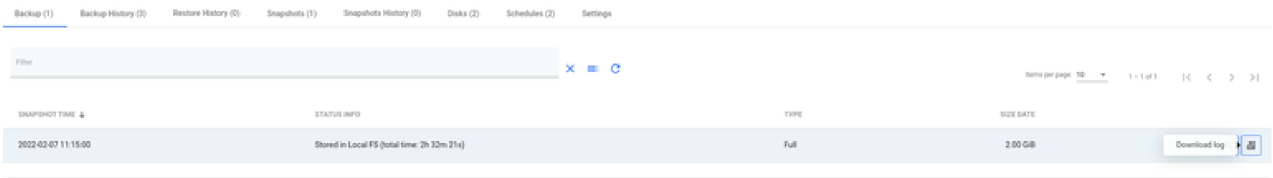

SNAPSHOT TIME	STATUS INFO	TYPE	SIZE DATE
2022-02-07 11:15:00	Stored in Local FS (total time: 2h 32m 21s)	Full	2.00 GiB

- As we can see in the logs, we do not need to correct the "offset" value. Data Protector for Cloud Workloads wants to start a backup from /dev/sdc, which is correct behavior because this disk device does not belong to Proxy VM.

```
[2021-04-08 14:51:40.959] INFO [Thread-47] IProxyVmProvider.waitForDevice
[ffc65c30-8952-4ffa-b5d5-eefcfe01f333] Checking if device '/dev/sdc' is

[2021-04-08 14:51:45.959] DEBUG [Thread-47] CommandExecutor.exec:75
[ffc65c30-8952-4ffa-b5d5-eefcfe01f333] Exec: [lsblk, -l, /dev/sdc]

[2021-04-08 14:51:45.969] DEBUG [Thread-47] CommandExecutor.exec:102
[ffc65c30-8952-4ffa-b5d5-eefcfe01f333] [lsblk, -l, /dev/sdc]
Return code: 0
output:
[NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sdc    8:32   0  20G  0 disk
sdc1   8:33   0   1G  0 part
sdc2   8:34   0  19G  0 part
]
error:
[]

[2021-04-08 14:51:45.970] INFO [Thread-47] IProxyVmProvider.waitForDevice
[ffc65c30-8952-4ffa-b5d5-eefcfe01f333] Device '/dev/sdc' is present

[2021-04-08 14:51:55.991] INFO [Thread-47] NutanixHypervisorManager.expor
[ffc65c30-8952-4ffa-b5d5-eefcfe01f333] Data export of scsi.0
(917a15a2-5815-4d20-b693-6fb77ea59293)[20 GiB]: '/dev/sdc' -> '/vprotect_
```

- If you meet with a situation, when Data Protector for Cloud Workloads want to back up its own disk device, read our knowledge base article: KB10037 How to change "Dynamically attached disks slot offset" parameter

Limitations

- Backup of VMs with vTPM enabled is not supported.

Red Hat Openshift Virtualization

Red Hat Openshift Virtualization

Data Protector for Cloud Workloads supports backup for OpenShift using OADP operator. Metadata of Virtual Machines is exported using OADP operator, volume data is exported using side pod using custom OpenShift Virtualization Plugin docker image. The backup supports both full and incremental types. Incremental backup does not require previous snapshots to remain in OpenShift.

> Prior to adding OpenShift as a new Hypervisor Manager, you must install the OADP operator, version 1.3 or higher, from the Operator Hub within the OpenShift cluster.

Adding Openshift Hypervisor Manager

Log in to the web interface and add a new OpenShift Hypervisor Manager:

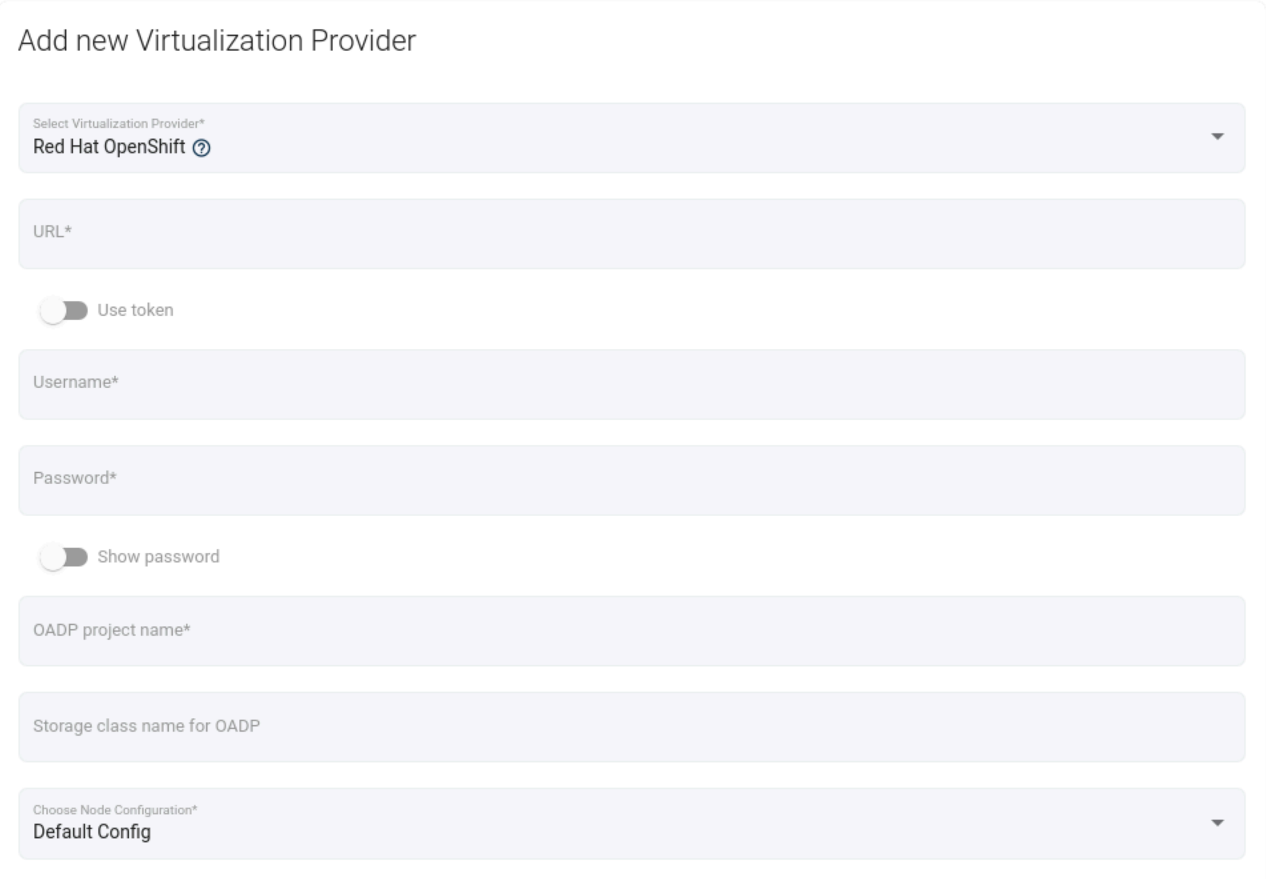

Add new Virtualization Provider

Select Virtualization Provider*
Red Hat OpenShift

URL*

Use token

Username*

Password*

Show password

OADP project name*

Storage class name for OADP

Choose Node Configuration*
Default Config

- **URL** - URL of the Openshift API e.g. `api.your.cluster.local:6443`
- **Username** - login of a user with the cluster-admin role
- **OADP project name** - project name where OADP Operator was installed (`openshift-adp` by default)
- **Storage class name for OADP** - specify storage class that will be used for OADP setup, if this field is empty, default storage class will be used (optional)

The Openshift Nodes should appear in Data Protector for Cloud Workloads after indexing the cluster.

Using own image registry for OpenShift Virtualization Plugin

Data Protector for Cloud Workloads use quay.io as default image registry for OpenShift Virtualization Plugin docker image. You can use your own registry to store the plugin image.

1. Download the Data Protector for Cloud Workloads package from the Micro Focus download page.

2. Extract your package and find plugin file in **addons** directory.
3. Upload it to your image registry host.
4. Import image to your registry. Example:

   ```
   gunzip sbr-openshift-virtualization-plugin-jvm-x.x.x.x.tar.gz
   docker load -i sbr-openshift-virtualization-plugin-jvm-x.x.x.x.tar
   ```

5. Edit `/opt/vprotect/node.properties` file and change value for `openshift.virtualization.sidepod.image` parameter. Example:

   ```
   openshift.virtualization.sidepod.image=<Registry IP>:5000/sbr-openshif
   ```

6. Restart vprotect-node service.

   ```
   systemctl restart vprotect-node
   ```

Limitations

- For a successful backup, Virtual Machine should have an **app** label assigned appropriately.
- Hot-plugged disks are not supported.
- Backup of disks: CDROM and LUN is not supported.
- Storage class used for disk should support snapshots.

Red Hat Virtualization

Red Hat Virtualization

General

For RHV 4+ environments you can use API v4 for invoking all backup-related tasks.

Import/export mode defines the way the backups and restores are done. Red Hat Virtualization (with API v4) supports 4 modes:

1. **Disk attachment**, which exports VM metadata (in OVF format) with separate disk files (in RAW format) via the Proxy VM with the Node installed.
   - supports RHV 4.0+
   - no incremental backup
   - proxy VM required in each cluster - used for the disk attachment process
2. **Disk image transfer**, which exports VM metadata (in OVF format) with disk snapshot chains as separate files (QCOW2 format):
   - supports RHV 4.2+/oVirt 4.2.3+
   - supports incremental backup
   - disk images are transferred directly from API (no Proxy VM required)
3. **SSH Transfer,** this method assumes that all data transfers are directly from the hypervisor - over SSH
4. **Change Block Tracking,** this method backup only blocks with changes and skip zeroed sectors.
   - supports oVirt 4.4+ (with Libvirt 6+, qemu-kvm 4.2+ and vdsm 4.40+)
   - supports incremental backup

> **Note:** When using backup APIs - Red Hat highly recommends updating the RHV environment to the most recent version (4.4 - at the time of writing) - refer to this article for more information.

When adding RHV 4.0+ hypervisor managers, use a URL similar to the following:

```
https://RHV_MGR_HOST/ovirt-engine/api
```

> **Note:** a username for RHV environments needs to be provided in the **user@domain** format - for example **admin@internal**. This user must have all permissions related to managing snapshots, creating/removing VMs, operating disks, and exporting data.

Backup Strategies

Red Hat Virtualization environments can be protected in several ways.

Note: Different strategies require a node to be installed either as a VM on the environment that you back up or installed separately.

Note: All live snapshots are attempted with quiescing enabled. If the snapshot command fails because there is no compatible guest agent present, the live snapshot is re-initiated without the use-quiescing flag.

Disk attachment with Proxy VM

In this strategy, you have a VM called “Proxy VM” that invokes commands on your hypervisor manager to snapshot and attach drives of a specific VM to itself (Proxy VM). The Proxy VM is able to read the data from the attached disk snapshots and forward them to the backup provider.

This strategy allows you to exclude drives from the backup that you do not need. Remember that you need to install 1 Proxy VM per cluster so that the drives the node tries to attach are reachable.

Drawback - no incremental backup for now.

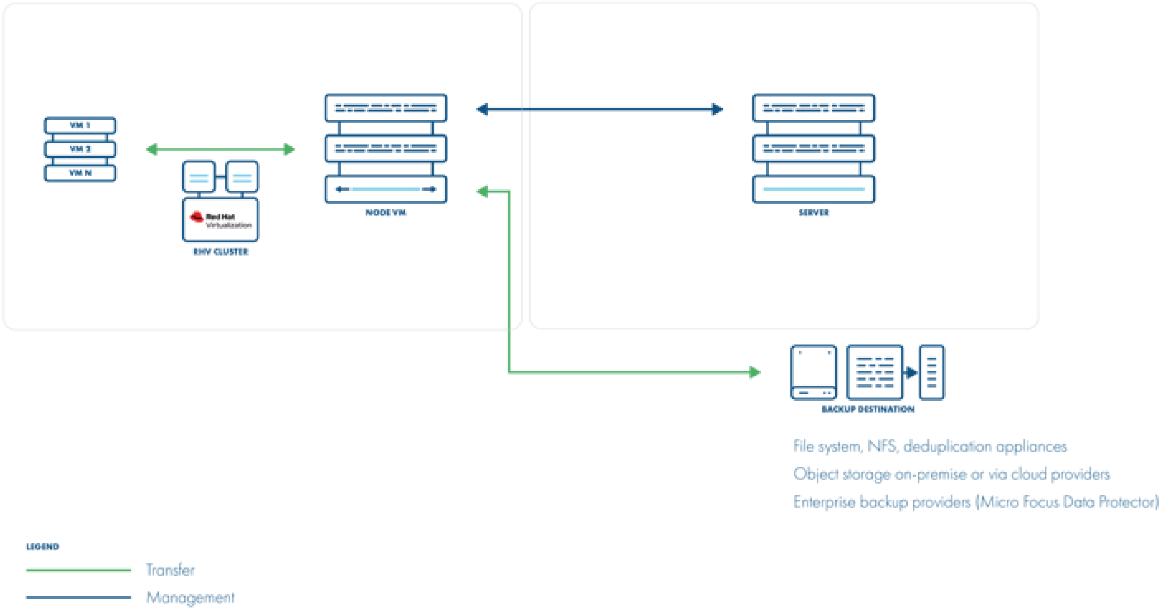
VM 1
VM 2
VM N
Red Hat Virtualization
RHV CLUSTER
NODE VM
SERVER
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot using hypervisor's API
- optionally FS freeze can be executed before snapshot can be executed (FS thaw once the snapshot is completed) if enabled and guest tools installed inside
- optional application consistency using pre/post snapshot command execution
- metadata exported from API
- snapshot disks are mounted one by one to the Proxy VM
- data read directly on the Proxy VM
- incremental backups are _**_not supported
- restore creates empty disks on the Proxy VM, imports merged data then recreates VM and reattaches volumes to the target VM

Note: RHV API v4 environments require Data Protector for Cloud Workloads Node to be installed in one of the VMs residing on the RHV cluster. Data Protector for Cloud Workloads should automatically detect the VM with Data Protector for Cloud Workloads during the index operation.

Disk attachment mode requires `Virtio-SCSI` to be enabled on the Data Protector for Cloud Workloads Node VM (which can be enabled in `VM settings` → `Resource Allocation` → `VirtIO-SCSI Enabled` at the bottom).

During backup/restore operations, disks are transferred by attaching them to the proxy VM. This approach does not require an export storage domain to be set up.

Please make sure that you follow these steps: LVM setup on Data Protector for Cloud Workloads Node for disk attachment backup mode.

Disk image transfer API

This API appears in RHV 4.2 and allows the export of individual snapshots directly from the RHV manager. So instead of having to install multiple Proxy VMs, you can have a single external Node installation, which just invokes APIs via the RHV manager.

This strategy supports incremental backups. Assuming you have RHV 4.2 or newer – just add your manager to Data Protector for Cloud Workloads and setup is done. From a network perspective, it requires two additional ports to be open - 54322 and 54323 - and your data to be pulled from the hypervisor manager.

Unfortunately, there are a few problems with the current architecture of this solution. The biggest issue is that all traffic passes via the RHV manager, which may impact the transfer rates that you can achieve during the backup process. To put this into perspective – in disk attachment, you can basically read data as if it is a local drive, where it could potentially be deduplicated even before transferring it to the backup destination.

> **Note:** From RHV version 4.4.3, data is transferred directly from/to hosts.

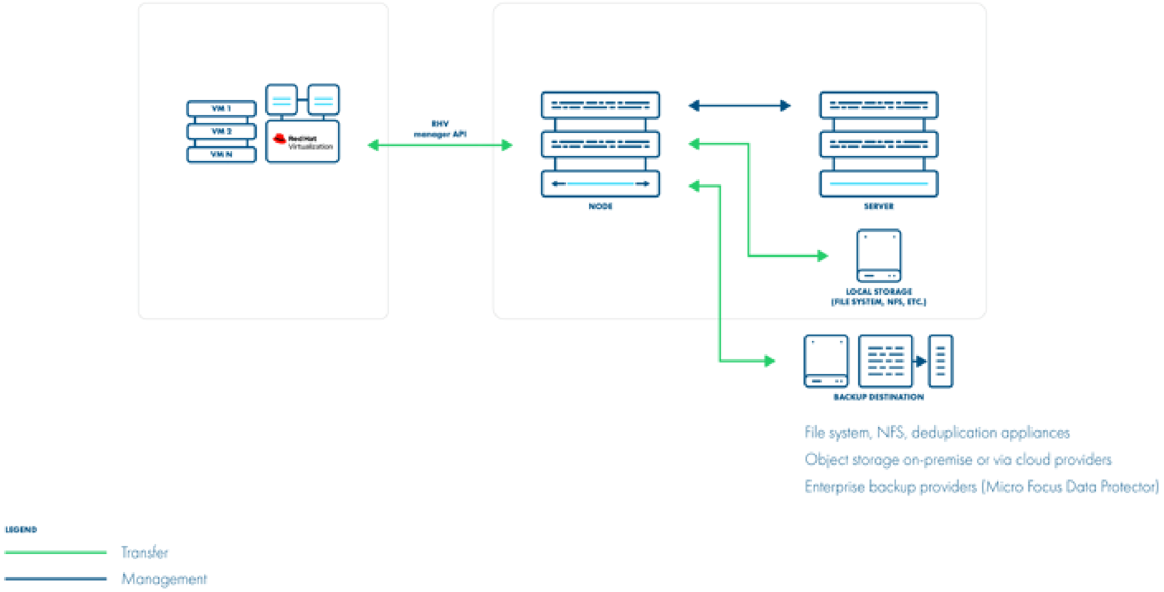
VM 1
VM 2
VM N
Red Hat Virtualization
RHV manager API
NODE
SERVER
LOCAL STORAGE (FILE SYSTEM, NFS, ETC.)
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot using hypervisor's API
- optionally FS freeze can be executed before snapshot can be executed (FS thaw once the snapshot is completed) if enabled and guest tools installed inside
- optional application consistency using pre/post snapshot command execution
- supported for oVirt/RHV/OLVM 4.3+
- metadata exported from API
- data transfer initiated on the manager and actual data exported from the hypervisor using imageio API
- incremental backups use the same APIs, but requests for changed blocks only
- the last snapshot kept on the hypervisor for the next incremental backup (if at least one schedule assigned to the VM has a backup type set to incremental)
- restore recreates VM from metadata using API and imports merged chain of data for each disk using imageio API

Disk image transfer mode exports data directly using RHV 4.2+ API. There is no need to set up an export storage domain or set up an LVM. This mode uses snapshot chains provided by RHV.

You may need to open communication for the additional port **54323** on the OLVM manager and **54322** on the RHV hosts - it needs to be accessible from Data

Protector for Cloud Workloads Node. Also, make sure that your **ovirt-imageio-proxy** services are running and properly configured (you can verify it by trying to upload images with OLVM UI).

Follow the steps in this section: Full versions of libvirt/qemu packages installation.

SSH transfer

This is an enhancement for the disk image transfer API strategy. It allows Data Protector for Cloud Workloads to use RHV API v4.2+ (HTTPS connection to RHV manager) only to collect metadata. Backup is done over SSH directly from the hypervisor (optionally using netcat for transfer), import is also using SSH (without netcat option). No need to install a node on the RHV environment. This method can boost backup transfers and supports incremental backups.

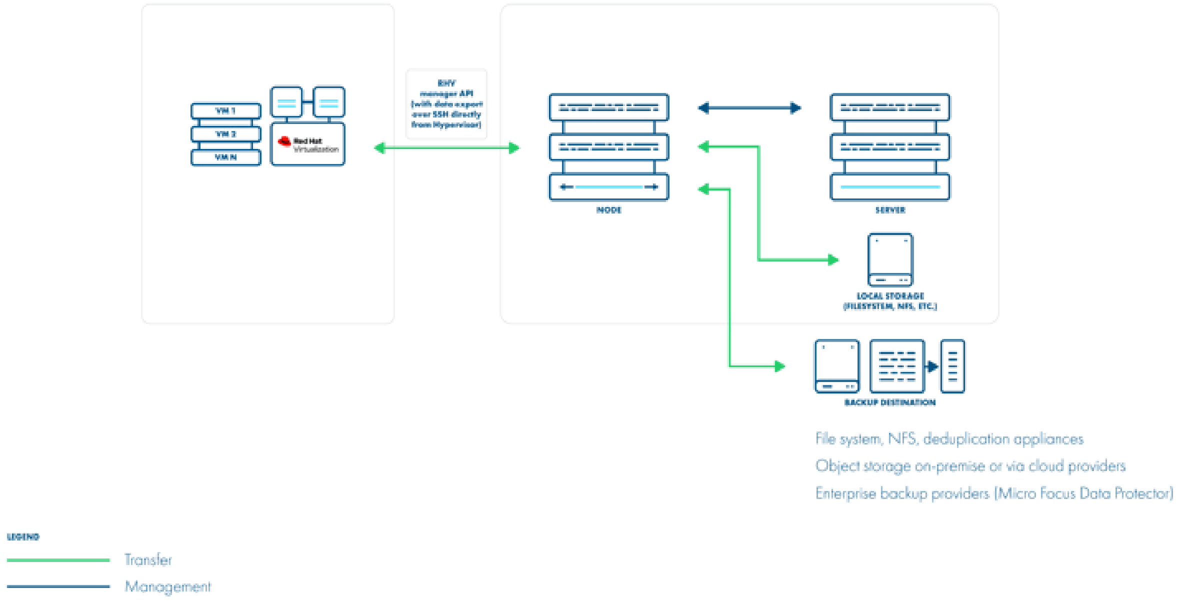
VM 1
VM 2
VM N
Red Hat Virtualization
RHV manager API (with data export over SSH directly from Hypervisor)
NODE
SERVER
LOCAL STORAGE (FILESYSTEM, NFS, ETC.)
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot using hypervisor's API
- optionally FS freeze can be executed before snapshot can be executed (FS thaw once the snapshot is completed) if enabled and guest tools installed inside
- optional application consistency using pre/post snapshot command execution • metadata exported from API

- data transfer via SSH (optional using netcat) - the full chain of disk snapshot files for each disk o if LVM-based storage is used, then node activates volumes if necessary to read data o if Gluster FS is used, then disk files are copied directly
- incremental backup export just sub-chain of QCOW2-deltas snapshots since last stored snapshot
- the last snapshot kept on the hypervisor for the next incremental backup (if at least one schedule assigned to the VM has the backup type set to incremental)
- restore recreates VM with empty storage from metadata using API and imports merged data over SSH to appropriate location on the hypervisor

This method assumes that all data transfers are directly from the hypervisor - over SSH. This means that after adding the RHV manager and detecting all available hypervisors - **you also need to provide SSH credentials or SSH keys for each of the hypervisors**. You can also use SSH public key authentication.

Change Block Tracking

This is a new method that is possible thanks to changes in RHV 4.4. It uses information about zeroed and changed blocks to reduce data size and make the process faster.

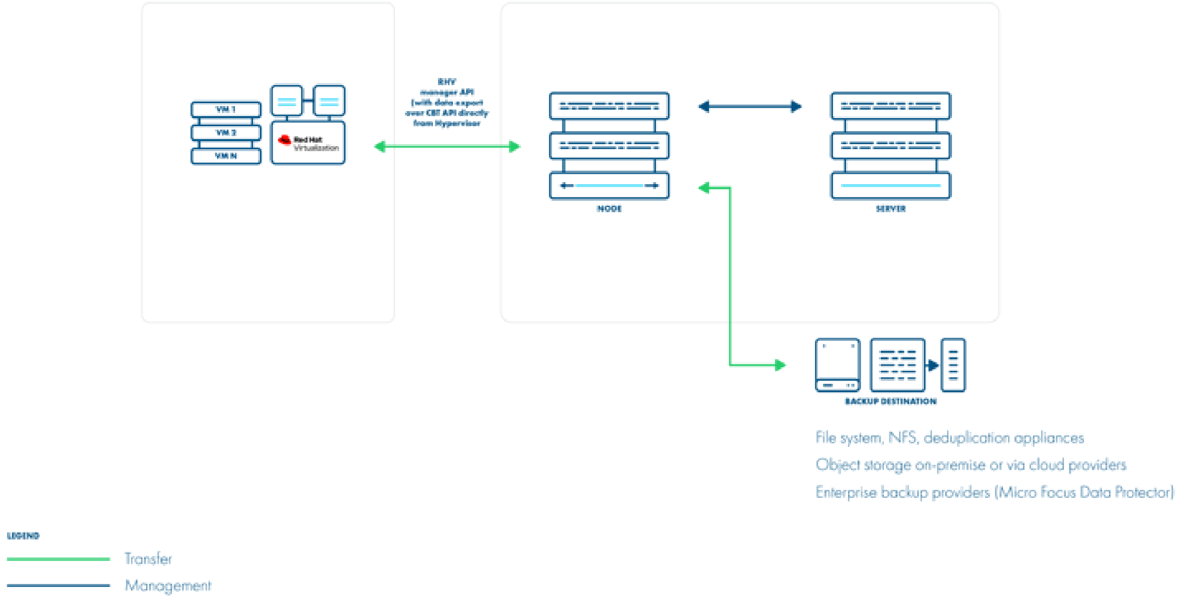
VM 1
VM 2
VM N
Red Hat Virtualization
RHV manager API (with data export over CBT API directly from Hypervisor
NODE
SERVER
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

This strategy supports incremental backups.

The QCOW2 format is required for incremental backups so that disks enabled for incremental backup use the QCOW2 format instead of the raw format.

Also, this strategy doesn't need snapshots in the backup process. Instead, every incremental backup uses a checkpoint that is a point in time that was created after the previous backup.

Export storage domain (API v3)

This setup requires you to create a storage domain used for VM export. The export storage domain should also be accessible by Data Protector for Cloud Workloads Node in its staging directory. This implies that the storage space doesn't have to be exported by Data Protector for Cloud Workloads Node - it can be mounted from an external source. The only requirement is to have it visible from both the RHV host and the Node itself. Keep in mind that ownership of the files on the share should allow both Data Protector for Cloud Workloads and RHV to read and write files.

The backup process requires that once the snapshot is created, it will be cloned and exported (in fact to Data Protector for Cloud Workloads Node staging). The reason for additional cloning is that RHV doesn’t allow you to export a snapshot directly. The Node can be outside of the environment that you back up.

This strategy is going to be deprecated, as Red Hat may no longer support it in future releases.

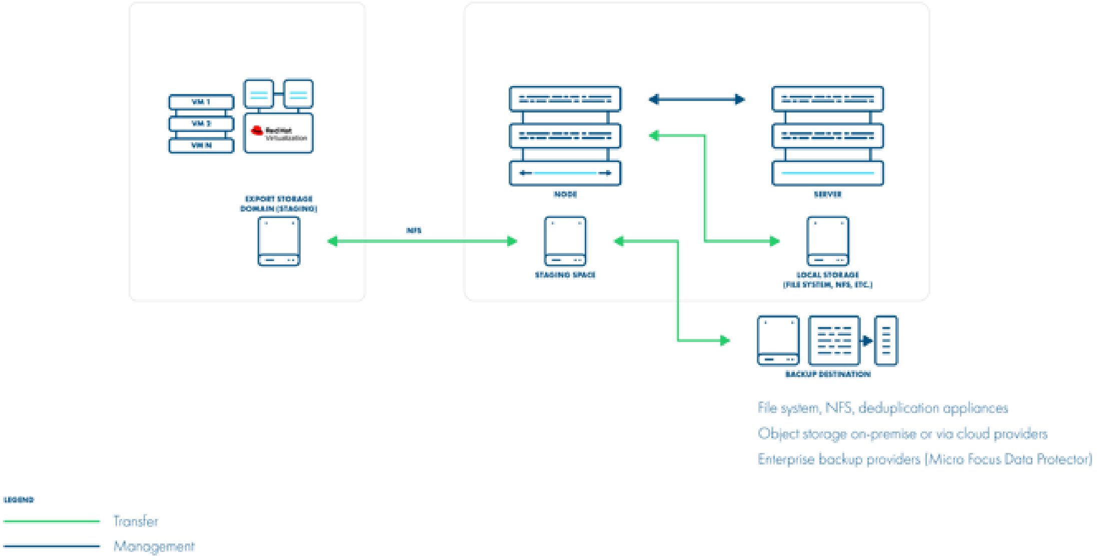
VM 1
VM 2
VM N
Red Hat Virtualization
EXPORT STORAGE DOMAIN (STAGING)
NFS
NODE
SERVER
STAGING SPACE
LOCAL STORAGE (FILE SYSTEM, NFS, ETC.)
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot is taken via API
- optional application consistency using pre/post snapshot command execution
- initial VM clone of the snapshot to the local repository is created
- cloned VM (data+metadata) exported by the manager to the Data Protector for Cloud Workloads staging space (visible as the export Storage Domain in managers UI)
- full backup only is supported
- restore is done to the export Storage Repository, the administrator needs to import the VM using manager UI

How to set up a backup with an export storage domain

RHV 3.5.1+ environments (using API v3) require an export storage domain to be set up.

1. Add a backup storage domain in RHEV (which points to the NFS export on Data Protector for Cloud Workloads Node)
    - If you have multiple data centers, you need to enable the **Multi DC export** a checkbox in the node configuration
        - Remember that you need to use named data centers in your RHV environment to avoid name conflicts
        - An RHV datacenter may use only one export storage domain, which is why you need to create subdirectories for each data center in the export path for example `/vprotect_data/dc01`, `/vprotect_data/dc02`, and use each sub-directory as NFS share for each data center export domain (separate NFS exports)
        - The export (staging) path in the above-mentioned scenario is still `/vprotect_data`, while `dc01` and `dc02` are data center names
        - Older versions of RHV (3.5.x) require you to specify a mapping between DC names and export storage domains - you need to provide pairs of a DC name and a corresponding SD name in the node configuration (section `Hypervisor`)
    - If you have only one data center and don't want to use the multiple data centers export feature in the future, you can use the default settings and set

up the NFS export pointing to the staging path (e.g. `/vprotect_data`)

- Note that export must be set to use the UID and GID of the `vprotect` user
- Example export configuration in `/etc/exports` to a selected hypervisor in the RHV cluster:

  ```
  /vprotect_data    10.50.1.101(fsid=6,rw,sync,insecure,all_squash,ar
  ```

  where `anonuid=993` and `anongid=990` should have the correct UID and GID returned by command:

  ```
  [root@vProtect3 ~]# id vprotect
  uid=993(vprotect) gid=990(vprotect) groups=990(vprotect)
  ```

2. Both import and export operations will be done using this NFS share – restore will be done directly to this storage domain, so you can easily import the backup into RHV (shown below)
   - backups must be restored to the export path (the node automatically changes names to the original paths that are recognized by the RHV manager).
3. When adding RHV 4.0+ hypervisor managers, make sure you have a URL like the following:

   ```
   https://RHV_MGR_HOST/ovirt-engine/api/v3
   ```

> **Note:** Restore to RHV using SPARSE disk allocation format is not supported if backup files are in RAW format and destination storage domain type in either Fibre Channel or iSCSI. If such configuration is detected, then disk allocation format is automatically switched to PREALLOCATED

Instant restore

To use an instant restore feature, backup destination from which VM will be restored, has to be of a synthetic type. The restore process creates a NFS share on the Data Protector for Cloud Workloads node, later this share is attached to the RHV as a new storage domain. Then it creates a new virtual machine and attaches the disks from the newly created storage domain to it. To use instant restore you have to click the restore button in the instances list and choose the option **instant restore**.

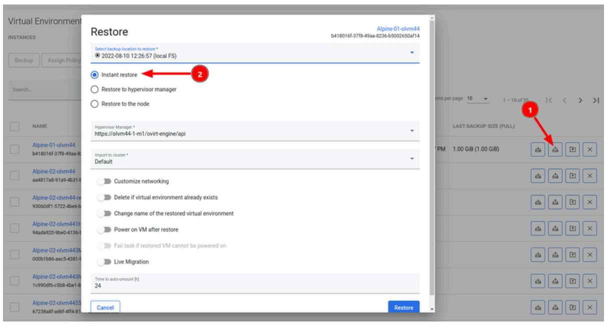

Restore

Alpine-01-olvm44
b418016f-37f8-49aa-8236-b5002650af14

Select backup location to restore *
● 2022-08-10 12:26:57 (local FS)

● Instant restore
○ Restore to hypervisor manager
○ Restore to the node

Hypervisor Manager *
https://olvm44-1-m1/ovirt-engine/api

Import to cluster *
Default

Customize networking
Delete if virtual environment already exists
Change name of the restored virtual environment
Power on VM after restore
Fail task if restored VM cannot be powered on
Live Migration

Time to auto-umount [h]
24

Cancel
Restore

Live migration

You can enable the live migration option during instant restore. It will automatically start the disks migration to the chosen storage after the VM is restored and powered on.

oVirt

oVirt

General

For oVirt 4+ environments you can use API v4 for invoking all backup-related tasks.

Import/export mode defines the way the backups and restores are done. oVirt (with API v4) supports 4 modes:

1. **Disk attachment**, which exports VM metadata (in OVF format) with separate disk files (in RAW format) via Proxy VM with the Node installed.
   - supports oVirt 4.0+
   - no incremental backup
   - proxy VM required in each cluster - used for the disk attachment process
2. **Disk image transfer**, which exports VM metadata (in OVF format) with disk snapshot chains as separate files (QCOW2 format):
   - supports oVirt 4.2+/oVirt 4.2.3+
   - supports incremental backup
   - disk images are transferred directly from API (no Proxy VM required)
3. **SSH Transfer,** this method assumes that all data transfers are directly from the hypervisor - over SSH
4. **Change Block Tracking,** this method backs up only blocks with changes and skip zeroed sectors.
   - supports oVirt 4.4+ (with Libvirt 6+, qemu-kvm 4.2+ and vdsm 4.40+)
   - supports incremental backup
   - only disks marked with "enable incremental backup" in ovirt will be backed up

> **Note:** When using backup APIs - Red Hat highly recommends updating the oVirt environment to the most recent version (4.4 - at the time of writing) - refer to this article for more information.

When adding oVirt 4.0+ hypervisor managers, use a URL similar to the following:

```
https://oVirt_MGR_HOST/ovirt-engine/api
```

> **Note:** a username for oVirt environments needs to be provided in the **user@domain** format - for example **admin@internal**. This user must have all permissions related to managing snapshots, creating/removing VMs, operating disks, and exporting data.

Backup Strategies

oVirt environments can be protected in several ways.

> **Note:**
>
> Different strategies require a node to be installed either as a VM on the environment that you back up or installed separately.
>
> All live snapshots are attempted with quiescing enabled. If the snapshot command fails because there is no compatible guest agent present, the live snapshot is re-initiated without the use-quiescing flag.

Disk attachment with Proxy VM

In this strategy, you have a VM called “Proxy VM” that invokes commands on your hypervisor manager to snapshot and attach drives of a specific VM to itself (Proxy VM). Proxy VM is able to read the data from the attached disk snapshots and forward them to the backup provider.

This strategy allows you to exclude drives from a backup that you do not need. Remember that you need to install 1 Proxy VM per cluster so that the drives the node tries to attach are reachable.

Drawback - no incremental backup for now.

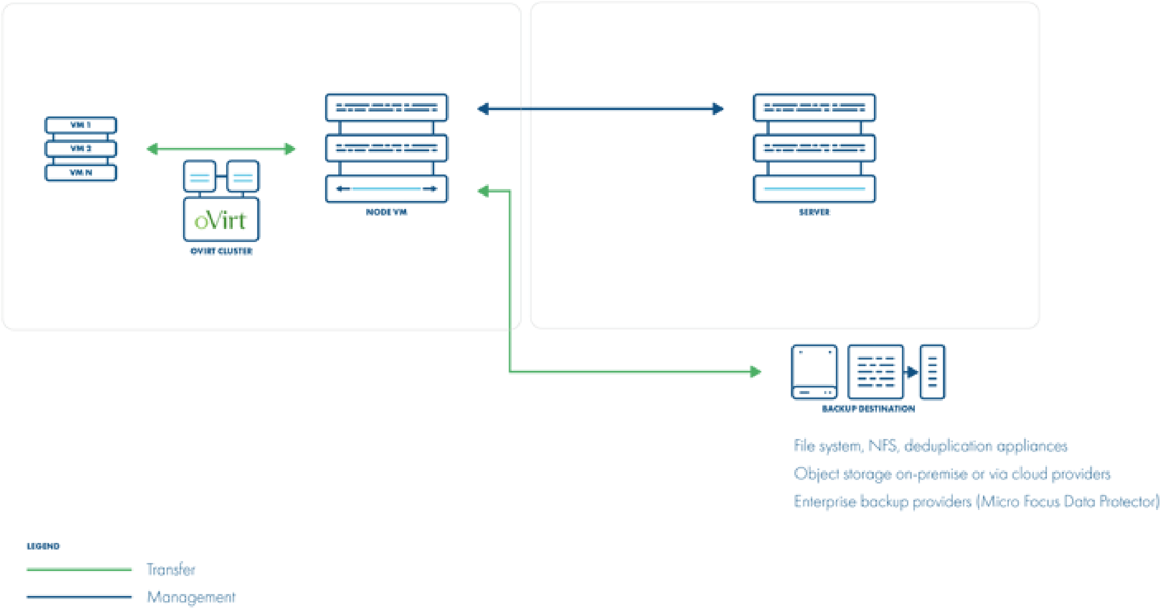
VM 1
VM 2
VM N
oVirt
OVIRT CLUSTER
NODE VM
SERVER
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot using hypervisor's API
- optionally FS freeze can be executed before snapshot can be executed (FS thaw once the snapshot is completed) if enabled and guest tools installed inside
- optional application consistency using pre/post snapshot command execution
- metadata exported from API
- snapshot disks are mounted one by one to the Proxy VM
- data read directly on the Proxy VM
- incremental backups are _**_not supported
- restore creates empty disks on the Proxy VM, imports merged data then recreates VM and reattaches volumes to the target VM

> **Note**: oVirt API v4 environments require Data Protector for Cloud Workloads Node to be installed in one of the VMs residing in the oVirt cluster. Data

> Protector for Cloud Workloads should automatically detect the VM with Data Protector for Cloud Workloads during the index operation.

Disk attachment mode requires `Virtio-SCSI` to be enabled on the Data Protector for Cloud Workloads Node VM (which can be enabled in `VM settings` → `Resource Allocation` → `VirtIO-SCSI Enabled` at the bottom).

During backup/restore operations, disks are transferred by attaching them to the proxy VM. This approach does not require an export storage domain to be set up.

Make sure you follow these steps: LVM setup on Data Protector for Cloud Workloads Node for disk attachment backup mode.

Disk image transfer API

This API appeared in oVirt 4.2 and allowed the export of individual snapshots directly from the oVirt manager. So instead of having to install multiple Proxy VMs, you can have a single external Node installation, which just invokes APIs via the oVirt manager.

This strategy supports incremental backups. Assuming you have oVirt 4.2 or newer – just add your manager to Data Protector for Cloud Workloads and the setup is done. From a network perspective, it requires two additional ports to be opened - 54322 and 54323 - and your data to be pulled from the hypervisor manager.

Unfortunately, there are a few problems with the current architecture of this solution. The biggest issue is that all traffic passes via the oVirt manager, which may impact the transfer rates that you can achieve during the backup process. To put that into perspective – in disk attachment, you can basically read data as if it is a local drive, where it could potentially be deduplicated even before it is transferred to the backup destination.

> **Note:** From oVirt version 4.4.3, data is transferred directly from/to hosts.

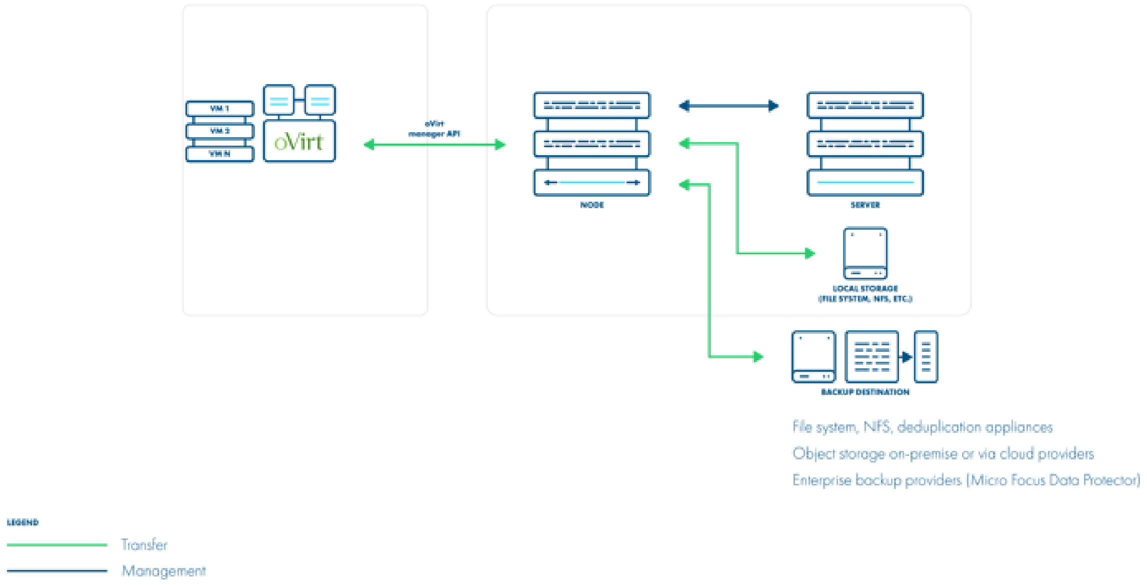
VM 1
VM 2
VM N
oVirt
oVirt manager API
NODE
SERVER
LOCAL STORAGE (FILE SYSTEM, NFS, ETC.)
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot using hypervisor's API
- optionally FS freeze can be executed before snapshot can be executed (FS thaw once the snapshot is completed) if enabled and guest tools installed inside
- optional application consistency using pre/post snapshot command execution
- supported for oVirt/RHV/OLVM 4.3+
- metadata exported from API
- data transfer initiated on the manager and actual data exported from the hypervisor using imageio API
- incremental backups use the same APIs, but requests for changed blocks only
- the last snapshot kept on the hypervisor for the next incremental backup (if at least one schedule assigned to the VM has the backup type set to incremental)
- restore recreates VM from metadata using API and imports merged chain of data for each disk using imageio API

Disk image transfer mode exports data directly using oVirt 4.2+ API. There is no need to set up an export storage domain or setup LVM. This mode uses snapshot chains provided by oVirt.

You may need to open communication for the additional port **54323** on the OLVM manager and **54322** on the oVirt hosts - it needs to be accessible from Data

Protector for Cloud Workloads Node. Also, make sure that your **ovirt-imageio-proxy** services are running and properly configured (you can verify it by trying to upload images with OLVM UI).

Follow the steps in this section: Full versions of libvirt/qemu packages installation.

SSH transfer

This is an enhancement for disk image transfer API strategy. It allows Data Protector for Cloud Workloads to use oVirt API v4.2+ (HTTPS connection to oVirt manager) only to collect metadata. Backup is done over SSH directly from the hypervisor (optionally using netcat for transfer), import is also using SSH (without the netcat option). There is no need to install a node in the oVirt environment. This method can significantly boost backup transfers and supports incremental backups.

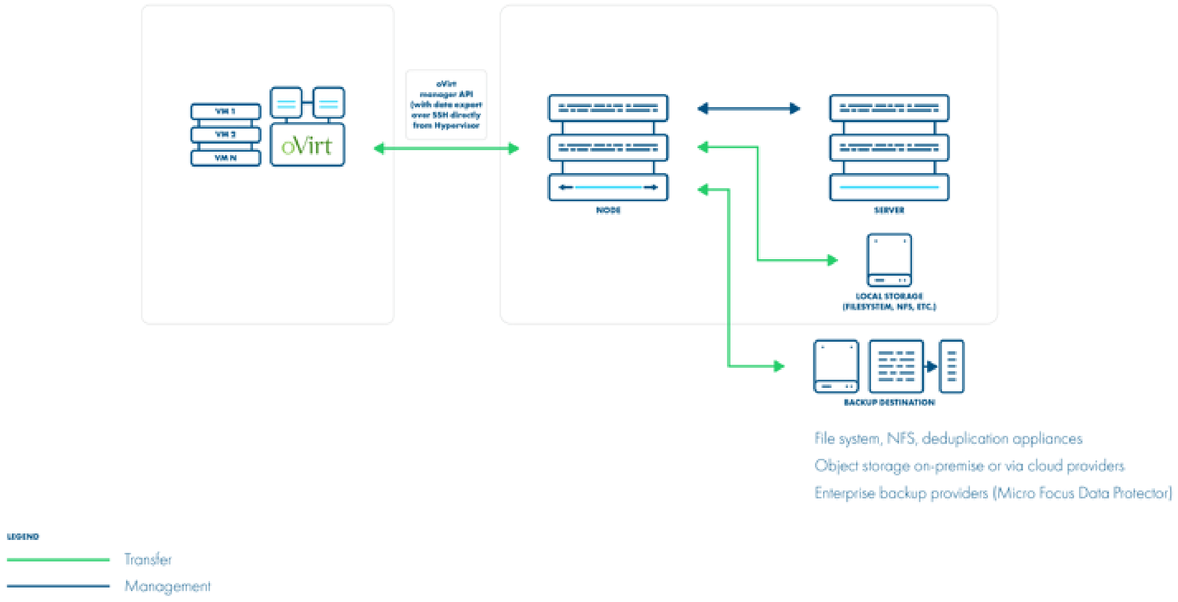
VM 1
VM 2
VM N
oVirt
oVirt manager API (with data export over SSH directly from Hypervisor
NODE
SERVER
LOCAL STORAGE (FILESYSTEM, NFS, ETC.)
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot using hypervisor's API
- optionally FS freeze can be executed before snapshot can be executed (FS thaw once the snapshot is completed) if enabled and guest tools installed inside
- optional application consistency using pre/post snapshot command execution • metadata exported from API

- data transfer via SSH (optional using netcat) - the full chain of disk snapshot files for each disk o if LVM-based storage is used, then node activates volumes if necessary to read data o if Gluster FS is used, then disk files are copied directly
- incremental backup export just sub-chain of QCOW2-deltas snapshots since last stored snapshot
- the last snapshot kept on the hypervisor for the next incremental backup (if at least one schedule assigned to the VM has the backup type set to incremental)
- restore recreates VM with empty storage from metadata using API and imports merged data over SSH to appropriate location on a hypervisor

This method assumes that all data transfers are directly from the hypervisor over SSH. This means that after adding oVirt manager and detecting all available hypervisors - **you also need to provide SSH credentials or SSH keys for each of the hypervisors**. You can also use SSH public key authentication.

Change Block Tracking

This is a new method which is possible thanks to changes in oVirt 4.4. It uses information about zeroed and changed blocks to reduce data size and make the process faster.

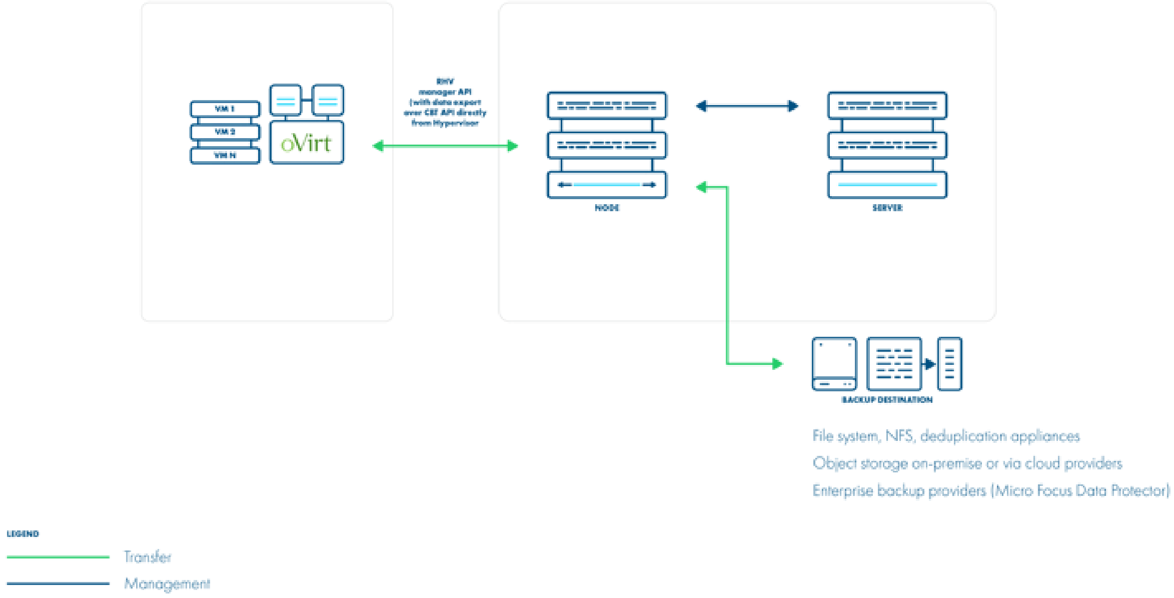
VM 1
VM 2
VM N
oVirt
RHV manager API (with data export over CBT API directly from Hypervisor
NODE
SERVER
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

This strategy supports incremental backups.

The QCOW2 format is required for incremental backups, so disks enabled for the incremental backup will use the QCOW2 format instead of the raw format.

Also, this strategy doesn't need snapshots in the backup process. Instead, every incremental backup uses a checkpoint that is a point in time that was created after the previous backup.

Export storage domain (API v3)

This setup requires you to create a storage domain used for VM export. The export storage domain should also be accessible to Data Protector for Cloud Workloads Node in its staging directory. This implies that storage space doesn't have to be exported by Data Protector for Cloud Workloads Node - it can be mounted from an external source. The only requirement is to have it visible from both the oVirt host and the Node itself. Keep in mind that ownership of the files on the share should allow both Data Protector for Cloud Workloads and oVirt to read and write files.

The backup process requires that once a snapshot is created it will be cloned and exported (in fact to Data Protector for Cloud Workloads Node staging). The reason for additional cloning is that oVirt doesn't allow you to export snapshots directly. The Node can be outside of the environment that you back up.

This strategy is going to be deprecated, as oVirt may no longer support it in future releases.

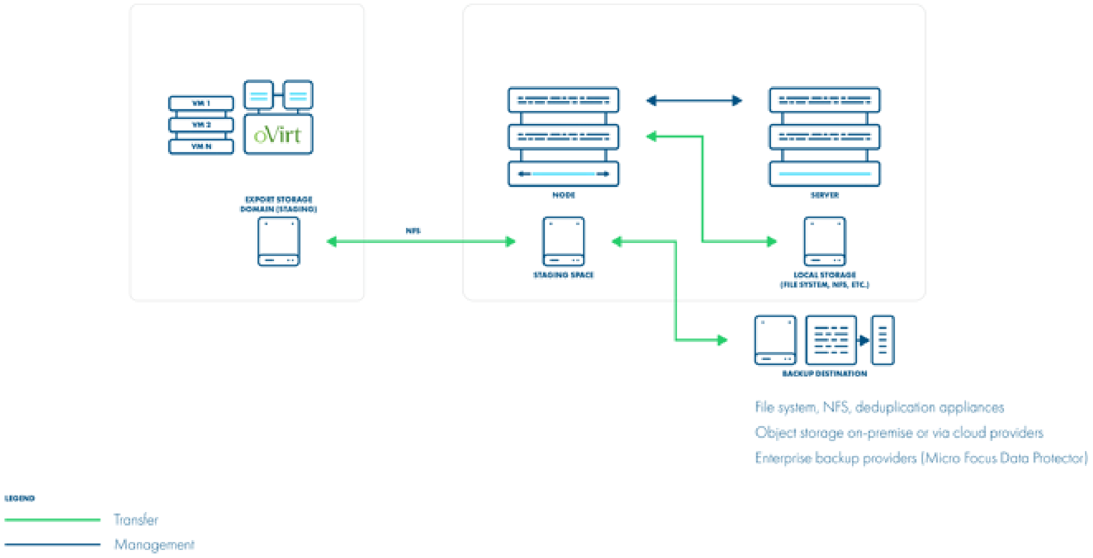
VM 1
VM 2
VM N
oVirt
EXPORT STORAGE DOMAIN (STAGING)
NFS
NODE
STAGING SPACE
SERVER
LOCAL STORAGE (FILE SYSTEM, NFS, ETC.)
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot is taken via API
- optional application consistency using pre/post snapshot command execution
- initial VM clone of the snapshot to the local repository is created
- cloned VM (data+metadata) exported by the manager to the Data Protector for Cloud Workloads staging space (visible as the export Storage Domain in managers UI)
- full backup only is supported
- restore is done to the export Storage Repository, the administrator needs to import the VM using manager UI

How to set up a backup with an export storage domain

oVirt 3.5.1+ environments (using API v3) require an export storage domain to be set up.

1. Add a backup storage domain in oVirt (which points to the NFS export in Data Protector for Cloud Workloads Node)
   - If you have multiple data centers, you need to enable the **Multi DC export** a checkbox in node configuration
     - Remember that you need to use named data centers in your oVirt environment to avoid name conflicts
     - An oVirt data center may use only one export storage domain, that is why you need to create sub-directories for each data center in the export path for example `/vprotect_data/dc01`, `/vprotect_data/dc02`, and use each sub-directory as NFS share for each data center export domain (separate NFS exports)
     - The export (staging) path in the above-mentioned scenario is still `/vprotect_data`, while `dc01` and `dc02` are data center names
     - Older versions of oVirt (3.5.x) require you to specify the mapping between DC names and export storage domains - you need to provide pairs of a DC name and a corresponding SD name in the node configuration (section `Hypervisor`)
   - If you have only one data center and don't want to use the multiple data centers export feature in the future, you can use the default settings and

setup NFS export pointing to the staging path (e.g. `/vprotect_data`)

- Note that the export must be set to use the UID and GID of `vprotect` user
- Example export configuration in `/etc/exports` to a selected hypervisor in the oVirt cluster:

  ```
  /vprotect_data    10.50.1.101(fsid=6,rw,sync,insecure,all_squash,an
  ```

  where `anonuid=993` and `anongid=990` should have the correct UID and GID returned by command:

  ```
  [root@vProtect3 ~]# id vprotect
  uid=993(vprotect) gid=990(vprotect) groups=990(vprotect)
  ```

2. Both import and export operations will be done using this NFS share – restore will be done directly to this storage domain, so you can easily import the backup into oVirt (shown below)
   - backups must be restored to the export path (the node automatically changes names to the original paths that are recognized by the oVirt manager.
3. When adding oVirt 4.0+ hypervisor managers, make sure you have a URL like the following:

   ```
   https://oVirt_MGR_HOST/ovirt-engine/api/v3
   ```

Instant restore

To use an instant restore feature, backup destination from which VM will be restored, has to be of a synthetic type. The restore process creates a NFS share on the Data Protector for Cloud Workloads node, later this share is attached to the RHV as a new storage domain. Then it creates a new virtual machine and attaches the disks from the newly created storage domain to it. To use instant restore you have to click the restore button in the instances list and choose the option **instant restore**.

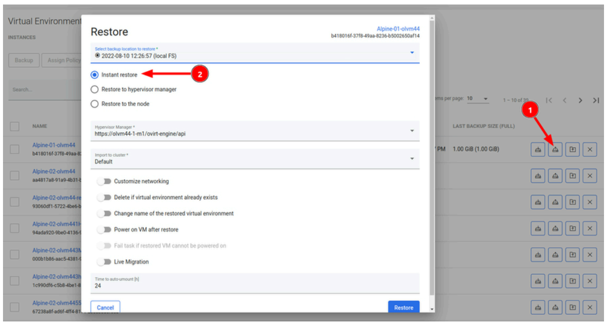

Restore

Alpine-01-olvm44
b418016f-37f8-49aa-8236-b5002650af14

Select backup location to restore *
● 2022-08-10 12:26:57 (local FS)

◉ Instant restore
○ Restore to hypervisor manager
○ Restore to the node

Hypervisor Manager *
https://olvm44-1-m1/ovirt-engine/api

Import to cluster *
Default

Customize networking
Delete if virtual environment already exists
Change name of the restored virtual environment
Power on VM after restore
Fail task if restored VM cannot be powered on
Live Migration

Time to auto-umount [h]
24

Live migration

You can enable the live migration option during instant restore. It will automatically start the disks migration to the chosen storage after the VM is restored and powered on.

Oracle Linux Virtualization Manager

Oracle Linux Virtualization Manager

General

For Oracle Linux Virtualization Manager (OLVM) 4+ environments you can use API v4 for invoking all backup-related tasks.

Import/export mode defines the way the backups and restores are done. OLVM (with API v4) supports 3 modes:

1. **Disk attachment**, which exports VM metadata (in OVF format) with separate disk files (in RAW format) via Proxy VM with the Node installed.
   - supports OLVM 4.0+
   - no incremental backup
   - proxy VM required in each cluster - used for the disk attachment process
2. **Disk image transfer**, which exports VM metadata (in OVF format) with disk snapshot chains as separate files (QCOW2 format):
   - supports OLVM 4.2+/oVirt 4.2.3+
   - supports incremental backup
   - disk images are transferred directly from the API (no Proxy VM required)
3. **SSH Transfer,** this method assumes that all data transfers are directly from the hypervisor over SSH

When adding OLVM hypervisor managers, use a URL similar to the following:

```
https://OLVM_MGR_HOST/ovirt-engine/api
```

> **Note:** a username for OLVM environments needs to be provided in the **user@domain** format - for example **admin@internal**. This user must have all

> permissions related to managing snapshots, creating/removing VMs, operating disks, and exporting data.

Backup Strategies

OLVM environments can be protected in several ways.

> **Note:**
>
> Different strategies require a node to be installed either as a VM in the environment that you back up or installed separately.
>
> All live snapshots are attempted with quiescing enabled. If the snapshot command fails because there is no compatible guest agent present, the live snapshot is re-initiated without the use-quiescing flag.

Disk attachment with Proxy VM

In this strategy, you have a VM called “Proxy VM” that invokes commands on your hypervisor manager to snapshot and attach drives of a specific VM to itself (Proxy VM). The proxy VM is able to read the data from the attached disk snapshots and forward them to the backup provider.

This strategy allows you to exclude drives from a backup that you do not need. Remember that you need to install 1 Proxy VM per cluster so that the drives the node tries to attach are reachable.

Drawback - no incremental backup for now.

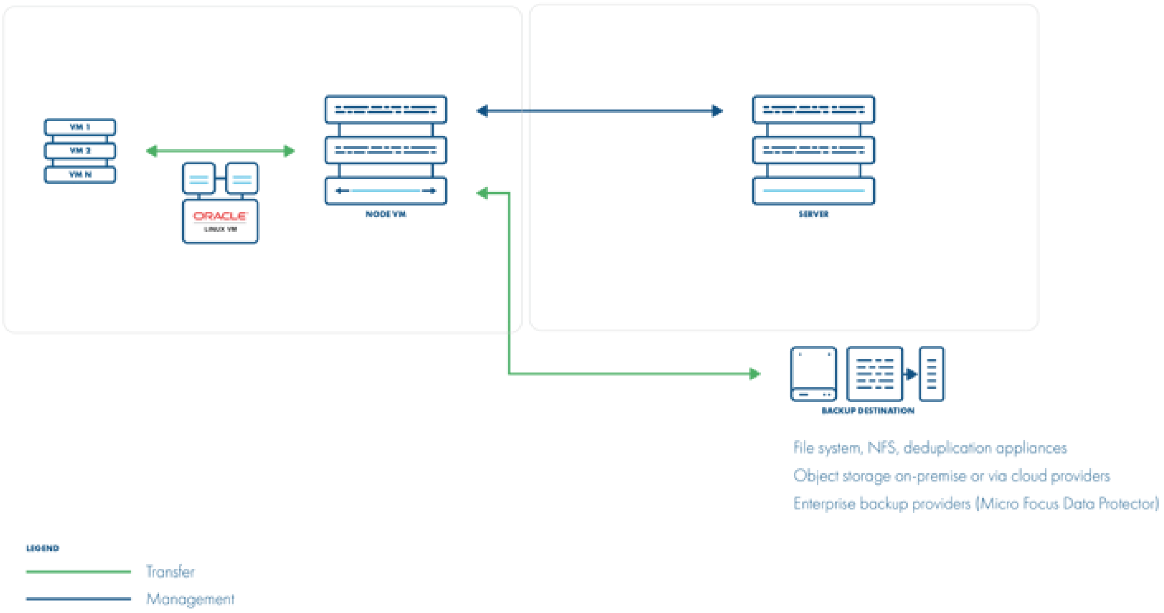
VM 1
VM 2
VM N
ORACLE
LINUX VM
NODE VM
SERVER
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot using hypervisor's API
- optionally FS freeze can be executed before snapshot can be executed (FS thaw once the snapshot is completed) if enabled and guest tools installed inside
- optional application consistency using pre/post snapshot command execution
- metadata exported from API
- snapshot disks are mounted one by one to the Proxy VM
- data read directly on the Proxy VM
- incremental backups are _**_not supported
- restore creates empty disks on the Proxy VM, imports merged data then recreates VM and reattaches volumes to the target VM

> **Note**: OLVM API v4 environments require Data Protector for Cloud Workloads Node to be installed in one of the VMs residing in the OLVM cluster. Data Protector for Cloud Workloads should automatically detect the VM with Data Protector for Cloud Workloads during the index operation.

The disk attachment mode requires `Virtio-SCSI` to be enabled on the Data Protector for Cloud Workloads Node VM (which can be enabled in `VM settings` → `Resource Allocation` → `VirtIO-SCSI Enabled` at the bottom).

During backup/restore operations, disks are transferred by attaching them to the proxy VM. This approach does not require an export storage domain to be set up.

Make sure you follow these steps: LVM setup on Data Protector for Cloud Workloads Node for disk attachment backup mode.

Disk image transfer API

This API appeared in OLVM 4.2 and allowed the export of individual snapshots directly from the OLVM manager. So instead of having to install multiple Proxy VMs, you can have a single external Node installation, which just invokes APIs via the OLVM manager.

This strategy supports incremental backups. Assuming you have OLVM 4.2 or newer – just add your manager to Data Protector for Cloud Workloads and setup is done. From a network perspective, it requires two additional ports to be opened - 54322 and 54323 - and your data to be pulled from the hypervisor manager.

Unfortunately, there are a few problems with the current architecture of this solution. The biggest issue is that all traffic passes via the OLVM manager, which may impact the transfer rates that you can achieve during the backup process. To put that into perspective – in disk attachment, you can basically read data as if it is a local drive, where it could potentially be deduplicated even before it is transferred to the backup destination.

> **Note:** From OLVM version 4.4.3, data is transferred directly from/to hosts.

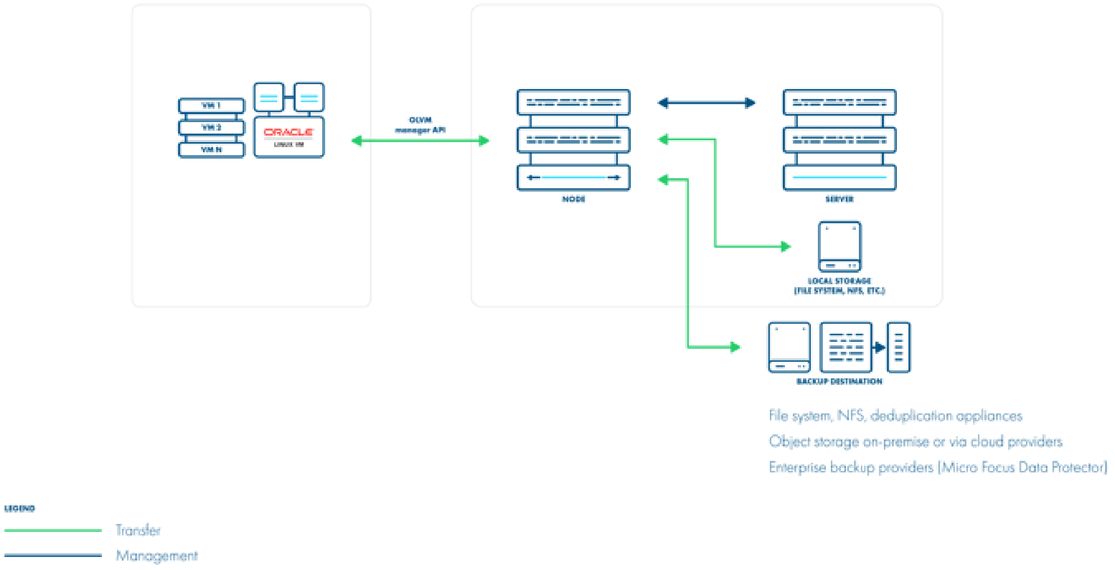
VM 1
VM 2
VM N
ORACLE
LINUX VM
OLVM
manager API
NODE
SERVER
LOCAL STORAGE
(FILE SYSTEM, NFS, ETC.)
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot using hypervisor's API
- optionally FS freeze can be executed before snapshot can be executed (FS thaw once the snapshot is completed) if enabled and guest tools installed inside
- optional application consistency using pre/post snapshot command execution
- supported for oVirt/RHV/OLVM 4.3+
- metadata exported from API
- data transfer initiated on the manager and actual data exported from the hypervisor using imageio API
- incremental backups use the same APIs, but requests for changed blocks only
- the last snapshot kept on the hypervisor for the next incremental backup (if at least one schedule assigned to the VM has a backup type set to incremental)
- restore recreates VM from metadata using API and imports merged chain of data for each disk using imageio API

Disk image transfer mode exports data directly using OLVM 4.2+ API. There is no need to set up an export storage domain or setup LVM. This mode uses snapshot chains provided by OLVM.

You may need to open communication for the additional port **54323** on the OLVM manager and **54322** on the OLVM hosts - it needs to be accessible from Data Protector for Cloud Workloads Node. Also, make sure that your **ovirt-imageio-proxy** services are running and properly configured (you can verify it by trying to upload images with OLVM UI).

Follow the steps in this section: [Full versions of libvirt/qemu packages installation](#).

SSH transfer

This is an enhancement to the disk image transfer API strategy. It allows Data Protector for Cloud Workloads to use OLVM API v4.2+ (HTTPS connection to OLVM manager) only to collect metadata. Backup is done over SSH directly from the hypervisor (optionally using netcat for transfer), import is also using SSH (without the netcat option). There is no need to install a node on the OLVM environment. This method can significantly boost backup transfers and supports incremental backups.

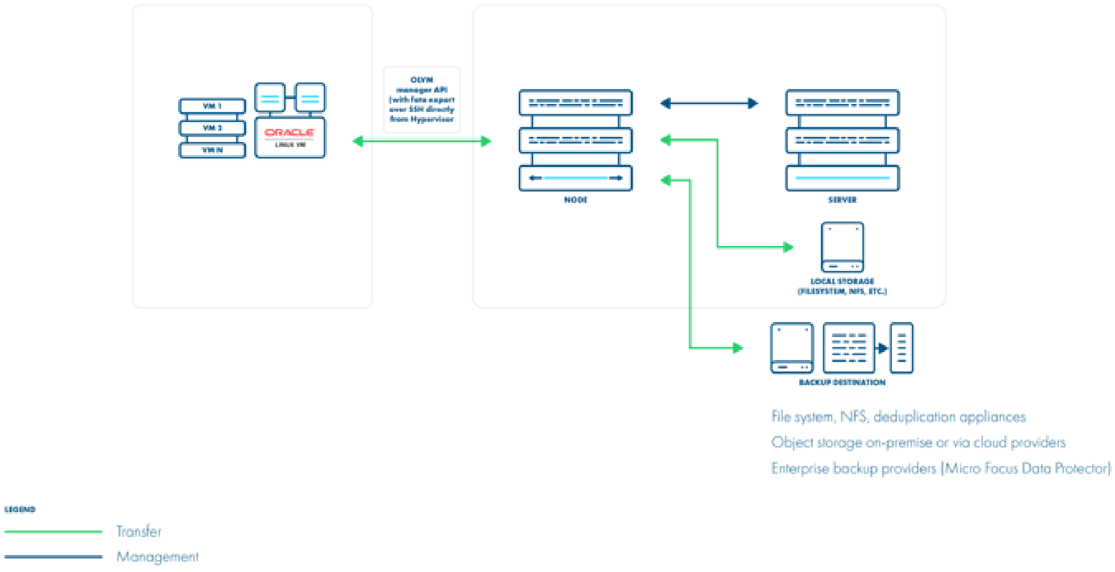
VM 1
VM 2
VM N
ORACLE
LINUX VM
OLVM manager API (with fata export over SSH directly from Hypervisor
NODE
SERVER
LOCAL STORAGE (FILESYSTEM, NFS, ETC.)
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot using hypervisor's API

- optionally FS freeze can be executed before snapshot can be executed (FS thaw once the snapshot is completed) if enabled and guest tools installed inside
- optional application consistency using pre/post snapshot command execution • metadata exported from API
- data transfer via SSH (optional using netcat) - the full chain of disk snapshot files for each disk o if LVM-based storage is used, then node activates volumes if necessary to read data o if Gluster FS is used, then disk files are copied directly
- incremental backup export just sub-chain of QCOW2-deltas snapshots since last stored snapshot
- the last snapshot kept on the hypervisor for the next incremental backup (if at least one schedule assigned to the VM has a backup type set to incremental)
- restore recreates VM with empty storage from metadata using API and imports merged data over SSH to appropriate location on a hypervisor

This method assumes that all data transfers are directly from the hypervisor over SSH. This means that after adding OLVM manager and detecting all available hypervisors - **you also need to provide SSH credentials or SSH keys for each of the hypervisors**. You can also use SSH public key authentication.

Instant restore

To use an instant restore feature, backup destination from which VM will be restored, has to be of a synthetic type. The restore process creates a NFS share on the Data Protector for Cloud Workloads node, later this share is attached to the RHV as a new storage domain. Then it creates a new virtual machine and attaches the disks from the newly created storage domain to it. To use instant restore you have to click the restore button in the instances list and choose the option **instant restore**.

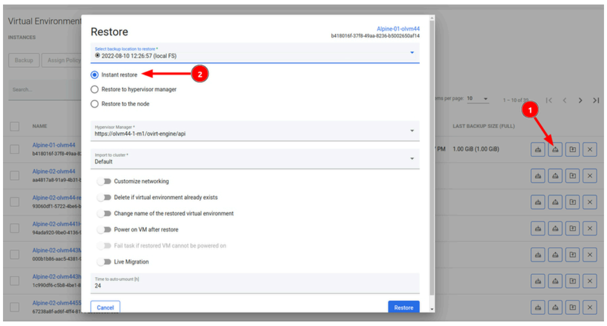

Restore

Alpine-01-olvm44
b418016f-37f8-49aa-8236-b5002650af14

Select backup location to restore *
● 2022-08-10 12:26:57 (local FS)

● Instant restore
○ Restore to hypervisor manager
○ Restore to the node

Hypervisor Manager *
https://olvm44-1-m1/ovirt-engine/api

Import to cluster *
Default

Customize networking
Delete if virtual environment already exists
Change name of the restored virtual environment
Power on VM after restore
Fail task if restored VM cannot be powered on
Live Migration

Time to auto-umount [h]
24

Live migration

You can enable the live migration option during instant restore. It will automatically start the disks migration to the chosen storage after the VM is restored and powered on.

Oracle VM

Oracle VM

Deployment in Oracle VM environment

The Oracle VM environment requires you to create storage used for VM export. The export storage repository should also be accessible by Data Protector for Cloud Workloads Node in its staging directory. This implies that storage space doesn't have to be exported by Data Protector for Cloud Workloads Node - it can be mounted from an external source. The only requirement is to have it visible from both OVM hosts and Node itself. Keep in mind that ownership of the files on the share should allow both Data Protector for Cloud Workloads and OVM to read and write files.

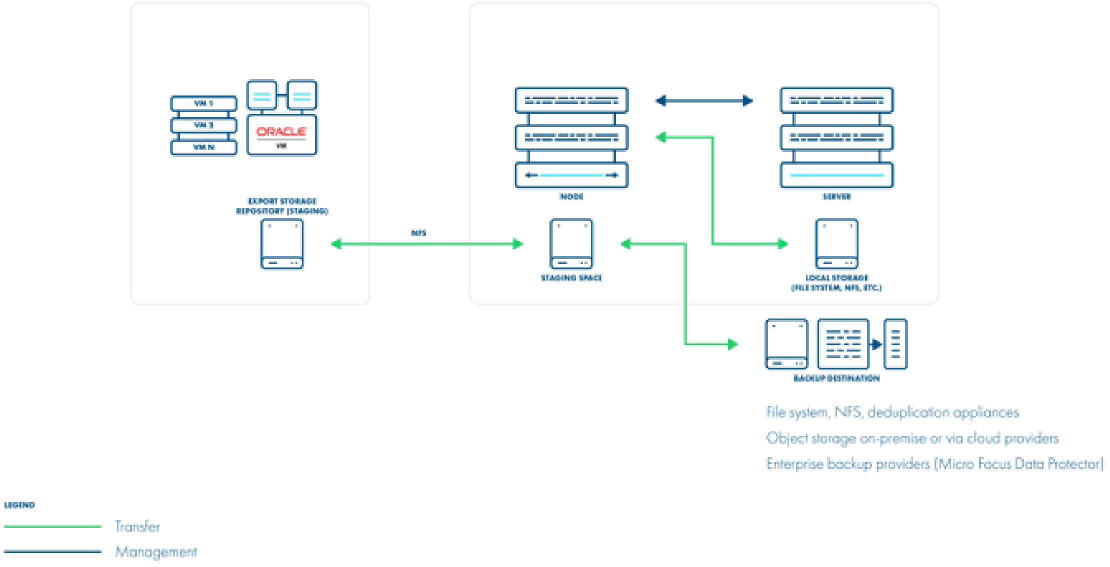
VM 1
VM 2
VM N
ORACLE
VM
EXPORT STORAGE REPOSITORY (STAGING)
NFS
NODE
SERVER
STAGING SPACE
LOCAL STORAGE (FILE SYSTEM, NFS, ETC.)
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- crash-consistent snapshot is taken by OVM during hot-clone of the VM

- data + metadata exported by the manager to the Data Protector for Cloud Workloads staging space (visible as the export Storage Repository in OVM)
- full backup only is supported
- OVM exports are limited to 1 task per Storage Repository being used by VM (this is the _**_limitation of OVM)
- restore is done to the export Storage Repository, the administrator needs to clone the VM using manager UI

> **Note:**
>
> - If the virtual machines are running on NFS storage, you must shut down the Oracle virtual machines to perform the backup
> - Make sure the NFS share have the Data Protector for Cloud Workloads user UID and GID
> - The directory under / vprotect_data needs to be the same name as the OVS server pool name
> - Oracle VM needs to disable services (nfslock , rpcbind**)**
> - Restore of VM is multi-step
> - Restore to the staging space on vNode
> - Move VM from the staging space to the Oracle protection repository
> - Migrate the VM into the Oracle server pool

Oracle VM environments require storage repositories to be defined for each server pool and must be mounted on Data Protector for Cloud Workloads Node.

1. Create a repository from NFS share on Data Protector for Cloud Workloads Node
   - One server pool should have a separate subdirectory in the export path for example `/vprotect_data/pool01`, `/vprotect_data/pool2` - each subdirectory is a separate NFS share
   - The export (staging) path in the above-mentioned scenario is still `/vprotect_data`, while `pool01` and `pool02` are server pool names
   - Specify mapping between server pool names and storage repository names in the hypervisor manager configuration

- Note that the export must be set to use the UID and GID of the `vprotect` user
- Example export configuration in `/etc/exports` to the selected hypervisor in the RHV cluster:

  ```
  /vprotect_data/pool01    10.50.1.101(fsid=6,rw,sync,insecure,
  all_squash,anonuid=993,anongid=990)
  /vprotect_data/pool02    10.50.1.102(fsid=7,rw,sync,insecure,
  all_squash,anonuid=993,anongid=990)
  ```

  where `anonuid=993` and `anongid=990` should have the correct UID and GID returned by command:

  ```
  [root@vProtect3 ~]# id vprotect
  uid=993(vprotect) gid=990(vprotect) groups=990(vprotect)
  ```

2. Both import and export operations will be done using these NFS shares – restore will be done directly to this storage domain, so you can easily import the backup into the Oracle VM environment
   - Backups must be restored to the export path (the node automatically changes the names to the original paths that are recognized by the OVM manager.

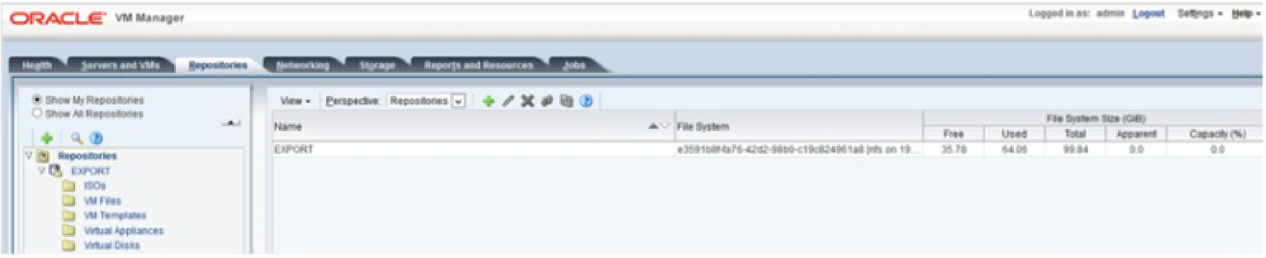

Name	File System	File System Size (GiB)				
		Free	Used	Total	Apparent	Capacity (%)
EXPORT	e3591b8f-fa76-42d2-98b0-c19c824961a8 (nfs on 19...	35.78	64.06	99.84	0.0	0.0

Example - How to configure OVM protection with PowerProtect DD

- Create a DDBoost device

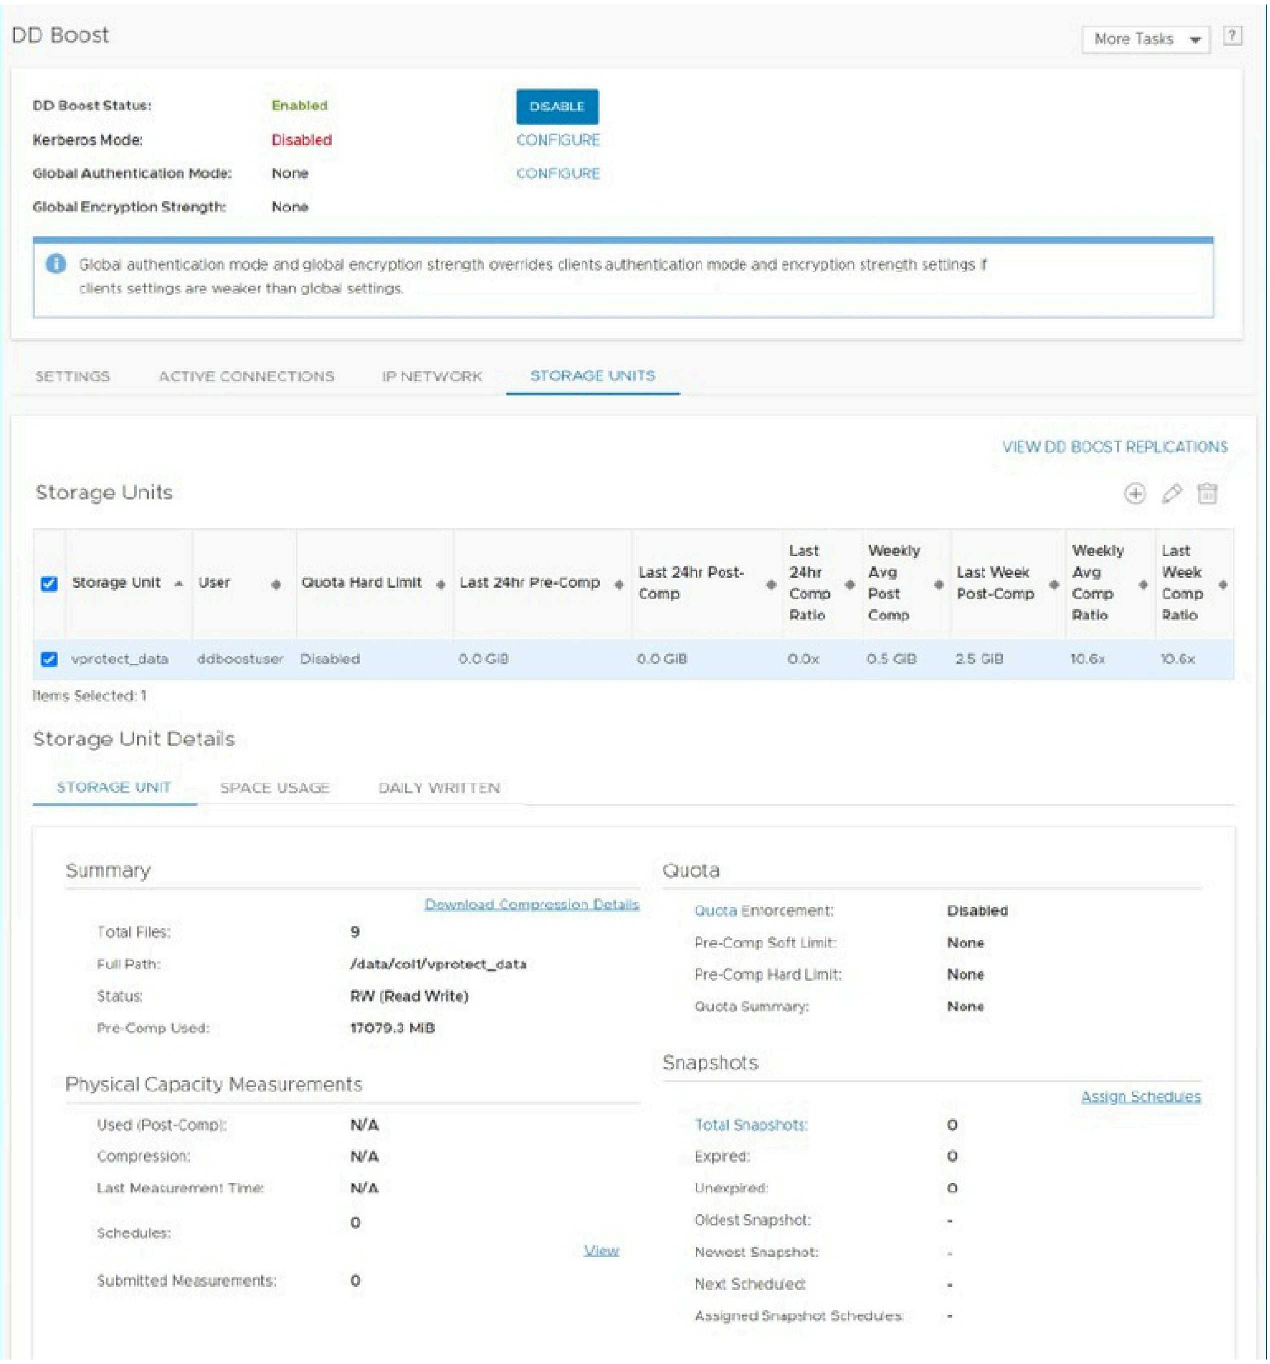

DD Boost

DD Boost Status:	Enabled	DISABLE
Kerberos Mode:	Disabled	CONFIGURE
Global Authentication Mode:	None	CONFIGURE
Global Encryption Strength:	None	

Global authentication mode and global encryption strength overrides clients authentication mode and encryption strength settings if clients settings are weaker than global settings.

Storage Units

☑	Storage Unit	User	Quota Hard Limit	Last 24hr Pre-Comp	Last 24hr Post-Comp	Last 24hr Comp Ratio	Weekly Avg Post Comp	Last Week Post-Comp	Weekly Avg Comp Ratio	Last Week Comp Ratio
☑	vprotect_data	ddboostuser	Disabled	0.0 GiB	0.0 GiB	0.0x	0.5 GiB	2.5 GiB	10.6x	10.6x

Items Selected: 1

Storage Unit Details

Summary

Total Files:	9
Full Path:	/data/col1/vprotect_data
Status:	RW (Read Write)
Pre-Comp Used:	17079.3 MiB

Physical Capacity Measurements

Used (Post-Comp):	N/A
Compression:	N/A
Last Measurement Time:	N/A
Schedules:	0
Submitted Measurements:	0

Quota

Quota Enforcement:	Disabled
Pre-Comp Soft Limit:	None
Pre-Comp Hard Limit:	None
Quota Summary:	None

Snapshots

Total Snapshots:	0
Expired:	0
Unexpired:	0
Oldest Snapshot:	-
Newest Snapshot:	-
Next Scheduled:	-
Assigned Snapshot Schedules:	-

- Create NFS share

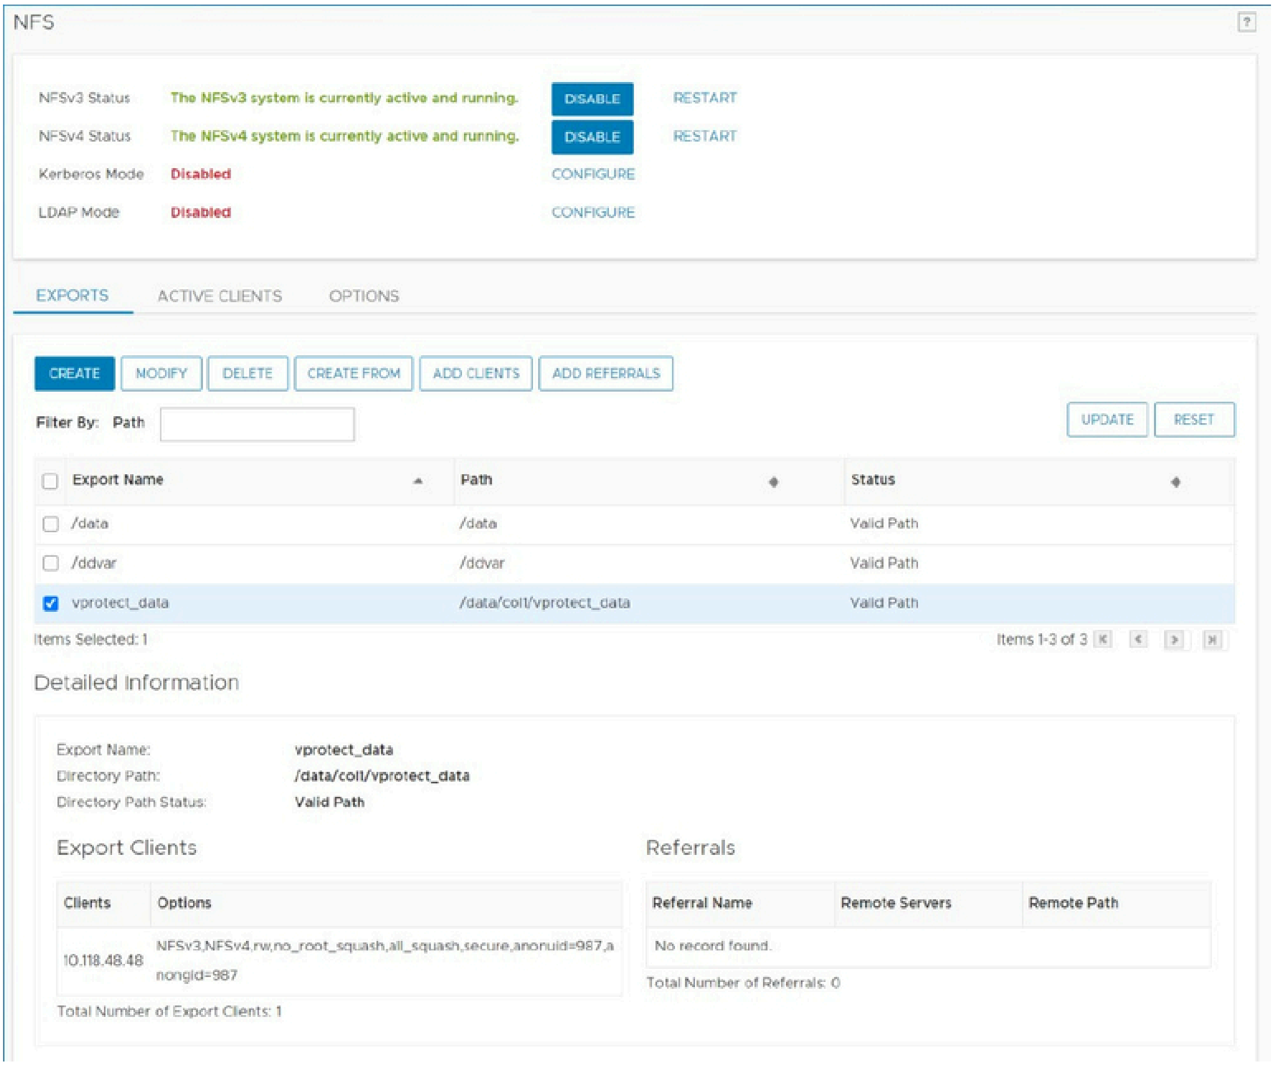

NFSv3 Status The NFSv3 system is currently active and running.

NFSv4 Status The NFSv4 system is currently active and running.

Kerberos Mode Disabled

LDAP Mode Disabled

Filter By: Path

☐	Export Name	Path	Status
☐	/data	/data	Valid Path
☐	/ddvar	/ddvar	Valid Path
☑	vprotect_data	/data/col1/vprotect_data	Valid Path

Items Selected: 1

Items 1-3 of 3

Detailed Information

Export Name: vprotect_data
Directory Path: /data/col1/vprotect_data
Directory Path Status: Valid Path

Export Clients

Clients	Options
10.118.48.48	NFSv3,NFSv4,rw,no_root_squash,all_squash,secure,anonuid=987,anongid=987

Total Number of Export Clients: 1

Referrals

Referral Name	Remote Servers	Remote Path
No record found.		

Total Number of Referrals: 0

- Mount the DDBoost device on Data Protector for Cloud Workloads Node

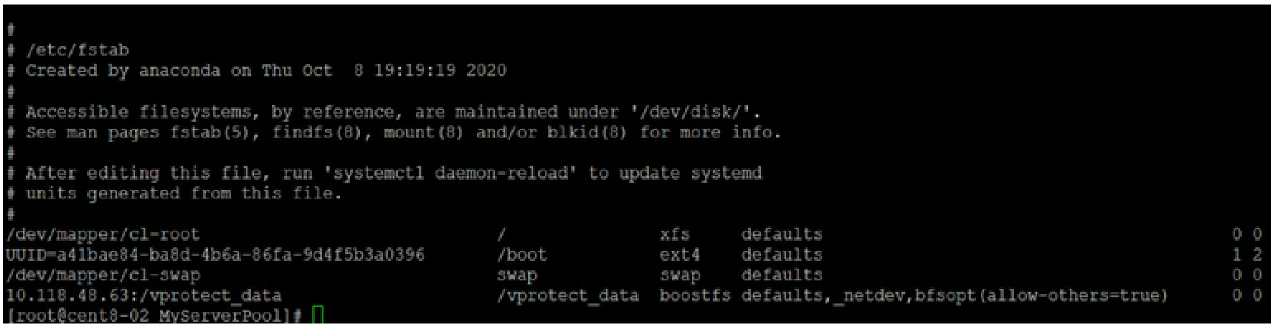

```
#
# /etc/fstab
# Created by anaconda on Thu Oct  8 19:19:19 2020
#
# Accessible filesystems, by reference, are maintained under '/dev/disk/'.
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info.
#
# After editing this file, run 'systemctl daemon-reload' to update systemd
# units generated from this file.
#
/dev/mapper/cl-root                          /              xfs     defaults                                         0 0
UUID=a41bae84-ba8d-4b6a-86fa-9d4f5b3a0396    /boot          ext4    defaults                                         1 2
/dev/mapper/cl-swap                          swap           swap    defaults                                         0 0
10.118.48.63:/vprotect_data                  /vprotect_data boostfs defaults,_netdev,bfsopt(allow-others=true)       0 0
[root@cent8-02 MyServerPool]# 
```

- Create symbolic links for the OVM Pool name (MyServerPool) to BoostFS mount point command, for example: ln -s /vprotect_data /MyServerPool

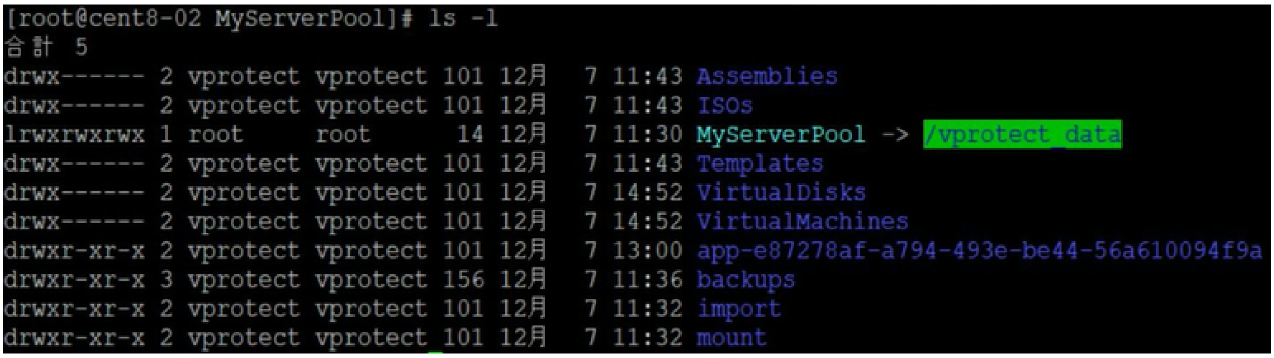

```
[root@cent8-02 MyServerPool]# ls -l
合計 5
drwx------ 2 vprotect vprotect 101 12月  7 11:43 Assemblies
drwx------ 2 vprotect vprotect 101 12月  7 11:43 ISOs
lrwxrwxrwx 1 root     root      14 12月  7 11:30 MyServerPool -> /vprotect_data
drwx------ 2 vprotect vprotect 101 12月  7 11:43 Templates
drwx------ 2 vprotect vprotect 101 12月  7 14:52 VirtualDisks
drwx------ 2 vprotect vprotect 101 12月  7 14:52 VirtualMachines
drwxr-xr-x 2 vprotect vprotect 101 12月  7 13:00 app-e87278af-a794-493e-be44-56a610094f9a
drwxr-xr-x 3 vprotect vprotect 156 12月  7 11:36 backups
drwxr-xr-x 2 vprotect vprotect 101 12月  7 11:32 import
drwxr-xr-x 2 vprotect vprotect 101 12月  7 11:32 mount
```

- Create a Storage Server for DD NFS Share

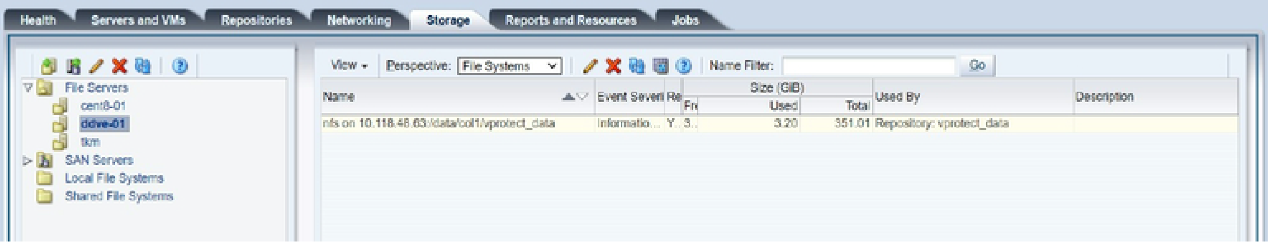

Name	Event Severi	Re	Size (GiB)			Used By	Description
			Fr	Used	Total		
nfs on 10.118.48.63:/data/col1/vprotect_data	Informatio...	Y..	3..	3.20	351.01	Repository: vprotect_data	

- Create a Repository using DD

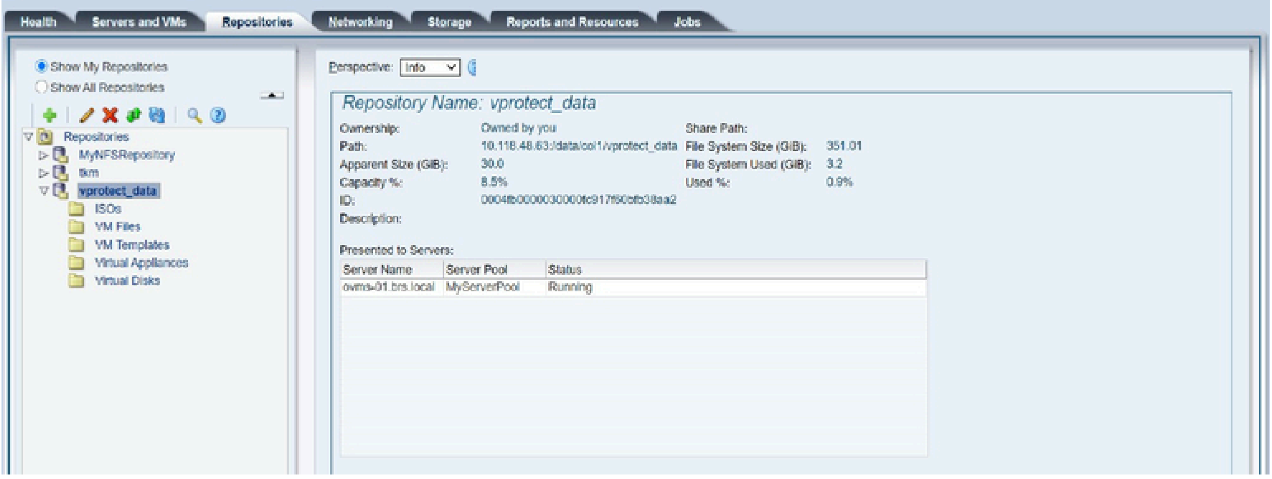

Perspective: Info

Repository Name: vprotect_data

Ownership: Owned by you

Path: 10.118.48.63:/data/col1/vprotect_data

Apparent Size (GiB): 30.0

Capacity %: 8.5%

ID: 0004fb0000030000fc917f60bfb38aa2

Description:

Share Path:

File System Size (GiB): 351.01

File System Used (GiB): 3.2

Used %: 0.9%

Presented to Servers:

Server Name	Server Pool	Status
ovms-01.brs.local	MyServerPool	Running

- Add the OVM Hypervisor Manager to Data Protector for Cloud Workloads

> ***Note:*** *You can get the "Storage Repository ID" from the "OVM repositories" menu shown in the previous step*

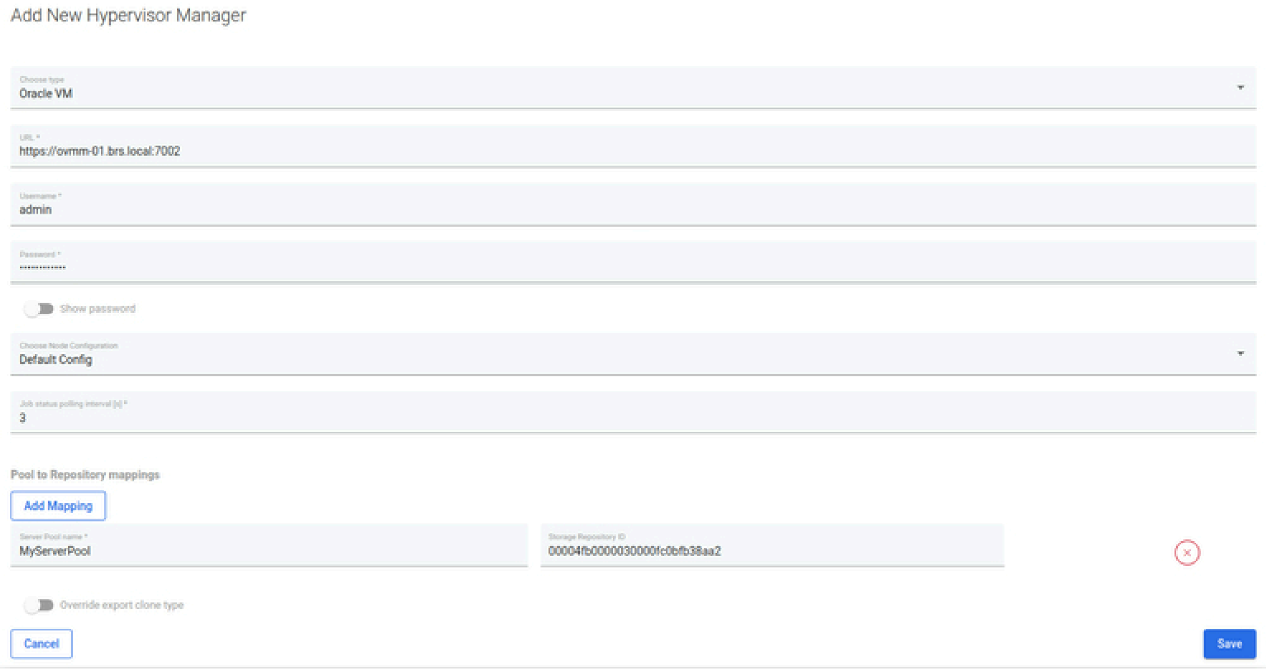

Add New Hypervisor Manager

Choose type
Oracle VM

URL *
https://ovmm-01.brs.local:7002

Username *
admin

Password *
••••••••••••••

Show password

Choose Node Configuration
Default Config

Job status polling interval [s] *
3

Pool to Repository mappings

Server Pool name *
MyServerPool

Storage Repository ID
00004fb0000030000fc0bfb38aa2

Override export clone type

Proxmox VE

Proxmox VE

SSH Transfer

SSH Transfer strategy:

- supports Proxmox 5.0+
- supports only QCOW2 disk images
- supports incremental backups
- supports over iSCSI

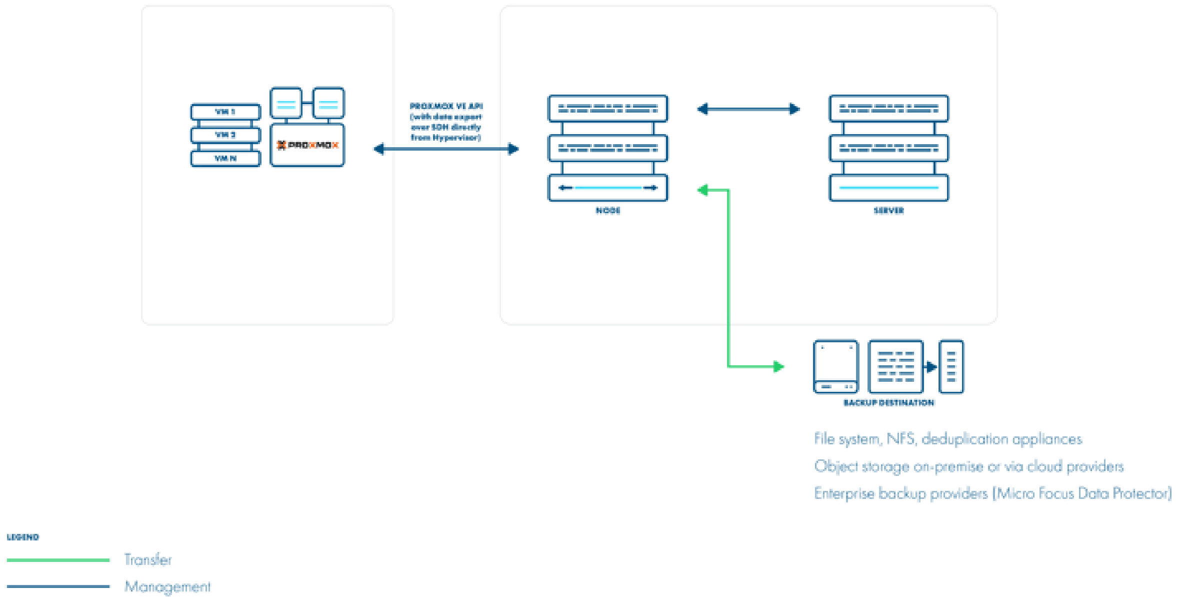
VM 1
VM 2
VM N
PROXMOX
PROXMOX VE API (with data export over SSH directly from Hypervisor)
NODE
SERVER
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup is done by transferring QCOW2 disk images over SSH directly from the hypervisor (optionally using Netcat for transfer). Metadata is backed up only in the full backup. This method supports incremental backups where the last snapshot is required for the next incremental backups. The resulting backup has separate files for each disk + metadata, so you have the option to exclude specific drives as well.

Backup Process

- QCOW2 - file-based storage only
- crash-consistent snapshots created using hypervisor CLI over SSH
- optionally FS freeze can be executed before snapshot can be executed (FS thaw once the snapshot is completed) if enabled and guest tools installed inside
- QCOW2 snapshots mounted locally on the hypervisor and exported via SSH (optionally with netcat)
- for incremental backups, both last and currently created snapshots are mounted and block-difference is sent via SSH
- metadata exported via SSH • restore imports metadata and overwrites empty disks with data from a merged backup over SSH

Export storage repository

The Proxmox virtual environment requires you to create storage used for VM export. Export storage should also be accessible to Data Protector for Cloud Workloads Node in its staging directory. This implies that storage space doesn't have to be exported by Data Protector for Cloud Workloads Node - it can be mounted from an external source. The only requirement is to have it visible from both Proxmox VE hosts and the Node itself. Keep in mind that ownership of the files on the share should allow both Data Protector for Cloud Workloads and Proxmox VE to read and write files.

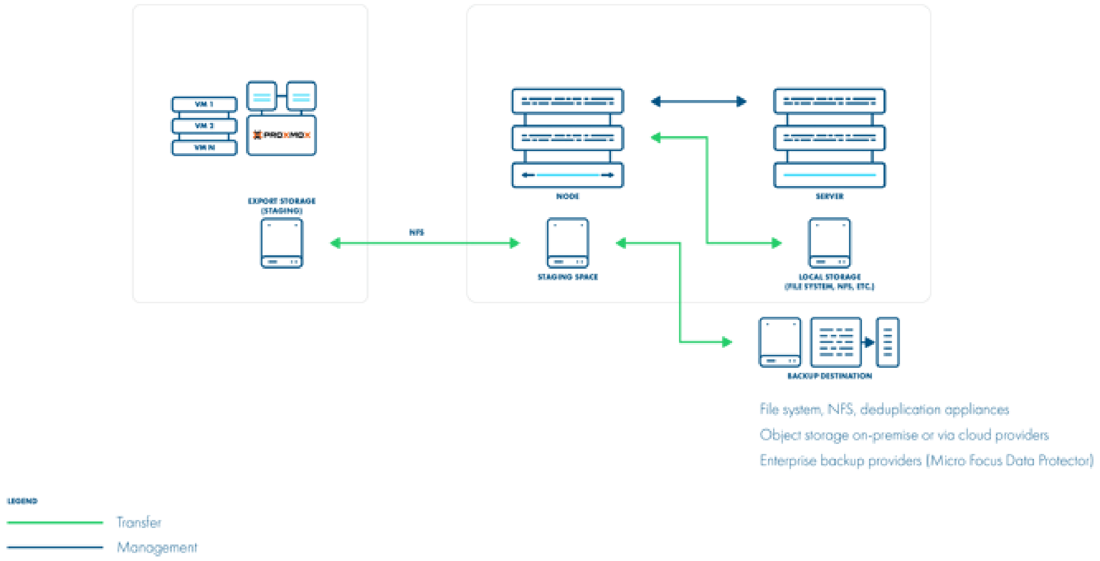
VM 1
VM 2
VM N
PROXMOX
EXPORT STORAGE (STAGING)
NFS
NODE
SERVER
STAGING SPACE
LOCAL STORAGE (FILE SYSTEM, NFS, ETC.)
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- native VM export is done to the prepared NFS storage (staging space on Data Protector for Cloud Workloads) using SSH access to the hypervisor
- crash-consistency - FS freeze used for VMs, LVM snapshot for containers
- optional application consistency using pre/post export command execution for VMs (pre/post snapshot) for containers
- data and metadata are in a single VMA image
- only full backups are supported • restore imports VMA image to the hypervisor

How to set up export storage repository backup

Proxmox virtual environments require backup storage to be defined on each server. This storage must be a location accessible from Data Protector for Cloud Workloads Node (the simplest setup, when you use only 1 node, is to create NFS share for the staging path on Data Protector for Cloud Workloads Node)

1. Create storage from NFS share (Content-type: **only VZDump**)

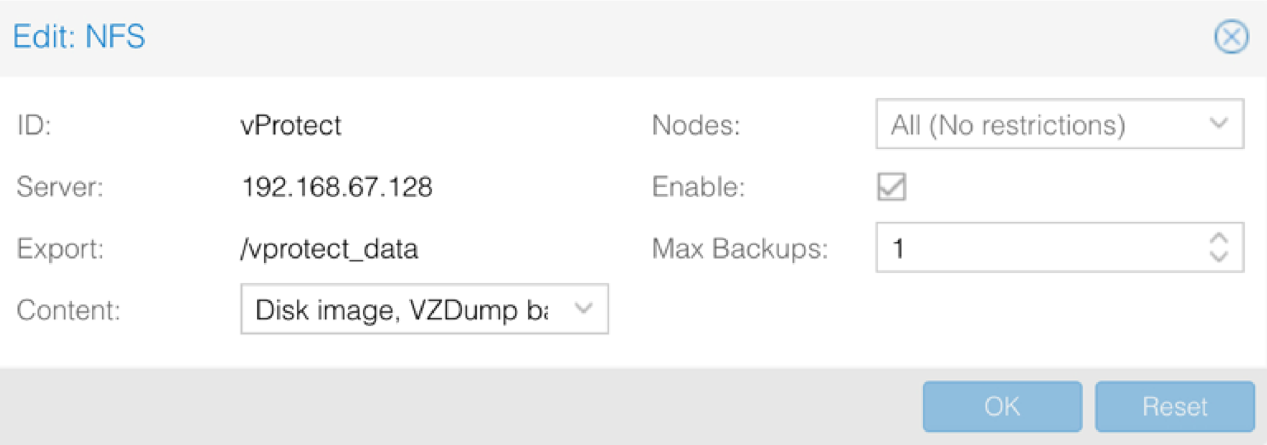

Edit: NFS

ID:	vProtect	Nodes:	All (No restrictions)
Server:	192.168.67.128	Enable:	☑
Export:	/vprotect_data	Max Backups:	1
Content:	Disk image, VZDump ba		

- Export share must be set to use the UID and GID of the `vprotect` user
- Example export configuration in `/etc/exports` to the selected hypervisor in the cluster:

  ```
  /vprotect_data     PROXMOX_HOSTS(fsid=6,rw,sync,insecure,all_squash,
  anonuid=993,anongid=990)
  ```

  where `anonuid=993` and `anongid=990` should have correct UID and GID returned by command:

  ```
  [root@vProtect3 ~]# id vprotect
  uid=993(vprotect) gid=990(vprotect) groups=990(vprotect)
  ```

- Both import and export operations will be done using these NFS shares – restore will be done directly to this storage domain, so you can easily import the backup into Proxmox VE
  - backups must be restored to the export path (the node automatically changes names to the original paths that are recognized by Proxmox VE.
- A name for storage must be provided later in the **Virtual Environments** → **Infrastructure** → **Hypervisors**

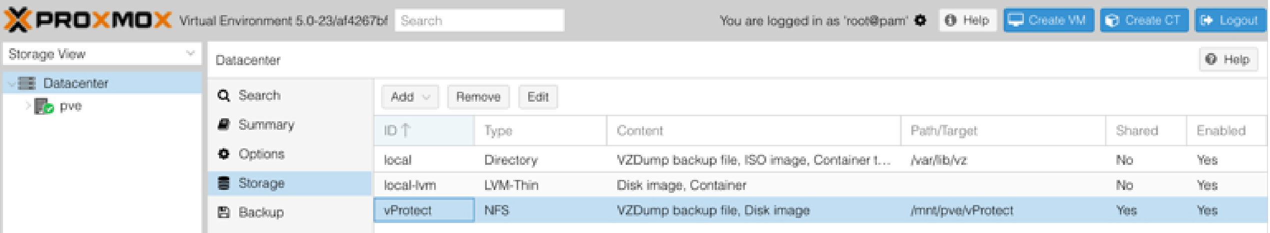

ID	Type	Content	Path/Target	Shared	Enabled
local	Directory	VZDump backup file, ISO image, Container t...	/var/lib/vz	No	Yes
local-lvm	LVM-Thin	Disk image, Container		No	Yes
vProtect	NFS	VZDump backup file, Disk image	/mnt/pve/vProtect	Yes	Yes

File-level restore support for VMA images

Prepare the VMA extractor on Data Protector for Cloud Workloads Node:

- build VMA extractor like this (requires Internet on the **node**):

  ```
  cd /opt/vprotect/scripts/vma
  ./setup_vma.sh
  ```

KVM/Xen

KVM/Xen

Data Protector for Cloud Workloads access KVM/Xen (stand-alone libvirt) hosts over SSH. The Data Protector for Cloud Workloads Node can be installed outside of the environment.

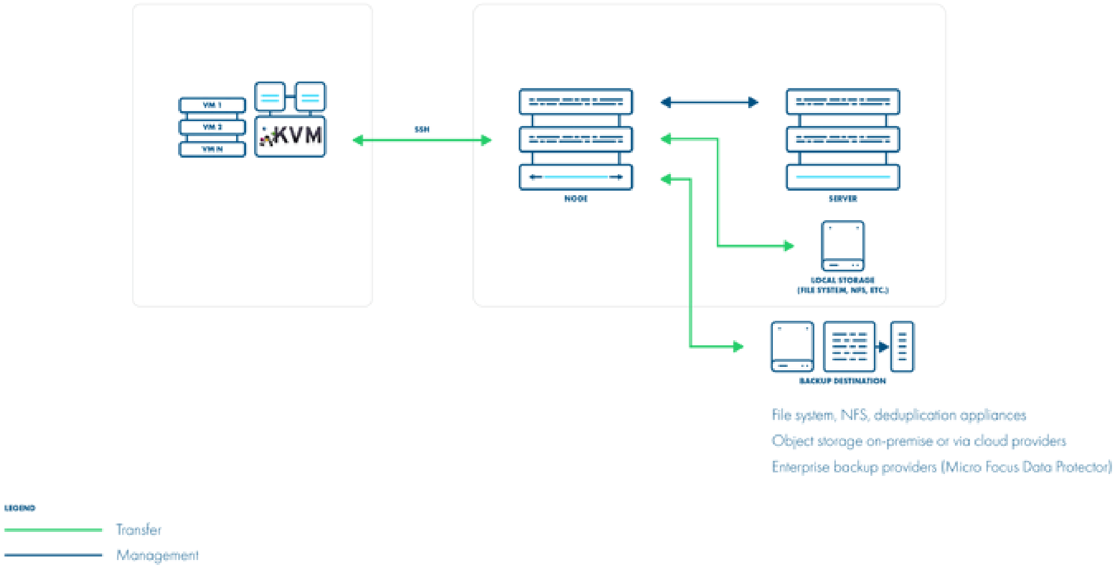
VM 1
VM 2
VM N
KVM
SSH
NODE
SERVER
LOCAL STORAGE
(FILE SYSTEM, NFS, ETC.)
BACKUP DESTINATION
File system, NFS, deduplication appliances
Object storage on-premise or via cloud providers
Enterprise backup providers (Micro Focus Data Protector)
LEGEND
Transfer
Management

Backup Process

- direct access to the hypervisor over SSH
- crash-consistent snapshot taken directly using virsh (QCOW2/RAW file), lvcreate (LVM), rbd snapshot for Ceph (separate call for each storage backend) For QCOW2/RAW file, virsh snapshot-create-as is used when VM is running - otherwise, qemu-img create is used
- optional application consistency using pre/post snapshot command execution
- QCOW2/RAW-file/LVM data exported over SSH (optionally with netcat)
- Ceph RBD data exported using rbd export or RBD-NBD when incremental is used If last stored snapshot is not missing, snapshot diffs are downloaded using