opentext™

Data Protector for Cloud Workloads 7.2.0

July 2025

Table of Contents

        - Nutanix Volume Groups
    - Backup & Restore
- Administration
    - Dashboard
    - Virtual Environments
        - Instances
            - Backup on-demand
            - Restore on-demand
            - Snapshot Management
        - Infrastructure
        - Backup SLAs
            - Policies
            - Schedules
        - Snapshot SLAs
            - Policies
            - Schedules
        - Recovery Plans
            - Policies
            - Schedules
        - Mounted Backups (File-level Restore)
    - Storage Providers
        - Instances
            - Backup on-demand
            - Restore on-demand
        - Infrastructure
        - Backup SLAs
            - Policies
            - Schedules
        - Snapshot SLAs
            - Policies

What's new in this release

Enhanced Filtering & Insights

Enjoy advanced filtering options in list views and gain deeper insights into backup attempts, particularly when using the retry option.

Reports enhancements

Grouped backup retries ensure only the final backup status is included in email reports and dashboards.

Removed support for standalone KVM

Starting with this release, standalone KVM environments are no longer supported. Existing configurations will continue to function as expected, but the addition of new standalone KVM instances is now disabled.

Removed support for Huawei Fusion Compute

Starting with this release, Huawei Fusion Compute environments are no longer supported. Existing configurations will continue to function as expected, but the addition of new Huawei Fusion Compute instances is now disabled.

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

Data Protector for Cloud Workis a data protection solution for virtual environments, storage, M365, and endpoints. It provides a centralized and automated solution for data protection. It can be deplo

Supported backup platforms

Virtual Machines:

- Nutanix
- HC3/Scale
- Red Hat Virtualization
- oVirt
- Oracle Linux Virtualization Manager
- Nutanix Acropolis Hypervisor (AHV)
- OpenNebula
- XCP-ng with CBT support
- Virtuozzo
- Proxmox VE
- Oracle VM
- OpenStack
- OpenShift Virtualization

Containers:

- Kubernetes (deployment-level protection for Persistent Volumes)
- Red Hat OpenShift (deployment-level protection for Persistent Volumes)
- Proxmox VE

Cloud solutions:

- Amazon EC2
- Google Cloud Platform / Google Computer Engine
- Azure Cloud
- Microsoft 365

Storage:

- Ceph RBD (with snapshot difference support)
- Nutanix Files (with Changed-File Tracking)
- Nutanix Volume Groups (with Changed-Region Tracking)

Applications:

- the generic backup mechanism for the custom backup processes
- templates for commonly used applications

Advanced backup features:

- Snapshot Management (Copy Data Management)
- Snapshot consistent technology (quiesced/application-consistent snapshots or filesystem freeze)
- Pre/post snapshot remote command execution on VM to enable operations such as database quiesce
- Change Block Tracking (CBT) and Change file tracking (CFT) for faster incremental backups
- VM disks exclusion option
- Backup schedules
- 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:
    - scalability
    - automatic task load balancing
    - suitable for geographically distributed environments
- Built-in Storware Backup & Recovery database backup

Recovery features

- File-level restore using mountable backups
    - directly via a web browser
    - transfer to the remote host over SSH and WinRM
- Backup mount - RAW disks shareable over iSCSI (for direct block-access to your backup data)
- Recovery plans for automated Disaster Recovery
    - on-demand restore of multiple instances
    - automated schedules for recovery testing
- Customizable networking and disk layout during a restore
- oVirt/RHV/OLVM instant restore
- Individual disk recovery

Supported backup destinations

- File-based:
    - A synthetic backup provider using XFS, NFS 4.2
    - Any mounted file system (local or remote, especially GlusterFS for replication, CephFS, NFS, SMB, and many more)

    - Rubrik Managed Volumes
- 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
- 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 features

- Role-based access control (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

Other features

- Web-based (HTML5) central management portal
- Advanced reporting
    - administrative portal
    - e-mail

- Event notifications using e-mail, Slack, or custom API call
- Command Line Interface (CLIfor advanced administrators)
- Open API for 3rd party software integration (REST API)
- LDAP authentication
- OpenStack Horizon plugin
- oVirt/RHV/OLVM console integration

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

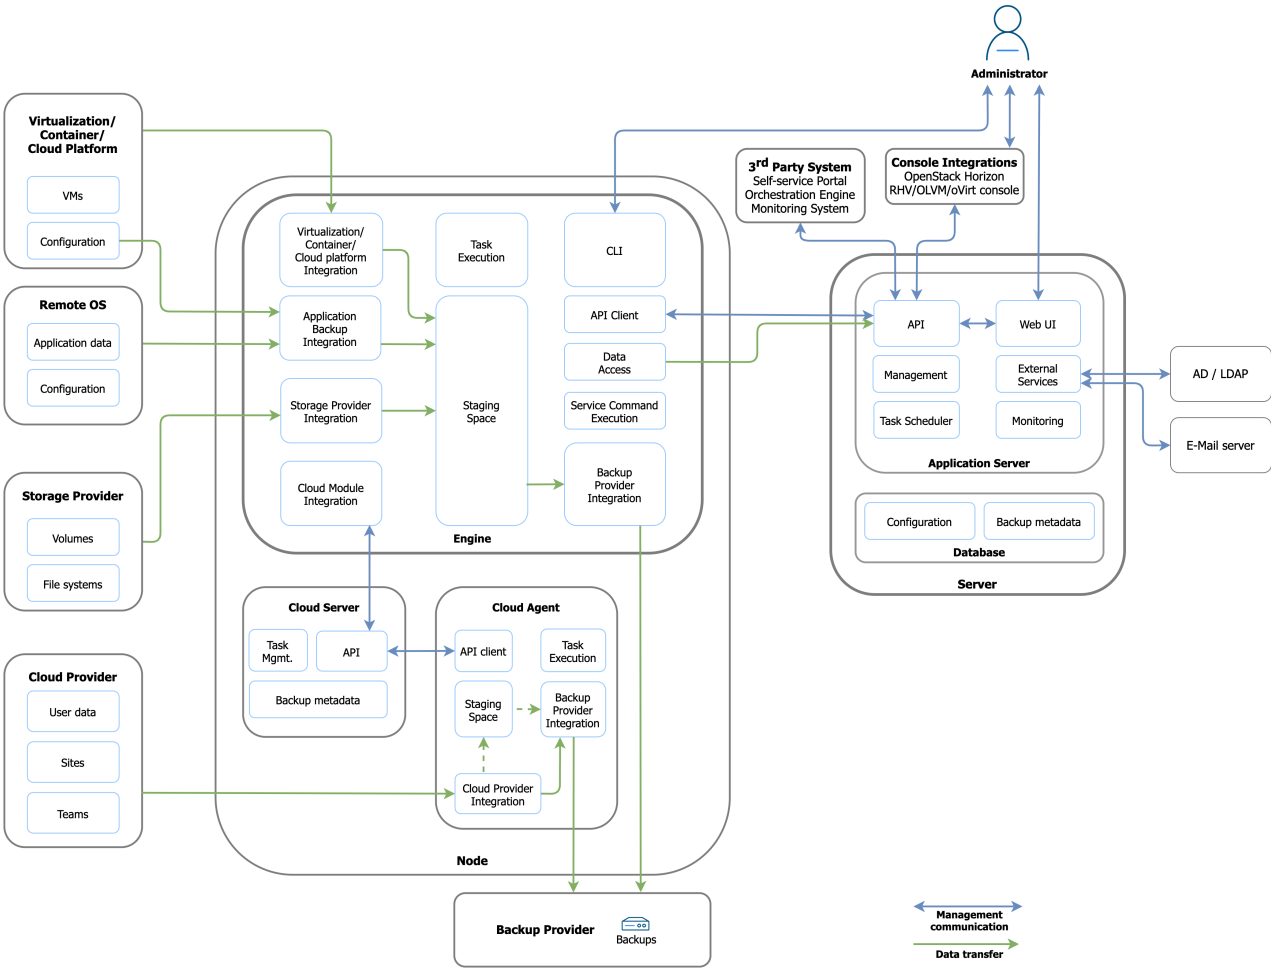

Data Protector for Cloud Workloads consists of 2 main components:

- **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

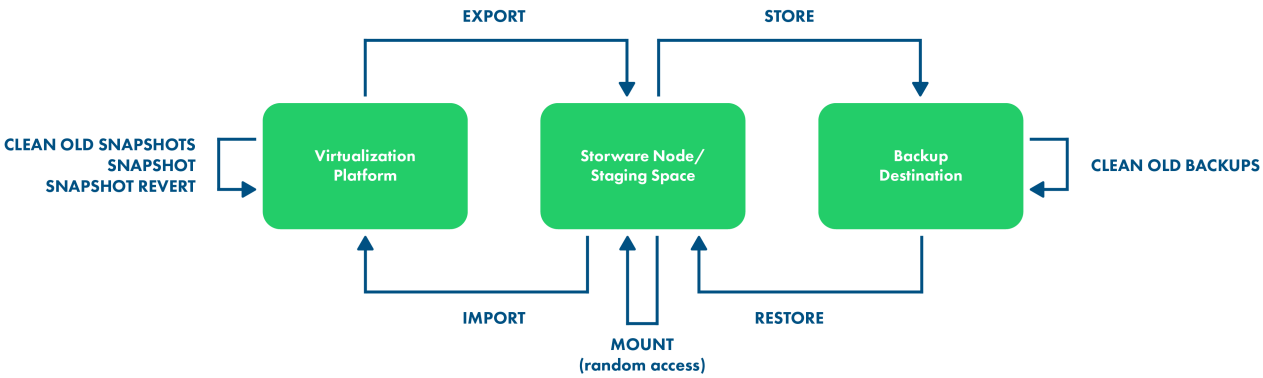
EXPORT
STORE
CLEAN OLD SNAPSHOTS
SNAPSHOT
SNAPSHOT REVERT
Virtualization Platform
Storware Node/ Staging Space
Backup Destination
CLEAN OLD BACKUPS
IMPORT
MOUNT
(random access)
RESTORE

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

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 restore	Not supported (always on)	Not supported (always on)	Not supported (always on)

** 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
Proxy VM required	No	Yes

	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
Proxy VM required	No	No

	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 required	No	Yes
Proxy VM required	No	No

	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
Power-on VM after restore	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
Local restore (raw data)	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
Pre/post command execution	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

- CentOS Linux Stream 8
- CentOS Linux Stream 9

Red Hat Enterprise Linux

- Red Hat Enterprise Linux 8.8
- Red Hat Enterprise Linux 8.9
- Red Hat Enterprise Linux 8.10
- Red Hat Enterprise Linux 9.0
- Red Hat Enterprise Linux 9.1
- Red Hat Enterprise Linux 9.2
- Red Hat Enterprise Linux 9.3
- Red Hat Enterprise Linux 9.4
- Red Hat Enterprise Linux 9.5

SUSE Linux Enterprise Server

- SUSE Linux Enterprise Server 15 SP3
- SUSE Linux Enterprise Server 15 SP4
- SUSE Linux Enterprise Server 15 SP5

- Using Red Hat Enterprise Linux requires an active subscription.

- Minimal installation is required.

Supported browsers

Administrative UI supports the following browsers. We recommend that you use the most up-to-date browser that's compatible with your operating system.

- Chrome (latest version)

MariaDB

Data Protector for Cloud Workloads 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)
		111/tcp, 111/UDP, 2049/tcp, 2049/UDP, ports	

Server	Node	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 Ceph

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
oVirt/RHV/OLVM host selected in export storage domain configuration	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 oVirt/RHV/OLVM host selected in export storage domain configuration	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

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`	if staging space (export storage repository) is hosted on the Node - NFS access

		(UDP), otherwise check the documentation of your NFS storage provider	
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 data transfer over NBD

Proxmox VE

Export storage repository

Source	Destination	Ports	Description
Node	Hypervisor	22/tcp	SSH access
		If Node is hosting staging space: 111/tcp, 111/UDP, 2049/tcp, 2049/UDP, ports specified in	

Hypervisor	Node	`/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/* /root/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/* /root/DP-for-Cloud-Workloads-repo
   ```

8. If your network is using proxy server update 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
vprotect
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 CERT_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/* /root/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_server/meta/collections.yml
ansible-galaxy install -r
~/.ansible/roles/xe0nic.ansible_vprotect_node/meta/collections.yml
```

- 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:///root/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 }}/{{
mariadb_distro }}"
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 /root/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
virtual environments repository
[DP-for-Cloud-Workloads]
name = Data Protector for Cloud Workloads
baseurl = file:///root/DP-for-Cloud-Workloads-repo
gpgcheck = 0
```

For Red Hat Enterprise Linux 9 and compatible

```
# Data Protector for Cloud Workloads - Enterprise backup solution for
virtual environments repository
[DP-for-Cloud-Workloads]
name = Data Protector for Cloud Workloads
baseurl = file:///root/DP-for-Cloud-Workloads-repo
gpgcheck = 0
```

For SUSE Linux Enterprise Server 14 and compatible

```
# Data Protector for Cloud Workloads - Enterprise backup solution for
virtual environments repository
[DP-for-Cloud-Workloads]
name = Data Protector for Cloud Workloads
baseurl = file:///root/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> --
   password=<user password> --apiurl='http://<server address>:
   <port>/api'
   ```

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
--apiurl='http://localhost:8080/api'
```

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 Nodes hosts - downgrade nodes with
```

```
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
CERT_ALIAS
```

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
- 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:

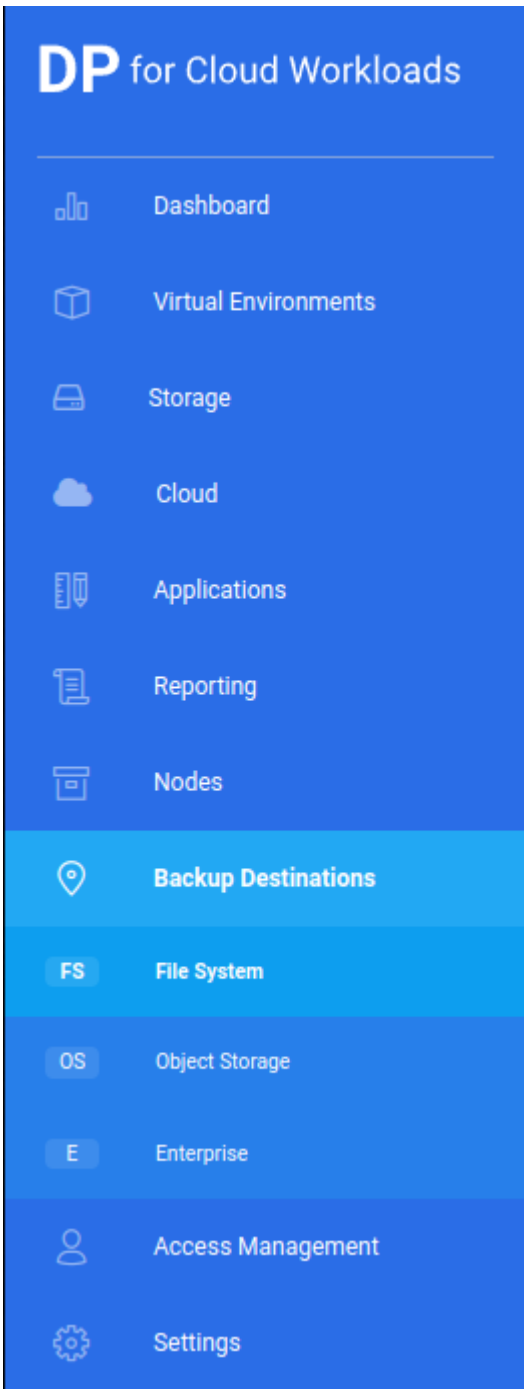

- Open your Backup Destination (click on the name)
- Provide pre/post access command arguments (the first argument is the command executed locally on the **node**):

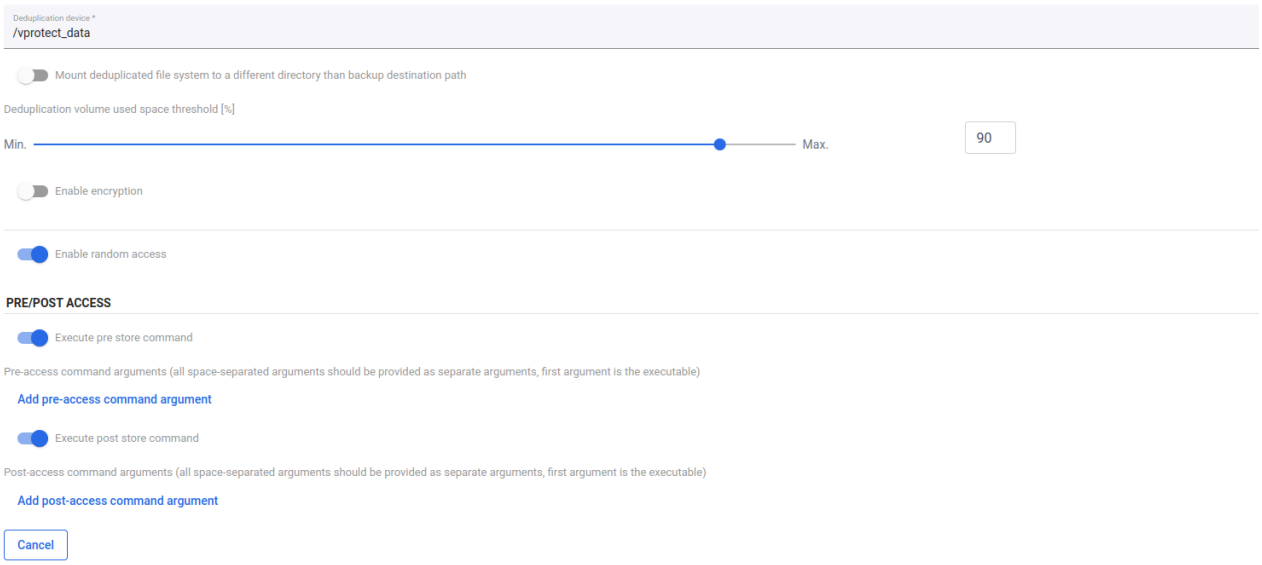

Encryption

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
		Data Protector	Generated based on metadata in the		

MS Azure Blob Storage	Yes	for Cloud Workloads Node	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
Micro Focus Data	Provider dependent	n/a	n/a	n/a	n/a

Protector					

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
sectors
Disk /dev/mapper/centos-swap: 3221 MB, 3221225472 bytes, 6291456
sectors
```

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
   sectors
   Disk /dev/mapper/centos-swap: 3221 MB, 3221225472 bytes, 6291456
   sectors
   ```

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.

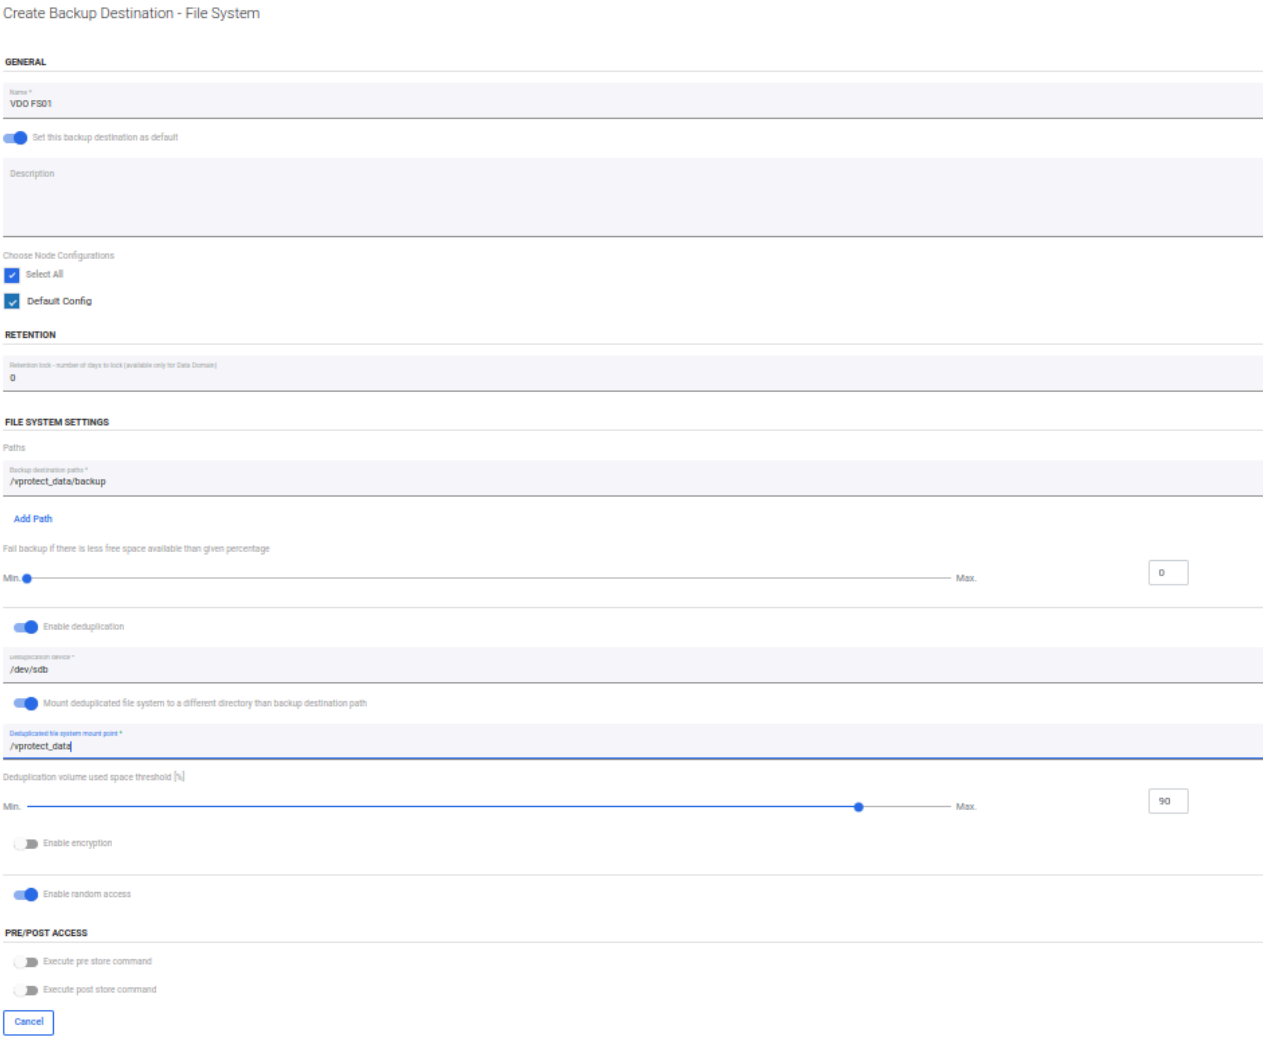

> **Note**:
>
> Only one file system backup destination with deduplication using VDO pointing to a specific directory can be used. If you want to add another backup destination using the same VDO device, but just a different subdirectory, create it without deduplication enabled.

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 bytes
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
volume WITHOUT formating.
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
vdovg/vdolvol VDO volume.
```

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/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 =
backupdestination.guid
set filesystembackupdestination.dedupvolume = CONCAT('/dev/',
REGEXP_REPLACE(backupdestination.name,'\\W','_'), '/',
REGEXP_REPLACE(backupdestination.name,'\\W','_'))
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
MOUNTPOINTS
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,

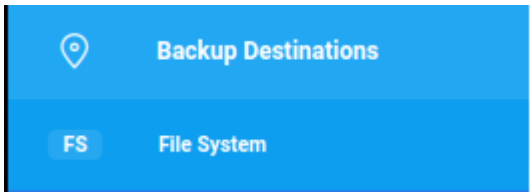

2. Select Create Backup Destination → File System (Synthetic),

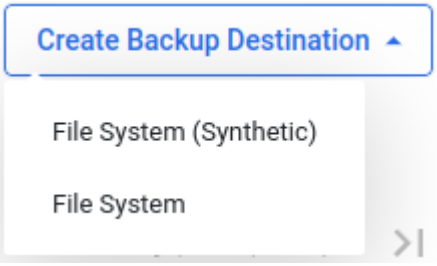

3. The rest of the configuration is similar to a regular File system.

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

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.

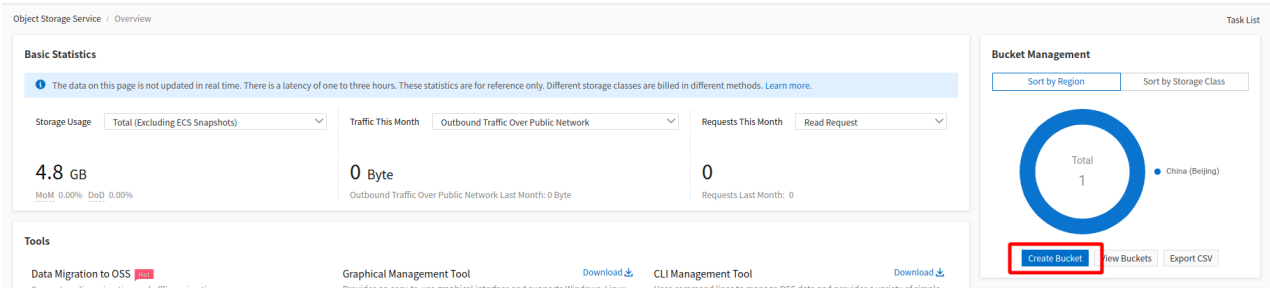

Provide necessary details for your bucket and enable versioning.

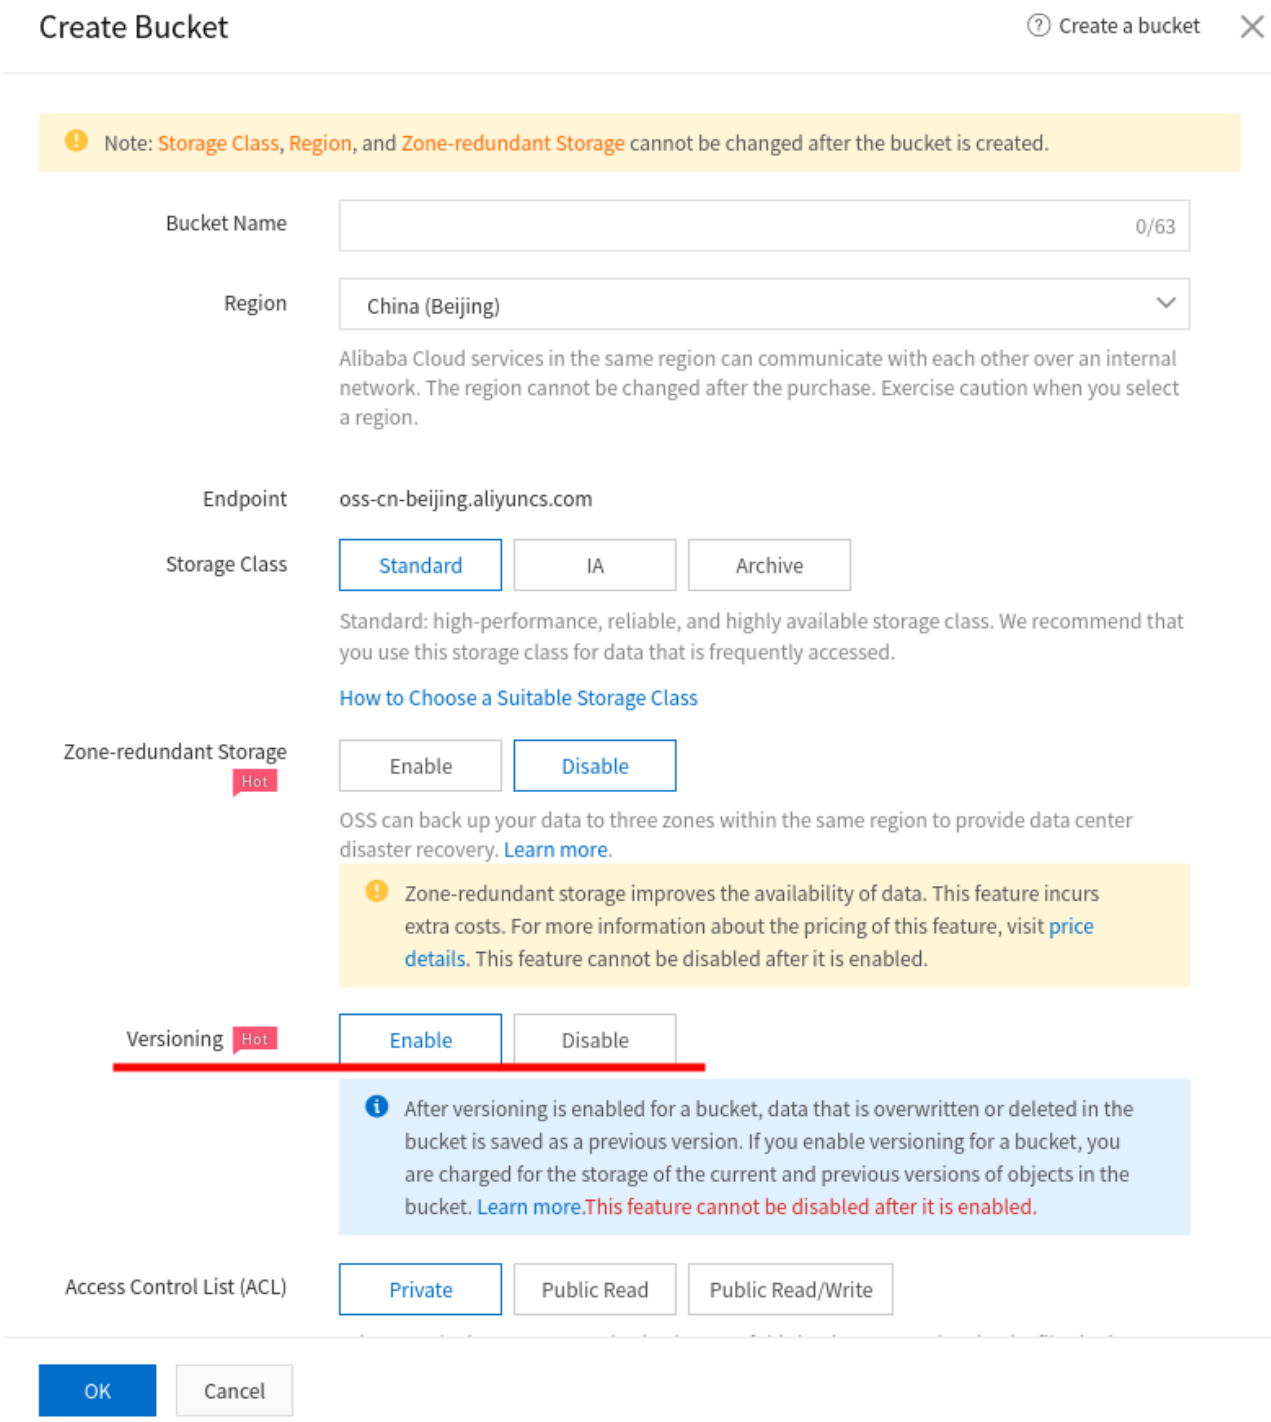

Next, go to Manage AccessKey Management and create new AccessKey

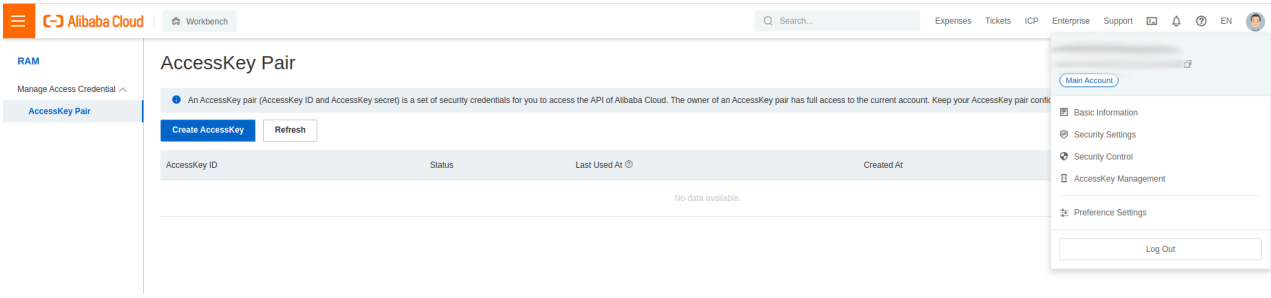

Now go to the Backup destination tab on the Data Protector for Cloud Workloads dashboard and change the sub-tab to object storage. Provide the bucket name and key credentials, and then configure the remaining options according to your requirements:

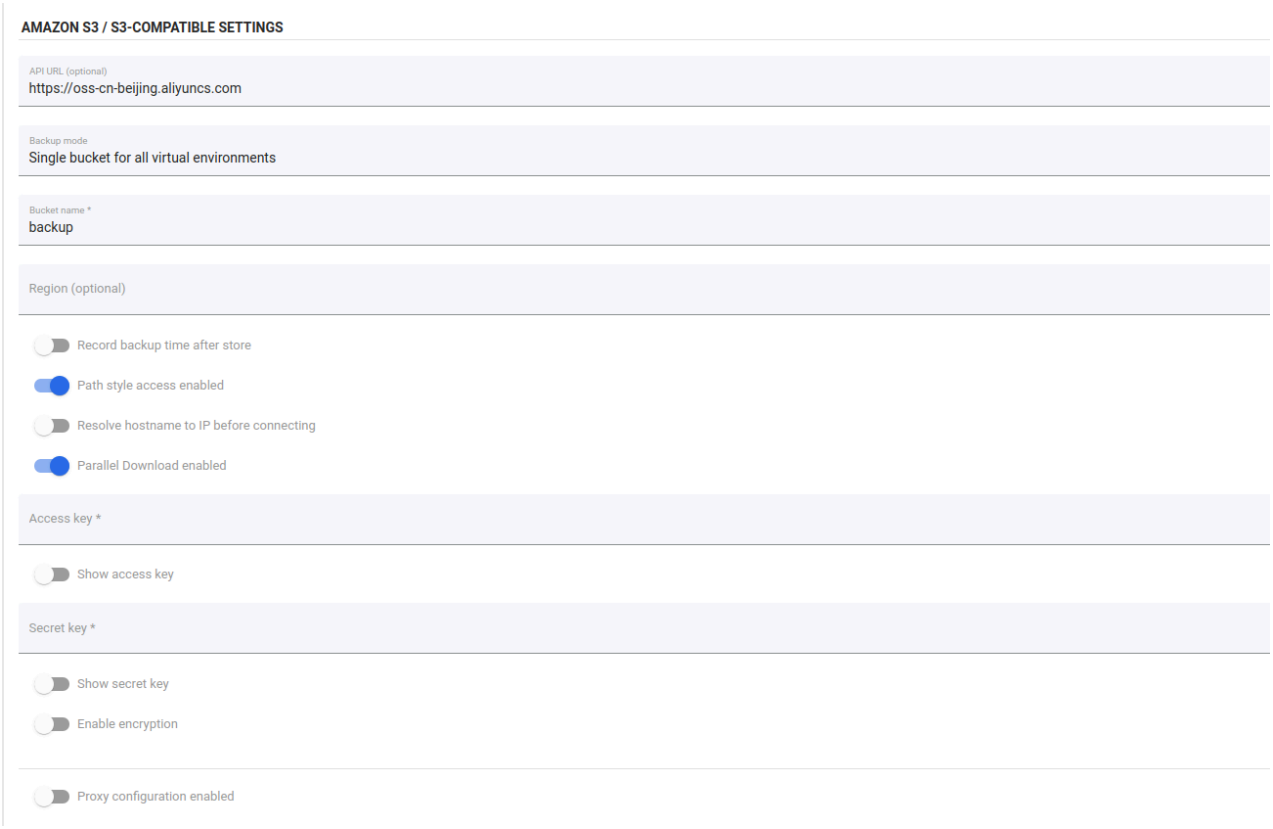

AWS S3 or S3-compatible

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:

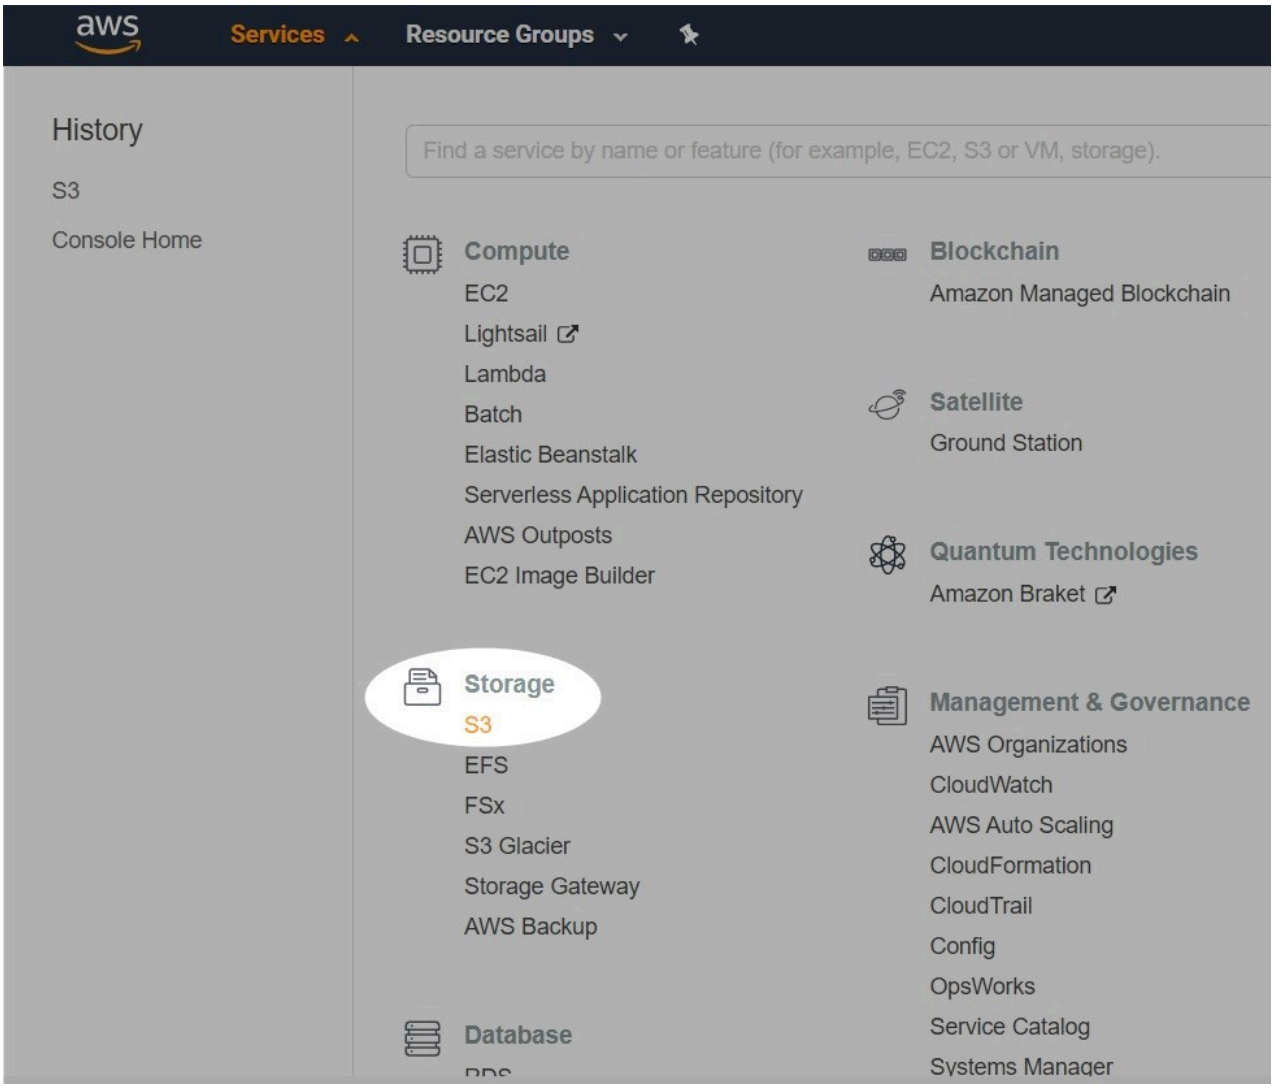

Now create a new bucket for your backups:

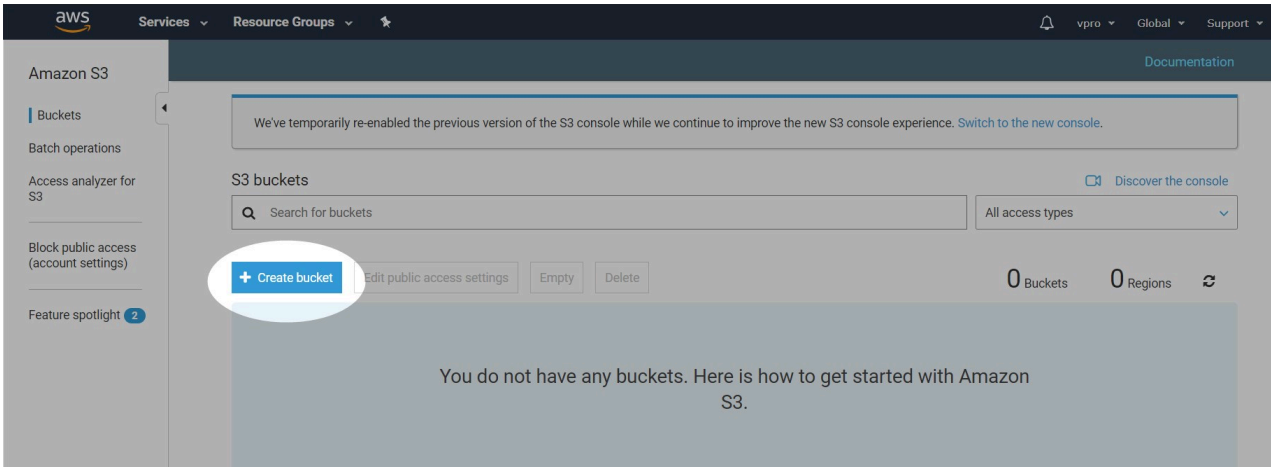

In "Configure options" activate versioning:
(In all other tabs, you can leave the default settings)

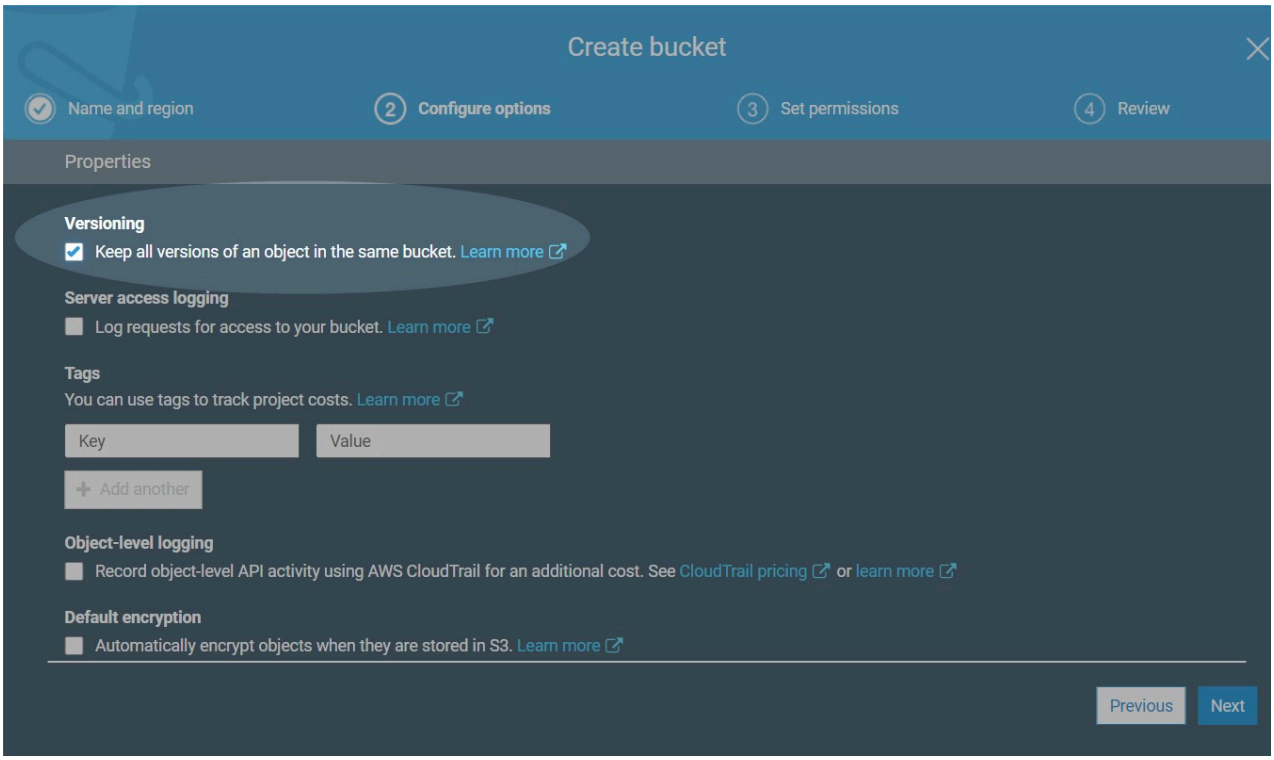

After creating a bucket, we need to create a new user with appropriate permissions:

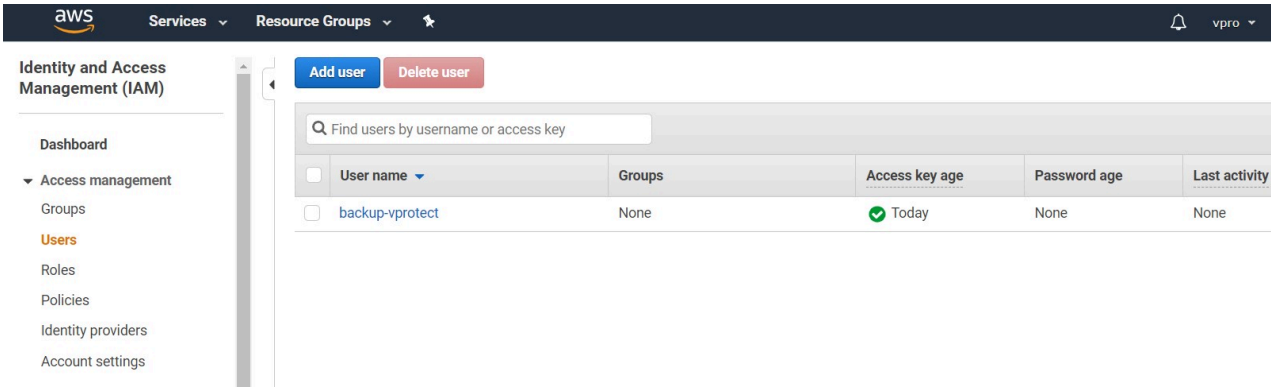

Remember to choose the "Programmatic access" account type:

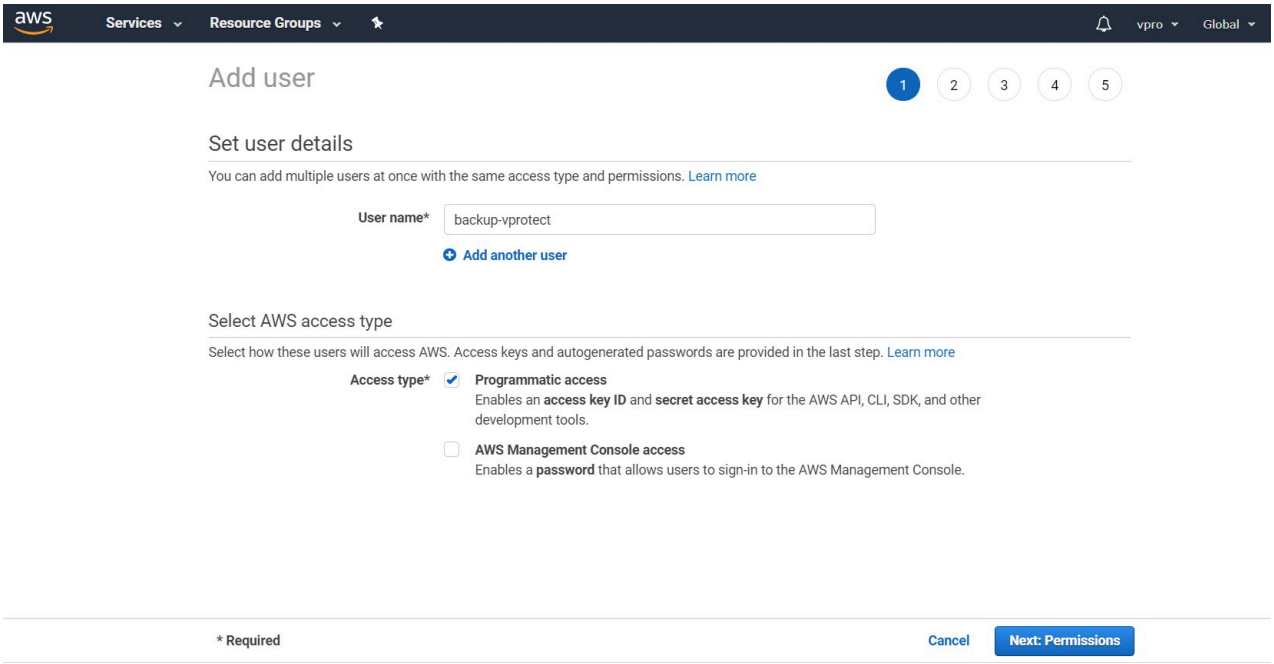

From the predefined roles, you can choose "AmazonS3FullAccess" or you can create a new one as described in the Permissions section:

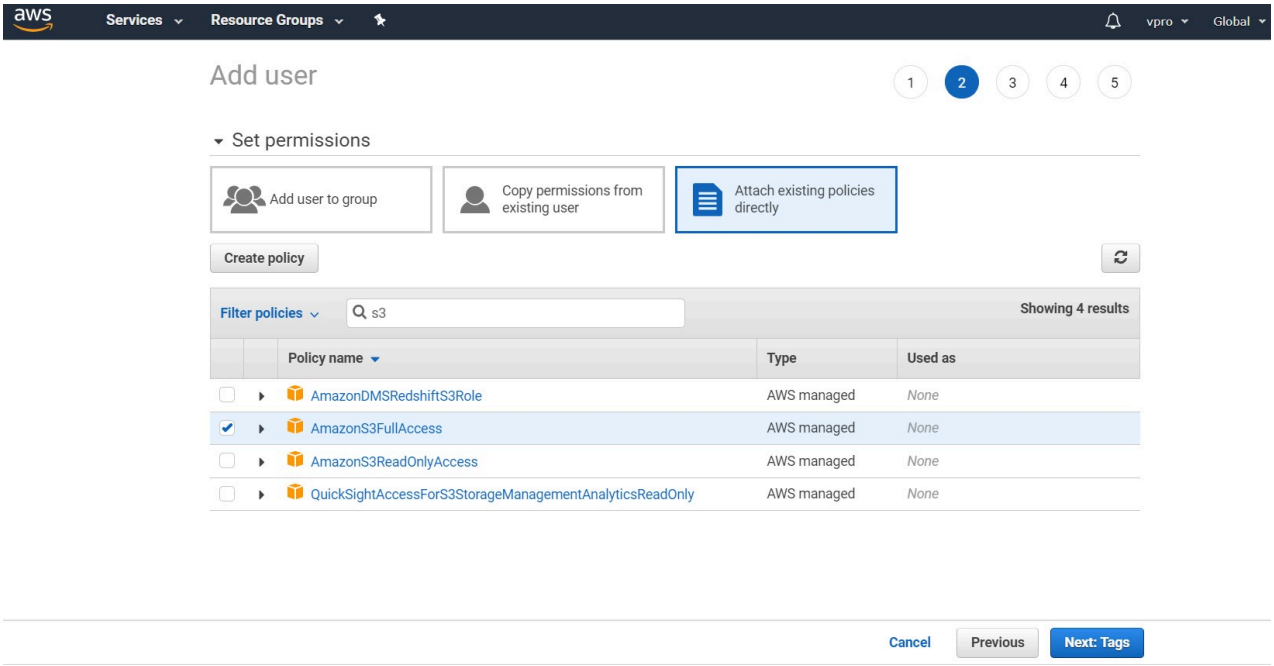

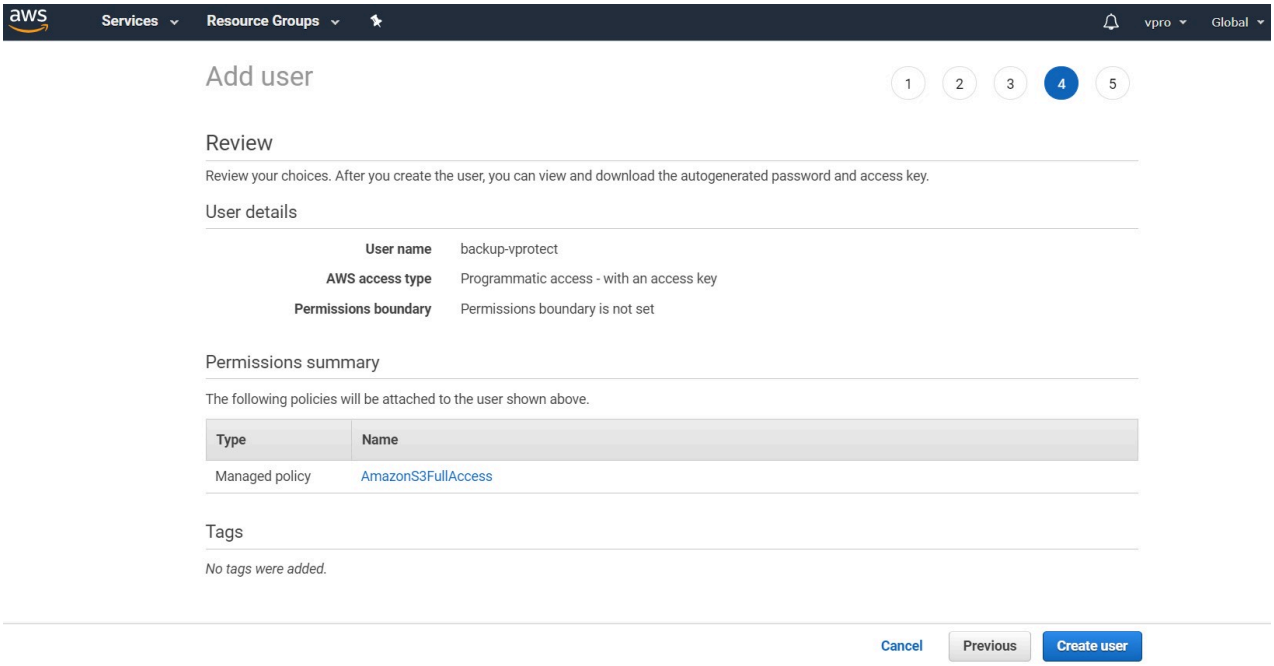

> Remember to download the .csv or copy the key credentials manually:

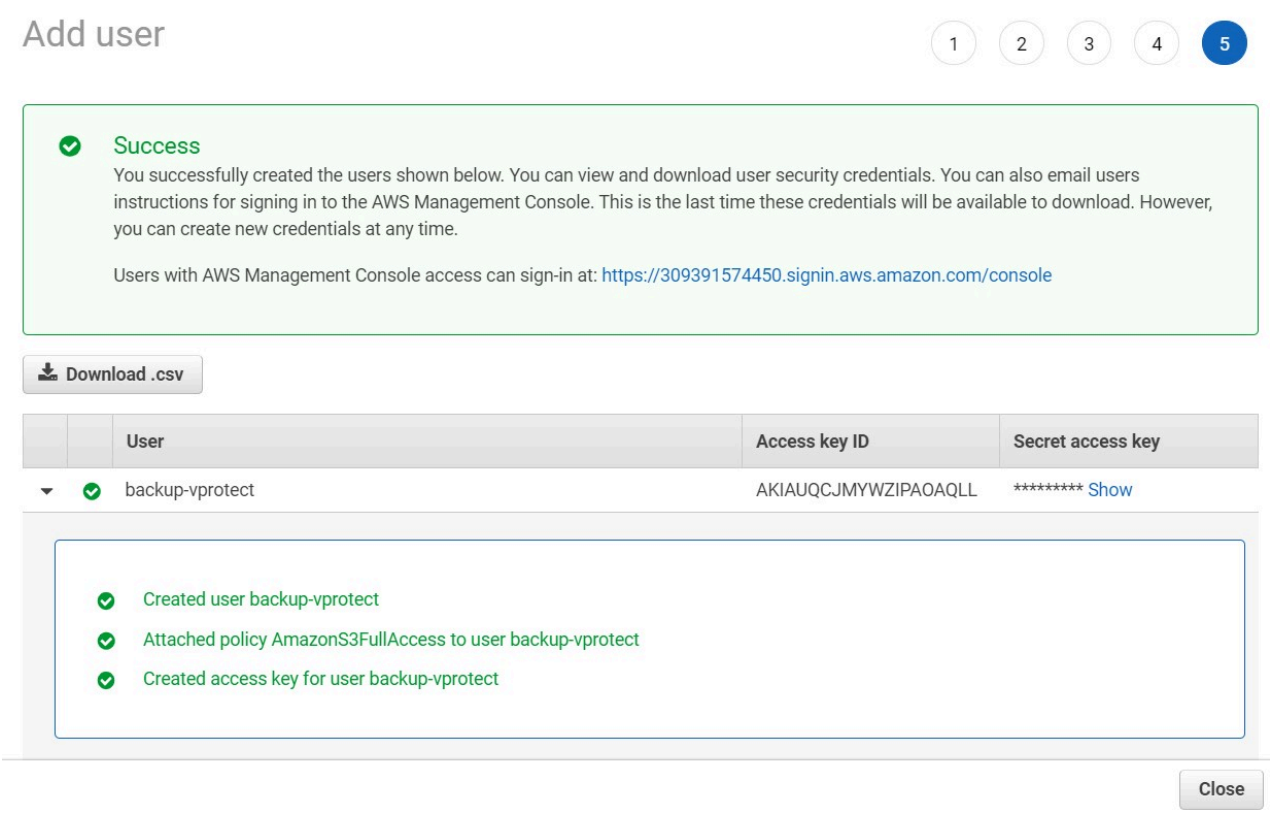

Now go to the Backup destination tab on the Data Protector for Cloud Workloads dashboard and change the sub-tab to object storage. Provide the bucket name and key credentials, and then configure the remaining options according to your requirements:

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".

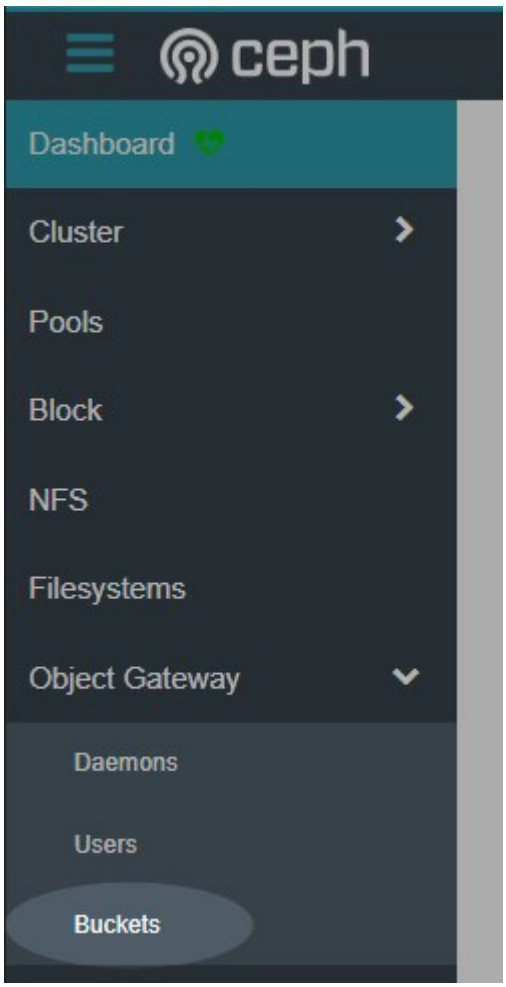

Then click on the "Create" button.

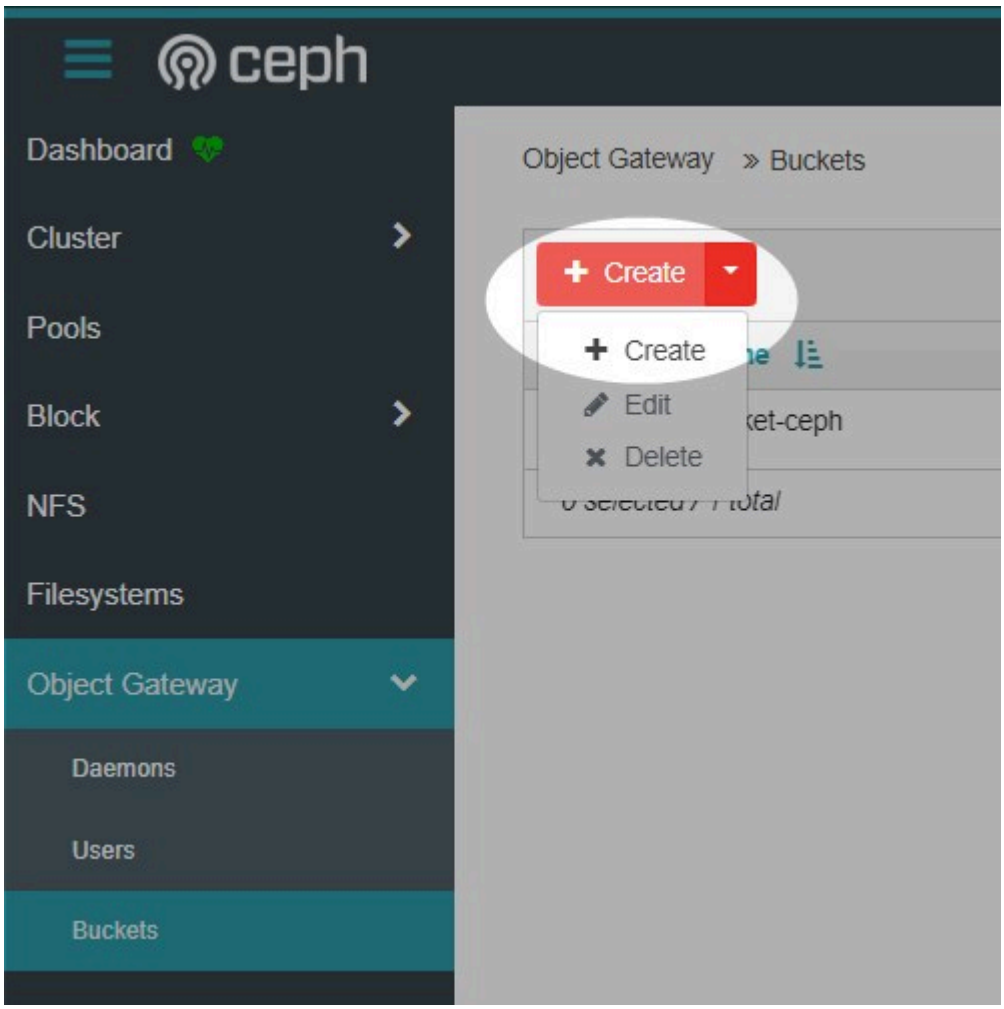

Fill in the required fields.

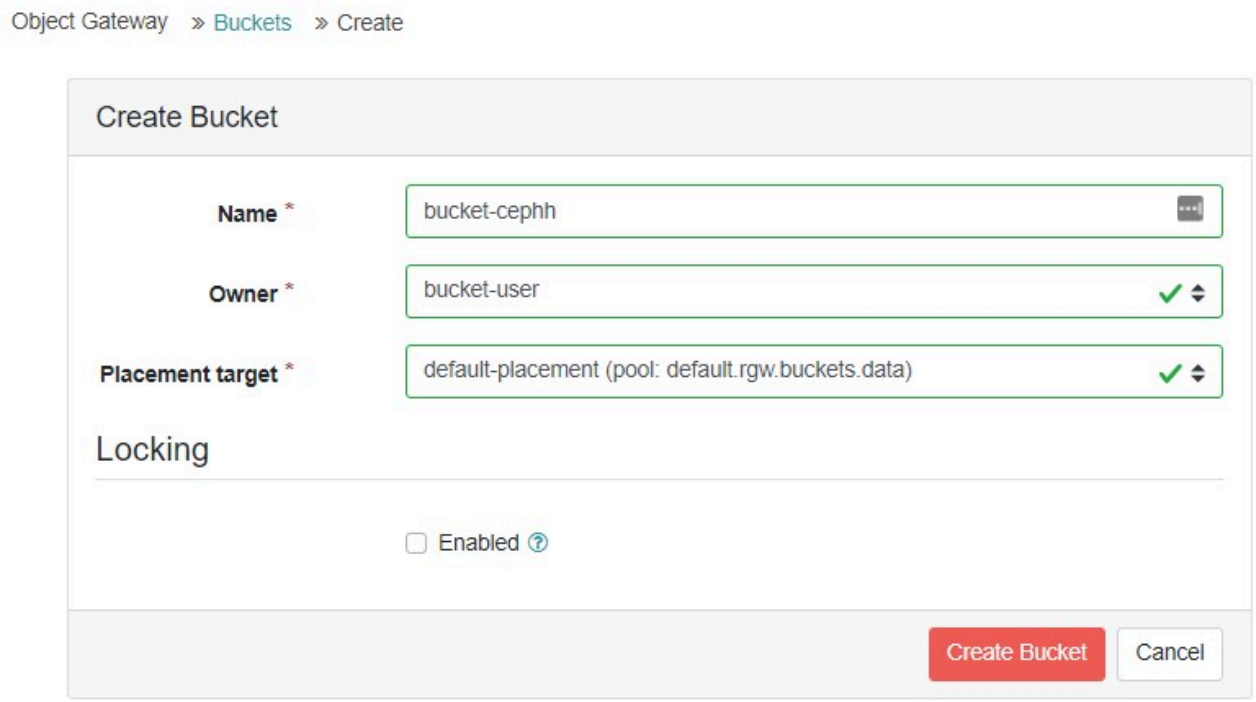

Now create a dedicated access account for the backup destination. Open the Users tab under the object gateway menu.

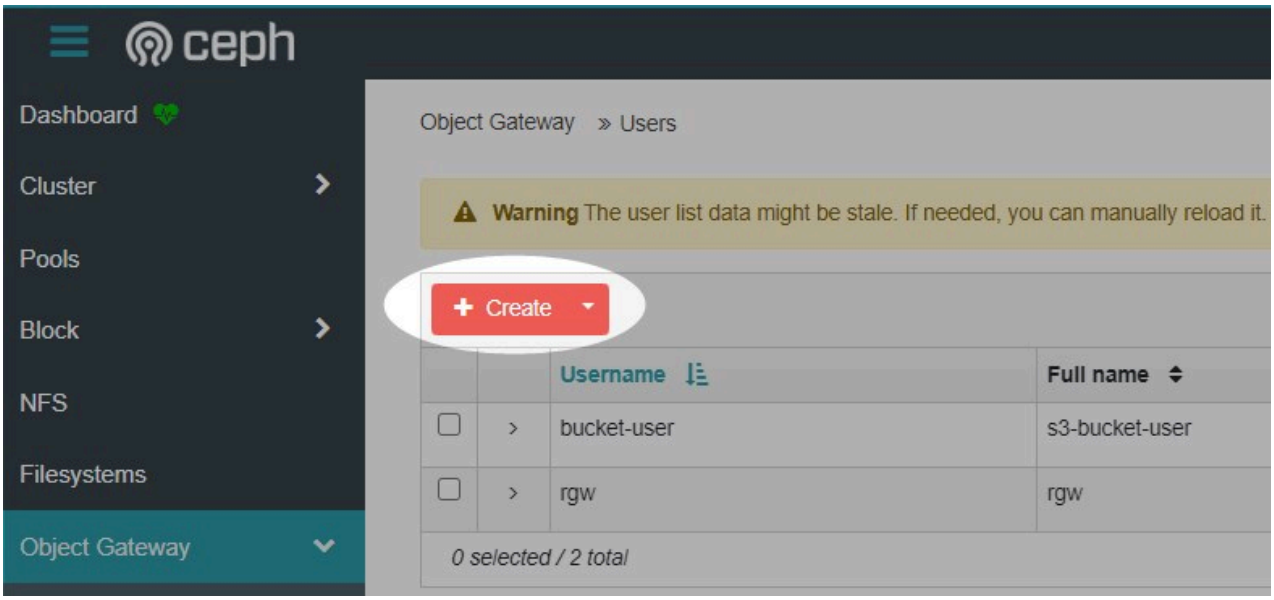

Fill in the username field, you can leave the other settings as default.

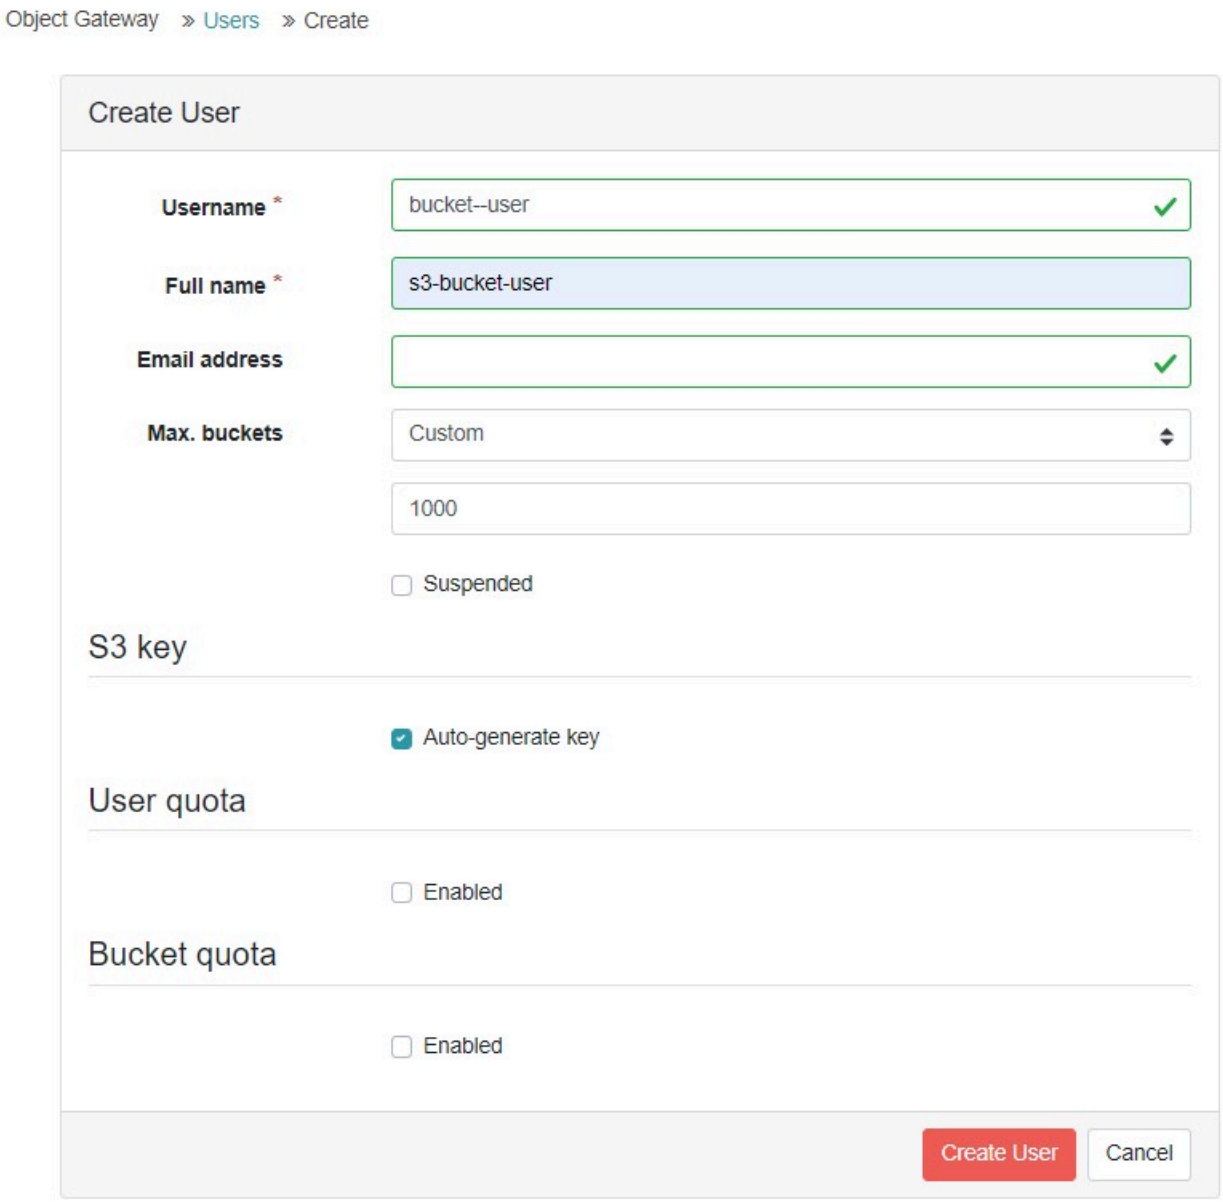

To see the account key and secret key, expand the user details and open the keys tab, click on the key, and then on the show button.

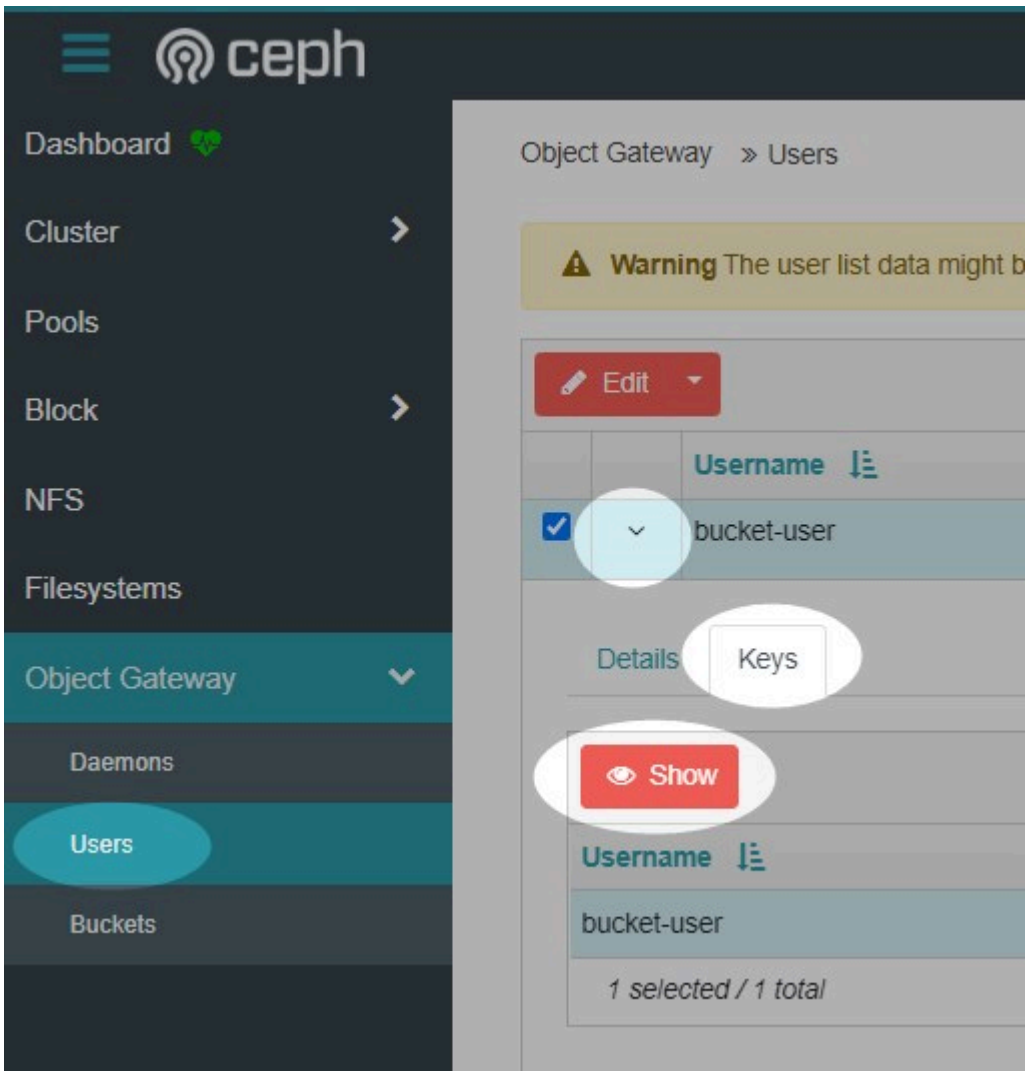

The access key and secret key will be needed to create a backup destination in Data Protector for Cloud Workloads.

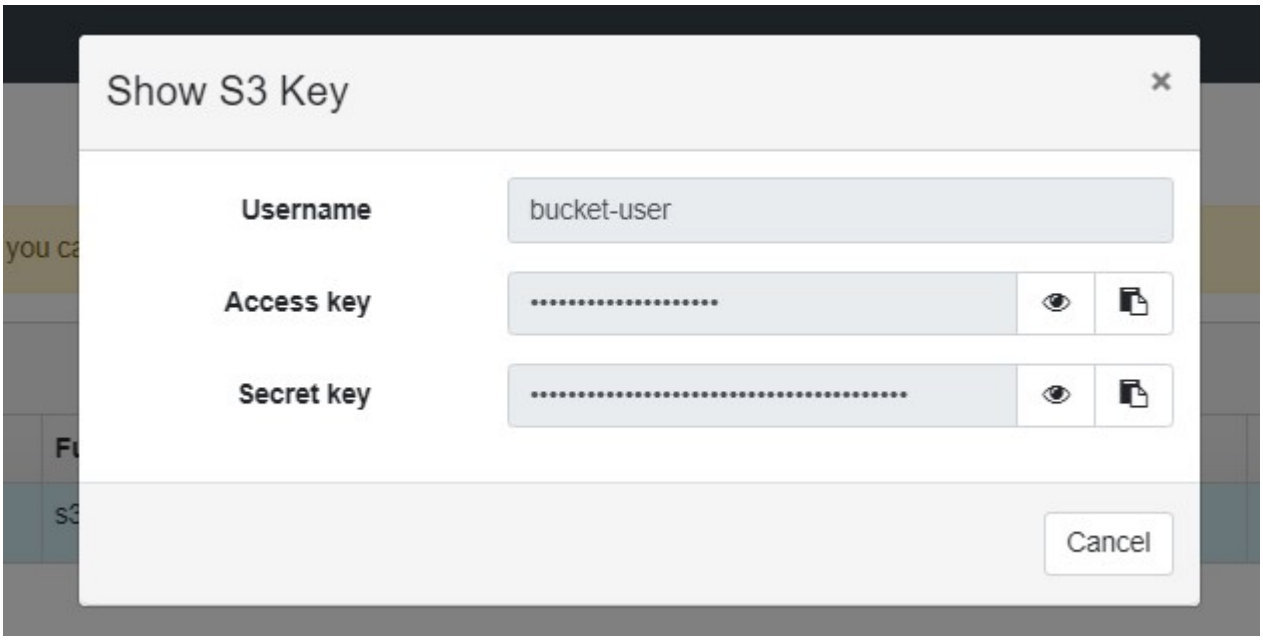

Now we can go to the Data Protector for Cloud Workloads Dashboard. Open the "Backup Destination" tab from the left side menu and choose "Amazon S3 / S3-compatible" as the new type of backup destination.

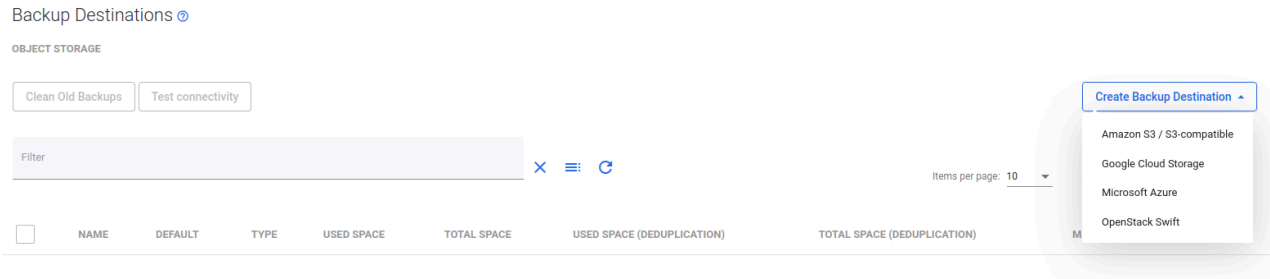

By default, Ceph provides S3 via port 8000. Also, remember to enable the "record backup time after store" option.

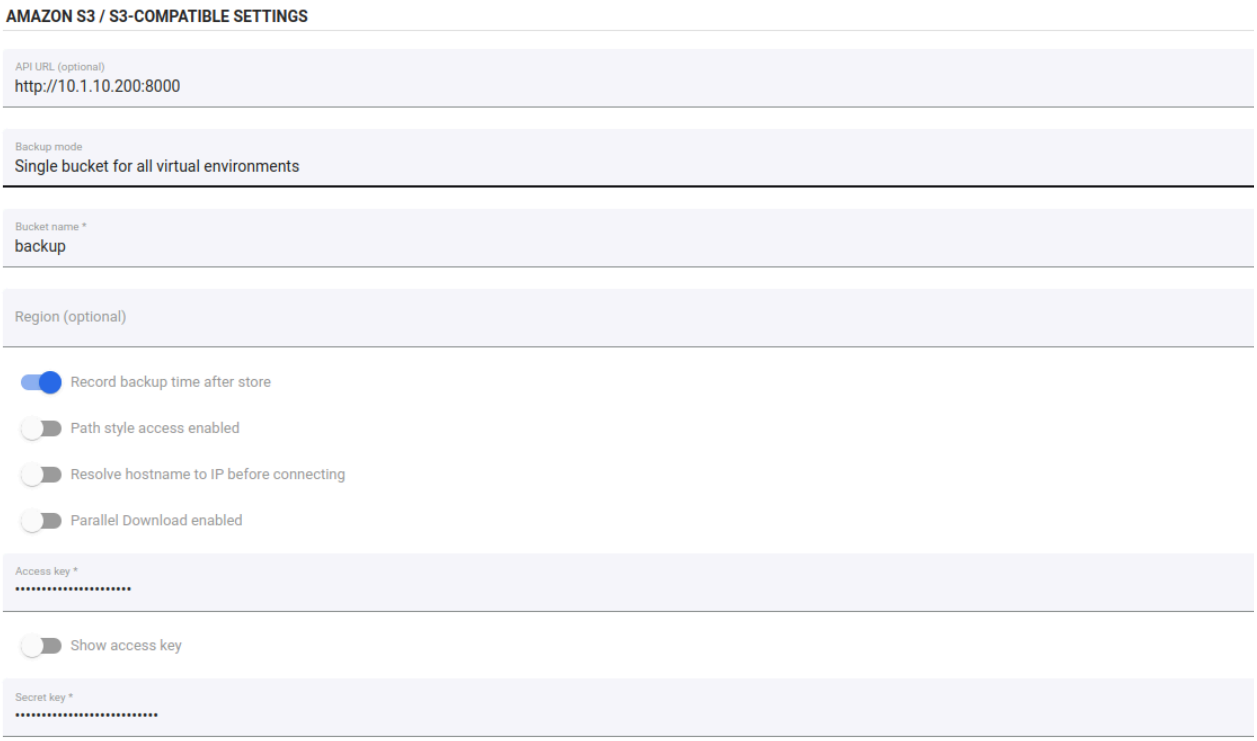

Cloudian S3

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

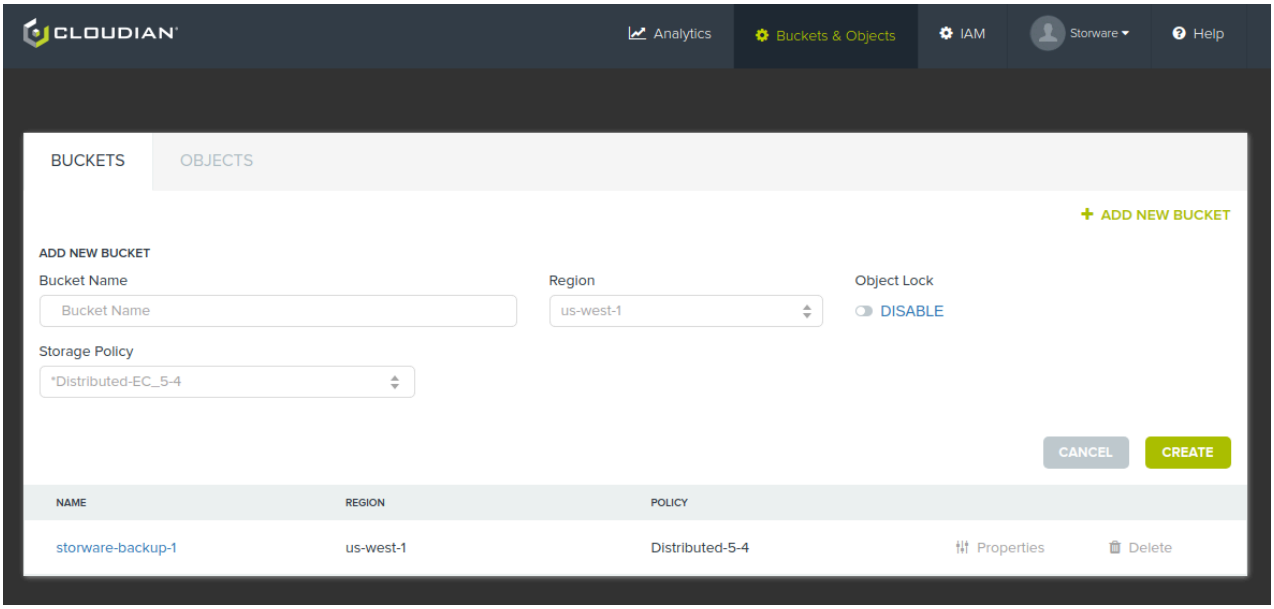

Next, go to security credentials and generate a new access key.

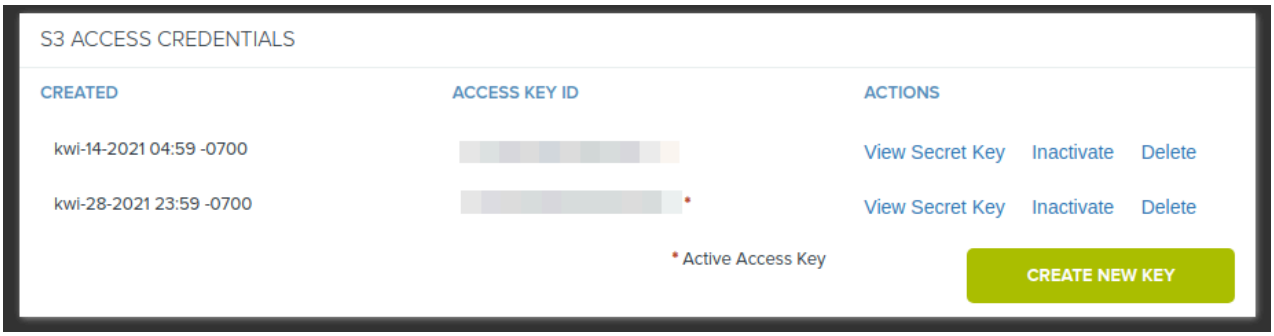

Now go to the Backup destination tab on the Data Protector for Cloud Workloads dashboard and change the sub-tab to object storage. Provide the bucket name and key credentials, and then configure the remaining options according to your requirements. Also, enable `Path style access enabled` option:

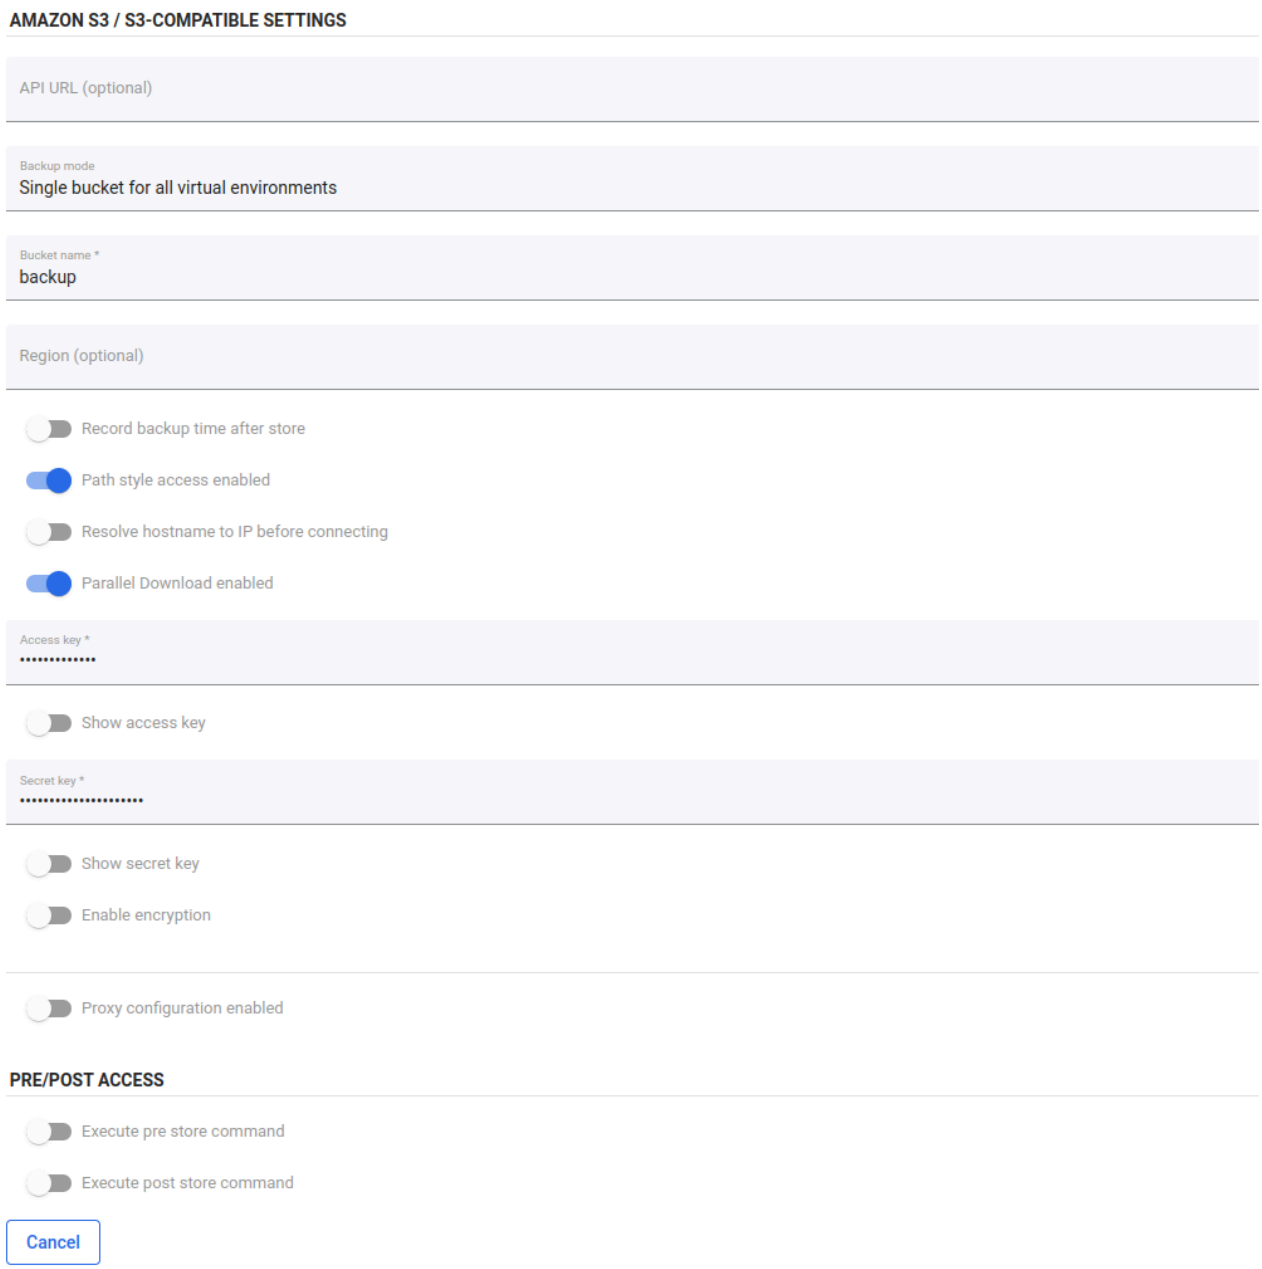

Google Cloud Storage

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

CONTINUE

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)

CONTINUE

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

CONTINUE

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

CONTINUE

Advanced settings (optional)

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**.

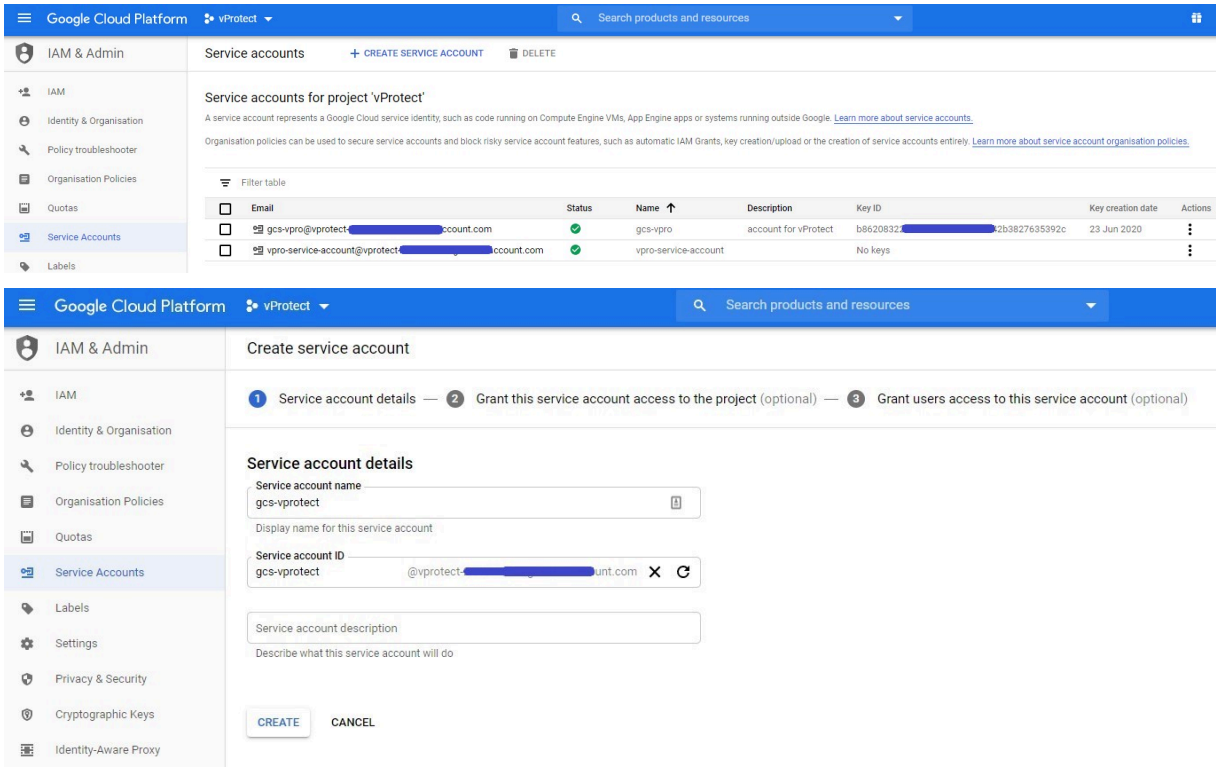

- You can leave the third tab - Grant users access to this service account (optional).\
- To generate an account key, click on the "three-dot" button next to your service account and then click on "create key". You should then see the

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

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.

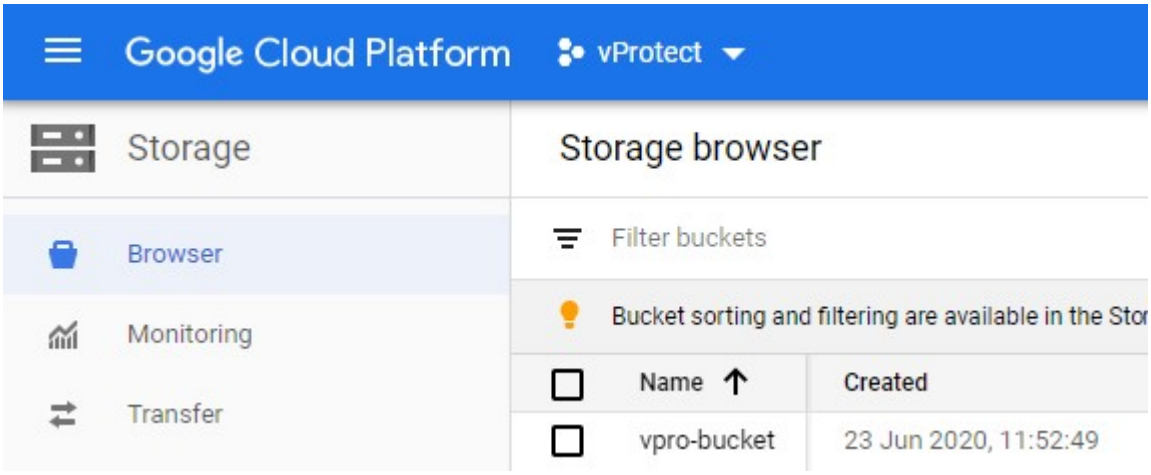

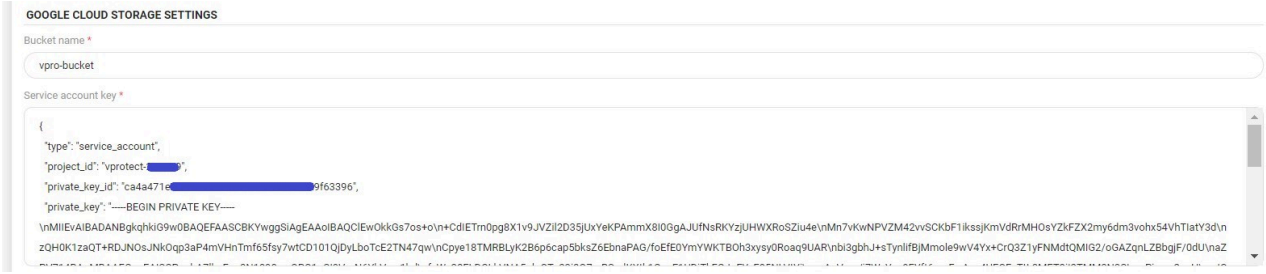

Now, you can store Data Protector for Cloud Workloads backups on Google Cloud Storage.

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.

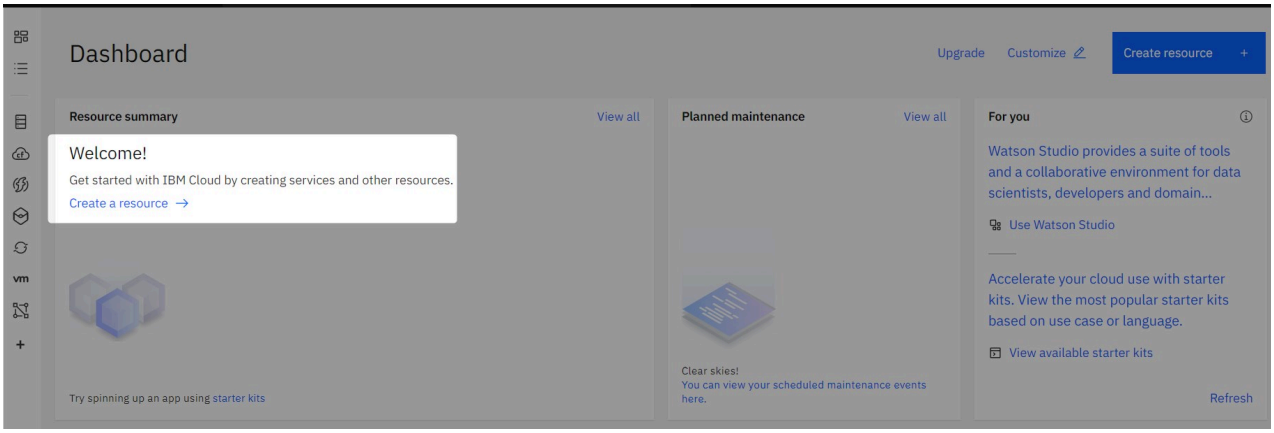

On the next screen, search for a resource named "Object Storage" and click on it.

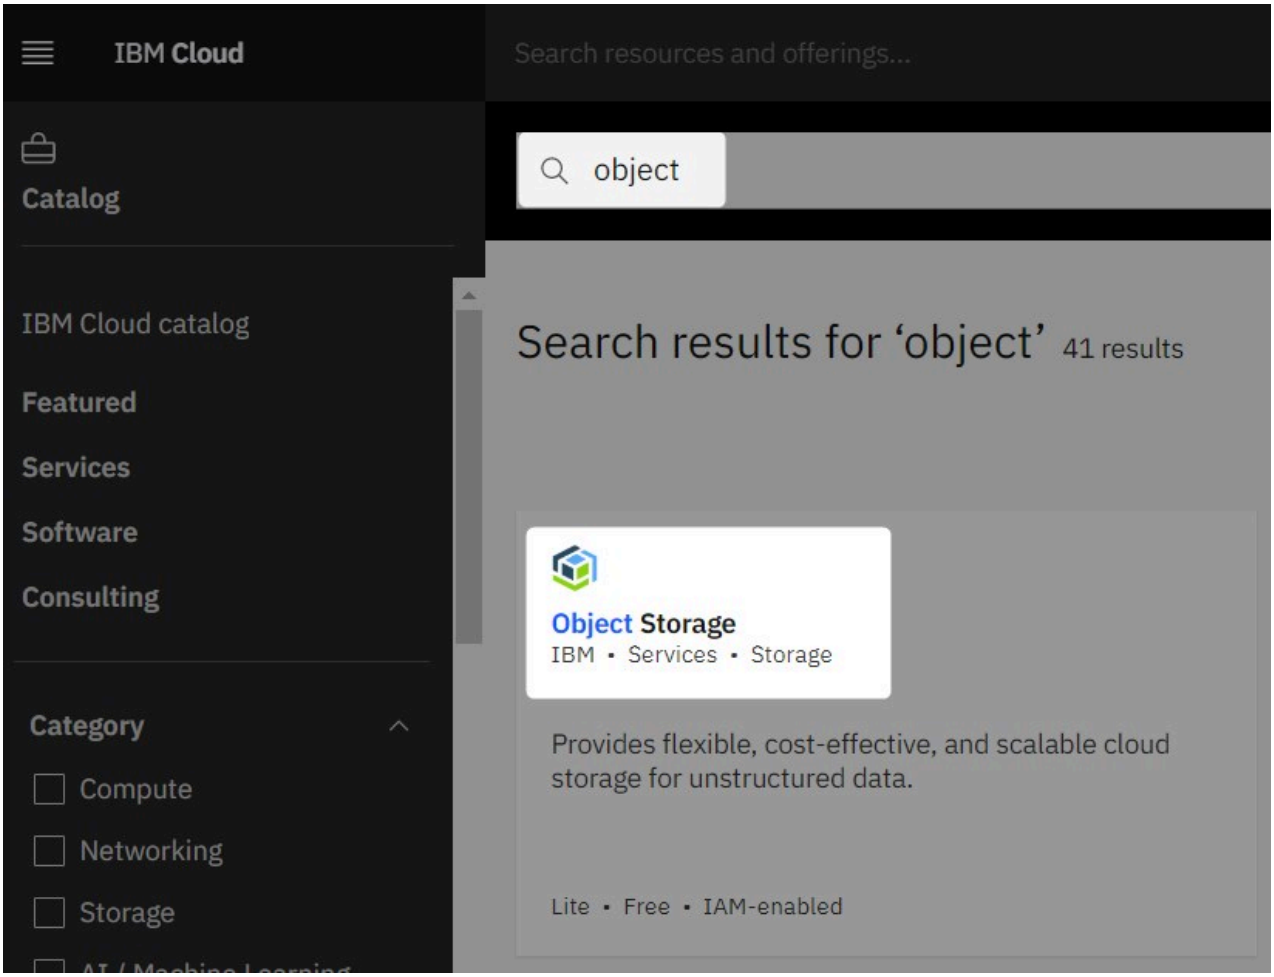

On the next screen, you can choose piercing plans, etc. Select the options according to your requirements.

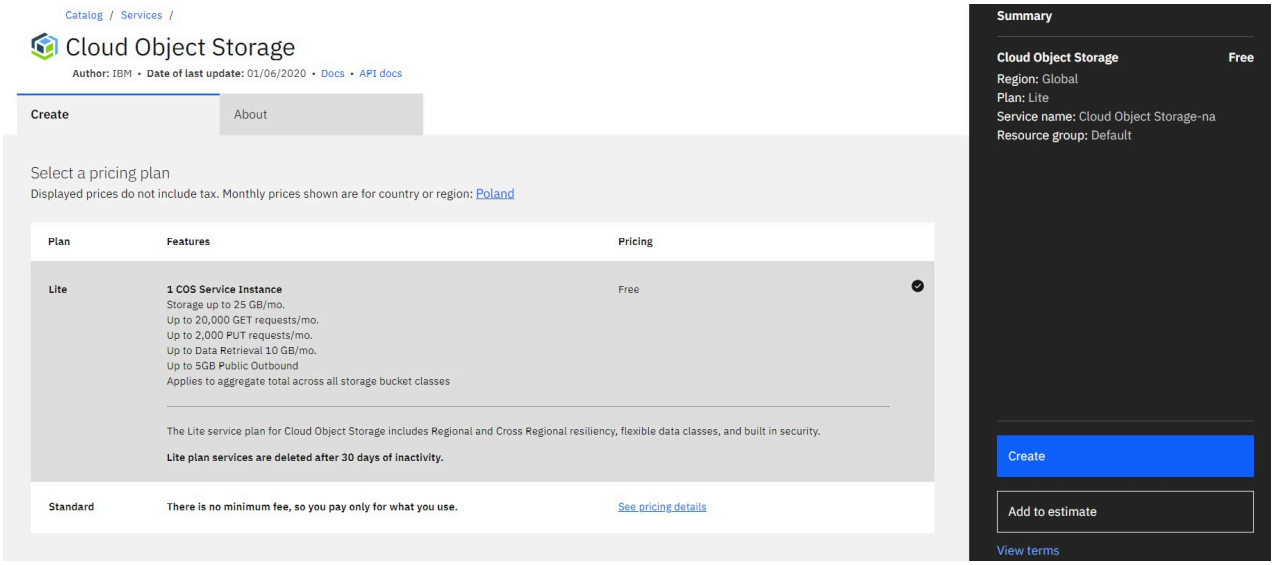

After creating a storage resource we need to create a bucket.

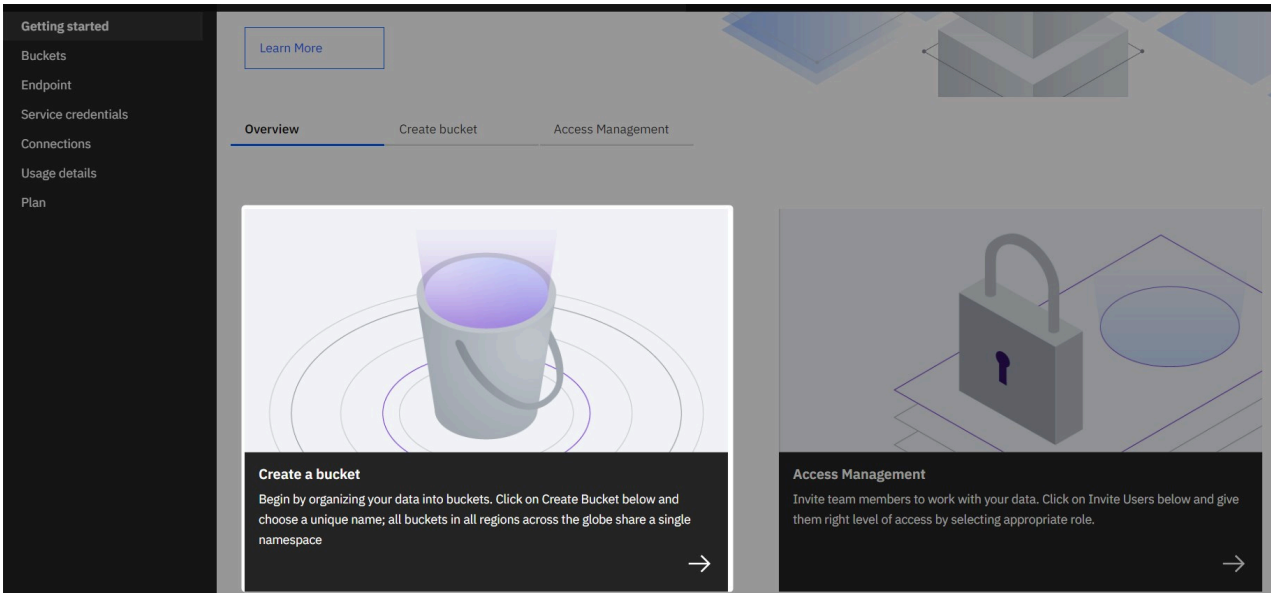

You can choose predefined templates or select the option to create a bucket with your own settings. In this example, we will choose "Custom bucket".

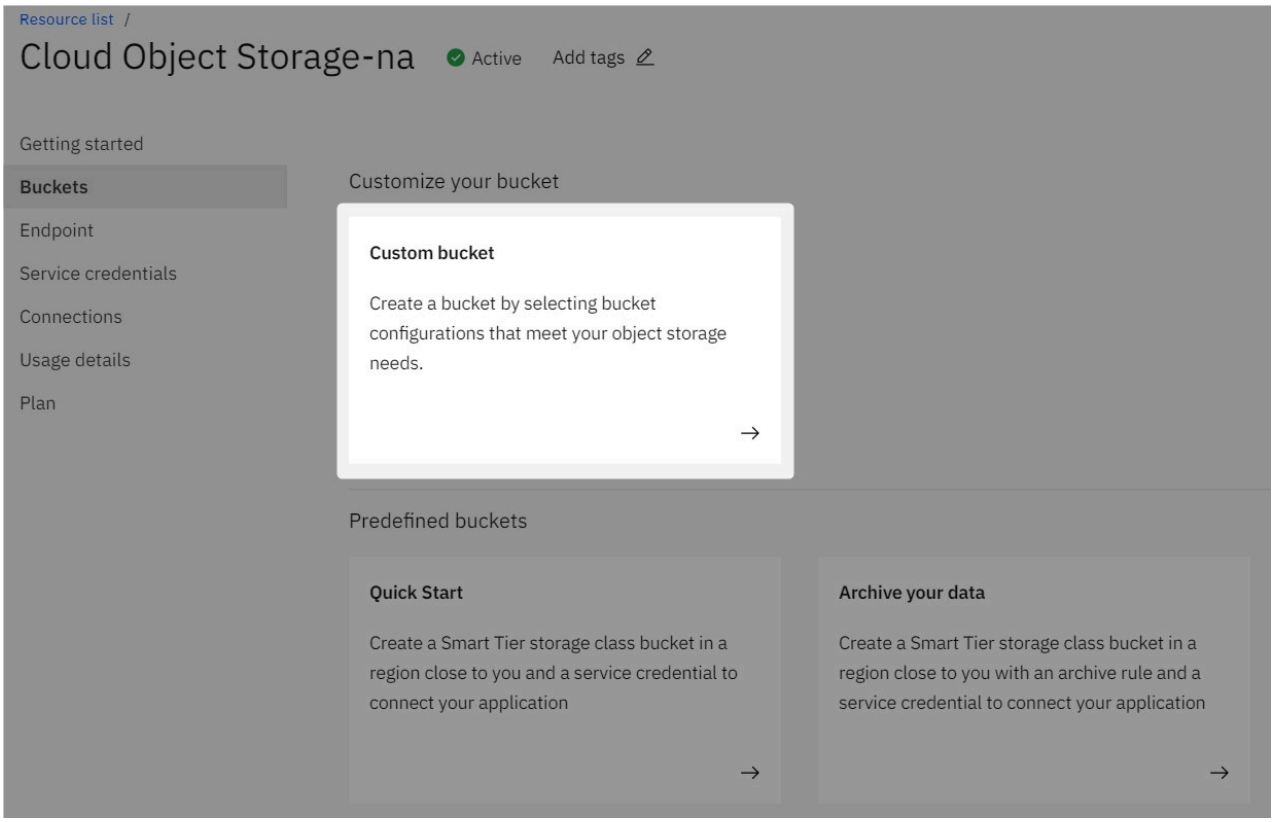

Data Protector for Cloud Workloads has no special requirements for the bucket, all options can be configured according to customer needs.

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

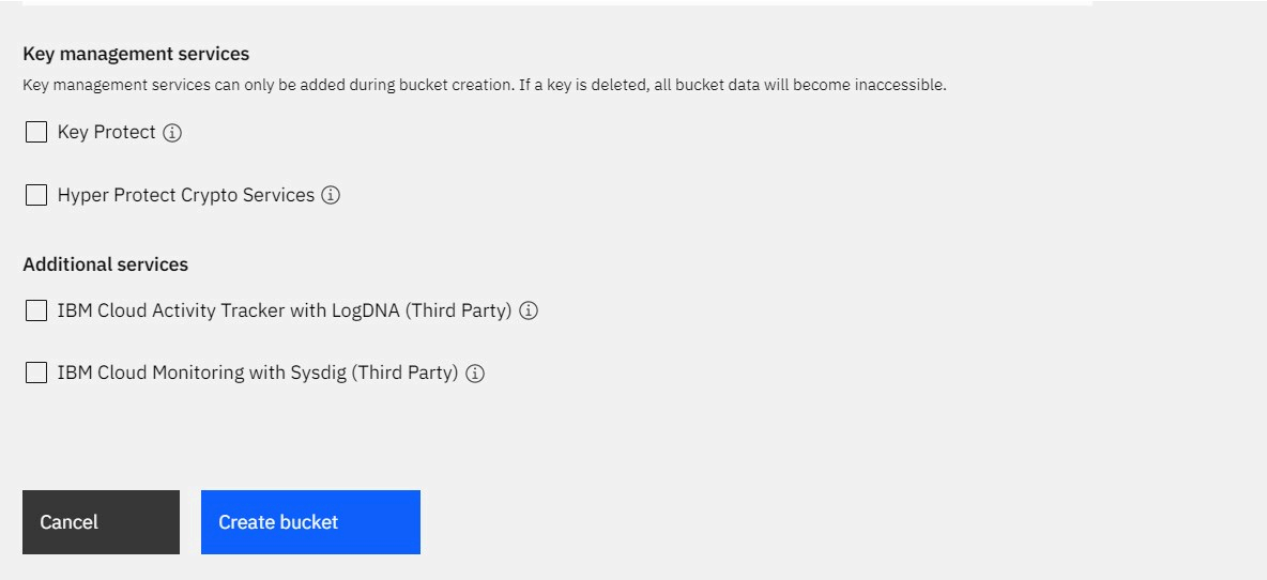

After creating the bucket, you'll see the objects page. From the menu on the left select the configuration tab. You will see a summary of the resource you have created. To create a backup destination you will need the ***"public" address from the endpoints section*** from here.

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-bucket
```

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

Add Activity Tracker

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

Monitoring

Add Monitoring

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

Archive rule

Create

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

Add

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

Add rule +

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.

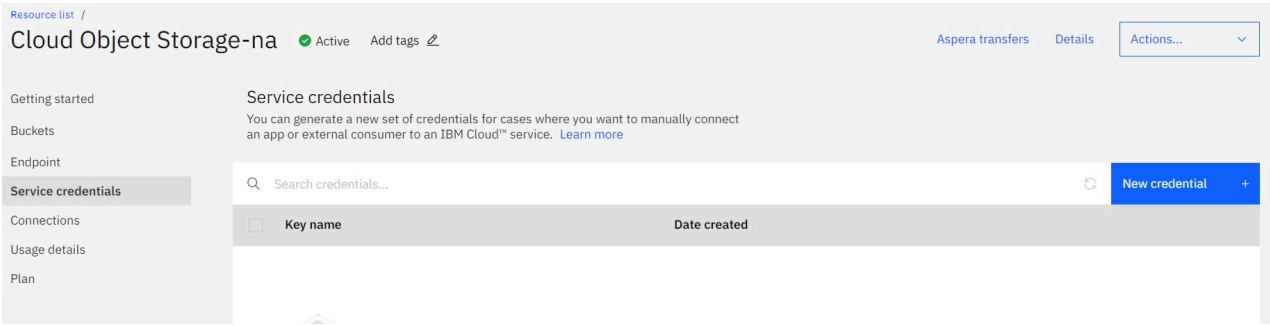

There are two important options on this screen. You must select the appropriate role (for Data Protector for Cloud Workloads it is the "Writer" role) and select the option "Include HMAC credential".

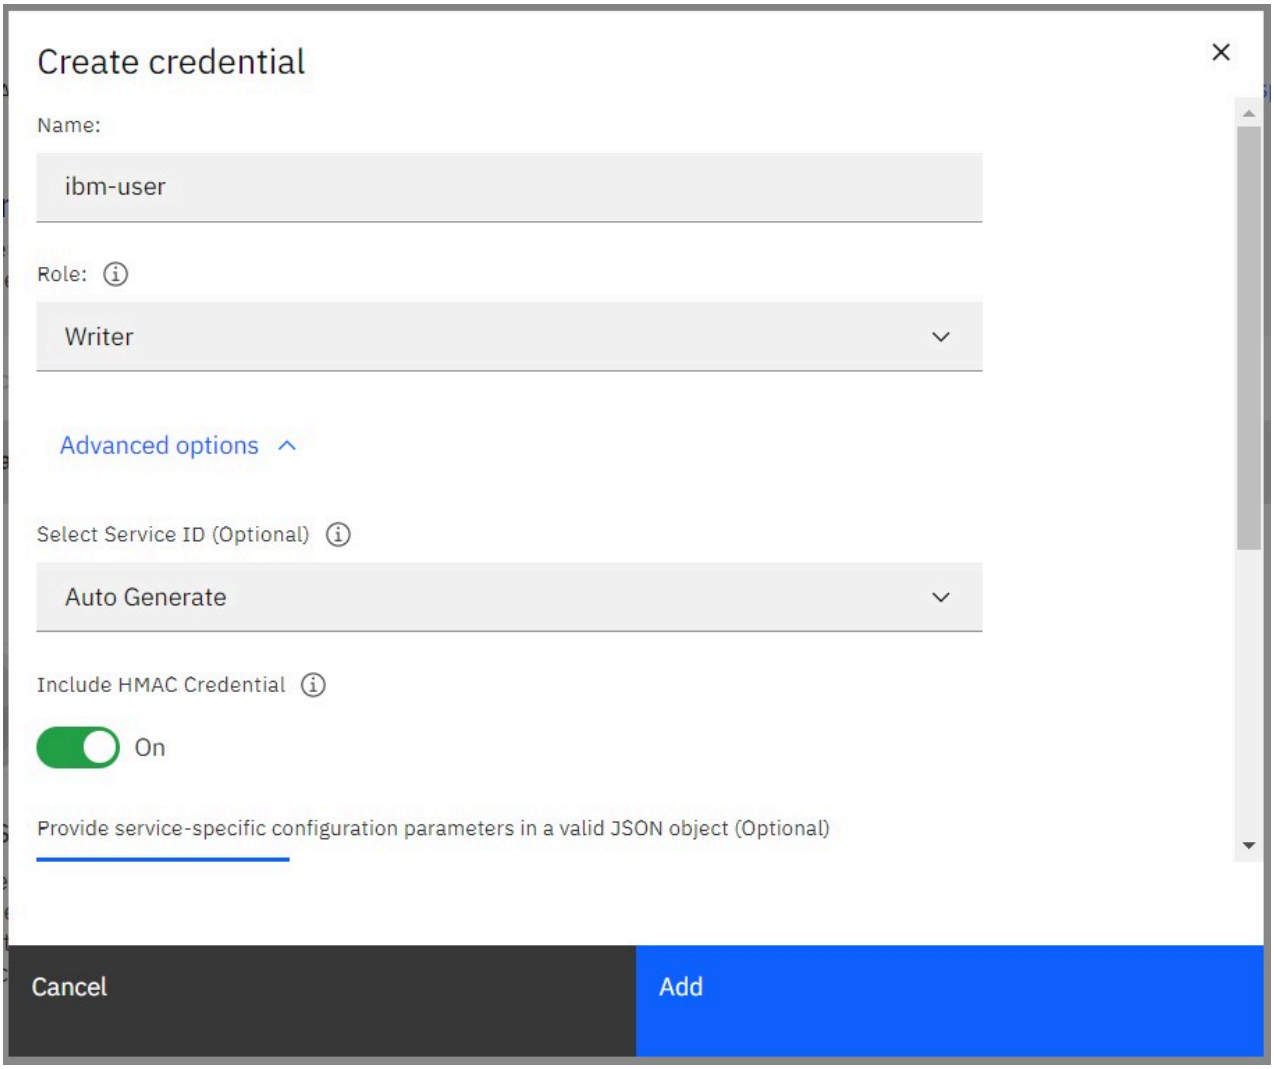

Now expand the detailed information about the created credentials by clicking on the arrow next to the name. What we need is "access_key_id" and "secret_access_key".

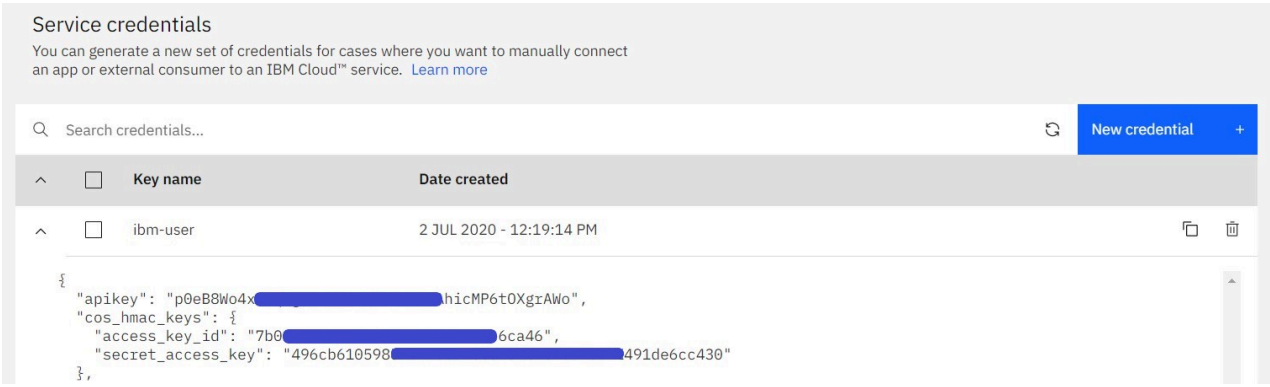

Now we can log in to the Data Protector for Cloud Workloads Dashboard and create a backup destination. Go to the backup destination tab on the left side menu and then choose "Amazon S3 / S3-compatible".

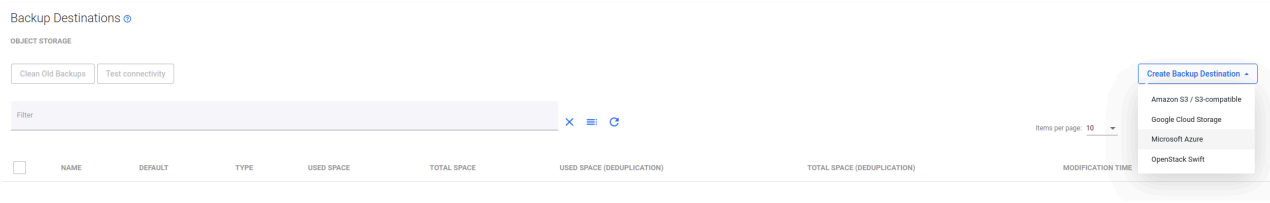

As IBM cloud storage is compatible with Amazon-S3, many settings will be very similar. However, remember to enter the API URL (remember about "https://" at the beginning), select the "Record backup time after store" option, and enter the region.

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

Cancel

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/ ↗.

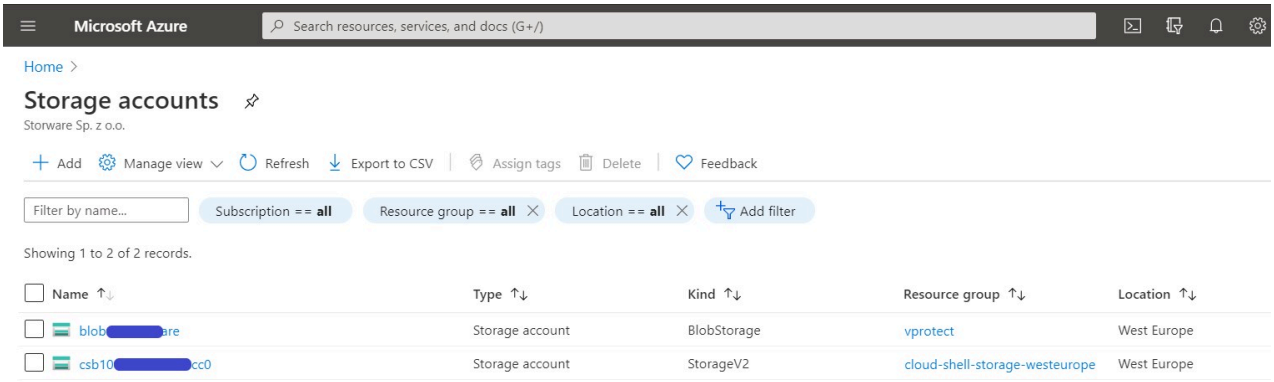

To configure Azure as a backup destination for Data Protector for Cloud Workloads, we just need:

- The storage account name
- One of the account keys

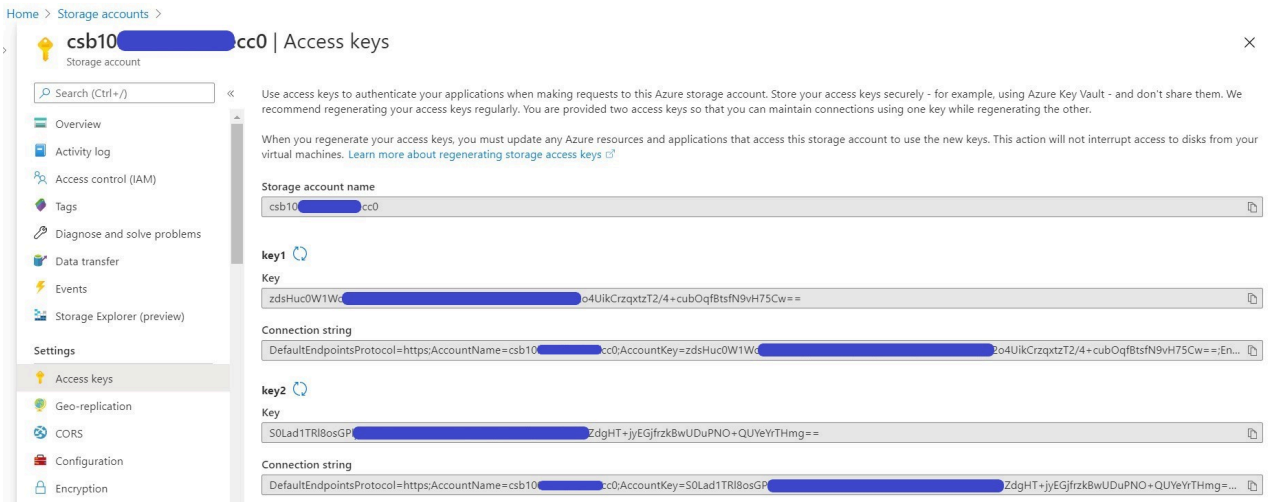

Now you can go to the backup destinations tab in Data Protector for Cloud Workloads and create a new Microsoft Azure backup destination.

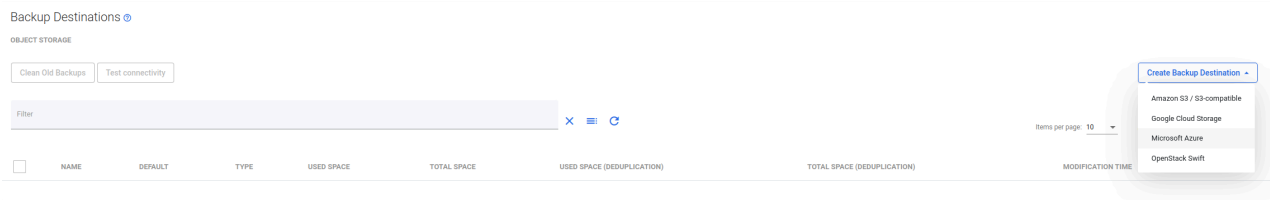

You just need to provide an account name, bucket name and key.

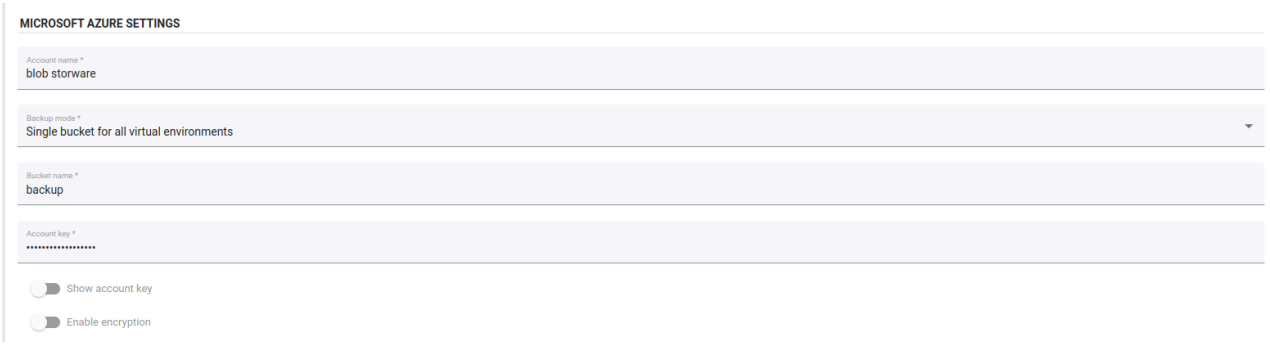

And that's all. As you see, in a few minutes you can integrate Data Protector for Cloud Workloads with Azure Blob storage to securely store your backups

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.

> 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.

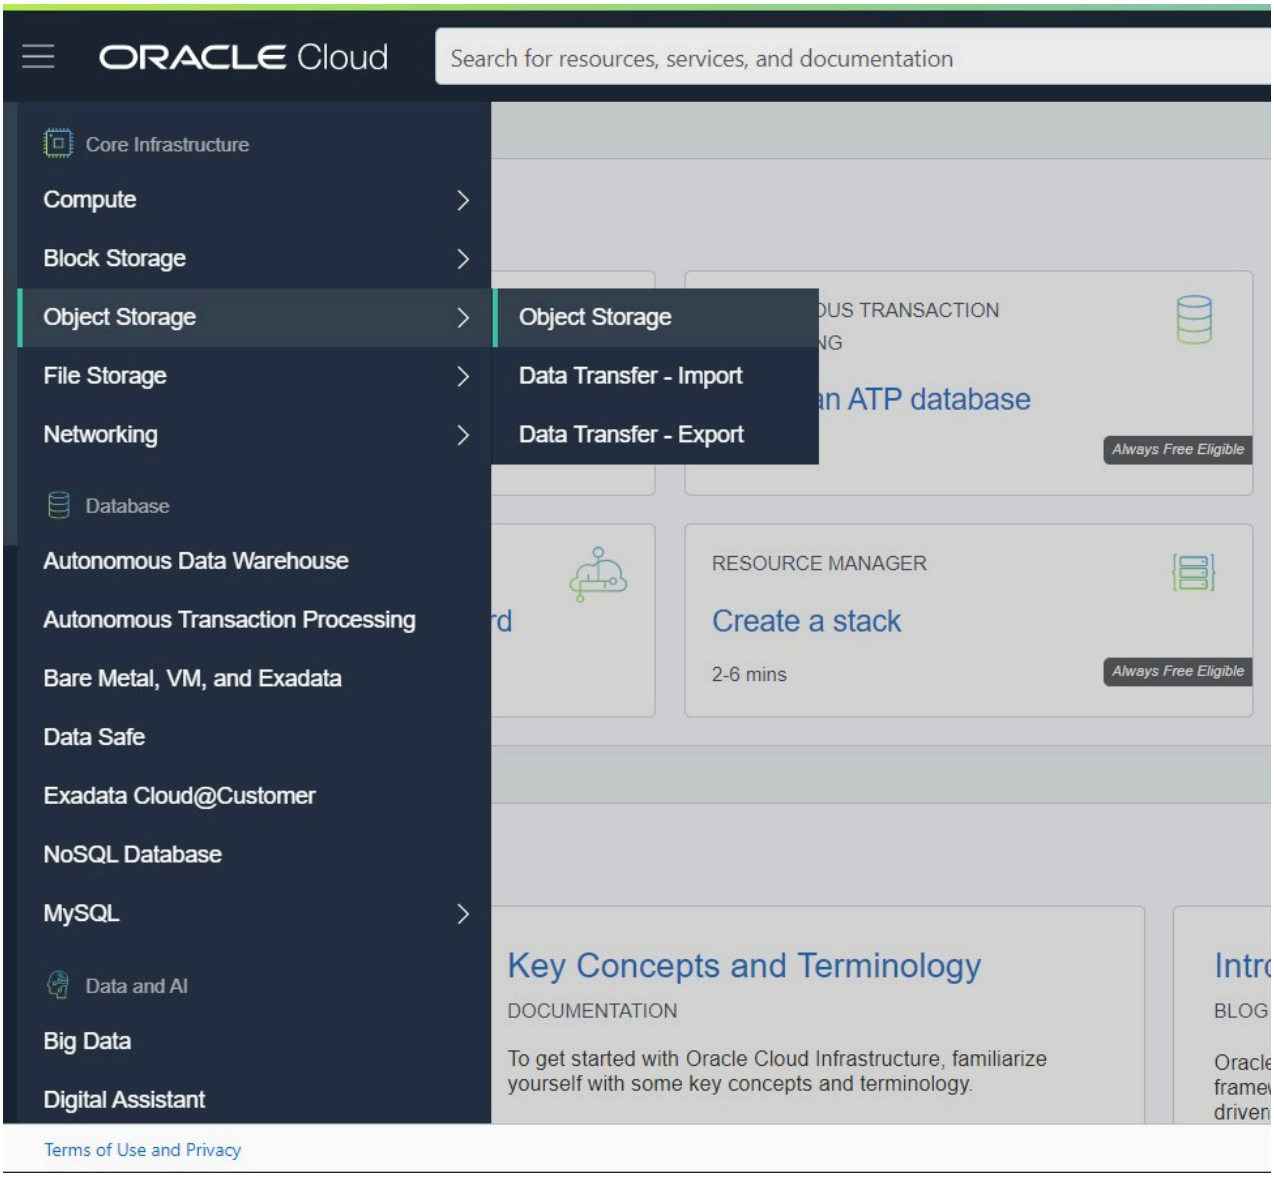

Now let's create a new bucket.

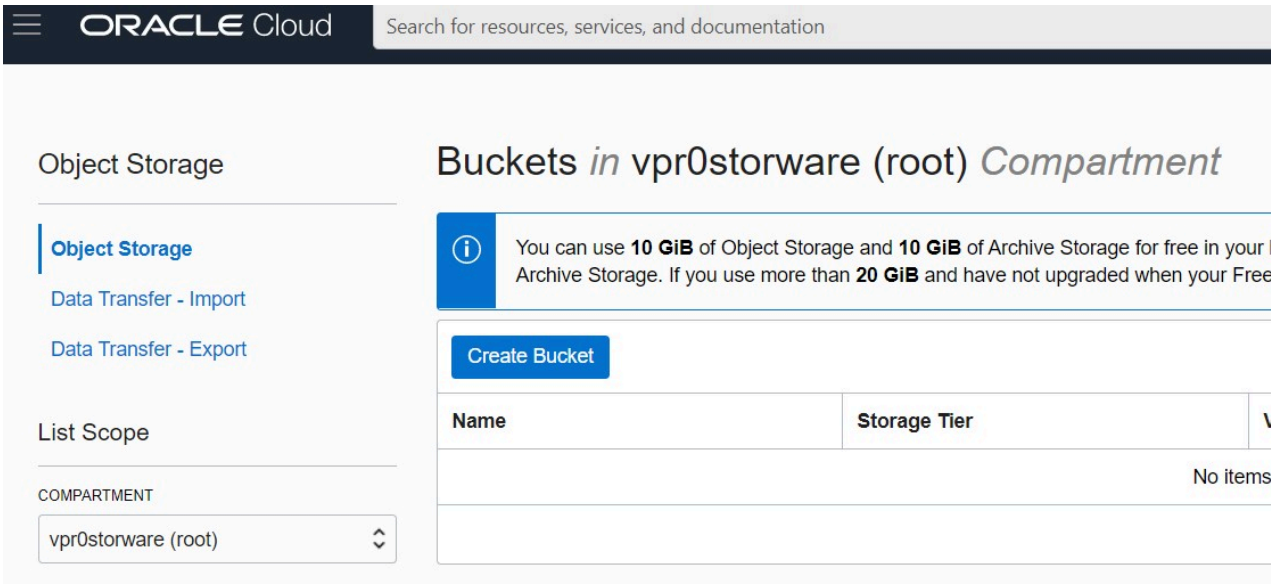

We do not require specific bucket settings for Data Protector for Cloud Workloads. The bucket name will be needed when we want to create a backup destination in Data Protector for Cloud Workloads.

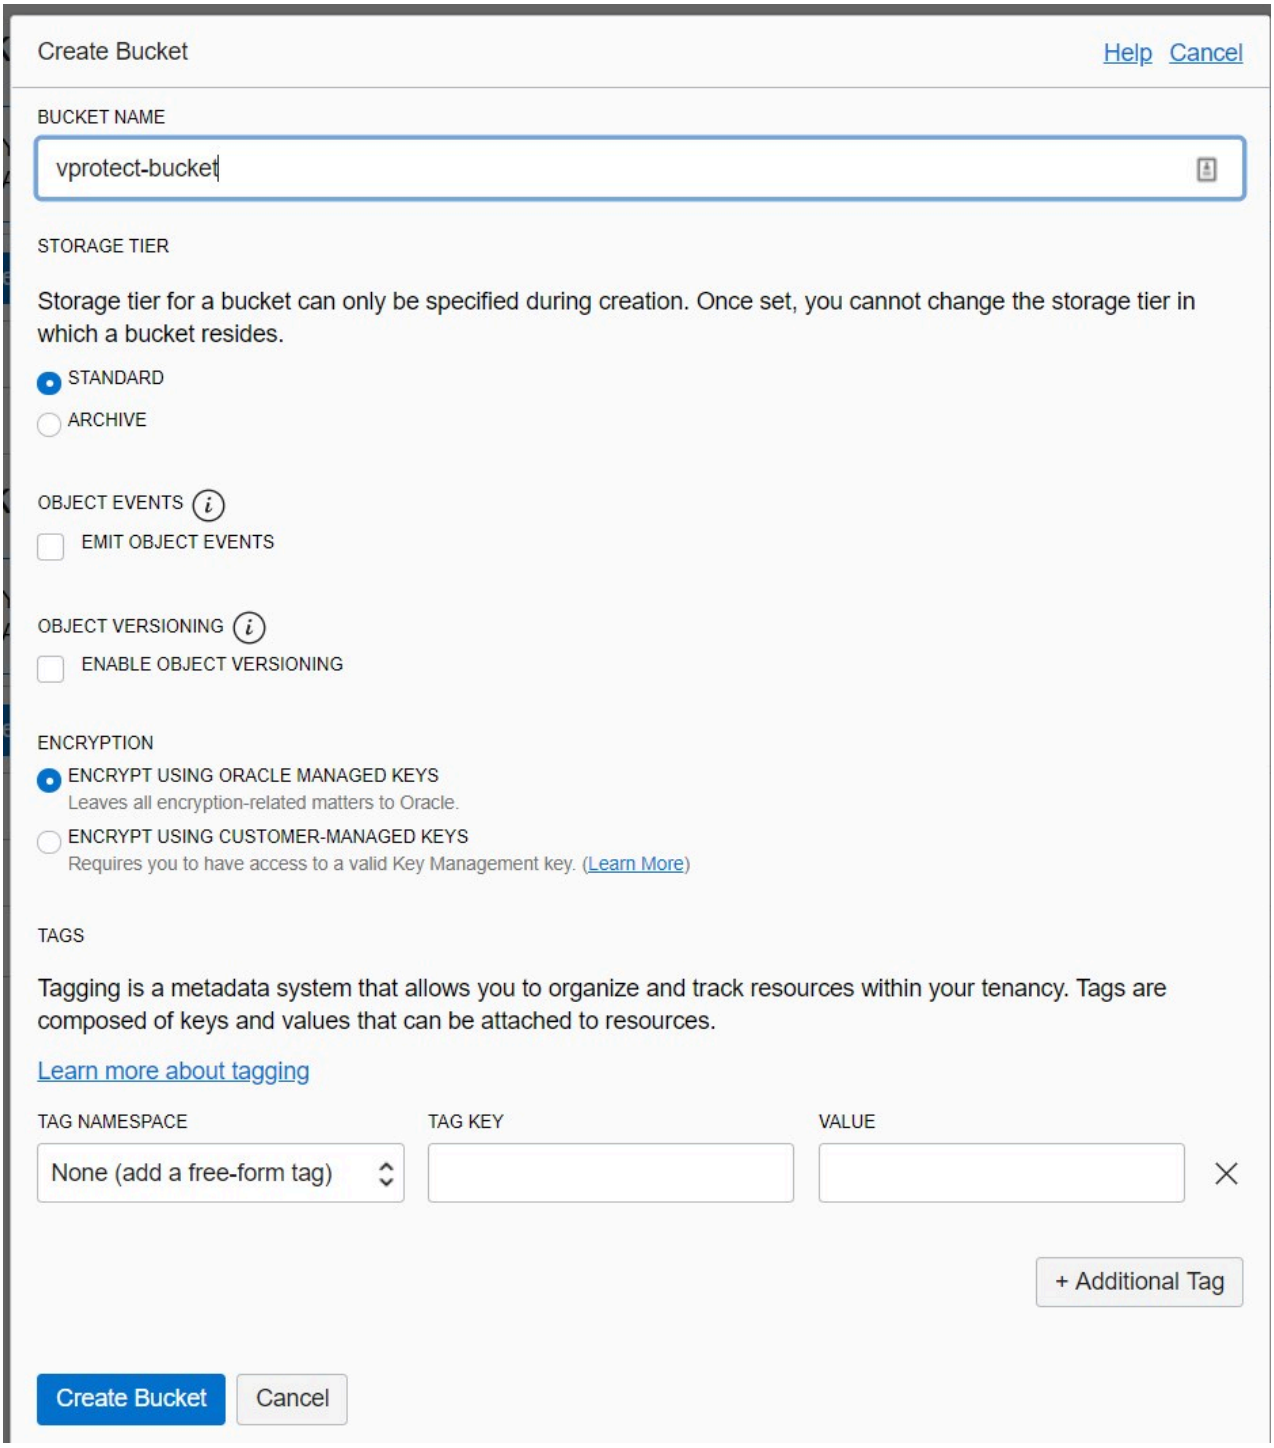

After creating the bucket, you'll see a list of buckets. Click on the name to view the details of the object. Remember the "namespace", we also need it when creating a backup destination.

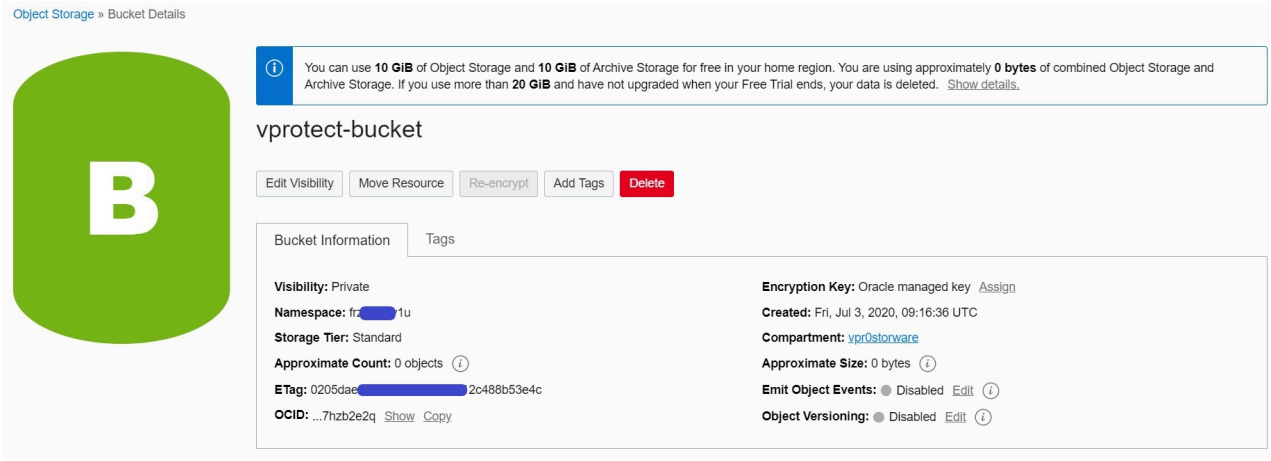

Now we need to create a user that we will use to authenticate our backup destination. Go to the Users tab under the Identity tab in the menu on the left.

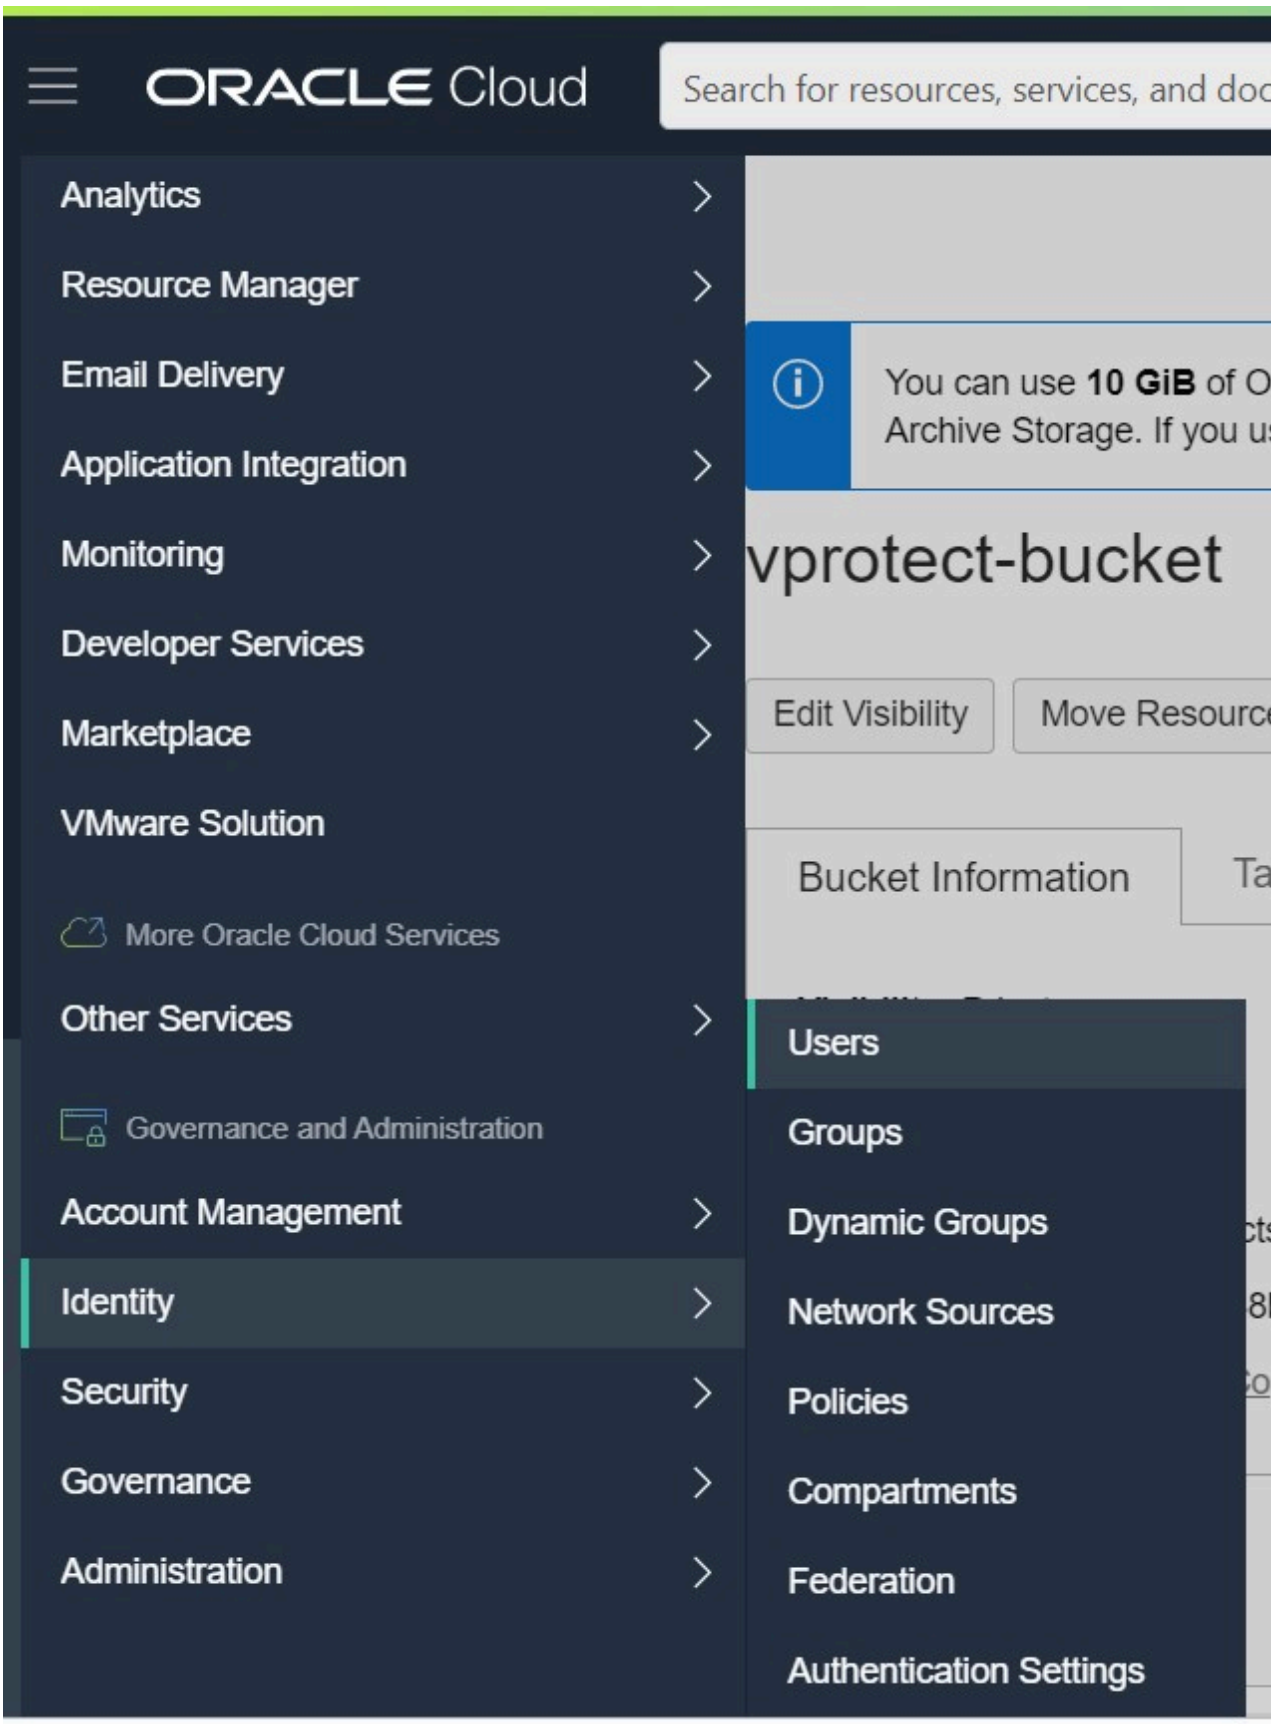

Now create a new user.

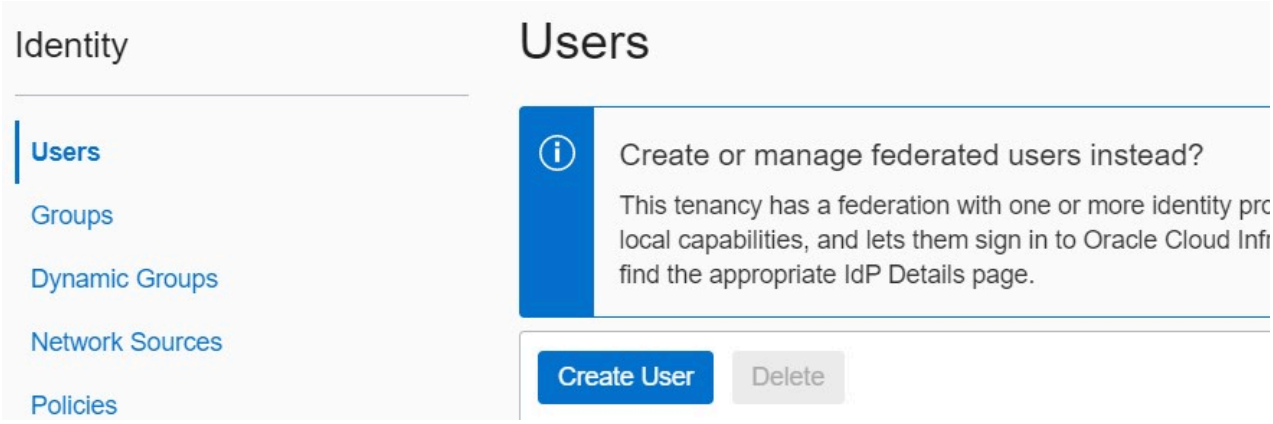

Fill in the required fields.

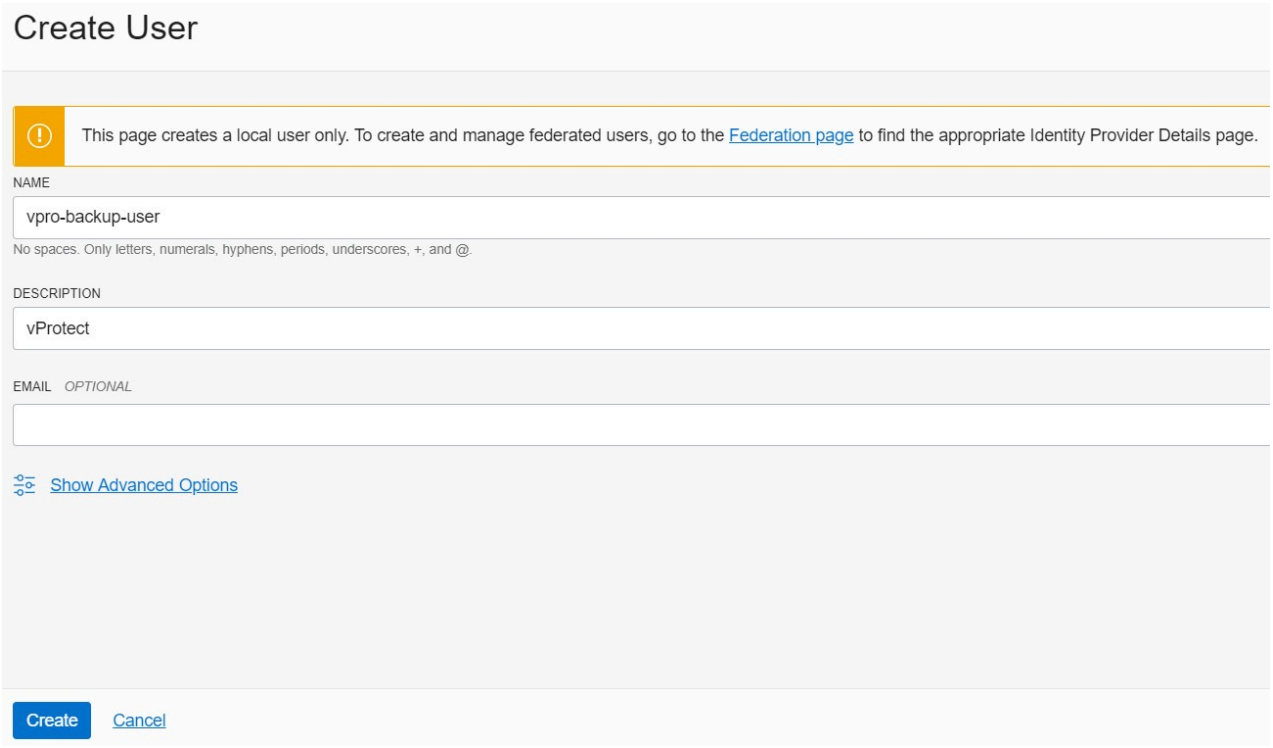

Then go to the Groups page, which you can also find under the identity tab in the left side menu.

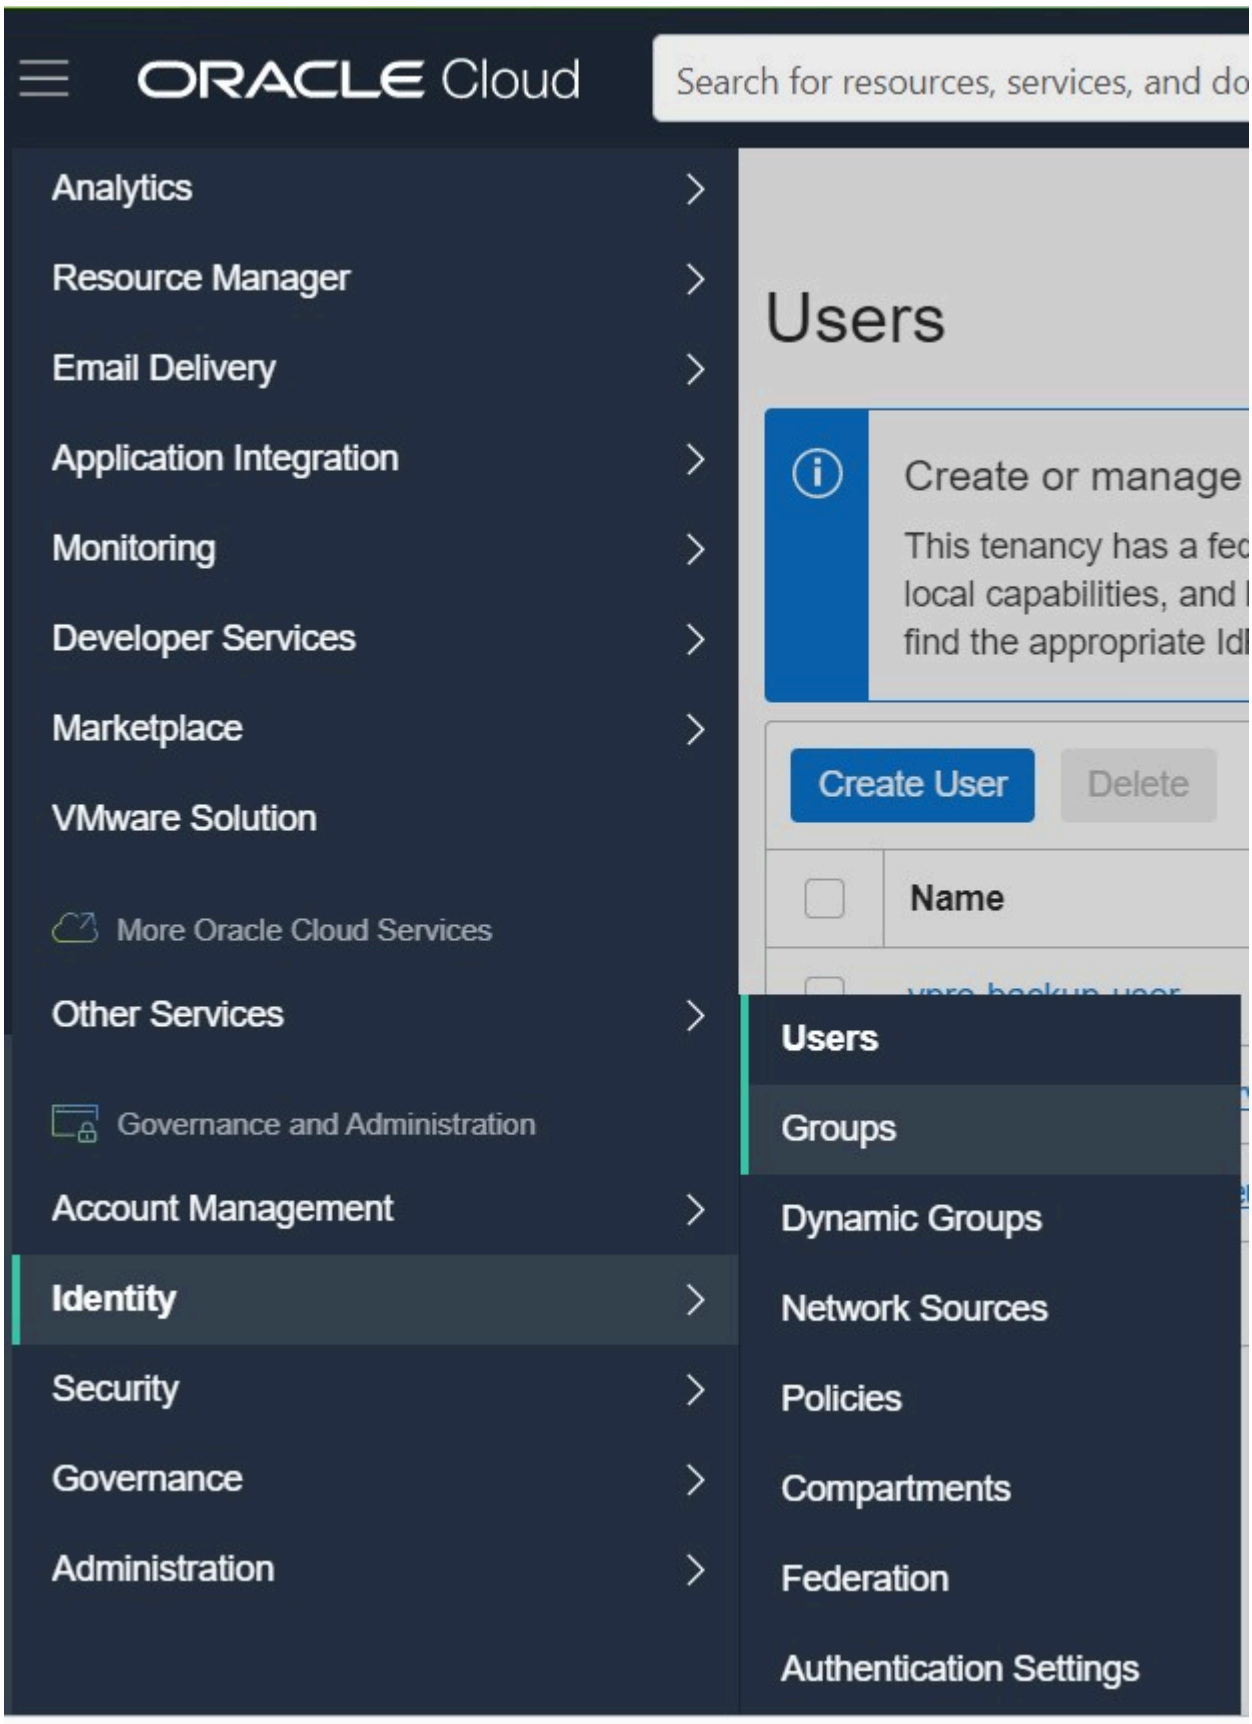

Now click on the existing group "Administrators".

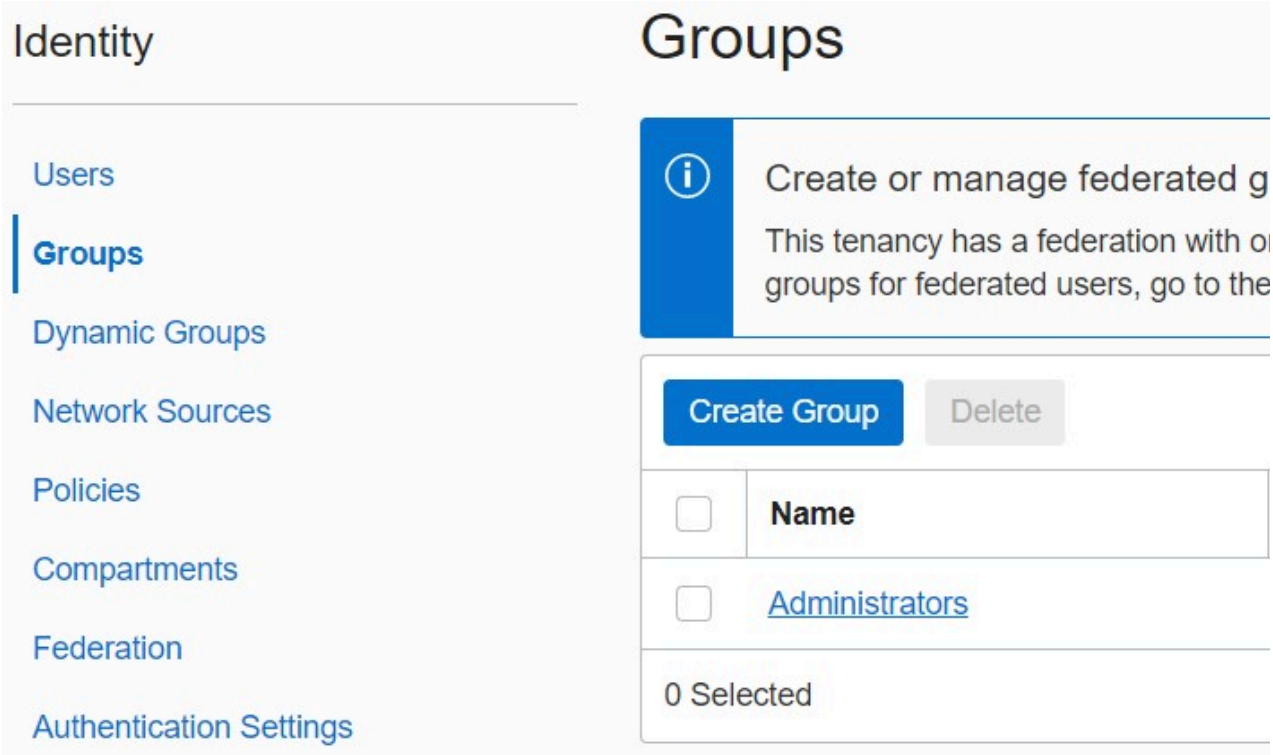

Now click on "Add User to Group" and choose the user you created previously.

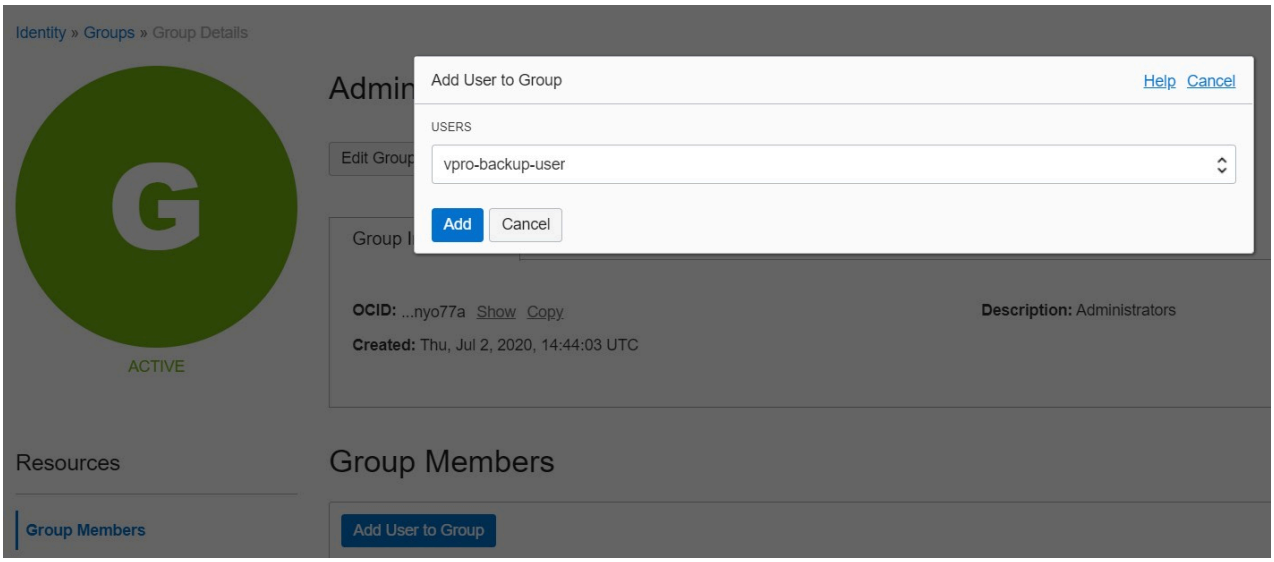

Go back to the Users page and go to the details page of our user.

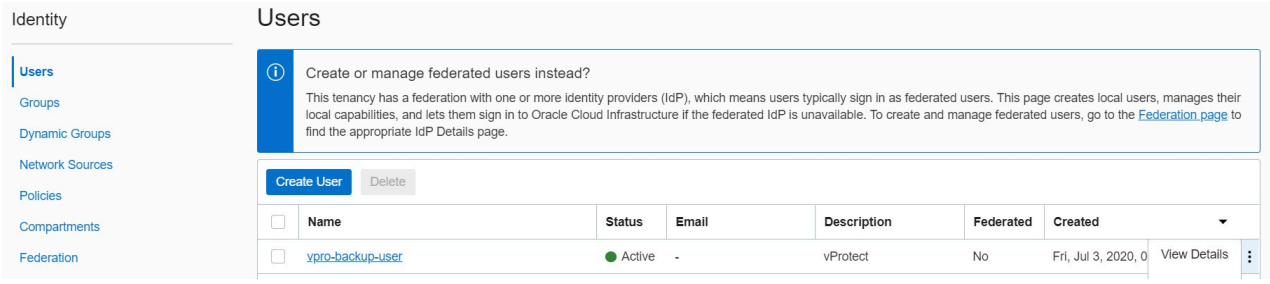

Scroll down and open the "Customer Secret Keys" tab. Click on "Generate Secret Key".

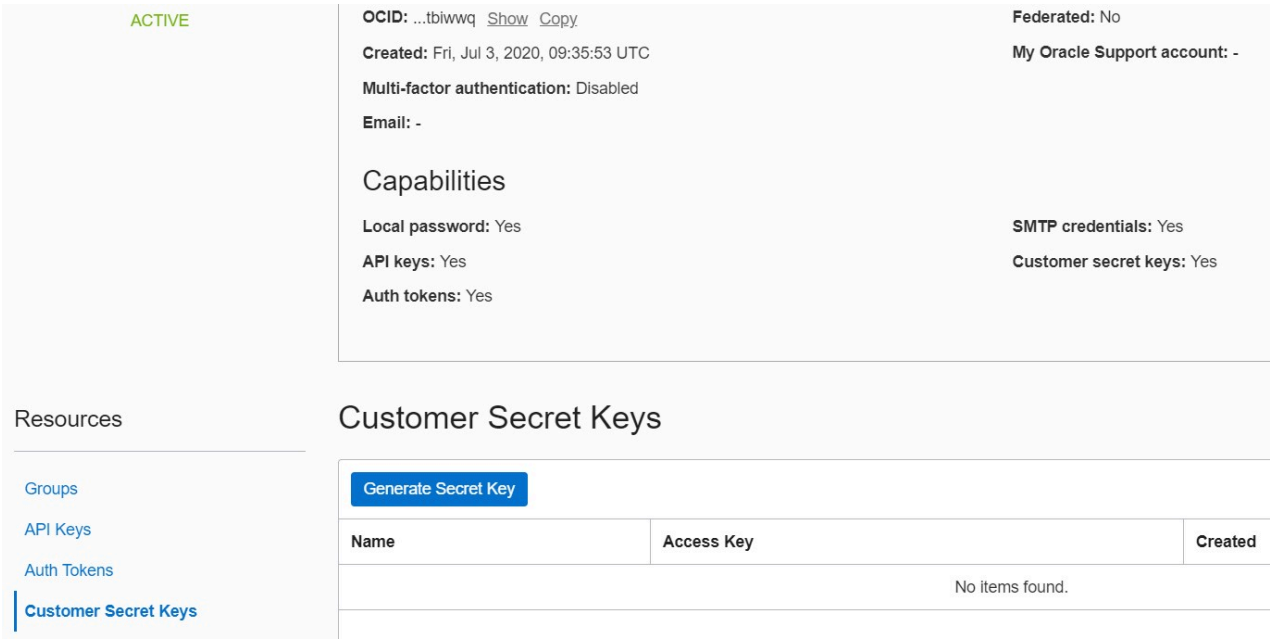

Enter any name.

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

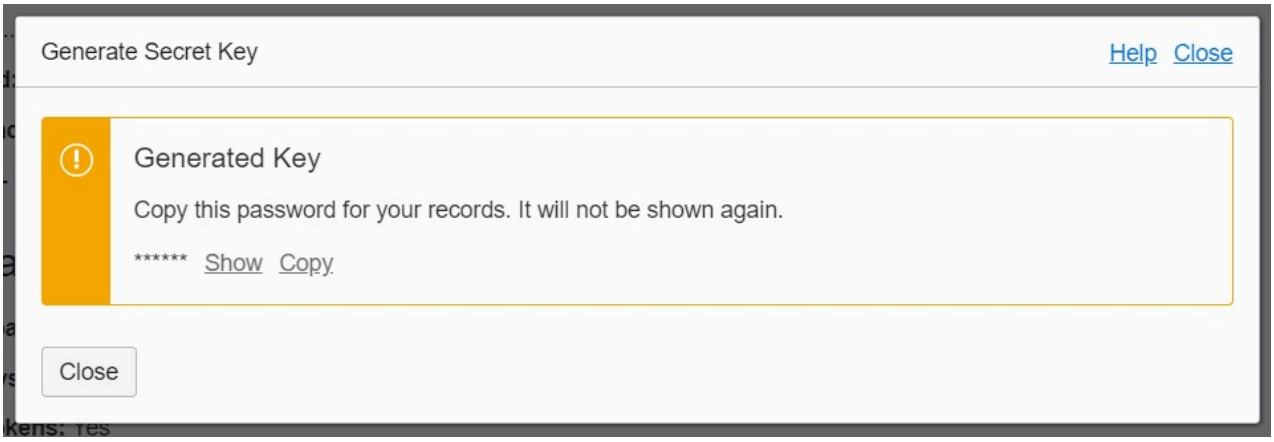

After generating the secret key, you can view the access key, just move the mouse over it.

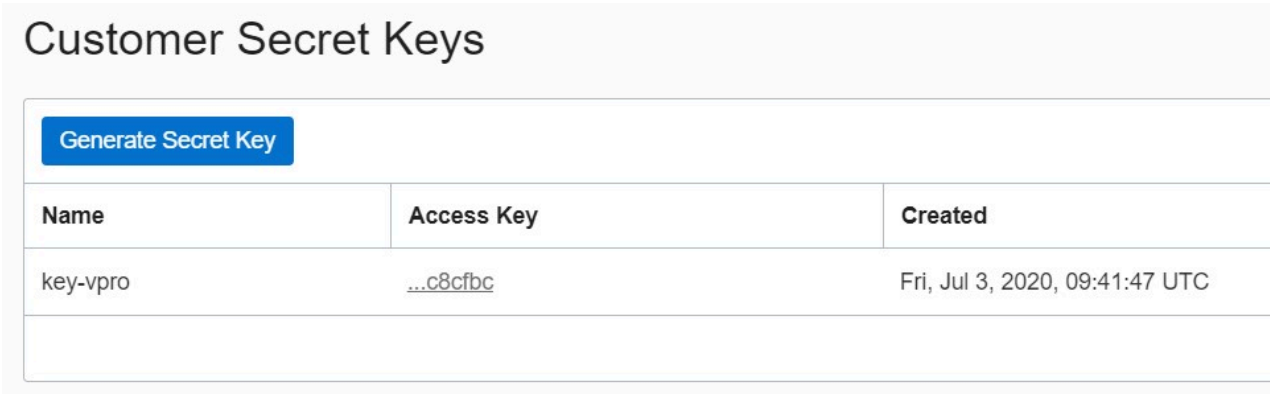

Now we can go to the Data Protector for Cloud Workloads Dashboard. Open the "Backup Destination" tab from the left side menu, then the sub-tab "Object Storage" and choose "Amazon S3 / S3-compatible" as the new type of backup destination.

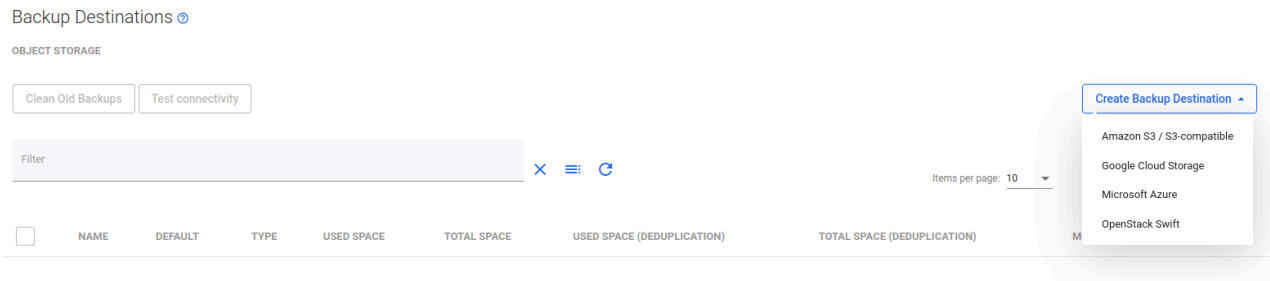

First, let's focus on the "S3-Compatible" section.
To generate an API URL, you will need this site: https://docs.cloud.oracle.com/en-us/iaas/api/#/en/s3objectstorage/20160918/ ↗
As We mentioned earlier, you will need an object storage namespace (choose the API URL from the list according to your region).
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

Cancel

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.

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".

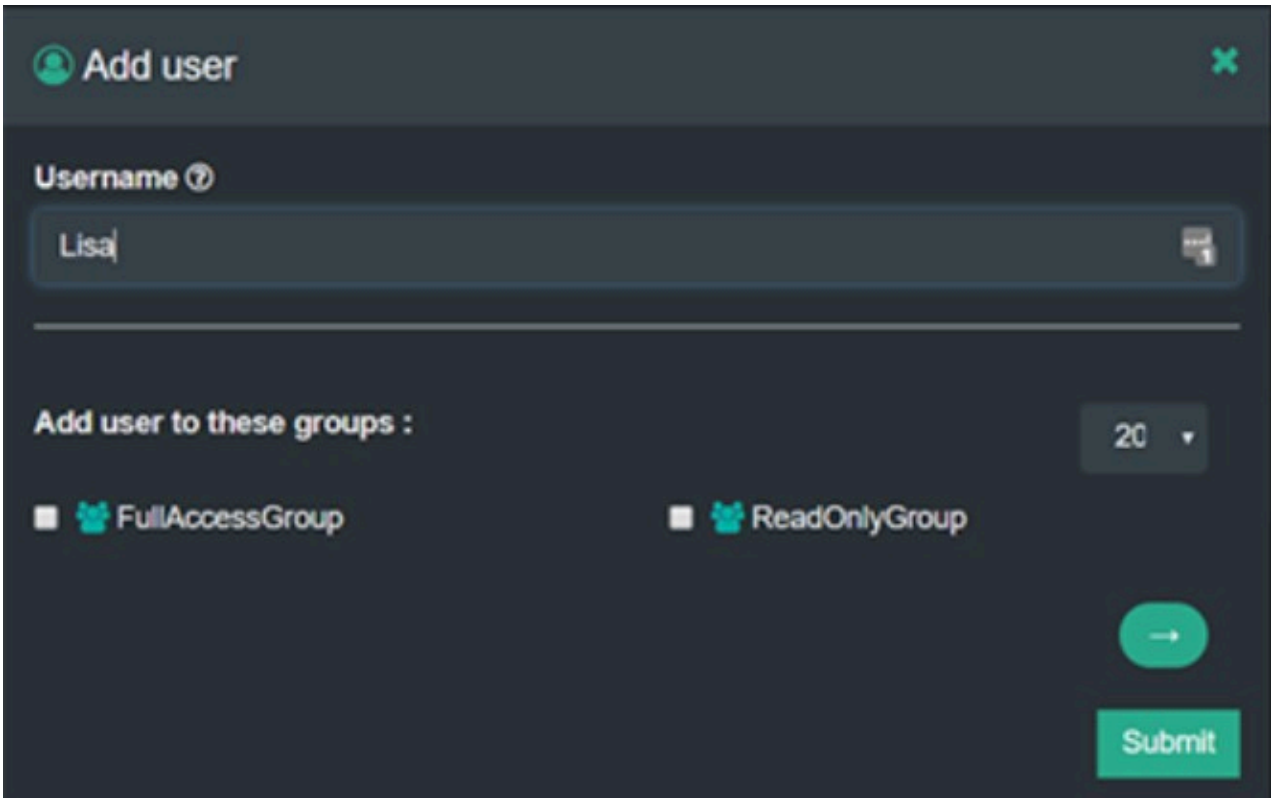

The user management panel displays the user name and the Amazon Resource Name (ARN).
Now we will generate the access and secret keys for the user.
Click on the key icon in the Actions column of the user row.

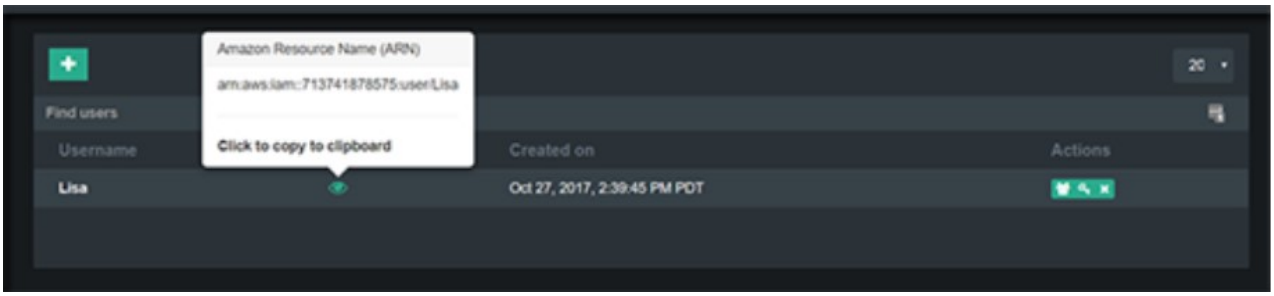

Click on Generate a new key.

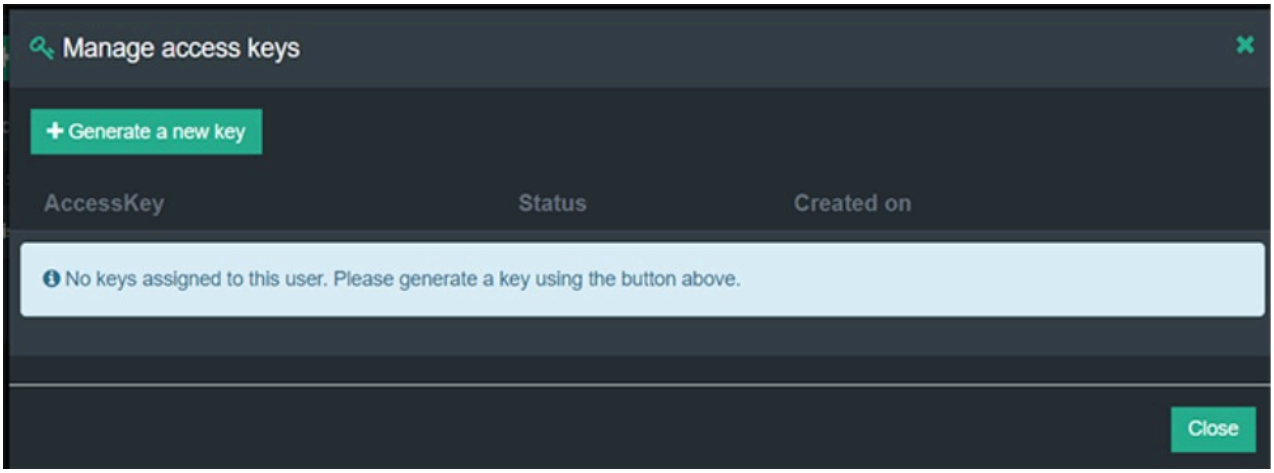

Click on Proceed to generate the user's AccessKey and SecretAccessKey.

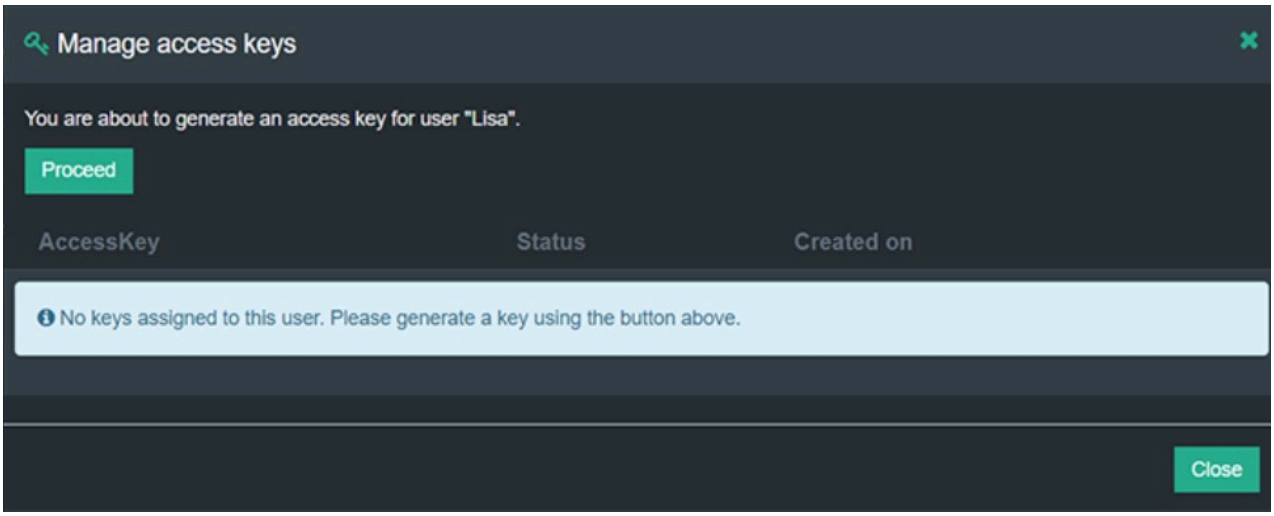

Copy and save the SecretAccessKey to a secure location. It is not shown again and cannot be recovered later.

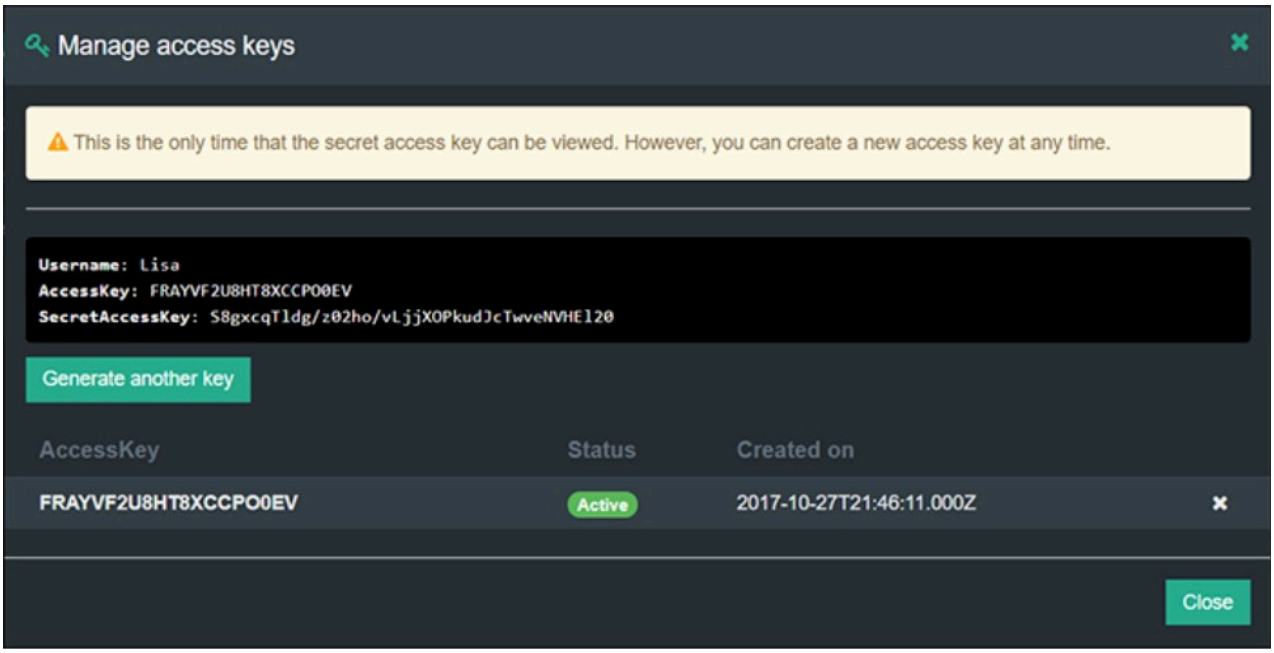

Now we can go to bucket creation. Go to the S3 Browser interface.

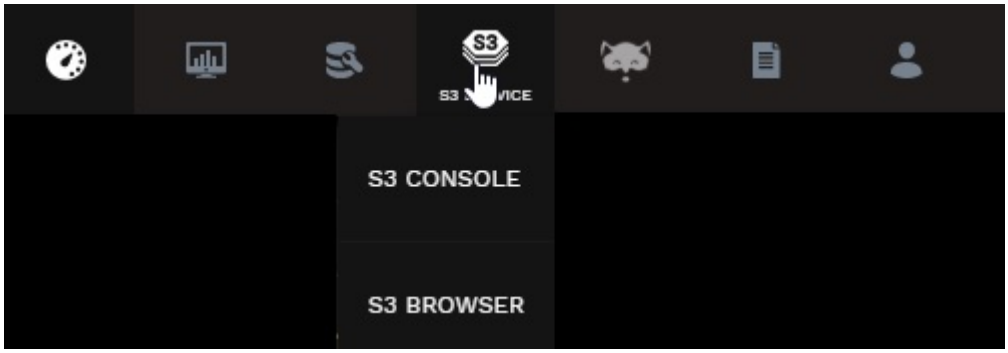

The S3 Browser opens the main window, from which one can see the entire roster of buckets.
Click the Create Bucket button in the top left of the main window.

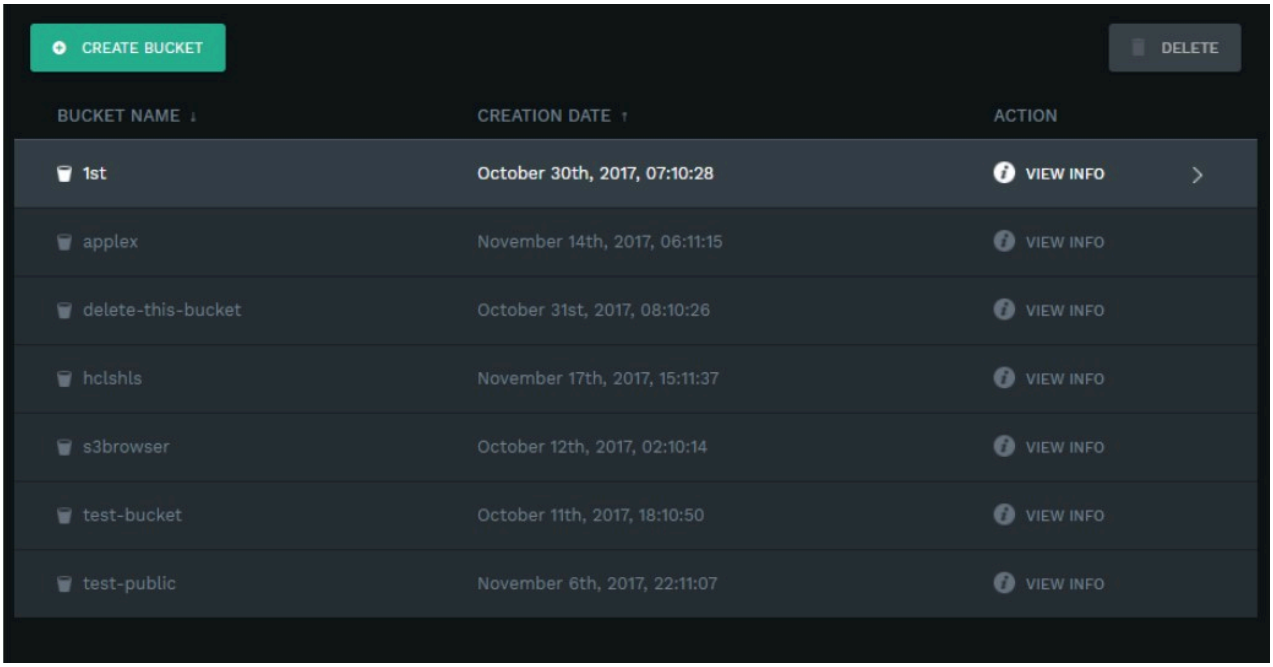

Enter a name for the new bucket and click on Create button.

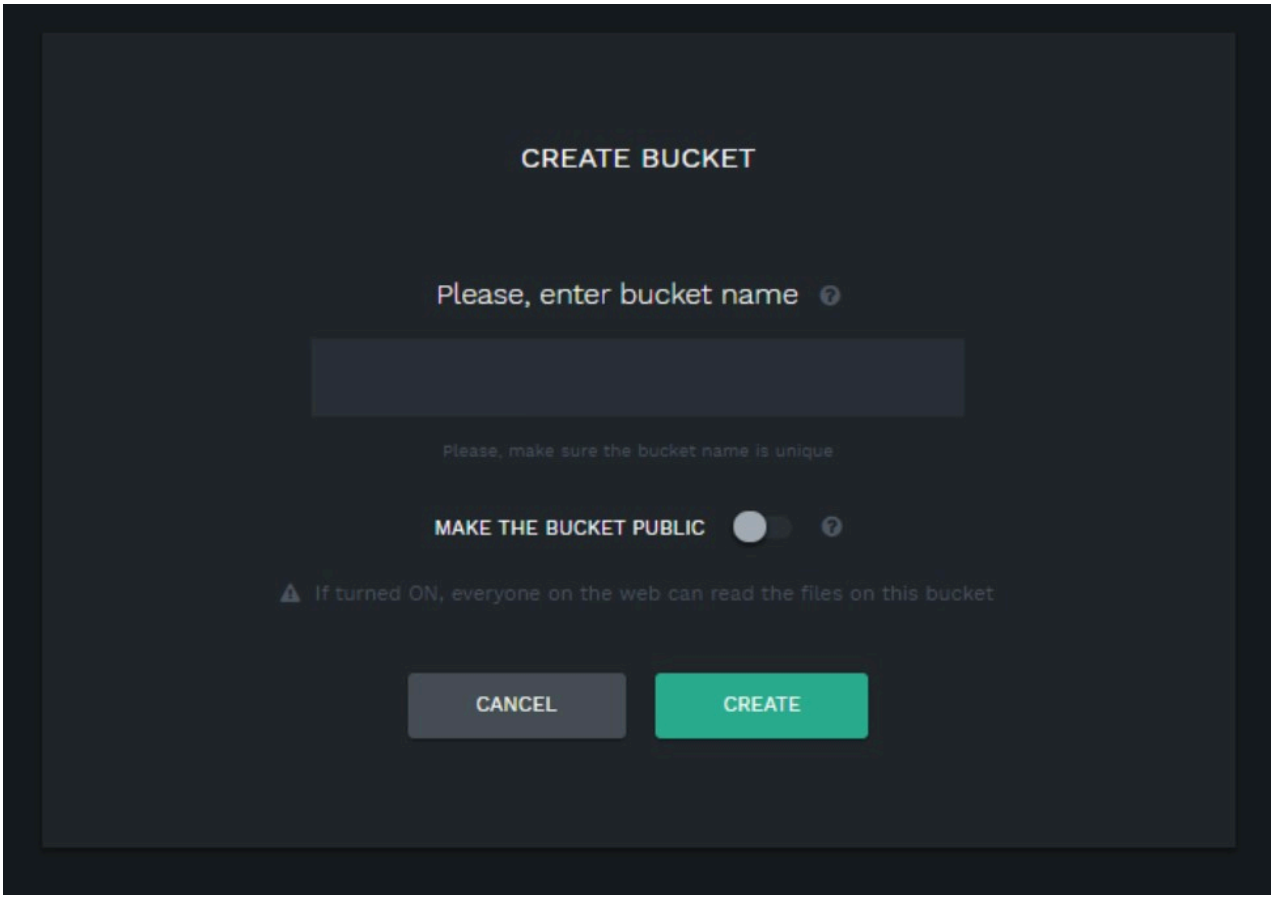

That's all on the Scality side. Now we can go to Data Protector for Cloud Workloads.

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

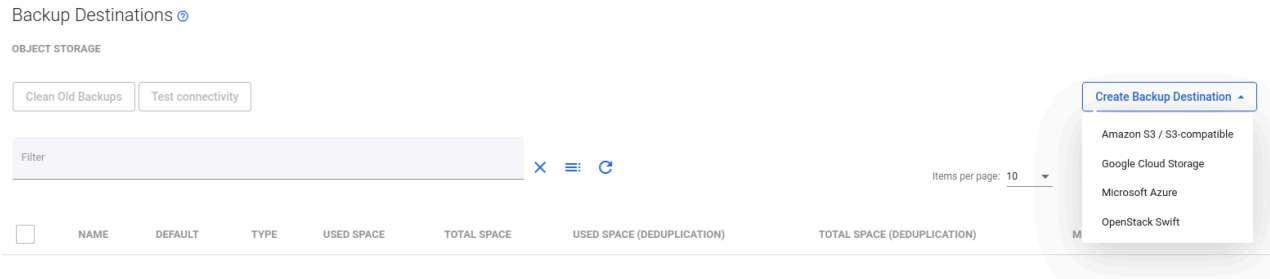

Like in other S3-compatible backup destinations, you have to fill in the fields below and provide the access and secret key.

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
Device Type                           Pool Name
========================================================================
========================================================================
=========================
dyskd_gw1                         win-srv-proxy
Backup To Disk StoreOnce software de dyskd_MediaPool
vp_gw2                            protector11.lab.local
Backup To Disk StoreOnce software de vp_MediaPool
vp_protectorlab_gw1          protectorlab.lab.local
Backup To Disk StoreOnce software de vp_protectorlab_MediaPool
vp_protectorvp_gw1                protectorvp.lab.local
Backup To Disk StoreOnce software de vp_protectorvp_MediaPool
========================================================================
========================================================================
=========================
```

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.

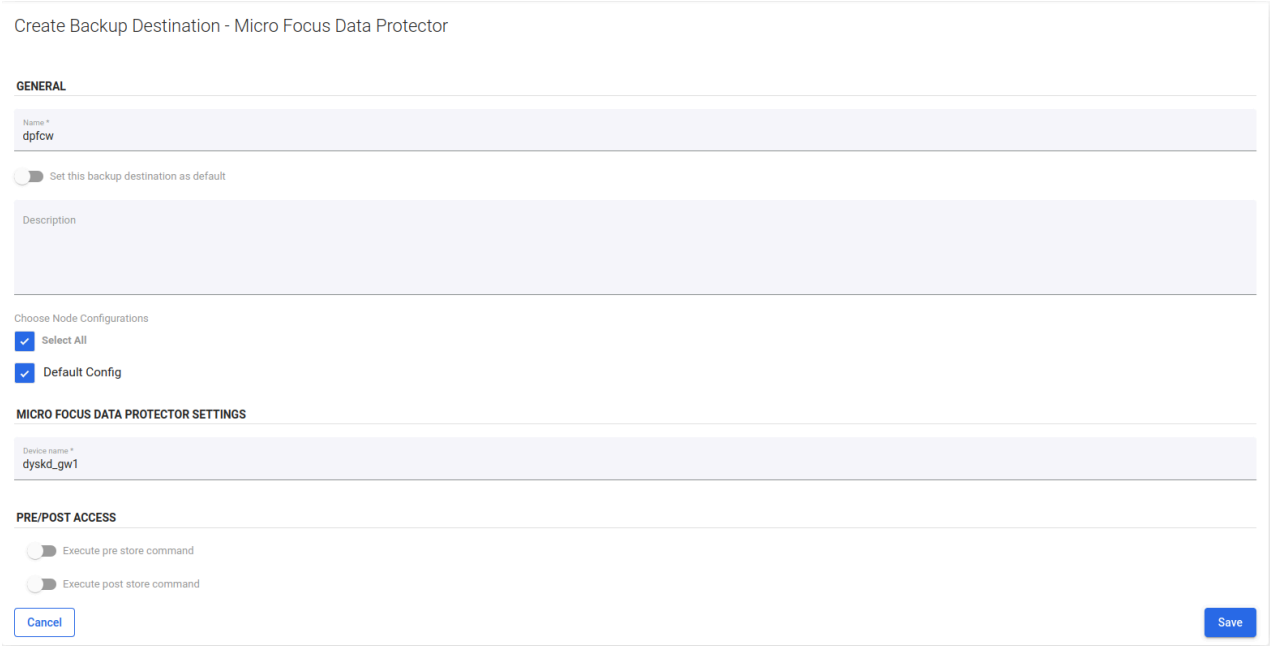
Create Backup Destination - Micro Focus Data Protector
GENERAL
Name *
dpfcw
Set this backup destination as default
Description
Choose Node Configurations
Select All
Default Config
MICRO FOCUS DATA PROTECTOR SETTINGS
Device name *
dyskd_gw1
PRE/POST ACCESS
Execute pre store command
Execute post store command
Cancel
Save

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.

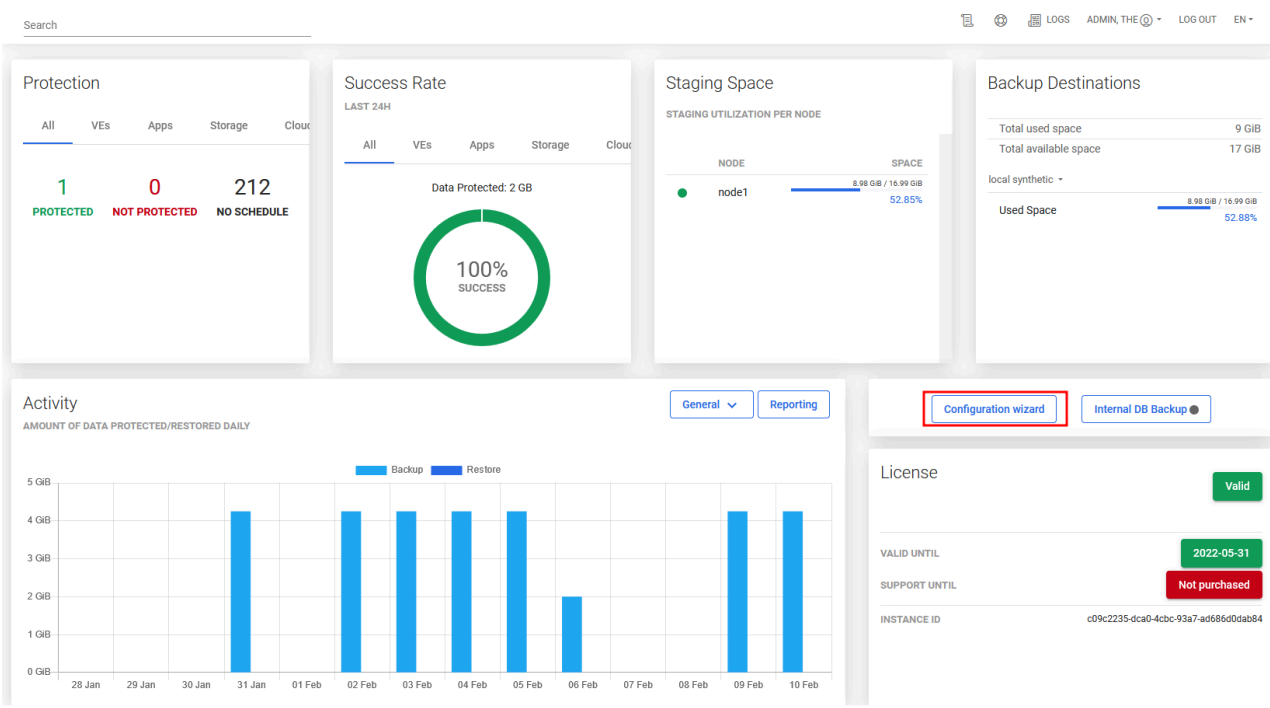

Welcome page - nodes

- 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.

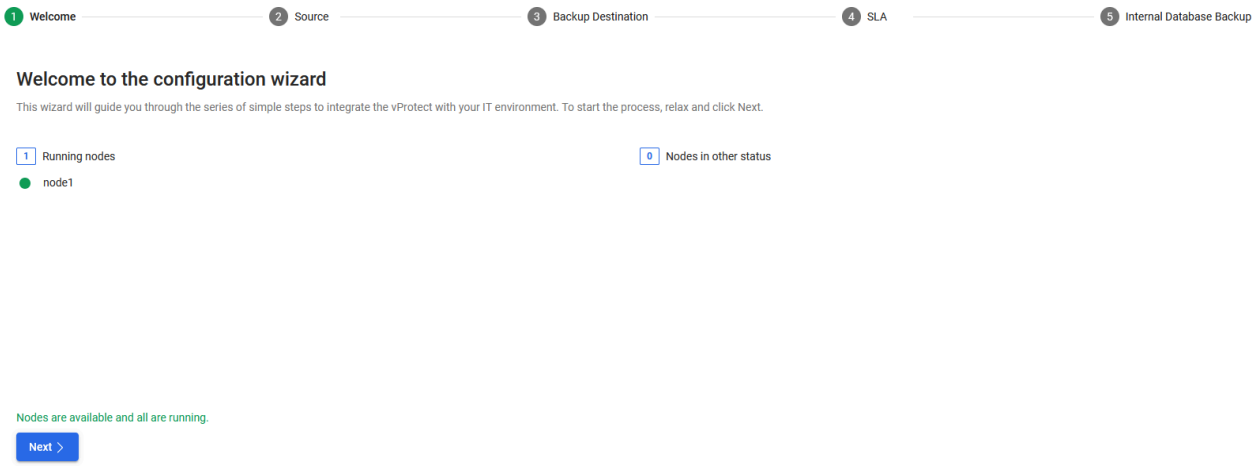

Add a hypervisor

- 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.

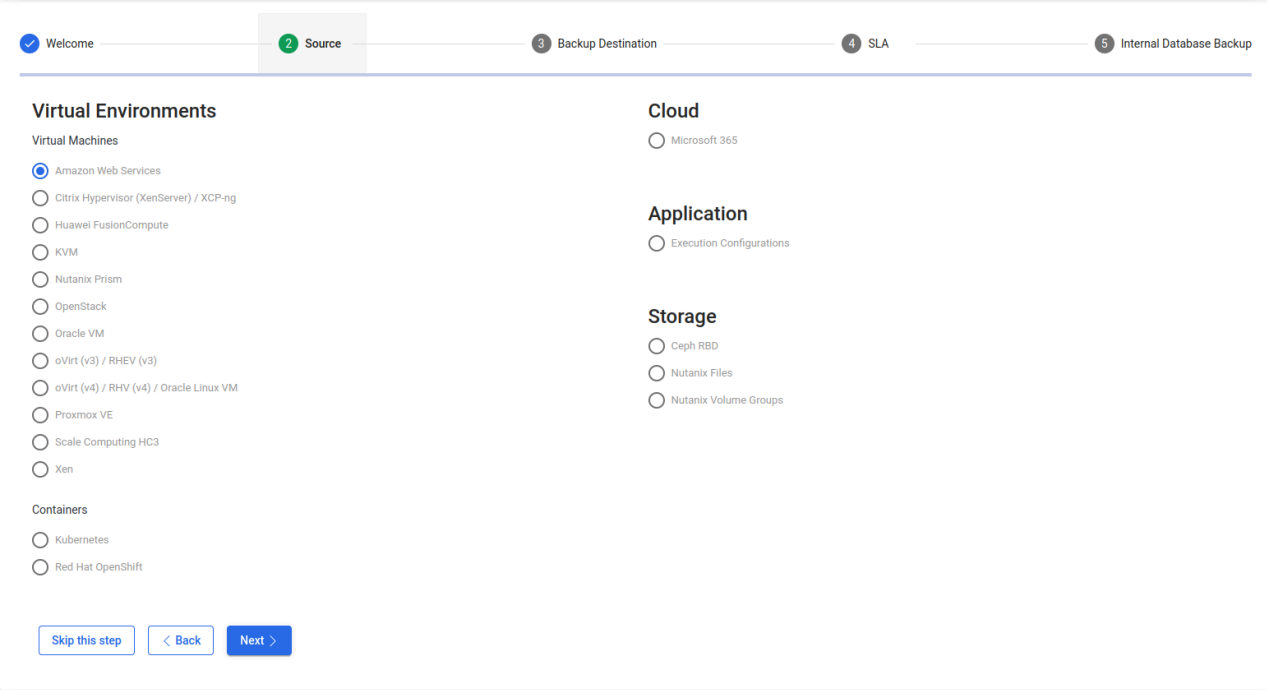

- For the Citrix hypervisor (as an example) you have to enter the following parameters

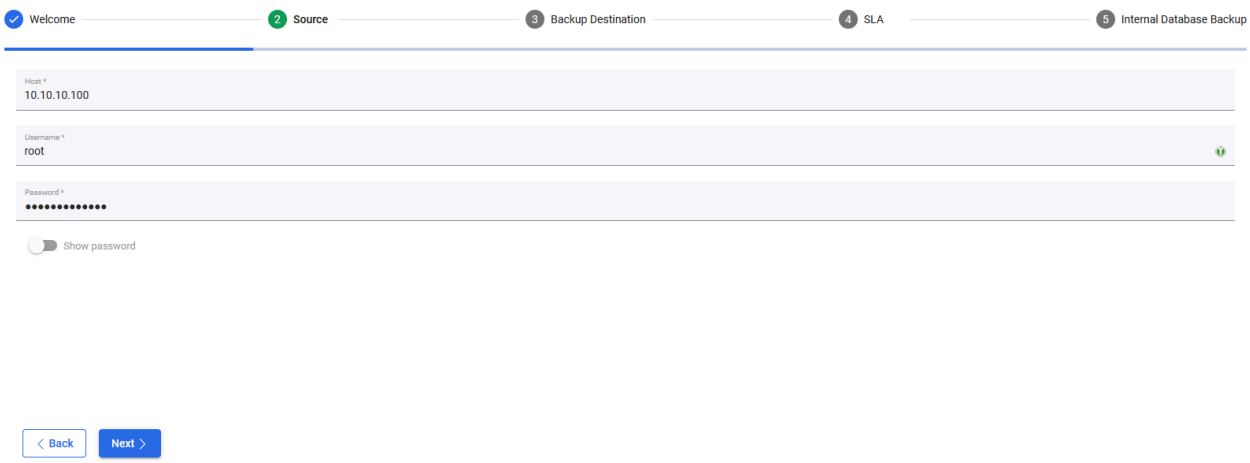

- Choose node configuration

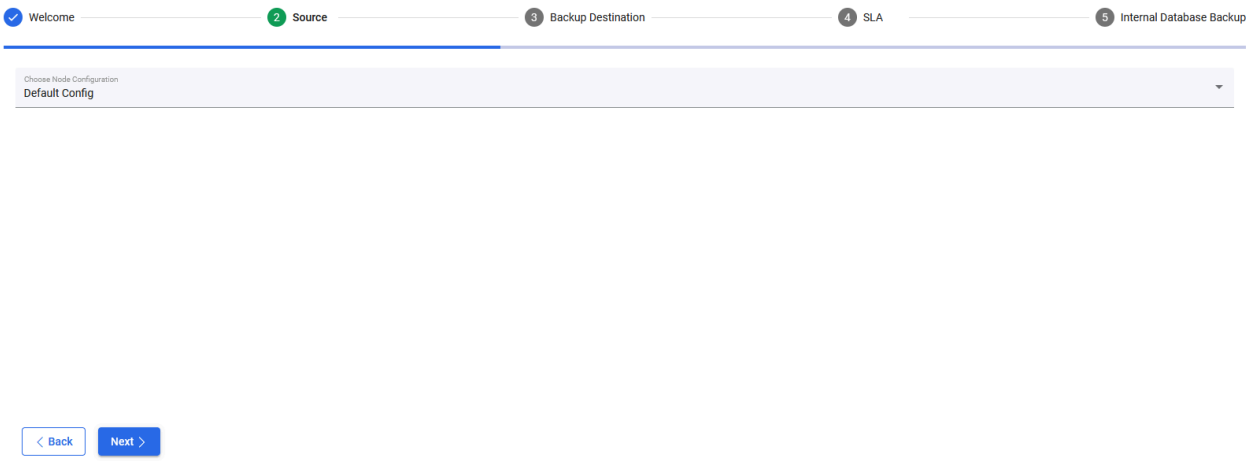

Select a backup strategy for your hypervisor

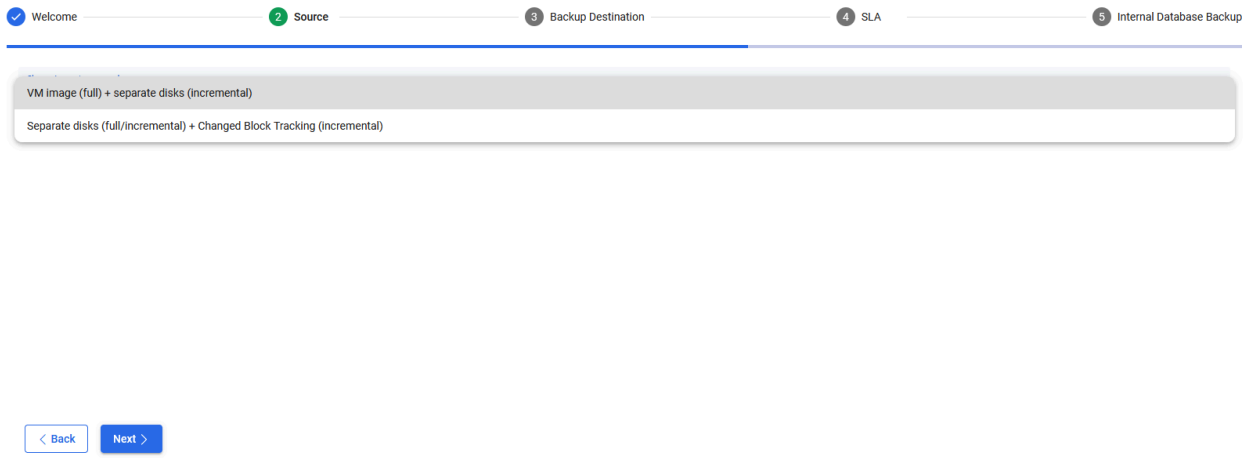

- Optionally, you can add an additional NIC for transfer purposes (provide IP address)

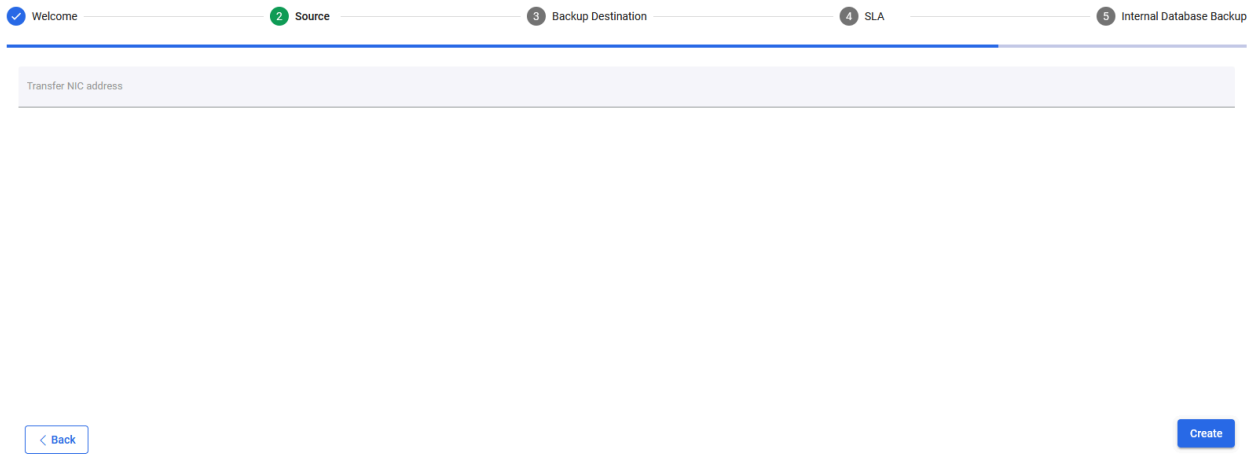

- At the end, you will see a popup window that allows you to run inventory synchronization. After that, you should see all the virtual machines from that hypervisor.

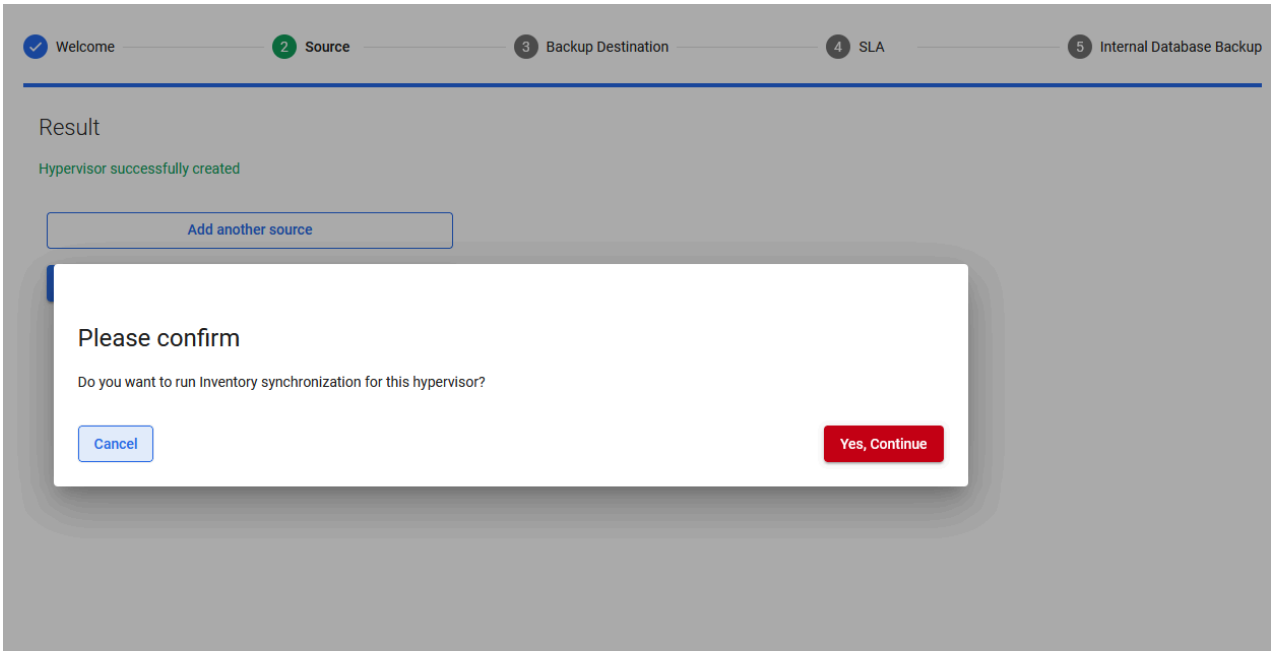

Add backup destination

- 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)

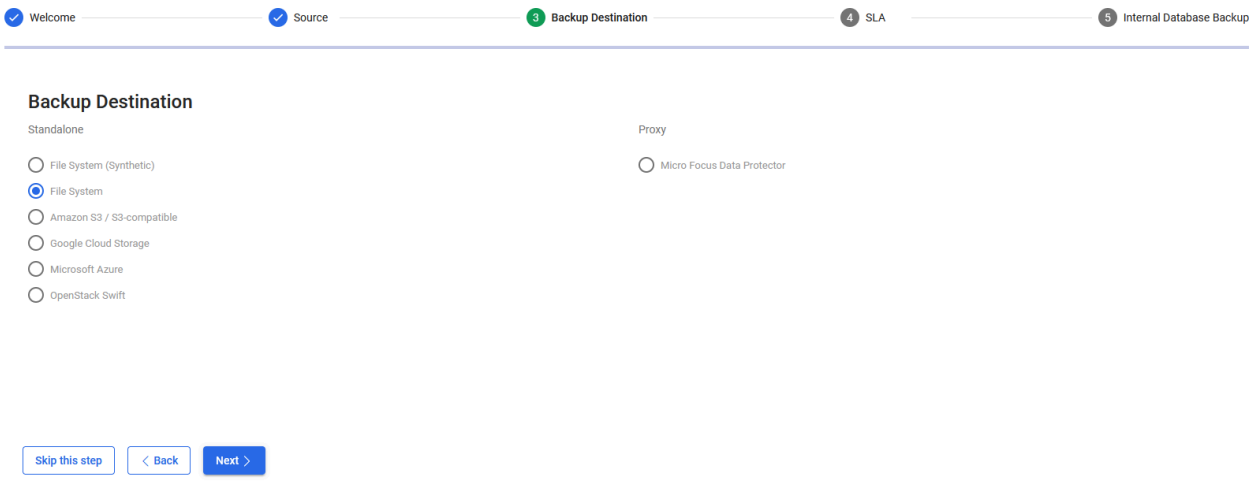

- First, enter a name for your backup destination

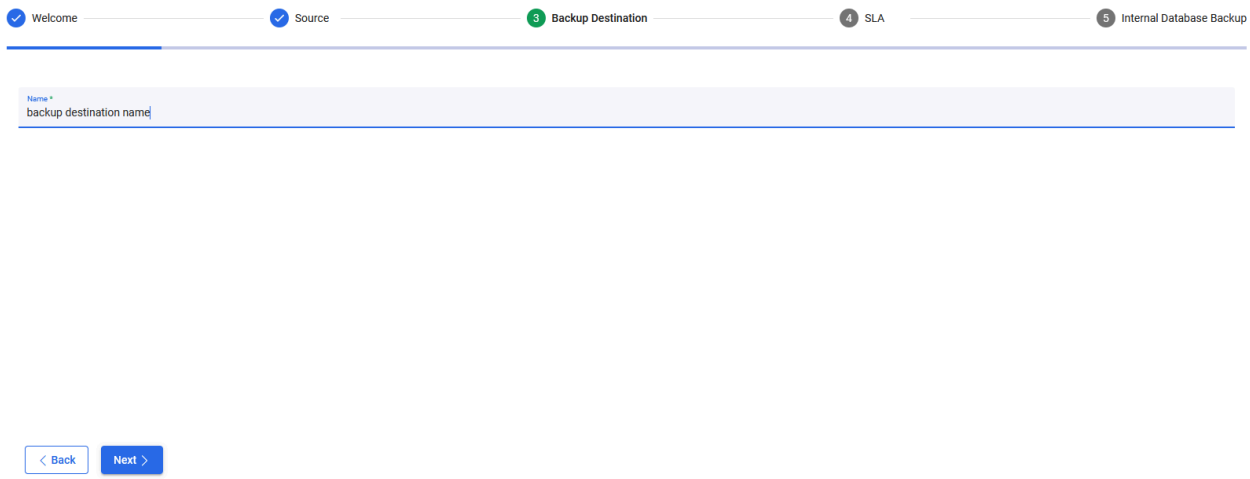

- Choose, if you want to use deduplication based on Virtual Data Optimizer (VDO)

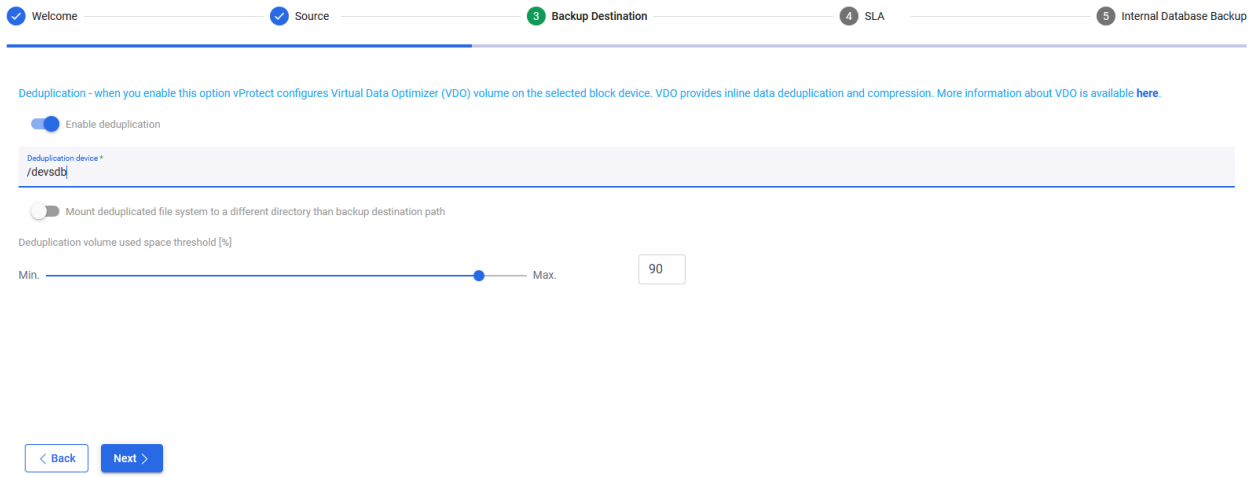

- Set up a storage path, where your data should be stored

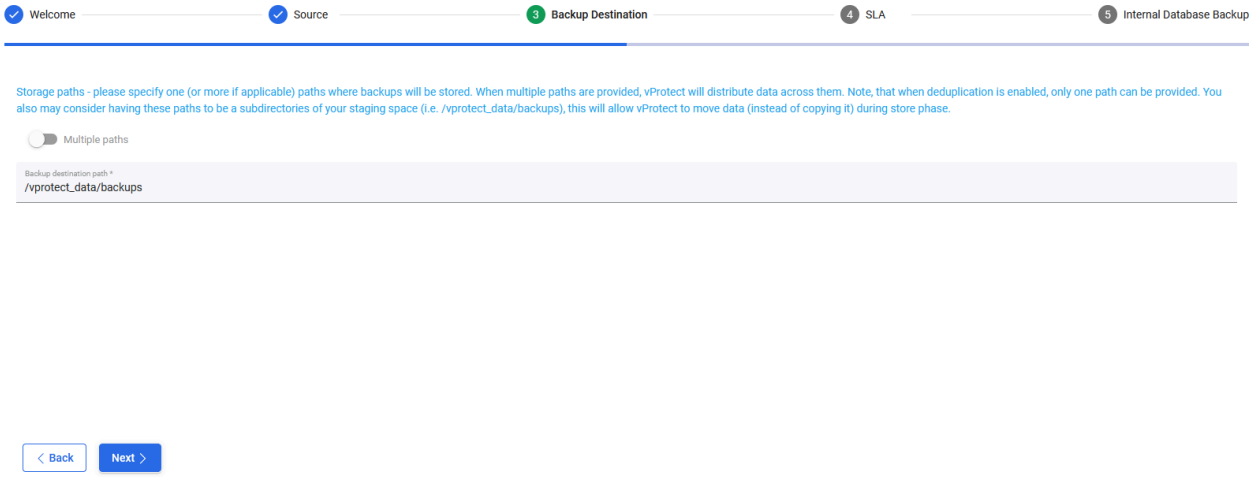

- Optionally you can enable encryption (AES-256 algorithm) - if you enable it, remember that you will not benefit from deduplication.

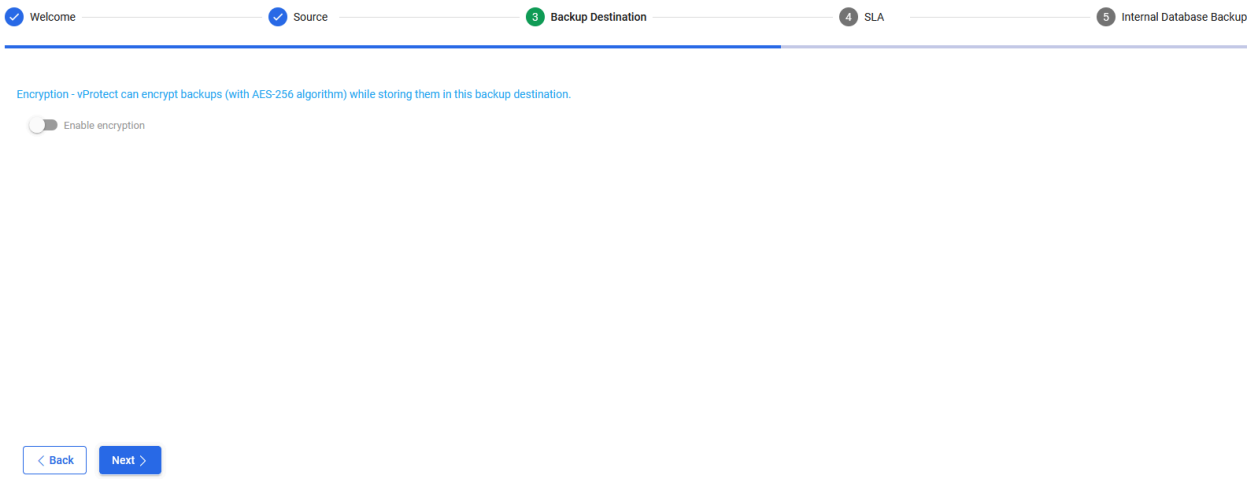

- configuration for pre/post execution command. If you use a File System with VDO, skip this step.

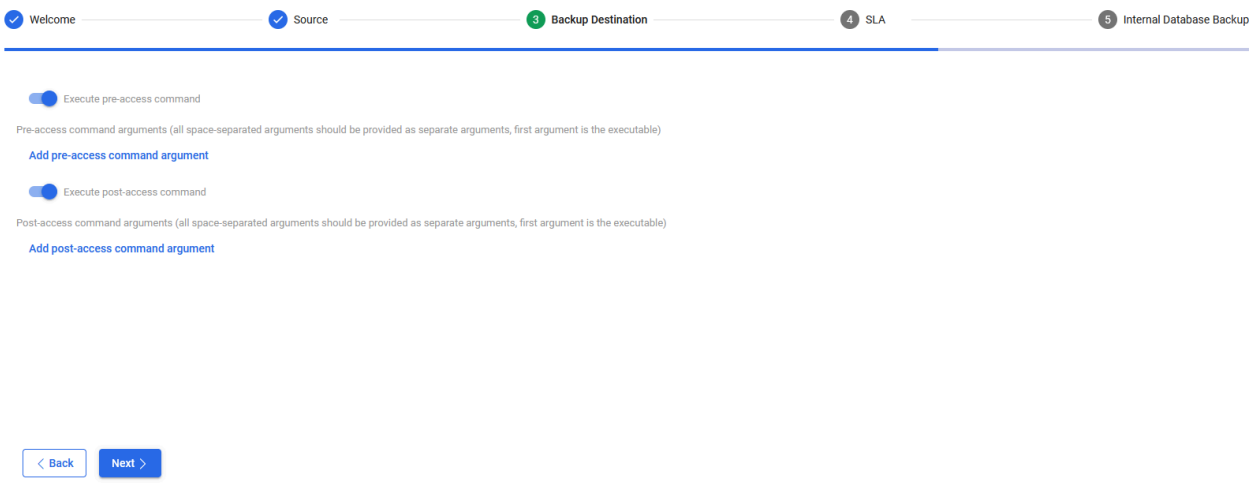

- Decide if you want to set up this backup destination as the default one.

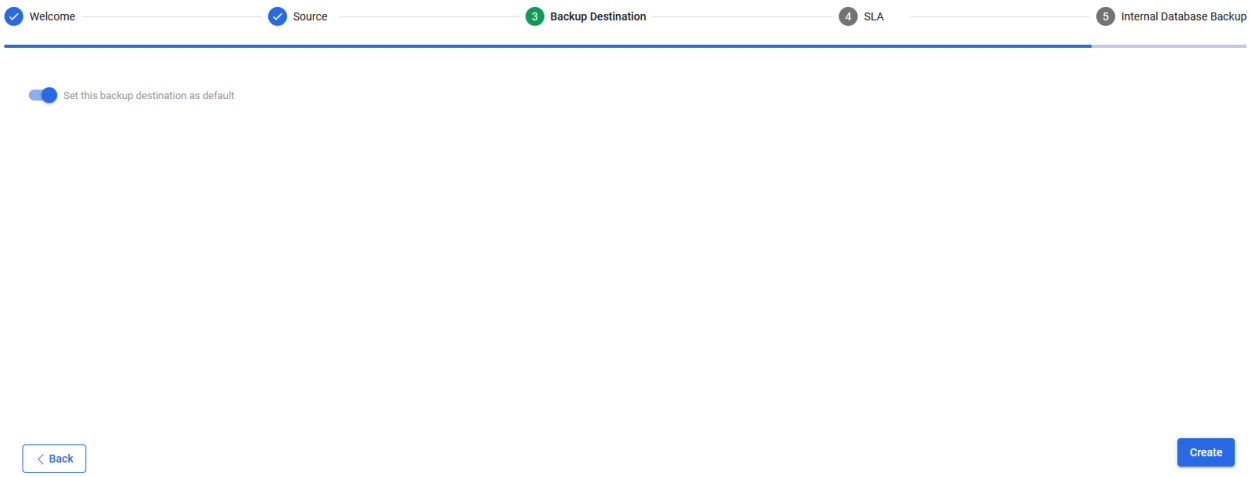

- Finish this step by going to the next section or adding another backup destination

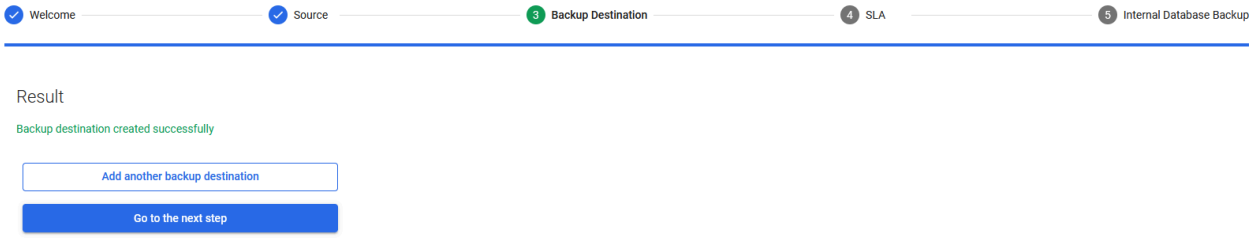

Add SLA

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

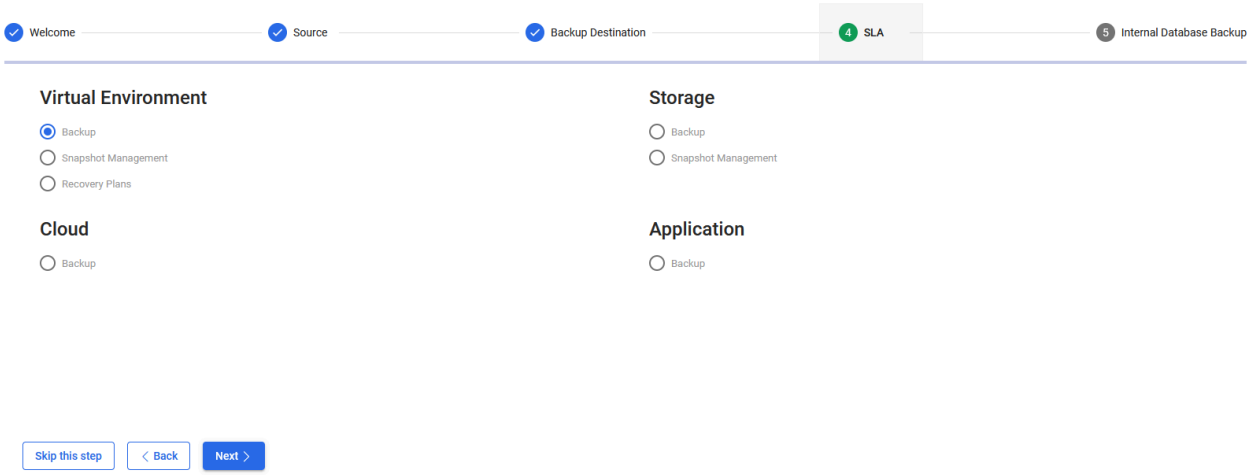

Add policy

- 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

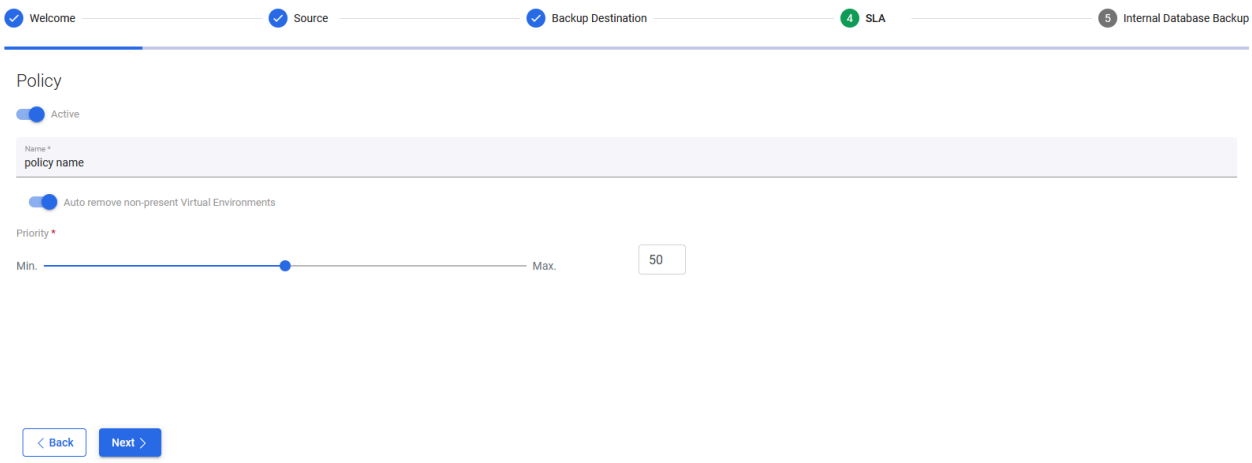

- Choose if you want to use auto assign mode based on tags and regular expressions (matched against the VM name, `.*` matches all characters 0 or more times)

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

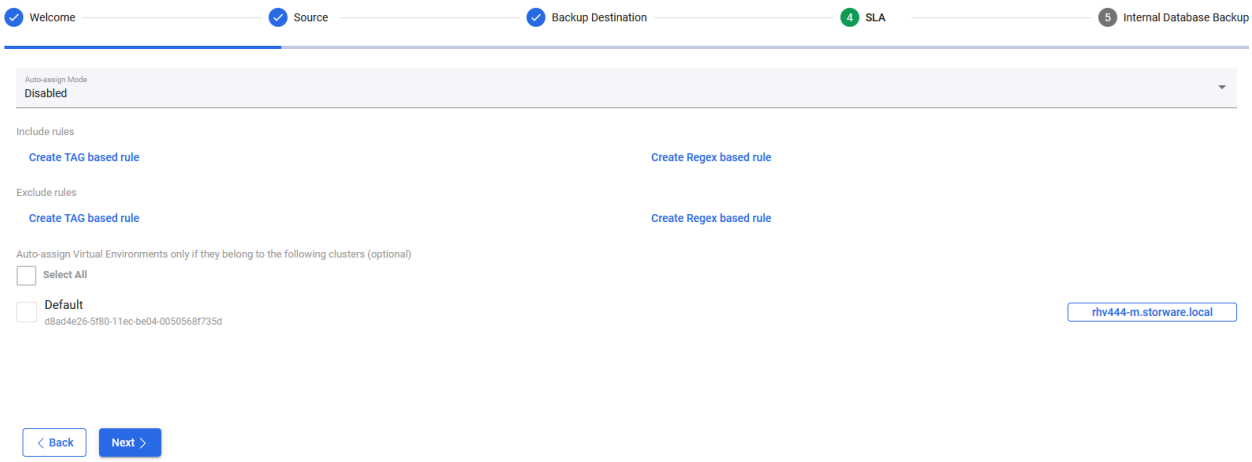

- Manually add the VMs if you do not want to use the auto-assignment mode

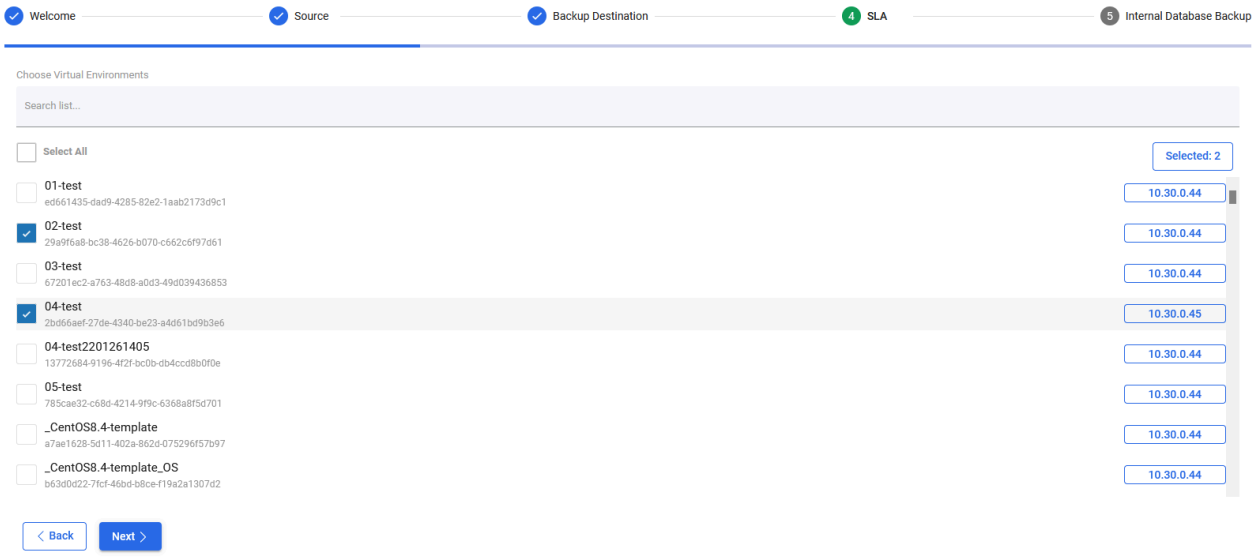

- Choose a backup destination target for this policy

> **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.

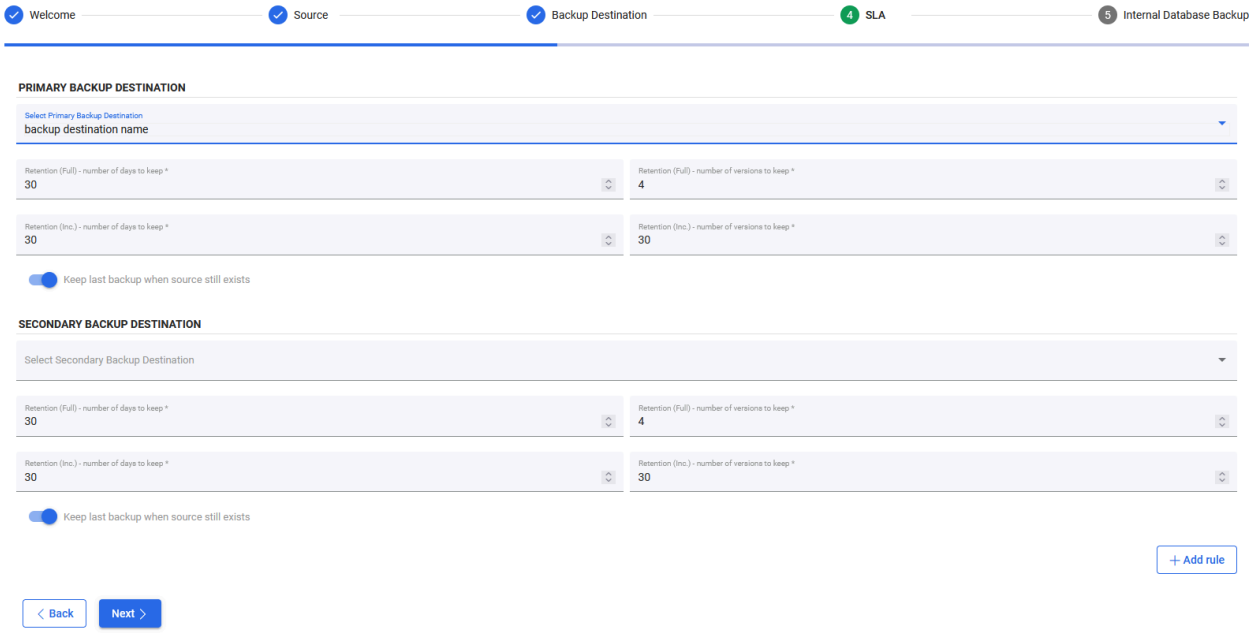

Configure the following thresholds:

- 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

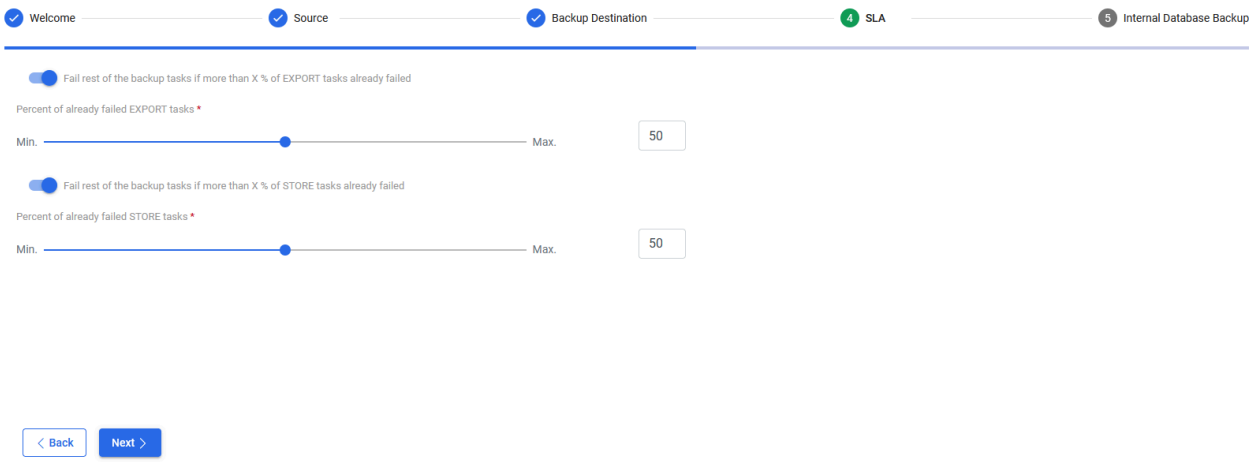

Add schedule

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

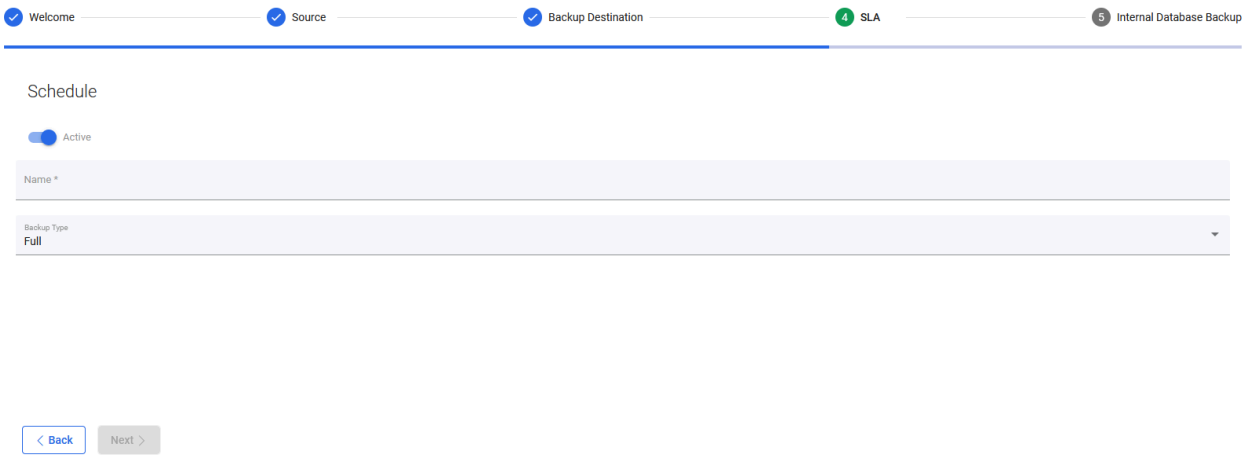

- Define the execution type:
  - time
  - interval
- Define the start window length
- Choose the time of day for backup

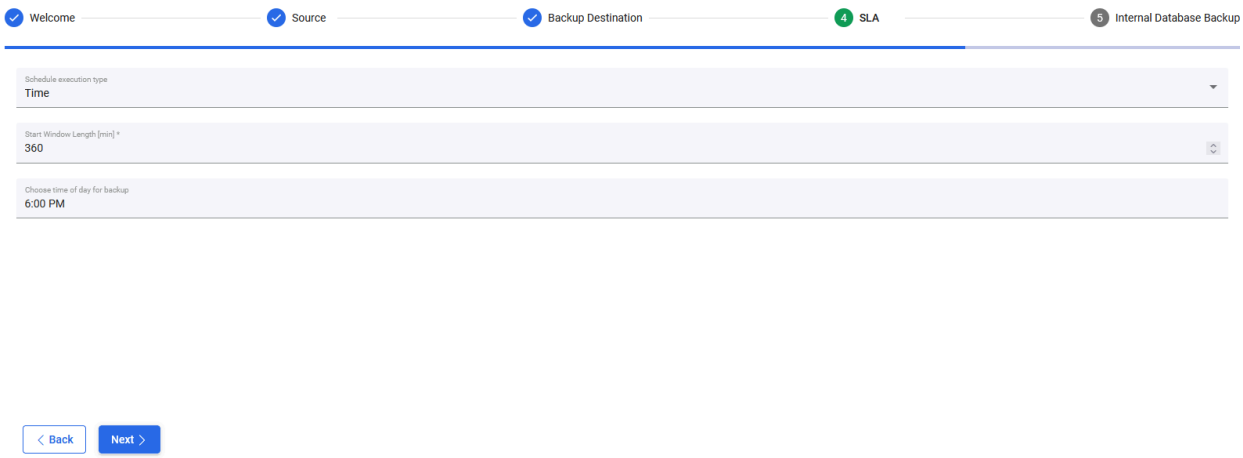

- Choose
  - days (required).
  - day of week occurrence (optional)
  - selected months (optional)

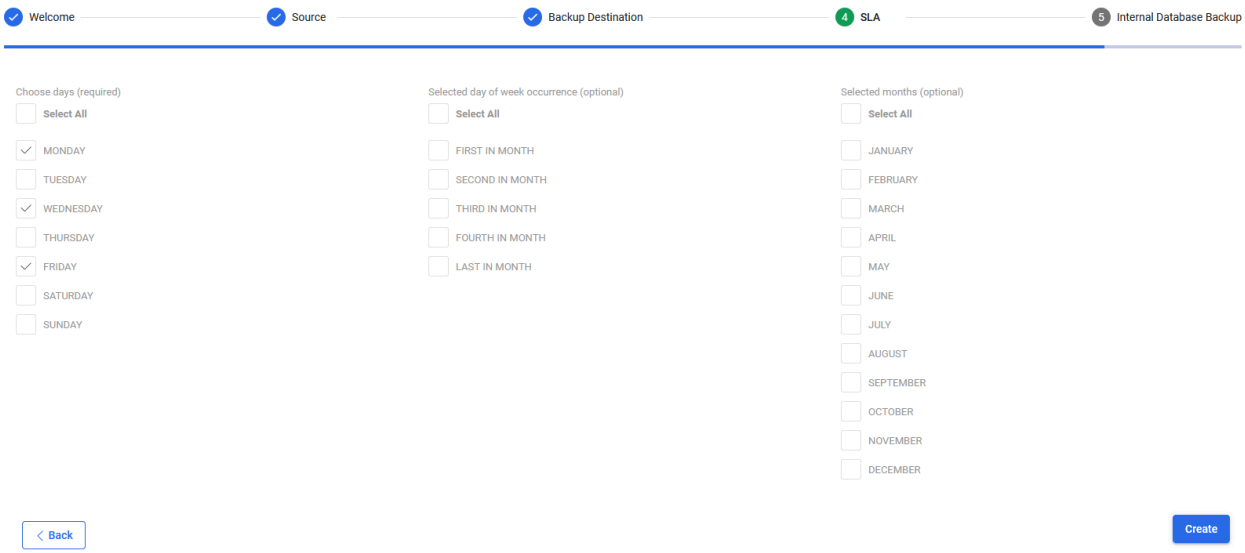

- Finish this step by going to the next section or adding another SLA.

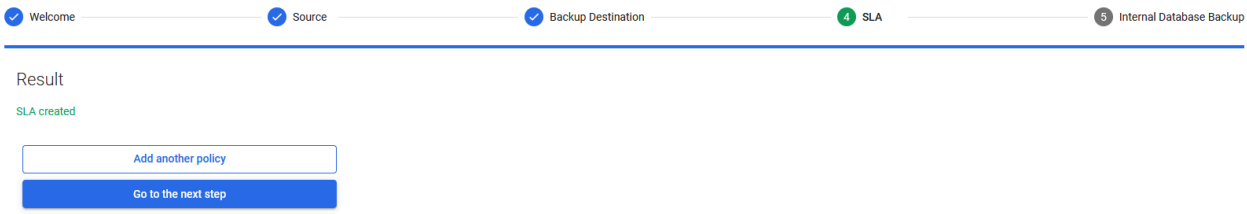

Add internal DB backup

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

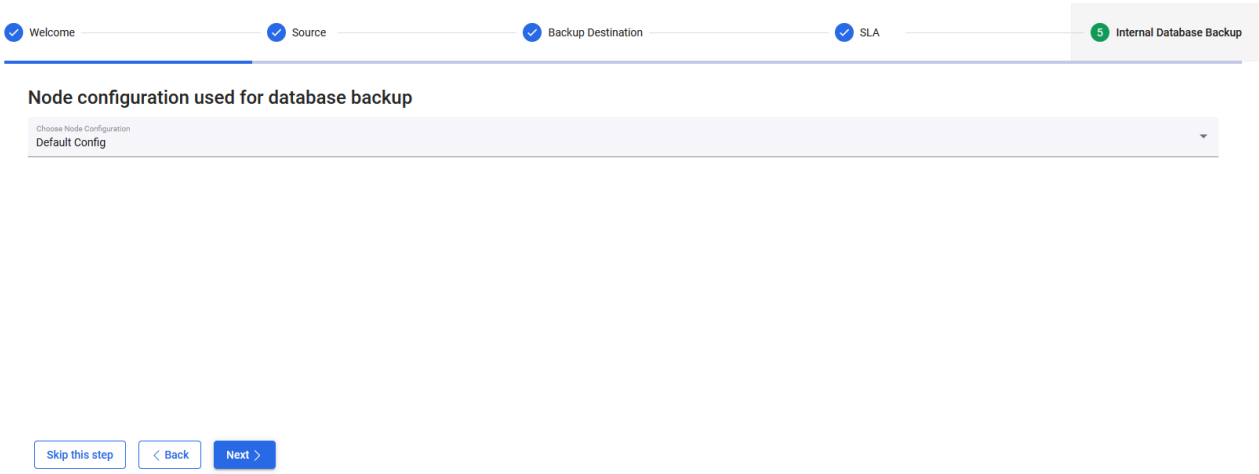

- Choose the backup destination for the DB backup

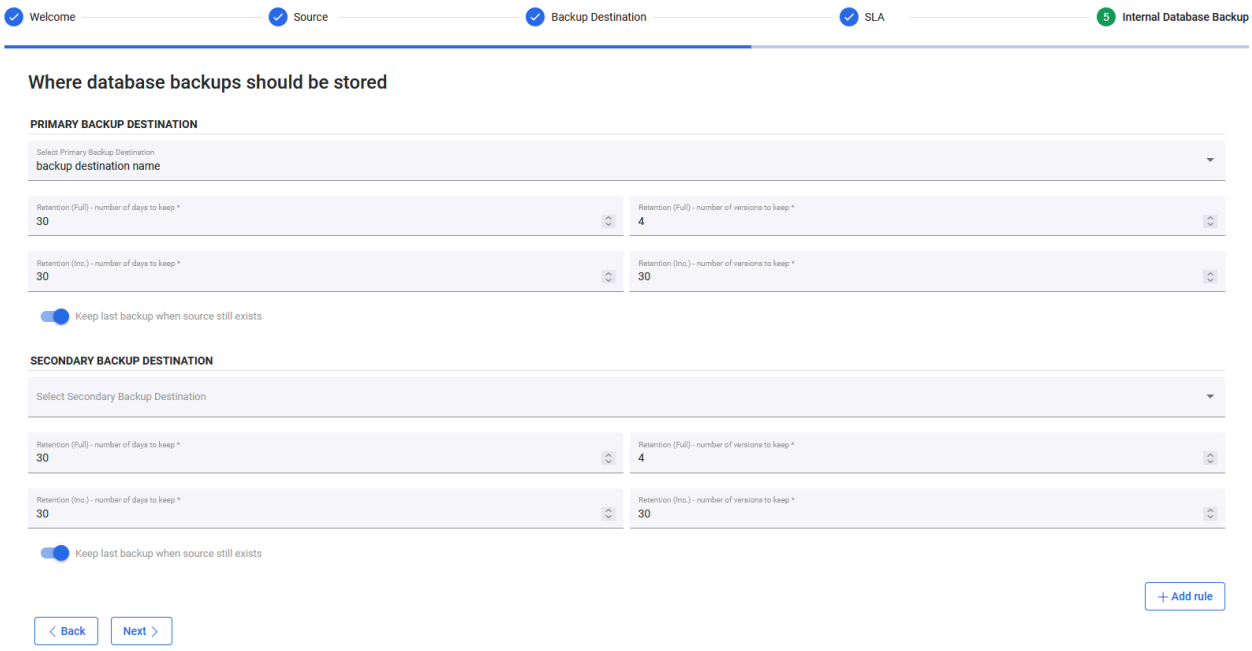

Choose when the DB backup should be run (daily basis)

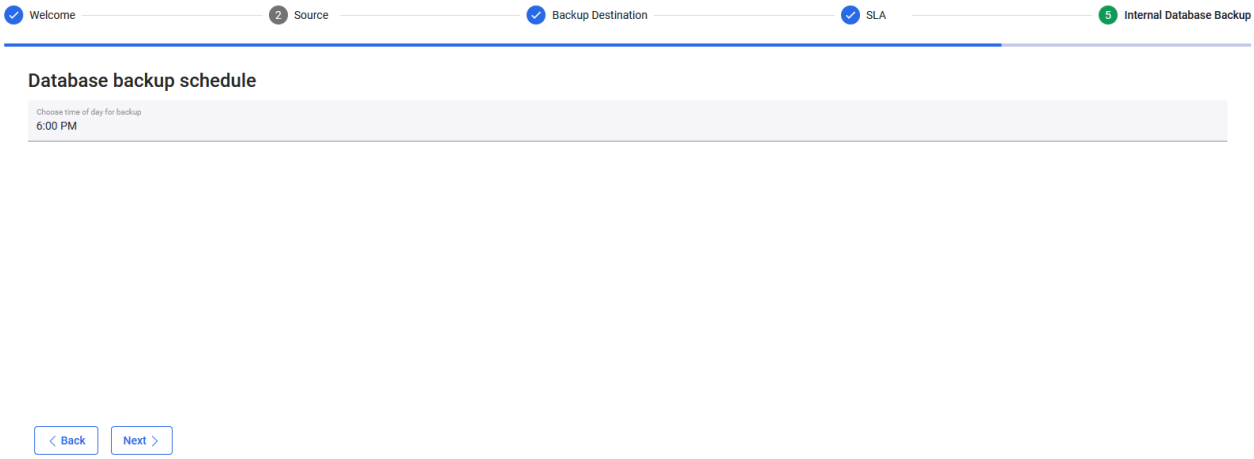

- Finalize the configuration and/or run the backup manually (on demand)

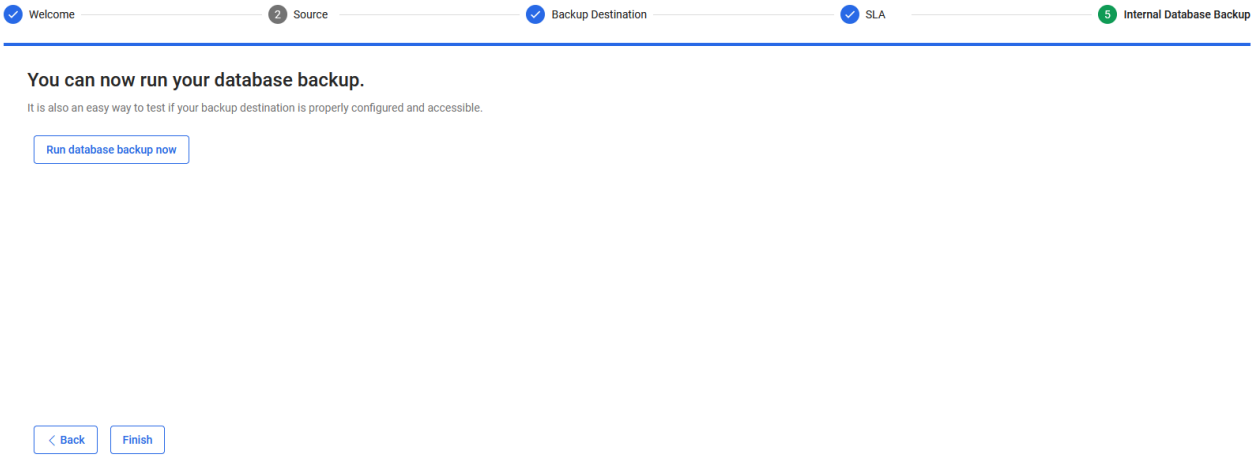

- you are ready to go!

Well done!

Now let's launch some backup jobs.

Go back to dashboard

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

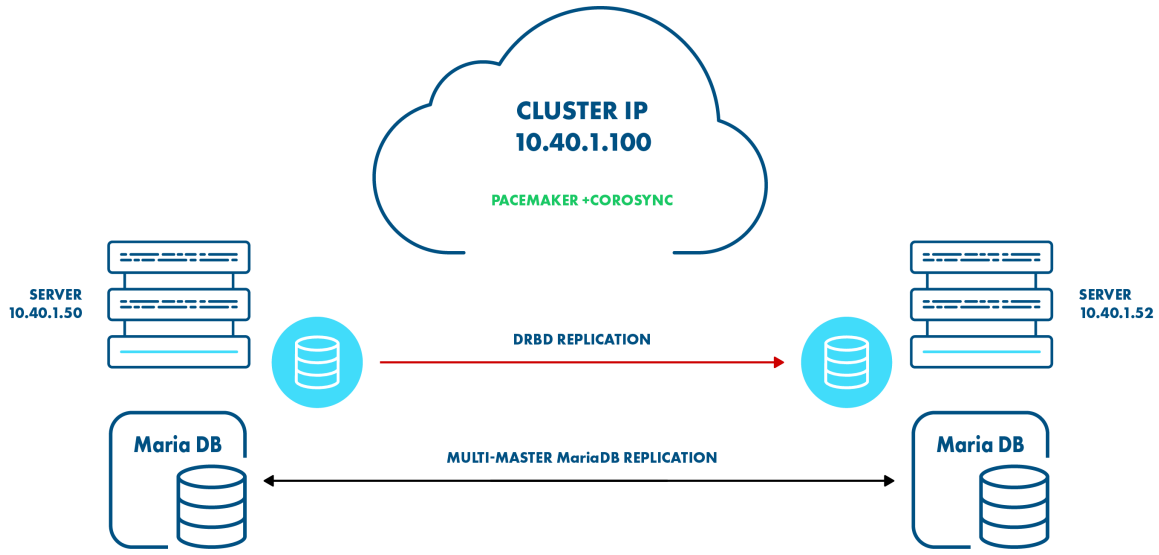

HA cluster setup

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/config
# 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
vprotect2
```

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 -L
```

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:IPaddr2 ip=10.40.1.100 cidr_netmask=22 op monitor
interval=30s
```

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:

The next step is to

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

```
[root@vprotect1 ~]# pcs resource create ping ocf:pacemaker:ping
dampen=5s multiplier=1000 host_list=10.40.0.1 clone
[root@vprotect1 ~]# pcs constraint location Failover_IP rule score=-
INFINITY pingd lt 1 or not_defined pingd
```

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-node op monitor timeout=300s on-fail="stop" --group
vProtect-group
[root@vprotect1 ~]# pcs resource create "vProtect-server"
service:vprotect-server op start on-fail="stop" timeout="300s" op stop
timeout="300s" on-fail="stop" op monitor timeout="300s" on-fail="stop"
--group vProtect-group
```

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

- **Resource colocation**

```
[root@vprotect1 ~]# pcs constraint colocation add Failover_IP with
vProtect-group
```

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=INFINITY
[root@vprotect1 ~]# pcs constraint location vProtect-group prefers
vprotect1=INFINITY
```

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.rpm
Retrieving https://www.elrepo.org/elrepo-release-7.0-
4.el7.elrepo.noarch.rpm
Preparing...                          ################################# 
[100%]
Updating / installing...
   1:elrepo-release-7.0-4.el7.elrepo  ################################# 
[100%]
```

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
kmod-drbd84.x86_64 0:8.4.11-1.1.el7_6.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
-R vprotect:vprotect /vprotect_data
```

- **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 -
-vdoLogicalSize=400G --compression=enabled --deduplication=enabled
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=26214400 blks
    =                       sectsz=4096  attr=2, projid32bit=1
    =                       crc=1        finobt=0, sparse=0
data     =                       bsize=4096   blocks=104857600,
imaxpct=25
    =                       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/ &&
chown -R vprotect:vprotect /vprotect_data
```

- **Copy the VDO config to the second node**

```
[root@vprotect1 ~]# scp /etc/vdoconf.yml
root@vprotect2:/etc/vdoconf.yml
```

- **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 \
         drbd_resource=replicate op monitor interval=10s --group
fs_group

[root@vprotect1 ~]#  pcs -f drbd_cfg resource master replicateClone
replicate \
         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:heartbeat:vdo-vol volume=vdo_data --group fs_group
[root@vprotect1 ~]#  pcs -f drbd_cfg resource create fs_resource
ocf:heartbeat:Filesystem device=/dev/mapper/vdo_data
directory=/vprotect_data fstype=xfs  --group fs_group
[root@vprotect1 ~]#  pcs cluster cib-push drbd_cfg --config

[root@vprotect1 ~]#  pcs constraint colocation add vdo_resource with
replicateClone
[root@vprotect1 ~]#  pcs constraint order start vdo_resource then
fs_resource
[root@vprotect1 ~]#  pcs constraint order start replicateClone then
vdo_resource
[root@vprotect1 ~]#  pcs constraint colocation add vProtect-group with
fs_group
[root@vprotect1 ~]#  pcs constraint colocation add vdo_resource with
replicateClone INFINITY with-rsc-role=Master
[root@vprotect1 ~]#  pcs constraint order promote replicateClone then
start fs_group
```

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.

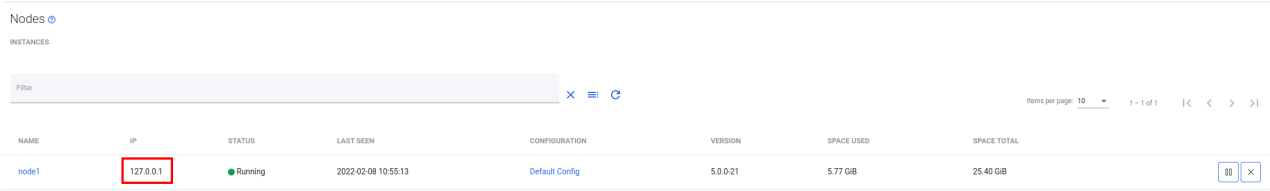

If the node's IP is different than 127.0.0.1, delete the node and re-register it using

```
[root@vprotect1 ~]# vprotect node -e <Node_Name> admin
http://127.0.0.1:8080/api
```

- 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_USER =
'replicator',MASTER_PASSWORD='R3pLic4ti0N',MASTER_LOG_FILE =
'vprotect1-bin.000007',MASTER_LOG_POS=46109;
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='10.40.1.52', master_user='replicator',
master_password='R3pLic4ti0N',MASTER_LOG_FILE = 'vprotect2-bin.000002',
master_log_pos=501051;
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

Enabling HTTPS connectivity for nodes

LVM setup on Data Protector for Cloud Workloads Node for disk attachment backup mode

Full versions of libvirt/qemu packages installation

SSH public key authentication

Enabling HTTP(S) Proxy for Data Protector for Cloud Workloads

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
sectors
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.key -out dp4cw.crt -days 365
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
incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name
or a DN.
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.key -out dp4cw.p12 -name 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/vprotect/keystore.jks -srckeystore dp4cw.p12 -srcstoretype
PKCS12 -alias dp4cw
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/cacerts -storepass changeit -noprompt
   -alias CA_ALIAS -file ca.crt
   ```

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|", "r|.*|" ]
            global_filter = [ "a|^/dev/sda|", "a|^/dev/sda1|",
   "a|^/dev/sda2|", "r|.*|" ]
   }
   ```

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
-y
```

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@HYPERVISOR
```

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@dkvm
Last failed login: Mon Jan 29 17:53:01 CET 2018 from 10.50.1.107 on
ssh:notty
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

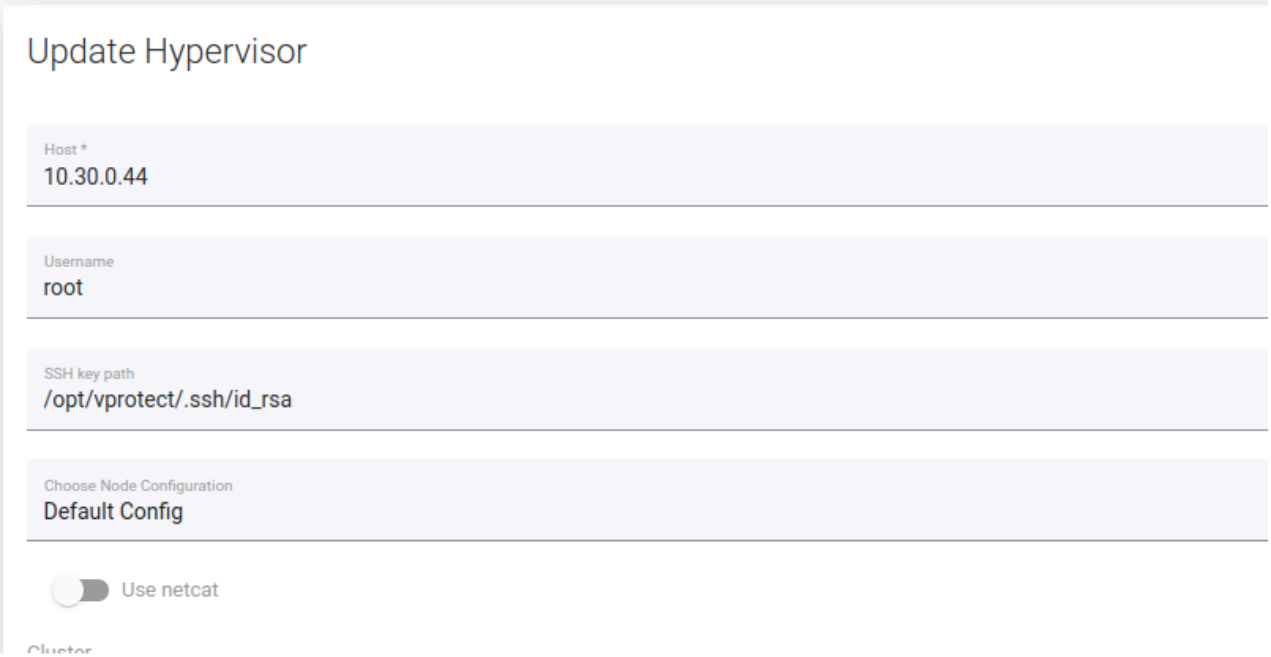

Enabling HTTP(S) Proxy for Data Protector for Cloud Workloads

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
- OpenStack
- OpenNebula
- Virtuozzo
- Citrix Hypervisor (XenServer)
- XCP-ng
- 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.

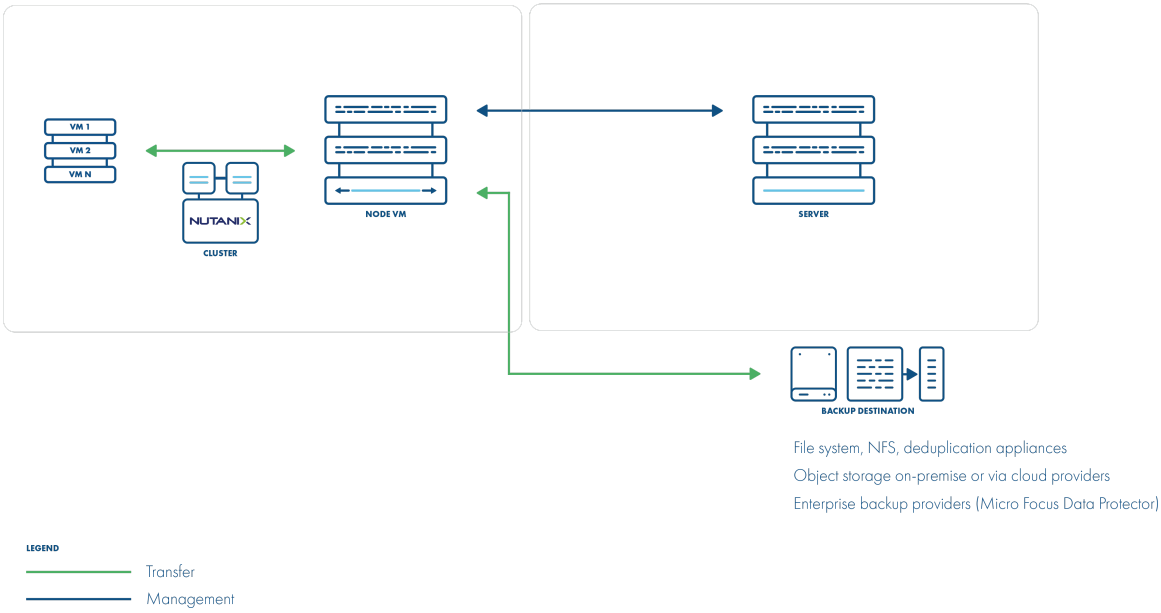

Backup Process

- 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)

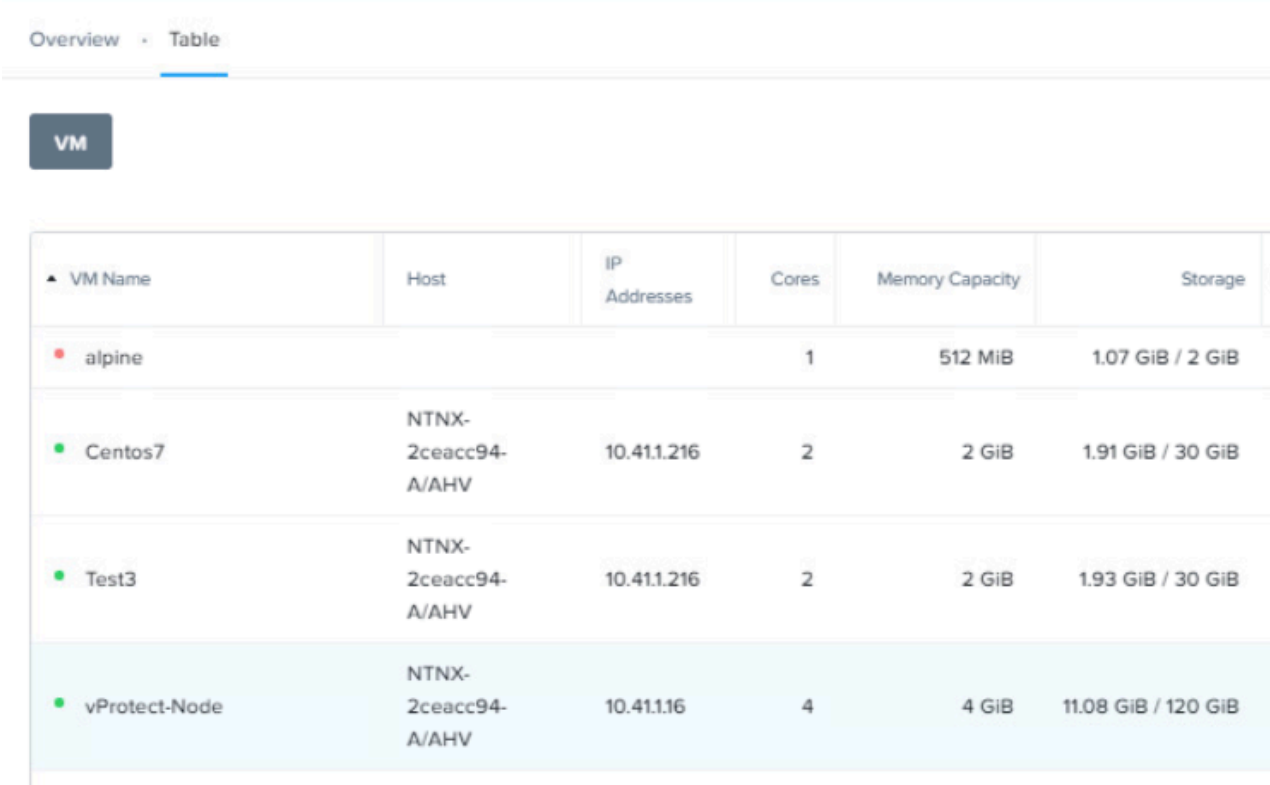

VM Name	Host	IP Addresses	Cores	Memory Capacity	Storage
alpine			1	512 MiB	1.07 GiB / 2 GiB
Centos7	NTNX-2ceacc94-A/AHV	10.41.1.216	2	2 GiB	1.91 GiB / 30 GiB
Test3	NTNX-2ceacc94-A/AHV	10.41.1.216	2	2 GiB	1.93 GiB / 30 GiB
vProtect-Node	NTNX-2ceacc94-A/AHV	10.41.1.16	4	4 GiB	11.08 GiB / 120 GiB

- 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)*

- Run inventory synchronization task, after that you should see all Nutanix hosts under the hypervisor tab

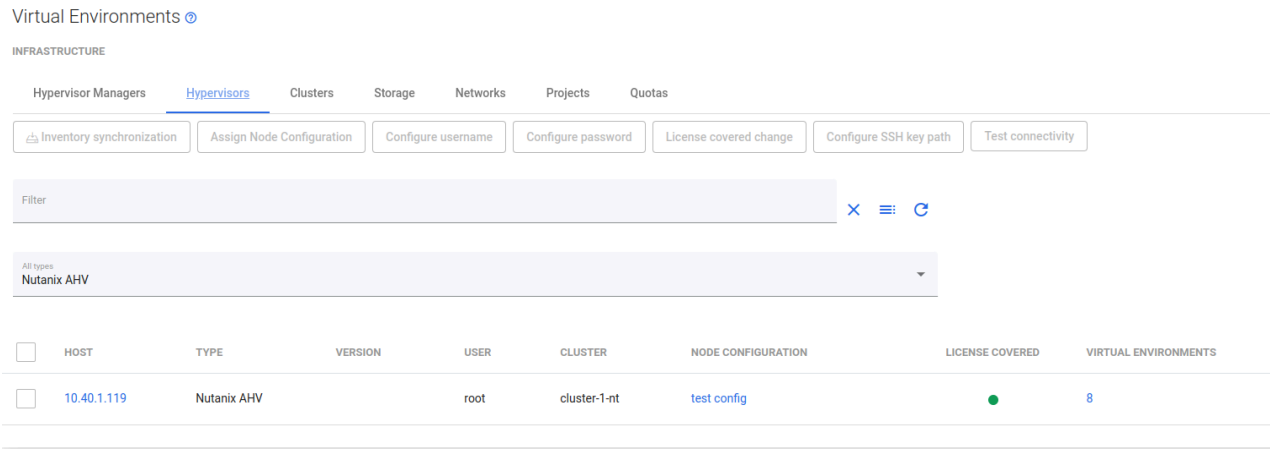

- As we describe above, we can back up Nutanix VMs thanks to the disk attachment backup strategy. As this is one of the most demanding methods, at this point we recommend that you perform a few easy tests to make sure that the backup you are going to perform is correct.
- Connect via SSH to the Proxy 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

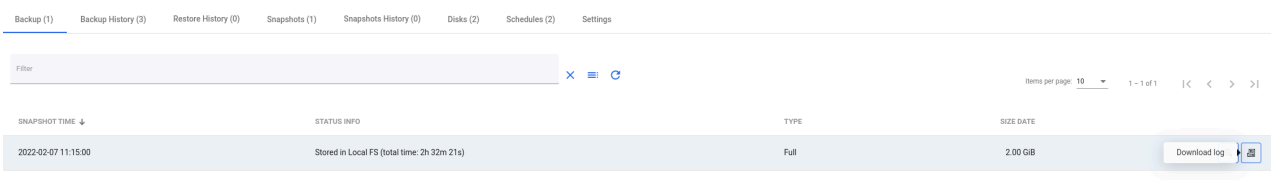

- As we can see in the logs, we do not need to correct the "offset" value. Data Protector for Cloud Workloads wants to start a backup from /dev/sdc, which is correct behavior because this disk device does not belong to Proxy VM.

```
[2021-04-08 14:51:40.959] INFO [Thread-47]
IProxyVmProvider.waitForDevice:38
[ffc65c30-8952-4ffa-b5d5-eefcfe01f333] Checking if device '/dev/sdc' is
present...

[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:45
[ffc65c30-8952-4ffa-b5d5-eefcfe01f333] Device '/dev/sdc' is present

[2021-04-08 14:51:55.991] INFO [Thread-47]
NutanixHypervisorManager.exportData:895
[ffc65c30-8952-4ffa-b5d5-eefcfe01f333] Data export of scsi.0
(917a15a2-5815-4d20-b693-6fb77ea59293)[20 GiB]: '/dev/sdc' ->
'/vprotect_data/vProtect-node__fb96db59/scsi.0.raw'...
```

- 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:

- **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-
openshift-virtualization-plugin:x.x.x.x
```

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.

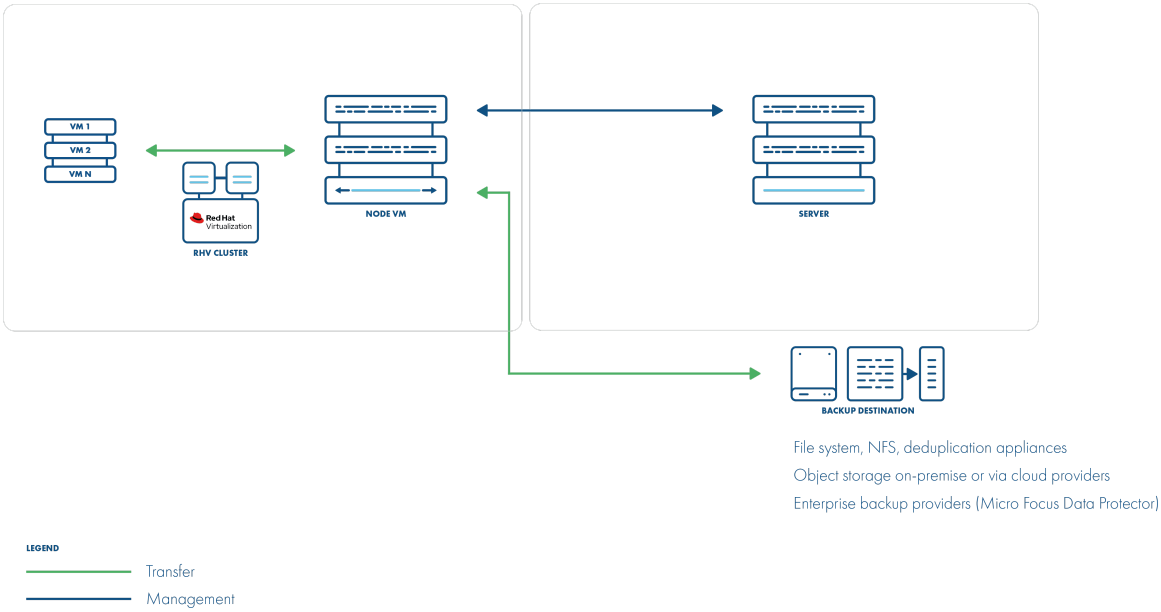

Backup Process

- 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.

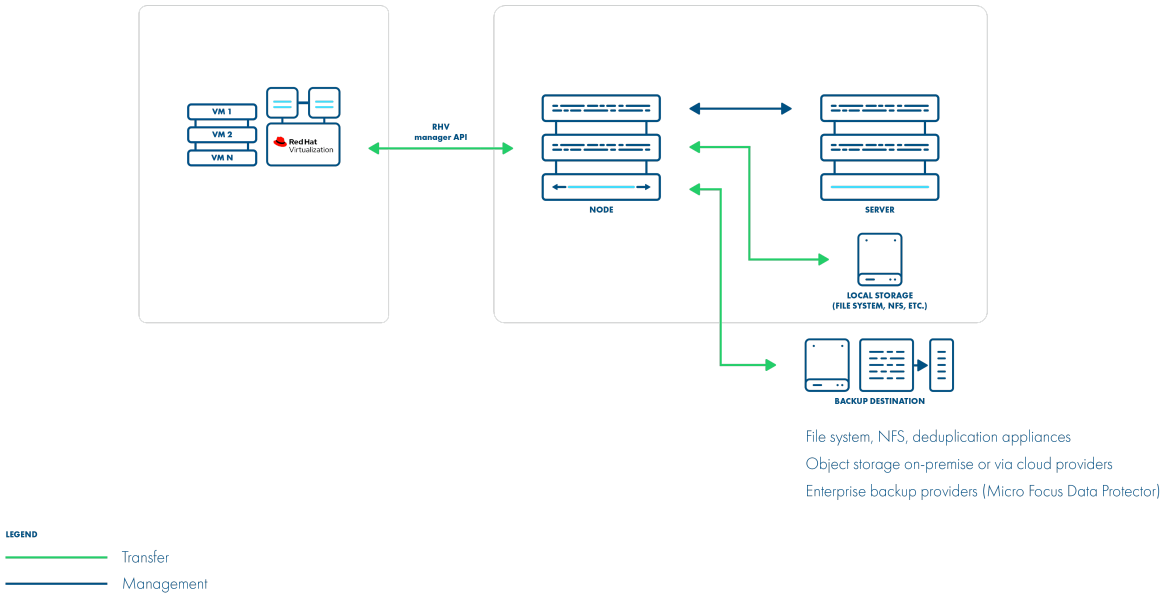

Backup Process

- 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.

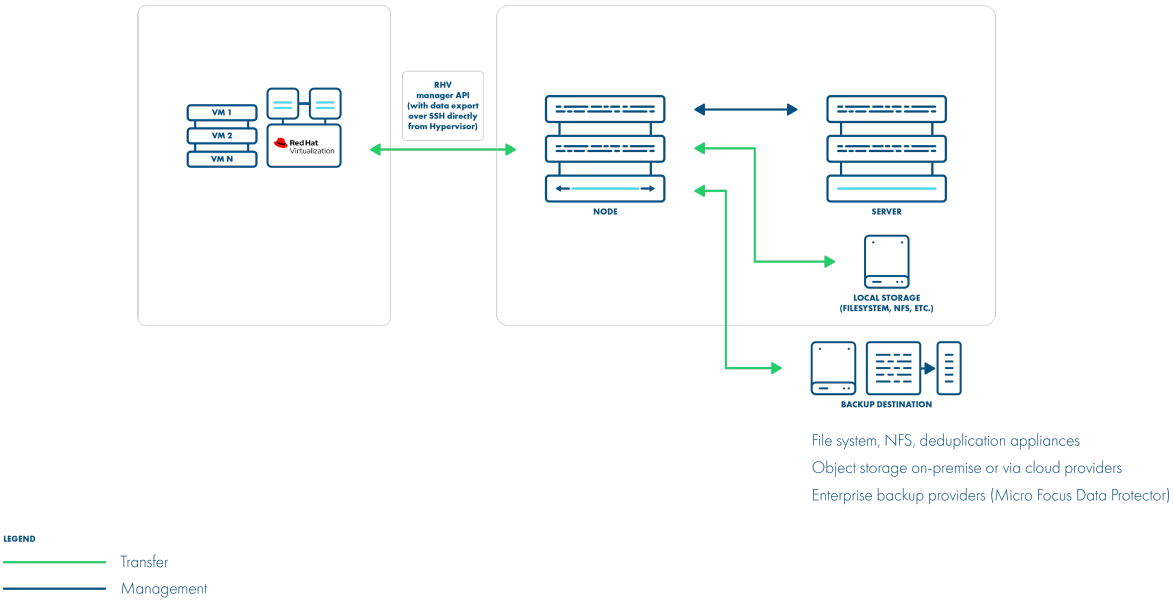

Backup Process

- 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.

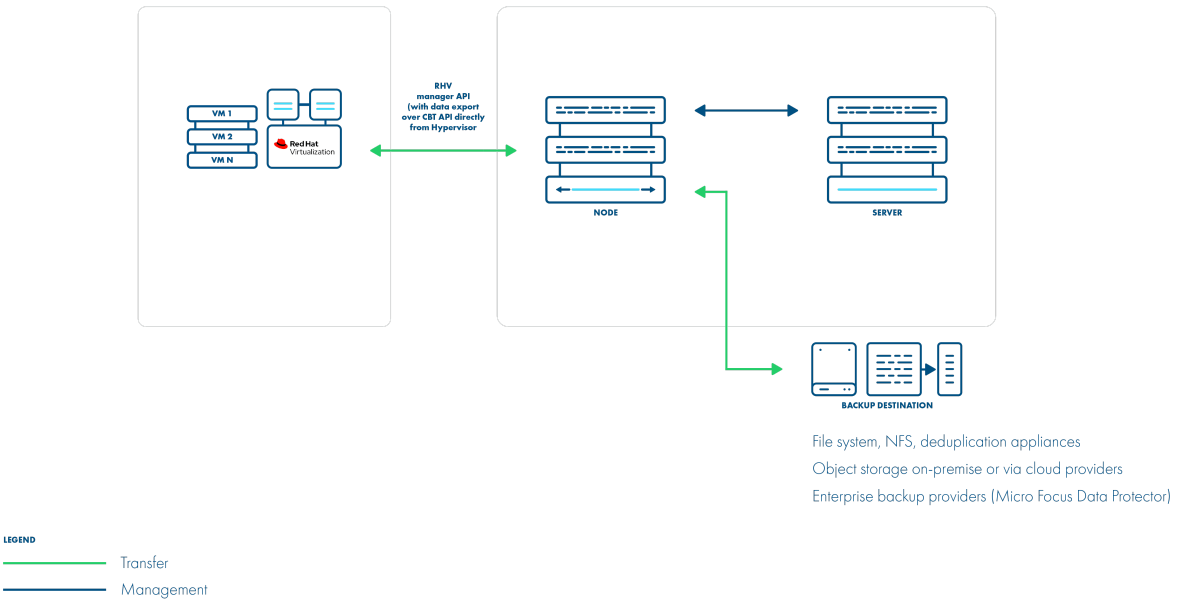

This strategy supports incremental backups.

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.

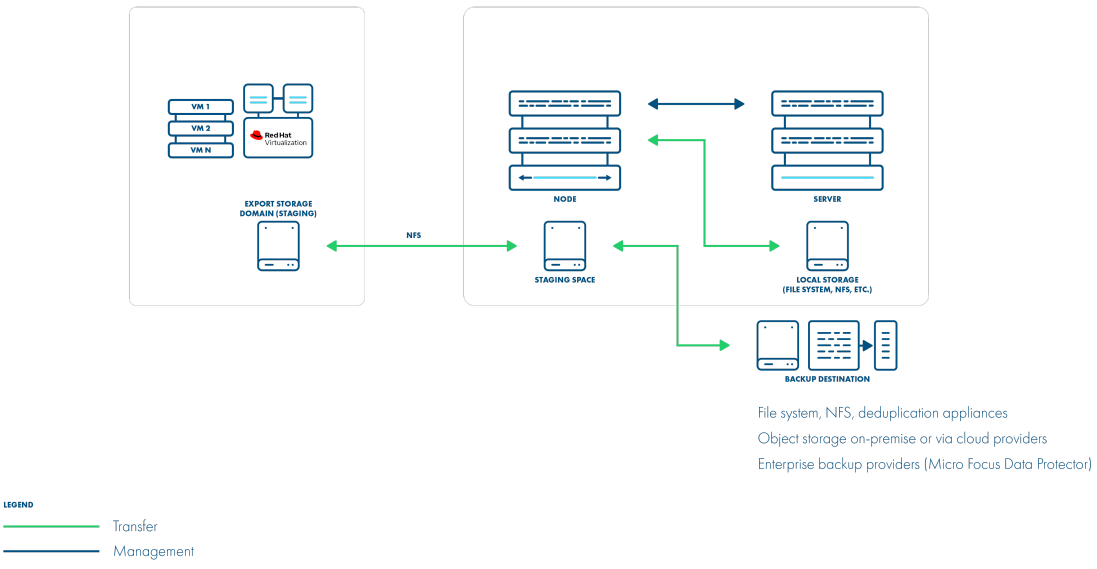

Backup Process

- 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,anonuid=993,anongi
d=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 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**.

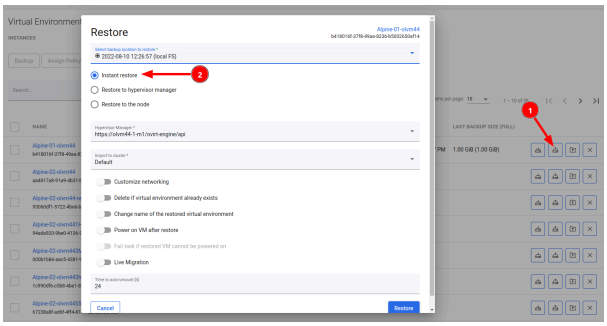

Live migration

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.

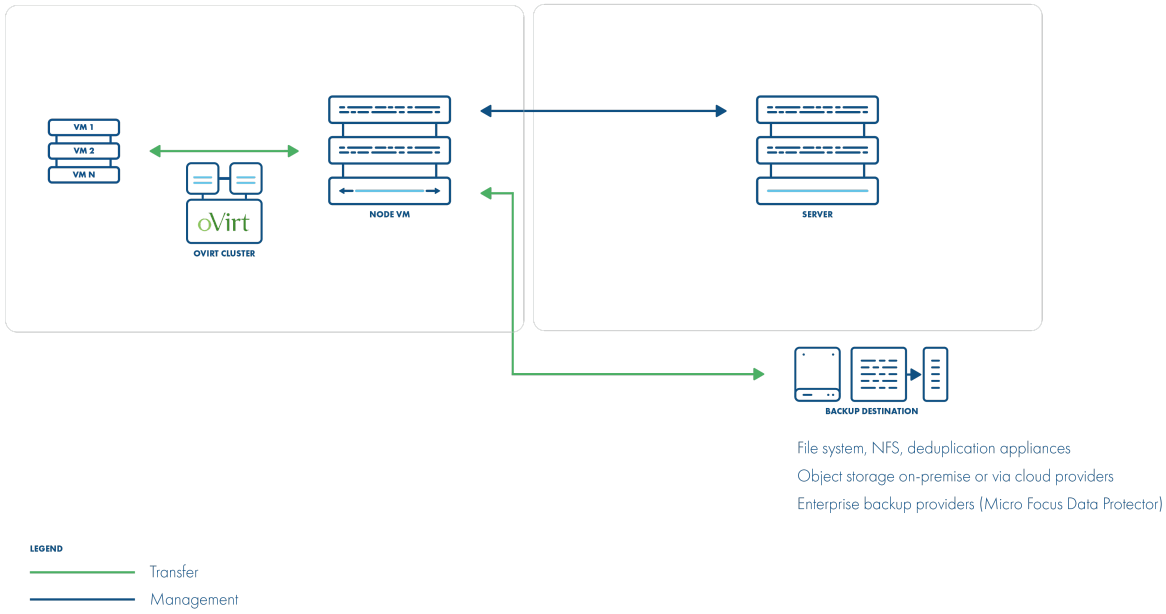

Backup Process

- 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.

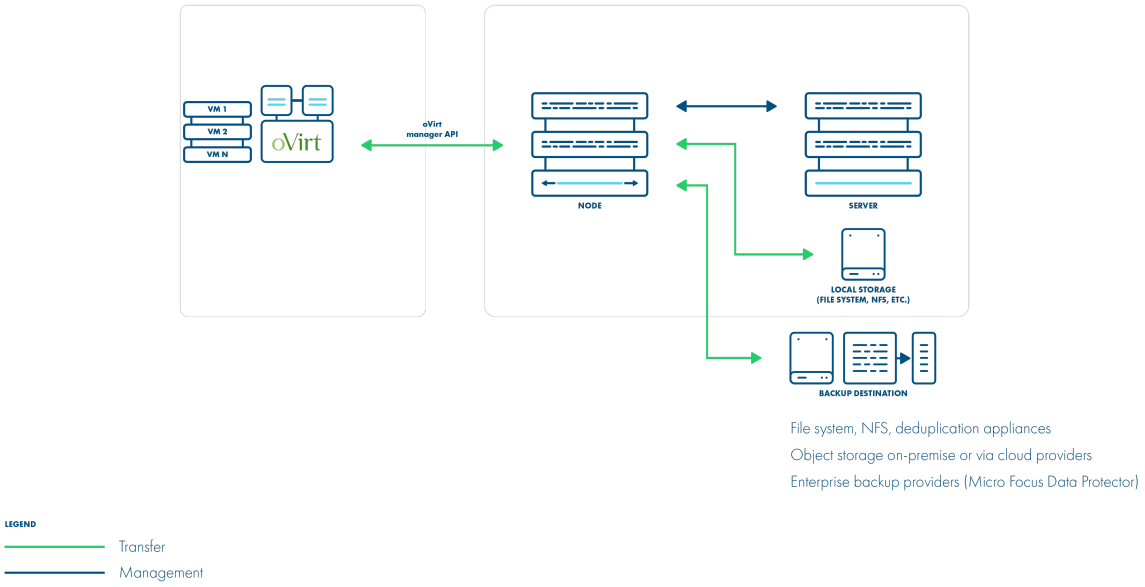

Backup Process

- 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.

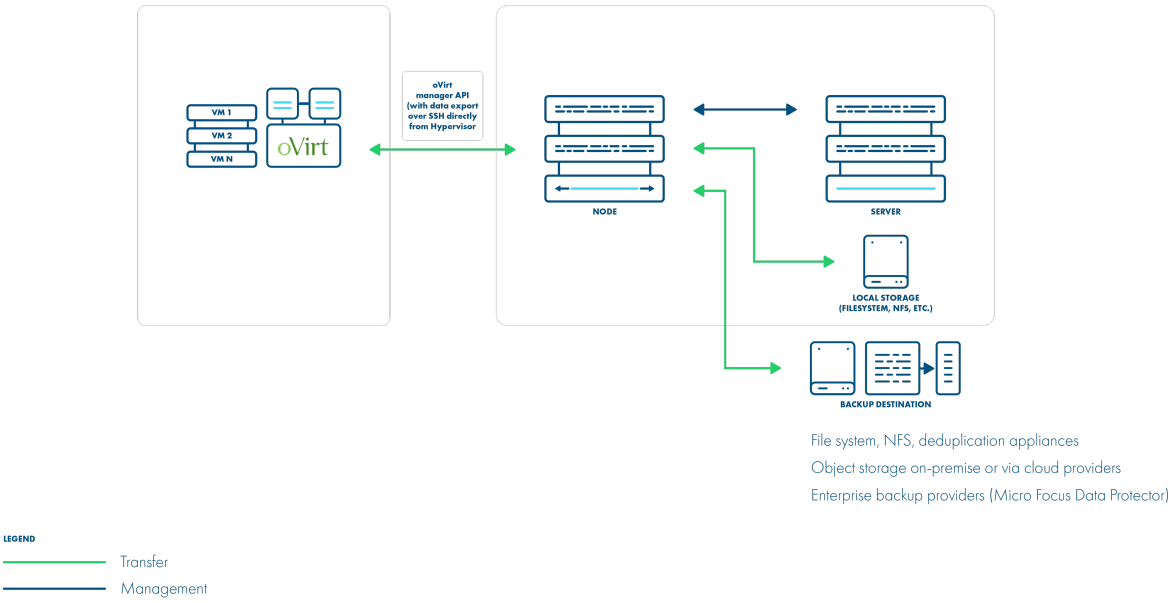

Backup Process

- 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.

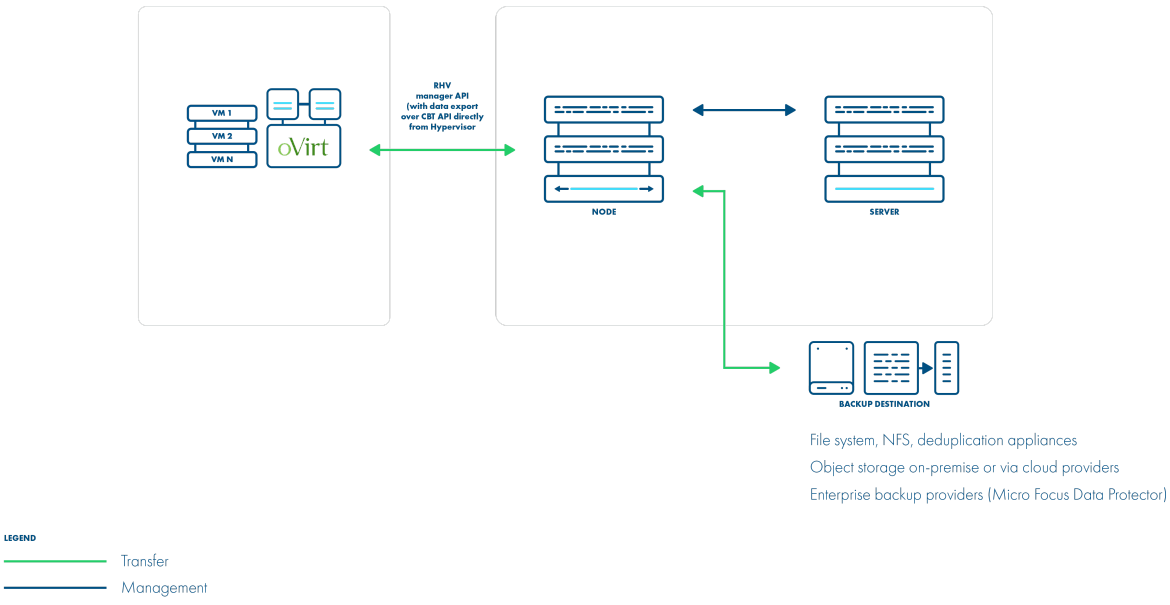

This strategy supports incremental backups.

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.

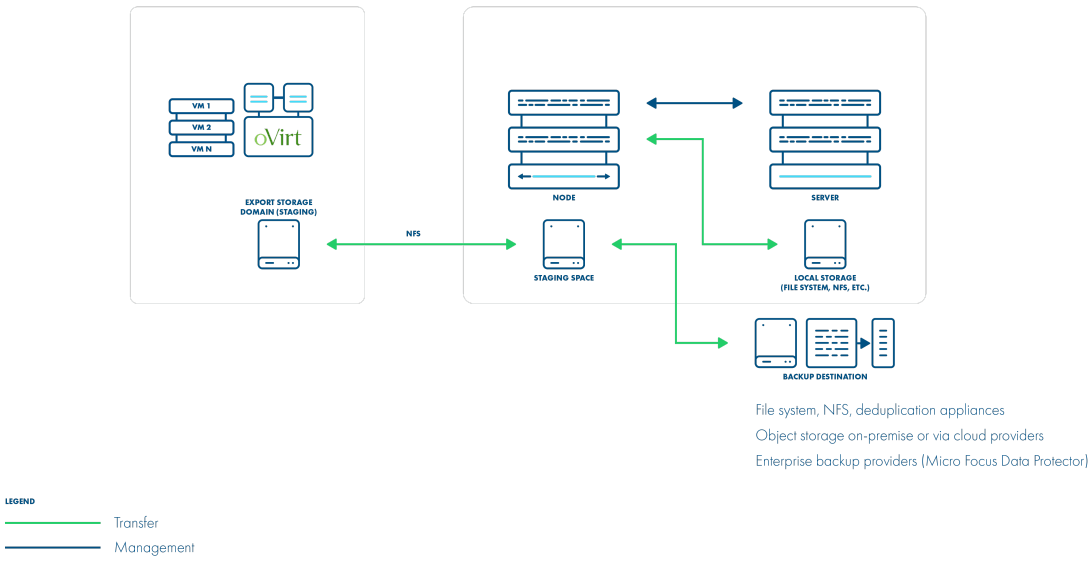

Backup Process

- 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,anonuid=993,anongi
d=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 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**.

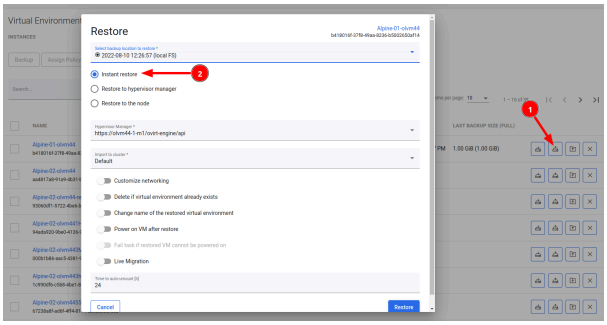

Live migration

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.

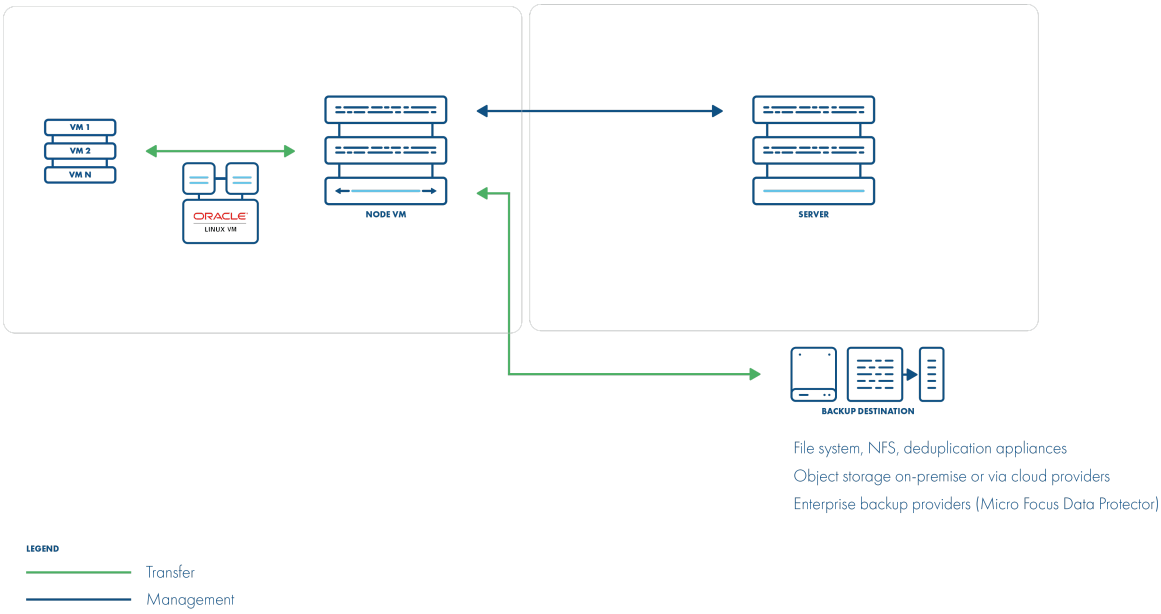

Backup Process

- 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.

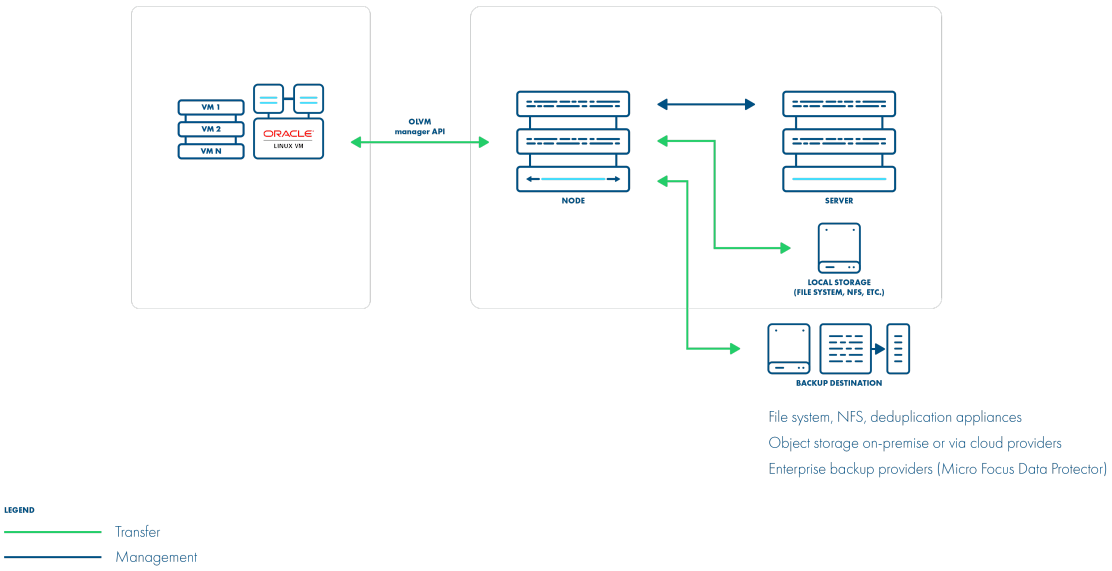

Backup Process

- 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.

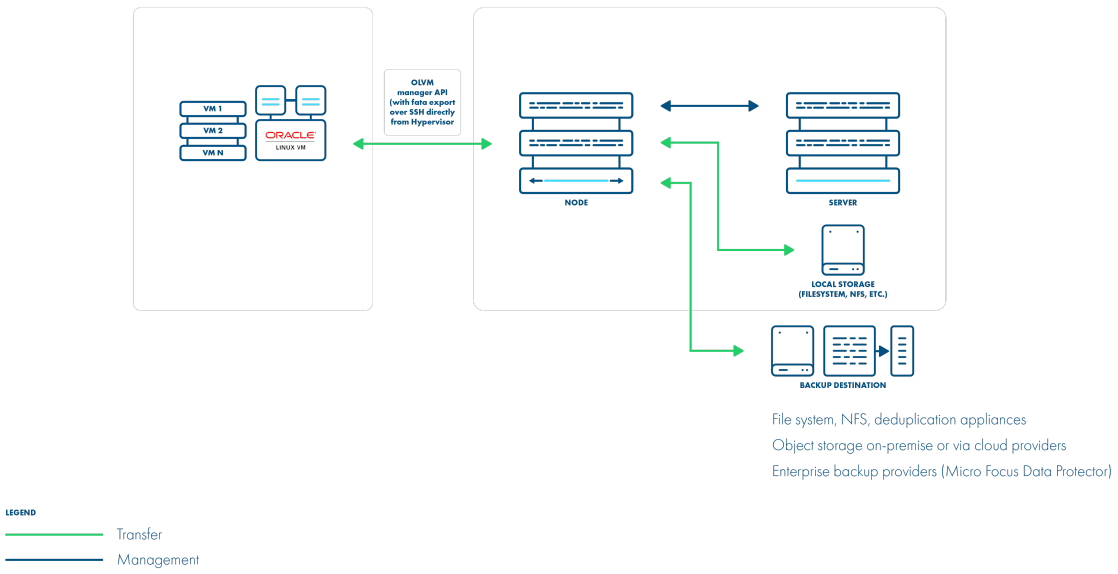

Backup Process

- 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**.

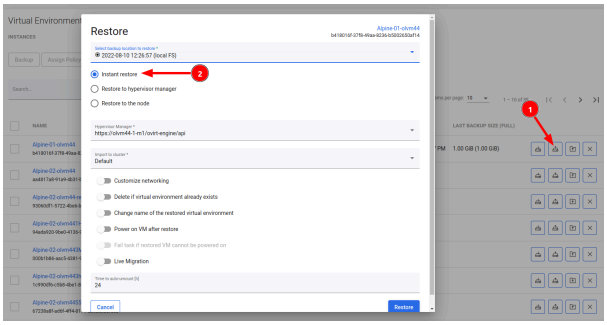

Live migration

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.

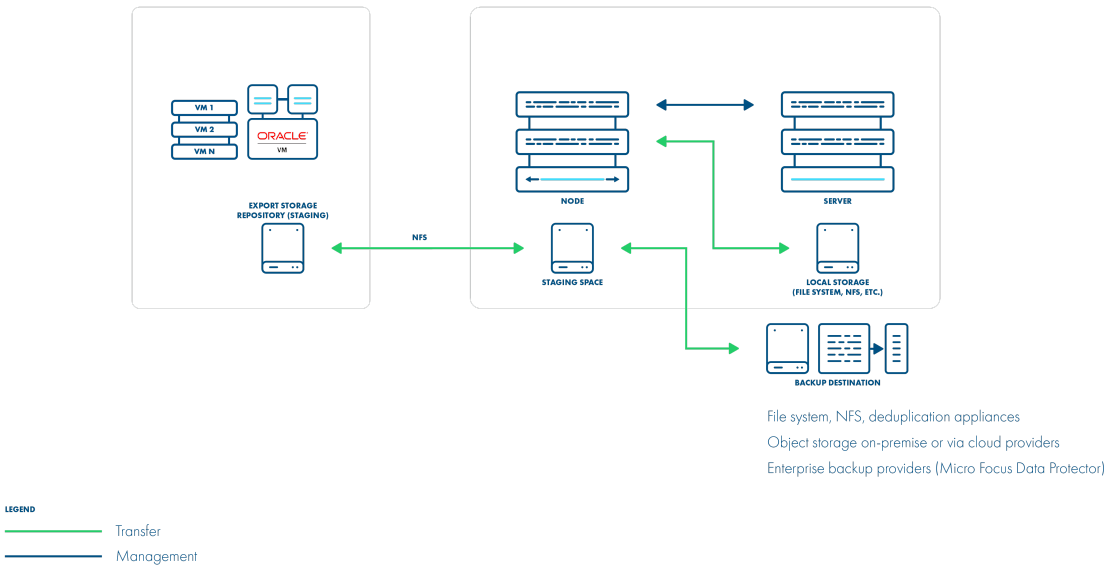

Backup Process

- 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.

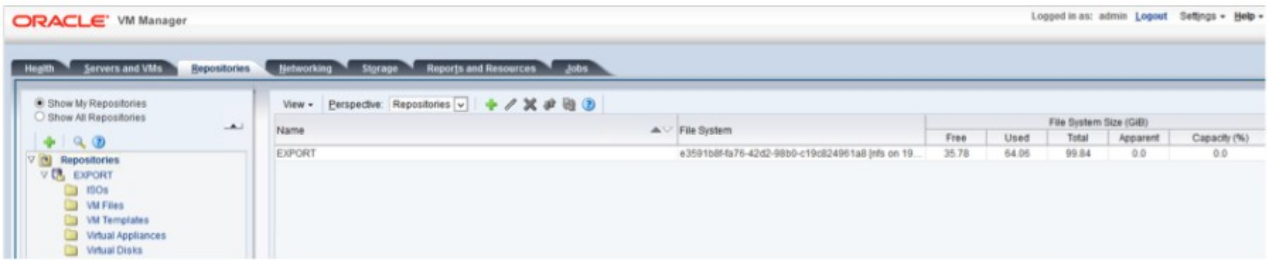

Example - How to configure OVM protection with PowerProtect DD

- Create a DDBoost device

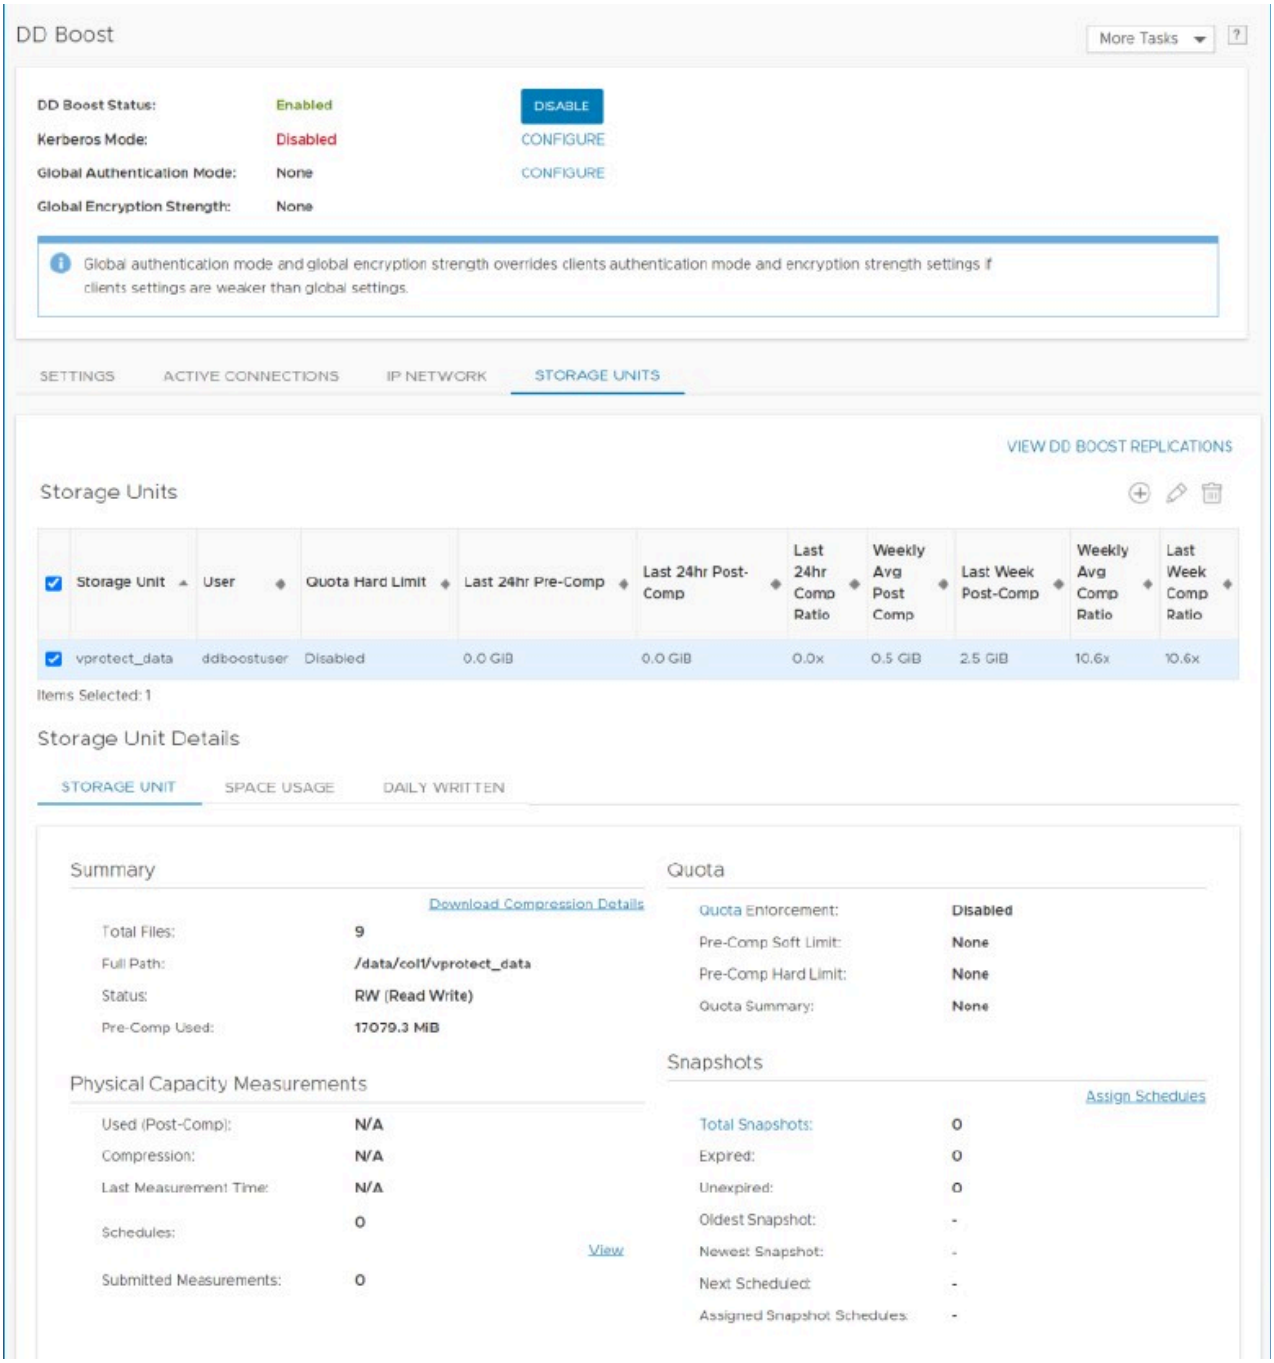

- Create NFS share

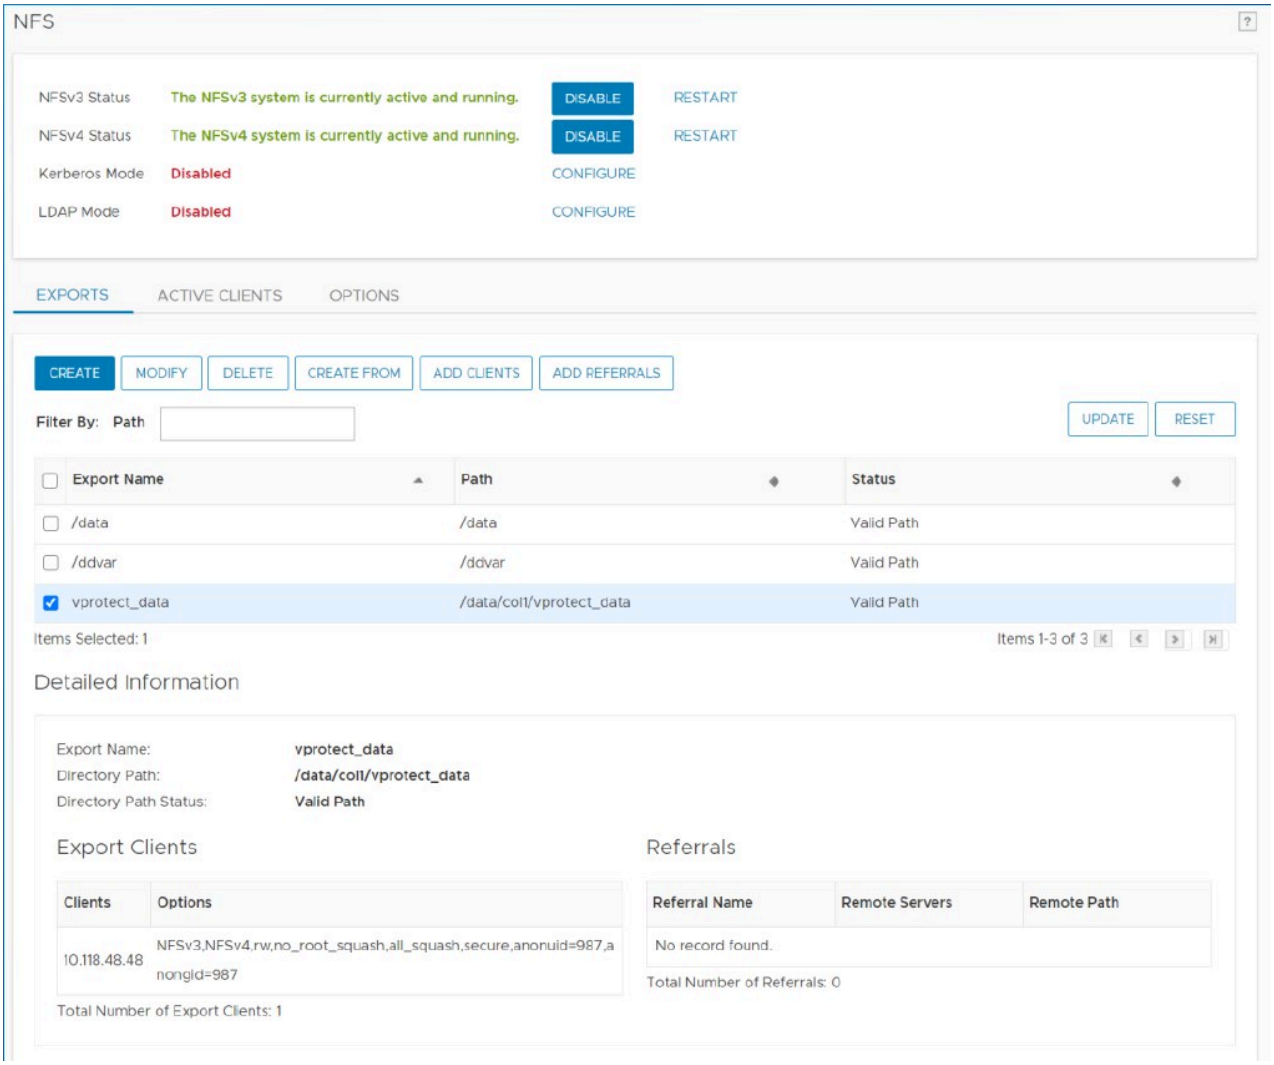

- Mount the DDBoost device on Data Protector for Cloud Workloads Node

```
#
# /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

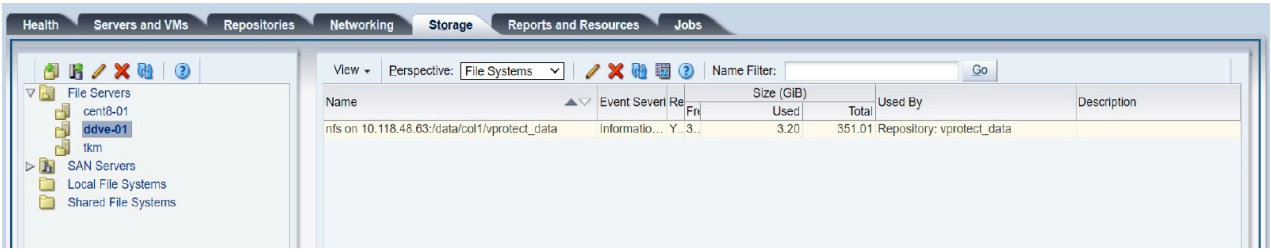

- Create a Repository using DD

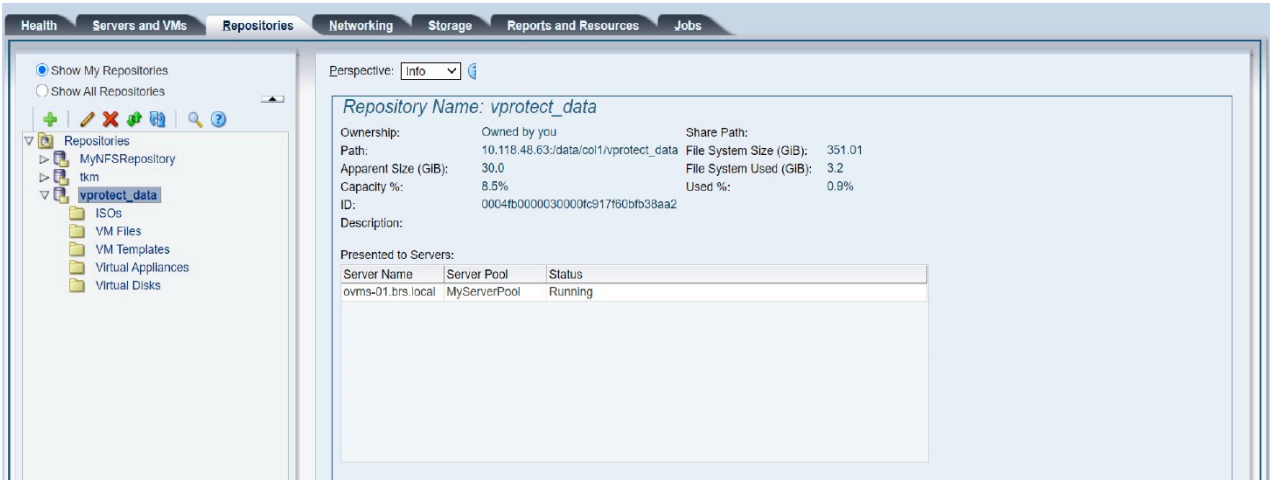

- Add the OVM Hypervisor Manager to Data Protector for Cloud Workloads

> ***Note:*** *You can get the "Storage Repository ID" from the "OVM repositories" menu shown in the previous step*

Add New Hypervisor Manager

Proxmox VE

Proxmox VE

SSH Transfer

SSH Transfer strategy:

- supports Proxmox 5.0+
- supports only QCOW2 disk images
- supports incremental backups
- supports over iSCSI

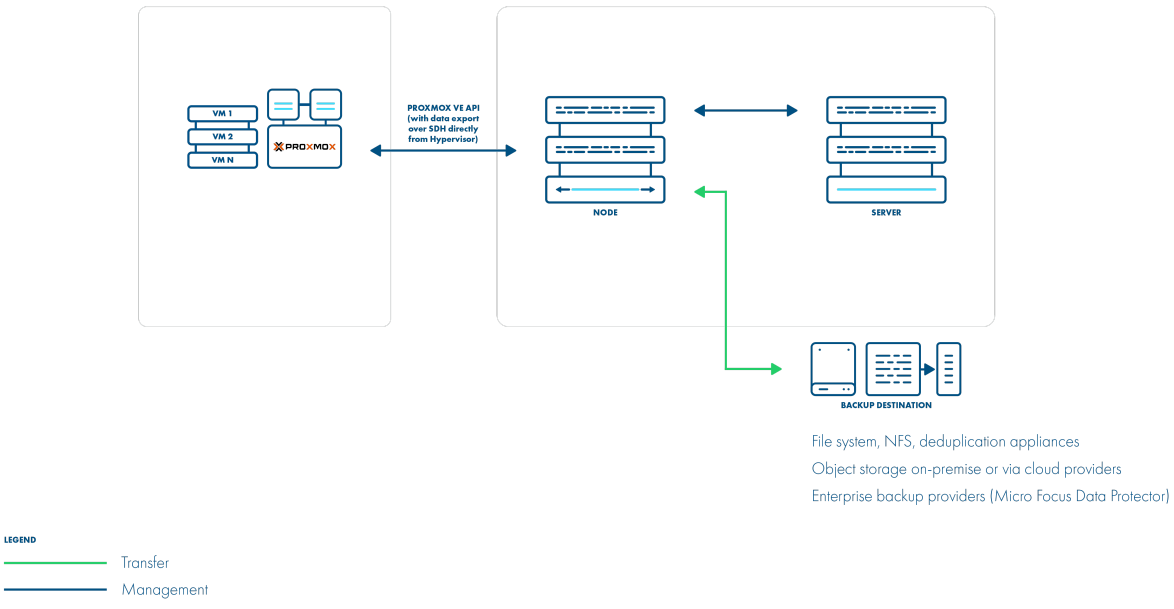

Backup is done by transferring QCOW2 disk images over SSH directly from the hypervisor (optionally using Netcat for transfer). Metadata is backed up only in the full backup. This method supports incremental backups where the last snapshot is required for the next incremental backups. The resulting backup 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.

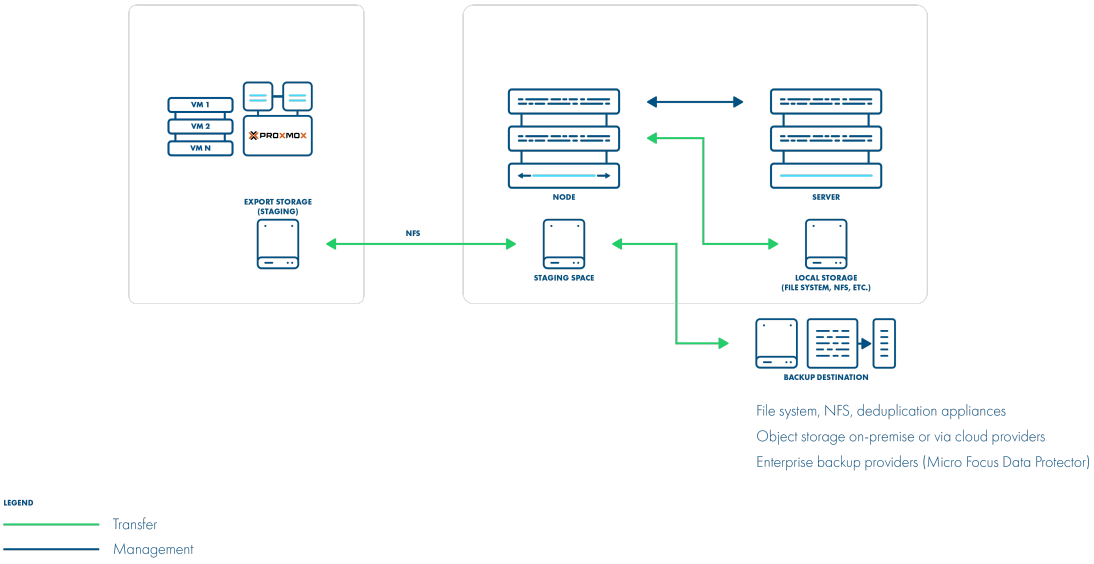

Backup Process

- 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**)

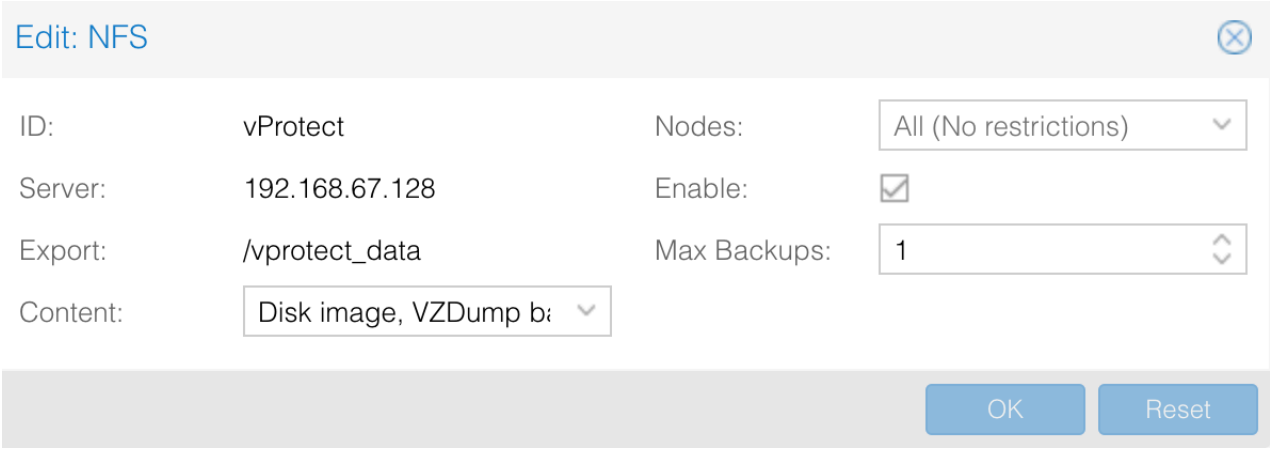

- Export share must be set to use the UID and GID of the `vprotect` user
- Example export configuration in `/etc/exports` to the selected hypervisor in the cluster:

  ```
  /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**

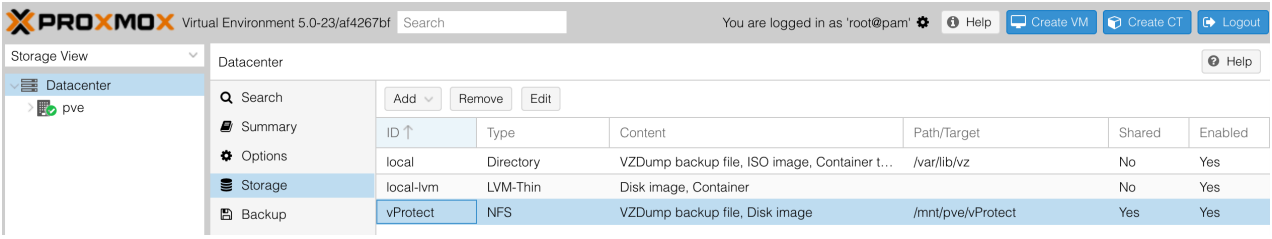

File-level restore support for VMA images

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
```

OpenStack

OpenStack

Data Protector for Cloud Workloads supports backup for OpenStack:

- Disk attachment through Cinder with changed block tracking (**preferred**):
  - supports all hypervisors and storages
  - supports incremental backup
  - proxy VM is required - used for the disk attachment process.
- Disk image transfer - for KVM hypervisors with VMs using QCOW2
- Volumes or Ceph-based storage:
  - supports incremental backup
  - disk images are transferred directly from API (no Proxy VM required)
- Disk attachment through Cinder:
  - supports all hypervisors and storages
  - no incremental backup
  - proxy VM is required - used for the disk attachment process.

Backup Strategies

Libvirt strategy

Data Protector for Cloud Workloads supports OpenStack environments that use KVM hypervisors and VMs running on QCOW2 or RAW files. Data Protector for Cloud Workloads communicates with OpenStack APIs such as Nova and Glance to collect metadata and for the import of the restored process. However, the actual backup is done over SSH directly from the hypervisor. The process is exactly the same as in Deployment in the KVM/Xen environment. Data Protector for Cloud Workloads Node can be installed anywhere - it just needs to have access to the

OpenStack APIs and hypervisor SSH via a network. Both full and incremental backups are supported.

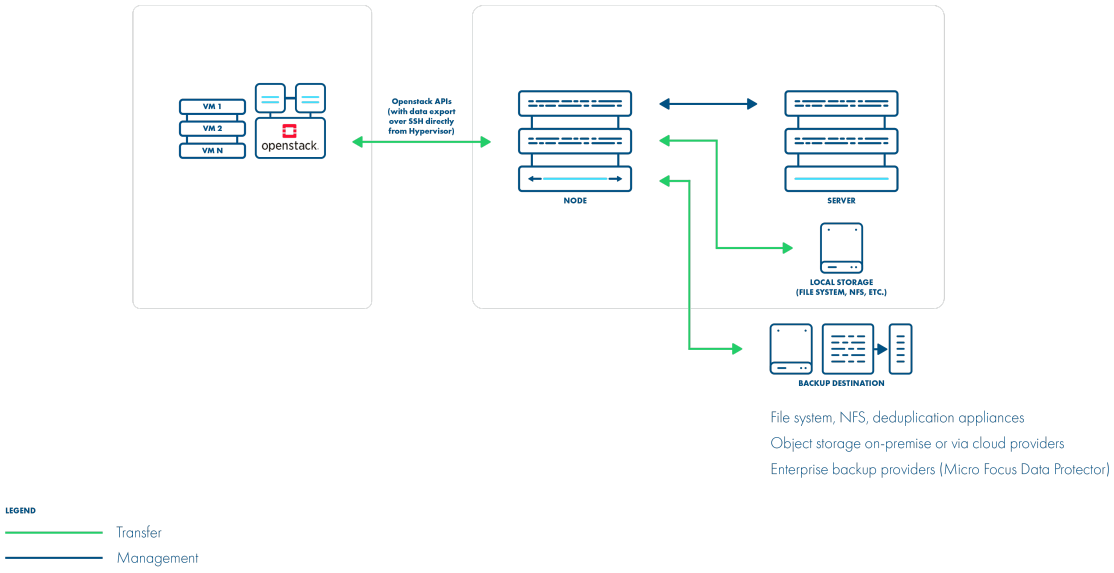

Backup Process

- direct access to the hypervisor over SSH
- crash-consistent snapshot taken directly using virsh (QCOW2/RAW file), rbd snapshot for Ceph (separate call for each storage backend)
- optional application consistency using pre/post snapshot command execution • QCOW2/RAW-file data exported over SSH (optionally with netcat)
- Ceph RBD data exported using rbd export or RBD-NBD when incremental is used
- metadata exported from OpenStack APIs (nova, glance, cinder)
- the last snapshot kept on the hypervisor for the next incremental backup (if at least one schedule assigned to the VM has backup type set to incremental)
- restore recreates files/volumes according to their backend (same transfer mechanism as used in backup) and then defines VM on the hypervisor

Disk attachment

Data Protector for Cloud Workloads also supports the disk-attachment method using cinder. This should allow you to use cinder-compatible storage and still allow Data Protector for Cloud Workloads to create backups. Incremental backup is supported in disk attachment changed block tracking (which has higher CPU overhead). Data Protector for Cloud Workloads needs to communicate OpenStack service's API to attach drives to the proxy VM with Data Protector for Cloud Workloads Node installed.

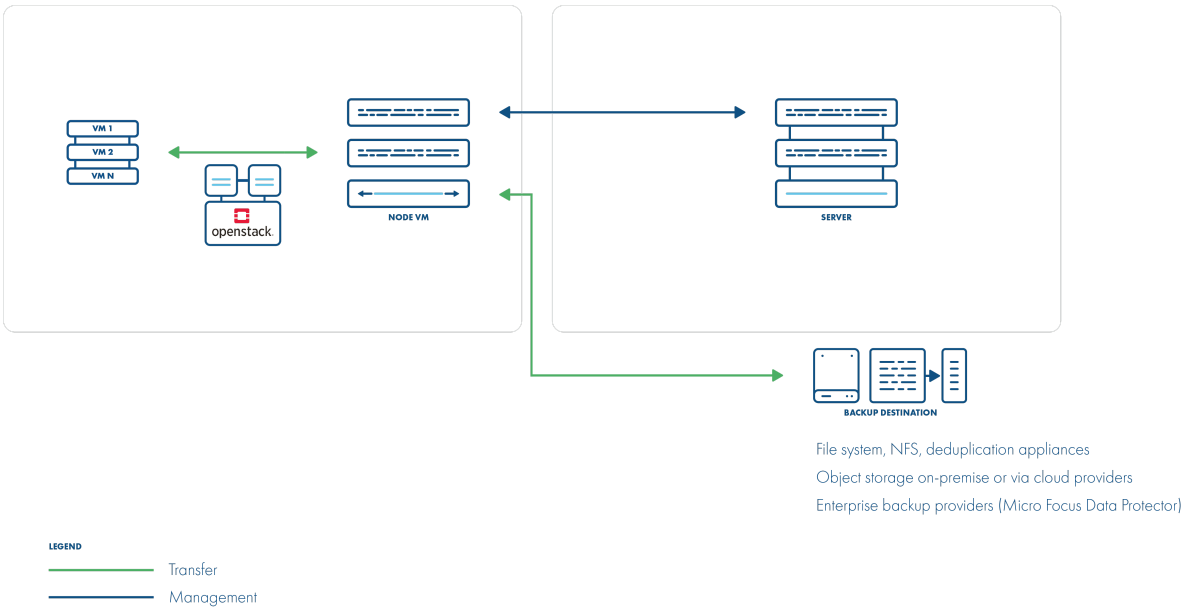

Backup Process

- crash-consistent snapshot using cinder API
- optional application consistency using pre/post snapshot command execution
- metadata exported from API
- volumes created from snapshotted disks are mounted one by one to the Proxy VM
- data read directly on the Proxy VM
- incremental backups supported for Ceph RBD - a list of the changed blocks are fetched from the monitors, and only these blocks are read from the attached disk on the Proxy VM
- if an instance is created from the glance image and "download image from glance" option is enabled data is downloaded from glance API, instance is

created from the instance metadata and the images which is fetched from the glance API

- restore creates empty disks on the Proxy VM, imports merged data then recreates the VM using these volumes, it will try to use the image from a glance if present in the target environment or it will upload the image to the glance and register it with the restored VM

Ceph RBD storage backend

Data Protector for Cloud Workloads also supports deployments with Ceph RBD as a storage backend. Data Protector for Cloud Workloads communicates directly with Ceph monitors using RBD export/RBD-NBD when used with the Libvirt strategy or - when used with the Disk-attachment method - only during incremental backups (snapshot difference).

Libvirt strategy

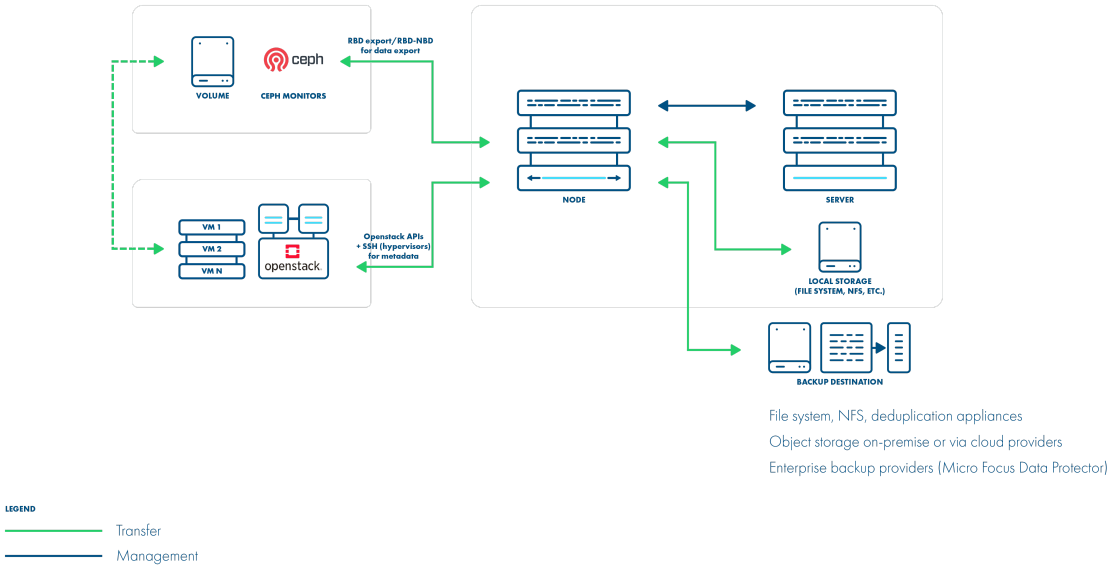

Disk attachment strategy

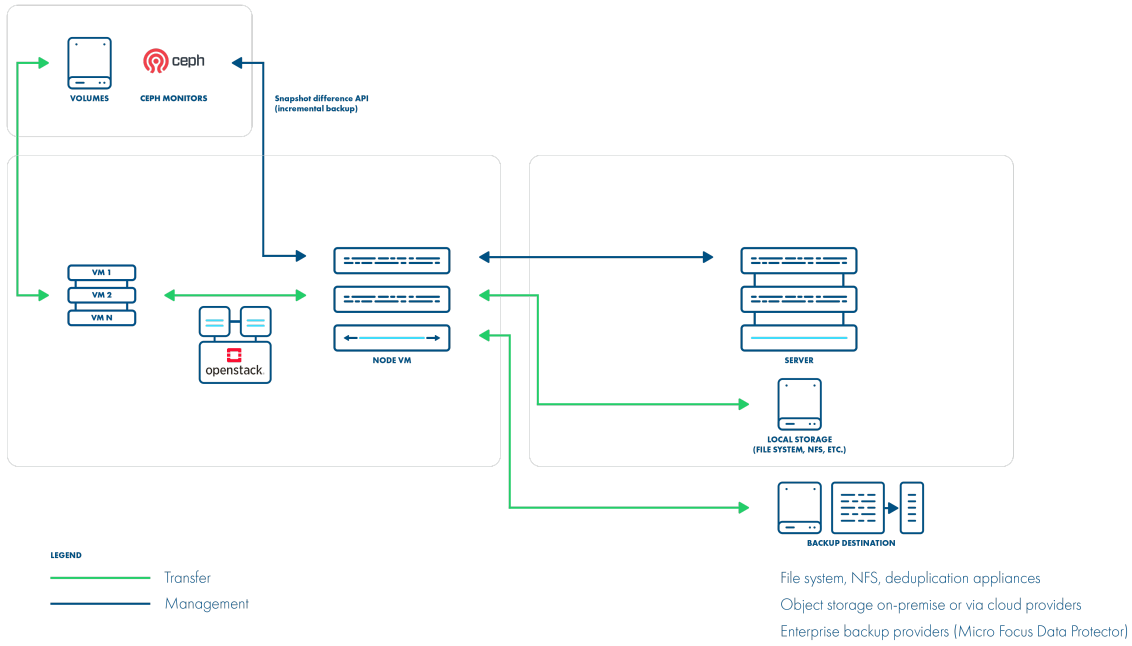

Data Protector for Cloud Workloads supports OpenStack with Ceph RBD volumes. Here is an example of a typical (expected) section that needs to be added in **cinder.conf**for Ceph in the OpenStack environment:

```
[rbd]
volume_backend_name = rbd
volume_driver = cinder.volume.drivers.rbd.RBDDriver
rbd_pool = volumes
rbd_ceph_conf = /etc/ceph/ceph.conf
rbd_flatten_volume_from_snapshot = false
rbd_max_clone_depth = 5
rbd_store_chunk_size = 4
rados_connect_timeout = -1
glance_api_version = 2
rbd_user = volumes
rbd_secret_uuid = ce6d1549-4d63-476b-afb6-88f0b196414f
```

A good article on how to set up Ceph with OpenStack can be found here ↗.

To set up the Openstack HVM with Ceph RBD volumes in Data Protector for Cloud Workloads:

1. Add Ceph storage as described here
2. Add the hypervisor manager as described here.

3. Go to **Virtual Environments**→ **Infrastructure**→ **Clusters**and select cluster that is used by Openstack.

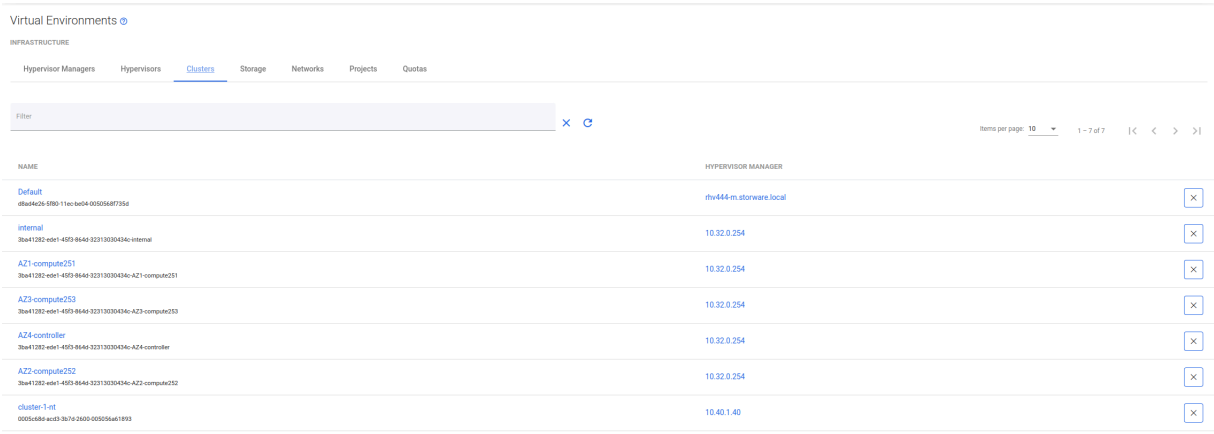

4. In **Storage Provider** field select previously added Ceph storage.

Now you can save and sync the inventory - if Ceph communication works properly you should be able to see Hypervisor Storage entries (in Hypervisors → Storage tab) representing your Ceph storage pools.

QCOW2 files on NFS storage

Example: scenario QCOW2 files residing on NFS

You can configure the NFS volume backend here:

https://docs.openstack.org/cinder/rocky/admin/blockstorage-nfs-backend.html ↗

Make sure the QCOW2 volumes are enabled.

For an NFS backend, it's recommended to set these values in **/etc/cinder/cinder.conf**:

```
default_volume_type=nfs nfs_sparsed_volumes = true nfs_qcow2_volumes = true volume_driver = cinder.volume.drivers.nfs.NfsDriver enabled_backends = nfs
```

Nova volumes

Data Protector for Cloud Workloads is able to backup nova volumes using libvirt strategy. In the hypervisor manager settings there is an option **Download image from glance**. When this option is enabled the original image from glance is downloaded. If it is disabled, then the image is not downloaded, however the nova volume created from it is backed up.

Adding hypervisor managers

When creating the hypervisor manager, provide the following data in the fields:

URL - Keystone API URL, e.g. **http://10.201.32.40:5000/v3**

Authentication domain:

- name - name of domain
- domainId - optional domain id
- user - OpenStack user.
- password - password for that user.
- default project - name of default project in domain.

Scope VMs to Domain - you can create one or more Authentication Domains based on this setting, as described in the **Authentication Domains** section below.

Download image from glance - allows Data Protector for Cloud Workloads to use images from glance as described in the disk attachment strategy

When you index the hypervisor manager, **make sure you provide the correct SSH credentials** for each hypervisor that appeared on the Hypervisors tab. You can also use SSH public key authentication.

> **Note:** When restoring the instances, make sure that the provided user is a member of the target tenant.

Authentication Domains

Data Protector for Cloud Workloads supports OpenStack environments with multiple domains. Each OpenStack Hypervisor Manager needs to have at least one Authentication Domain provided.

Data Protector for Cloud Workloads supports two types of domain authorization:

- Unscoped - single credentials to multiple domains
- Scoped - single credentials to single domain

Single credentials to single domain

Scoping VMs to Domain option needs to be turned on.

In that setup user can create Authentication Domains for every Domain in OpenStack environment. Projects and Virtual Machines are only scanned in provided Authentication Domains.

Single credentials to multiple domains

Scoping VMs to Domain option needs to be turned off.

In that setup user need to create only one Authentication Domain. Projects and Virtual Machines are scanned in every domain that provided user has access to.

Openstack tags

To list tags for specific instance:

```
root@c254:~# nova show d6787375-ea0c-49fd-878b-35b71747c62a |grep tags
| tags                                 | ["test"]
```

Tags for Openstack requires nova API version >= 2.26.

Openstack Access Keys

During Inventory Synchronization, Data Protector for Cloud Workloads scans all Keypairs (to which a user has access) and stores them as Access Keys. When restoring an instance, the user can specify the Access Key.

> **Note:** If the access key selected for restore is no longer present in the Openstack environment, it will be restored.

Openstack Flavor

During Inventory Synchronization, Data Protector for Cloud Workloads scans all Flavors and saves their configuration. When restoring an instance, the user can specify the flavor.

> **Note:** If the flavor selected for restore is no longer present in the Openstack environment, it will be restored with the same configuration and availability as the original flavor.

Instant restore

Node configuration

First step to initiate the configuration process for OpenStack Instant Restore service, is creating a directory that will be accessible for mounting by the NFS server. Create a specific directory on your Node machine that will be used as the target space for Instant Restore's shared resources.

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

Next, create an NFS share that will allow access to the **/vprotect_data/** directory from other machines in the network. Sharing this directory will enable OpenStack clients to use it for virtual machine restoration.

```
echo '/vprotect_data/ *
(fsid=0,no_subtree_check,rw,sync,no_root_squash,insecure)' >>
/etc/exports
exportfs -arv
```

OpenStack configuration

> Paths and commands may vary depending on the version of OpenStack you are using.

After creating the NFS share on the node, you need to configure the NFS backend in the OpenStack environment. This step will allow OpenStack to access resources stored on the node via NFS. Edit the **/etc/cinder/cinder.conf** file and add this section at the end of the file:

```
[nfs-instant-restore]
volume_backend_name=nfs-instant-restore
volume_driver=cinder.volume.drivers.nfs.NfsDriver
nfs_shares_config=/etc/cinder/nfs_instant_restore
nfs_snapshot_support=True
nfs_qcow2_volumes=True
nfs_sparsed_volumes=true
nfs_mount_options=vers=4
```

Create the file you provided in the configuration as the value of **nfs_shares_config** parameter

```
vi /etc/cinder/nfs_instant_restore
```

and path to the NFS share:

```
DP4CW_node_ip:/instant_restore
```

After creating the NFS server and configuring cinder, restart the cinder volume service. Please note that the name of this service may be different depending on the OpenStack version.

```
systemctl restart openstack-cinder-volume
```

Data Protector for Cloud Workloads web UI configuration

After completing the inventory synchronization of the OpenStack, in Node edition window you can select the storage with the NFS backend configuration. Details about the NFS backend should be supplied by the OpenStack administrator.

You should now be able to select "Instant restore" for backed up virtual machines.

Cross restore

Data Protector for Cloud Workloads allows restore of virtual machines between OpenStack and Virtuozzo environments.

Before you start

- You need to have OpenStack/Virtuozzo provider added to your environment
  - You need to have a completed backup
- You need to have OpenStack/Virtuozzo provider added to your environment
  - The import/export mode must be configured as Disk Attachment
  - The inventory synchronization process must be completed

Restore process

To initiate the restore process, start by accessing the restore interface and selecting the appropriate virtualization manager from the list, ensuring that OpenStack or Virtuozzo appear as options if they have been properly configured in your environment.

Limitations

- Data Protector for Cloud Workloads does not backup and restores keypairs that user used in Data Protector for Cloud Workloads doesn't have access to. The restored instance will have no keypairs assigned. In such a case, the keypairs have to be backed up and restored manually under the same name before restoring the instance.
- For the libvirt strategy only, QCOW2/RAW files or Ceph RBD are supported as the backend.
- The disk attachment method with Ceph requires access to the monitors from the Proxy VM.

OpenNebula

OpenNebula

Data Protector for Cloud Workloads supports backup for OpenNebula in Disk Attachment Strategy with changed block tracking

- supports KVM hypervisors
- supports incremental backup
- proxy VM is required - used for the disk attachment process

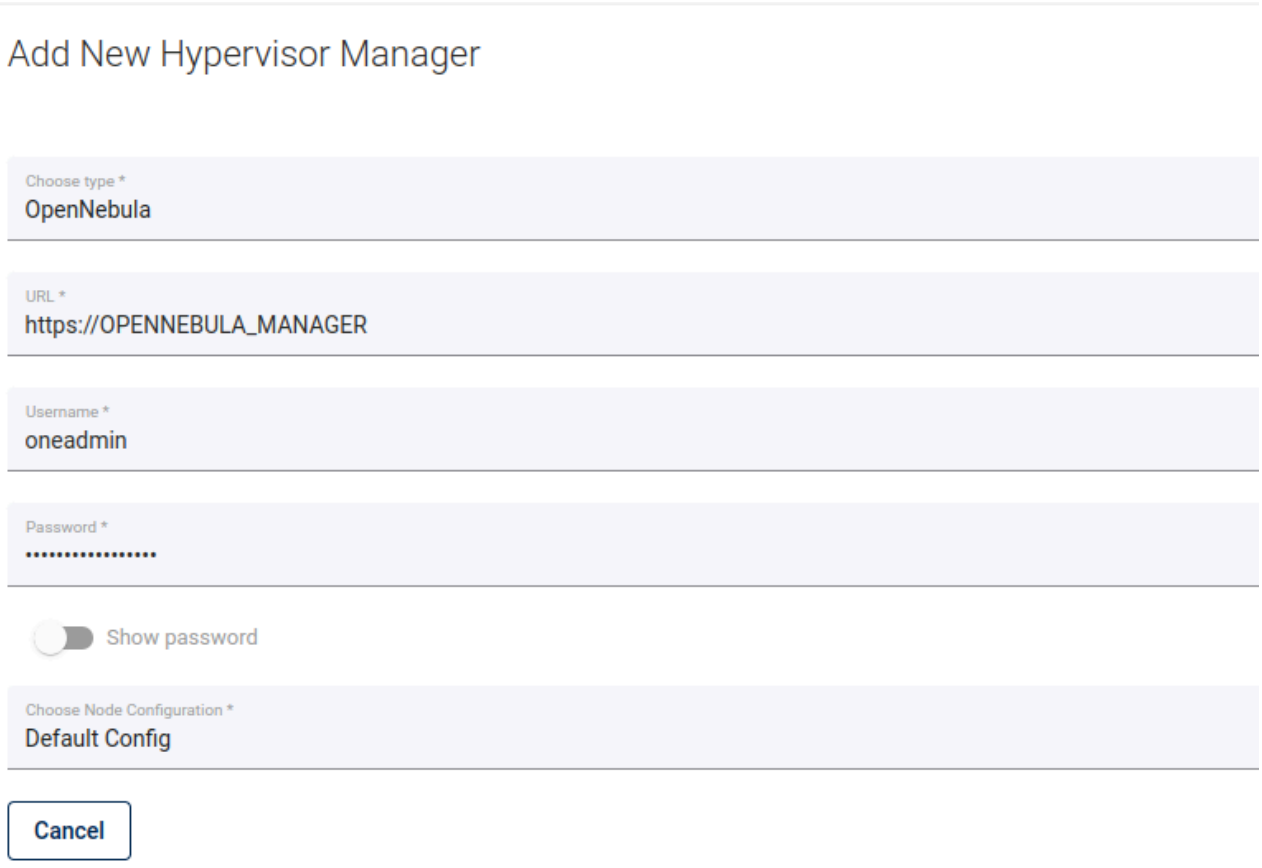

Backup Strategies

Disk Attachment CBT strategy

Backup process

- Both full and incremental backup is supported.
- OpenNebula SDK is used for all operations on hypervisor manager using XML-RPC API.
- Crash consistent snapshots are taken of every non excluded disk.
- For every exported disk, Data Protector for Cloud Workloads fetch its snapshot and create a new image from it. Later, the new disk is attached to proxy VM from newly created image. After copying data from disk to new backup file, disk is detached and image removed.
- While exporting disks Data Protector for Cloud Workloads is scanning for changes based on checksums. This allows us to later perform incremental backup using these checksums and we don't require last snapshot to be present on OpenNebula.
- VM metadata is retrieved using API and saved to the backup file.
- Used images metadata is retrieved using API and saved to the separate backup files.

> Our recommendation is also to configure LVM filters on Proxy VM. You need to add all OS disks and partitions, please follow these steps: LVM setup on Data Protector for Cloud Workloads Node for disk attachment backup mode

Restore process

- When restoring VM, Data Protector for Cloud Workloads restore disks first and then create from saved metadata new VM, assigning already restored disks to this VM.
- When restoring disks that were excluded from backup, Data Protector for Cloud Workloads create a new empty image.
- When restoring non excluded disks, Data Protector for Cloud Workloads create a new empty image, then attach new disk to proxy VM from this image. After copying data from the backup file to the attached disk, a new image is created from the data populated disk. This disk is later detached, and the empty image is removed.
- In order to create a restored VM on OpenNebula, Data Protector for Cloud Workloads first create a template of it, and then instantiate it. After the new VM is created, the template is removed.

- VM is always restored already in RUNNING state.
- When restoring VM we can choose SYSTEM datastore in which VM will start running after restore.
- When restoring VM we can choose IMAGE datastore in which images for VM will be created.

Snapshots

- Snapshot management is supported.
- OpenNebula does not allow reverting snapshot when VM is still running. Because of that, for the duration of reverting all disk snapshots, VM is suspended.
- Creating and reverting snapshots for every disk is not performed in parallel, because OpenNebula does not allow it.

Limitations

- Currently only KVM hosts are supported.
- Export of volatile disks (VM disks that were not created from image and hold their content for the duration of VM's deployment) is not supported
- QCOW2 based disks are not supported.

Virtuozzo

Virtuozzo Hybrid Infrastructure

Data Protector for Cloud Workloads supports backup for Virtuozzo Hybrid Infrastructure:

- Disk image transfer - for KVM hypervisors with VMs using QCOW2
- Volumes storage:
  - supports incremental backup
  - disk images are transferred directly from API (no Proxy VM required)

Libvirt Backup Strategy

Data Protector for Cloud Workloads supports Virtuozzo environments that use KVM hypervisors and VMs running on QCOW2 or RAW files. Data Protector for Cloud Workloads communicates with Virtuozzo APIs such as Nova and Glance to collect metadata and for the import of the restored process. However, the actual backup is done over SSH directly from the hypervisor. The process is exactly the same as in Deployment in the KVM/Xen environment. Data Protector for Cloud Workloads Node can be installed anywhere - it just needs to have access to the Virtuozzo APIs and hypervisor SSH via a network. Both full and incremental backups are supported.

Backup Process

- direct access to the hypervisor over SSH
- crash-consistent snapshot taken directly using virsh (QCOW2/RAW file)
- optional application consistency using pre/post snapshot command execution • QCOW2/RAW-file data exported over SSH (optionally with netcat)
- metadata exported from Virtuozzo APIs (nova, glance, cinder)

- the last snapshot kept on the hypervisor for the next incremental backup (if at least one schedule assigned to the VM has backup type set to incremental)
- restore recreates files/volumes according to their backend (same transfer mechanism as used in backup) and then defines VM on the hypervisor

Nova volumes

Data Protector for Cloud Workloads is able to backup nova volumes using libvirt strategy. In the hypervisor manager settings there is an option **Download image from glance**. When this option is enabled the original image from glance is downloaded. If it is disabled, then the image is not downloaded, however the nova volume created from it is backed up.

Adding hypervisor managers

When creating the hypervisor manager, provide the following data in the fields:

URL - Keystone API URL, e.g. `https://YOUR_VIRTUOZZO_IP_MGMT:5000/v3`

Authentication domain:

- name - name of domain
- domainId - optional domain id
- user - Virtuozzo user.
- password - password for that user.
- default project name - name of default project in domain.

Scope VMs to Domain - you can create one or more Authentication Domains based on this setting, as described in the **Authentication Domains** section below.

Download image from a glance - allows Data Protector for Cloud Workloads to use images from a glance.

When you index the hypervisor manager, **make sure you provide the correct SSH credentials** for each hypervisor that appeared on the Hypervisors tab. You can also

use SSH public key authentication.

> **Note:** When restoring the instances, make sure that the provided user is a member of the target tenant.

Virtuozzo Access Keys

During Inventory Synchronization, Data Protector for Cloud Workloads scans all Keypairs (to which a user has access) and stores them as Access Keys. When restoring an instance, the user can specify the Access Key.

> **Note:** If the access key selected for restore is no longer present in the Virtuozzo environment, it will be restored.

Virtuozzo Flavor

During Inventory Synchronization, Data Protector for Cloud Workloads scans all Flavors and saves their configuration. When restoring an instance, the user can specify the flavor.

> **Note:** If the flavor selected for restore is no longer present in the Virtuozzo environment, it will be restored with the same configuration and availability as the original flavor.

Cross restore

Data Protector for Cloud Workloads allows restore of virtual machines between OpenStack and Virtuozzo environments.

Before you start

- You need to have OpenStack/Virtuozzo provider added to your environment
    - You need to have a completed backup
- You need to have OpenStack/Virtuozzo provider added to your environment
    - The import/export mode must be configured as Disk Attachment
    - The inventory synchronization process must be completed

Restore process

To initiate the restore process, start by accessing the restore interface and selecting the appropriate virtualization manager from the list, ensuring that OpenStack or Virtuozzo appear as options if they have been properly configured in your environment.

Limitations

- Data Protector for Cloud Workloads does not backup and restores keypairs that user used in Data Protector for Cloud Workloads doesn't have access to. The restored instance will have no keypairs assigned. In such a case, the keypairs have to be backed up and restored manually under the same name before restoring the instance.
- Only QCOW2/RAW files are supported as the backend.

Citrix Hypervisor (XenServer)

Citrix Hypervisor (XenServer)

Backup Strategies

XVA-based

In this strategy, the VM is exported as a single XVA bundle containing all of the data. Incremental backup is also supported. Data is transferred directly from the XenServer API without the need to set up anything on the hosts.

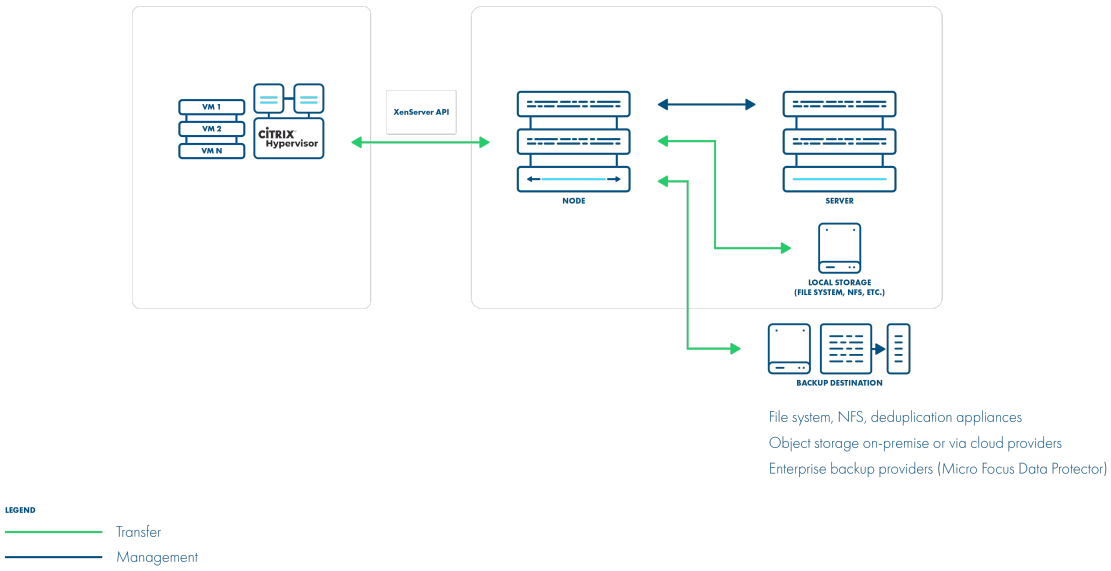

Backup Process

- crash-consistent snapshot using hypervisor's API only for full backups
- optionally quiesced snapshot can be done if enabled and guest tools installed inside - if quiesced snapshot has been failed we are doing regular one
- optional application consistency using pre/post snapshot command execution

- data export directly from the hypervisor using hypervisor's API - both full (XVA) and delta (VHD for each disk)
- full backup (XVA) contains metadata
- snapshot taken with full backup is kept on the hypervisor for the next incremental backup - if at least one schedule assigned to the VM has backup type set to incremental
- incremental backups are cumulative (all data since last full backup)
- restore recreates VM from XVA, and then applies changes from each incremental backup using Hypervisor APIs

Changed-Block Tracking

In this strategy, the VM is exported using XenServer API (full backup) and the Network Block Device service (NBD, incremental backups) on the XenServer hosts. The CBT feature in Citrix XenServer 7.3+ may require an additional license. The resulting backup has separate files for each disk + metadata, so you also have the option to exclude specific drives.

> **Note:** For full backups only you can still use this strategy without CBT enabled on the hypervisor.

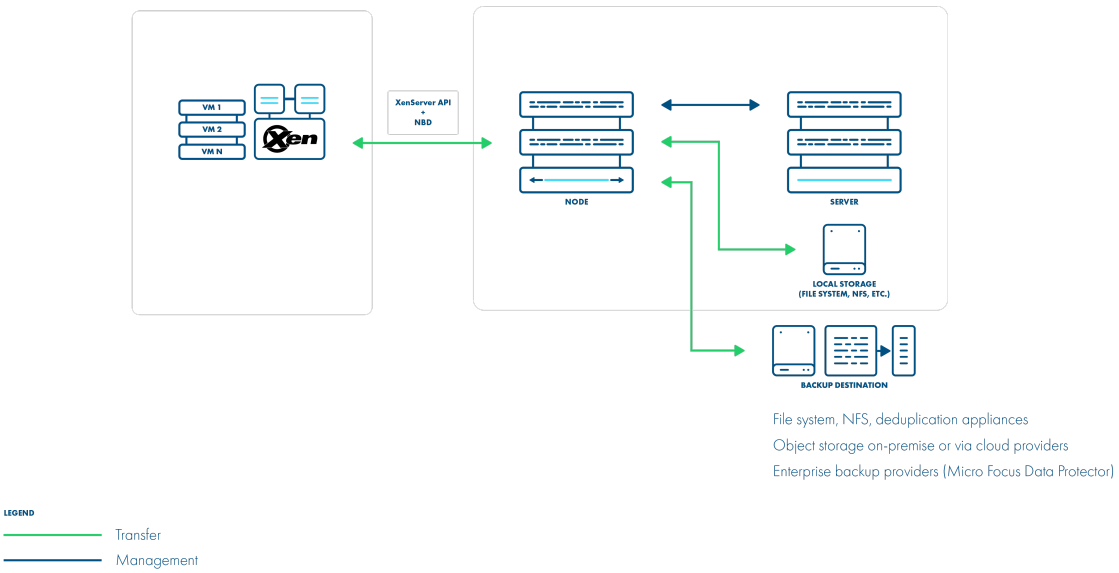

Backup Process

- crash-consistent snapshot using hypervisor's API
- optionally quiesced snapshot can be done if enabled and guest tools installed inside - if quiesced snapshot has been failed we are doing regular one
- optional application consistency using pre/post snapshot command execution
- CBT enabled during full backup on each disk if it wasn't done earlier
- metadata exported from API
- full backup - each disk exported from API (RAW format)
- incremental backup - each disk queried for changed blocks and which are exported over NBD
- last snapshot kept on the hypervisor for the next incremental backup - if at least one schedule assigned to the VM has backup type set to incremental
- restore recreates VM from metadata using API and imports merged chain of data for each disk using API

Change Block Tracking setup

Citrix introduced the CBT mechanism in XenServer 7.3. In order to enable CBT backups, the following requirements must be met:

1. Citrix XenServer 7.3 or above must be used - note that CBT is a licensed feature
2. The NBD server must be enabled on the hypervisor
3. The NBD client and NBD module must be installed on Data Protector for Cloud Workloads Node

Notes on restore

1. When image-based backups (XVA) are used - Data Protector for Cloud Workloads restore VMs as templates and renames them appropriately after the restore
2. When separate disk backups are used:
    - if there is already a VM in the infrastructure with the UUID of the VM being restored (check `present` flag in VM list) - Data Protector for Cloud

Workloads restores it as a new VM (MAC addresses will be generated)

- otherwise Data Protector for Cloud Workloads attempts to restore the original configuration including MAC addresses

NBD Server setup (on XenServer)

1. Get the Network UUID that you intend to use for communication with Data Protector for Cloud Workloads - run on the XenServer shell:

```
[root@xenserver-cbt ~]# xe network-list
uuid ( RO)            : e16b4e34-47d4-9a6e-371b-65beb7252d69
       name-label ( RW): Pool-wide network associated with eth0
 name-description ( RW):
           bridge ( RO): xenbr0
..........
uuid ( RO)            : 244a2fa7-ae7c-e45c-819a-44cecf51e8fa
       name-label ( RW): Host internal management network
 name-description ( RW): Network on which guests will be assigned a
private link-local IP address which can be used to talk XenAPI
           bridge ( RO): xenapi
```

   For example: `e16b4e34-47d4-9a6e-371b-65beb7252d69`

2. Enable the NBD service on your hypervisor:

```
xe network-param-add param-name=purpose param-key=nbd
uuid=<network-uuid>
```

NBD Client setup (on Data Protector for Cloud Workloads Node)

Data Protector for Cloud Workloads comes with a pre-built RPM and modules.

1. Go to the NBD directory:

```
cd /opt/vprotect/scripts/nbd
```

2. If your Linux does not have the NBD module installed you may try to build one yourself (there is a script for Red Hat based distributions that downloads the kernel, enables the NBD module, and builds it) or use an already provided module:

   - you can compile the module by running:

```
./compile_nbd_module.sh
```

3. Enable the module by invoking the script (the following command will either use a module in your kernel or copy the provided `nbd.ko`):

```
./enable_nbd.sh
```

4. Verify that you have `/dev/nbd*` devices available on your Data Protector for Cloud Workloads Node host:

```
[root@localhost nbd]# ls /dev/nbd*
/dev/nbd0  /dev/nbd1  /dev/nbd10  /dev/nbd11  /dev/nbd12  /dev/nbd13
/dev/nbd14  /dev/nbd15  /dev/nbd2  /dev/nbd3  /dev/nbd4  /dev/nbd5
/dev/nbd6  /dev/nbd7  /dev/nbd8  /dev/nbd9
```

5. Restart your Data Protector for Cloud Workloads Node:

```
systemctl restart vprotect-node
```

Limitations

- VM migration between clusters is not supported.
- Cloning VM to another cluster is not supported.

XCP-ng

XCP-ng

Backup Strategies

XVA-based

In this strategy, the VM is exported as a single XVA bundle containing all of the data. Incremental backup is also supported. Data is transferred directly from the XenServer API without the need to set up anything on the hosts.

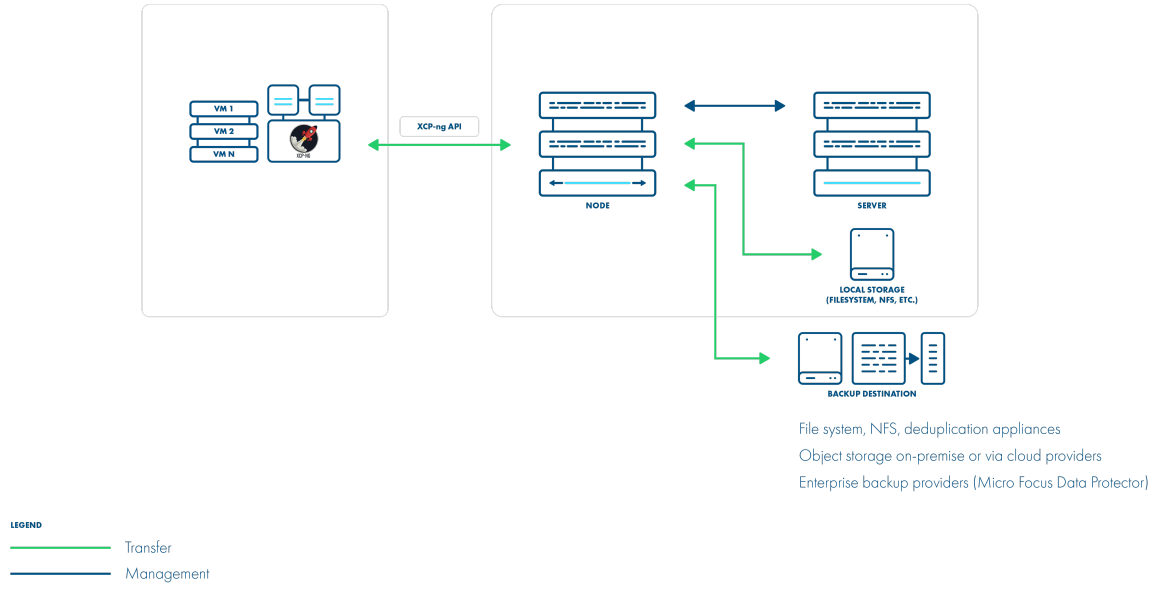

Backup Process

- crash-consistent snapshot using hypervisor's API only for full backups
- optionally quiesced snapshot can be done if enabled and guest tools installed inside - if the quiesced snapshot has been failed we are doing regular one
- optional application consistency using pre/post snapshot command execution

- data export directly from the hypervisor using hypervisor's API - both full (XVA) and delta (VHD for each disk)
- full backup (XVA) contains metadata
- snapshot taken with full backup is kept on the hypervisor for the next incremental backup - if at least one schedule assigned to the VM has backup type set to incremental
- incremental backups are cumulative (all data since last full backup)
- restore recreates VM from XVA, and then applies changes from each incremental backup using Hypervisor APIs

Changed-Block Tracking

In this strategy, the VM is exported using XenServer API (full backup) and the Network Block Device service (NBD, incremental backups) on the XenServer hosts. The CBT feature in Citrix XenServer 7.3+ may require an additional license. The resulting backup has separate files for each disk + metadata, so you also have the option to exclude specific drives.

> **Note:** For full backups only you can still use this strategy without CBT enabled on the hypervisor.

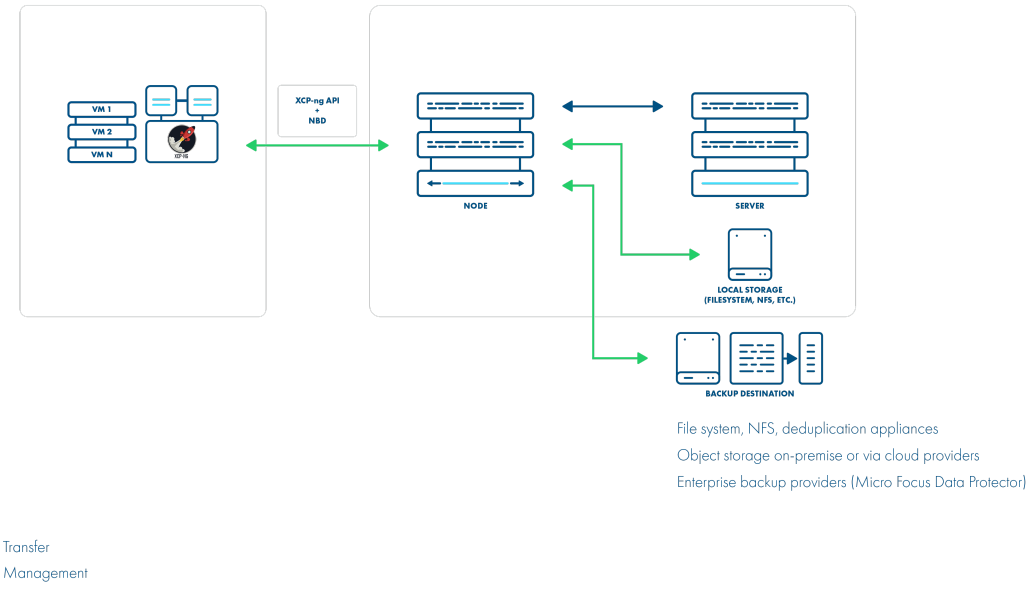

Backup Process

- crash-consistent snapshot using hypervisor's API
- optionally quiesced snapshot can be done if enabled and guest tools installed inside - if the quiesced snapshot has been failed we are doing regular one
- optional application consistency using pre/post snapshot command execution
- CBT enabled during full backup on each disk if it wasn't done earlier
- metadata exported from API
- full backup - each disk exported from API (RAW format)
- incremental backup - each disk queried for changed blocks and which are exported over NBD
- the last snapshot is 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 API

Change Block Tracking setup

Citrix introduced the CBT mechanism in XenServer 7.3. In order to enable CBT backups, the following requirements must be met:

1. Citrix Hypervisor 7.3 (XCP-ng 7.4) or above must be used - note that CBT is a licensed feature
2. The NBD server must be enabled on the hypervisor
3. The NBD client and NBD module must be installed on Data Protector for Cloud Workloads Node (Data Protector for Cloud Workloads should take care of this automatically during installation)

Notes on restore

1. When image-based backups (XVA) are used - Data Protector for Cloud Workloads restore VMs as templates and renames them appropriately after the restore
2. When separate disk backups are used:

- if there is already a VM in the infrastructure with the UUID of the VM being restored (check `present` flag in VM list) - Data Protector for Cloud Workloads restore it as a new VM (MAC addresses will be generated)
- otherwise Data Protector for Cloud Workloads attempts to restore the original configuration including MAC addresses

NBD Server setup (on XenServer)

1. Get the Network UUID that you intend to use for communication with Data Protector for Cloud Workloads - run on the XenServer shell:

   ```
   [root@xenserver-cbt ~]# xe network-list
   uuid ( RO)            : e16b4e34-47d4-9a6e-371b-65beb7252d69
          name-label ( RW): Pool-wide network associated with eth0
    name-description ( RW):
              bridge ( RO): xenbr0
   ..........
   uuid ( RO)            : 244a2fa7-ae7c-e45c-819a-44cecf51e8fa
          name-label ( RW): Host internal management network
    name-description ( RW): Network on which guests will be assigned a
   private link-local IP address which can be used to talk XenAPI
              bridge ( RO): xenapi
   ```

   For example: `e16b4e34-47d4-9a6e-371b-65beb7252d69`
2. Enable the NBD service on your hypervisor:

   ```
   xe network-param-add param-name=purpose param-key=nbd
   uuid=<network-uuid>
   ```

NBD Client setup (on Data Protector for Cloud Workloads Node)

> **Note:** This part is done by Data Protector for Cloud Workloads automatically during installation. The article may be helpful in case of problems with the NBD module.

Data Protector for Cloud Workloads comes with a pre-built RPM and modules.

1. Go to the NBD directory:

```
cd /opt/vprotect/scripts/nbd
```

2. If your Linux does not have the NBD module installed you may try to build one yourself (there is a script for Red Hat based distributions that downloads the kernel, enables the NBD module, and builds it) or use the already provided module:
   - you can compile the module by running:

     ```
     ./compile_nbd_module.sh
     ```

3. Enable the module by invoking the script (the following command will either use a module in your kernel or copy the provided `nbd.ko`):

   ```
   ./enable_nbd.sh
   ```

4. Verify that you have `/dev/nbd*` devices available on your Data Protector for Cloud Workloads Node host:

   ```
   [root@localhost nbd]# ls /dev/nbd*
   /dev/nbd0  /dev/nbd1  /dev/nbd10  /dev/nbd11  /dev/nbd12  /dev/nbd13
   /dev/nbd14  /dev/nbd15  /dev/nbd2  /dev/nbd3  /dev/nbd4  /dev/nbd5
   /dev/nbd6  /dev/nbd7  /dev/nbd8  /dev/nbd9
   ```

5. Restart your Data Protector for Cloud Workloads Node:

   ```
   systemctl restart vprotect-node
   ```

Limitations

- VM migration between clusters is not supported.
- Cloning VM to another cluster is not supported.

SC//Platform

SC//Platform

General

All the operations are using REST API to communicate with the HC3 cluster. Depending on the version of SC//Platform, additional configuration of the cluster may be required.

1. REST API should be enabled in the settings.

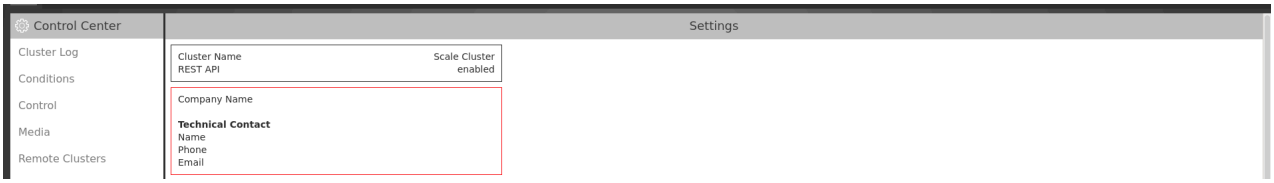

2. Before taking a snapshot, the current state of the `disableSnapshotting` flag of every disk in the virtual machine is checked. If even one of the disks has disabled snapshotting, the snapshot task ends with the error. All disks with disabled snapshotting are listed in the error message.


```
{
  "uuid": "7add7112-df10-499d-9c60-2f1462b7386e",
  "virDomainUUID": "121a4382-67b7-4bf2-8bcd-a4bcf45a10db",
  "type": "IDE_DISK",
  "cacheMode": "WRITETHROUGH",
  "capacity": 2000683008,
  "allocation": 1017118720,
  "physical": 0,
  "shareUUID": "",
  "path": "scribe/7add7112-df10-499d-9c60-2f1462b7386e",
  "slot": 1,
  "name": "",
  "disableSnapshotting": false,
  "tieringPriorityFactor": 8,
  "mountPoints": [],
  "createdTimestamp": 1631880304,
  "readOnly": false
},
```


Backup Strategies

Export Storage Domain Strategy

Export storage domain strategy performs the export using a Samba server running on a node machine. This strategy does not require a proxy VM.

> **Note:**
>
> - supported version: 8.9.x
> - no incremental backup
> - backup files are transferred to the SMB file share created on the node host (no Proxy VM required

In general, the export process is as follows:

• A snapshot of the virtual machine is taken

• Samba server is started on the node machine and export share is added

• Export of the VM to created SMB share is ordered

• Export share is removed and Samba server is stopped after Data Protector for Cloud Workloads finishes copying files

• Metadata of the excluded disks is exported

In general, the restore process is as follows:

• Samba server is started on the node machine and import share is added

• Import of the VM to the created SMB share is ordered

• Empty disks for excluded disks are added to the restored VM

• Import share is removed and Samba server is stopped

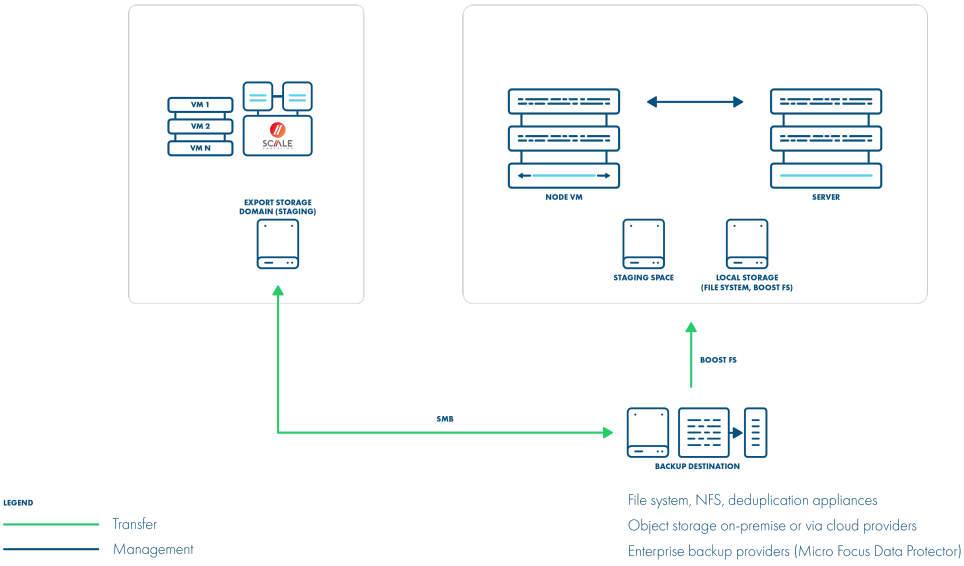

Disk Attachment Strategy

Disk attachment strategy performs the export using proxy VM running on the HC3 cluster.

> **Note:**
>
> - supported version: 8.9.x
> - incremental backup available
> - a proxy VM is required - used for the disk attachment process.

In general, the export process is as follows:

• A snapshot of the virtual machine is taken

• The disks from the taken snapshot that were not excluded are attached to the proxy VM

• If full export is being performed, then the whole content of the attached disks is saved to the backup files. If incremental export is being performed, then the list of changed blocks is first retrieved and then only the content of these blocks is saved.

• The disks are detached from the proxy VM

In general, the restore process is as follows:

- A new virtual machine is created using exported metadata
- Empty disks are attached to the proxy VM
- Content of the exported disk files is written to the attached disks
- The restored disks are then reattached from the proxy VM to the restored VM
- List of boot devices for the restored VM is set

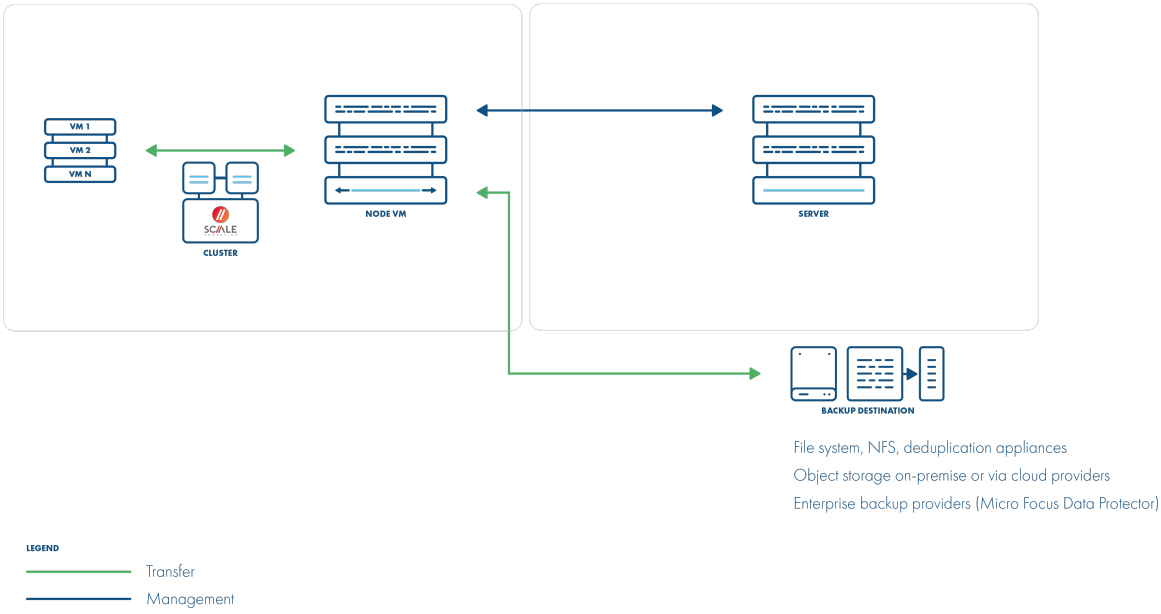

Adding HC3 as Hypervisor Manager

When adding a new Scale HC3 as Hypervisor Manager in Data Protector for Cloud Workloads, configure the following fields:

- URL: Address of HC3 hypervisor server: hostname or IP with HTTPS
- Fill in the admin username and password for HC3
- Select the export method

- The node configuration responsible for backup operations

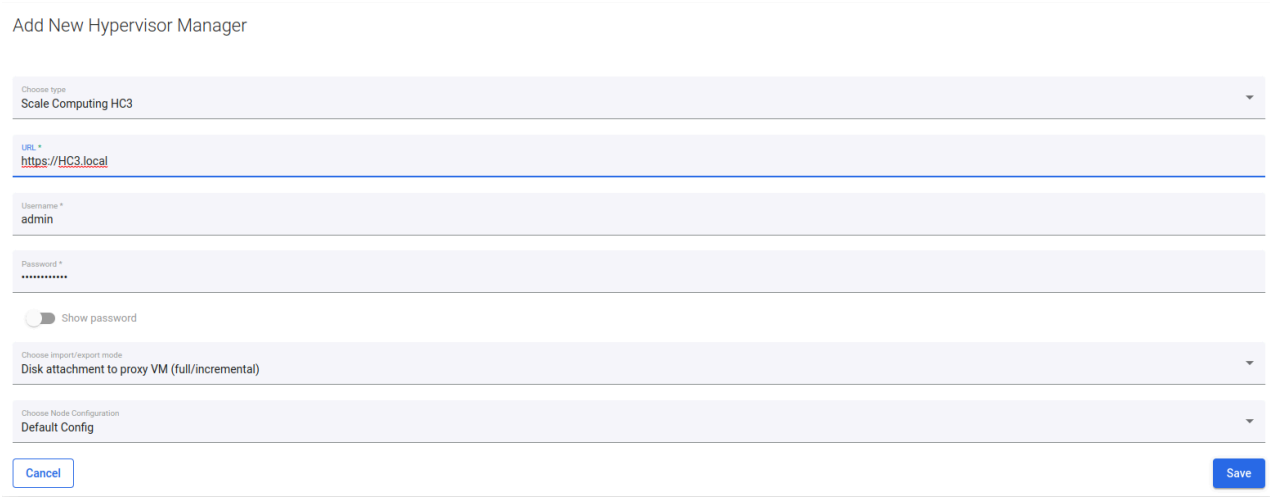

Cloud

Cloud

Data Protector for Cloud Workloads supports also cloud platforms. such as Amazon EC2, Microsoft 365 and GCP.

In this section, you will find information about the cloud compute backup strategy and specific steps needed to use them with Data Protector for Cloud Workloads.

- Amazon EC2
- GCP GCE
- Azure Cloud

> **Note:** Information about protecting Microsoft 365 you can find in the chapter Protecting Microsoft 365

Amazon EC2

Amazon EC2

Data Protector for Cloud Workloads supports the Amazon EC2 cloud platform by using a VM called “Proxy VM”. The node invokes commands on the AWS to snapshot and attach EBS 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 means that you need to create an EC2 instance (Proxy VM) in each zone.

The Data Protector for Cloud Workloads Server can be deployed anywhere, but keep in mind that Nodes need to be able to call the server over HTTP(S) on the port you have specified.

The AWS backup strategy allows you to exclude drives from the backup that you don’t need. Remember that you need to install 1 Proxy VM per AWS zone so that drives that the Node tries to attach are reachable.

All backup destinations can be used, but keep in mind that you may be charged for transferring data between regions, AZs and external backup providers.

Data Protector for Cloud Workloads Node has access to instances only in the **zone where it is hosted**.

Data Protector for Cloud Workloads Node requires the **account ID, access key** and **secret key** to connect to the AWS account.

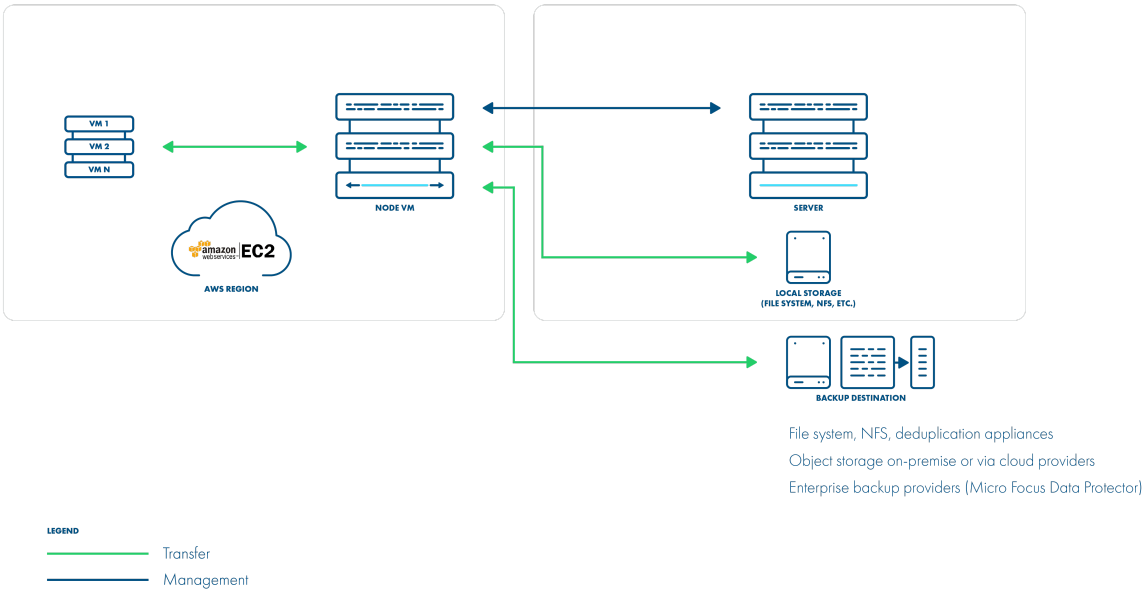

Typical use scenarios

There are several scenarios for AWS which may be suitable for your case:

- **Backup EC2 to S3** - in this case after dumping backup, Data Protector for Cloud Workloads can push them to the S3 bucket. You may consider using a VPC endpoint to boost your store operation performance.
- **Backup EC2 to EBS volume on the proxy** - you can use PowerProtect DD to deduplicate data and optimize your storage consumption significantly. Keep in mind that you may want to protect your EBS volume using EBS snapshots as well.
- **Backup EC2 to your local backup provider** - if you already have a central enterprise backup solution, you may want to use it as a backup provider for EC2 instances running in AWS. You should consider using Direct connect to have a higher bandwidth available.
- **Backup EC2 to your other cloud provider** - If you're using multiple clouds, you also may consider storing data in GCS or Azure backup providers.

> **Note**: In all cases, depending on your target, you may be charged for data transfers.

Setup considerations

It is assumed that you have working experience with Amazon EC2 to be able to deploy Data Protector for Cloud Workloads components. You also need to have an IAM user with permissions that allow you to deploy an instance and generate access/secret keys for Data Protector for Cloud Workloads.

Remember to use **CentOS 8 AMI** as a base image - both for the Server and Nodes. For a typical installation, we recommend 2 virtual processors and 4 GB of RAM. This means that **t3.medium** or **m5.large** should cover general use cases. For better performance, however, we recommend using storage optimized instances such as **i3.large** or bigger, where I/O intensive operations should perform better.

Both Data Protector for Cloud Workloads components are assumed to be deployed without HA (more precisely, all the nodes or server will probably be in separate AZs, and only need to communicate over HTTP). There is no requirement for multi-AZ deployment for now. While the Node is stateless and can be lost without data loss, the Server needs DB to be protected. Data Protector for Cloud Workloads provides a built-in automatic DB backup mechanism, which can be used to protect backup metadata. Refer to the Disaster recovery ↗ section for more details.

From a networking perspective, Data Protector for Cloud Workloads requires communicating with Amazon EC2 API, but it is still recommended to put in a private subnet and allow communication over a NAT Gateway.

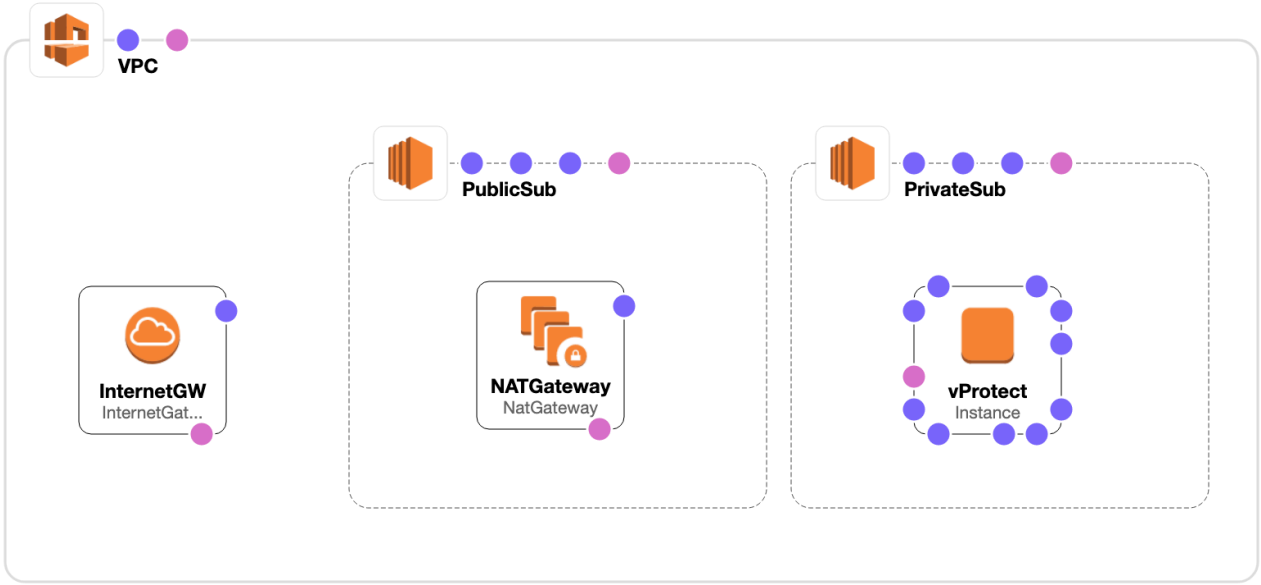

You add Amazon EC2 as a Hypervisor Manager. You need to provide the account ID and access/secret keys of a user that has permissions to handle snapshot, AMI and EBS volume operations, and EC2 instance creation.

On the same screen, you also specify if the AMIs of root volumes should be created during the backup process. For Windows instances, we recommend also keeping an AMI image with each backup to have the option to restore the original root volume as well. You can also skip AMI creation, but this means that during restore you need to specify the appropriate AMI ID that you want to boot from.

Permissions

Here are the IAM permissions that Data Protector for Cloud Workloads needs to have for backup/restore operations.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1565003475859",
      "Action": [
        "ec2:AttachNetworkInterface",
        "ec2:AttachVolume",
        "ec2:CreateImage",
        "ec2:CreateSecurityGroup",
        "ec2:CreateSnapshot",
        "ec2:CreateTags",
        "ec2:CreateVolume",
        "ec2:DeleteSnapshot",
        "ec2:DeleteVolume",
        "ec2:DeregisterImage",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeImages",
        "ec2:DescribeInstances",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSnapshots",
        "ec2:DescribeVolumes",
        "ec2:DetachVolume",
        "ec2:RegisterImage",
        "ec2:RunInstances",
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:ec2:*:*:instance/*",
        "arn:aws:ec2:*:*:volume/*",
        "arn:aws:ec2:*:*:snapshot/*",
        "arn:aws:ec2:*:*:image/*",
        "arn:aws:ec2:*:*:network-interface/*",
        "arn:aws:ec2:*:*:security-group/*"
      ]
    }
  ]
}
```

Adding a hypervisor manager

To properly configure your AWS account, go to Data Protector for Cloud Workloads → Virtual Environments → Infrastructure → add hypervisor manager

Enter parameters such as:

- Account ID

https://docs.aws.amazon.com/IAM/latest/UserGuide/console_account-alias.html#FindingYourAWSId"

- Access key
- Secret Key
- Enable/disable Windows, Linux image

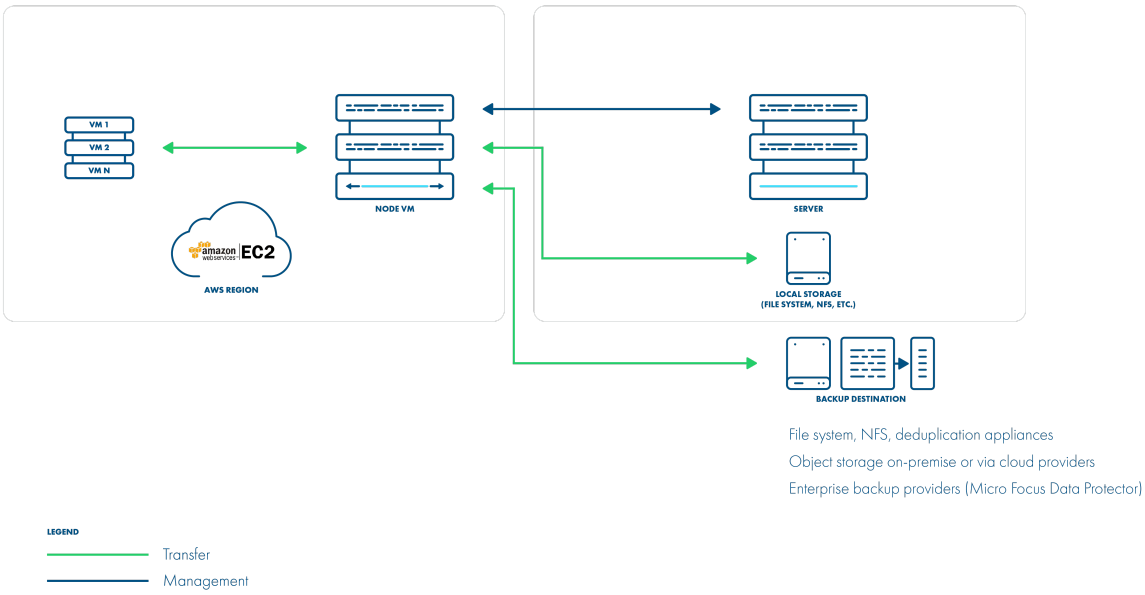

> **Note:** When Data Protector for Cloud Workloads creates a backup, some operating systems, such as Windows, may require an AMI for later restores in order to keep your OS settings. With this option, Data Protector for Cloud Workloads will keep the AMI necessary for future restores in your AWS account. Without this image, a new instance will have to be started with a fresh root device and additional volumes attached, which may not contain your OS-related settings, licenses, or data that were stored on the root device.

Backup modes

These settings regarding Windows or Linux images are required to define the way backups and restores are done. AWS supports 2 ways:

1. Using AMI and AWS snapshots:
    - During Export or Snapshot tasks an AMI is created for the root volume, other volumes are snapshotted.
    - The AMI is stored in AWS until Data Protector for Cloud Workloads snapshot or backup removal is initiated.
    - During Restore, a new instance is launched from a previously exported AMI, and imported non-root volumes are attached.
2. Using AWS snapshots:
    - During Export or Snapshot tasks, all volumes are snapshotted.
    - During Restore, an AMI is created from the imported root volume, a new instance is launched and imported non-root volumes are attached. The AMI is then removed.

In both cases, volume snapshots are kept in AWS only if the Data Protector for Cloud Workloads Snapshot task is completed.

> **Note:** A Windows AMI created from a snapshot is not launchable, hence it is recommended to **enable using AMI for the Windows platform**.

Backup strategies

Amazon EC2 supports two backup and restore strategies:

1. Disk attachment (full)
2. Disk attachment with changed block tracking (full/incremental)

Restore

It is possible to specify another AMI for attaching non-root volumes during restore. You can also specify an availability zone for a new instance.

Multi-zone configuration

To secure instances from multiple zones and regions, you need to register additional Data Protector for Cloud Workloads Nodes. For every zone you want to backup instances from, you need to create separate Node Configuration. Each Node Configuration needs to be assigned to correct Hypervisor Cluster (which reflects compute zone from AWS).

First synchronization task needs only one Node. After first scan, you can assign Node Configurations to Hypervisor Clusters, and run synchronization task once again to fetch all instances from other zones.

To assign Node Configuration to Hypervisor Cluster, go to Clusters list in Infrastructure tab. Next, click on selected Cluster to choose Node Configuration.

> **Note**: For backup and restore between zones, nodes and node configurations should have access to the same backup destination.

Costs

From the AWS perspective, you need to take inot account several additional costs that may be incurred:

- EC2 instance costs for the Data Protector for Cloud Workloads Server and Nodes:
  - depends on the number of nodes (assume at least one node per zone)
  - to reduce costs we recommend to use reserved instances for production use
  - https://aws.amazon.com/ec2/pricing/ ↗
- Backup destination and staging space storage on EBS:

    - staging space is necessary and we recommend it to be at least the size of the biggest VM multiplied by the number of export and store threads
    - if you want to store backups on EBS you also need to have additional storage
    - you can have both using the same EBS volume
    - we encourage you to use deduplication, as it may even result in over 95% of storage savings
    - https://aws.amazon.com/ebs/pricing/ ↗
- Data transfer costs:
    - if you upload data to external backup providers, or if a node needs to transfer a lot of data between AZs - this can be reduced by deploying one node per AZ
    - https://aws.amazon.com/ec2/pricing/on-demand/#Data_Transfer ↗

GCP GCE

Google Cloud Platform

Data Protector for Cloud Workloads supports the Google Compute Engine platform by using a VM called “Proxy VM”. The node invokes commands on the GCP to snapshot and attach temporary disks 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 means that you need to create a Compute Instance (Proxy VM) in each zone from which you want to secure instances.

The Data Protector for Cloud Workloads Server can be deployed anywhere, but keep in mind that Nodes need to be able to call the server over HTTP(S) on the port you have specified.

The GCP backup strategies allows you to exclude drives from the backup that you don’t need.

All backup destinations can be used, but keep in mind that you may be charged for transferring data between regions, zones, and external backup providers.

Data Protector for Cloud Workloads Node has access to instances only in the **projects where service account has access to**.

Data Protector for Cloud Workloads Node requires the **organization ID** and **service account key** to connect to the GCP account.

Typical use scenarios

There are several scenarios for GCP which may be suitable for your case:

- **Backup GCE to GCS** - in this case after dumping backup, Data Protector for Cloud Workloads can push them to the GCS bucket.
- **Backup GCE to your local backup provider** - if you already have a central enterprise backup solution, you may want to use it as a backup provider for GCE instances running in GCP.
- **Backup GCE to your other cloud provider** - If you're using multiple clouds, you also may consider storing data in S3 or Azure backup providers.

> **Note**: In all cases, depending on your target, you may be charged for data transfers.

Setup consideration

It is assumed that you have working experience with GCE to be able to deploy Data Protector for Cloud Workloads components. You also need to have an IAM user with permissions that allow you to deploy an instance and generate service account access key for Data Protector for Cloud Workloads.

Remember to use **CentOS 8** as a base image - both for the Server and Nodes. For a typical installation, we recommend 2 virtual processors and 8 GB of RAM. This means that **e2-standard-2** should cover general use cases.

Both Data Protector for Cloud Workloads components are assumed to be deployed without HA (more precisely, all the nodes or server will probably be in separate zones, and only need to communicate over HTTP). There is no requirement for multi-zone deployment for now. While the Node is stateless and can be lost without data loss, the Server needs DB to be protected. Data Protector for Cloud Workloads provides a built-in automatic DB backup mechanism, which can be used to protect backup metadata. Please refer to the Disaster recovery section for more details.

You add GCP as a Hypervisor Manager. You need to provide the organization ID and service account keys of a user that has required permissions.

Permissions

Cloud Resource Manager needs to be enabled in project where proxy vms are located.

Here are the IAM permissions that Data Protector for Cloud Workloads needs to have for backup/restore operations.

```
compute.disks.create
compute.disks.createSnapshot
compute.disks.delete
compute.disks.get
compute.disks.list
compute.disks.use
compute.disks.useReadOnly
compute.diskTypes.list
compute.globalOperations.get
compute.instances.attachDisk
compute.instances.create
compute.instances.delete
compute.instances.detachDisk
compute.instances.get
compute.instances.list
compute.instances.setMetadata
compute.instances.setServiceAccount
compute.instances.setTags
compute.machineTypes.list
compute.projects.get
compute.regions.list
compute.snapshots.create
compute.snapshots.delete
compute.snapshots.get
compute.snapshots.list
compute.snapshots.setLabels
compute.snapshots.useReadOnly
compute.subnetworks.list
compute.subnetworks.use
compute.subnetworks.useExternalIp
compute.zoneOperations.get
compute.zones.list
iam.serviceAccounts.actAs
resourcemanager.projects.get
```

Adding a hypervisor manager

To properly configure your GCP account, go to Data Protector for Cloud Workloads WebUI → Virtual Environments → Infrastructure → add hypervisor manager

Enter parameters such as:

- Organization ID
- Service Account Key

Note: If you want to backup instances in projects without organization provide 0 as the Organization ID.

Multi-zone configuration

To secure instances from multiple zones and regions, you need to register additional Data Protector for Cloud Workloads Nodes. For every zone you want to backup instances from, you need to create separate Node Configuration. Each Node Configuration needs to be assigned to correct Hypervisor Cluster (which reflects compute zone from GCP).

First synchronization task needs only one Node. After first scan, you can assign Node Configurations to Hypervisor Clusters, and run synchronization task once again to fetch all instances from other zones.

To assign Node Configuration to Hypervisor Cluster, go to Clusters list in Infrastructure tab. Next, click on selected Cluster to choose Node Configuration.

> **Note**: For backup and restore between zones nodes and node configurations should have access to the same backup destination.

Backup strategies

GCP supports two backup and restore strategies:

1. Disk attachment (full)
2. Disk attachment with changed block tracking (full/incremental)

In both cases, volume snapshots are kept in GCP only until the Data Protector for Cloud Workloads Export task is completed.

Restore

Data Protector for Cloud Workloads supports:

- restoring instances to different projects
- restoring instances to different regions and zones
- restoring instances to different machine types (that are available in selected regions)
- restoring instances to different subnetworks (that are available in selected project)
- restoring disks to different disk types

> **Note**: Data Protector for Cloud Workloads do not support backup/restore of VM with disk encrypted with an imported key. Currently, only VMs with key generated by GCP are supported.

Azure Cloud

Azure Cloud

Data Protector for Cloud Workloads supports the Azure Cloud platform by using a VM called “Proxy VM”. The node invokes commands on the Azure to snapshot and attach temporary disks 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 means that you need to create a Compute Instance (Proxy VM) in each zone and region without zone placement from which you want to secure instances. Remember that you need to install 1 Proxy VM in each location so that drives that the Node tries to attach are reachable.

The Data Protector for Cloud Workloads Server can be deployed anywhere, but keep in mind that Nodes need to be able to call the server over HTTP(S) on the port you have specified.

All backup destinations can be used, but keep in mind that you may be charged for transferring data between regions, zones, and external backup providers.

Data Protector for Cloud Workloads Node has access to instances only in the resource groups where service app has access to.

Setup considerations

It is assumed that you have working experience with Microsoft Azure to be able to deploy Data Protector for Cloud Workloads components. You also need an IAM user with the ability to change permissions on the resource group and creating new application registrations in Azure Active Directory.

Data Protector for Cloud Workloads Node requires the Application (client) ID, Client Secret, Subscription ID and Tenant ID to connect to the Microsoft Azure account.

Permissions

Data Protector for Cloud Workloads requires you to create a dedicated application registration in Azure Active Directory.

Register an application with the Microsoft identity platform ↗

For proper operation, the application must have access to the resource group where the VM with Data Protector for Cloud Workloads Node is located. To do this, add the app you created with Owner permissions to the Access Control list in the appropriate resource group.

Grant access to resource group ↗

Adding a hypervisor manager

To properly configure your Azure Cloud account, go to **Data Protector for Cloud Workloads WebUI → Virtual Environments → Infrastructure → Hypervisor Managers** and click **create** button.

Enter parameters such as:

- Client ID
- Tenant ID
- Client Secret

You can find these parameters in **Azure Active Directory → App registrations → Your App**. In the Overview tab you will find **Application (client) ID** and **Directory (tenant) ID**. In the **Certificates & secrets** tab you can generate a **Client Secret**.

- Subscription ID

Subscription ID can be found in the settings of your Azure subscription.

Multi-zone configuration

To secure instances from multiple zones and regions, you need to register additional Data Protector for Cloud Workloads Nodes. For every zone you want to backup instances from, you need to create separate Node Configuration. Each Node Configuration needs to be assigned to correct Hypervisor Cluster (which reflects compute zone from Azure).

First synchronization task needs only one Node. After first scan, you can assign Node Configurations to Hypervisor Clusters, and run synchronization task once again to fetch all instances from other zones.

To assign Node Configuration to Hypervisor Cluster, go to Clusters list in Infrastructure tab. Next, click on selected Cluster to choose Node Configuration.

> **Note**:
>
> - For backup and restore between zones nodes and node configurations should have access to the same backup destination.
> - In Azure, instances could be placed in Regions without specifying zone, for this use cases You need to assign Node configuration to Hypervisor Cluster with *(No Zone)*, for example Germany West Central (No Zone)

Backup strategies

Azure Cloud supports two backup and restore strategies:

1. Disk attachment to proxy VM (only full backup)
2. Disk attachment with changed block tracking (full and incremental backup)

In both cases, volume snapshots are kept in Azure only until the Data Protector for Cloud Workloads Export task is completed.

Access Keys

If no **access keys** are specified during the restore, the target instance will be restored with the same key as in the backup. If this key is no longer available, the only way to access the restored instance is to reset the password in the Azure Cloud Platform.

Containers

Containers

Data Protector for Cloud Workloads can back up data residing on Persistent Volumes in Red Hat OpenShift/Kubernetes environments or containers running on Proxmox VE (using its native backup mechanism).

- Kubernetes
- Red Hat OpenShift
- Proxmox VE

Kubernetes

Kubernetes

Data Protector for Cloud Workloads Node preparation

Data Protector for Cloud Workloads Node requires `kubectl` installed (you have to add Kubernetes repository to install `kubectl`) and `kubeconfig` with valid certificates (placed in `/home/user/.kube`) to connect to the Kubernetes cluster.

1. Check if your kubeconfig looks the same as below.

Example:

```
    current-context: admin-cluster.local
    kind: Config
    preferences: {}
    users:
    - name: admin-cluster.local
      user:
        client-certificate-data: <REDACTED>
        client-key-data: <REDACTED>
```

1. Copy configs to Data Protector for Cloud Workloads Node. (**Skip this and point 2 if you don't use Minikube**)
    - If you use Minikube, you can copy the following files to Data Protector for Cloud Workloads: `sudo cp /home/user/.kube/config /opt/vprotect/.kube/config sudo cp /home/user/.minikube/{ca.crt,client.crt,client.key} /opt/vprotect/.kube`
2. Modify the paths in `config` so they point to `/opt/vprotect/.kube` instead of `/home/user/.minikube`. Example:

```
- name: minikube
  user:
    client-certificate: /opt/vprotect/.kube/client.crt
    client-key: /opt/vprotect/.kube/client.key
```

1. Afterward, give permissions to the `vprotect` user:

```
chown -R vprotect:vprotect /opt/vprotect/.kube
```

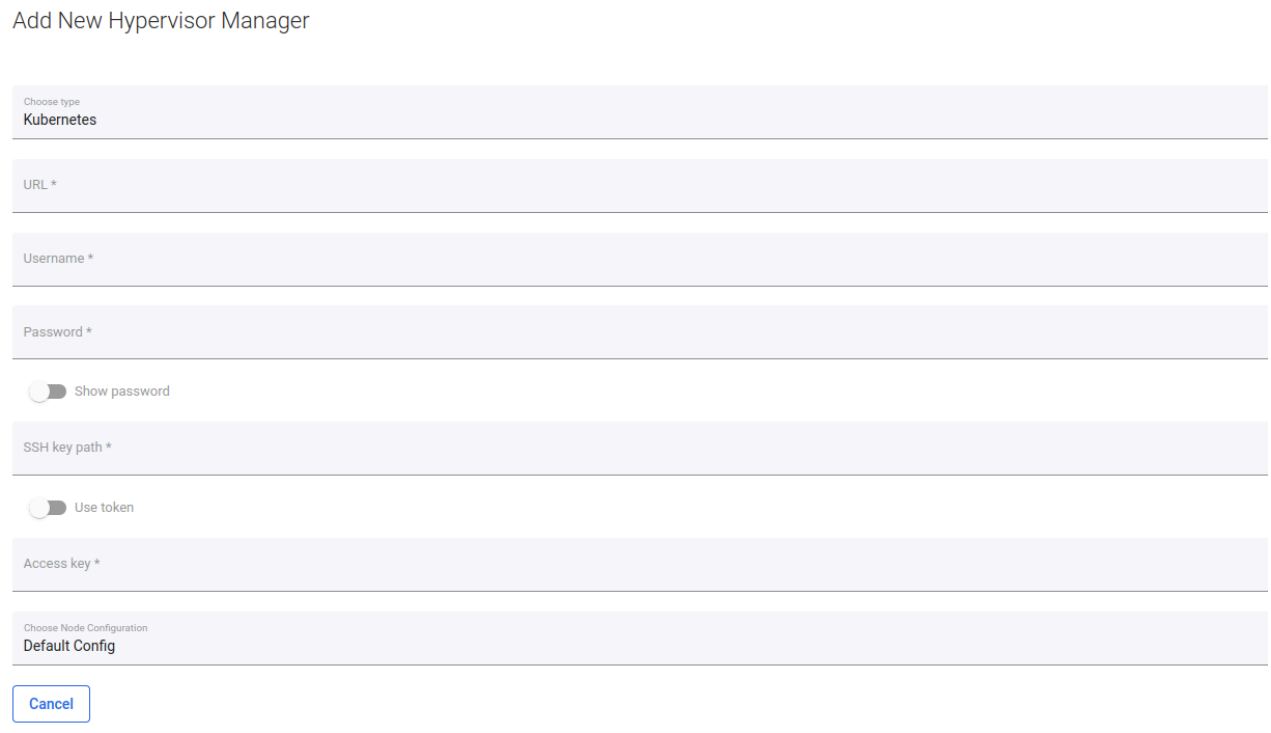

Kubernetes Nodes should appear in Data Protector for Cloud Workloads after indexing the cluster.

> **Note**: Provide the URL to the web console and SSH credentials to the master node when creating the OpenShift hypervisor manager in Data Protector for Cloud Workloads WebUI. You can also use SSH public key authentication. This is needed for Data Protector for Cloud Workloads to have access to your cluster deployments.

> **Note:** Valid SSH admin credentials should be provided **for every Kubernetes node** by the user (called *Hypervisor* in the Data Protector for Cloud Workloads

> WebUI). If Data Protector for Cloud Workloads is unable to execute docker commands on the Kubernetes node, it means that it is logged as a user lacking admin privileges. Make sure you added your user to sudo/wheel group (so it can execute commands with `sudo`).

> **Note:** If you want to use Ceph you must provide ceph keyring and configuration. Ceph requires ceph-common and rbd-nbd packages installed.

Persistent volumes restore/backup

There are two ways of restoring the volume content.

1. The user should deploy an automatic provisioner which will create persistent volumes dynamically. If Helm is installed, the setup is quick and easy https://github.com/helm/charts/tree/master/stable/nfs-server-provisioner↗.
2. The user should manually create a pool of volumes. Data Protector for Cloud Workloads will pick one of the available volumes with proper storage class to restore the content.

Limitations

- currently, we support only backups of Deployments/DeploymentConfigs (persistent volumes and metadata)
- **all deployment's pods will be paused during the backup operation** - this is required to achieve consistent backup data
- for a successful backup, every object used by the Deployment/DeploymentConfig should have an `app` label assigned appropriately
- a storage class must be defined in the Kubernetes environment for backup and restore operations to function properly

Red Hat OpenShift

Red Hat OpenShift

Data Protector for Cloud Workloads supports backup for OpenShift using OADP in backup process.

> 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:

- **URL** - URL of the Openshift API e.g. `api.your.cluster.local:6443`
- **Username** - login of user with 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.

> **Note**:
>
> - Metadata is backed up using OADP operator. All of the OADP resources required during backups and restores will be automatically created, such as: DataProtectionApplication, BackupStorageLocation, VolumeSnapshotLocation
> - Persistent volumes are primarily backed up using volume snapshots providing crash consistency.
> - Whether specific persistent volume can be snapshotted is determined by searching for existing VolumeSnapshotClass with a driver corresponding to the storage class of the persistent volume. Which means that if the storage class of the persistent volume does not use CSI driver or the proper volume snapshot class is not already present, export of this volume will be performed without snapshotting mechanism. Block volumes can be exported only with volume snapshots.

Persistent volumes restore

There are two ways of restoring the volume content.

- The user should deploy an automatic provisioner which will create persistent volumes dynamically. If Helm is installed, the setup is quick and easy https://github.com/helm/charts/tree/master/stable/nfs-server-provisioner↗.
- The user should manually create a pool of volumes. Data Protector for Cloud Workloads will pick one of the available volumes to restore the content.

Limitations

- Only backups of Deployments, DeploymentConfigs and StatefulSets (persistent volumes and metadata) are currently supported.
- If Storage Class used by volume do not support snapshots, all deployment pods will be paused during the backup operation - this is required to achieve consistent backup data.
- For a successful backup, Deployment/DeploymentConfig/StatefulSet should have an **app** label assigned appropriately.

Proxmox VE

Proxmox VE

Proxmox 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: VZDump)

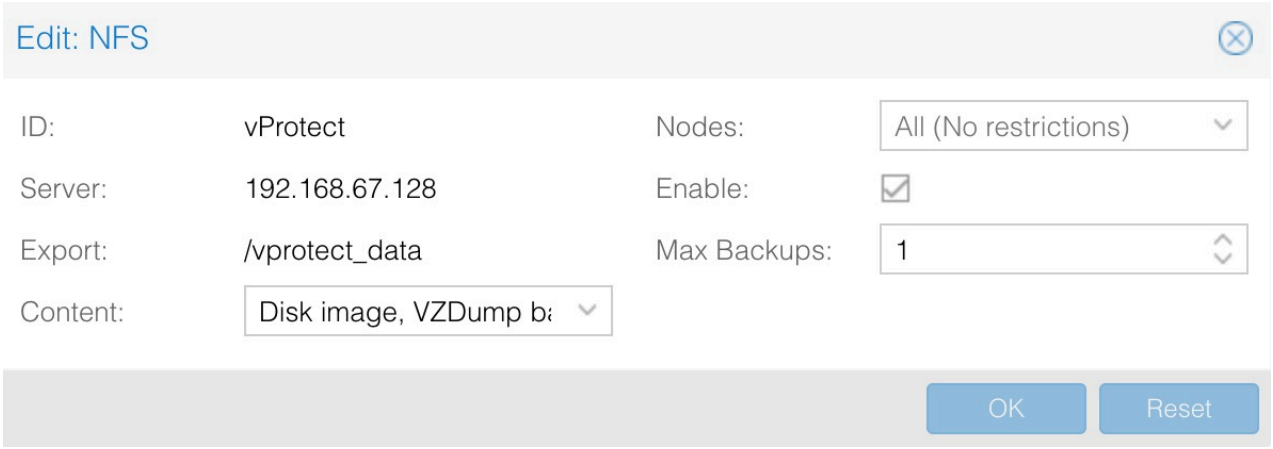

- The export share must be set to use the UID and GID of the `vprotect` user.
- Example export configuration in `/etc/exports` to the selected hypervisor in the RHV cluster:

  ```
  /vprotect_data
  PROXMOX_HOSTS(fsid=6,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)
  ```

- 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 Proxmox environment

- Backups must be restored to the export path (the node automatically changes the names to the original paths that are recognized by Proxmox).

- The name for storage must be provided later in the node configuration (`Hypervisor -> Proxmox`)section.

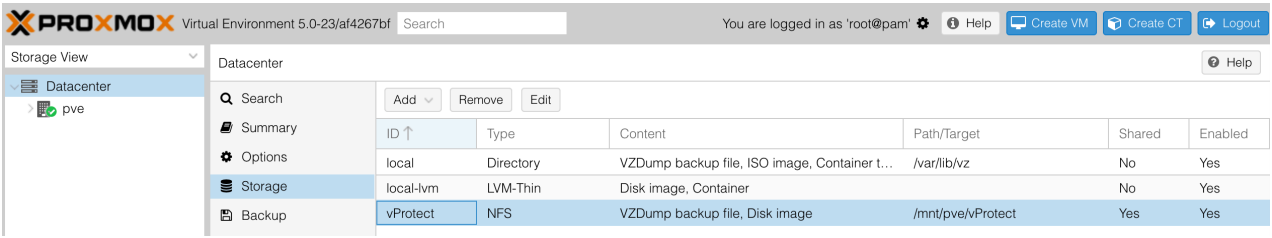

File-level restore support for VMA images

1. Prepare the VMA extractor on Data Protector for Cloud Workloads Node - you have 2 options:

   - Build a VMA extractor like this (requires Internet on the **node**):

     ```
     cd /opt/vprotect/scripts/vma
     ./setup_vma.sh
     ```

2.

   - Download the VMA extractor from the Micro Focus download page and install it.

     ```
     cd /opt/vprotect/scripts/vma
     ./setup_vma.sh PATH_TO_VMA_ARCHIVE
     ```

Public key authentication

The details are described in the SSH public key authentication section.

Backup & Restore

Backup & Restore

This section is a simple step-by-step guide on how to use Data Protector for Cloud Workloads to protect your Virtual Environments.

Manual backup

> **Note:** to perform a manual backup, a target VE instance must be assigned to Backup SLA

Search for VE instance you want to backup and click **Backup** the button next to it.

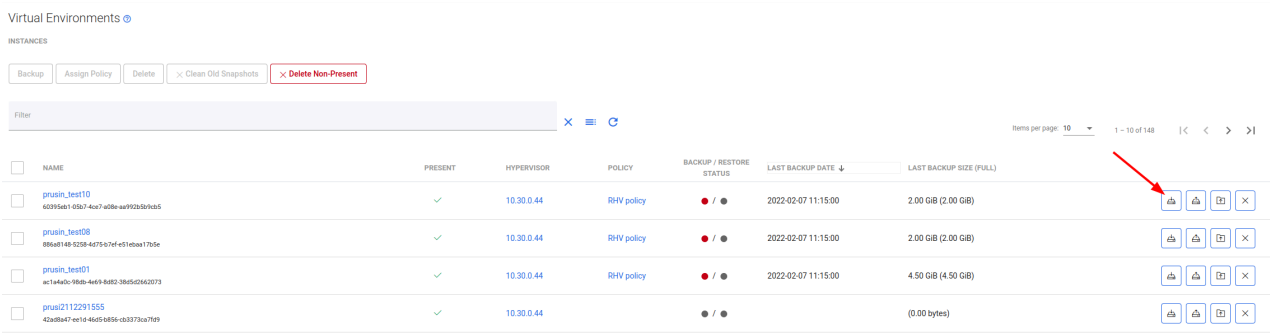

In the next window, choose backup type and click **Backup** button.

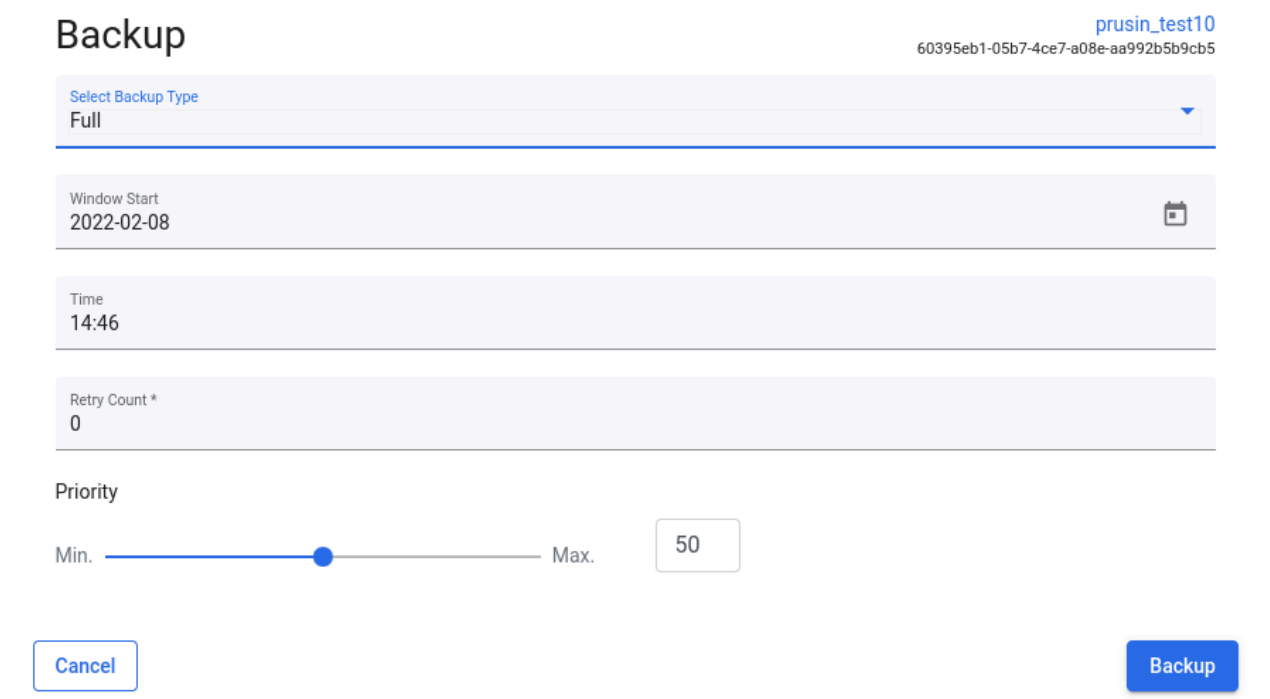

You can track the progress of the task in **tasks console**

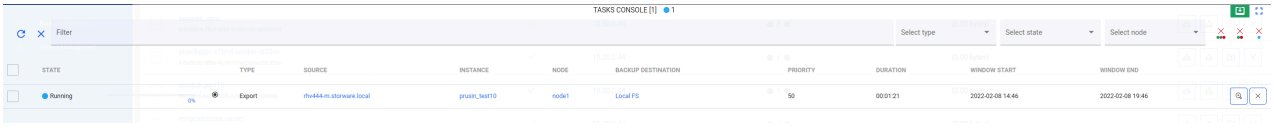

Manual restore

Search for VE instance you want to restore and click **Restore** the button next to it.

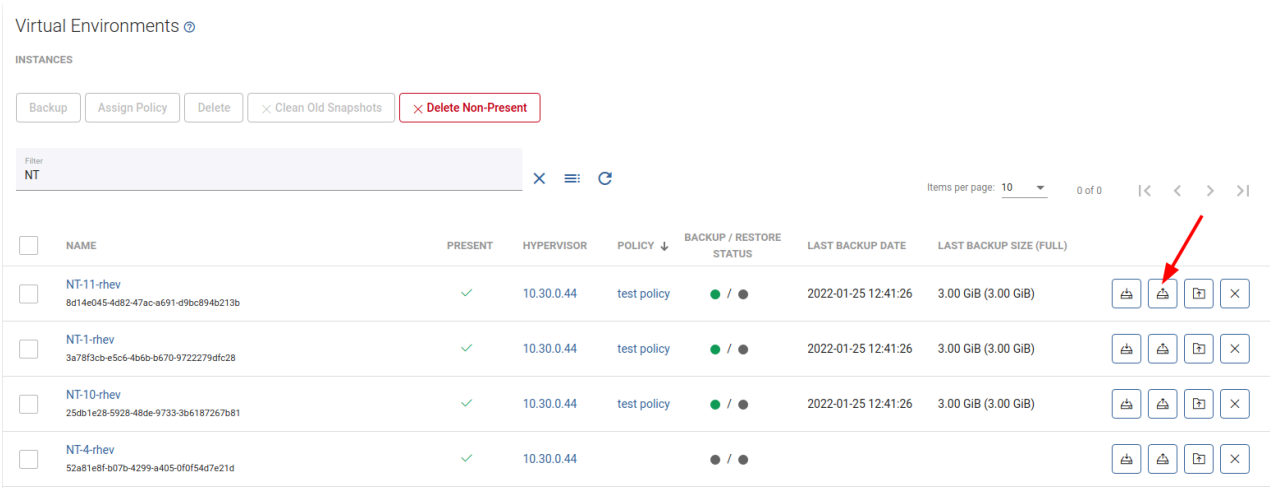

In the next window, choose if you want to restore to the file system or directly to the hypervisor (or hypervisor manager). Enter the rest parameters according to your needs and click **Restore** the button. The available parameters may differ depending on the hypervisor.

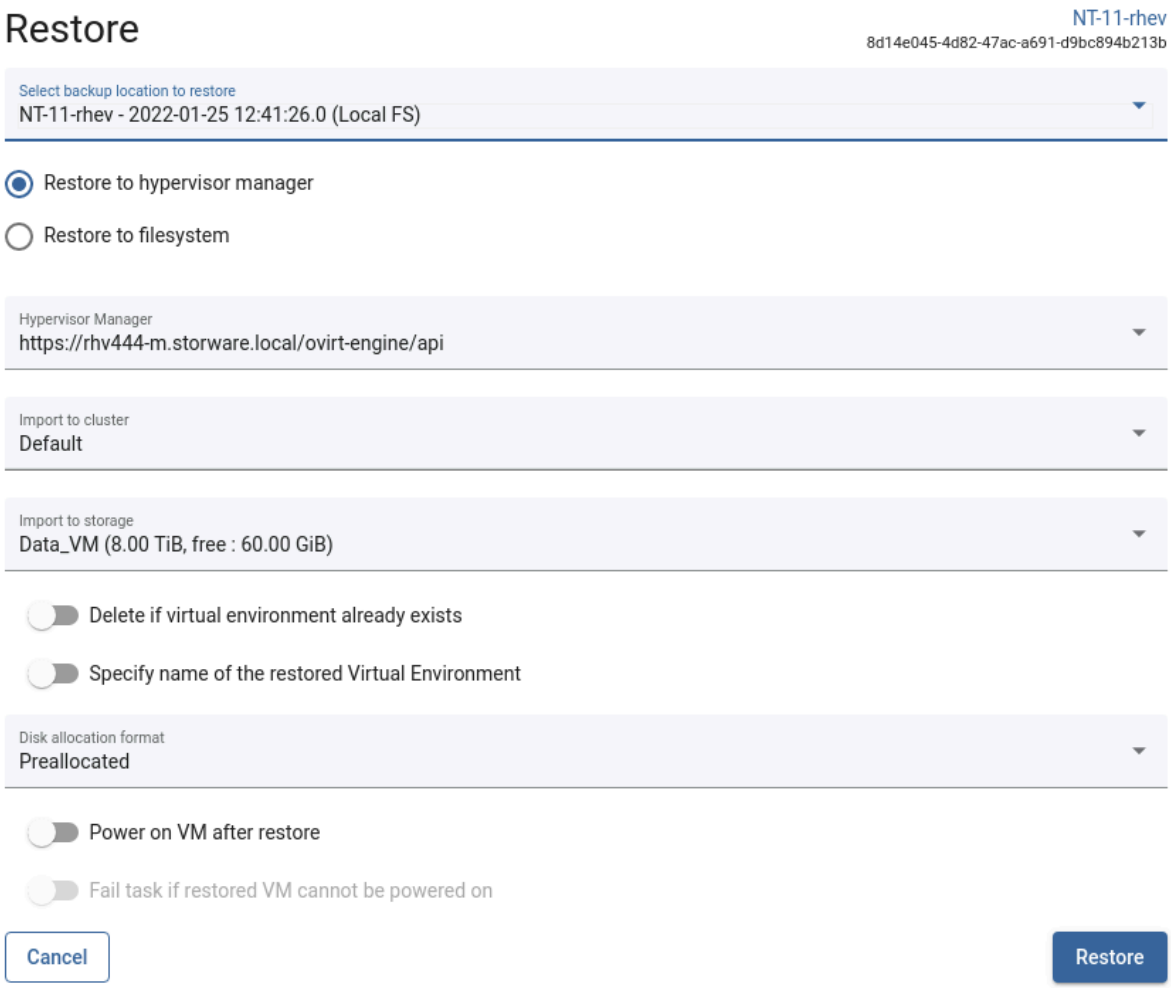

> **Note:** Staging cleanup task starts even if backup or restore process is failed or cancelled.

> **Note:** More information about Virtual Environments you can find in the administration chapter.

Protecting Microsoft 365

Protecting Microsoft 365

Data Protector for Cloud Workloads Server allows you to protect multiple Microsoft 365 organizations within a single server instance. Before you add the Microsoft 365 organization to Data Protector for Cloud Workloads Server, you have to configure all required permissions of your application in the Microsoft 365 service.

Go to the Microsoft 365 organization management section to learn how to configure Microsoft 365 organization permissions to be able to add them to Data Protector for Cloud Workloads Server.

After that, you can go to Backup & Restore to start protecting Yours Microsoft 365 organizations.

Microsoft 365 organization management

Microsoft 365 organization management

You can add a Microsoft 365 organization manually or using the Setup assistant.

To add the organization log into Data Protector for Cloud Workloads go to the **Cloud** menu and select **Service Providers**. Now you can choose one of the options: **Create** to add manually or use **Setup Assistant**.

How to add Microsoft 365 organization manually

How to add Microsoft 365 organization using the Setup Assistan

> **Note:** We recommend using the setup assistant to add Microsoft 365 organizations

Configure Microsoft 365 access

Configure Microsoft 365 access

Before you start to configure Backup SLAs, Data Protector for Cloud Workloads Server has to get access to your Microsoft 365 organization configuration first.

Access to data is performed via an application configured in your Microsoft 365 organization.

> **Note:**
>
> - You can skip this step if you want to add your Microsoft 365 organization to Data Protector for Cloud Workloads Server using the Setup Assistant. If not, proceed with the next steps.
> - Data Protector for Cloud Workloads does not store your Microsoft 365 administrative id and password.

Manually registering an application with Azure Active Directory

A new Microsoft 365 application has to be registered and configured in Azure Active Directory. When it's done, in the next step you can add the application (organization) to Data Protector for Cloud Workloads Server.

The instruction is as follow:

1. Go to the Azure portal (https://portal.azure.com/ ↗) page and sign in to your Microsoft account by using your Microsoft 365 administrative user ID and password.
2. In the home view, go to **Manage Azure Active Directory** (click the **View** button).
3. To open the Azure Active Directory admin center, in the left pane, click the ellipsis to expand the Show all menu, and then click **Admin centers** > **Azure**

Active Directory.

4. In the tenant dashboard menu, click **App registrations** and then click **New registration**.
5. Specify a user-facing name for the Microsoft 365 application, on the **Register an application** page by entering a name in the **Name** field.
6. Use the default options for the remaining fields, and click **Register**. The app registration is set up with the user-facing name that you entered.
7. To obtain the application (**client**) ID, and directory (**tenant**) ID string, go to **Azure Active Directory > tenant - App registrations > Owned applications**. Click the application name, copy the application ID string and directory ID. These strings will be required later when you register the Microsoft 365 service on Data Protector for Cloud Workloads Server.
8. To create a client secret for this application ID, click **Certificates & secrets** > **New client secret**.
9. On the "**Add a client secret**" pane, enter any user name in the **Description** field, and click **Add**. A client secret is generated, and the value is then displayed in the Client secrets pane.
10. Copy the client secret to the clipboard by using the copy icon next to the Client secret value field. This character string is also used for registration with Data Protector for Cloud Workloads Server.
11. To add permissions for this application ID, click **API permissions** > **Add permission**.
12. Specify permissions for each API in the following table by taking the following actions. Select the API name, for example, Azure Active Directory Graph.

API	Permission name	Permission type
Azure Active Directory Graph	Calendars.ReadWrite	Application
Microsoft Graph	Channel.Create	Application
Microsoft Graph	Channel.ReadBasic.All	Application
Microsoft Graph	ChannelMember.Read.All	Application
Microsoft Graph	ChannelMember.ReadWrite.All	Application

Microsoft Graph	ChannelMessage.Read.All	Application
Microsoft Graph	Chat.Create	Application
Microsoft Graph	Chat.Read.All	Application
Microsoft Graph	Chat.ReadBasic.All	Application
Microsoft Graph	Chat.ReadWrite.All	Application
Microsoft Graph	ChatMember.Read.All	Application
Microsoft Graph	ChatMember.ReadWrite.All	Application
Microsoft Graph	Contacts.ReadWrite	Application
Microsoft Graph	Directory.ReadWrite.All	Application
Microsoft Graph	Files.ReadWrite.All	Application
Microsoft Graph	Group.Create	Application
Microsoft Graph	Group.ReadWrite.All	Application
Microsoft Graph	GroupMember.Read.All	Application
Microsoft Graph	GroupMember.ReadWrite.All	Application
Microsoft Graph	Mail.ReadWrite	Application
Microsoft Graph	MailboxSettings.Read	Application
Microsoft Graph	MailboxSettings.ReadWrite	Application
Microsoft Graph	Member.Read.Hidden	Application
Microsoft Graph	Sites.FullControl.All	Application
Microsoft Graph	Sites.Manage.All	Application
Microsoft Graph	Sites.Read.All	Application
Microsoft Graph	Sites.ReadWrite.All	Application
Microsoft Graph	Team.Create	Application

Microsoft Graph	Team.ReadBasic.All	Application
Microsoft Graph	TeamMember.Read.All	Application
Microsoft Graph	TeamMember.ReadWrite.All	Application
Microsoft Graph	TeamMember.ReadWriteNonOwnerRole.All	Application
Microsoft Graph	TeamSettings.ReadWrite.All	Application
Microsoft Graph	TeamsTab.Create	Application
Microsoft Graph	TeamsTab.ReadWrite.All	Application
Microsoft Graph	TeamsTab.ReadWriteForChat.All	Application
Microsoft Graph	TeamsTab.ReadWriteForTeam.All	Application
Microsoft Graph	TeamsTab.ReadWriteForUser.All	Application
Microsoft Graph	User.Read.All	Application
Microsoft Graph	User.ReadWrite.All	Application
Office 365 Exchange Online	full_access_as_app	Application
SharePoint	Sites.FullControl.All	Application
SharePoint	Sites.Manage.All	Application
SharePoint	Sites.Read.All	Application
SharePoint	Sites.ReadWrite.All	Application
SharePoint	User.ReadWrite.All	Application

13. To set the permission "**full_access_as_app**" in the Office 365 Exchange Online API, click "**Add a permission**" option and in the "**Request API permissions**" window go to "**APIs my organization uses**" and search for "**Office 365 Exchange Online**", then select "**Application permissions**" and check "**full_access_as_app**" from "**Other permissions**".

14. To save the selected permissions, click **Grant admin consent for <your organization name>**.
15. Since you're granting tenant scoped permissions this granting can only be done via the **appinv.aspx** page on the tenant administration site. You can reach this site by typing the address: **https://tenantName-admin.sharepoint.com/_layouts/15/appinv.aspx**. (replace the **tenantName** with your tenant name). Once the page is loaded, do as follow:
    a. Enter your App Id (client ID) and click the **Lookup** button.
    b. Enter the App Domain name.
    c. In the "**App's Permission Request XML**" window enter the following lines:

```
<AppPermissionRequests AllowAppOnlyPolicy="true">
    <AppPermissionRequest Scope="http://sharepoint/content/tenant"
     Right="FullControl" />
 </AppPermissionRequests>
```

16. When you click on the **Create** button you'll be presented with a permission consent dialog. Press the **Trust It** button to grant the permissions.

17. Open Powershell command prompt and execute the command:

```
Install-Module -Name Microsoft.Online.SharePoint.Powershell
```

or download and install the module directly from this site ↗

then

```
Connect-SPOService https://tenantName-admin.sharepoint.com
Set-SPOTenant -LegacyAuthProtocolsEnabled $True
Set-SPOTenant -DisableCustomAppAuthentication $false
```

It enables the **LegacyAuthProtocolsEnabled** setting.

Add Microsoft 365 organization manually

Add Microsoft 365 organization manually

To add your Microsoft 365 organization to Data Protector for Cloud Workloads Server, do as follow:

1. Log into Data Protector for Cloud Workloads WebUi and select tab **Cloud**.
2. Go to the **Service Providers** menu.
3. Click the **Create** button.
4. In the **Microsoft 365 Configuration** window you have to enter information as follow:
    a. Tenant ID
    b. Client ID
    c. Client secret
5. Chose node configuration which will be assigned to this organization
6. Save your settings by clicking the **Save** button.
7. Click the **Synchronize** button to synchronize users/sites/teams from your organization with Data Protector for Cloud Workloads Server.
8. Now you can go to the **Instances** menu and see synchronize results.

Go to the Add Microsoft 365 organization using the Setup Assistant chapter to learn how to add Microsoft 365 organization using dedicated Setup Assistant.

> **Note:** Go to the Account auto-synchronization to learn how to synchronize Microsoft 365 accounts.

Add Microsoft 365 organization using the Setup Assistant

Add Microsoft 365 organization using the Setup Assistant

To add your Microsoft 365 organization to Data Protector for Cloud Workloads Server using Setup Assistant, do as follow:

1. Go to the **Cloud** menu → **Service Providers** → click the **Setup Assistant** button.
2. In the **Wizard** window read the information and click the **Next** button to go to the next step.
3. Copy the authorization code → click the link **https://microsoft.com/devicelogin** and enter it → sign in to your Microsoft 365 organization as a user with administrator's rights. Now close the tab and go back to **Setup Assistant**.
4. Provide number of application and click **Create** button. The application tenant is created. Click the **Next** button to go to the next step.
5. Follow the instructions in **Setup Assistant**. If the **Lookup** button is not working and fields are not automatically populated the user needs to put there the following:
   - for the Title field: DP4CW
   - for the App domain field: localhost
6. In the next step, click **Grant permissions** button. You will be redirected to the the page where you need to accept permissions.
7. Execute provided PowerShell commands and click **Save** button.
8. The configuration is now complete. Save your settings by clicking the **Save** button.

> We recommend creating 1 application for every 500 users.

Required permissions

1. Since you're granting tenant scoped permissions this granting can only be done via the **appinv.aspx** page on the tenant administration site. You can reach this site by typing the address: **https://tenantName-admin.sharepoint.com/_layouts/15/appinv.aspx**. (replace the **tenantName** with your tenant name). Once the page is loaded, do as follow:
2. Enter your App Id (client ID) and click the **Lookup** button.
3. Enter the App Domain name.
4. In the "**App's Permission Request XML**" window enter the following lines:

   ```
   <AppPermissionRequests AllowAppOnlyPolicy="true">
       <AppPermissionRequest Scope="http://sharepoint/content/tenant"
        Right="FullControl" />
    </AppPermissionRequests>
   ```

5. When you click on the **Create** button you'll be presented with a permission consent dialog. Press the **Trust It** button to grant the permissions.
6. Open Powershell command prompt and execute the command:

   ```
   Install-Module -Name Microsoft.Online.SharePoint.Powershell
   ```

   or download and install the module directly from this site ↗

   then

   ```
   Connect-SPOService https://tenantName-admin.sharepoint.com
   Set-SPOTenant -LegacyAuthProtocolsEnabled $True
   Set-SPOTenant -DisableCustomAppAuthentication $false
   ```

   It enables the **LegacyAuthProtocolsEnabled** setting.

> **Note:** Go to the Account auto-synchronization to learn how to synchronize Microsoft 365 accounts.

Account auto-synchronization

Account auto-synchronization

This feature allows Microsoft 365 service accounts to be automatically synchronized with the Data Protector for Cloud Workloads Server. It means, that every newly created Microsoft 365 account will be synchronized with the server.

To configure this feature on the server do as follow:

1. Log onto Data Protector for Cloud Workloads WebUI
2. Go to the **Settings** menu
3. Select the **Global Settings** tab
4. Set the **Periodic inventory synchronization** interval (by default is 8 hours)
5. Click the **Save** button

From now on all new Users, SharePoint sites, and Teams added to Microsoft 365 service are also available in the Protection menu on Data Protector for Cloud Workloads Server after automatic inventory synchronization.

Backup & Restore

Backup & Restore

This section is a simple step-by-step guide on how to use Data Protector for Cloud Workloads to protect your Microsoft 365.

Backup and restore of Microsoft 365

Manual backup

> **Note:** to perform a manual backup, a target account, site or team instance must be assigned to Backup SLA

Search for one of the cloud instances you want to backup and click **Backup** the button next to it.

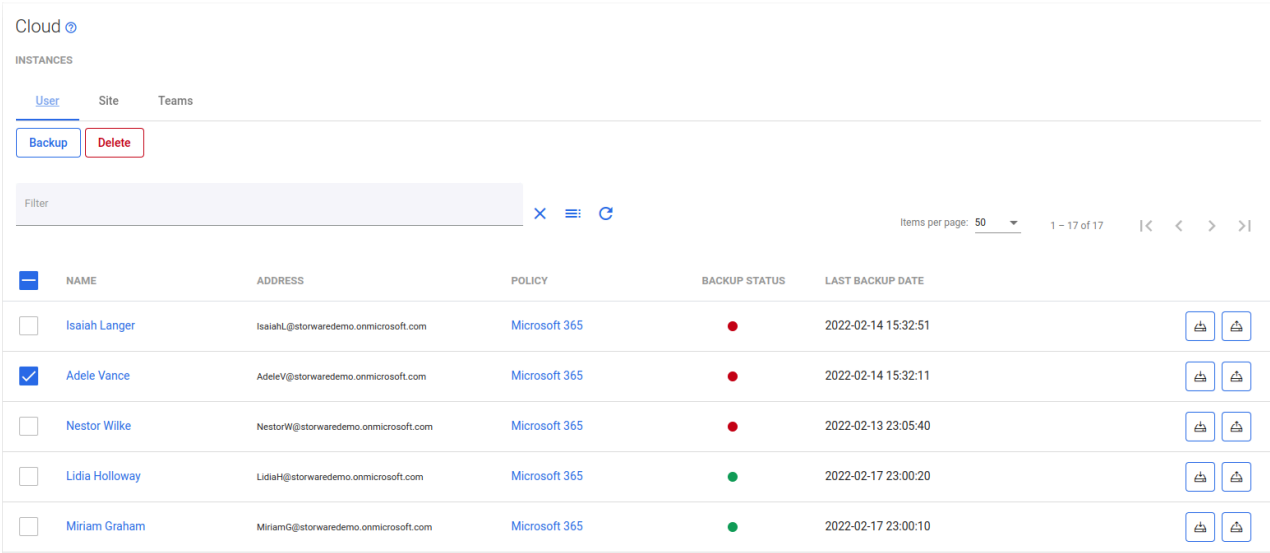

In the next window, choose backup type and what features you want to protect. If you want a backup of just one or more features use **manual selection** or use **according to the policy** to backup all features selected in Backup SLA then click **Backup** button.

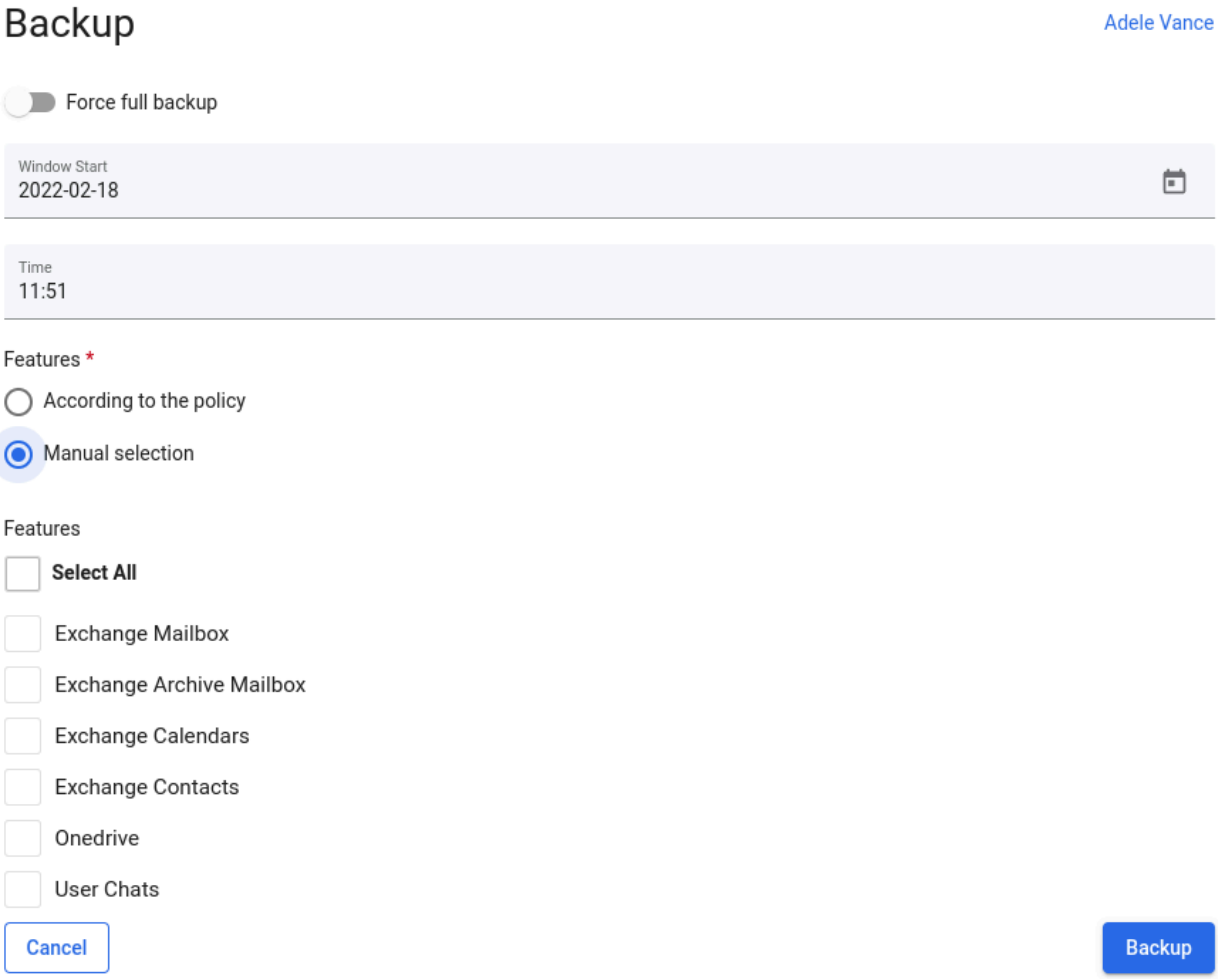

You can track the progress of the task in **tasks console**

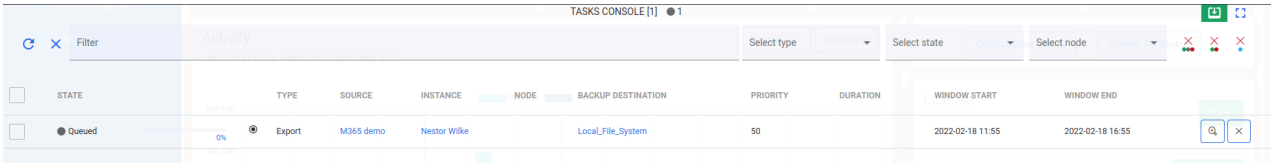

Manual restore

Search for cloud instance you want to restore and click **Restore** the button next to it.

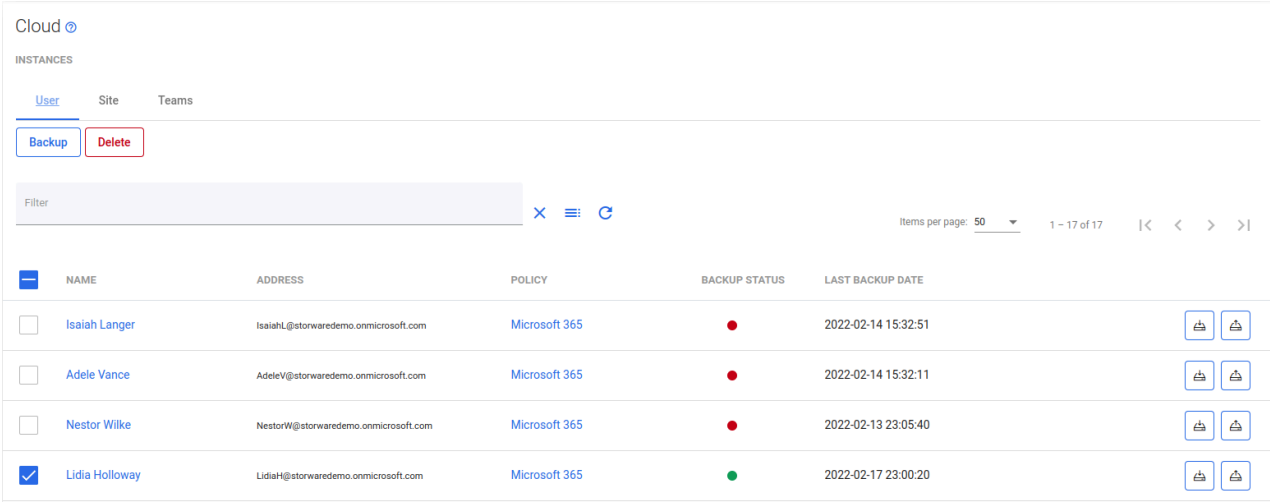

In the next window, choose if you want to restore **according to the policy** or use **manual selection** to restore one specific feature directly to the restored account and click **Restore** the button.

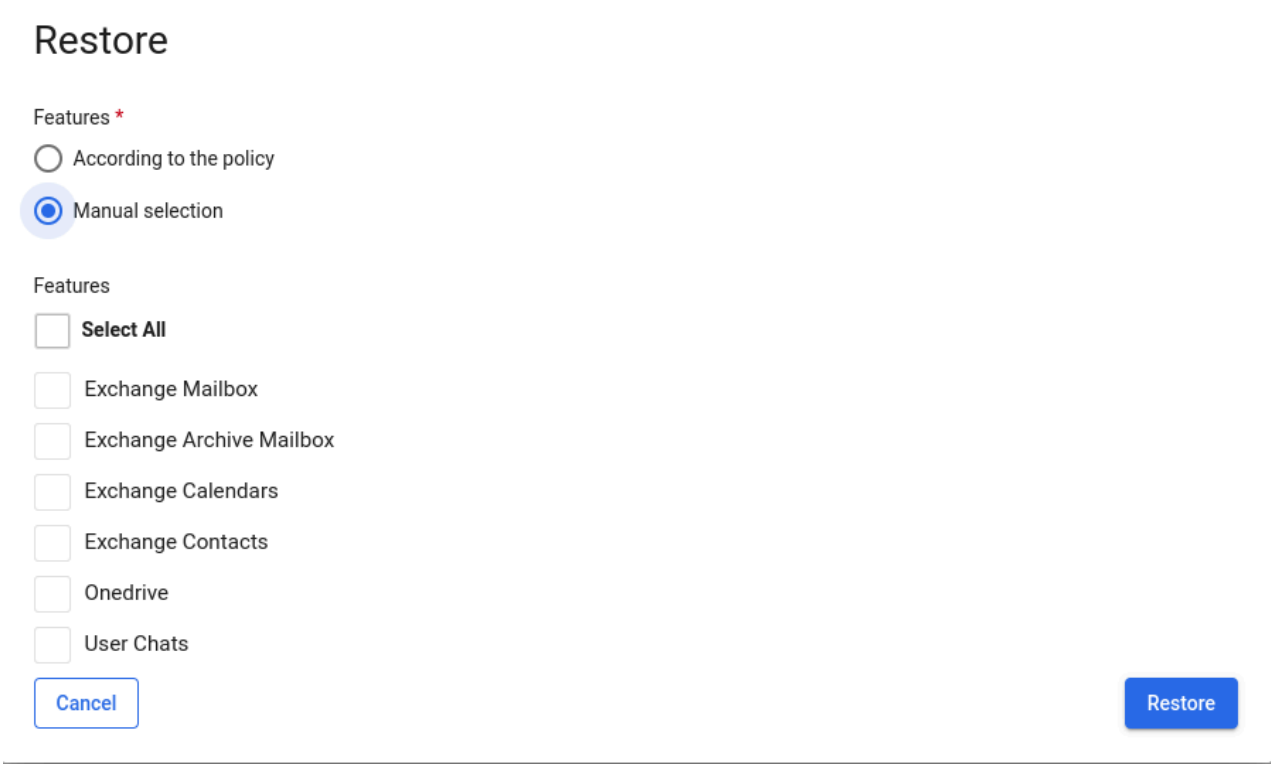

> **Note:** More information about Cloud you can find in the administration chapter.

Suppoted Sharepoint templates and limitations

Suppoted Sharepoint templates and limitations

Templates and apps supported for backup in Microsoft SharePoint Online

Supported site type:

- Teams site
- Communication site

Supported templates - Team site

	Event planning	Project management	Retail management team	Store collaboration	Training and courses	Tra an dev nt t
List	yes	yes	yes	yes	yes	yes
Document library	yes	yes	yes	yes	yes	yes
Page	yes	yes	yes	yes	yes	yes
Space	no	no	no	no	no	no
News post	yes (it is page)	yes (it is page)	yes (it is page)	yes (it is page)	yes (it is page)	yes pag
News link	yes (it is page)	yes (it is page)	yes (it is page)	yes (it is page)	yes (it is page)	yes pag

Plans	no	no	no	no	no	no
Event list	yes	yes	yes	yes	yes	yes
App	no	no	no	no	no	no
Link	no	no	no	no	no	no

Supported templates - Communication site

	Crisis management	Leadership connection	Learning central	Showcase	Topic	Vo ce
List	Yes	Yes	Yes	Yes	Yes	Yes
Document library	Yes	Yes	Yes	Yes	Yes	Yes
Page	Yes	Yes	Yes	Yes	Yes	Yes
Space	No	No	No	No	No	No
News post	Yes	Yes	Yes	Yes	Yes	Yes
New link	No	No	No	No	No	No
Event list	Yes	Yes	Yes	Yes	Yes	Yes
App	No	No	No	No	No	No

Supported elements on the site

- List
- List template
  - Issue tracker
  - Employee onboarding
  - Event itinerary
  - Asset manager

    - Recruitment tracker
    - Travel requests
    - Work progress tracker
    - Content scheduler
    - Gift ideas
    - Expense tracker
    - Recipe tracker
    - Blank
    - Progress tracker
    - Inventory list
- Document library
- Page
- News post
- Event list

Unsupported templates:

- Communication site template - Department
- Communication site template - New employee onboarding

Unsupported elements on templates:

- Space
- Plans
- App
- Link - restore
- List - template - Playlist

Limitations

- Sometimes, after restoring the site, it may not be set as the original home page.

- Theme and site template are not backed up and restored on subsites.
- Sometimes it is not possible to set the owner of the restored site.
- The order of items in the lists may be different after the restore.
- The menu items of the site after restore may have a different order.
- Hidden list cannot be restored.
- After restoring the pages on which the lists were placed, the connection between the page and the list is not restored. You need to re-add the link to the restored list in the webpart.
- Attachments are stored in base64 format. This encoding causes an overhead of 33–37%. This should be taken into account when restoring emails with attachments as large attachments may exceed the maximum message size (the estimated maximum size of the attachment that can be restored is 109 MB assuming the maximum message size is set to 150 MB).
- Restoring recurring event will only restore it as single event. To restore all occurrences, enable `Overwrite existing files` option.

Teams chat messages

Data Protector for Cloud Workloads is using Evaluation mode as the default model.

Payment models and licensing requirements for Microsoft Teams APIs - Microsoft Graph ↗

There is 500 messages limit (per month per tenant per app). If this model is not sufficient for your organization, you should enable metered APIs and services in Microsoft Graph:

Enable metered APIs and services in Microsoft Graph - Microsoft Graph ↗

> Enabling metered APIs and services in Microsoft Graph, will be billed for any metered charges. This association also allows you to use Azure Cost Management + Billing to understand and manage the costs of the application.
>
> Cost Management + Billing - Microsoft Cost Management ↗

Following are not supported

- Marking messages as read or not read is not supported. Restored messages are always marked as read.
- OneNote files are not supported.
- OneNote tabs are not supported.
- Sharepoint comments are not supported.
- List comments are not supported.
- Calendar Groups are not supported.
- Site comments and likes are not supported.
- Messages for a private channel are not supported.
- Wiki content is not supported.
- SharePoint site collection recycle bin is not supported
- External SharePoint lists are not supported.
- Categories in contacts are not supported.

Status Recovery Limitations

- If a document/item was in `Check Out` state when the backup was created, item's last version will not be restored to the target SharePoint and will be available for viewing only. Previous versions (if any) will be restored.
- If the `Declare this item as a record` action was originally applied to list item, the relevant status will not be preserved. Instead, the restored item status will be set in accordance with the target list/library content approval workflow.
- Original status `On Hold` will not be restored.
- Versioning settings of SharePoint lists are not preserved during restore.
- Restoring Generic List and Pages Library may fail with the "No content type 'XXX' found in web YYY" error.
- The `Created By` field of restored documents is updated with the account performing restore.
- Some `Rating` Settings of Discussion lists values are not restored.

Protecting Applications

Protecting Applications

This section describes Data Protector for Cloud Workloads **generic mechanism** for multiple scenarios where VM-level backup may not be enough.

With Data Protector for Cloud Workloads you can prepare a custom script or invoke any backup command that produces backup artifacts (or just initiates an external backup process) on a remote host and stores backups to your backup provider.

With Application the 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

Data Protector for Cloud Workloads internal DB backup mechanism also uses Application backup. The mechanism uses predefined artifacts that simply need to be configured according to your needs.

In the Application section, you can read how the **generic mechanism** works and know what application Data Protector for Cloud Workloads can protect.

In the Backup & Restore section, you can find information about using Data Protector for Cloud Workloads to protect your application.

Applications

Applications

Data Protector for Cloud Workloads can protect Applications and others likes:

- Relax and Recover - ReaR
- Git
- oVirt/RHV/OLVM
- Kubernetes/OpenShift etcd

Main concepts

There are 2 main concepts that Data Protector for Cloud Workloads uses to execute backups:

- Command Execution Configuration
- Application Definition

Command Execution Configuration

This describes **how** to perform a backup operation. That is how to execute a command that produces a backup artefact which Data Protector for Cloud Workloads later stores in a backup provider. Multiple Application definitions share Command Execution Configuration but with different parameter values.

Command Execution Configuration properties come in several sections:

1. **General:**
   - **Name** - Name of your configuration
   - **Execution type**:
     - Node - execute this command directly on the node

    - Remote SSH - execute this command over SSH using credentials provided in the Application definition
- **Timeout** - fail execution if a command doesn't complete within the time given
    - if you think that your backup should take longer, increase this value
    - this timeout is for whole command execution - if you have several steps in your script and you need additional timeouts for these steps - add them to your script

2. **Command arguments:**
    - add arguments that contain spaces as separate arguments
    - the first argument is the path to your executable
    - make sure this command is accessible on the remote host, and Data Protector for Cloud Workloads credentials will suffice to execute it
    - remote commands (over SSH) will invoke shell so you can use bash-style expressions (built-in commands such as `echo`, environmental variables or redirection) within the command argument
    - commands executed on the node are executed natively by OS, so if you want to use bash-style expressions (built-in commands such as `echo`, environmental variables or redirection) you need to split your command at least into 3 arguments: `/bin/bash`, `-c` and `your command > with some redirection`
3. **Data export:**
    - **Export data** - when enabled, Data Protector for Cloud Workloads will expect artifacts to be collected as a result of a command
    - **Source type:**
        - FILE - result will be a file, directory or path with `*` wildcard
        - STREAM - output of your command
    - **Source path:**
        - path to your artefacts that need to be collected
        - file, directory or path with `*` wildcard - more than 1 file on the source will result in files being stored as a single tar archive
    - **Remove files after export:**
        - if artefacts (files or source directory) need to be removed once exported

    - be careful when providing a path in the source directory, the whole directory will be removed when this setting is enabled

4. **Applications:**
    - select which applications will use this command execution config
5. **Parameters:**
    - this section allows you to define the parameters that will be expected to be entered in each application definition
    - each parameter will eventually become an environment variable in the application definition
    - each parameter has several properties
        - Name - Name of the resulting environmental variable
        - User-friendly hint - a hint what this parameter is to be shown later in the application definition
        - Default value - the default value, entered during initialization in the application definition form
        - Show in UI - if the value should be shown as dotted or not - useful for passwords
        - Obligatory - if we expect that its value should always be provided in the application definition form
6. **Error handling**
    - Standard error output stream handling (when non-empty):
        - Don't ignore it - will fail if anything is in the standard error output
        - Ignore without warning - will ignore it silently
        - Ignore with a warning - will ignore it but a warning indicator in the backup history will contain this output
    - Ignored Exit Codes:
        - error codes that should be ignored and not treated by Data Protector for Cloud Workloads as errors
        - by default, only 0 is assumed as a success

Application Definition

Once you have your command execution configuration defined (or you choose to use the predefined ones provided with Data Protector for Cloud Workloads), you should define the instances of your application.

There are a few parameters for application definition that come in several sections:

1. **General:**
    - Name - Name of your application instance
    - Choose node - which node is going to execute this command
    - Backup policy - optionally set policy for scheduled backups
    - Command execution configuration
        - configuration of your command used for this application
        - **Note:** when you create a definition for the first time, you select a configuration and click **Save** - you will be redirected to the Settings tab for additional details
2. **Environment variables**
    - shown only when the definition has been saved on the Settings tab
    - defines a list of environment variables that will be passed to your command/script during its invocation
    - parameters from the command execution config will be populated automatically
    - each parameter has several properties:
        - Key - name of the environmental variable
        - Value - Value of the environment variable
        - Show - if the value should be shown as dotted or not - useful for passwords
3. **SSH access:**
    - shown when Remote SSH is chosen as the execution type in command execution configuration
    - parameters:
        - SSH host - host where the command will be executed
        - SSH port - port on which the SSH service is running (by default 22)
        - SSH user - user used to connect via SSH

    - SSH key path:
        - path to your key - needs to be a file only accessible by Data Protector for Cloud Workloads with `400` permissions
        - alternatively, you can use the password access method

4. **Password:**
    - shown when Remote SSH is chosen as the execution type in command execution configuration
    - set your SSH password here if you're not using the public-key authentication method

Enabling WinRM on Windows machines

The Windows Remote Management (a.k.a. WinRM) interface is a network service that allows remote management access to computers via the network. It's used to allow remote management of computers via PowerShell. As a result, WinRM is not enabled by default on Windows Server.

There is an `enable_winrm.ps1` script in the `/opt/vprotect/scripts/winrm` directory

A Power Shell script performs the following steps:

- Automatically starts the WinRM service
- Adds all addresses to trusted hosts. This can be changed in line 7, replacing the aseterix symbol with the appropriate address, e.g.

Add all computers to the TrustedHosts list

```
Set-Item WSMan:\localhost\Client\TrustedHosts -Value *
```

Add all domain computers to the TrustedHosts list

```
Set-Item WSMan:\localhost\Client\TrustedHosts *.yourdomain.com
```

Add specific computers to the TrustedHosts list

```
Set-Item WSMan:\localhost\Client\TrustedHosts -Value DESKTOP-R88J8V5,
MacBookPro19
```

Add computers to the TrustedHosts list using the IP address

```
Set-Item WSMan:\localhost\Client\TrustedHosts -Value 192.168.100.69
```

- Adds an exception in Windows Firewall, which is used by WinRm over HTTPS (port 5986).
- Creates a self-signed certificate and creates Create HTTPS listener.

The `enable_winrm.ps1` script must be run on the Hyper-V server in the PowerShell console.

PowerShell for Linux must be installed on the machine where Node is installed. You can download it from GitHub at:

GitHub - PowerShell/PowerShell: PowerShell for every system!
GitHub

More about installation and versions for different Linux distributions here:

Install PowerShell on Linux - PowerShell
docsmsft

After the correct installation, we can test the connection. On Linux, run PowerShell with the pwsh command.

Then we connect to the Hyper-V server:

```
Enter-PSSession -ComputerName IP_ADDRESS -UseSSL -SessionOption (New-
PSSessionOption -SkipCNCheck -SkipCACheck) -Authentication Basic -
Credential (Get-Credential)
```

After providing the correct credentials, the PowerShell console will start on the remote machine. We end the session with the "exit" command. We can also try a test PowerShell script on a remote machine:

```
Invoke-Command -Session (New-PSSession -ComputerName SERVER_ADDRESS -
UseSSL -SessionOption (New-PSSessionOption -SkipCNCheck -SkipCACheck) -
Authentication Basic -Credential (Get-Credential)) -ScriptBlock {Get-
ChildItem Env:}
```

or

```
Enter-PSSession -ComputerName SERVER_ADDRESS -UseSSL -SessionOption
(New-PSSessionOption -SkipCNCheck -SkipCACheck) -Authentication Basic -
Credential (Get-Credential)
```

Relax and Recover - ReaR

Relax and Recover - ReaR

To create a new application for Relax and Recover database backup, go to the tab: **Applications → Instances**

Then select the **Create** button. The Creating an Application Definition section will be displayed, which needs to be completed.

This is a description of how your script is going to be invoked - you need to specify:

- **Name** - unique name of the application in the Data Protector for Cloud Workloads system
- **Choose Node Config** - select which node should perform the task
- **Backup policy** - specify which backup policy should be used for this application
- **Command execution configuration** - select a prepared template for backup or create a new one yourself

In the SSH access subtab, complete the following fields:

- **Host** - set the address of the host where the instance exists
- **Port** - set the host port

Next, select OS Credentials or create a new one and provide the following information:

- **Name** - unique name of the OS Credentials
- **User** - indicate a user for connecting to the ssh
- **Password** - enter the connection password
- **SSH key path** - alternatively you can specify the ssh key path for authorization

After **saving** the changes to the application, you need to configure the **environment variables** in the **settings section** of your application as needed.

When using the built-in script for database backup, define:

- **VP_REAR_OUTPUT** - Defines where the rescue image should be sent
- **VP_REAR_LOGFILE** - Path to the target server, where Data Protector for Cloud Workloads can place the log file from the backup job.
- **VP_REAR_MOUNTPATH** - Path where the NFS directory will be mounted
- **VP_REAR_STOREPATH** - Path to the Relax and Recover server, where Data Protector for Cloud Workloads can place the backup file.
- **VP_REAR_METHOD** - Relax and Recover backup method
- **VP_REAR_RETENTION** - Number of days to keep the copy
- **VP_REAR_NFSSERVER** - IP address of the NFS server
- **VP_REAR_SCRIPTPATH** - Path to the target server, where Data Protector for Cloud Workloads can place and execute the backup script.

> **Note:** After the new application is fully configured, **save** the changes and go to the Backup SLAs configuration.

Git

Git

To create a new application for Git Repository instances backup, go to the tab: **Applications → Instances** Then select the **Create** button. The 'Creating an Application Definition' section will be displayed, which needs to be completed. This is a description of how your script is going to be invoked - you need to specify:

- **Name** - unique name of the application in the Data Protector for Cloud Workloads system
- **Choose Node Config** - select which node should perform the task
- **Backup policy** - specify which backup policy should be used for this application
- **Command execution configuration** - select a prepared template for Git Repository backup or create a new one yourself

After **saving** the changes to the application, you need to configure the **environment variables** in the **settings section** of your application as needed. When using the built-in script for database backup, define:

- **VP_GITREPO_ADDRESS** - The git clone address (ssh/https) of the repository.
- **VP_GITREPO_USERNAME** - Git name of the user with access to the specified repository.
- **VP_GITREPO_PASSWORD** - Git password of the user with access to the specified repository.
- **VP_GITREPO_STOREPATH** - Path on the server with connection to the Git repository, where Data Protector for Cloud Workloads can place the temporary backup file. No other files should be stored in the given path since everything is removed after the backup is finished.
- **VP_GITREPO_SCRIPTPATH** - Path on Data Protector for Cloud Workloads node, where the backup script is located.

> **Note:** After the new application is fully configured, **save** the changes and go to

> the Backup SLAs configuration.

oVirt/RHV/OLVM

oVirt/RHV/OLVM database

To create a new application for oVirt/RHV/OLVM database backup, go to the tab: **Applications → Instances**

Then select the **Create** button. The Creating an Application Definition section will be displayed, which needs to be completed.

This is a description of how your script is going to be invoked - you need to specify:

- **Name** - unique name of the application in the Data Protector for Cloud Workloads system
- **Choose Node Config** - select which node should perform the task
- **Backup policy** - specify which backup policy should be used for this application
- **Command execution configuration** - select a prepared template for oVirt/RHV/OLVM backup or create a new one yourself

On the SSH access subtab, complete the following fields:

- **Host** - set the address of the host where the oVirt/RHV/OLVM engine exists
- **Port** - set the host port

Next, select OS Credentials or create a new one and provide the following information:

- **Name** - unique name of the OS Credentials
- **User** - indicate a user for connecting to the ssh
- **Password** - enter the connection password
- **SSH key path** - alternatively you can specify the ssh key path for authorization

After **saving** the changes to the application, you can change log path for the newly created instance in **environment variables** in the **settings section**.

Note: After the new application is fully configured, **save** the changes and go to the Backup SLAs configuration.

Kubernetes/OpenShift etcd

Kubernetes/OpenShift etcd

To create a new application for Kubernetes/OpenShift etcd database backup, go to the tab: **Applications → Instances**

Then select the **Create** button. The Creating an Application Definition section will be displayed, which needs to be completed.

This is a description of how your script is going to be invoked - you need to specify:

- **Name** - unique name of the application in the Data Protector for Cloud Workloads system
- **Choose Node Config** - select which node should perform the task
- **Backup policy** - specify which backup policy should be used for this application
- **Command execution configuration** - select a prepared template for Kubernetes/OpenShift etcd backup or create a new one yourself

On the SSH access subtab, complete the following fields:

- **Host** - set the address of the host where the Kubernetes/OpenShift node exists
- **Port** - set the host port

Next, select OS Credentials or create a new one and provide the following information:

- **Name** - unique name of the OS Credentials
- **User** - indicate a user for connecting to the ssh
- **Password** - enter the connection password
- **SSH key path** - alternatively you can specify the ssh key path for authorization

After **saving** the changes to the application, you need to configure the **environment variables** in the **settings section** of your application as needed.

When using the built-in script for database backup, define:

- **VP_ETCD_CLIENT_KEY** - client key path
- **VP_ETCD_CERT** - client cert path
- **VP_ETCD_CA_CERT** - CA cert path
- **VP_ETCD_ADDRESS** - address of etcd instance
- **ETCDCTL_API** - version of etcdctl api

Make sure you have etcdctl tool installed in node you are connecting to. You can find installation instructions on this page ↗

Note: After the new application is fully configured, **save** the changes and go to the Backup SLAs configuration.

Backup & Restore

Backup & Restore

This section is a simple step-by-step guide on how to use Data Protector for Cloud Workloads to protect your Applications.

Backup and restore of Application

Manual backup

Note: to perform a manual backup, a target Application instance must be assigned to Backup SLA

1. Search for Application instance you want to backup and click **Backup** the button next to it.

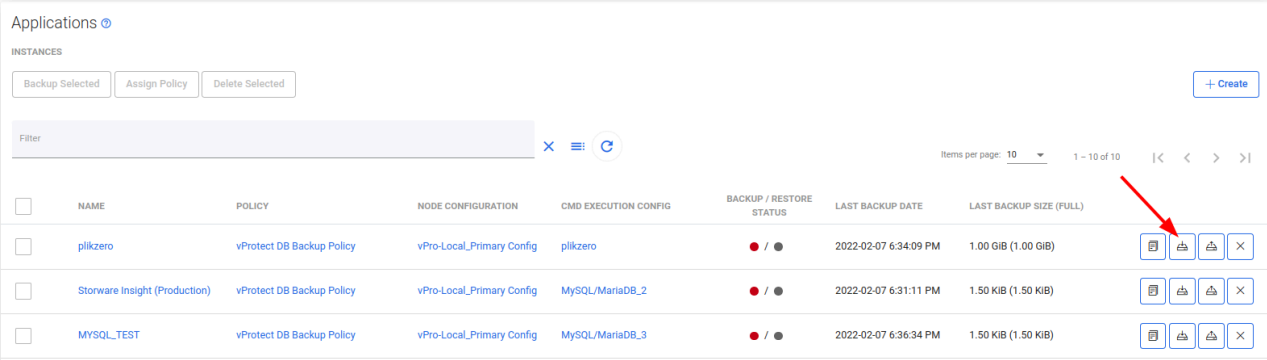

2. In the next window, choose backup type and click **Backup** button.

3. You can track the progress of the task in **tasks console**

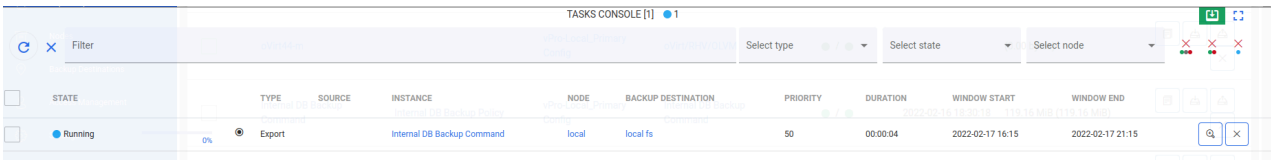

Manual restore

1. Search for Application instance you want to restore and click **Restore** the button next to it.

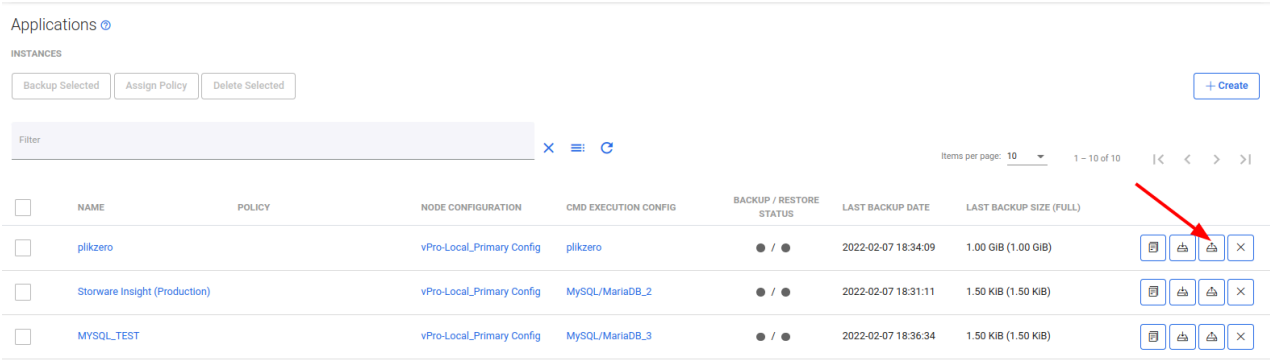

2. In the next window, choose if you want to restore to the file system or to the application instance host. Enter the rest parameters according to your needs and click **Restore** button.

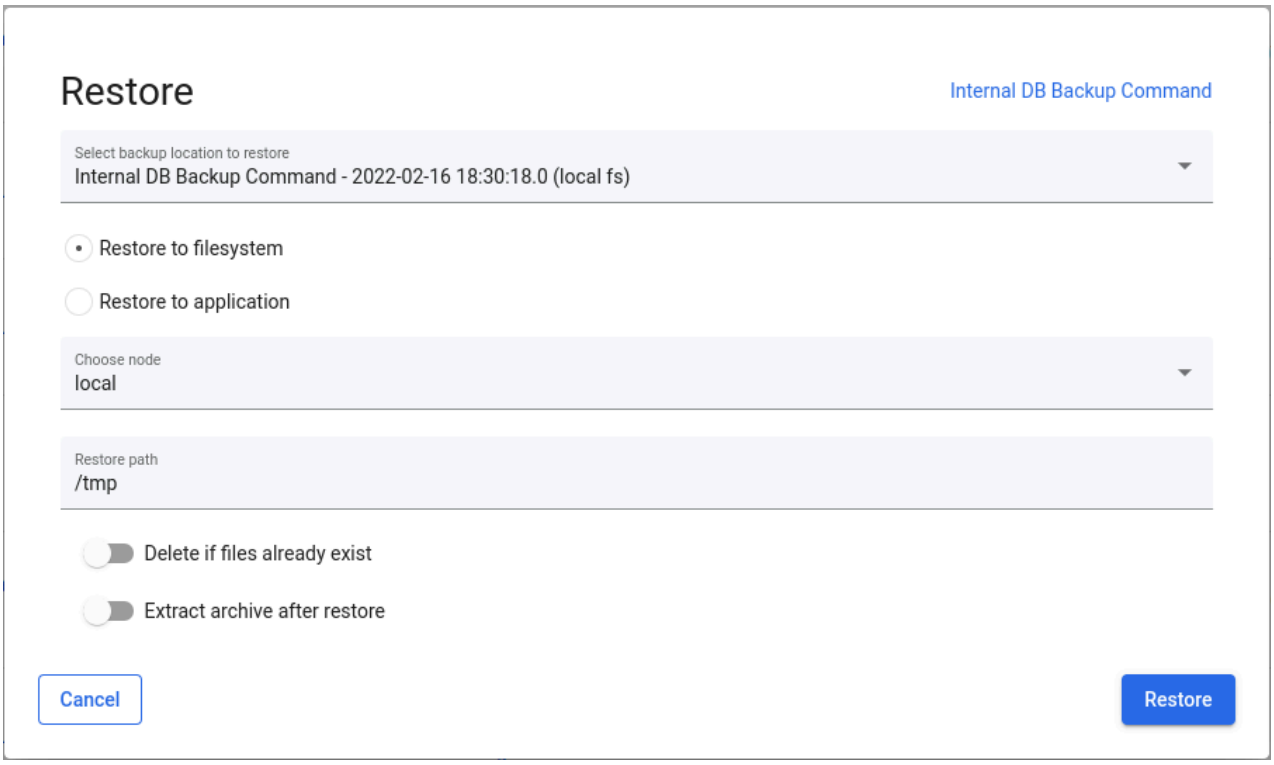

> **Note:** More information about protecting Applications you can find in the administration chapter.

Protecting Storage Providers

Protecting Storage Providers

Data Protector for Cloud Workloads supports data protection for several storage providers.

In the chapter Storage Providers you can read how to use Data Protector for Cloud Workloads to protect Storage Providers and how to protect them you can read in chapter Backup & Restore

Storage Providers

Protecting storage providers

This feature allows you to protect file systems mounted on Data Protector for Cloud Workloads Nodes, Ceph RBD volumes, and more.

- Ceph RBD
- Nutanix Files
- Nutanix Volume Groups

> **Note**: Synthetic backup destination is not available for Storage Providers

Ceph RBD

Ceph RBD

General

In order to connect to Ceph RBD you need to provide the keyring and configuration files. The Ceph RBD storage provider should detect the volumes and pools in the environment and allow you to assign backup policies. Data Protector for Cloud Workloads uses the RBD-NBD approach to mount a remote RBD snapshot over NBD and read data.

Note:

- Data Protector for Cloud Workloads needs access to the monitors specified in the Ceph configuration file.
- When creating Ceph RBD storage provider for the OpenStack environment, only the credentials specified in the storage provider form are used by the OpenStack backup process - the actual technique (RBD-NBD mount or cinder in disk-attachment strategy) and **node** for connecting and the backup volumes depend on the OpenStack hypervisor manager settings, not in the storage provider settings.

Example

Complete the following steps to add the Ceph RBD storage provider:

- Data Protector for Cloud Workloads Node supports Ceph RBD, for which you will need to install ceph libraries:
  - On Data Protector for Cloud Workloads **Node** enable the required repositories:

For Data Protector for Cloud Workloads node installed on RHEL 7:

```
sudo subscription-manager repo --enable=rhel-7-server-rhceph-4-tools-
rpms
```

For Data Protector for Cloud Workloads node installed on RHEL 8:

```
sudo subscription-manager repo --enable=rhceph-4-tools-for-rhel-8-
x86_64-rpms
```

For Data Protector for Cloud Workloads node installed on CentOS 7:

```
sudo yum install epel-release
sudo rpm --import 'https://download.ceph.com/keys/release.asc'
sudo yum install https://download.ceph.com/rpm-octopus/el7/noarch/ceph-
release-1-1.el7.noarch.rpm
```

For Data Protector for Cloud Workloads node installed on CentOS Stream 8:

```
sudo yum install epel-release
sudo rpm --import 'https://download.ceph.com/keys/release.asc'
sudo yum install https://download.ceph.com/rpm-octopus/el8/noarch/ceph-
release-1-1.el8.noarch.rpm
```

For Data Protector for Cloud Workloads node installed on CentOS Stream 9:

```
sudo yum install epel-releasem
```

Add Ceph repository

```
vi /etc/yum.repos.d/ceph.repo
```

```
[ceph]
name=Ceph packages for $basearch
baseurl=https://download.ceph.com/rpm-reef/el9/$basearch
enabled=1
priority=2
gpgcheck=1
gpgkey=https://download.ceph.com/keys/release.asc
```

- Install the rbd-nbd and ceph-common package, with all dependencies:

  ```
  yum install rbd-nbd ceph-common
  ```

- Go to **Storage** -> **Infrastructure** and click **Create** button
- Choose **Ceph RBD** as the type and select the node configuration responsible for backup operations
- Click **Upload keyring file** button and select **Ceph keyring file** which can be obtained from the Cinder host - for example in `/etc/ceph/ceph.client.admin.keyring**
- Provide **Ceph configuration file content**, for example:

  ```
  [global]
  cluster network = 10.40.0.0/16
  fsid = cc3a4e9f-d2ca-4fec-805d-2c40605723b3
  mon host = ceph-mon.domain.local
  mon initial members = ceph-00
  osd pool default crush rule = -1
  public network = 10.40.0.0/16
  [client.images]
  keyring = /etc/ceph/ceph.client.images.keyring
  [client.volumes]
  keyring = /etc/ceph/ceph.client.volumes.keyring
  [client.nova]
  keyring = /etc/ceph/ceph.client.nova.keyring
  ```

> **Note:** Remember, above content need to end with the new line sign.

- If you want to index only ceph pools of your choice, change **Storage pool management strategy** to **INCLUDE** and add storage pool names.
- Click **Save** - now you can initiate inventory synchronization (pop-up message) to collect information about available volumes and pools

    - later you can use the **Inventory Synchronization** button on the right of the newly created provider on the list.
- Your volumes will appear in the**Instances** section in the submenu on the left, from which you can initiate backup/restore/mount tasks or view volume backup history and its details.

Nutanix Files

Nutanix Files

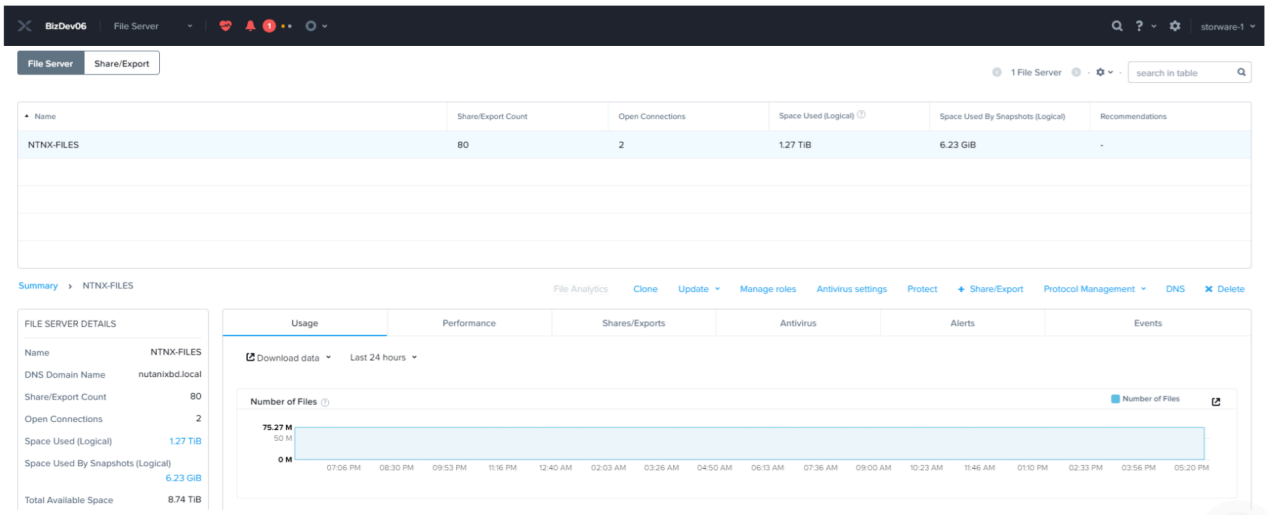

General

Nutanix Files offers file system storage in the form of network shares (SMB or NFS). With Data Protector for Cloud Workloads, you can easily back up and recover these shares.

> **Note:**
>
> - there must be an account defined in both Prism Central and File Server with the same credentials
> - each File Server is considered as a separate Storage Provider
> - Data Protector for Cloud Workloads tracks changes for each file system change using the Nutanix CFT mechanism
> - this shortens backup times because it does not need to perform a metadata scan across your file server, which could contain millions of files and directories.
> - only regular files, symlinks, and directories are backup

> - for file systems, Data Protector for Cloud Workloads currently build images for each backup before uploading data to the specific backup provider so that it doesn't have to upload objects one-by-one - these images will be merged automatically in restore tasks

In general, the process looks as follows:

- the user adds a new storage provider (Nutanix files)
- inventory synchronization retrieves all the information about storage (of two types: SMB and NFS)
- the user can then do a full backup (storage is mounted and all files are extracted from it)
- incremental backup uses CFT from the API, and based on this list Data Protector for Cloud Workloads only downloads changed files.

Example

Complete the following steps to add the Nutanix Files storage provider:

Go to **Storage** → **Infrastructure** and click on **Create** button

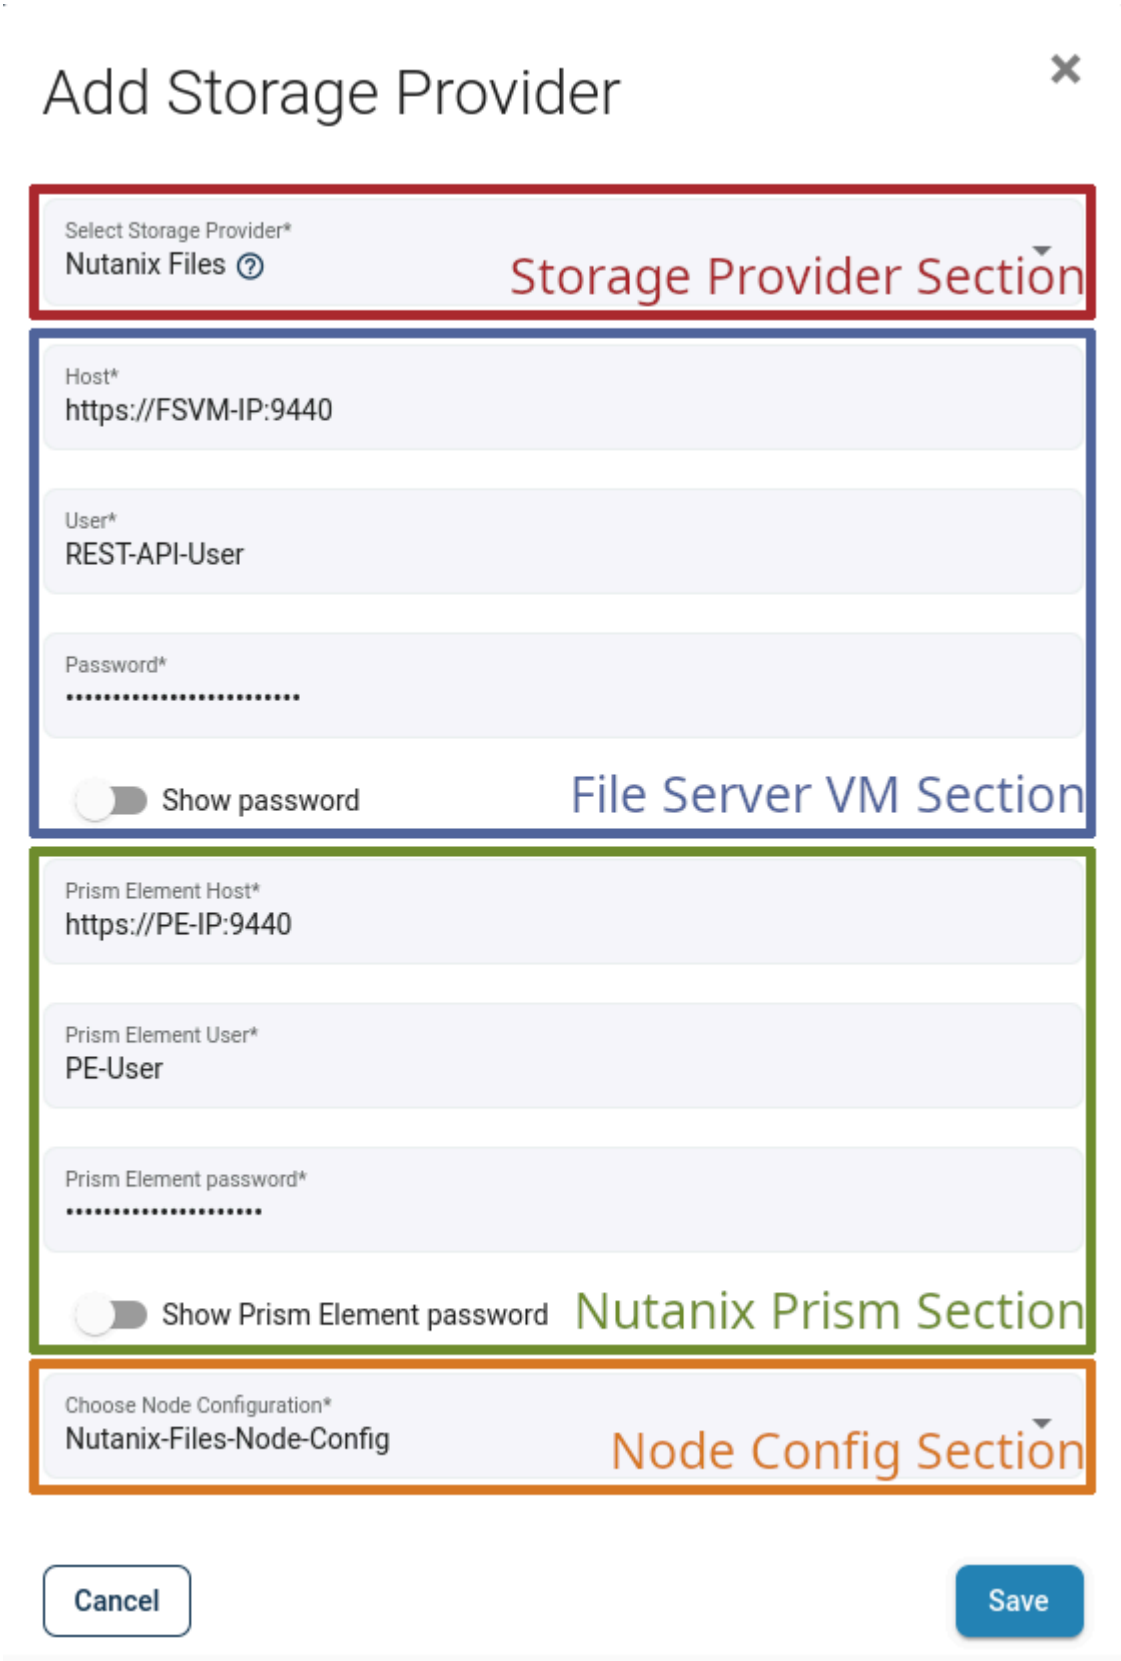

In the form, provide:

- **Nutanix Files** as the Storage Provider type
- File Server URL in the **Host** field (**https://FileServerVM-IP:9440**)

- Login and password for File Server VM
  - The **admin** account **cannot** be used
  - User has to be authorized as **REST API access user** in **File Server Console** > **Configuration** > **Manage Roles**
  - If you want to use domain account, enter it in the following format: **domain\username**
- The URL to the Prism Element Host (**https://PRISM_ELEMENT:9440**)
- User login and password to Prism Element Host
  - The **admin** account **cannot** be used
  - User has to have Cluster Admin privileges on the cluster
  - If you want to use domain account, enter it in the following format: **username@domain.local**
- Choose the node configuration responsible for operations.
  - The node has to have network communication over port 9440 to File Server VM Client Network IP and Nutanix Prism.

Click **Save** button - now you can run inventory synchronization to detect file systems (storage instances) - they will be visible in **Storage** → **Instances** view

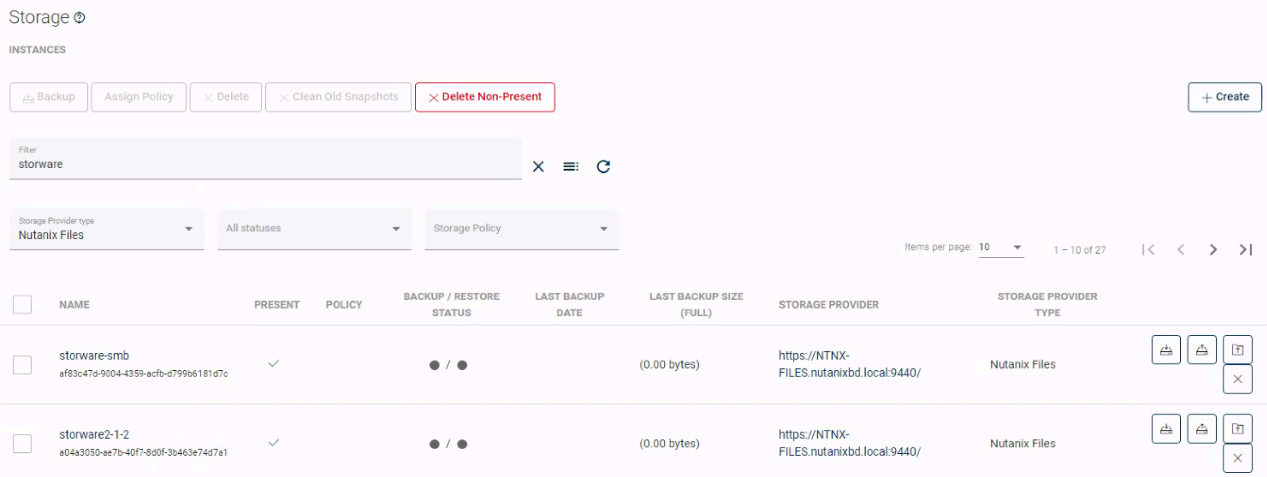

Now you can assign backup SLAs in **Storage** → Backup SLAs to schedule periodic backups of these shares. You are also able to mount backups to restore individual files using Mounted Backups (File-level Restore).

Limitations

- Backup/Restore of hidden shares on Nutanix Files with SMB are not supported.

Nutanix Volume Groups

Nutanix Volume Groups

General

Data Protector for Cloud Workloads supports volume groups for the Nutanix platform. Disk snapshots are being attached to the host with Data Protector for Cloud Workloads Node using iSCSI. Thanks to this backups and restores can be performed remotely without Proxy VM.

This storage provider also allows excluding disks for backups.

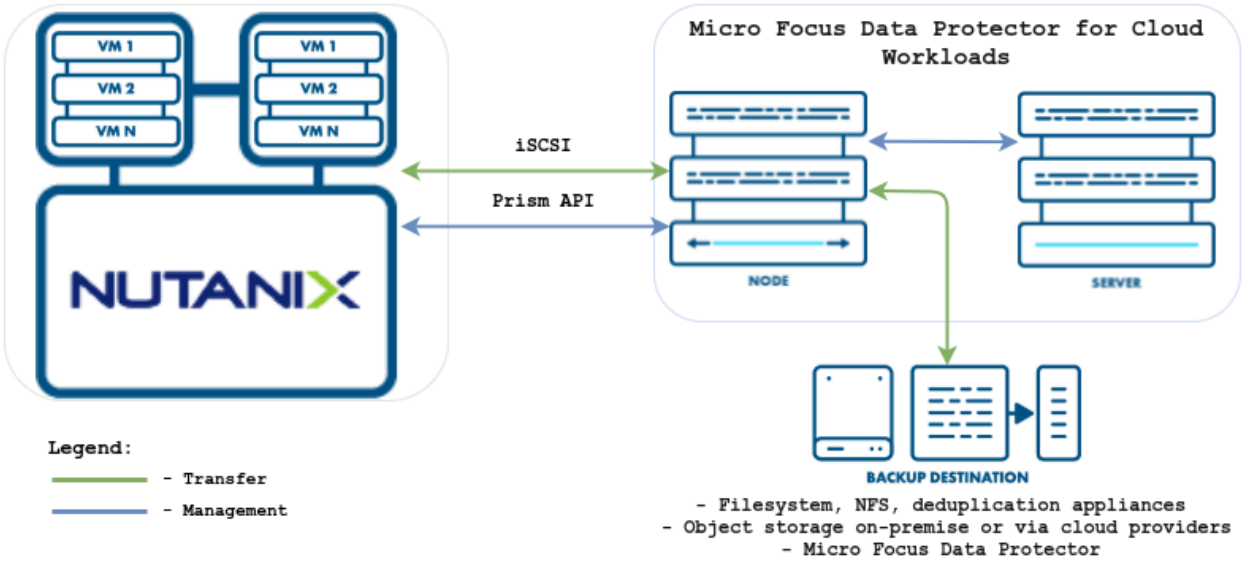

In general, the export process looks as follows:

- A Snapshot of the volume group is taken
- A Temporary volume group with disk snapshots is created
- Disks from the temporary volume group are attached to Data Protector for Cloud Workloads Node host using iSCSI
- Disks content is exported to RAW files
- Disks are then detached and temporary volume group is deleted

In general, the restore process looks as follows:

- A new volume group is created with yet empty disks
- Disks from a new volume group are attached to Data Protector for Cloud Workloads Node host using iSCSI
- Content from RAW files of backup is imported to attached disks
- Disks are detached from Data Protector for Cloud Workloads Node host

> **Note:** iSCSI uses `iSCSI Data Services IP` defined on the Nutanix platform to discover targets, which means that exception will be thrown if this value is missing.

Example

Complete the following steps to add the Nutanix Volume Groups storage provider:

Go to **Storage** → **Infrastructure** and add click on **Create**

In the form provide:

- `Nutanix Volume Groups` as a type
- URL to Prism Element Host (`https://PRISM_ELEMENT:9440`)
- Name of the storage provider
- Login and password
- The node configuration responsible for backup operations

Click **Save** - now you can run inventory synchronization to detect volume groups - they will be visible in **Storage** → **Instances** view

Now you can create a backup task for the chosen volume group.

Backup & Restore

Backup & Restore

This section is a simple step-by-step guide on how to use Data Protector for Cloud Workloads to protect your Storage Providers.

Backup and restore of Storage Provider

Manual backup

Note: to perform a manual backup, target Storage Provider instance must be assigned to Backup SLA

1. Search for Storage Provider instance you want to backup and click the **Backup** button next to it.

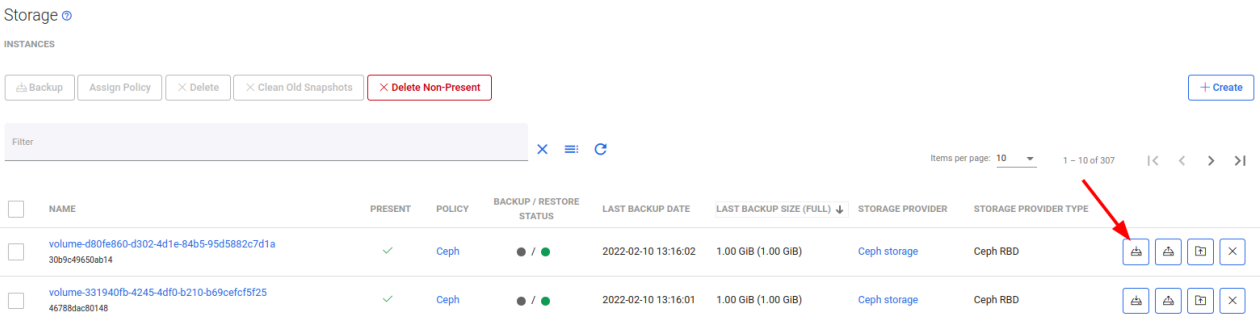

2. In the next window, choose backup type and click **Backup** button.

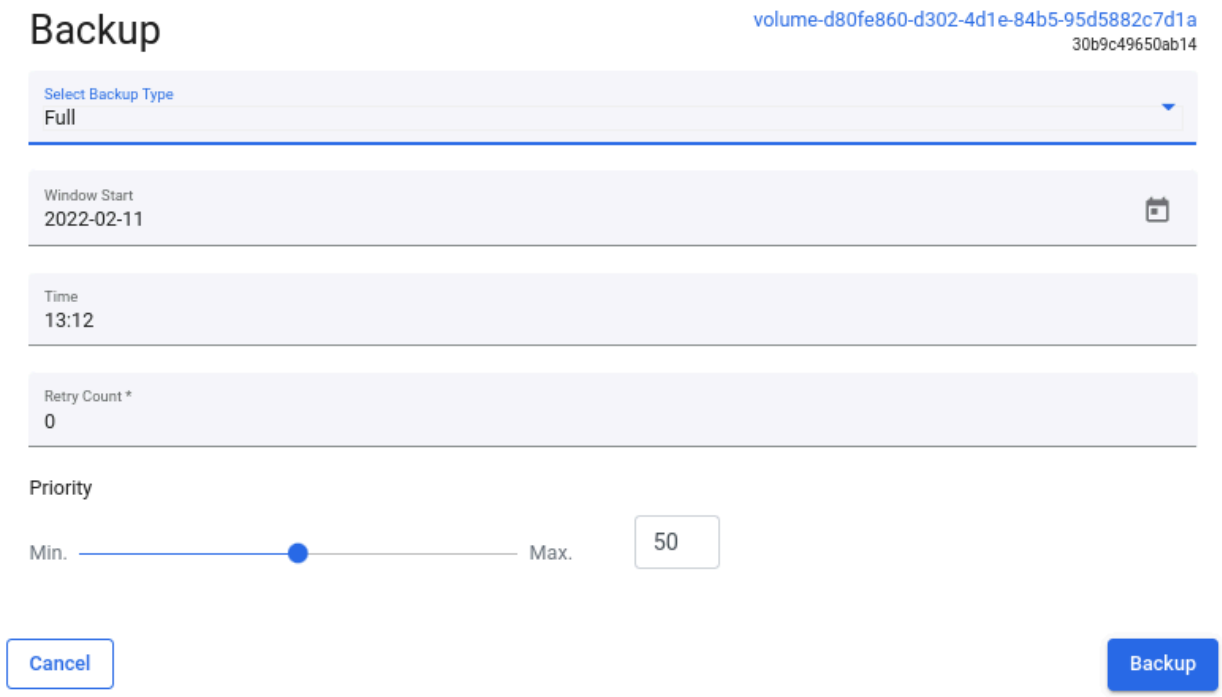

3. You can track the progress of the task in **tasks console**

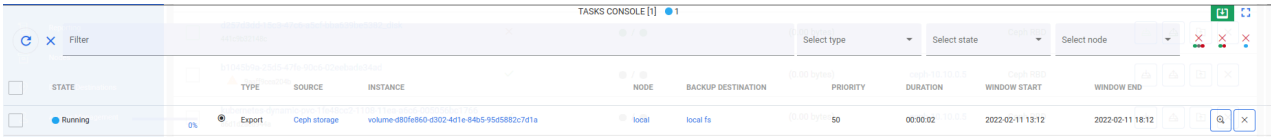

Manual restore

1. Search for Storage Provider instance you want to restore and click the **Restore** button next to it.

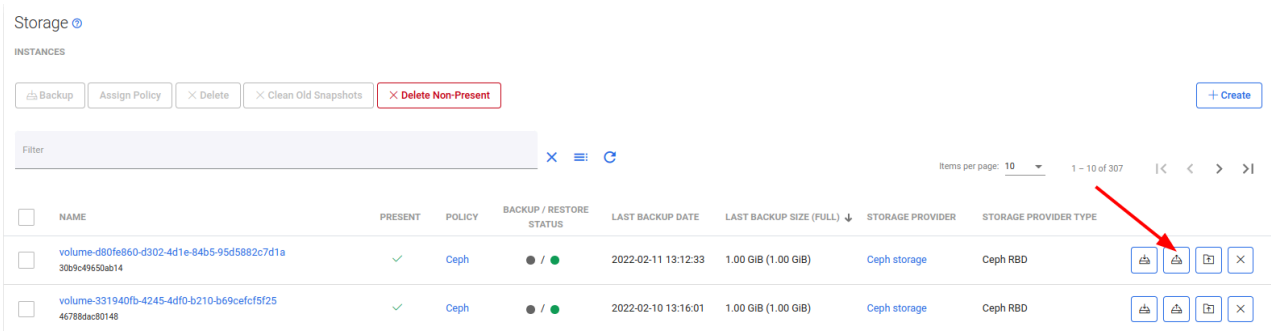

2. In the next window, choose if you want to restore to File System or directly to the storage provider. Enter the rest parameters according to your needs and

click the **Restore** button. The available parameters may differ depending on the storage provider.

> **Note:** More information about protecting Storage Providers you can find in the administration chapter.

Administration

Administration

This section provides information about administrative tasks such as how to backup, restore, or manage key elements using Web UI of Data Protector for Cloud Workloads like:

- Dashboard
- Virtual Environments
- Storage Providers
- Cloud
- Applications
- Reporting
- Nodes
- Access Management
- Settings

We also recommend reading the Disaster Recovery section carefully to be sure that you understand how to recover Data Protector for Cloud Workloads installation in case of losing the node or even server.

Dashboard

Dashboard

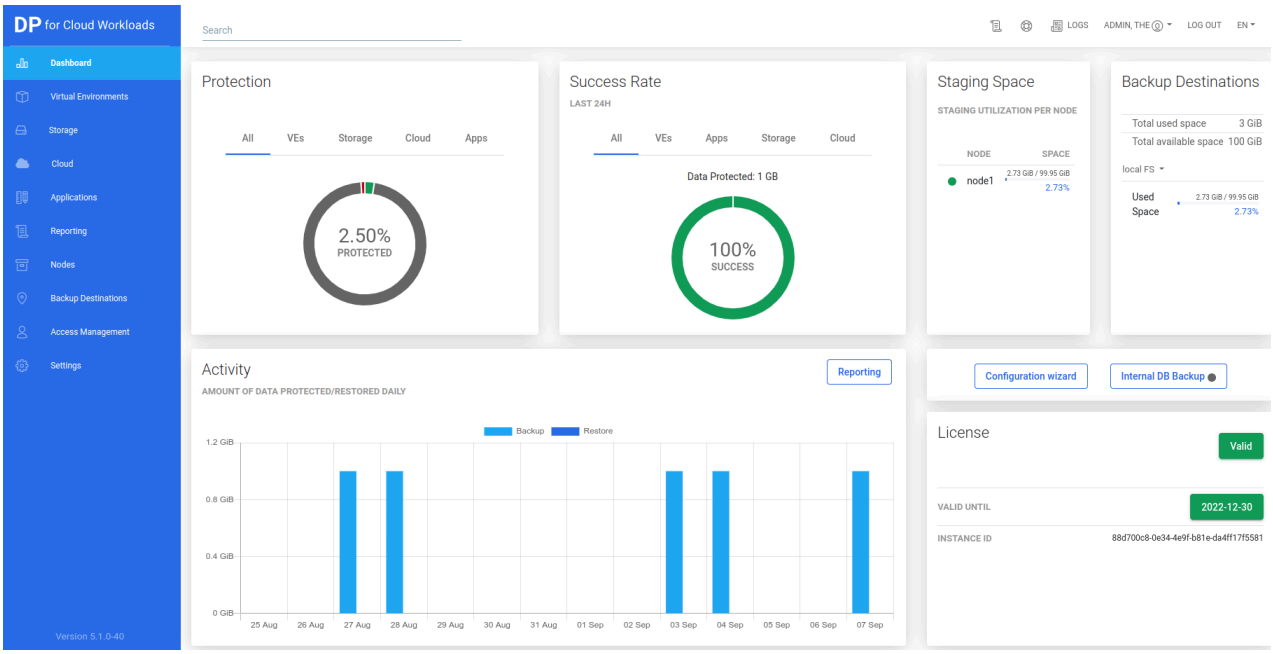

Overview

The primary Data Protector for Cloud Workloads interface is the WEB UI accessible via a web browser.

Divided into a few sections, it makes it possible to view and set the most vital options related to management, monitoring and reporting.

The left pane contains the main menu (see below for a more detailed description).

The **Protection** field present the overall statistics about protection of your environment. The instance is considered protected when it is attached to Backup SLA and the backup has been done in the last 24 hours.

The upper right corner provides access to documentation, support and system logs.

The lower part provides the option to view the (sliding) task console.

The menu on the left provides access to the most important sections:

- Dashboard - the main screen with a general summary and the configuration wizard
- Virtual Environments - page where you can add and protect your Environments
- Storage - page where you can add and protect your Storage
- Cloud - page where you can add and protect your Microsoft 365 organization
- Applications - page where you can add and protect your Applications
- Reporting - page where you can see all reports from Data Protector for Cloud Workloads
- Nodes - node management and node configurations
- Backup Destination - create and manage backup destinations
- Access Management - create and manage your accounts (Language, time zones, etc.)
- Settings - From here, you can manage global settings, licenses, email, authentication, and internal DB backup.

Virtual Environments

Virtual Environments

This section provides information about administrative tasks like:

Instances - list of currently known virtual machines and access to the details page of each object.

Infrastructure - access configuration for hypervisors and hypervisor managers, basic info about the inventoried environment, such as clusters and storage.

Backup SLAs - allows you to set up a correlation between virtual machines, the backup destination, and schedules (Policies tab). It also allows you to configure policy schedules (schedules tab).

Snapshot SLAs - allows you to set up a correlation between virtual machines, snapshot retention, and schedules (Policies tab). It also allows you to configure policy schedules (schedules tab).

Recovery Plans - this allows you to set up a number of rules to restore your virtual machines into a hypervisor or filesystem to test your backups (Policies tab). It also allows you to configure policy schedules (schedules tab).

Mounted Backups (File-level Restore) - browse and download files from mounted backups.

Instances

Instances

General

A list of currently known virtual machines and access to the details page of each object. From this place, you can also perform on-demand actions such as backup, restore and file-level restore.

The field lets the user filter the virtual machines by their:

- Name
- UUID
- GUID
- Tags
- Url of a linked hypervisor manager
- Node configuration name linked via hypervisor manager
- Host field of a linked hypervisor
- Node configuration name linked via a hypervisor
- Name field of a linked VM Backup Policy
- Name field of a linked Snapshot Management Policy

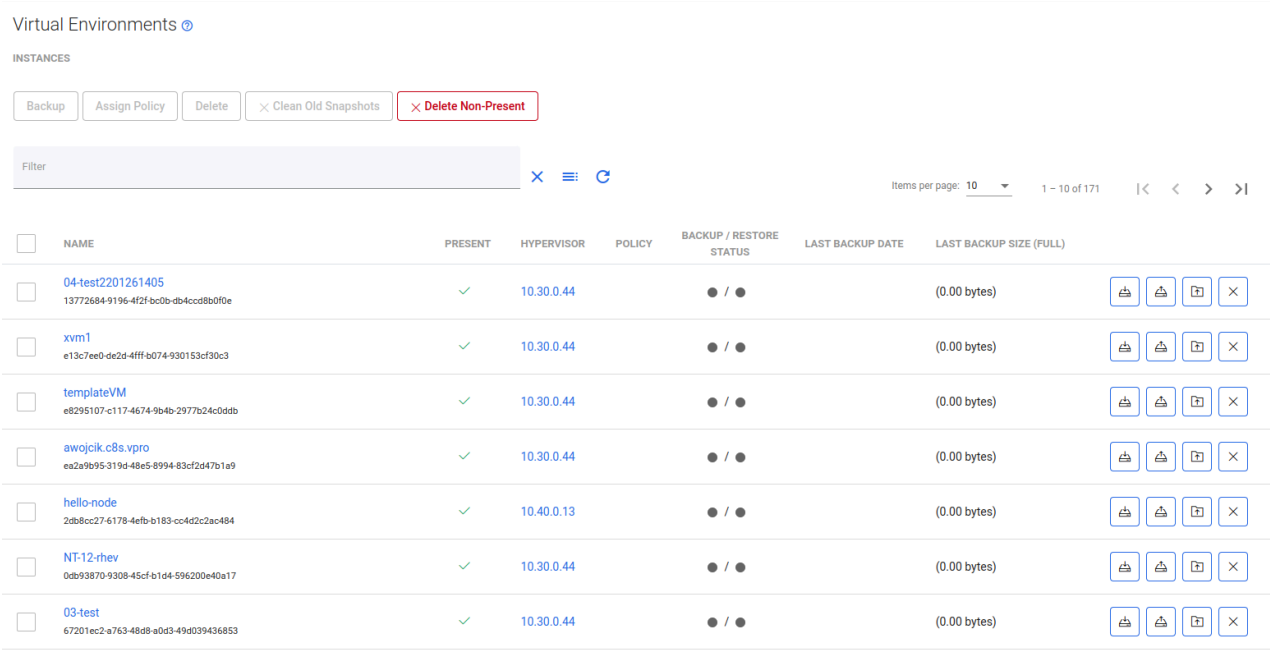

Returning to the VM details page, this is what it looks like:

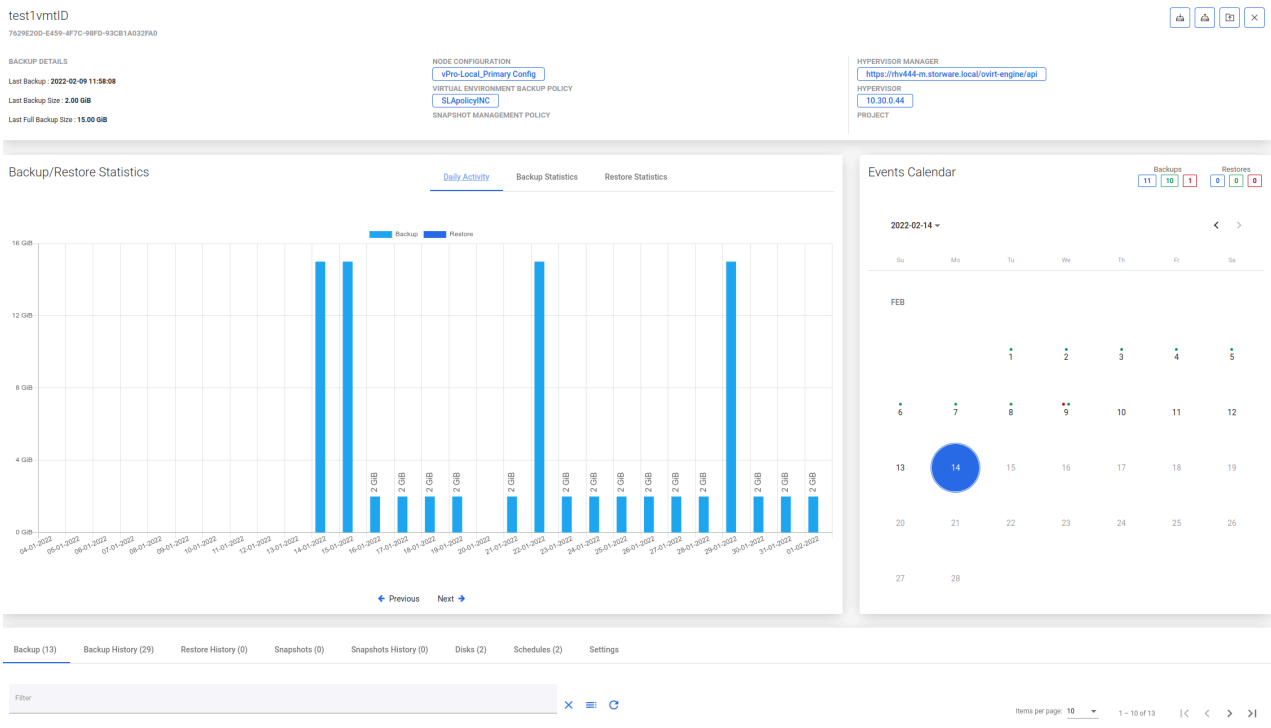

As you can see, the window has been divided into several areas:

VM Summary

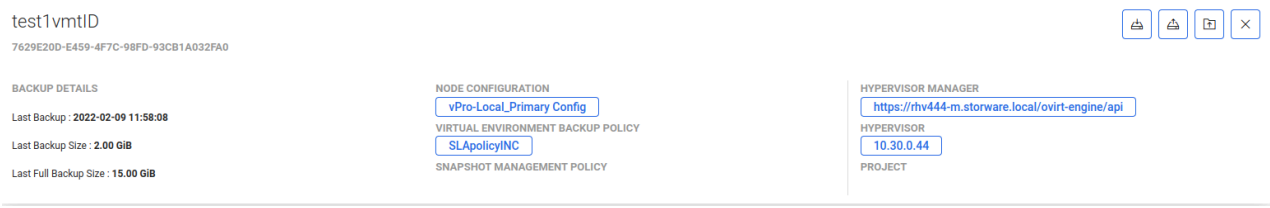

At the top, you can see summarized pieces of information about the VM, such as:

- the ID of the VM object into Data Protector for Cloud Workloads
- to which hypervisor the VM belongs
- which node is backing up this virtual machine
- short information about the last backup actions
- whether the virtual machine has tags or policies assigned to it

You can also use several function buttons:

- refresh
- back to list
- backup
- restore
- mount
- delete

Backup/Restore Statistics

Daily activity

First, you'll see a daily summary of the backup and restore operations for the last month. This view is called "Daily Summary" and is the default view. You can switch the report between four views.

Backup Size

This view shows separate columns for each backup made to the virtual machine. Thanks to this, you can easily determine what data increase occurs on a given machine.

Backup Time

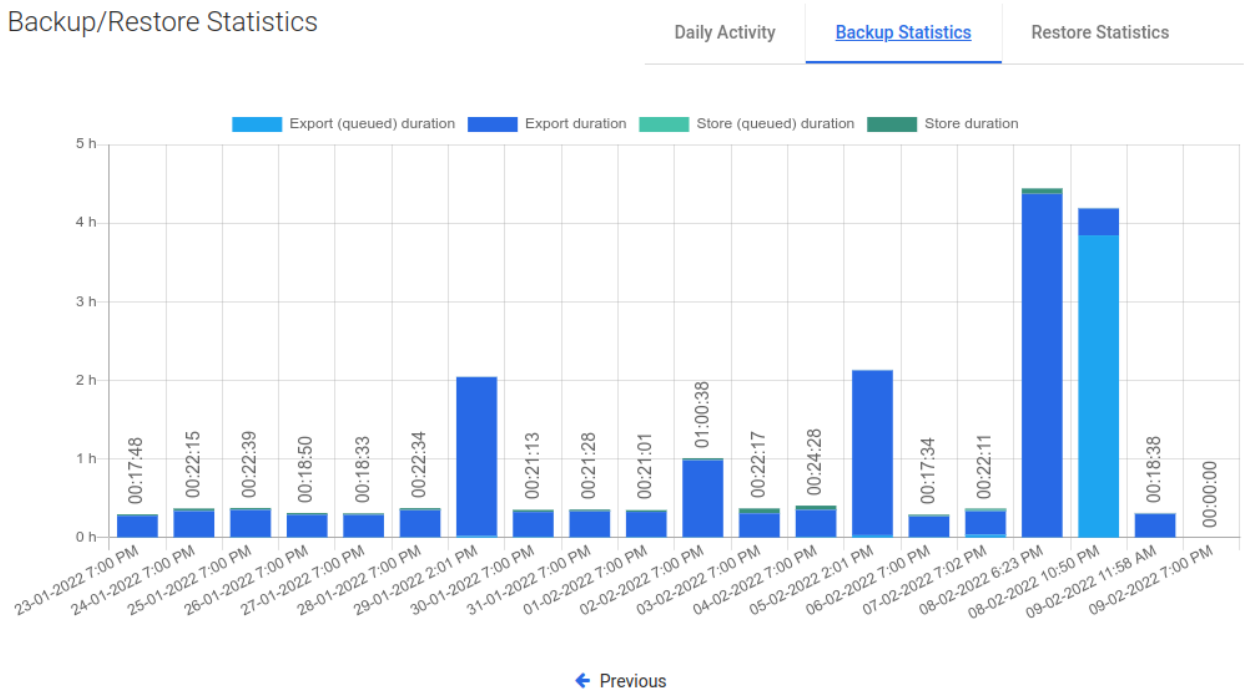

A very useful report. It allows you to determine the required window length for backups or, based on the time of individual phases, it is easy to deduce the cause of slow backups.

Transfer Rate

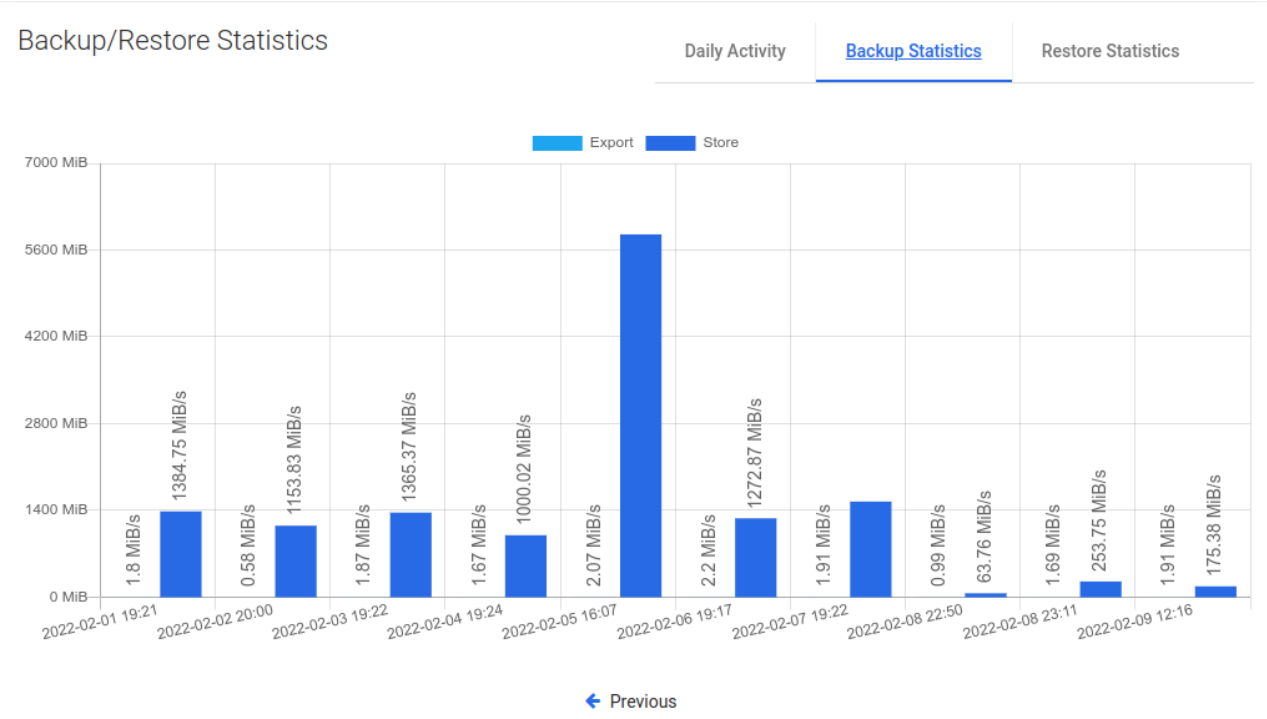

One of the latest reports, now you can easily see how fast data transfer is.

Restore Time

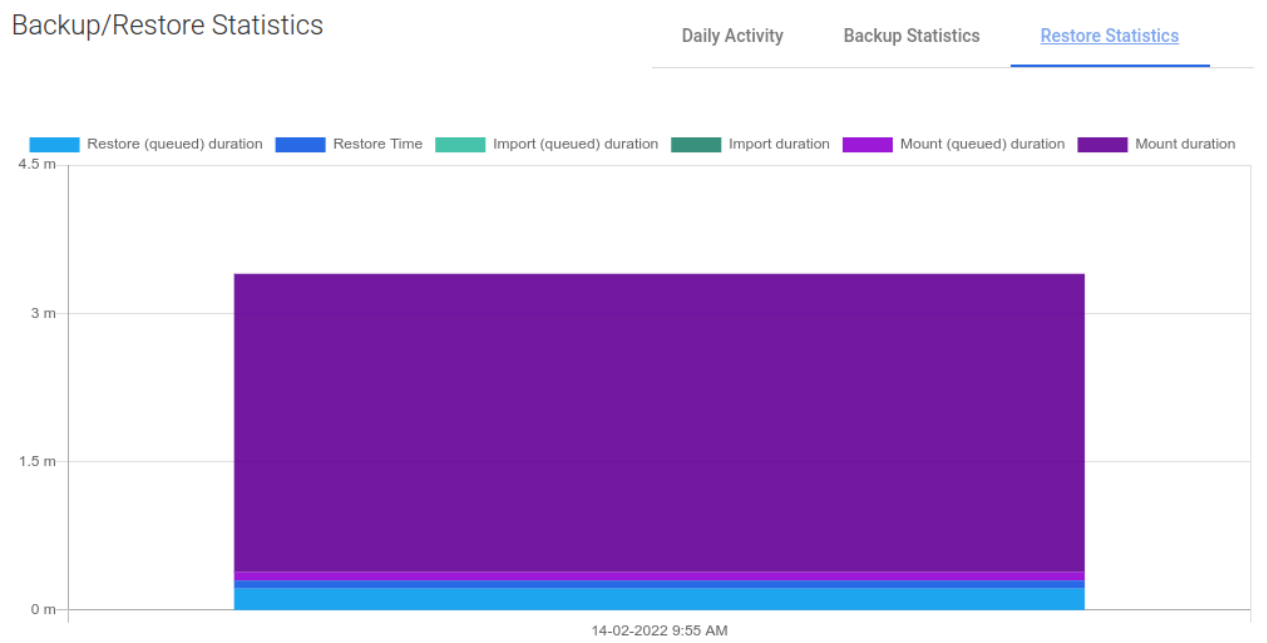

A view with the same properties as "Backup Time". It allows us to estimate how long it will take to restore the machine in the event of a failure.

Transfer Rate

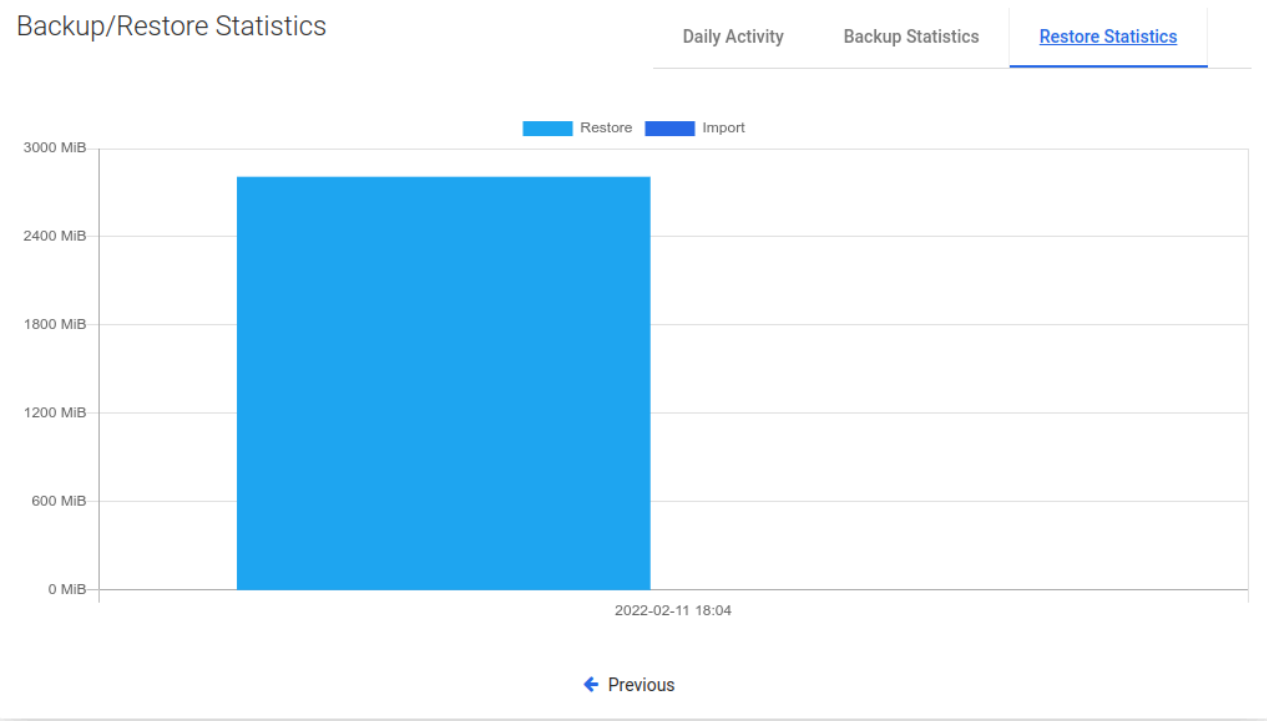

As in the previous case, we also have the transfer speed for the restore job.

Events Calendar

The calendar extends the possibilities of adjacent statistics. It allows you to neatly define the range of days you want to see, additionally makes a quick summary of the number of backups and restores (top right corner).

Blue - the sum of all backups, **Green** - the sum of successes, **Red** - the sum of failures.

Bottom menu

In the bottom menu, you can find a large number of tabs, each of which will present different information or will allow you to change the configuration of this particular virtual machine.

Backup

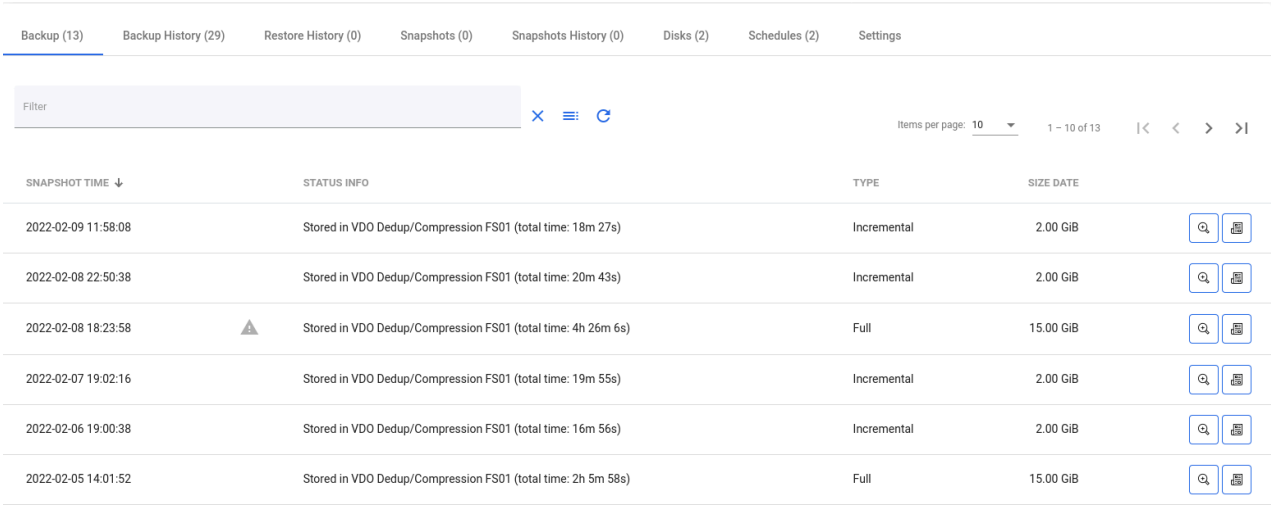

The first tab shows all virtual machine backups that are currently available and all basic information about them in a list. After clicking on the magnifying glass button, you will see additional information. The button next to it allows you to download logs in the form of a .txt file.

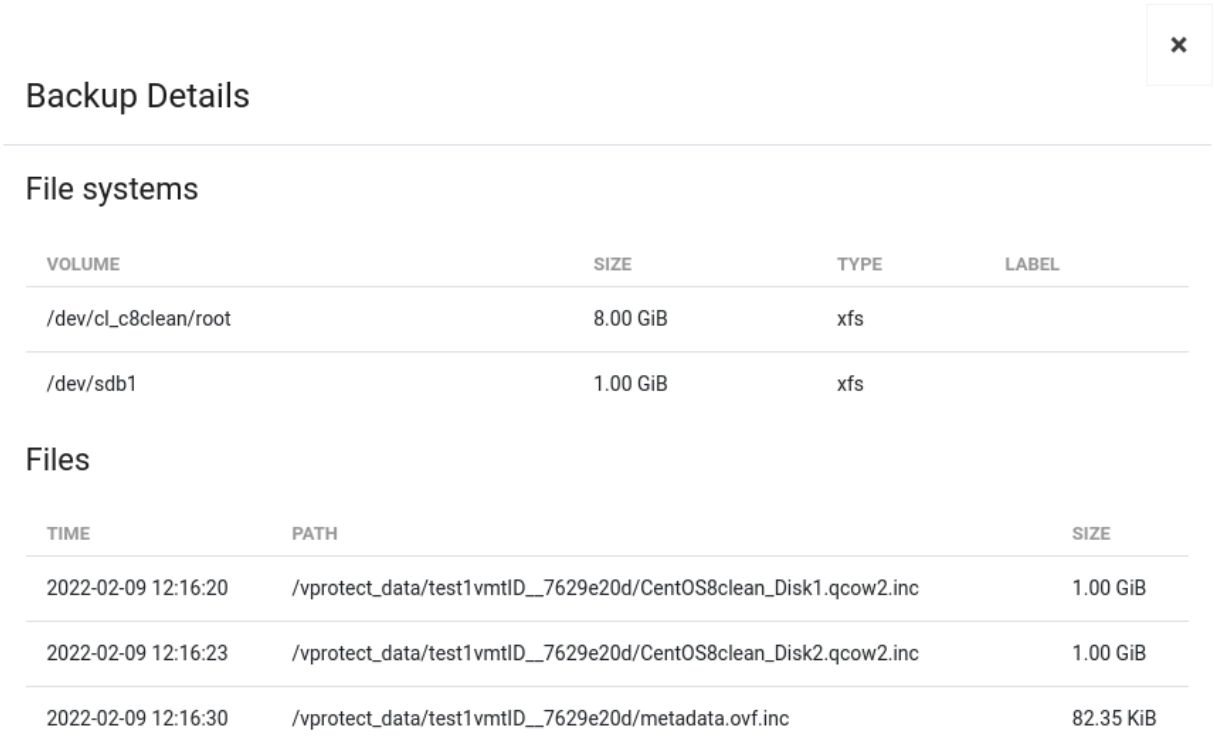

Backup History

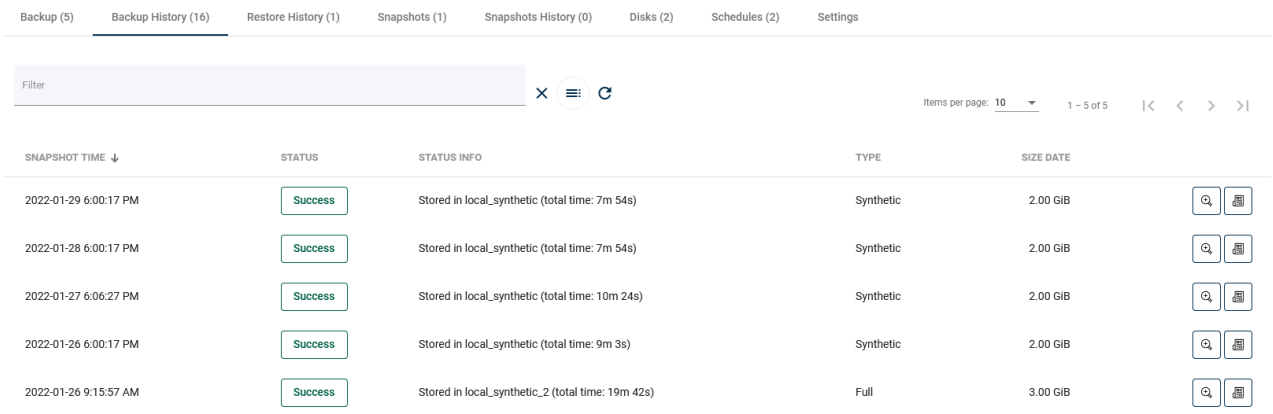

This tab shows information about all backups made for this virtual machine, as well as information about failed, removed (because of retention), or currently executing backups.

Restore History

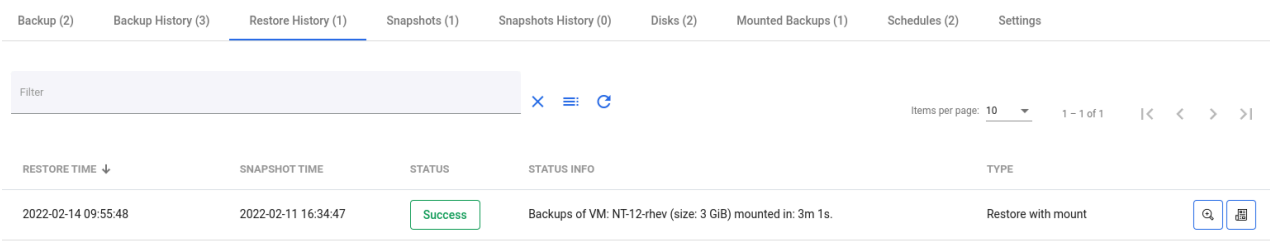

This tab is similar to "Backup History". This is a list with basic information about the virtual machine restores performed. When you open the details of the selected restore, you will see more detailed information.

Restore details

Node	local
Protected Entity	NT-12-rhev
Restore Time	2022-02-14 09:55:48
Status	Success
Status Info	Backups of VM: NT-12-rhev (size: 3 GiB) mounted in: 3m 1s.
Backup Type	Full
Restore Type	Restore and mount

Task Time Stats

Mount duration	03:01 s
Mount (queued) duration	00:05 s

Backup

Backup Time	2022-02-11 18:04:22
Protected Entity	NT-12-rhev
Size	3.00 GiB
Snapshot time	2022-02-11 16:34:47
Status	Success
Status Info	Stored in local fs (total time: 1h 29m 35s)
Type	Full

Snapshots

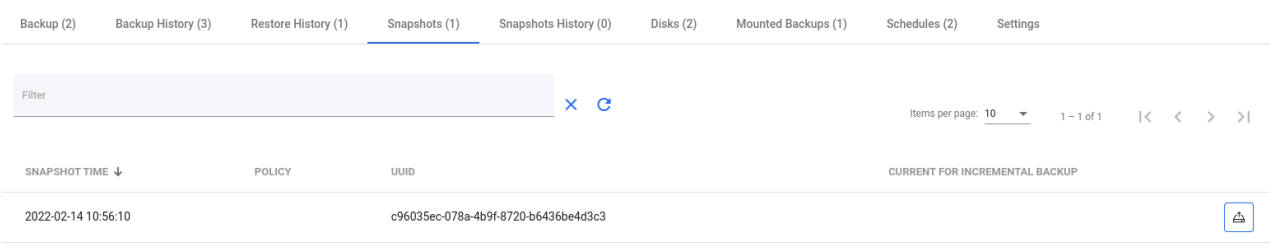

This tab shows virtual machine snapshots (remember - snapshots are stored on the hypervisor). The snapshot can be divided into two categories:

1. As you can see in the list above, there is a green dot next to the snapshot. This means that this snapshot is created for incremental backup purposes. This is an automatic operation and we only keep the last snapshot.
2. The second one on the list is a snapshot created at the user's request (scheduled or manual).

Next to the snapshot is a button that allows you to restore the virtual machine. It actually creates a new virtual machine and keeps the old one (security considerations to protect against the human factor).

Disks

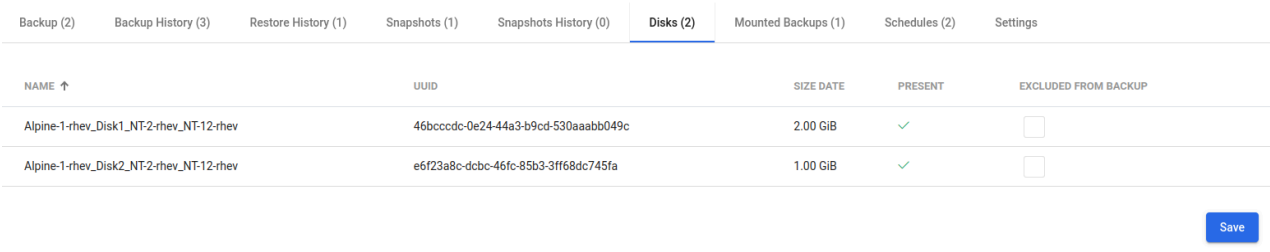

It is worth remembering that if such a virtual machine is restored, the excluded disks will be created from scratch and connected to the machine.

Mounted Backups

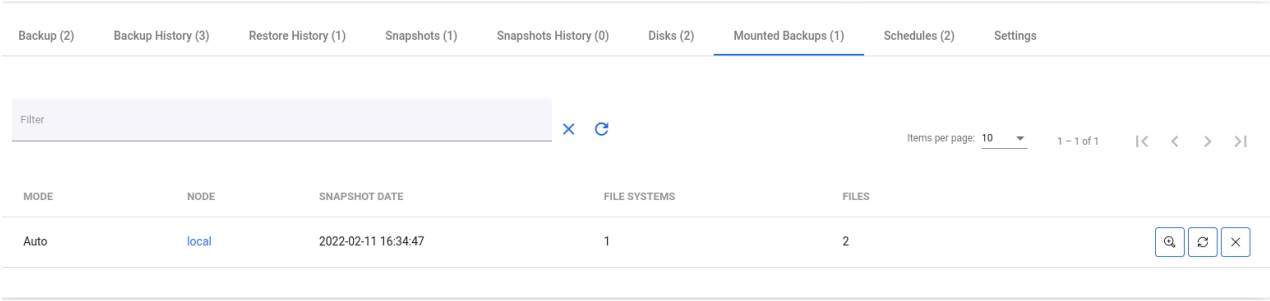

This tab shows a list of mounted backups for this virtual environment.

Schedules

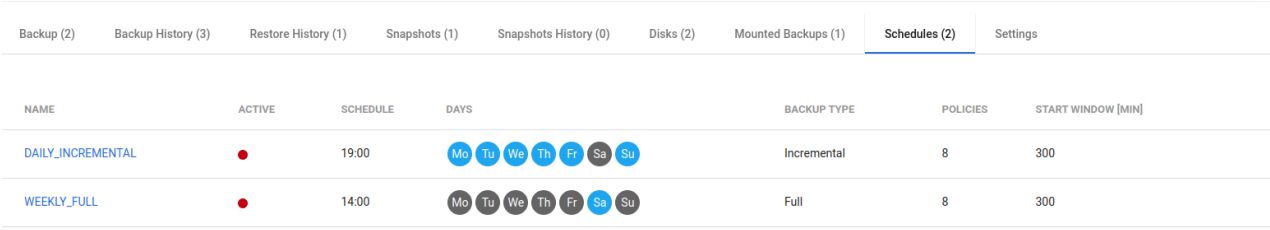

In this tab, you can see all the schedules assigned to the virtual machine.

Settings

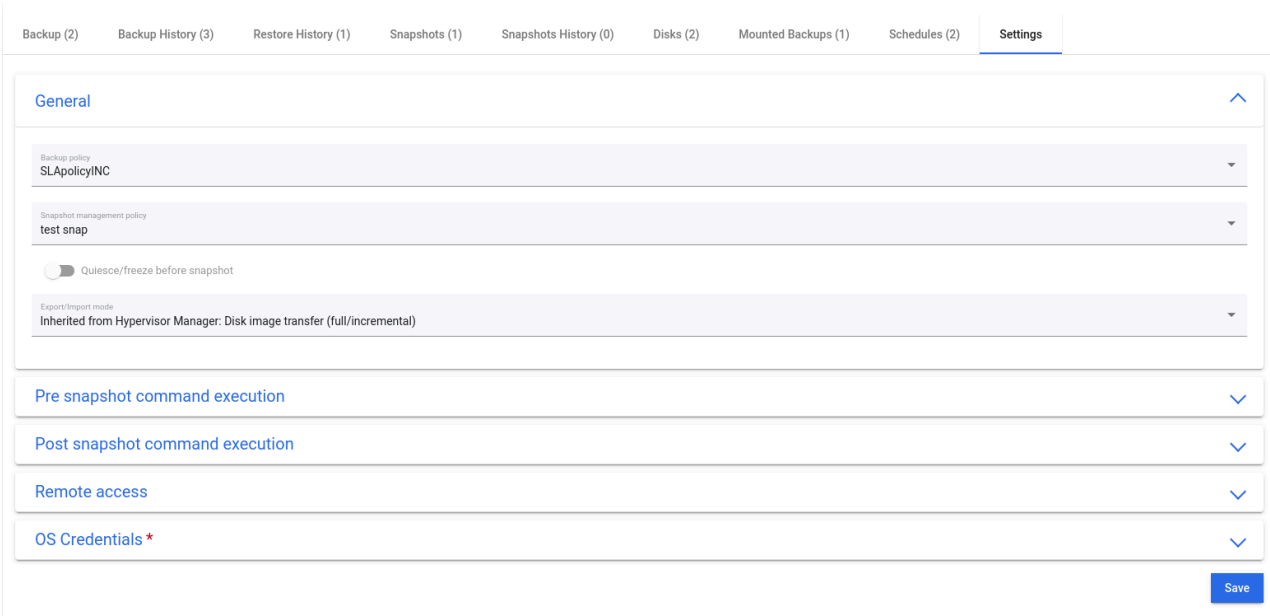

Finally, the last tab. The first two options allow you to change the policies assigned to the virtual machine. The third is a toggle to turn on or off the "Snapshot

consistent technology" feature. Below You can change the transfer mode which is used for this VM.

Performing pre/post snapshot commands is a function intended for advanced users. As the name implies, it allows us to execute scripts via an ssh connection, either before or after taking a snapshot.

Backup on-demand

Backup on-demand

Virtual Environments

To perform on-demand backup go to the Instances tab under the Virtual Environment section.

You can click on  icon to backup one VM

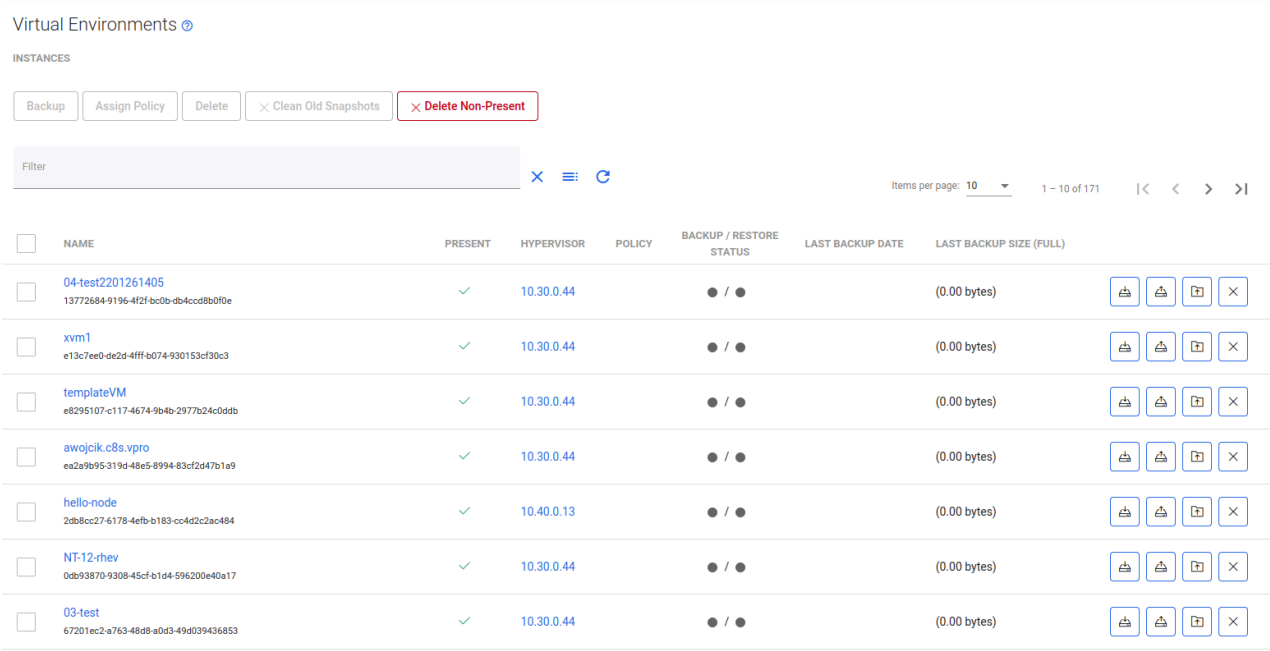

or select multiple virtual machines and click on  icon to backup it.

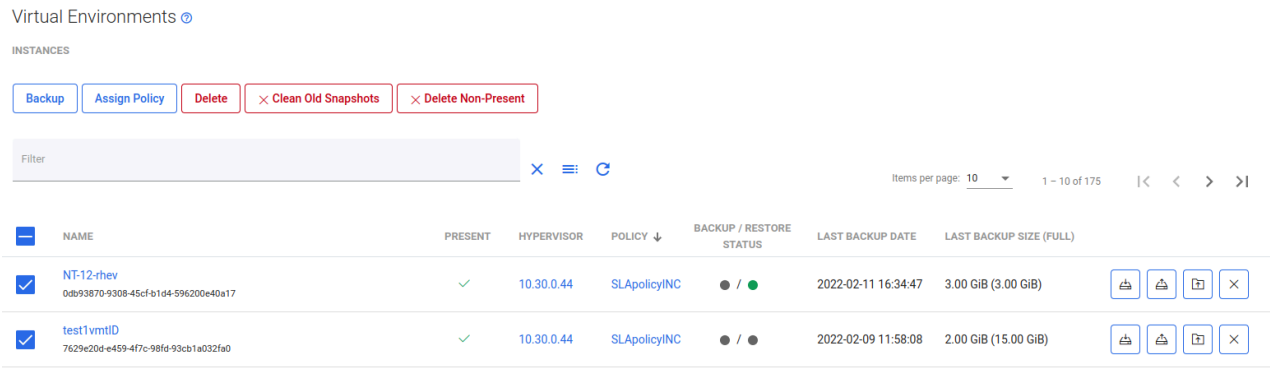

Finally select a backup type, backup destination, when a task should start, priority and click on **Backup** button.

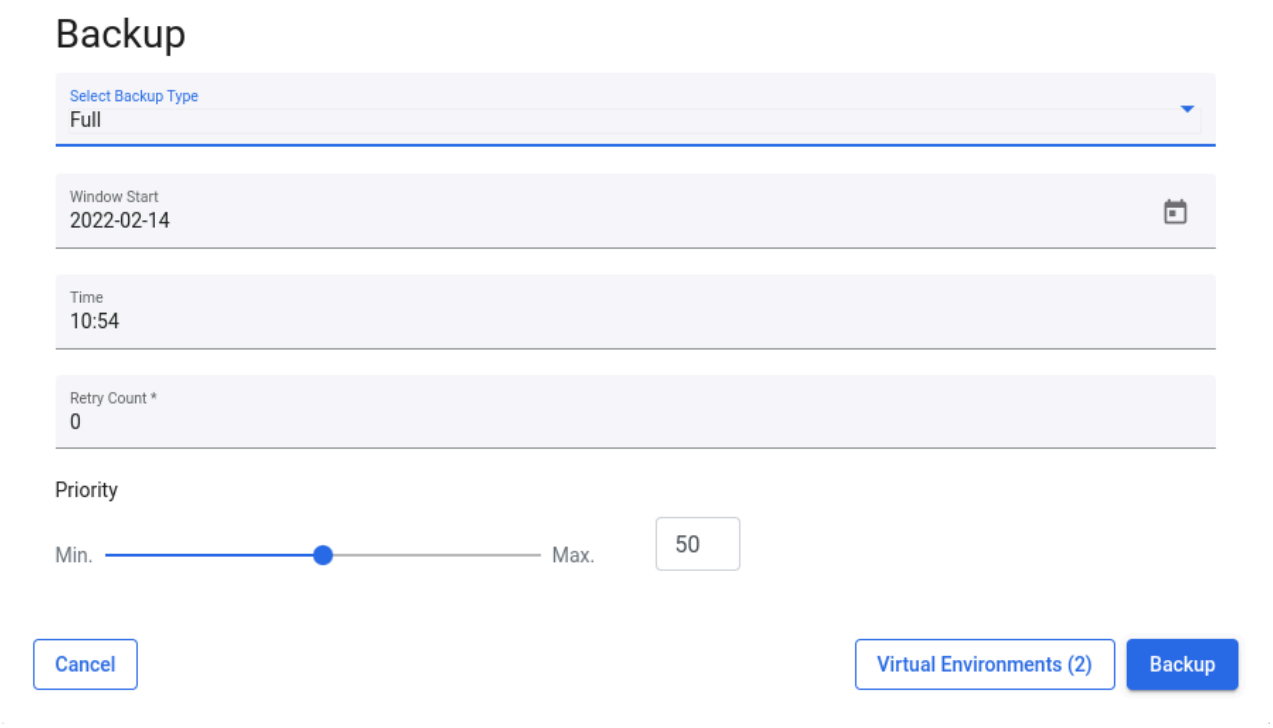

On the Tasks Console, you can see the progress of the backup task.

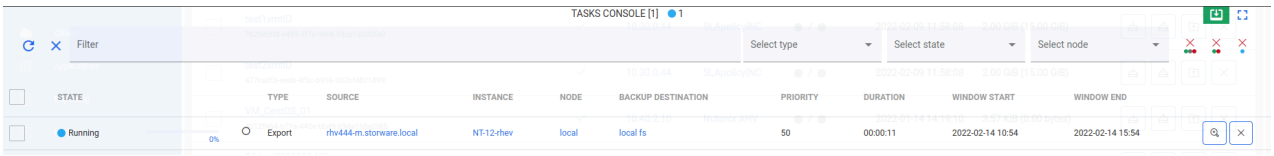

> **Note:** You can also perform the same action thanks to the CLI interface: CLI Reference ↗

Restore on-demand

Restore on-demand

Restore from virtual environment menu

To restore a single virtual machine on-demand, go to the instances tab under Virtual Environment section. Click on the restore icon next to the virtual machine 

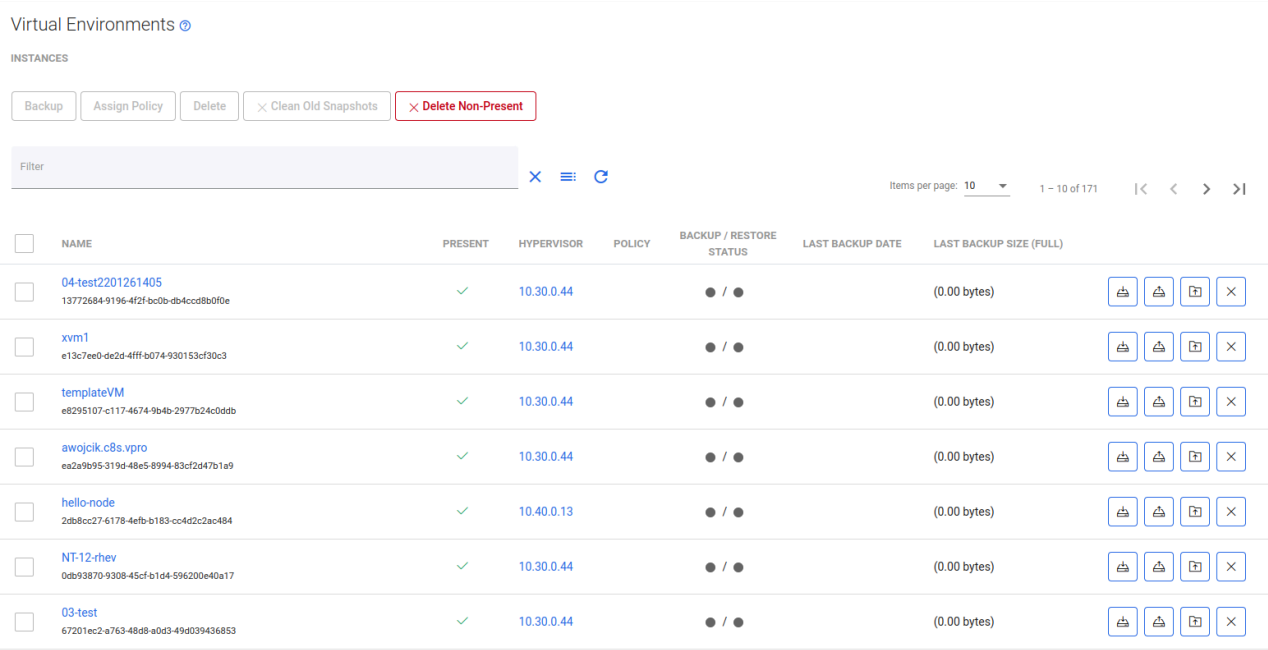

Now you should see a popup window where you might customize restore settings.

Restoring to the filesystem is an option to restore VMs directly to the Data Protector for Cloud Workloads Node storage.
(Remember, the Data Protector for Cloud Workloads user must have the appropriate write permissions to the given path.)

Restoring to a hypervisor or hypervisor manager is allowed for several providers, but not for all (for detailed info go to Data Protector for Cloud Workloads Support Matrix).

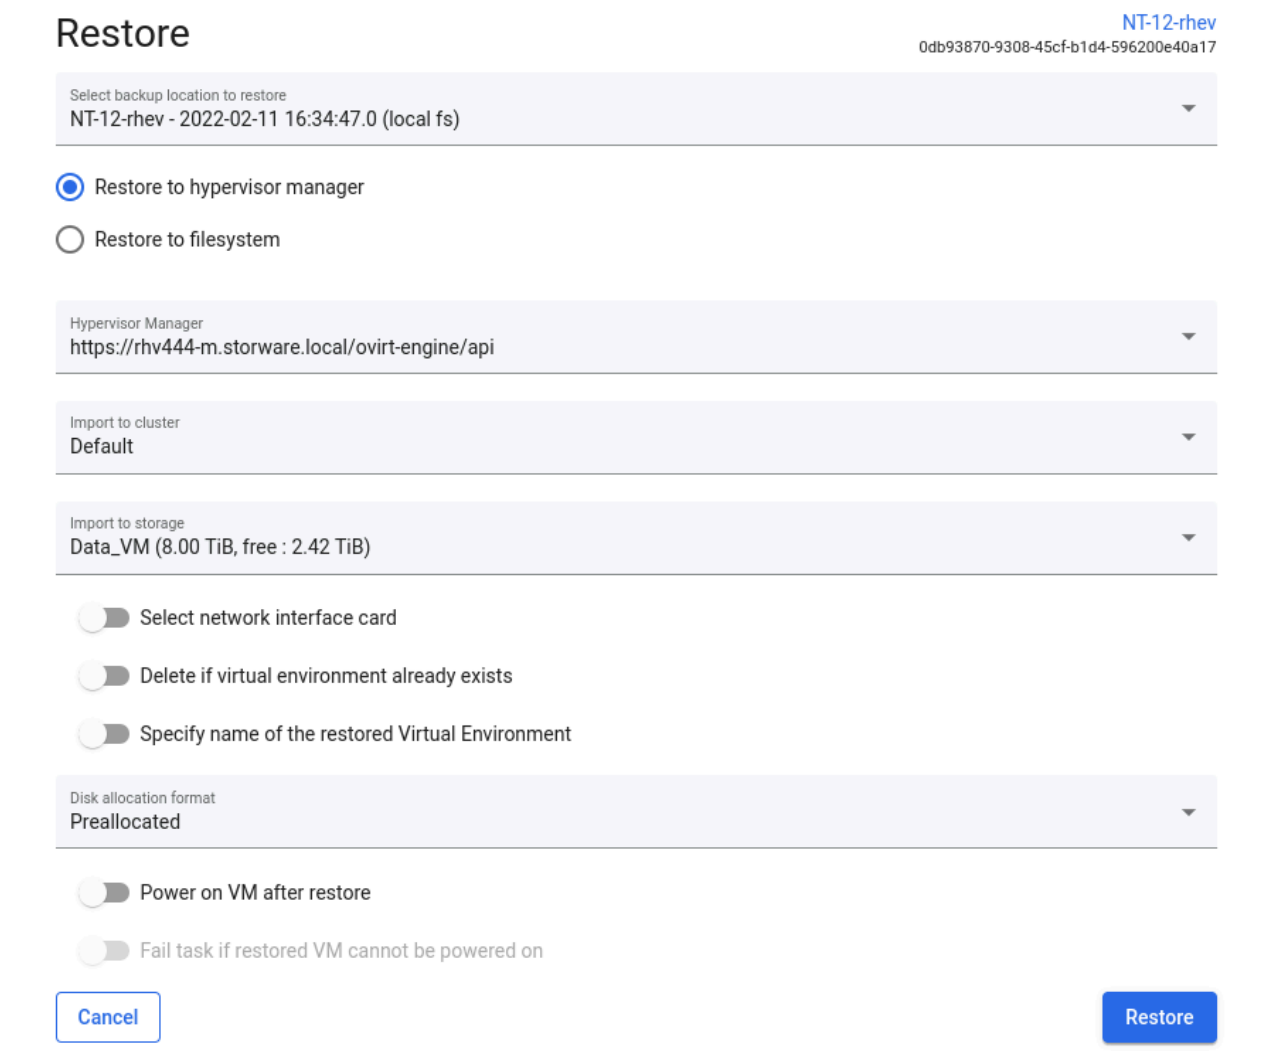

When you restore a virtual machine, disks of this machine can be:

- Restored to specified destination. It can be the same or diffrent datastore/volume type for each disk.
- Restored with changed disk name

Supported platforms: oVirt/OLVM/RHV, Openstack, Citrix and Nutanix AHV.

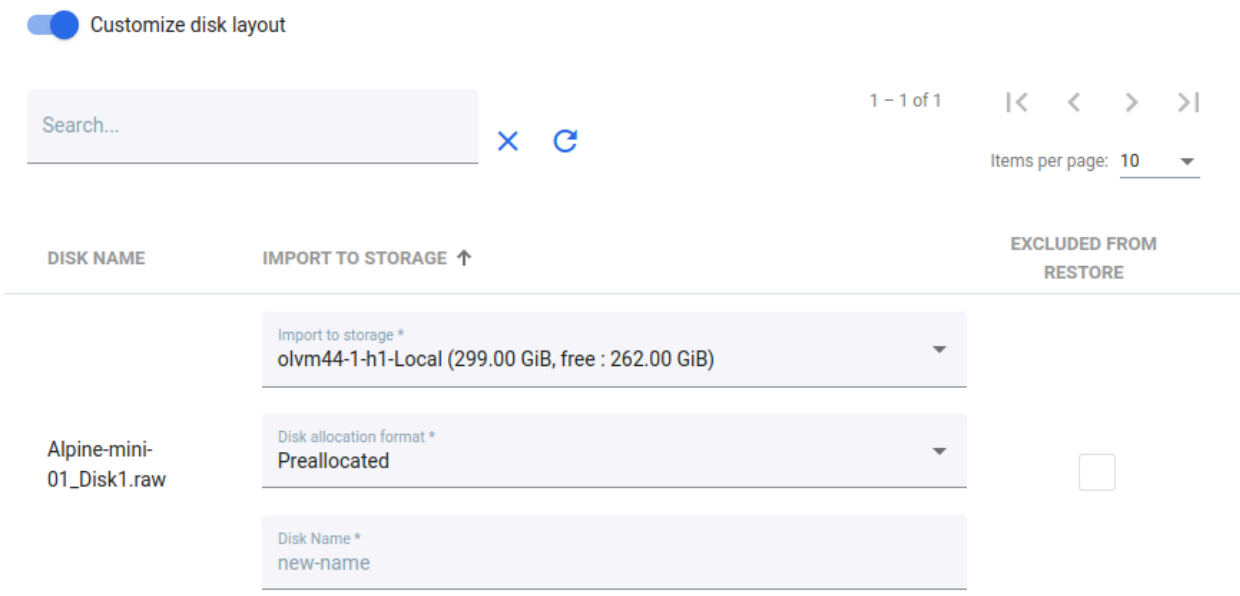

Finally, after customizing the restore, click the restore button.

> **Note:** every platform has some restrictions imposed on the VM name, such as length or characters that can be used. Verify these limits before restoring with a custom name.

Restore on-demand using recovery plans

To restore several virtual machines, you have to use recovery plans. As you can see on the below screenshot, next to the recovery plan policy you can find the same icon that allows you to restore virtual machines 

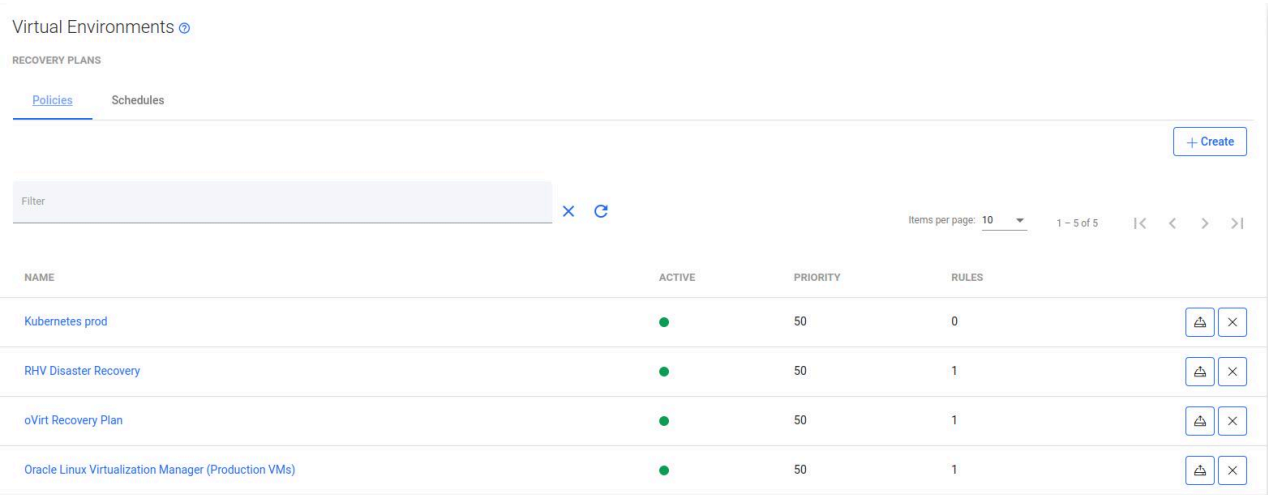

After clicking on it you will see a summary window showing what will be restored.

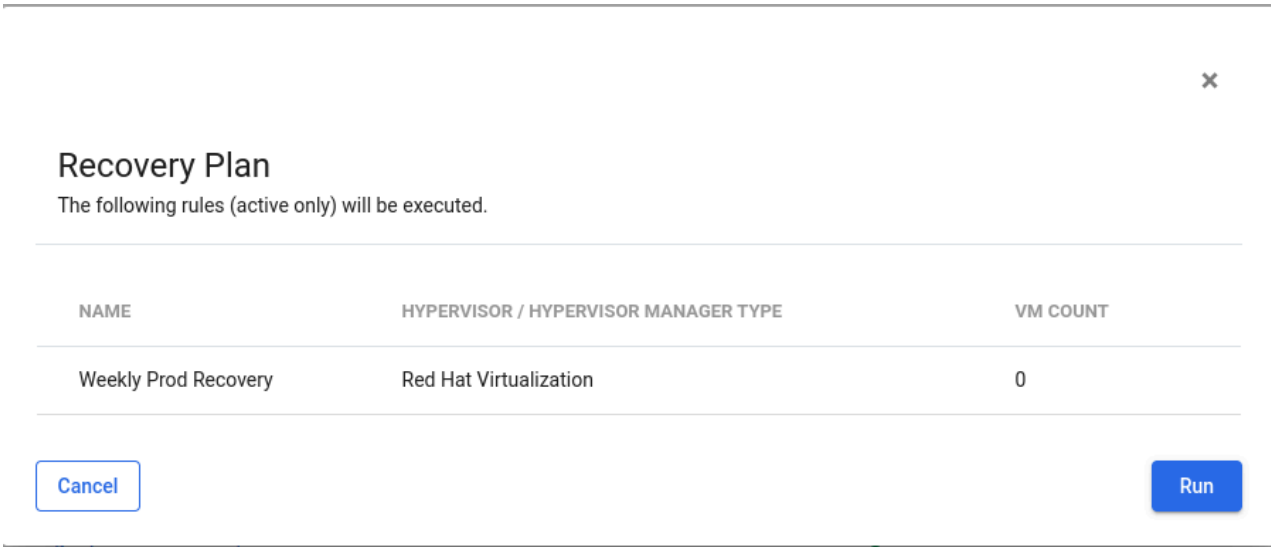

Click Run to start the restore process.

Snapshot Management

Snapshot Management

Data Protector for Cloud Workloads can periodically create snapshots and keep several of them on every VM. To see which hypervisor support this feature, go to Data Protector for Cloud Workloads Support Matrix.

Snapshots are kept in the virtualization platform and are not exported with the backup. It allows quick recovery of a VM without the need to restore and import data backup to the hypervisor. Snapshots cannot replace backup, but allow to increase RPO when used together with backups. Snapshot chains may impact VM performance and occupy significant additional space depending on the storage setup in the environment, so we recommend keeping low values for the number of kept snapshots in the policy, that's less than 3, and check the documentation of your environment and storage setup for aspects related to possible implications of keeping snapshots.

> **Note:**
>
> Snapshot management for on-demand or scheduled operations is not available until the virtual machine has a snapshot management policy assigned to it.
>
> In order to enable snapshot management for VM, you need to follow the steps in this article: Snapshot SLAs

Once VM is assigned, you'll have new buttons that allow on-demand operations. After opening the virtual machine details page, you can see the camera button in the upper right corner.

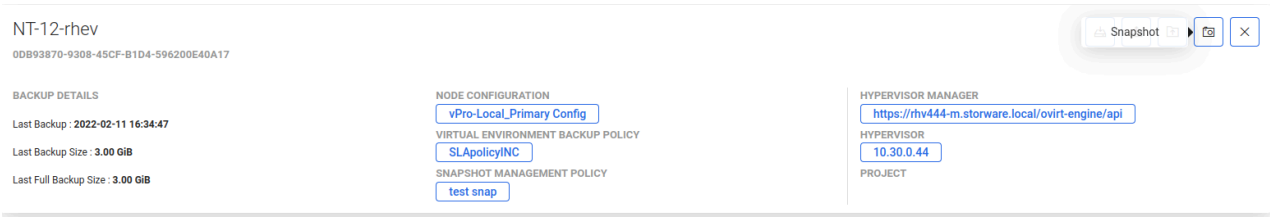

After pressing the button, the snapshot will be taken immediately, without additional confirmation.

You can also revert the snapshot from the virtual machine details page.

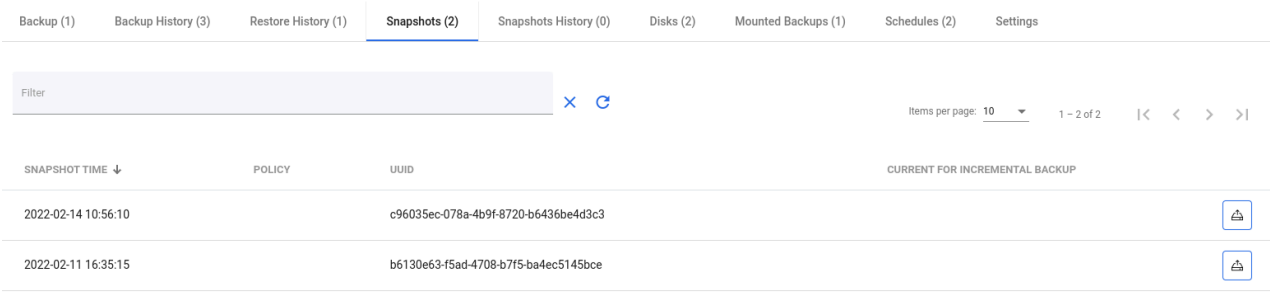

Next to the snapshot is a button that allows you to restore the virtual machine. It actually creates a new virtual machine and keeps the old one (security considerations to protect against the human factor).

Infrastructure

Infrastructure

This section describes how to manage hypervisors and their managers in Data Protector for Cloud Workloads. Inventory that Data Protector for Cloud Workloads needs first to be populated. The first step is always to add a hypervisor manager (if it virtualization platform supports a dedicated manager) or individual hypervisors (if these are not managed, but are stand-alone).

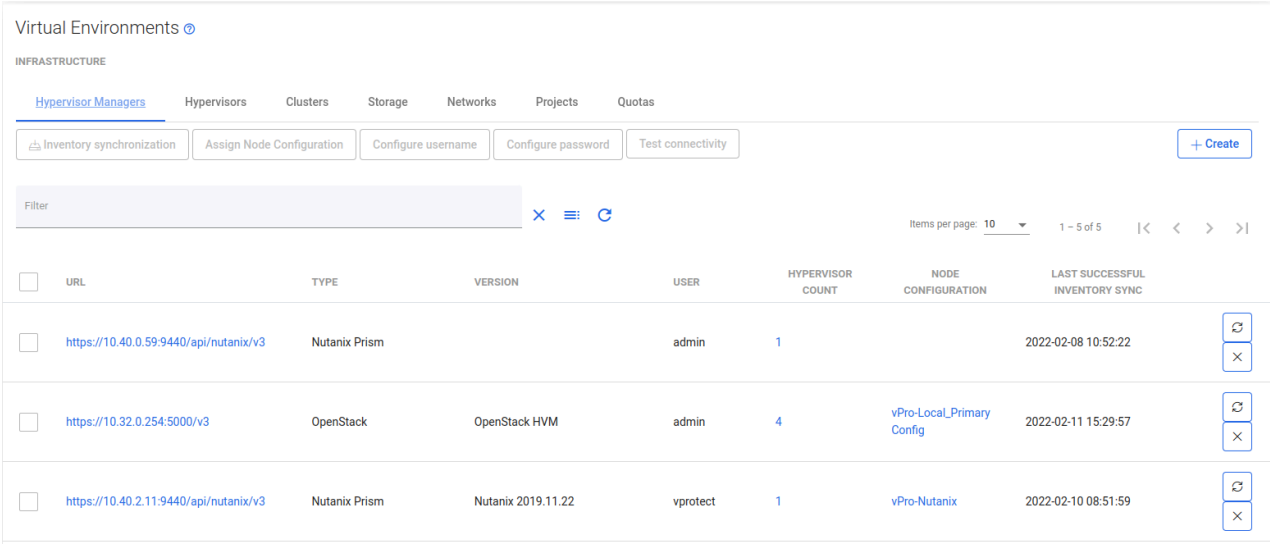

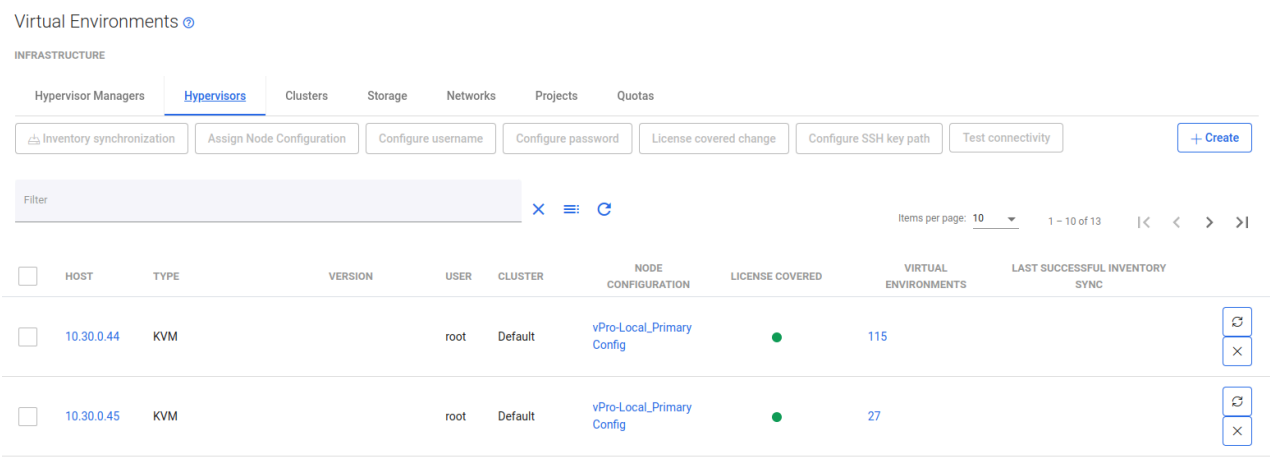

You also can verify if your Hypervisor Storage (datastores/storage repositories/storage domains, depending on how different platforms call it) or Hypervisor Clusters (which corresponds to server pools/ clusters/availability zones) that have been detected.

Click Add Hypervisor Manager / Add Hypervisor to add entries and in general, you always need to provide:

- URL (hypervisor manager - valid URLs are described in the sections describing the setup of a particular virtual platform type) or hostname/IP (hypervisor)
- the node which is responsible for executing tasks in this environment
- backup strategy - if available for a particular platform

Then synchronize inventory (either automatically - there will be a dialog box shown just after saving the form or manually with the button on the right of hypervisor or manager). If Inventory Synchronization tasks (visible in the console at the bottom) completes successfully it also proves that connection was successful, credentials are correct and all of the inventory items have been collected successfully.

Check Hypervisor Storage, Hypervisor Clusters tabs, as well as Virtual Environments → Instances to see the results of inventory synchronization.

> **Note:**
>
> - inventory synchronization executed on the manager level assigns the same node as used for the manager to all hypervisors - you can override it in the Hypervisor tab and assign a different node to handle VMs that reside on a specific hypervisor - this is especially important for scalability and when environments are divided into multiple clusters (disk-attachments strategies may not be able to access disks from different clusters)
> - if you use **disk-attachment** strategy - always execute **inventory synchronization at least once from each hypervisor** - the end result will be the same from the inventory point of view, but each node needs to detect its own Proxy VM ID in the environment to attach disks to the correct VM

Hypervisor clusters and storages

When using Openstack, for each hypervisor cluster and hypervisor storage you can select projects for which it is going to be available in horizon plugin restore dialog.

Selecting "Visible for all projects" will make a cluster/storage visible for all projects, regardless of their time of creation (including ones added after editing this setting).

Deselecting "Visible for all projects" option and selecting some (or all) of existing projects will result in cluster/storage being visible just for selected, already existing projects. The cluster/storage will not be automatically visible for any newly added projects.

This setting will also affect filtering clusters and storages by project selected in restore modal dialog or while editing restore setting for recovery plan rule.

Hypervisor SSL certificates management

> **Note:** SSL certificate management is available only for:
>
> - Red Hat Virtualization
> - Citrix Hypervisor (XenServer)
> - VMware vSphere

When you first time connect to the hypervisor host, the certificate will be fetched and stored. The certificate is used to validate the authenticity of the hypervisor host during the inventory, backup, and restore operations.

If the certificate will change the connection to the hypervisor will be failed, and in the console and logs, you will find proper information.

When you will have a new trusted certificate deployed on the hypervisor host, you can remove an old one from the product, and during the next connection, the new certificate will be fetched and used.

Validating the certificate

To validate fetched certificate:

1. Go to **Virtual Environments** → **Infrastructure**
2. Click on Hypervisor which certificate you would like to validate
3. Click Certificates tab

In this view you can check vertificate fetched from the selected hypervisor host.

Automatically trust all certificates

You can skip the certificate validation and automatically trust all certificates.

1. Go to **Virtual Environments** → **Infrastructure**
2. Click on Hypervisor which certificate you would like to validate
3. Click Certificates tab
4. Turn on switch "Trust all certificates"
5. Confirm your action

From this moment the certifiactes will be not validated.

Removing certificate

To remove stored certificate:

1. Go to **Virtual Environments** → **Infrastructure**
2. Click on Hypervisor which certificate you would like to validate
3. Click Certificates tab
4. Click on Clear button

5. Confirm your action

Certificate will be removed and a new certificate will be fetched during next connection to hypervistor host.

Quotas

Quotas manage the number of VM backups and restores in projects. Quotas uses user-defined rules to control the number of backups. Each Rule has two thresholds: **SOFT** and **HARD**. **SOFT** limits only warn you when a certain limit is exceeded. **HARD** limits prevent the execution of tasks that have exceeded the specified limits. When any of the rules are exceeded, the task fails or warnings are placed on the VM and on backup or restore. For each limit type you need to specify time frame in which rule will be applied. Quotas can be activated or deactivated at any moment, additionally each rule can be activated or deactivated.

To create new Quota, open **Infrastructure** tab under **Virtual Environments** section and go to **Quotas** then click on **Create** button on the right.

Backup SLAs

Backup SLAs

Backup SLAs allow you to group backup policies and schedules for multiple VMs. In general, a VM can have at most only one backup policy to always have easy-to-interpret configuration.

Each policy can have multiple schedules assigned so that you can define more complex schedules in which backups are executed even multiple times each day and with a different backup type.

> **Note:**
>
> - schedules define the type of backup - full or incremental
> - the approach we highly recommend is to create a schedule for **periodic full** backup and always assign at least 1 such schedule in backup SLAs
> - in order to create incremental backups, you always need to have **at least 1 incremental backup schedule** and run at least one full backup
> - in most virtualization platforms supported by Data Protector for Cloud Workloads, it is required to **keep the last snapshot** for future incremental backups

Policies

Policies

Policies allows you to group virtual machines in many ways. For example, based on the type of hypervisor.

To create a new backup policy, open the Backup SLAs tab under the Virtual Environments section and click on **Create** the button on the right.

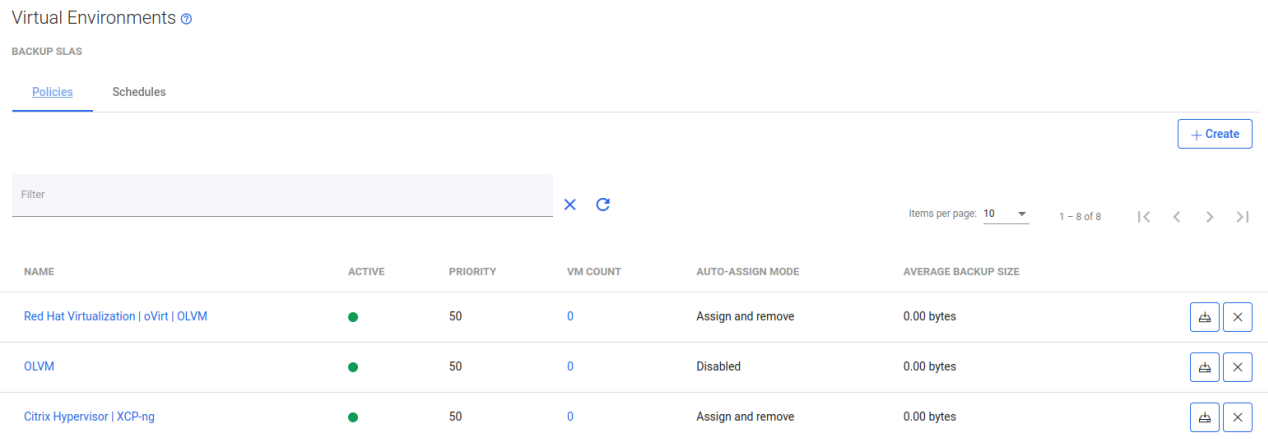

Now you should see the policy wizard with 5 main sections.

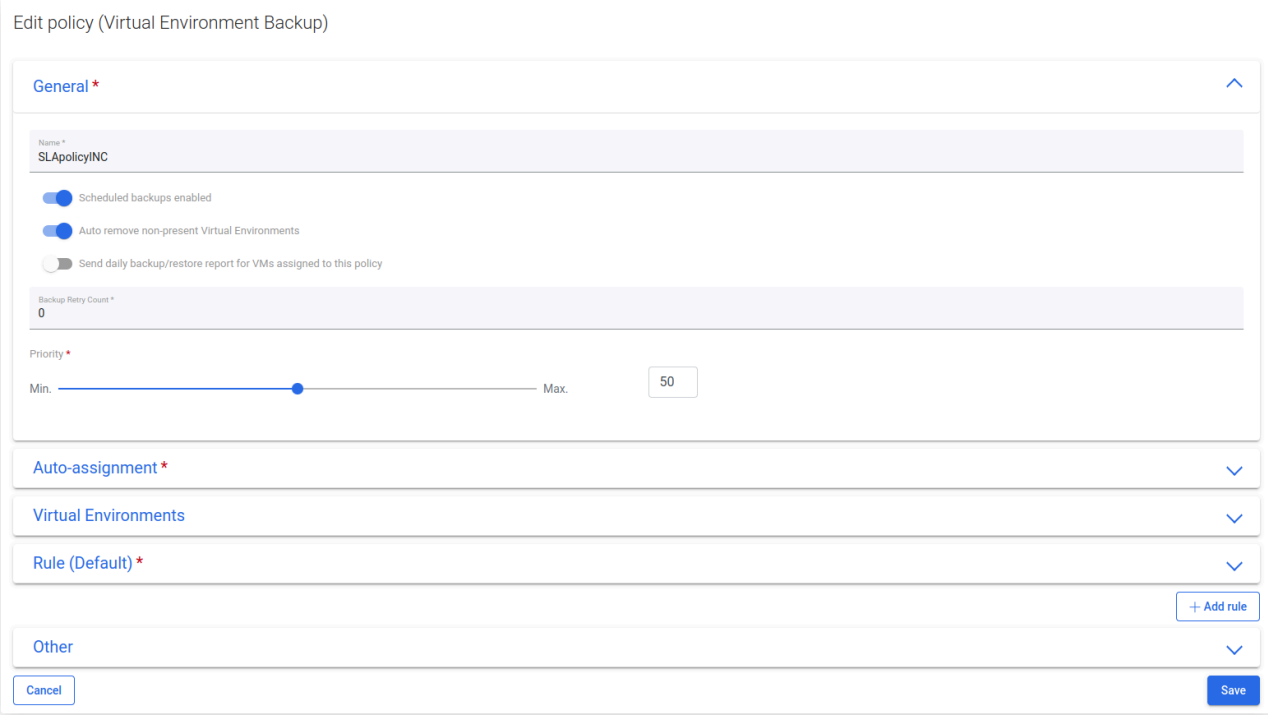

General

Under this section you can set up:

- The policy name
- Switch on/off auto-remove non-present virtual environments
- Set the priority for tasks
- Backup retry count

Auto-assignment

In this section you can configure automatic policy assignment based on certain criteria:

- Mode
  - Disabled
  - Assign only
  - Assign and remove
- Include or exclude rules based on hypervisor tags or regular expressions matching the VM name:
  - regular expression examples:
    - `.*` match any character any number of times
    - `vm-[0-9][0-9][0-9]` - match the name that starts with `vm-` and 3 digits
    - `(prod|uat|dev)-[0-9][0-9][0-9][a-z]?` - match the name that starts with `prod` or `uat` or `dev` prefix, then `-`, then 3 digits and an optional lower-case letter (matching is case-sensitive)
  - exclude rules always take precedence over include rules
  - VMs will not be reassigned to a different policy if they already have a matching policy assigned
  - VMs will be reassigned to a different policy only if the mode is `Assign and remove`, the current policy assignment rules don't match, and other's policy

rules match

- rules are joined with the OR operator, so
  - if **any** rule (tag or matched regular expression) excludes the VM - it will be excluded
  - if **no** rule (tag or matched regular expression) excludes the VM, and **any** rule (tag or matched regular expression) includes the VM - it will be included

- You can also select clusters to match only VMs that belong to them.

Virtual Environments

Here you can easily select virtual machines manually.

Rule

This section is used to select the backup destination.

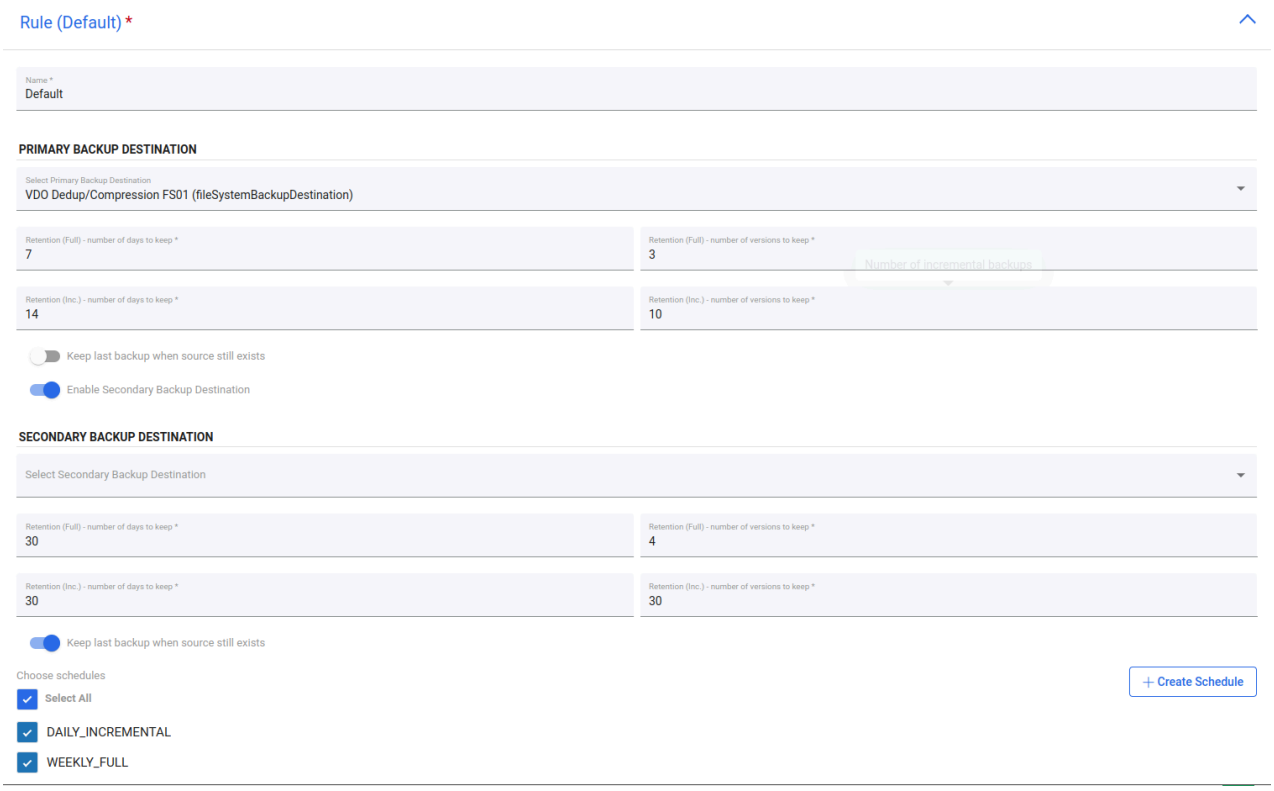

> **Note:** You can select Primary and Secondary Backup Destintanion in one Rule

You can also set here **Retention** settings for your backups. You can use a number of days and versions for full and incremental backups.

If you have already created a schedule, you can also select it or Create New Schedule.

Retention

Data Protector for Cloud Workloads handles retention for all backup destinations. There are 4 properties that define how long backup should be kept in the backup destination:

- `Retention (Full) - no. of versions to keep` - number of full backups
- `Retention (Inc.) - no. of versions to keep` - number of incremental backups
- `Retention (Full) - no. of days to keep` - number of days to keep a full backup
- `Retention (Inc.) - no. of days to keep` - number of days to keep an incremental backup

> If you are using Synthetic File System backup destination, you have only two options for retention:
>
> - `Retention - no. of versions to keep` - number of full backups
> - `Retention - no. of days to keep` - number of days to keep a full backup

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.

Immutable backup

An immutable backup is a feature that is used to prevent modification or deletion of backups for a predetermined period. That is, immutable backups are read-only until their retention period expires.

Settings

This is an optional section with the following option:

- Quiesce/freeze before snapshot
- Fail the rest of the backup tasks if more than xx% of the EXPORT tasks have already failed
- Fail the rest of the backup tasks if more than xx% of the STORE tasks have already failed
- Visibility for all projects

At the end, save the settings.

Schedules

Schedules

Schedules allow you to invoke specific policies periodically. This allows you to back up multiple VMs automatically.

A schedule defines when and on which days VMs should be backed up. To define a new schedule, open Backup SLAs under the Virtual Environments section and go to the Schedules tab, then use the + Create button.

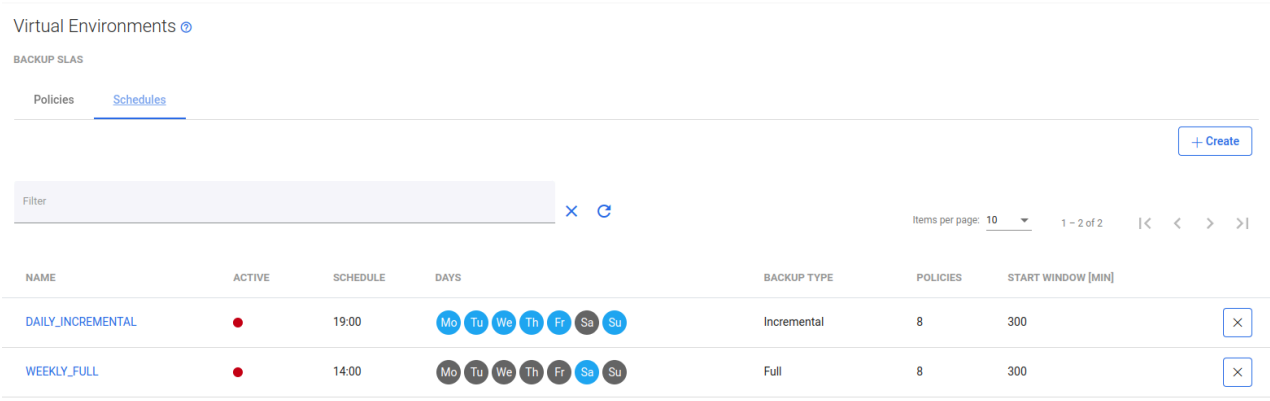

Now enter the properties:

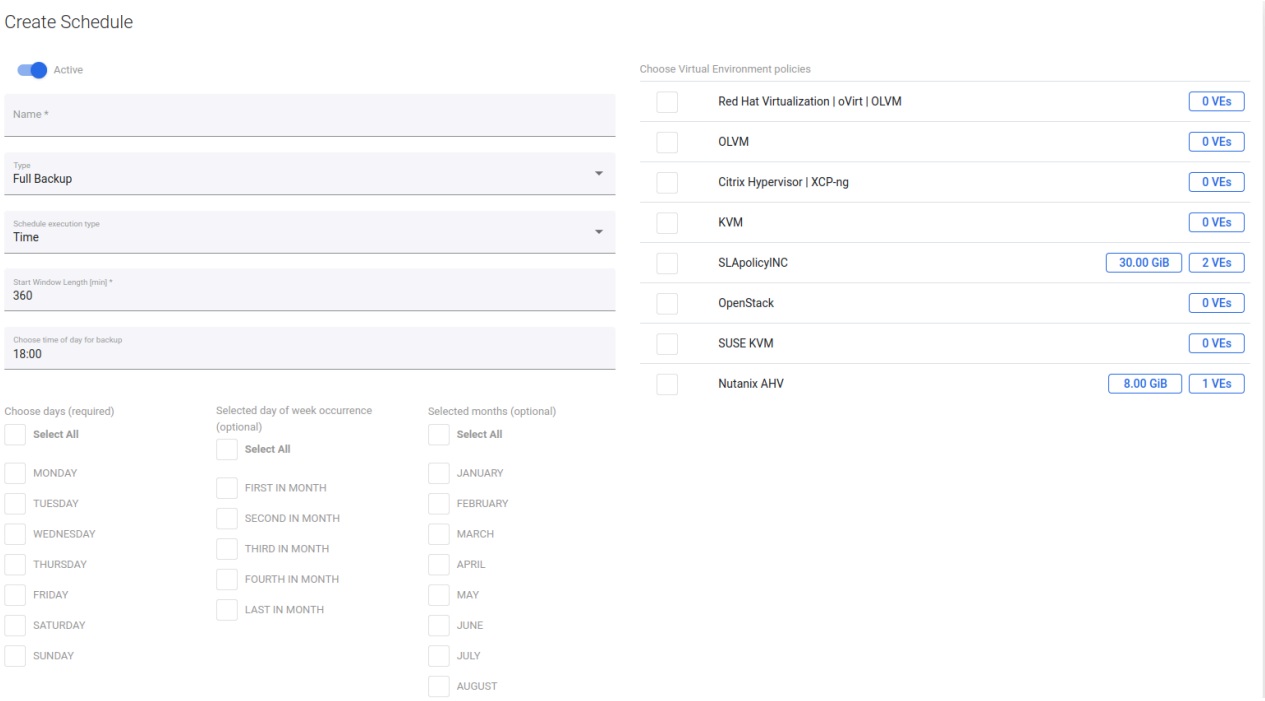

- `Schedule Active` - enable or disable executing schedule
- `Name` - schedule name
- `Backup Type` - defines the backup type: full or incremental
- `Execution Type` - choose the time or interval mode
- `Start Window` - defines for how long since the task start time scheduled tasks are allowed to be executed
- `Choose time of day` - for the time execution mode, this defines when the task should be added to the queue
- `Choose time of interval start` - for the interval execution mode, this defines when tasks should start
- `Choose time of interval end` - for the interval execution mode, this defines when tasks should end
- `Frequency` - defines how often the task will be executed during the interval
- `Choose days` - the last required parameter, select the days of the week on which the task will be performed

You can also use optional parameters to further personalize the backup time or select a virtual environment policy if it has been previously created.

> When you set the time with a user in a certain time zone, you specify a point in time at which you want the schedule to start. Changing the timezone doesn't change this point in time, it's converted to your timezone. The time displayed to the user is calculated based on the server time.

Snapshot SLAs

Snapshot SLAs

Data Protector for Cloud Workloads can periodically create snapshots and keep several of them for every VM. To see which hypervisor support this feature, go to Data Protector for Cloud Workloads Support Matrix.

> **Note:** Snapshot management for on-demand or scheduled operations is not available until the virtual machine has a snapshot management policy assigned to it.

Policies

Policies

In order to enable snapshot management for VM you need to do the following steps:

Go to Snapshot SLAs under Virtual Environment section and create a new Snapshot Management policy 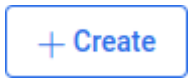

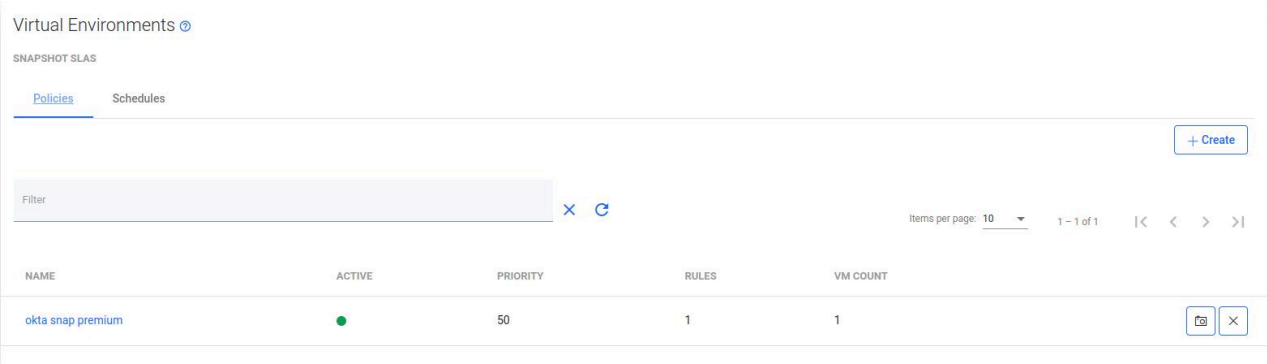

As well as other types of Policies, you'll also find 4 main sections here.

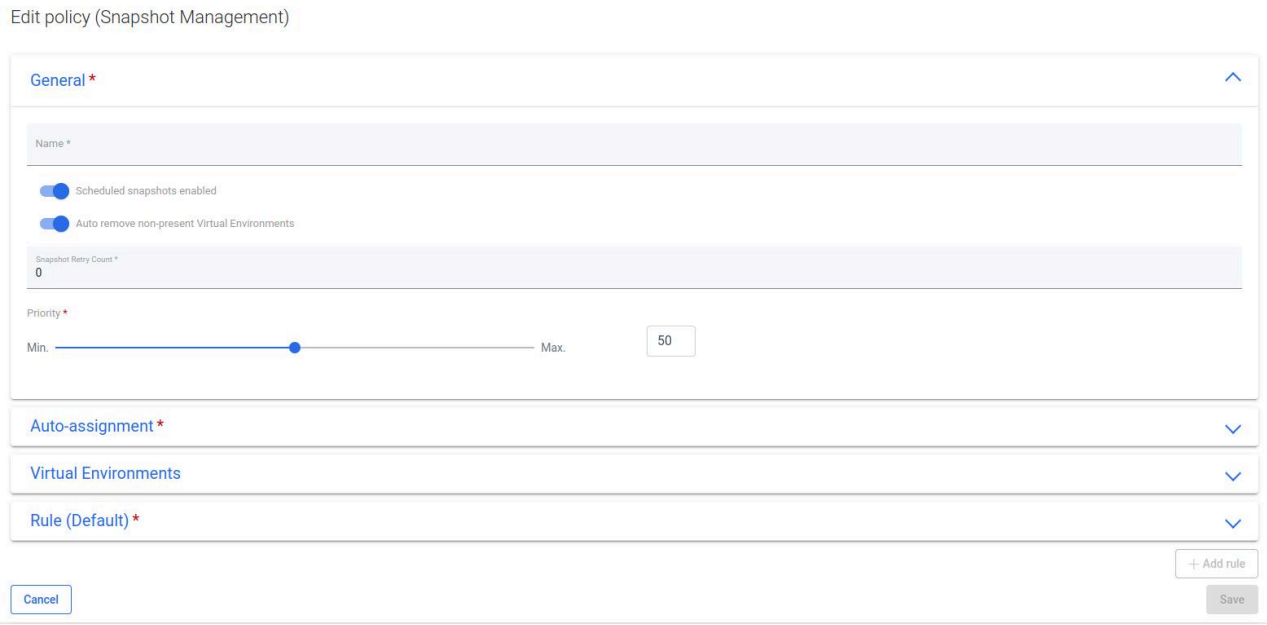

General

Under this section you can set up:

- Name of policy
- Switch on/off auto remove non-present virtual environments
- Set priority for tasks

Auto-assignment

In this section you can set up:

- Mode
  - Disabled
  - Assign only
  - Assign and remove
- Include or exclude rules based on hypervisor tag's or regular expression matching VM name:
  - regular expression examples:
    - `.*` match any character any number of times
    - `vm-[0-9][0-9][0-9]` - match name that starts with `vm-` and 3 digits
    - `(prod|uat|dev)-[0-9][0-9][0-9][a-z]?` - match name that starts with `prod` or `uat` or `dev` prefix, then `-`, then 3 digits and optional lower-case letter (matching is case-sensitive)
  - exclude rules always take precedence over include rules
  - VMs may will not be reassigned to the different policy if they already have matching policy assigned
  - VMs may will be reassigned to the different policy only if mode is `Assign and remove`, current policy assignment rules don't match, and other's policy rules match
  - rules are joined with OR operatorator, so
    - if **any** rule (tag or matched regular expression) excludes VM - it will be excluded
    - if **no** rule (tag or matched regular expression) excludes VM, and **any** rule (tag or matched regular expression) includes VM - it will be included

- You can also select clusters to match only VMs that belong to them

Virtual Environments

In this place, you can select virtual machines manually in a simple way.

Rule

Provide retention settings - how many snapshots (created by this policy) will be kept and for how long. If you have already created a schedule, you can also select it.

Schedules

Schedules

Schedule defines when and on which days snapshots should be created. To define new schedule use the `+ Create` button.

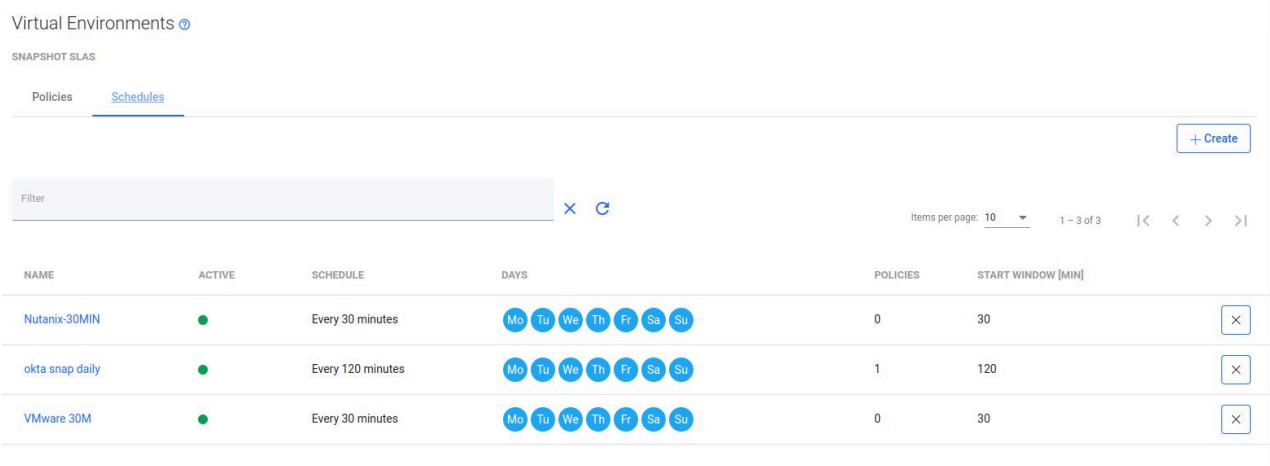

Now provide properties:

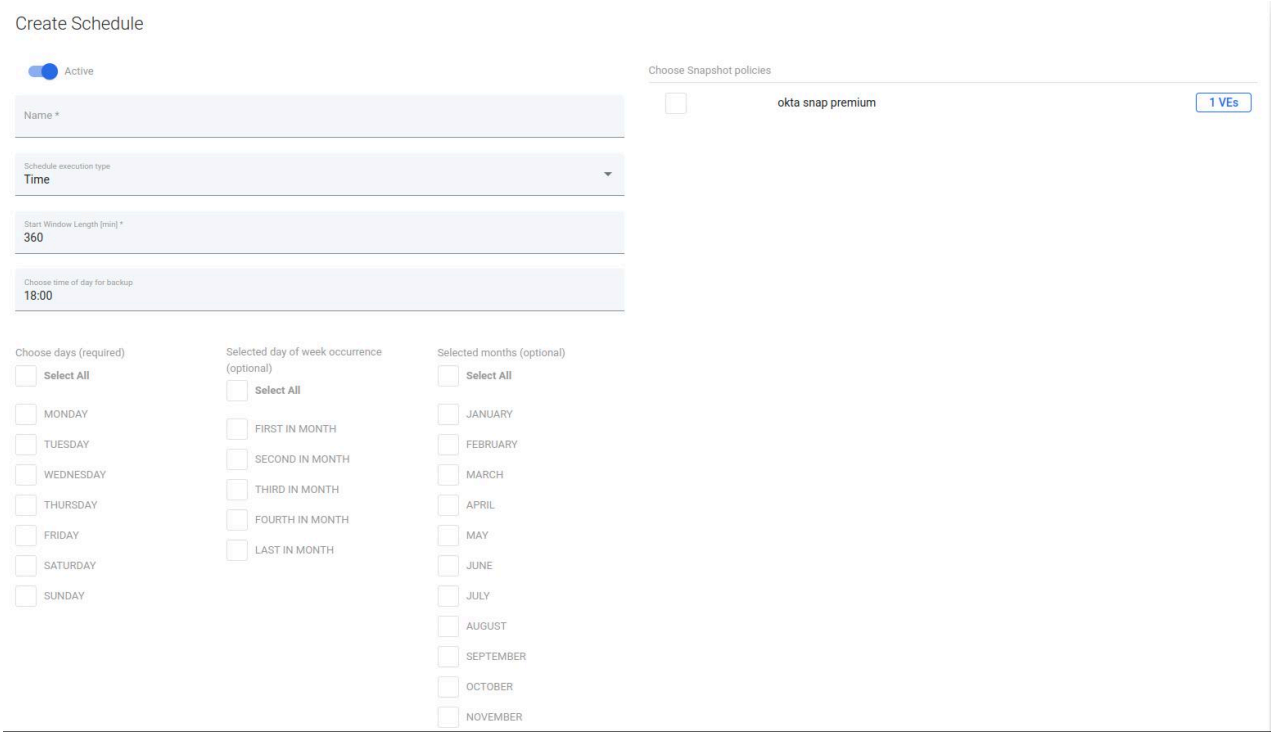

- `Schedule Active` - enable or disable executing schedule

- `Name` - schedule name
- `Execution Type` - choose time or interval mode
- `Start Window` - defines for how long since task start time scheduled tasks are allowed to be executed
- `Choose time of day` - for time execution mode defines when the task should be added to queue
- `Choose time of interval start` - for interval execution mode defines when tasks should start
- `Choose time of interval end` - for interval execution mode defines when tasks should end
- `Frequency` - defines how often task will be executed during the interval
- `Choose days` - last required parameter, select days of the week on which the task will be performed

You can also use optional parameters to further personalize the backup time or select a virtual environment policy if it has been previously created.

> When you set the time with a user in a certain time zone, you specify a point in time at which you want the schedule to start. Changing the timezone doesn't change this point in time, it's converted to your timezone. The time displayed to the user is calculated based on the server time.

Recovery Plans

Recovery Plans

Recovery plans are used to automate DR process so that Data Protector for Cloud Workloads executes multiple restore operations to the target environment with preconfigured settings.

Recovery plans can be executed on-demand or on a scheduled basis (for instance to test the recovery process periodically). Recovery plans consist of rules, each one for a particular virtualization platform, which specify VMs, restore settings, and optionally schedules. Only rules marked as active are executed.

Restore operations will be performed according to the restore/import task limits set in Node Configuration. In target environment new VMs will be created. Names are generated by default (with optional prefix/suffix), optionally original name can be used. For periodic restores, it is common to replace previous VM, which typically means that options to use original VM name and deletion of virtual environment with the same name are checked.

Policies

Policies

To schedule a VM restore using a recovery plan (or execute a recovery plan manually), you must first create a policy.

Go to the Recovery Plans from the left menu under the Virtual Environments section.

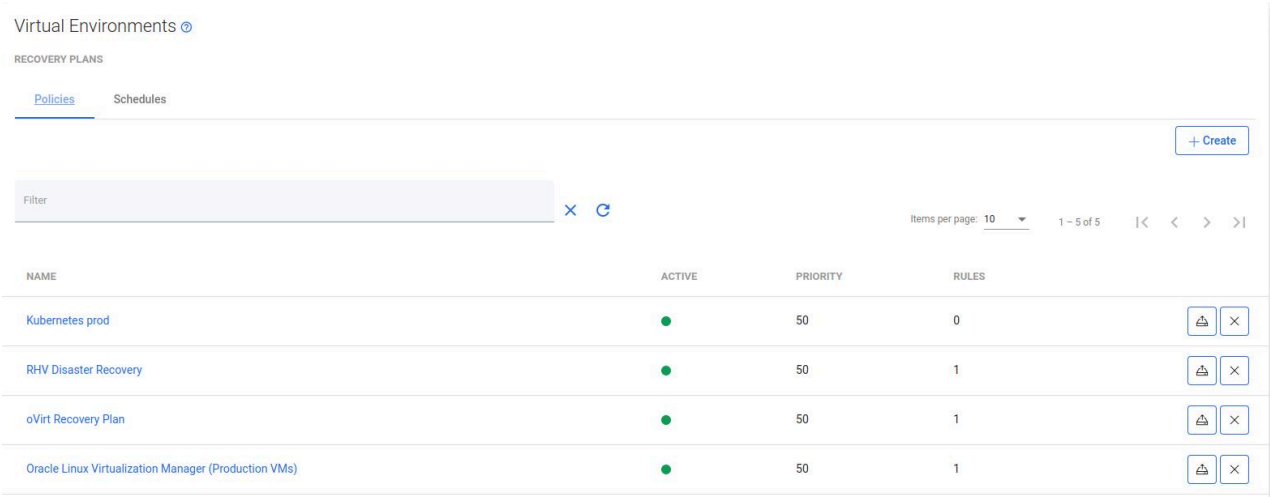

After clicking on + Create provide the name of the policy and set priority. After saving you will be able to add rules using the new button on the right.

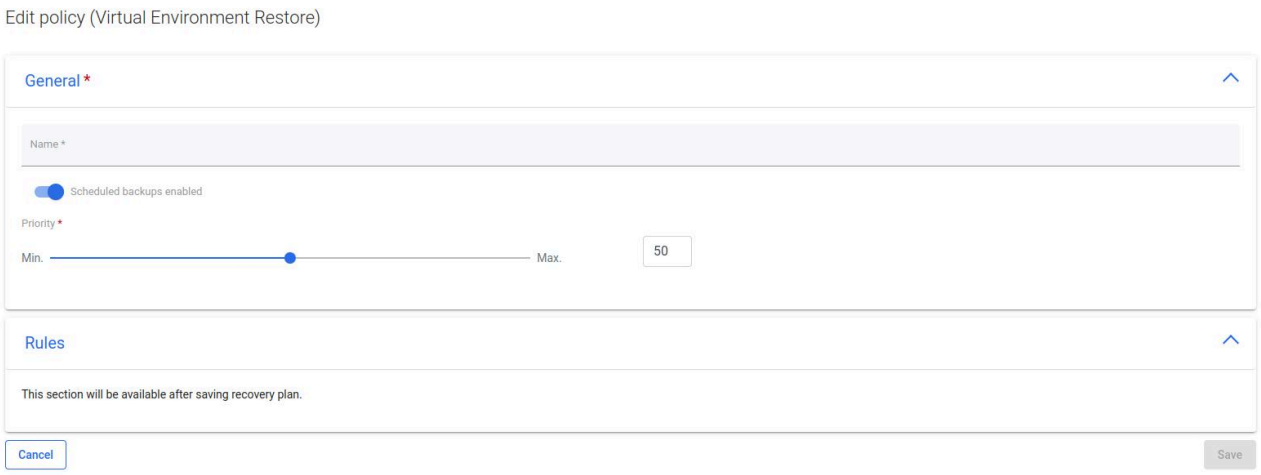

Click on it and customize restore settings.

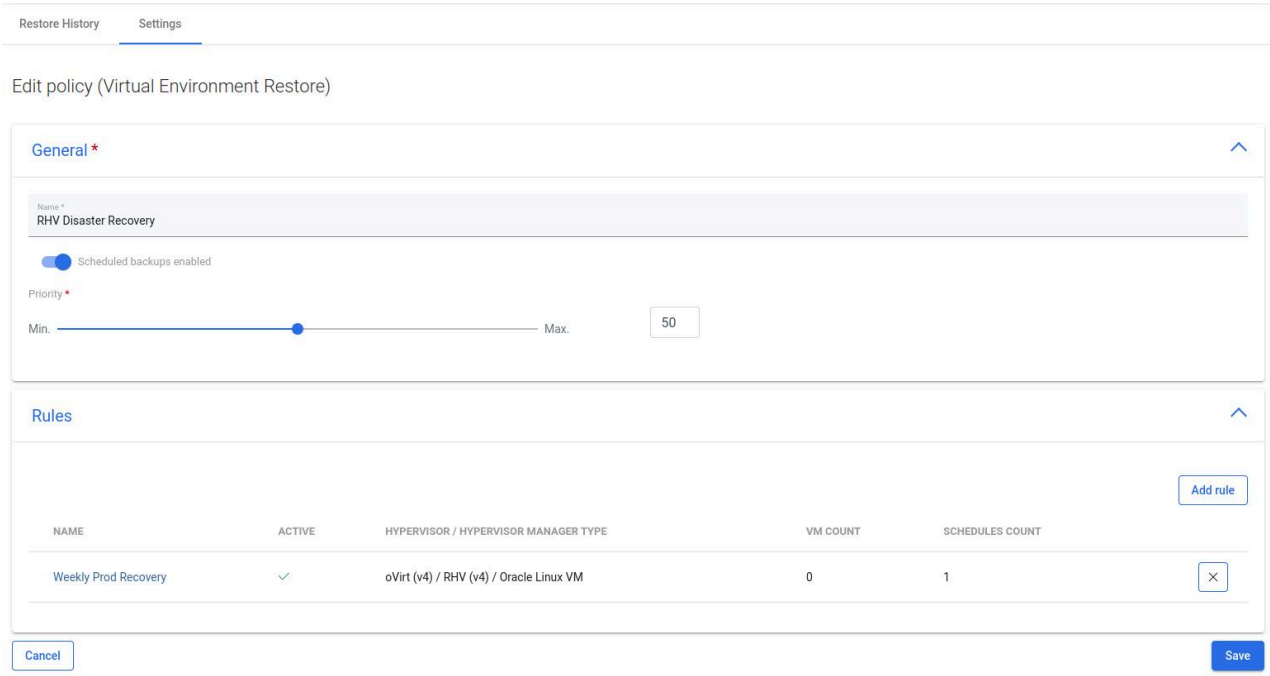

Each rule requires a **name** for easier identification later

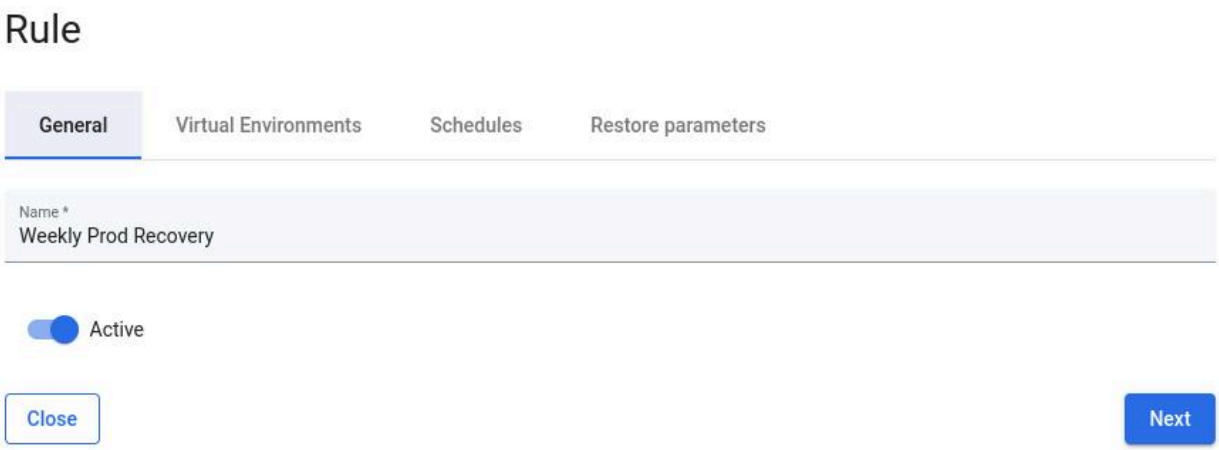

In the **Virtual Environments** tab, you need to select **Hypervisor type** for this rule and corresponding **virtual environments** of this type

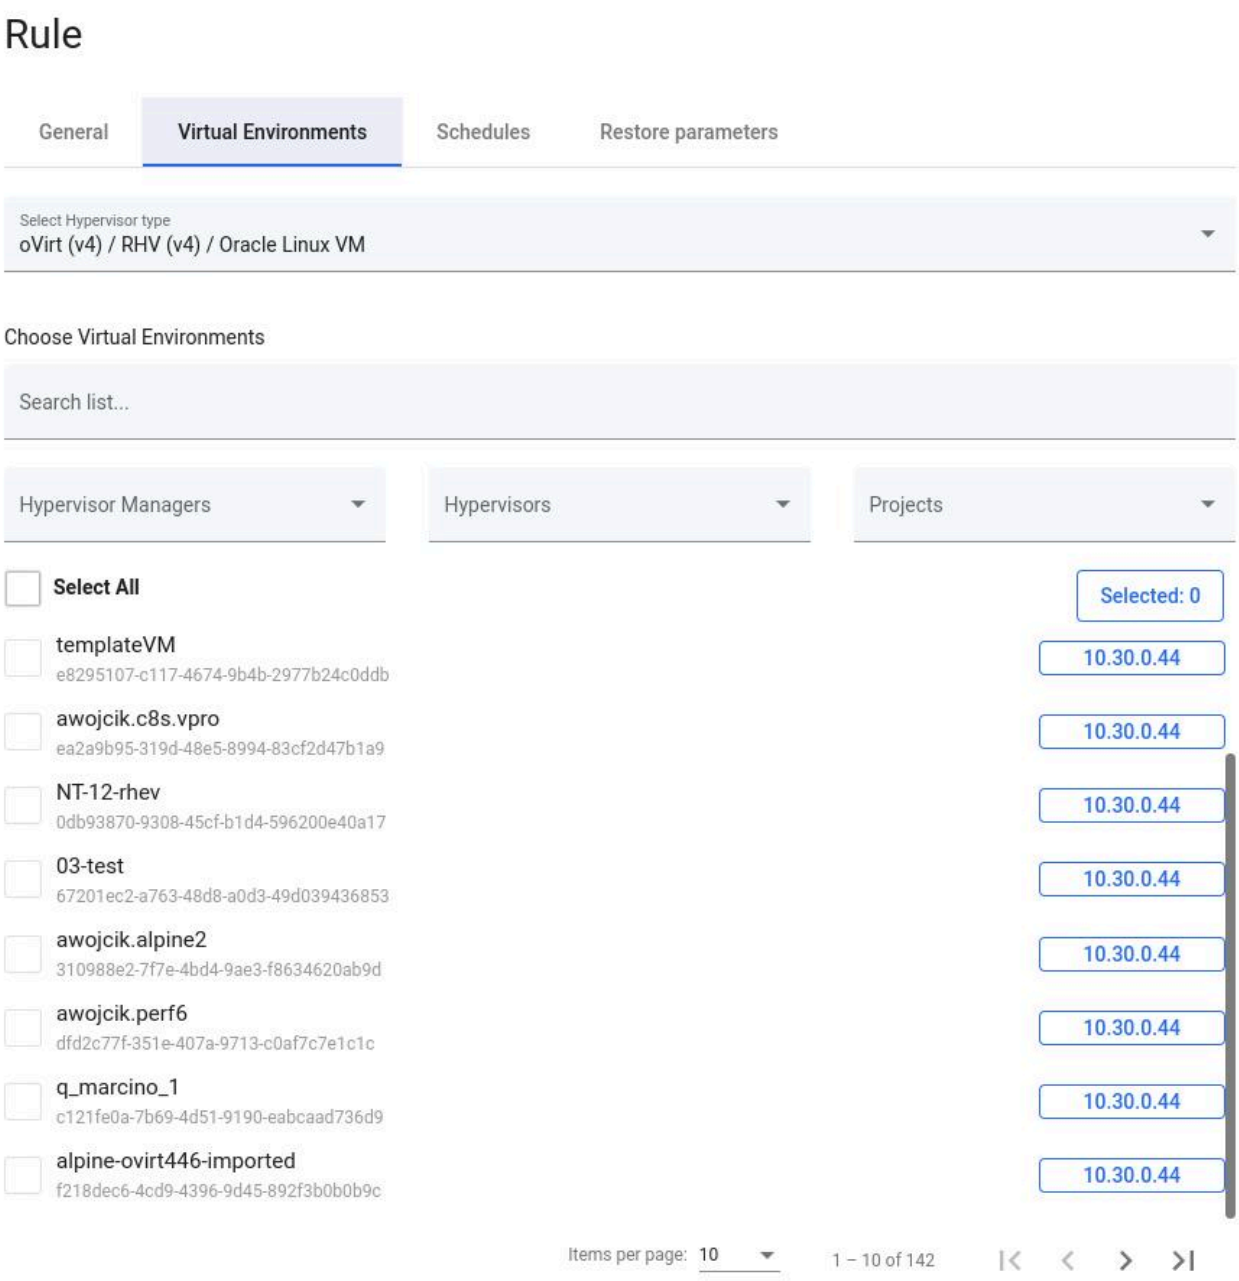

If you previously defined any schedules for recovery plans you can select them in **Schedules** tab

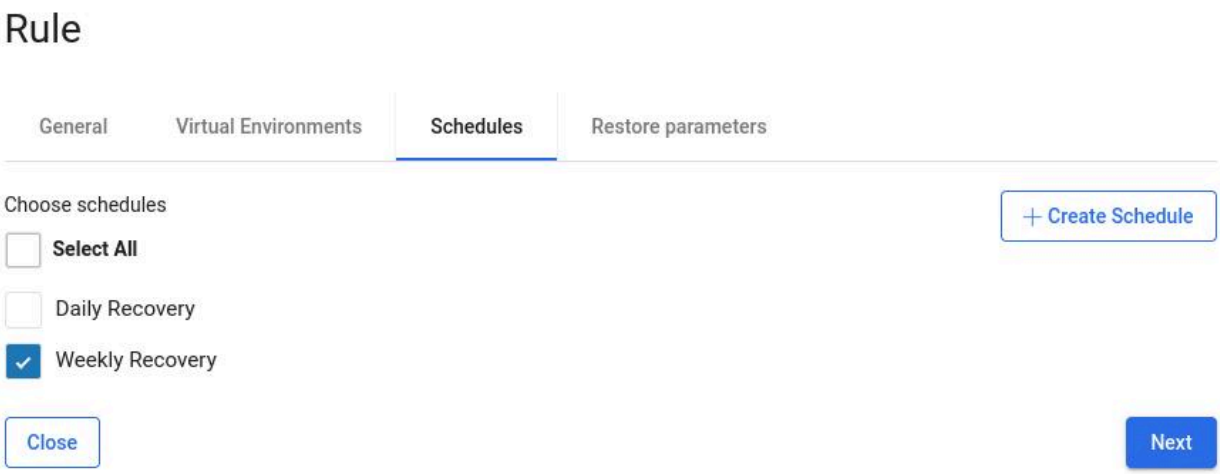

In **Restore Parameters** tab you specify where VMs are going to be restored - compared to regular restore parameters provided in manual restore window, notice that:

- you need to choose ****which backup to restore -** last (regardless of status) **or** last successful
- you may want to use **Delete if Virtual Environment already exists** - which allows Data Protector for Cloud Workloads to remove VM with the same name as the one being restored

Rule

General | Virtual Environments | Schedules | **Restore parameters**

Select backup to restore
Last successful

- (•) Restore to hypervisor manager
- ( ) Restore to filesystem

Hypervisor Manager
https://rhv444-m.storware.local/ovirt-engine/api

Import to cluster
Default

Import to storage
Data_VM (8.00 TiB, free : 2.42 TiB)

- [ ] Generate VM name automatically

Prefix for the name of the virtual machine being restored
DR_Weekly

Suffix for the name of the virtual machine being restored
Webinar

- [ ] Delete if virtual environment already exists

Disk allocation format
PREALLOCATED

- [ ] Power on VM after restore
- [ ] Fail task if restored VM cannot be powered on

Schedules

Schedules

The Schedule defines when and on which days recovery plans should run. To define a new schedule use the + Create

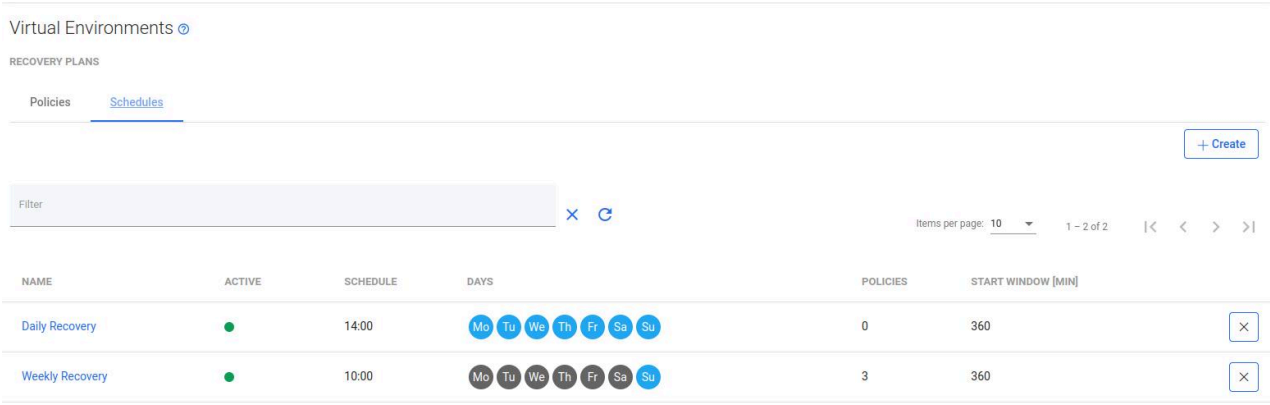

Now provide properties:

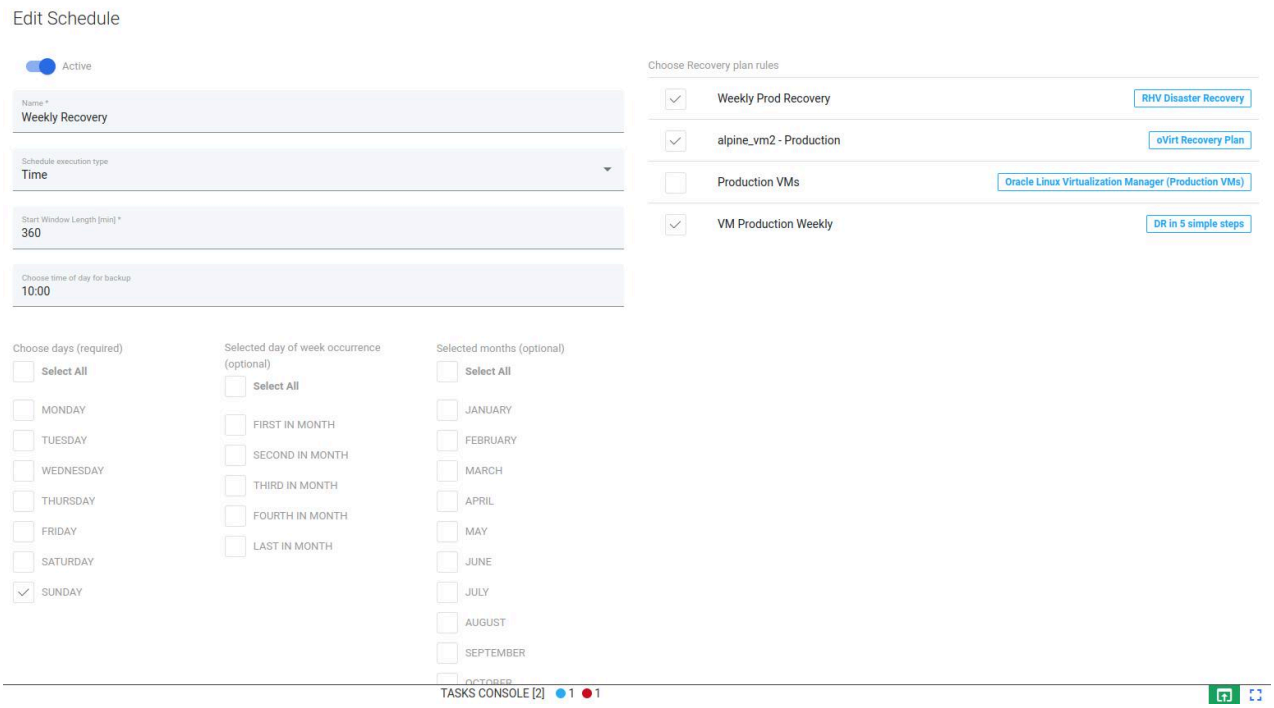

- `Schedule Active` - enable or disable executing schedule
- `Name` - schedule name

- `Execution Type` - choose time or interval mode
- `Start Window` - defines for how long since task start time scheduled tasks are allowed to be executed
- `Choose time of day` - for time execution mode defines when the task should be added to the queue
- `Choose time of interval start` - for interval execution mode defines when tasks should start
- `Choose time of interval end` - for interval execution mode defines when tasks should end
- `Frequency` - defines how often task will be executed during the interval
- `Choose days` - last required parameter, select days of the week on which the task will be performed

You can also use optional parameters to further personalize the backup time or select a virtual environment policy if it has been previously created.

> When you set the time with a user in a certain time zone, you specify a point in time at which you want the schedule to start. Changing the timezone doesn't change this point in time, it's converted to your timezone. The time displayed to the user is calculated based on the server time.

Mounted Backups (File-level Restore)

Mounted Backups (File-level Restore)

> **Note:** To see which hypervisor support this feature, go to Data Protector for Cloud Workloads Support Matrix.

To mount backup go to the Instances tab under Virtual Environment section on the left side menu, then click on the mount icon next to a chosen virtual machine 

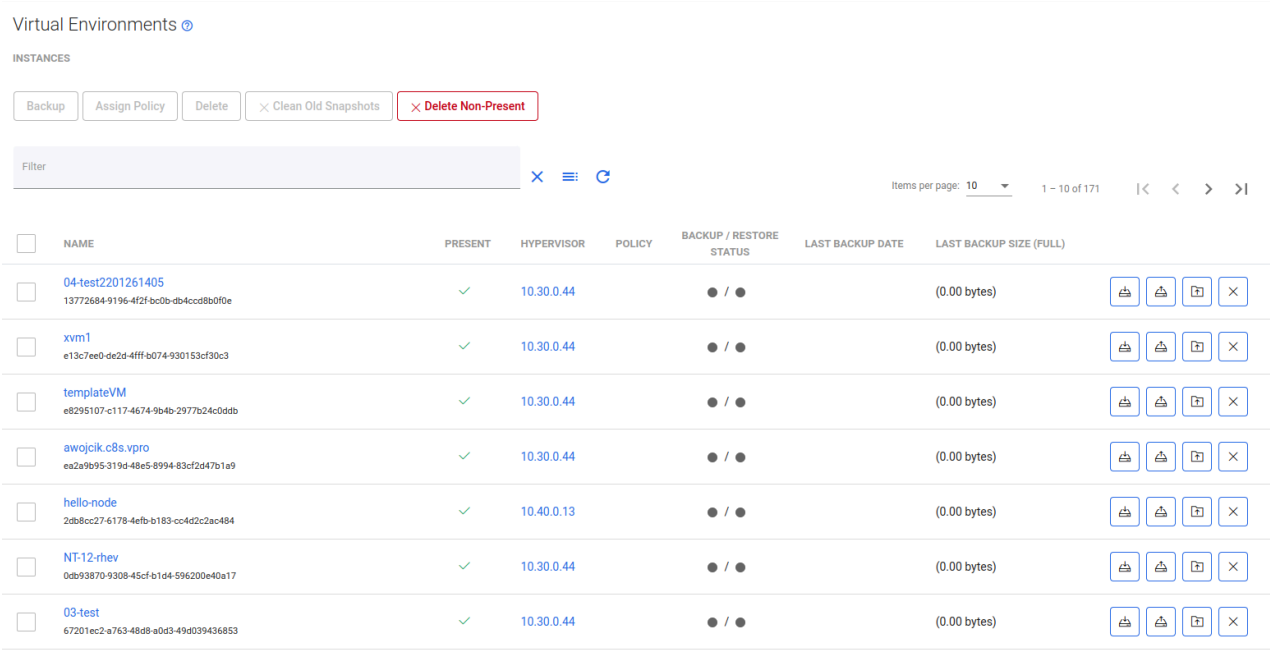

On the popup window, you can select which backup you want to mount and on which node. You can also change the mounting method, but we recommend leaving the default setting "Mount filesystem automatically".

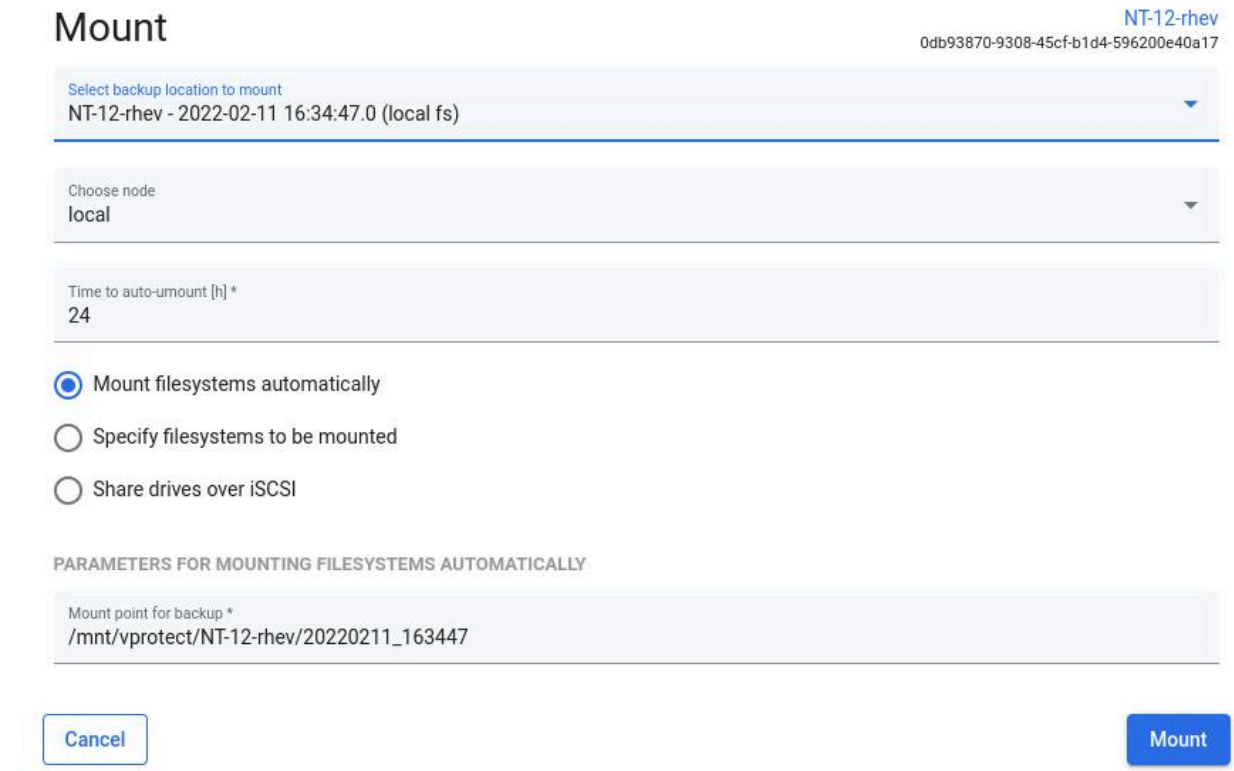

The Mounted Backups tab show mounted Virtual Machine backups on the Data Protector for Cloud Workloads Node.

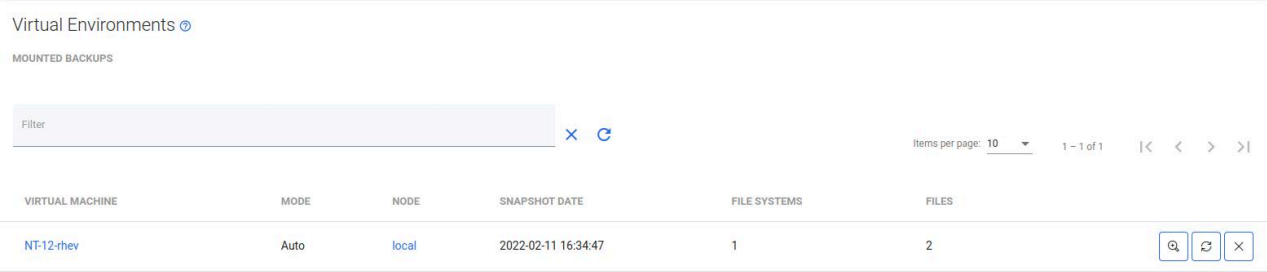

- `VIRTUAL MACHINE` - mounted virtual machine name
- `MODE` - Auto - Data Protector for Cloud Workloads auto-detect filesystems and mount it on path "/mnt/vprotect/". In Manual mode, the user chose a mount point for selected filesystems.
- `NODE` - Data Protector for Cloud Workloads Node responsible for mount job.
- `SNAPSHOT DATE` - date of mounted backup of the VM.
- `FILE SYSTEMS` - a number of mounted filesystems.
- `FILES` - a number of mounted virtual disk images.

Next to every mounted backup you can see three buttons: To unmount backup click on 

To remount backup click on 

To go to the details page of mounted backup click on 

On the details page, you can view some basic information or go deeper and browse files.

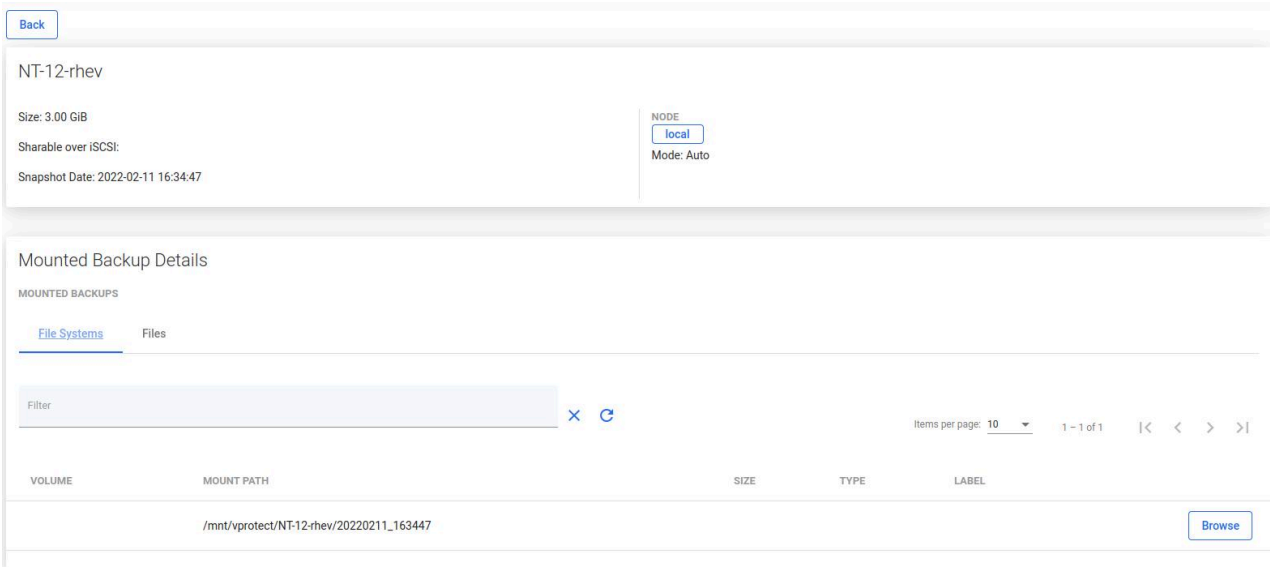

With a web browser, you can obtain even a single file from inside of your virtual machine backup.

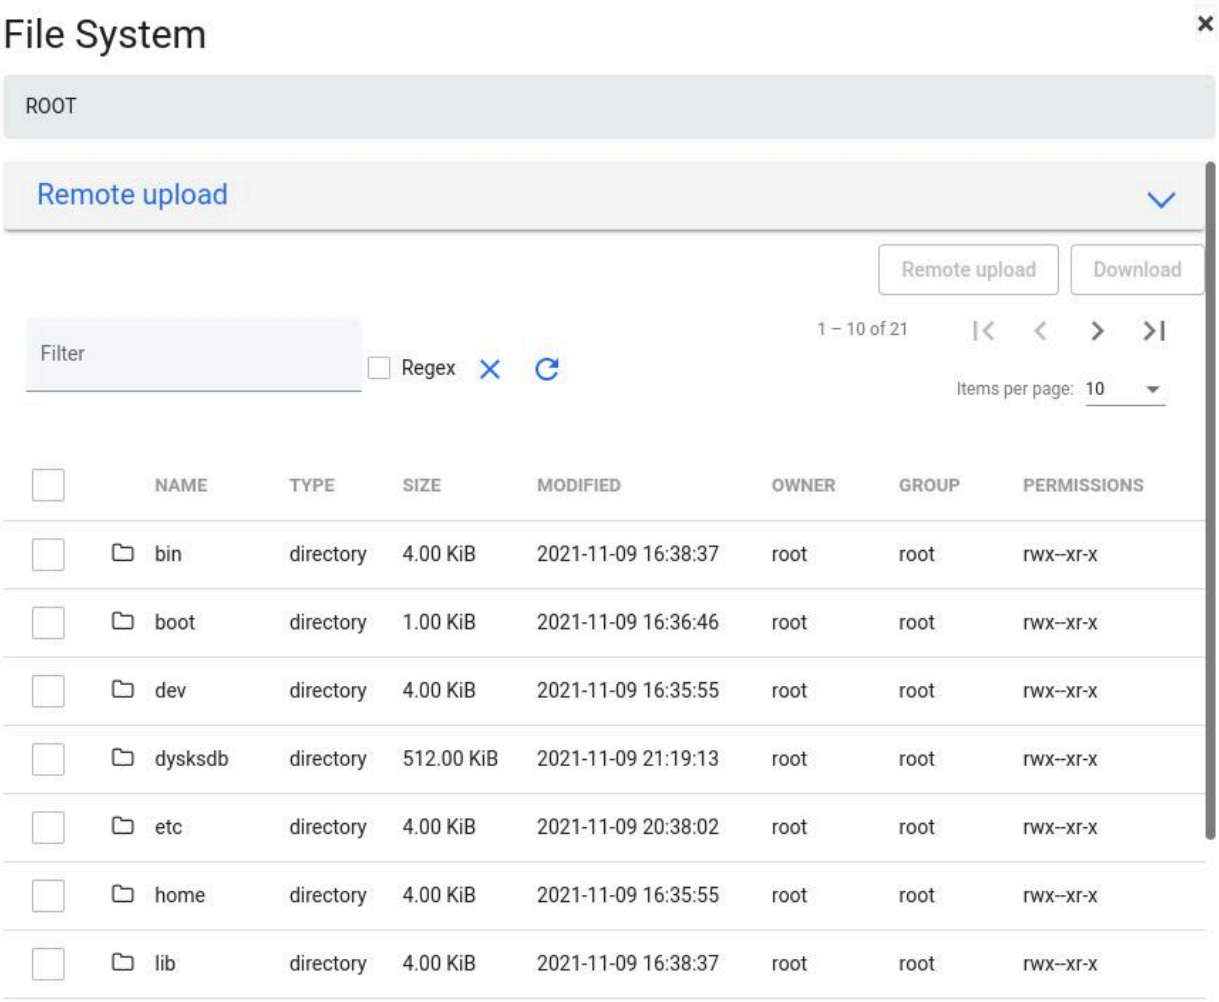
File System
ROOT
Remote upload
Remote upload
Download
Filter
Regex
1 – 10 of 21
Items per page: 10
NAME TYPE SIZE MODIFIED OWNER GROUP PERMISSIONS
bin directory 4.00 KiB 2021-11-09 16:38:37 root root rwx--xr-x
boot directory 1.00 KiB 2021-11-09 16:36:46 root root rwx--xr-x
dev directory 4.00 KiB 2021-11-09 16:35:55 root root rwx--xr-x
dysksdb directory 512.00 KiB 2021-11-09 21:19:13 root root rwx--xr-x
etc directory 4.00 KiB 2021-11-09 20:38:02 root root rwx--xr-x
home directory 4.00 KiB 2021-11-09 16:35:55 root root rwx--xr-x
lib directory 4.00 KiB 2021-11-09 16:38:37 root root rwx--xr-x

Storage

Storage Providers

Data Protector for Cloud Workloads allows backup of multiple storage providers. You can protect Ceph RBD volumes, plain file systems, Nutanix Files (AFS), and Nutanix Volume Groups.

You can execute full and incremental backups, and recover individual files using mounted backups or share them over iSCSI.

Instances

Instances

General

A list of currently known storage instances and access to the details page of each object. From this place, you can also perform on-demand actions such as backup, restore and file-level restore.

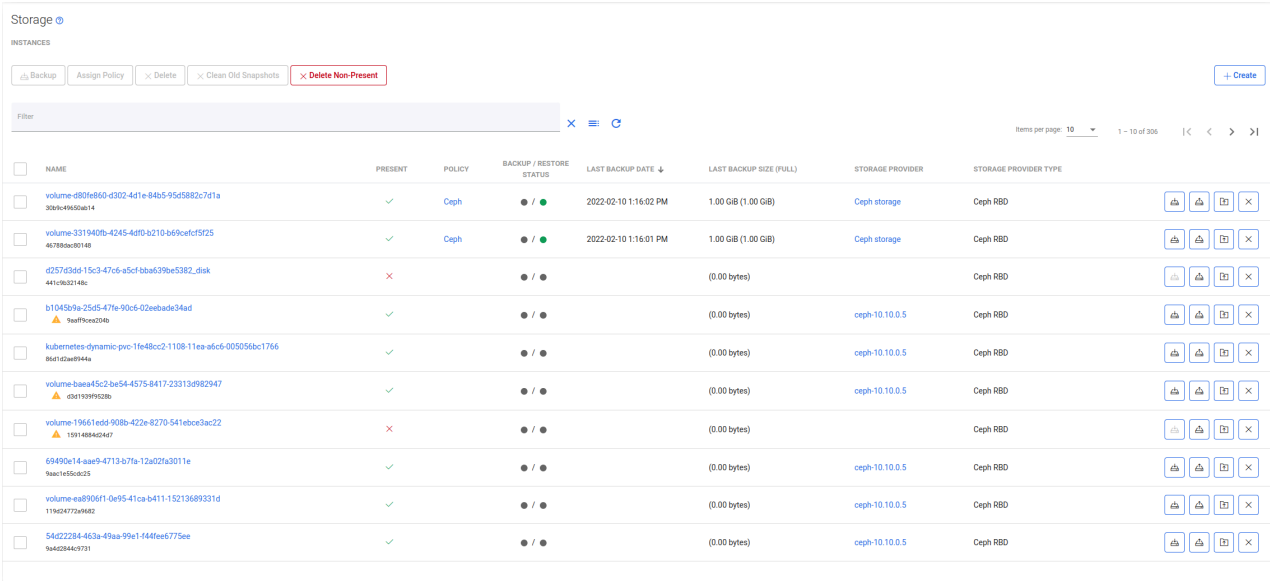

Going back to the Storage details page, this is what it looks like:

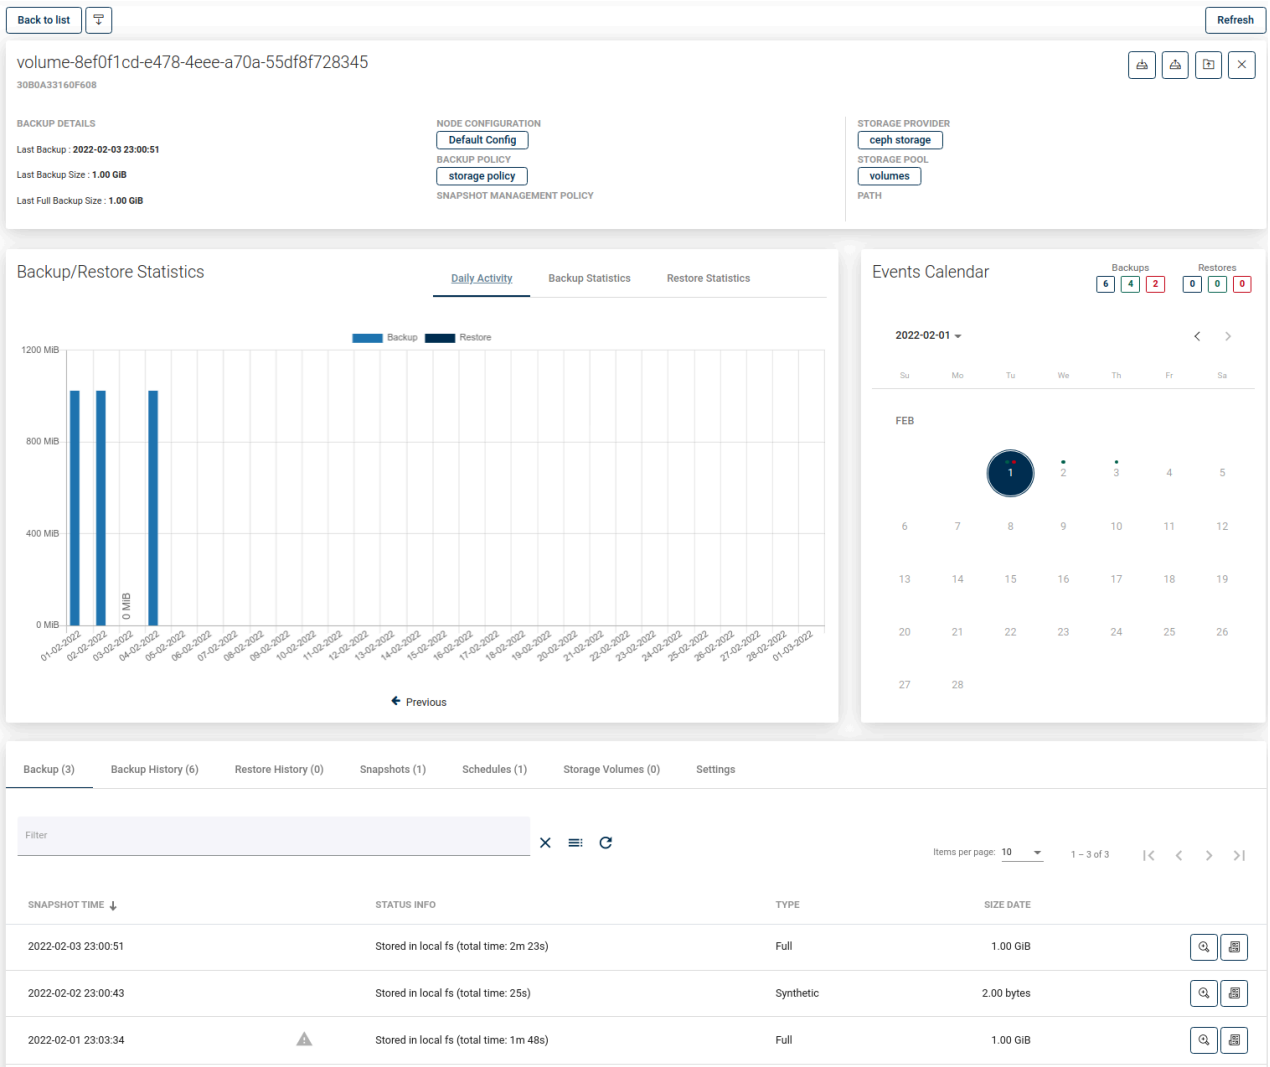

As you can see, the window has been divided into several areas:

Storage instance summary

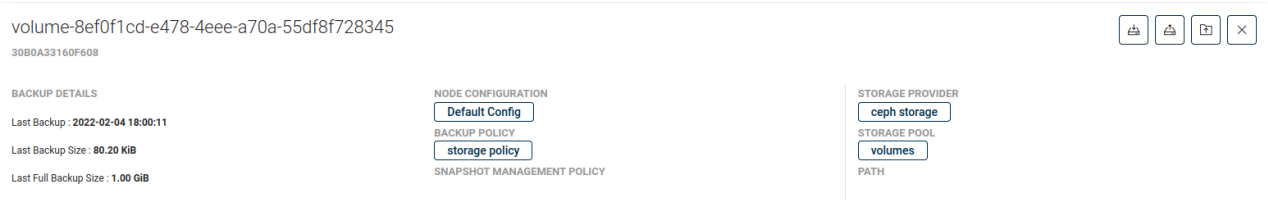

At the top, you can see summarized pieces of information, such as:

- ID of Storage instance into Data Protector for Cloud Workloads
- to which Provider the instance belongs
- to which Pool the instance belongs

- which node is responsible for backup
- short information about the last backup actions
- whether the storage instance has policies assigned to it

You can also use several function buttons, such as:

- refresh
- back to list
- change section order
- backup
- restore
- mount
- delete

Backup/Restore Statistics

Daily activity

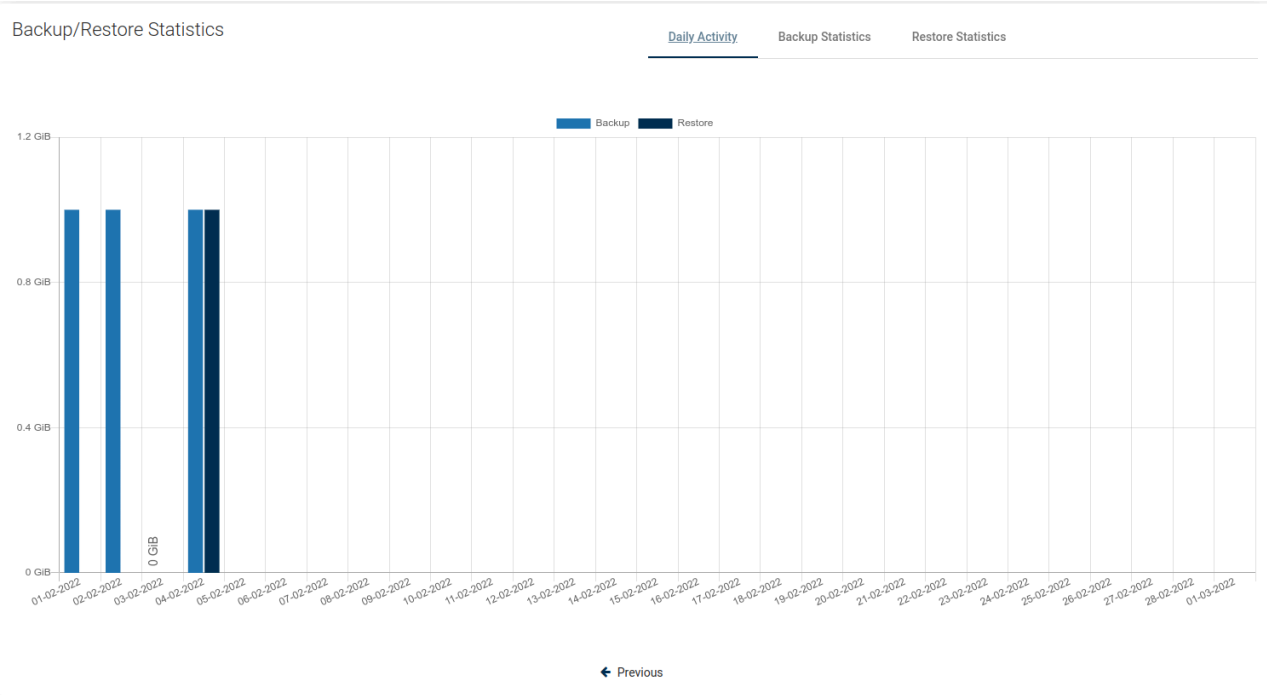

First, you'll see a daily summary of the backup and restore operations for the last month. This view is called "Daily Summary" and is the default view. You can switch

the report between multiple views.

Backup Size

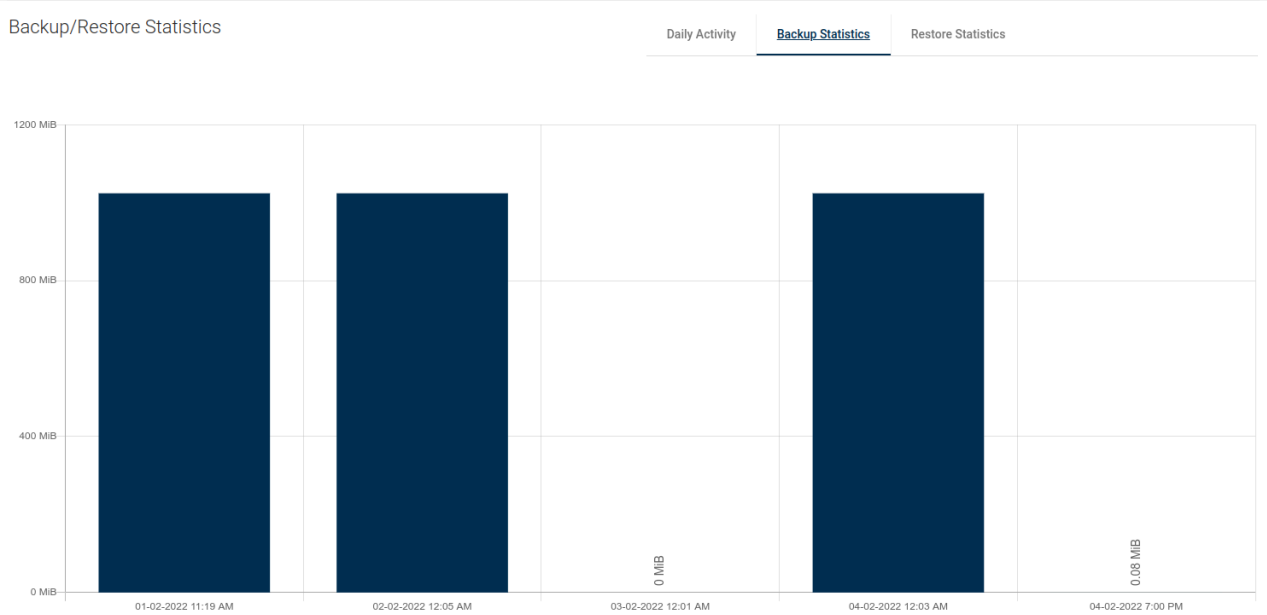

This view shows separate columns for each backup made. Thanks to this, you can easily determine what data increase occurs on a given machine.

Backup Time

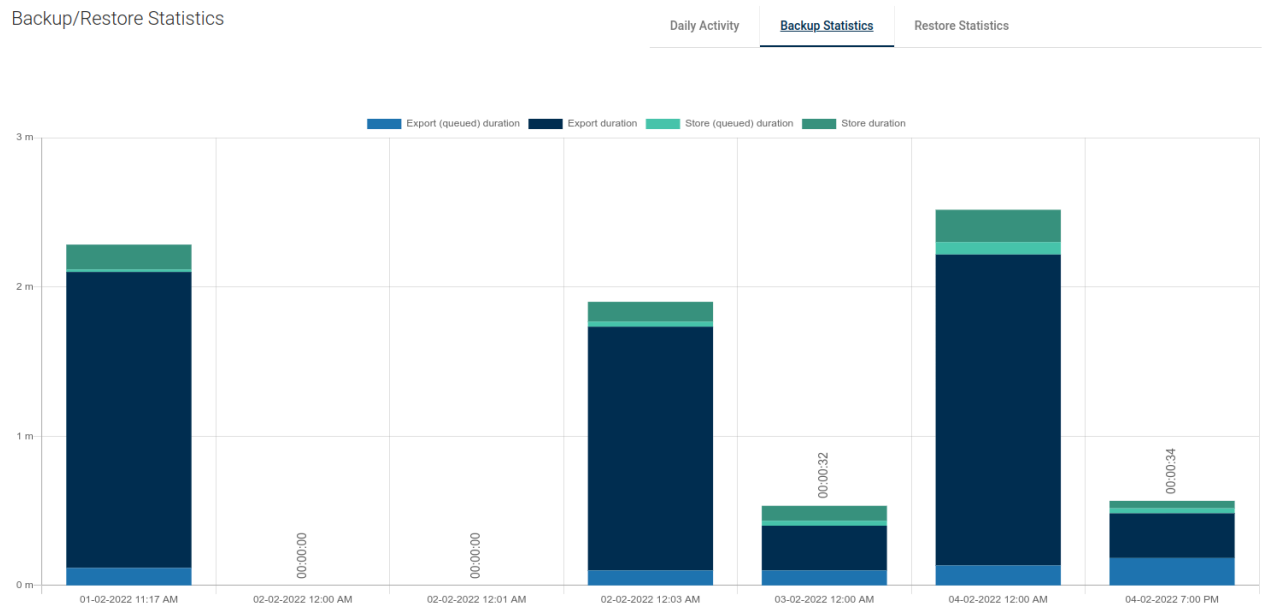

A very useful report. It allows you to determine the required window length for backups or, based on the time of individual phases, it is easy to deduce the cause of slow backups.

Transfer Rate

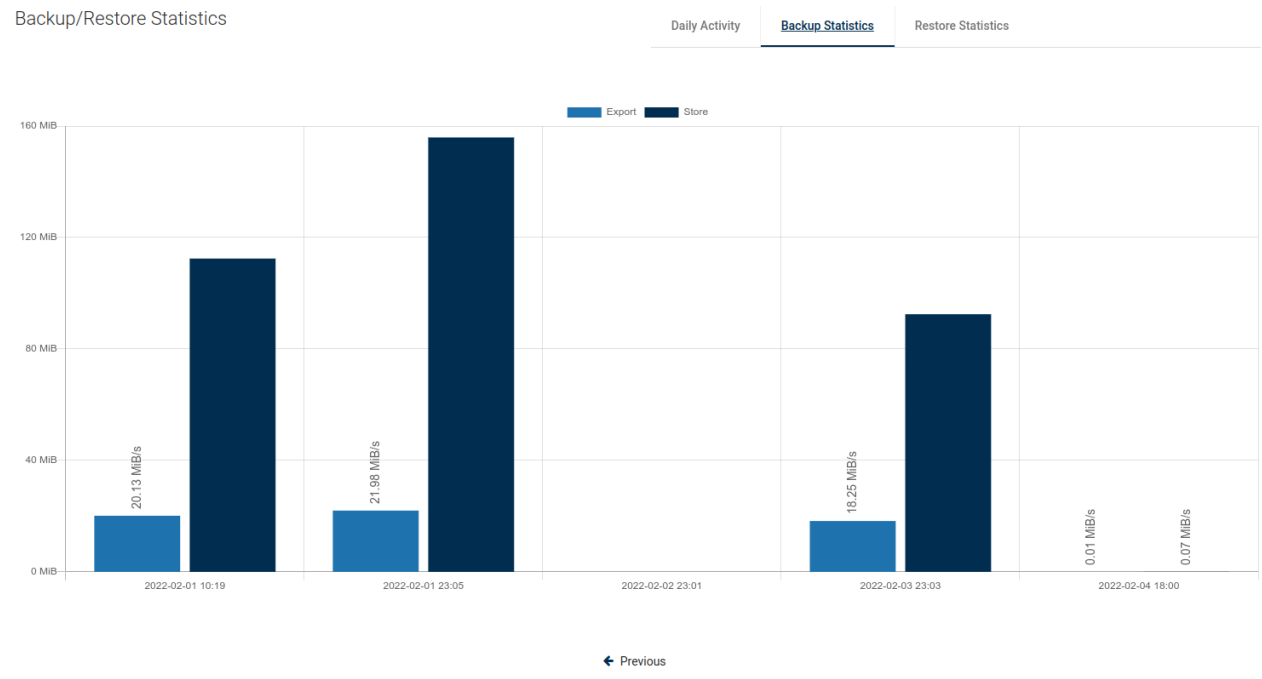

One of the latest reports, now you can easily see how fast data transfer is.

Restore Duration

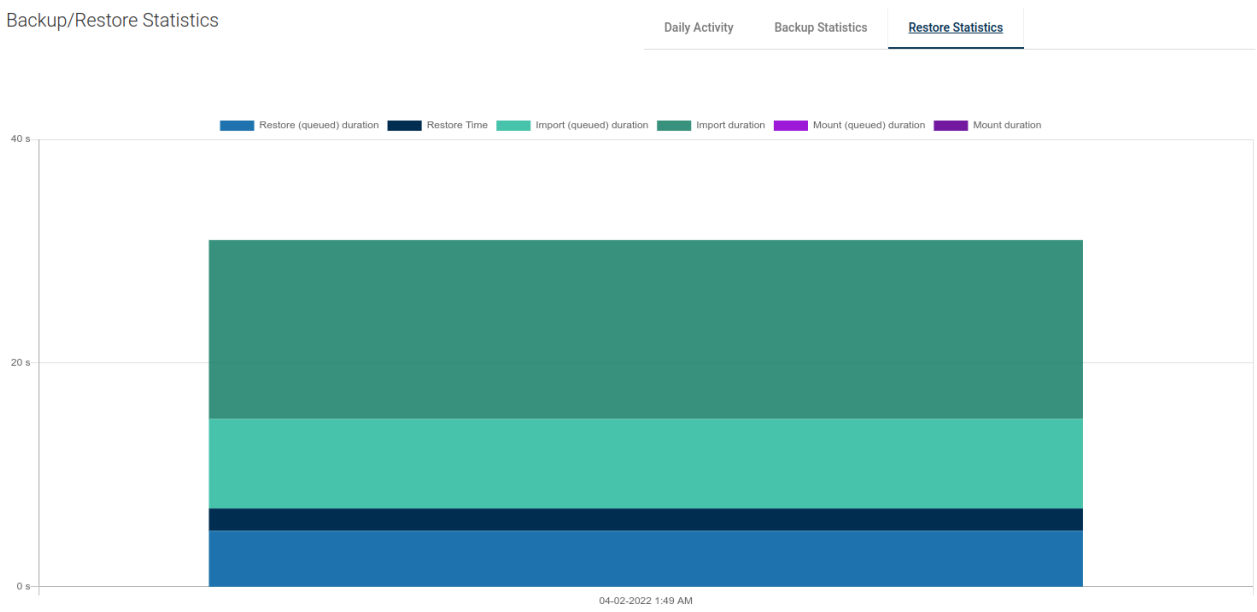

A view with the same properties as "Backup Time". It allows us to estimate how long it will take to restore the storage instance in the event of a failure.

Restore Rate

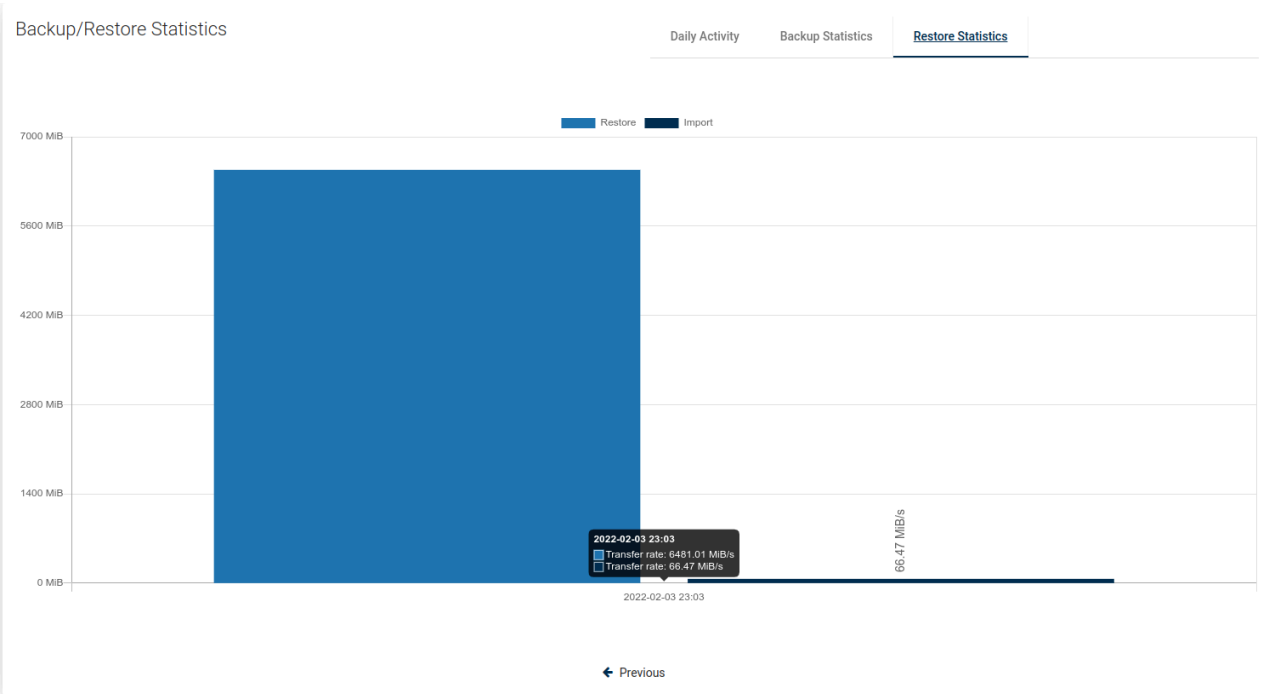

As in the previous case, we also have the transfer speed for the restore job.

Events Calendar

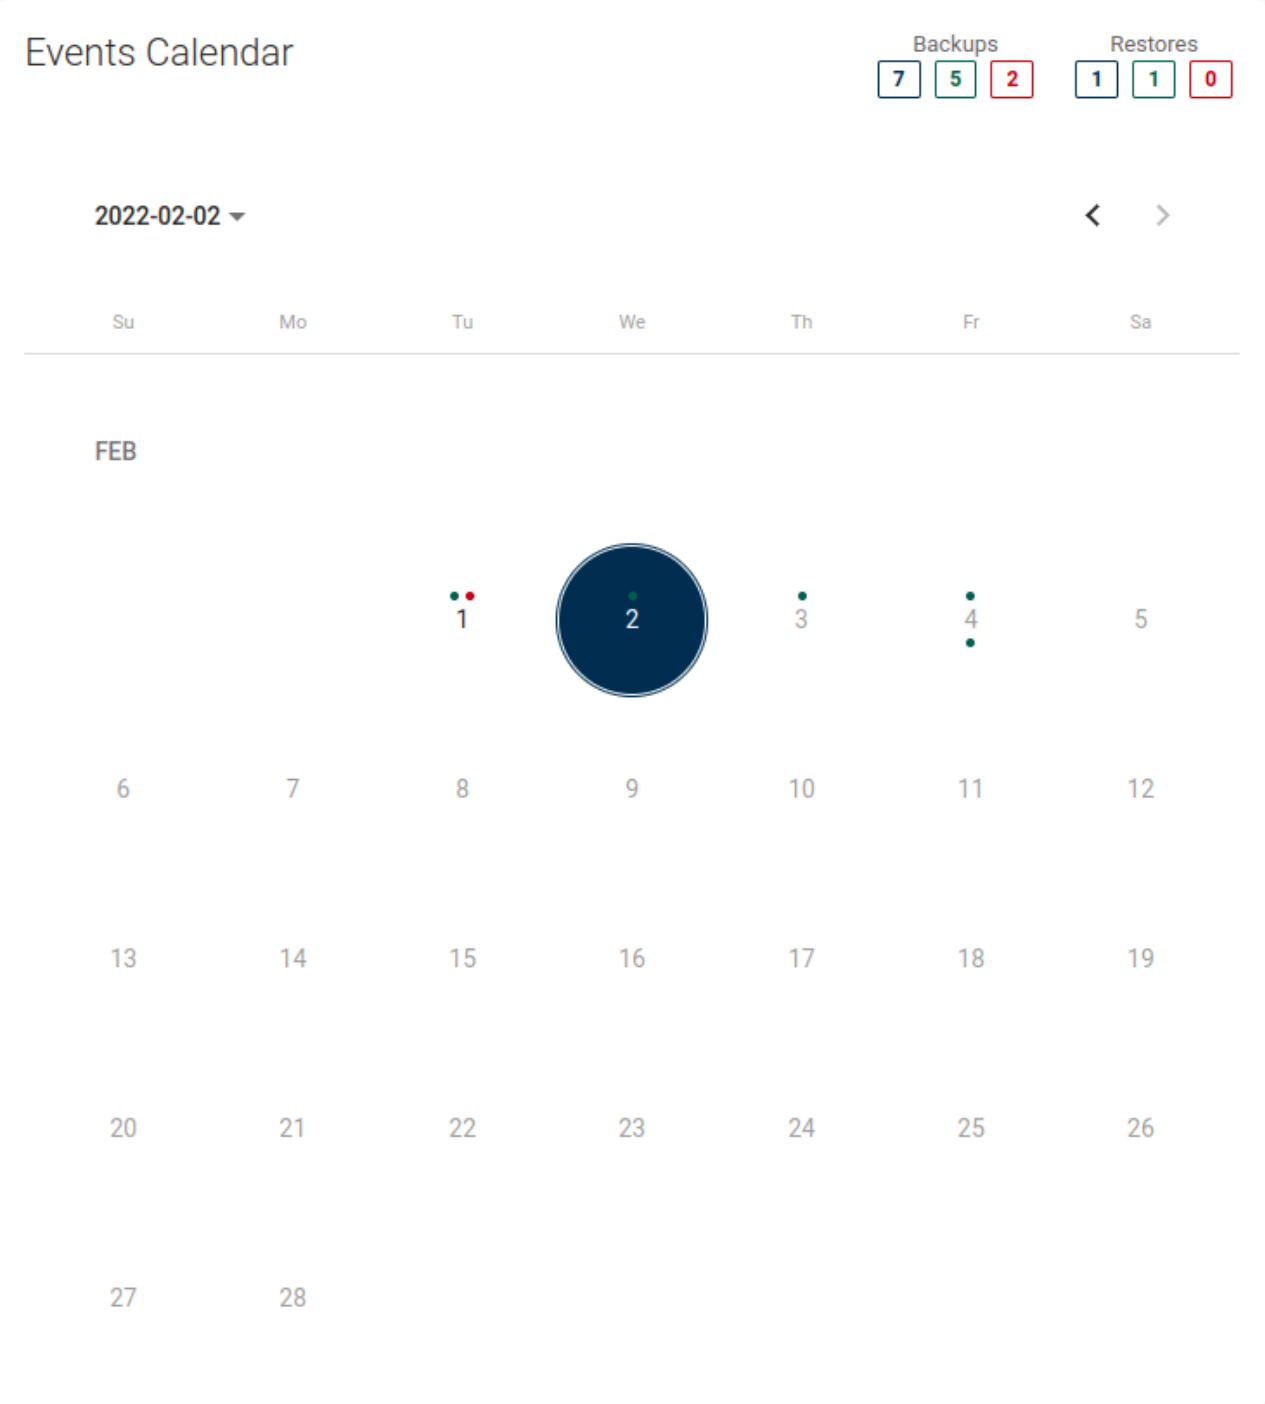

The calendar extends the possibilities of adjacent statistics. It allows you to neatly define the range of days you want to see, additionally makes a quick summary of the number of backups and restores (top right corner).
Blue - the sum of all backups, **Green** - the sum of successes, **Red** - the sum of failures.

Bottom menu

In the bottom menu, you can find a large number of tabs, each of which will present different information or will allow you to change the configuration of this particular storage instance.

Backup

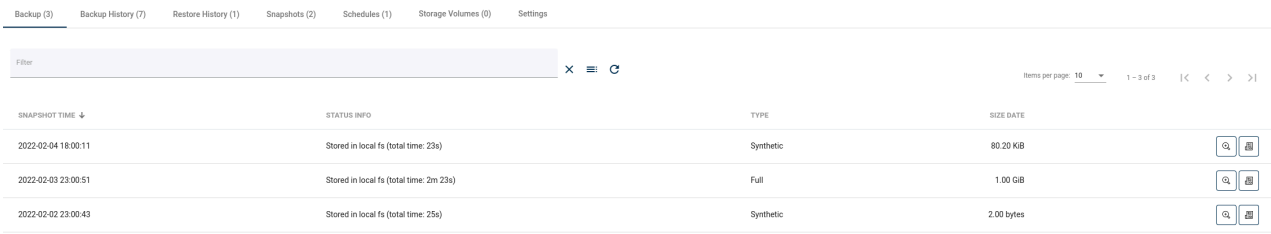

The first tab shows all backups that are currently available and all the basic information about them in a list. After clicking on the magnifying glass button, you will see additional information. The button next to it allows you to download logs in the form of a .txt file.

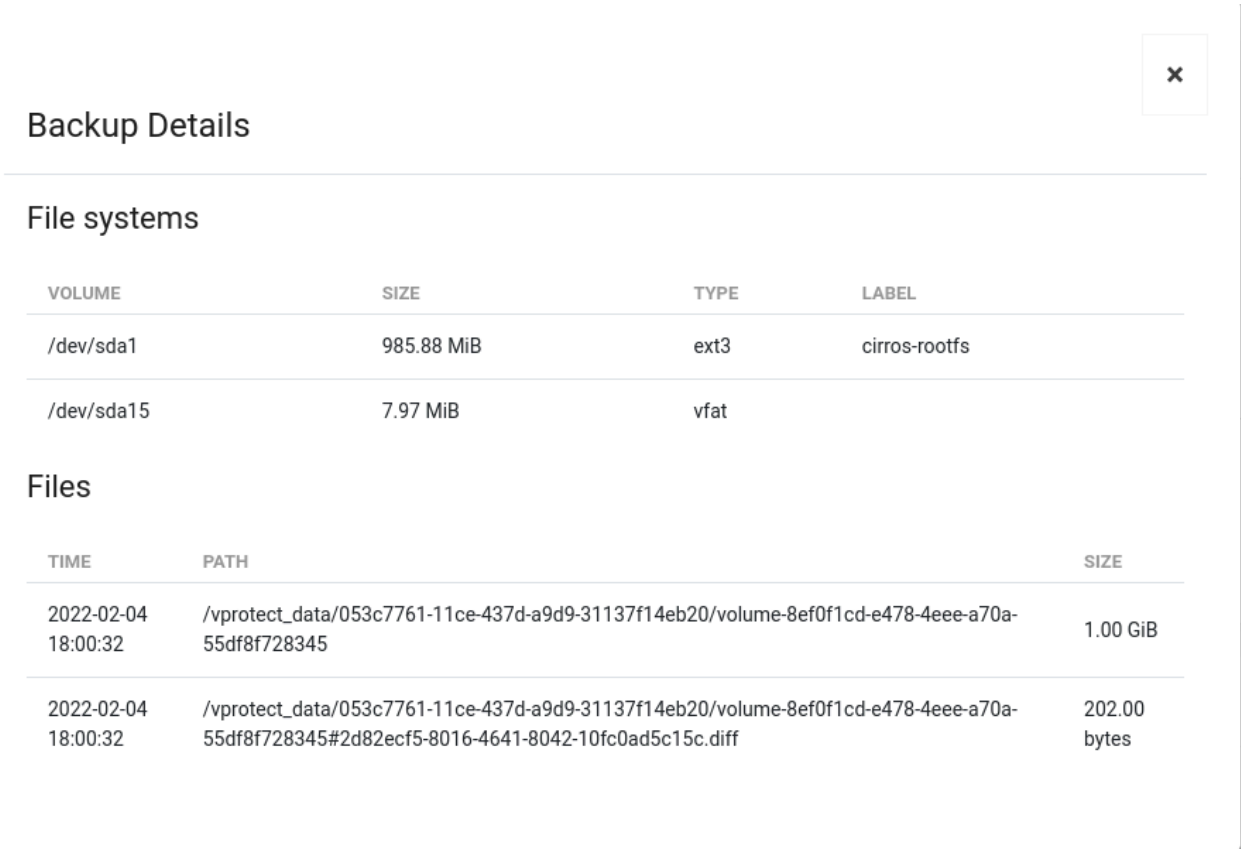

Backup History

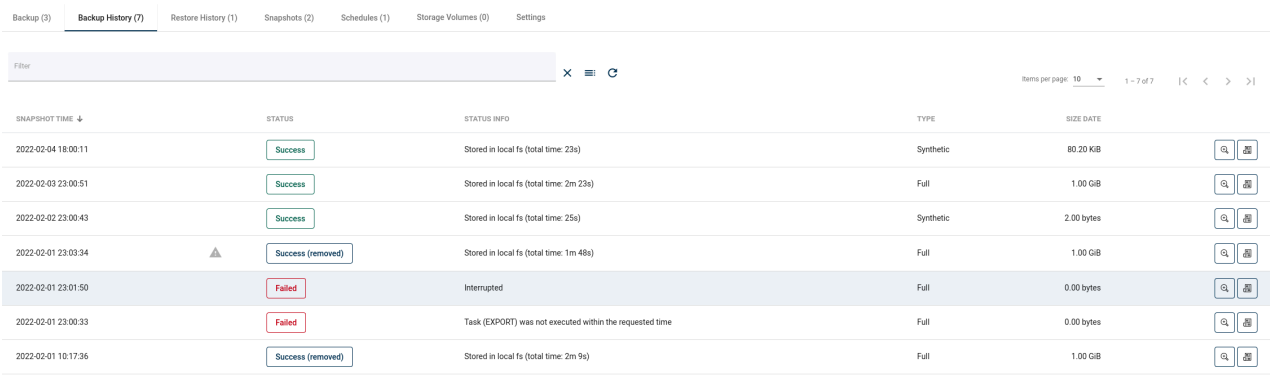

This tab shows information about all backups made. Also about those that failed, were removed (because of retention) or are currently executing.

Restore History

This tab is similar to "Backup History". This is a list with basic information about the storage instance restores performed. When you open the details of the selected restore, you will see much more detailed information.

Restore details

Node	node1
Protected Entity	volume-8ef0f1cd-e478-4eee-a70a-55df8f728345
Restore Time	2022-02-04 00:49:47
Status	Success
Status Info	Backups of STORAGE: volume-8ef0f1cd-e478-4eee-a70a-55df8f728345 (size: 1 GiB) imported in: 15s.
Backup Type	Full
Restore Type	Restore and import

Task Time Stats

Import duration	00:16 s
Import (queued) duration	00:08 s

Backup

Backup Time	2022-02-03 23:03:14
Protected Entity	volume-8ef0f1cd-e478-4eee-a70a-55df8f728345
Size	1.00 GiB
Snapshot time	2022-02-03 23:00:51
Status	Success
Status Info	Stored in local fs (total time: 2m 23s)
Type	Full

Backup - task time stats

Export time	02:05 s
Queued Export Time	00:08 s
Queued Store Time	00:05 s
Store time	00:13 s

Snapshots

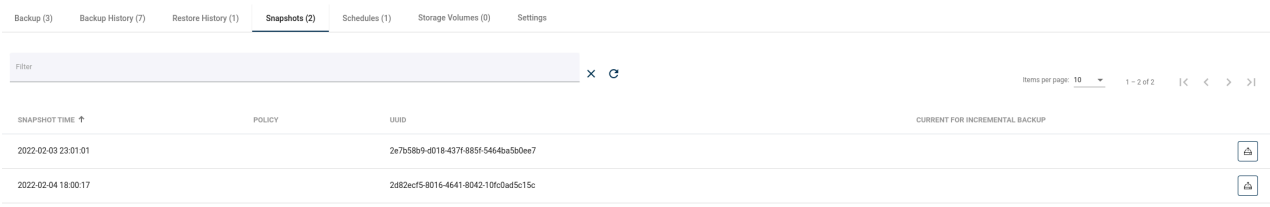

This tab shows the storage instance snapshot - the tab is visible only for the ceph and nutanix Storage Provider.
As you can see in the list above, there is a green dot next to the snapshot. This means that the snapshot is created for incremental backup purposes. This is an automatic operation and we only keep the last snapshot.

Mounted Backups

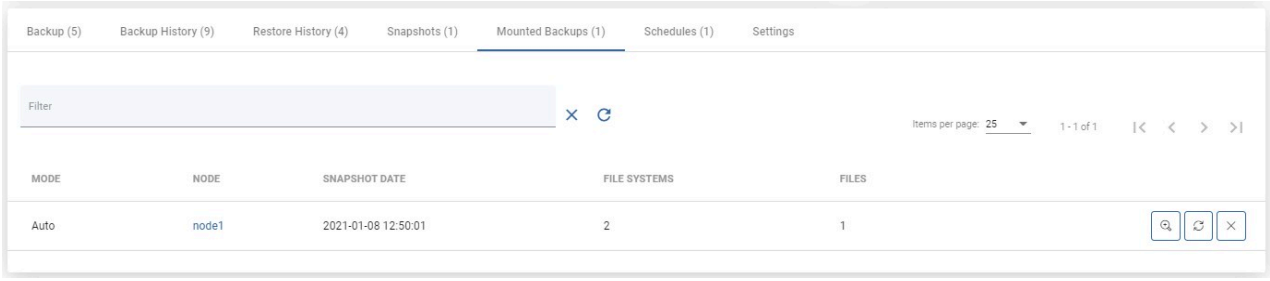

This tab lists all mounted backups of this particular storage instance. With the buttons on the right, you can browse/remount/delete it.

Schedules

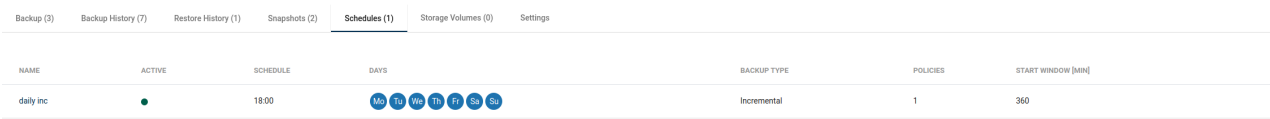

On this tab, you can see all the schedules assigned to the instance.

Settings

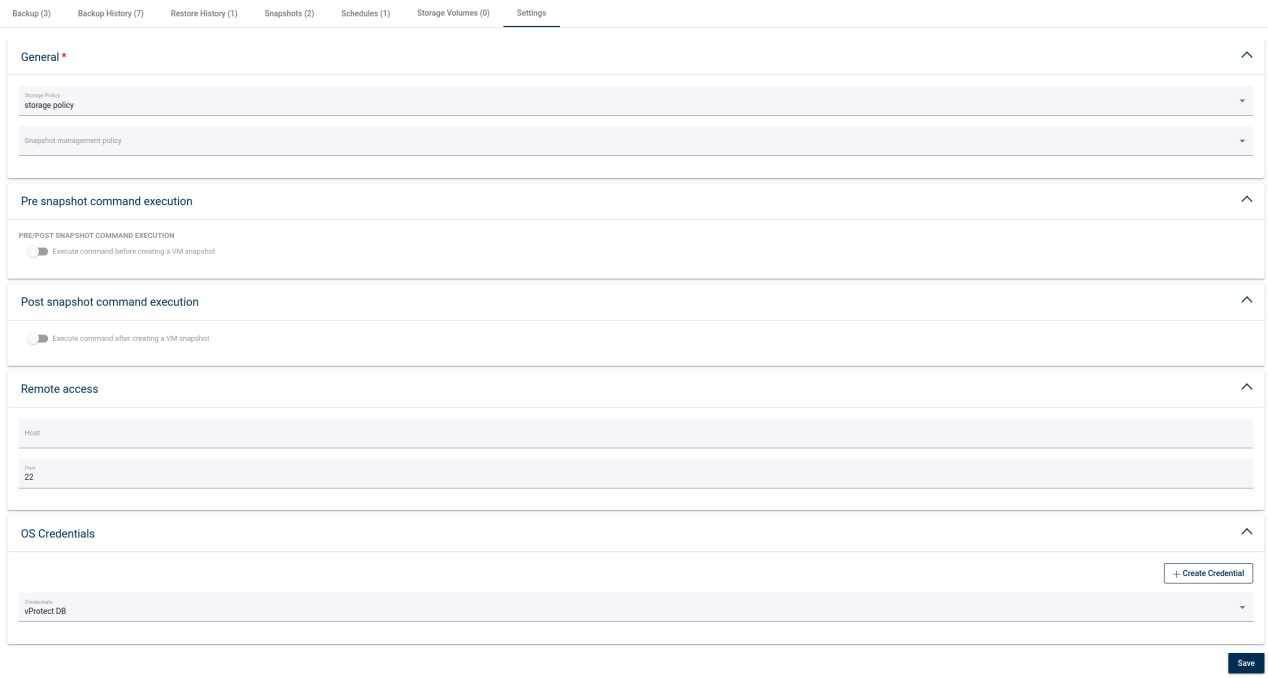

Finally, the last tab. The first option allows you to change the policies assigned to the storage instance.

Performing pre/post snapshot commands is a function intended for advanced users. As the name implies, it allows us to execute scripts via an ssh connection, either before or after taking a snapshot.

Backup on-demand

Backup on-demand

Storage List

To perform on-demand backup, go to the instances tab under the Storage Providers section.

You can click on the  icon to back up a single storage instance

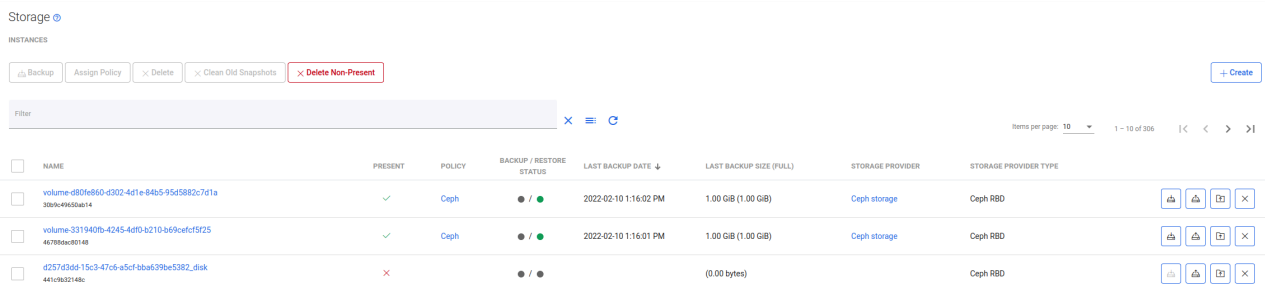

Or select multiple instances and click on the 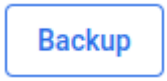 icon to back them up.

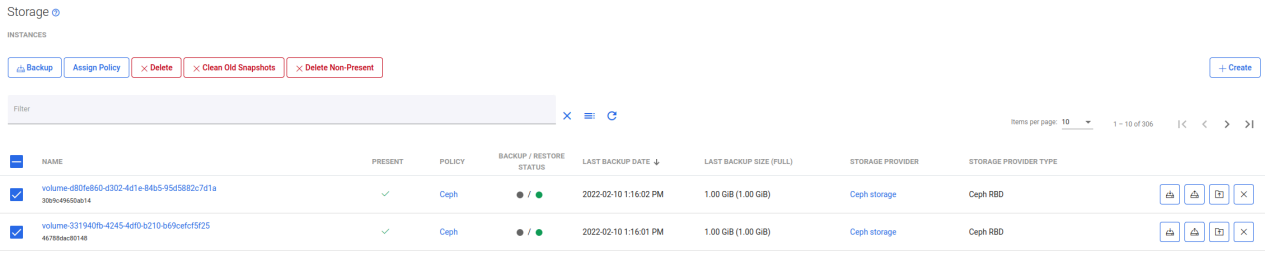

Finally select the backup type, backup destination, when a task should start, the priority, and click on the backup button.

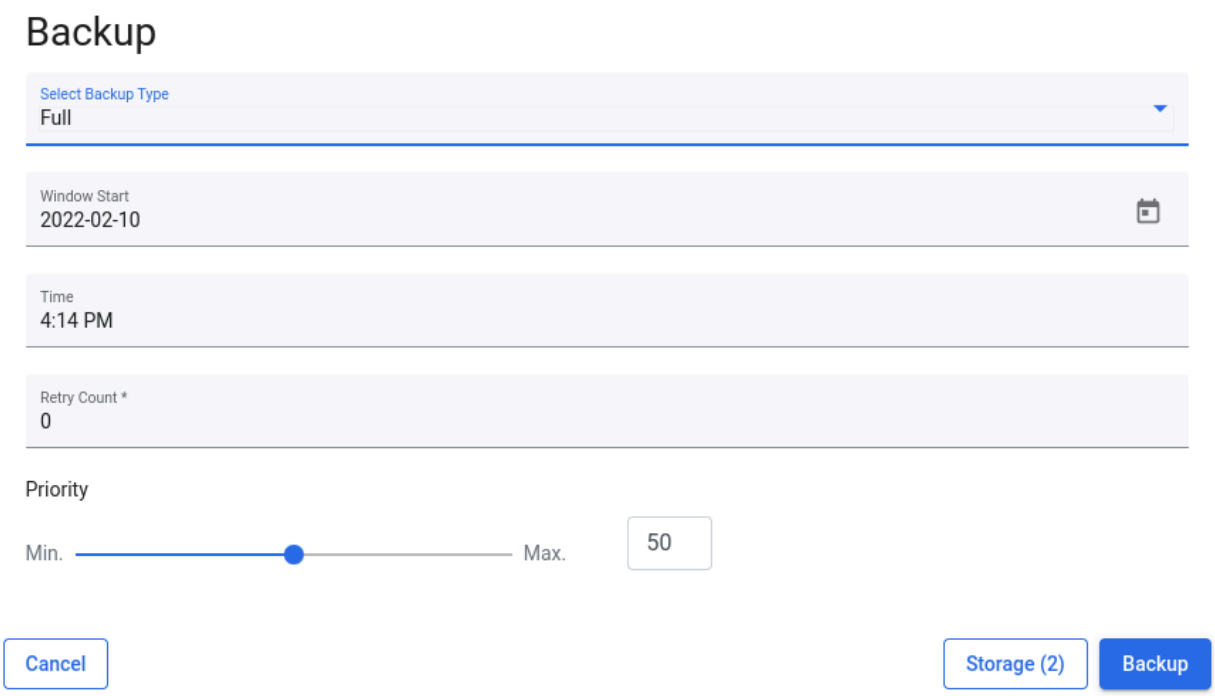

> **Note:** You can also perform the same action thanks to the CLI interface: CLI Reference ↗

Restore on-demand

Restore on-demand

Restore from Storage Providers menu

To restore a single storage instance on-demand, go to the instances tab under the Storage Providers section. Click on the restore icon next to the object 

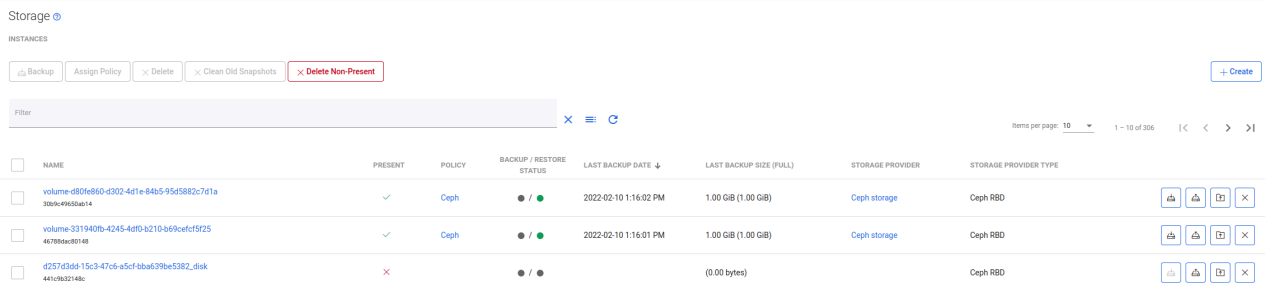

Now you should see a popup window where you can customize the restore settings.

Restoring to the filesystem is an option to restore directly to the Data Protector for Cloud Workloads Node storage.
(Remember, the Data Protector for Cloud Workloads user must have the appropriate write permissions to the given path.)

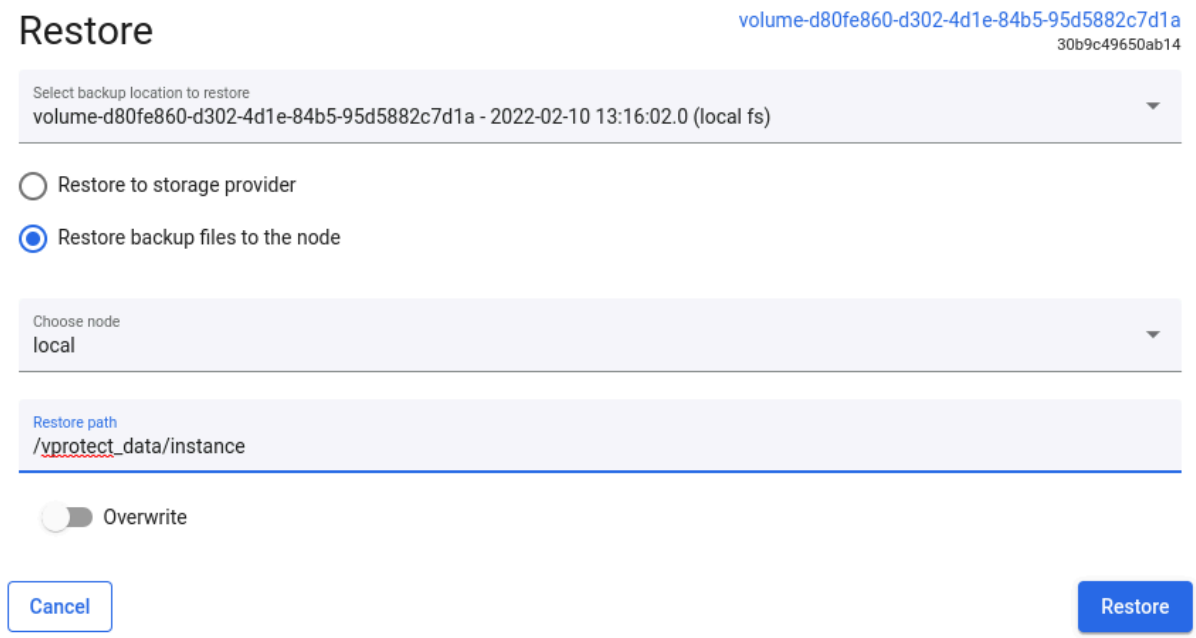

You can also restore the backup to the storage provider from which the backup was created.

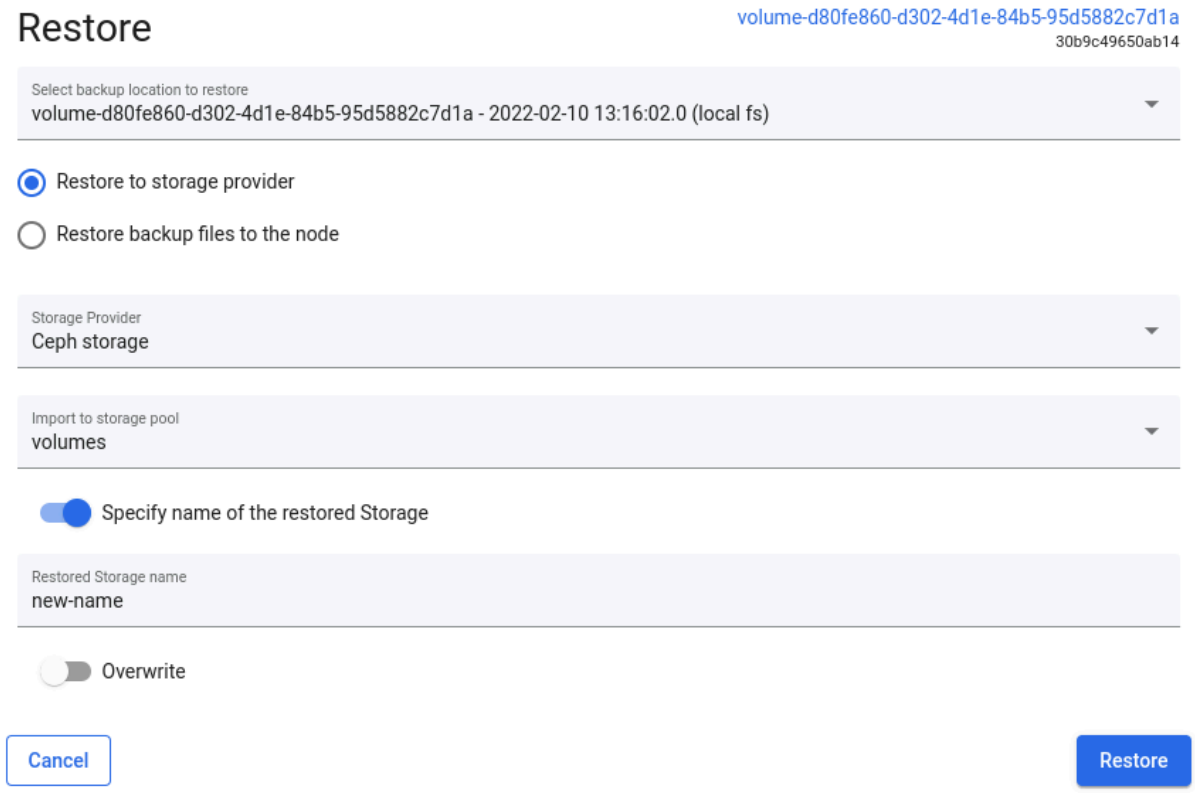

Finally, after customizing the restore, click the restore button.

> **Note:**
>
> - every platform has some restrictions imposed on the storage instance name, such as length or characters that can be used. Check these limits before restoring with a custom name.
> - You can also perform the same action thanks to the CLI interface: CLI Reference ↗

Infrastructure

Infrastructure

This section describes how to manage Storage Providers in Data Protector for Cloud Workloads. The inventory that Data Protector for Cloud Workloads needs first has to be populated. The first step is always to add a storage provider or file system visible on the Data Protector for Cloud Workloads Node.

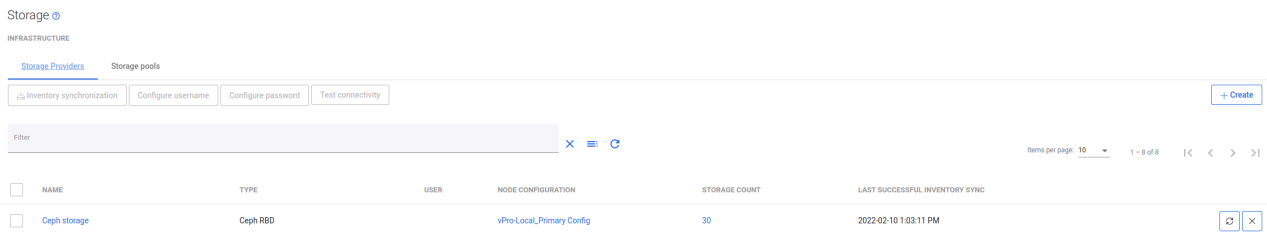

Click Add Storage Provider to add entries

Then synchronize the inventory (either automatically - a dialog box will be shown just after saving the form, or manually with the button on the right of the storage provider).

If the inventory synchronization tasks (visible on the console at the bottom) were completed successfully, this also proves that the connection was successful, authentication is correct and all of the inventory items have been collected successfully.

Check the storage pools as well as Storage → Instances to see the results of the inventory synchronization.

Storage

INFRASTRUCTURE

Storage Providers | Storage pools

Filter: ceph

Items per page: 10 | 1 – 7 of 7

NAME	STORAGE PROVIDER
images 31c5e1de-1ce6-4954-acbb-5014e59e33ed_3	Ceph storage
images 690494d1-6553-11eb-a6fe-005056a63972_1	ceph-10.10.0.5
kube 690494d1-6553-11eb-a6fe-005056a63972_4	ceph-10.10.0.5
vms 31c5e1de-1ce6-4954-acbb-5014e59e33ed_1	Ceph storage
vms 690494d1-6553-11eb-a6fe-005056a63972_3	ceph-10.10.0.5
volumes 31c5e1de-1ce6-4954-acbb-5014e59e33ed_2	Ceph storage
volumes 690494d1-6553-11eb-a6fe-005056a63972_2	ceph-10.10.0.5

Backup SLAs

Backup SLAs

Backup SLAs allow you to group backup policies and schedules for multiple storage instances. In general, a storage instance should have only one backup policy so as to always have an easy-to-interpret configuration.

Each policy can have multiple schedules assigned so that you can define more complex schedules so that backups are executed even multiple times each day and with different backup types.

> **Note:**
>
> - schedules define the type of backup - full or incremental
> - the approach we highly recommend is to create a schedule for **periodic full** backup and always assign at least 1 such schedule in the backup SLAs
> - in order to create incremental backups, you need to always have **at least 1 incremental backup schedule** and run at least one full backup

Policies

Policies

Policies allow you to group storage instances in many ways. For example, based on the type of storage provider.

To create a new backup policy, open the Backup SLAs tab under the Storage Provider section and click on the + Create button on the right.

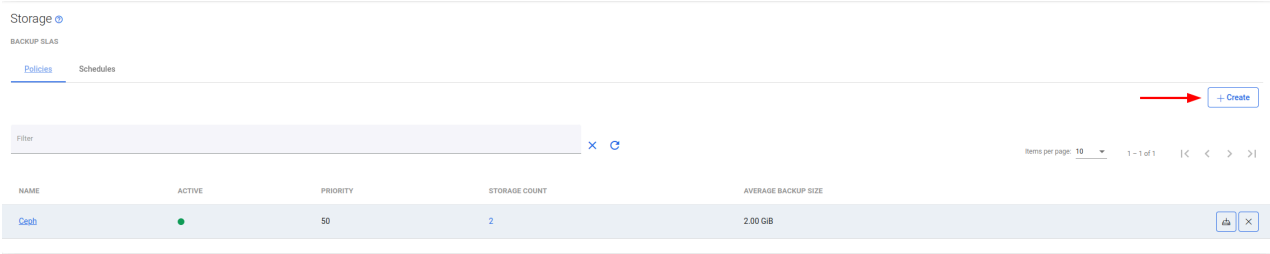

Now you should see the policy wizard with 5 main sections.

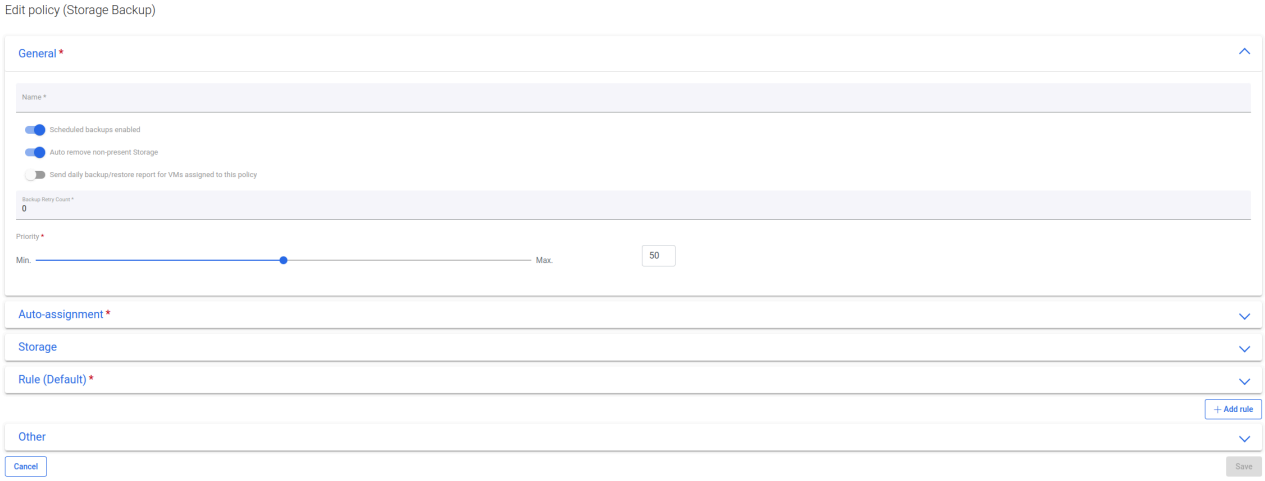

General

Under this section you can set up:

- Name of policy

- Switch on/off auto-remove non-present virtual environments
- Set the priority for tasks

Auto-assignment

In this section you can configure automatic policy assignment based on certain criteria:

- Mode
  - Disabled
  - Assign only
  - Assign and remove
- Include or exclude rules based on regular expressions matching storage instance names:
  - regular expression examples:
    - `.*` match any character any number of times
    - `st-[0-9][0-9][0-9]` - match names that start with `st-` and 3 digits
    - `(prod|uat|dev)-[0-9][0-9][0-9][a-z]?` - match names that start with the `prod` or `uat` or `dev` prefix, then `-`, then 3 digits and an optional lower-case letter (matching is case-sensitive)
  - exclude rules always take precedence over include rules
  - objects may not be reassigned to a different policy if they already have a matching policy assigned
  - objects may be reassigned to a different policy only if the mode is `Assign and remove`, the current policy assignment rules don't match, and the other policy's rules do match
  - rules are joined with the OR operator, so
    - if **any** rule (tag or matched regular expression) excludes the storage instance - it will be excluded
    - if **no** rule (tag or matched regular expression) excludes the storage instance, and **any** rule (tag or matched regular expression) includes the VM - it will be included
- You can also select clusters to match only VMs that belong to them

Storage

Here you can easily select storage instances manually.

Rule

This section is used to select the backup destination.

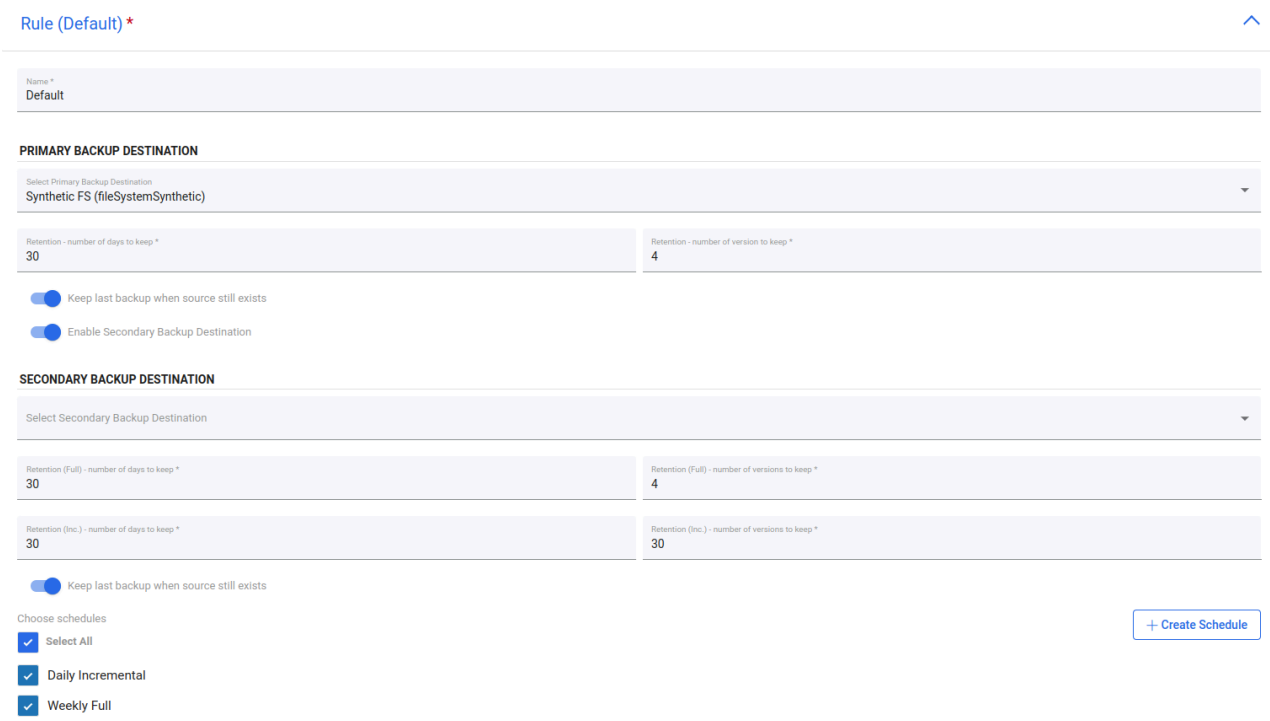

> **Note:** You can select Primary and Secondary Backup Destintanion in one Rule

You can also set here **Retention** settings for your backups. You can use a number of days and versions for full and incremental backups.

Retention

Data Protector for Cloud Workloads handles retention for all backup destinations. There are 4 properties that define how long backup should be kept in the backup destination:

- `Retention (Full) - no. of versions to keep` - number of full backups
- `Retention (Inc.) - no. of versions to keep` - number of incremental backups
- `Retention (Full) - no. of days to keep` - number of days to keep a full backup
- `Retention (Inc.) - no. of days to keep` - number of days to keep an incremental backup

> **Note:** If you are using Synthetic File System backup destination, you have only two options for retention:
>
> - `Retention - no. of versions to keep` - number of full backups
> - `Retention - no. of days to keep` - number of days to keep a full backup

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.

Other

This is an optional section with two switches:

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

Here are two examples of when using switches is very useful: It is very likely that if 30% of the backup tasks fail, the remaining tasks will also fail because the environment has failed. Or, if you are backing up a set of storage instances, and if even one is not secured, there is no point in backing up the rest.

At the end, save the settings.

Schedules

Schedules

Schedules allow you to invoke specific policies periodically. This allows you to backup multiple storage instances automatically.

A schedule defines when and on which days instances should be backed up. To define a new schedule, open Backup SLAs under the Storage section and go to the Schedules tab, then use the **Create** button.

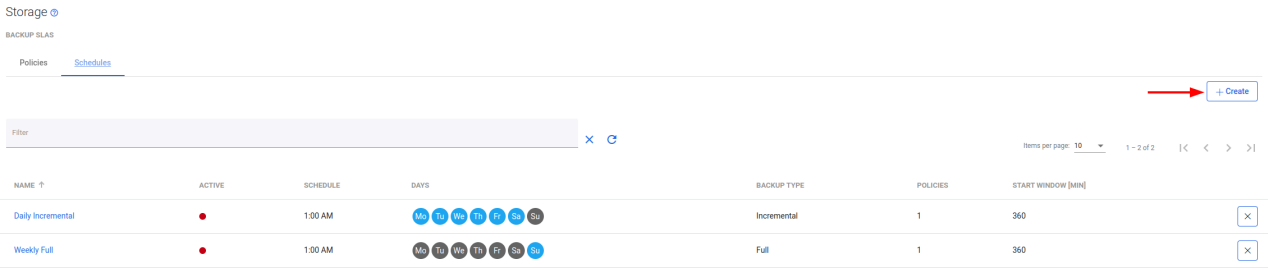

Now enter the properties:

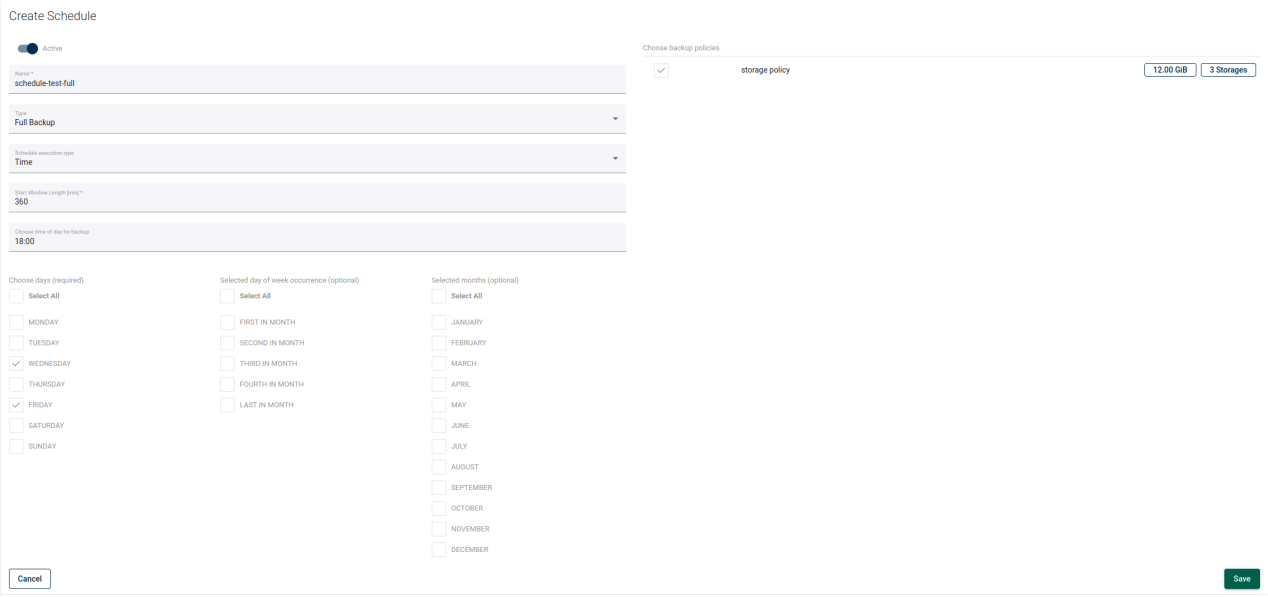

- `Schedule Active` - enable or disable executing schedule
- `Name` - schedule name
- `Backup Type` - defines backup type: full or incremental

- `Execution Type` - choose time or interval mode
- `Start Window` - defines for how long since the task start time scheduled tasks are allowed to be executed
- `Choose time of day` - for the time execution mode, this defines when the task should be added to the queue
- `Choose time of interval start` - for the interval execution mode, this defines when tasks should start
- `Choose time of interval end` - for the interval execution mode, this defines when tasks should end
- `Frequency` - defines how often the task will be executed during the interval
- `Choose days` - last required parameter, select days of the week on which the task will be performed

You can also use optional parameters to further personalize the backup time or select a storage instance policy if it has been previously created.

> When you set the time with a user in a certain time zone, you specify a point in time at which you want the schedule to start. Changing the timezone doesn't change this point in time, it's converted to your timezone. The time displayed to the user is calculated based on the server time.

Snapshot SLAs

Snapshot SLAs

Data Protector for Cloud Workloads can periodically create snapshots and keep several of them for every Storage Provider instance. To see which Storage Provider support this feature go to Data Protector for Cloud Workloads Support Matrix.

> **Note:**
> Snapshot management for on-demand or scheduled operations is not available until the storage instance has a snapshot management policy assigned to it

Policies

Policies

In order to enable snapshot management for Storage Instance you need to do the following steps:

Go to Snapshot SLAs under the Storage section and create a new Snapshot Management policy:

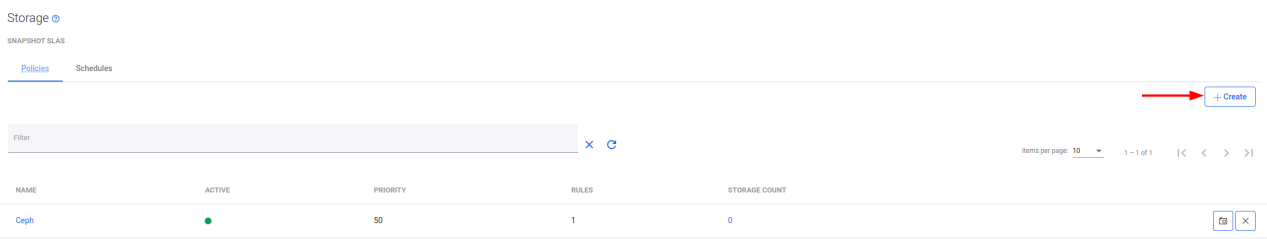

As well as other types of Policies, you'll also find 4 main sections here:

General

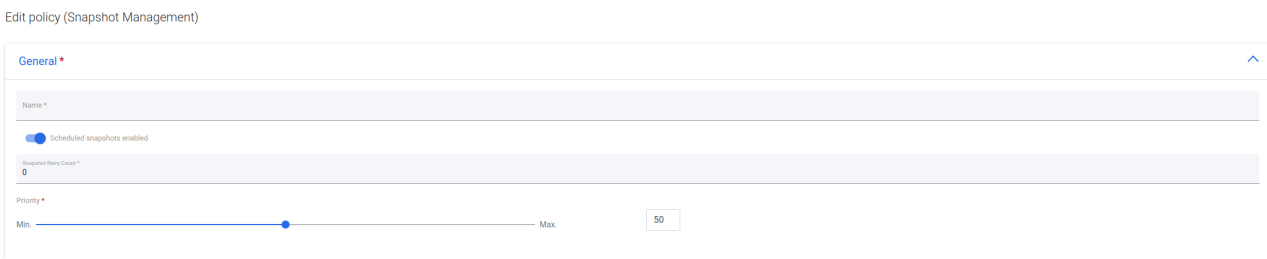

Under this section you can set up:

- Name of policy
- Enable/disable this policy
- **Snapshot Retry Count** - how many times, Data Protector for Cloud Workloads will try to repeat the snapshot task
- Set priority for the task

Auto-assignment

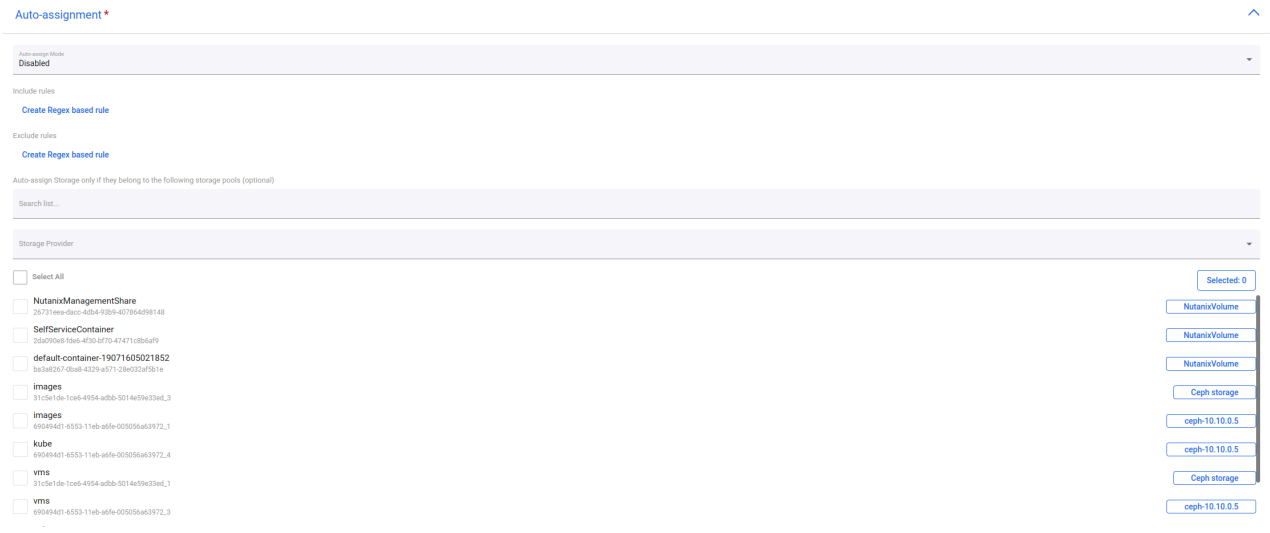

In this section you can set up:

- Mode
  - Disabled
  - Assign only
  - Assign and remove
- Include or exclude rules based on regular expression matching the Storage Instance name:
  - regular expression examples:
    - **.*** match any character any number of times
    - **storage-[0-9][0-9][0-9]** - match name that starts with storage- and 3 digits
    - **(prod|uat|dev)-[0-9][0-9][0-9][a-z]?** - match the name that starts with prod or uat or dev prefix, then -, then 3 digits and optional lower-case letter (matching is case-sensitive)
  - exclude rules always take precedence over include rules
  - Storage Instances may not be reassigned to the different policy if they already have a matching policy assigned
  - Storage Instances may be reassigned to the different policy only if mode is assigned and remove, current policy assignment rules don't match, and other's policy rules match

- rules are joined with OR operator, so
  - if any rule (tag or matched regular expression) excludes Storage Instance - it will be excluded
  - if no rule (tag or matched regular expression) excludes Storage Instance, and any rule (tag or matched regular expression) includes Storage Instance - it will be included
- You can also select Storage Pools to match only Storage Instances that belong to them

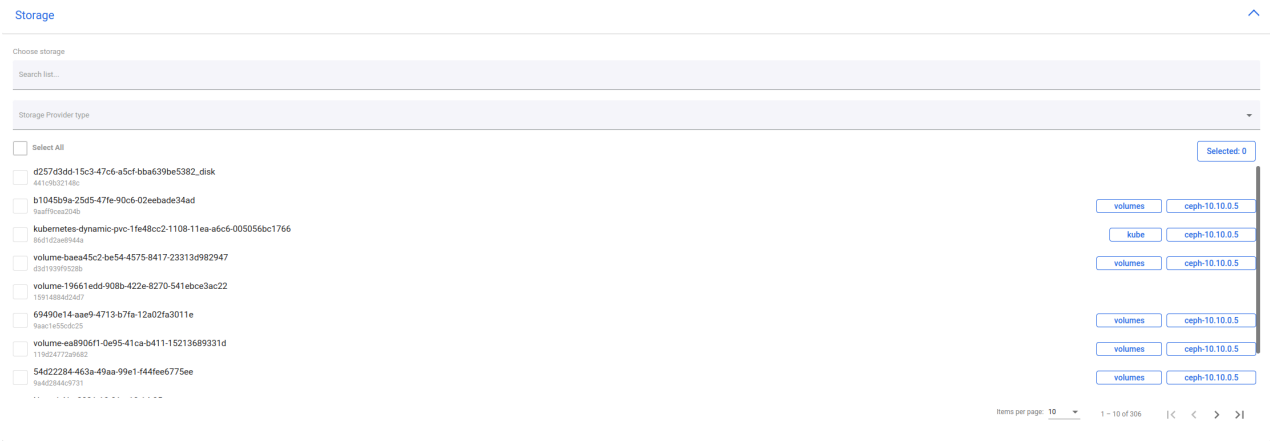

In this place, you can select storage instances manually in a simple way.

Provide retention settings - how many snapshots (created by this policy) will be kept and for how long. If you have already created a schedule, you can also select it. You can also create another rule with different retention settings with different schedules.

Schedules

Schedules

The schedule defines when and on which days snapshots should be created. To define a new schedule click Create Schedule button.

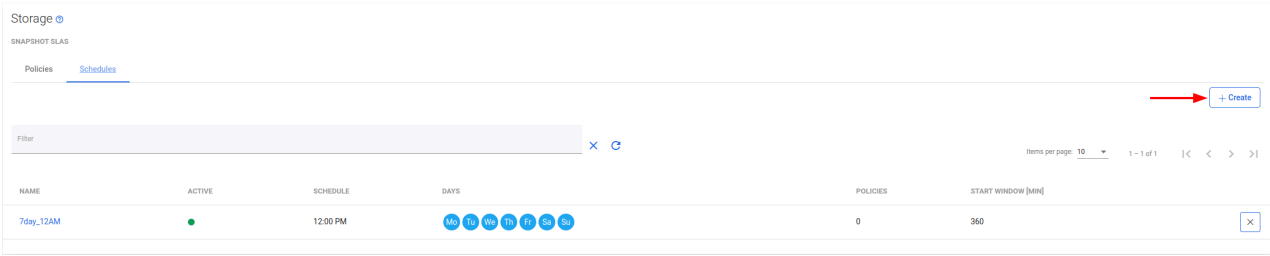

Now provide properties:

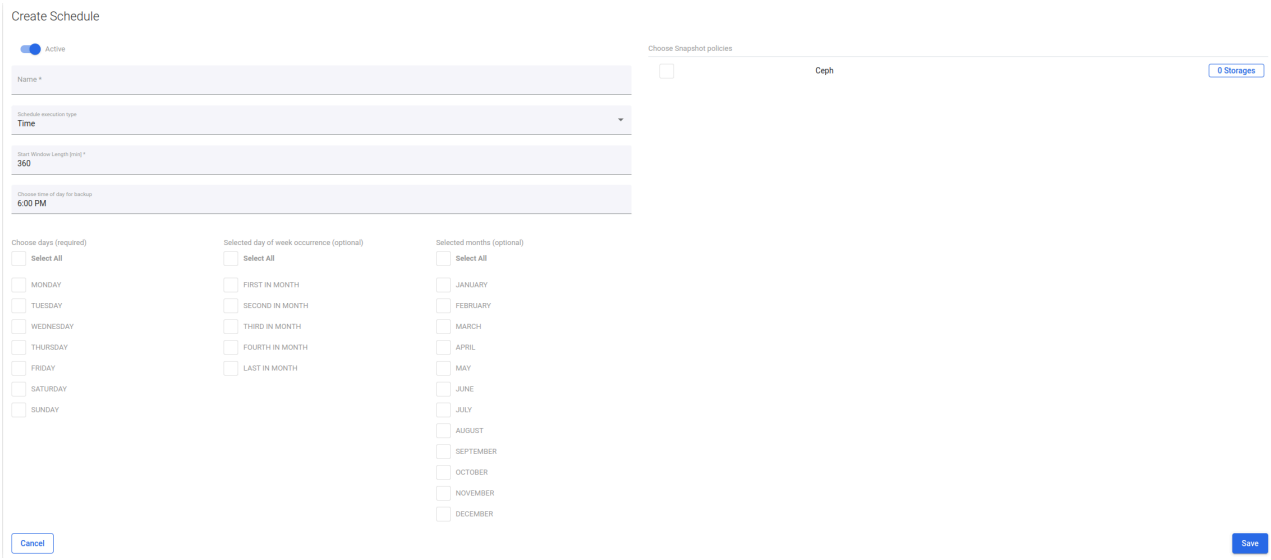

- Schedule Active - enable or disable executing schedule
- Name - schedule name
- Execution Type - choose time or interval mode
- Start Window - defines for how long since task start time scheduled tasks are allowed to be executed
- Choose a time of day - for time execution mode defines when the task should be added to the queue

- Choose a time of interval start - for interval execution mode defines when tasks should start
- Choose a time of interval end - for interval execution mode defines when tasks should end
- Frequency - defines how often the task will be executed during the interval
- Choose days - last required parameter, select days of the week on which the task will be performed

You can also select a Storage Snapshot policy if it has been previously created

> When you set the time with a user in a certain time zone, you specify a point in time at which you want the schedule to start. Changing the timezone doesn't change this point in time, it's converted to your timezone. The time displayed to the user is calculated based on the server time.

Mounted Backups (File-level Restore)

Mounted Backups (File-level Restore)

To mount backup, go to the Instances tab under the Storage Providers section on the left side menu, then click on the mount icon next to the chosen storage instance

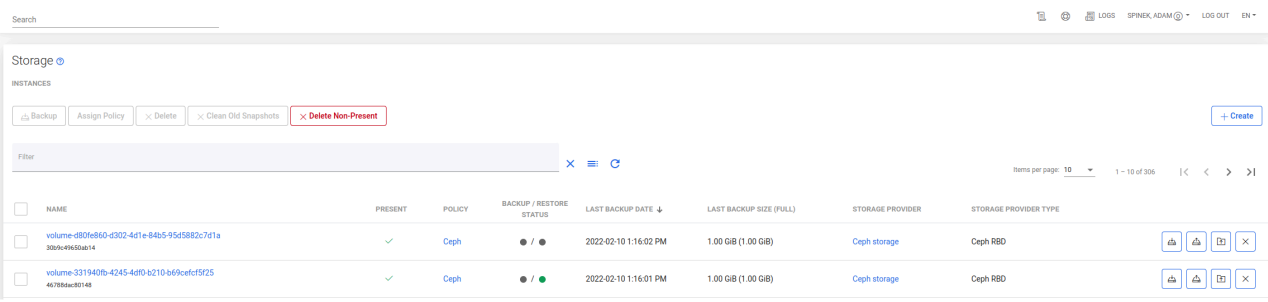

In the popup window, you can select which backup you want to mount and on which node.

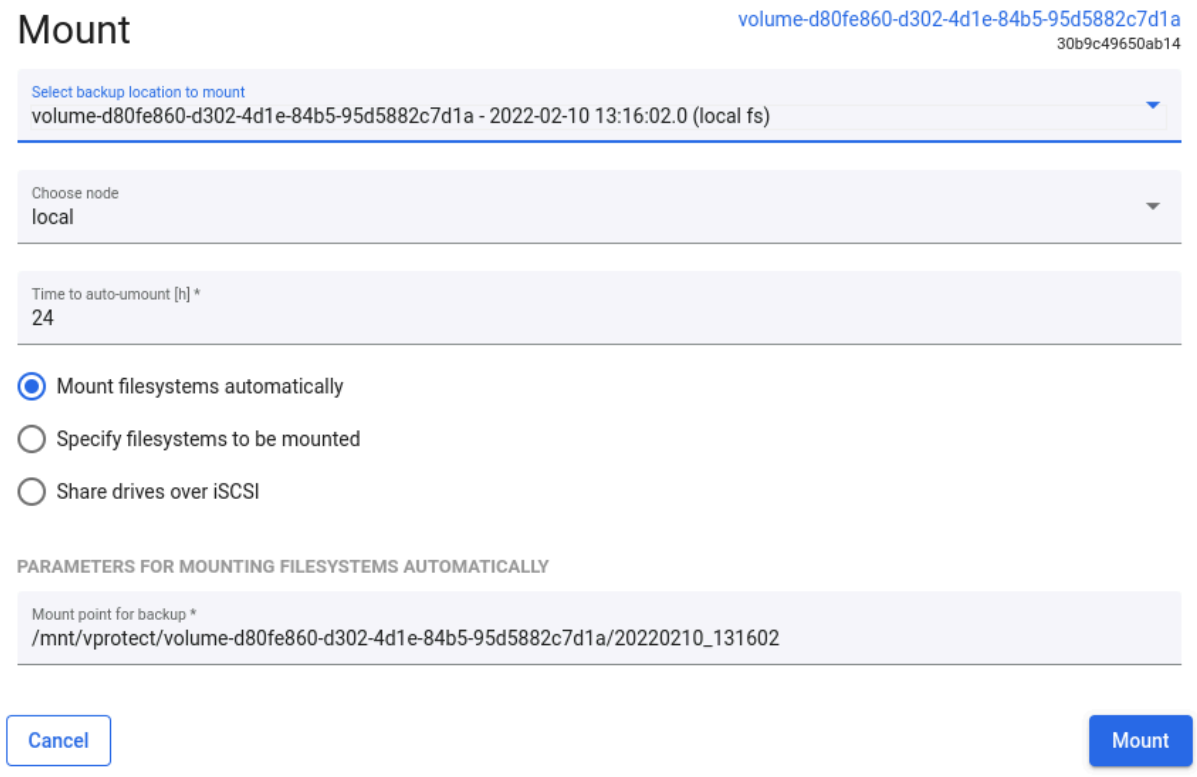

The Mounted Backups tab shows mounted storage-instance backups on the Data Protector for Cloud Workloads Node.

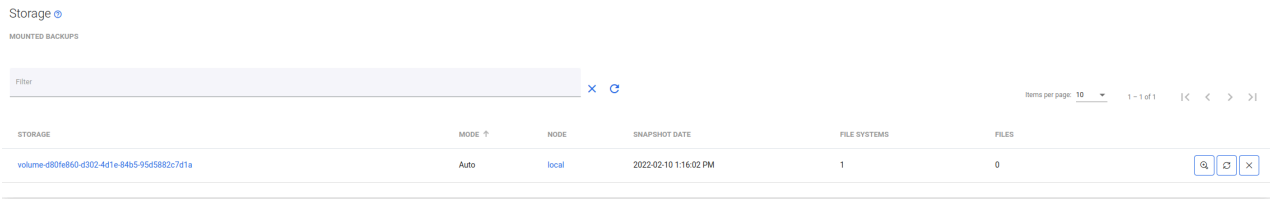

- `Storage` - mounted instance name
- `MODE` - Auto - Data Protector for Cloud Workloads auto-detects filesystems and mounts them on the path "/mnt/vprotect/". In Manual mode, the user chooses the mount point for selected filesystems.
- `NODE` - the Data Protector for Cloud Workloads Node is responsible for the mounting job.
- `SNAPSHOT DATE` - the date when the backup was created.
- `FILE SYSTEMS` - the number of mounted filesystems.
- `FILES` - the number of mounted virtual disk images.

Next to every mounted backup you can see three buttons: To unmount backup click on 

To remount backup click on 

To go to the details page of mounted backup click on 

On the details page, you can view basic information or you can go deeper and browse the files.

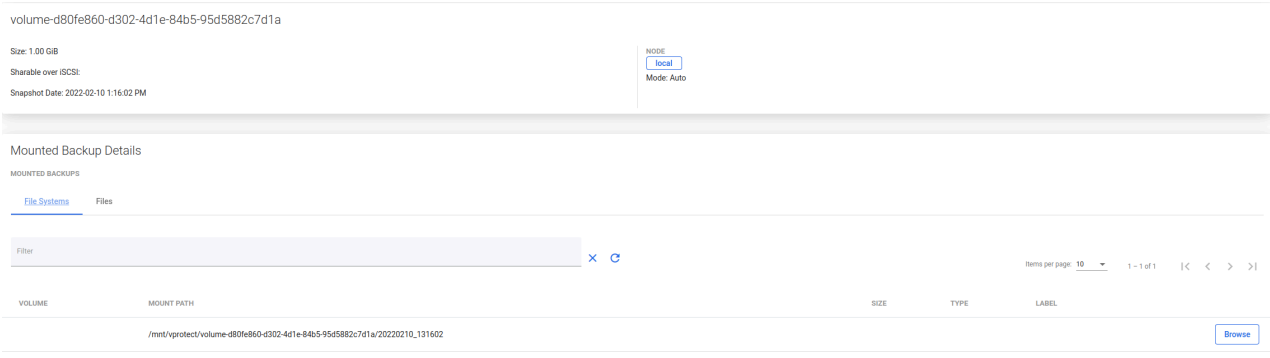

With a web browser, you can even obtain a single file from your storage instance backup.

Cloud

Cloud

This section provides information about administrative tasks like:

Instances - list of currently known entities and access to the details page of each object

Service Providers - access configuration for Microsoft365 organization.

Backup SLAs - allows you to setup a correlation between protected entities, the backup destination and schedules (Policies tab). It also allows you to configure policy schedules (schedules tab).

Instances

Instances

In this view, you can see all users and sites, and teams synchronized (imported) from the defined organization (Microsoft 365)

Click **Synchronize** button on cloud service providers or use account auto-synchronization to import users or sites and teams from Microsoft 365 organization.

Switch to the **Users, Grooups, Sites,** or **Teams** view to see only selected instances.

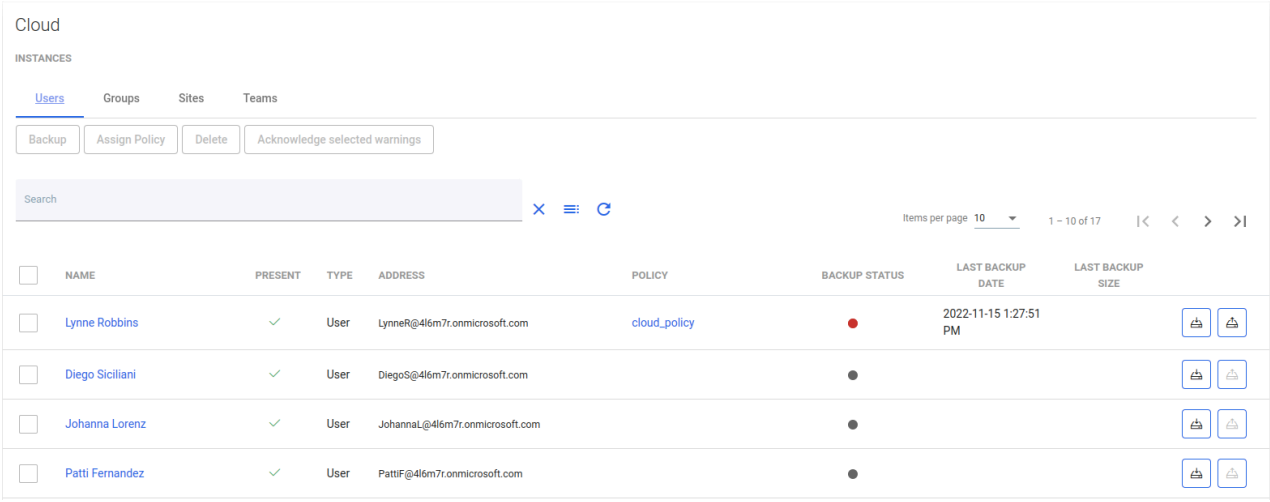

You can do the following activities on the protected entities (users or sites):

- **Backup** - to start a backup process for the selected entities
- **Delete** - to delete selected entities
- **Restore** - to restore entities data

> **Note:** Go to the Backup SLAs section to learn about how to configure a backup of your organization's data.

You can go into User/Group/Site/Teams to see the details page. As you can see, the window has been divided into several areas:

Summary

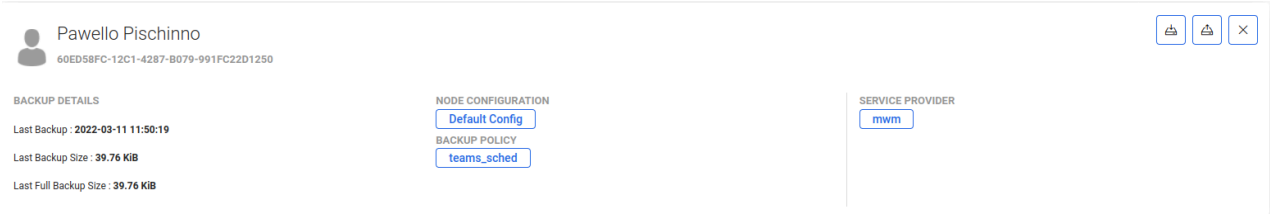

At the top, you can see summarized pieces of information about the selected entity, such as:

- the ID of the entity object in Data Protector for Cloud Workloads
- to which service provider belongs
- which node is assign to backup this entity
- short information about the last backup actions
- whether the entity has tags or policies assigned to it

You can also use several function buttons:

- refresh
- back to list
- change view
- backup
- restore
- delete

Backup/Restore Statistics

Daily activity

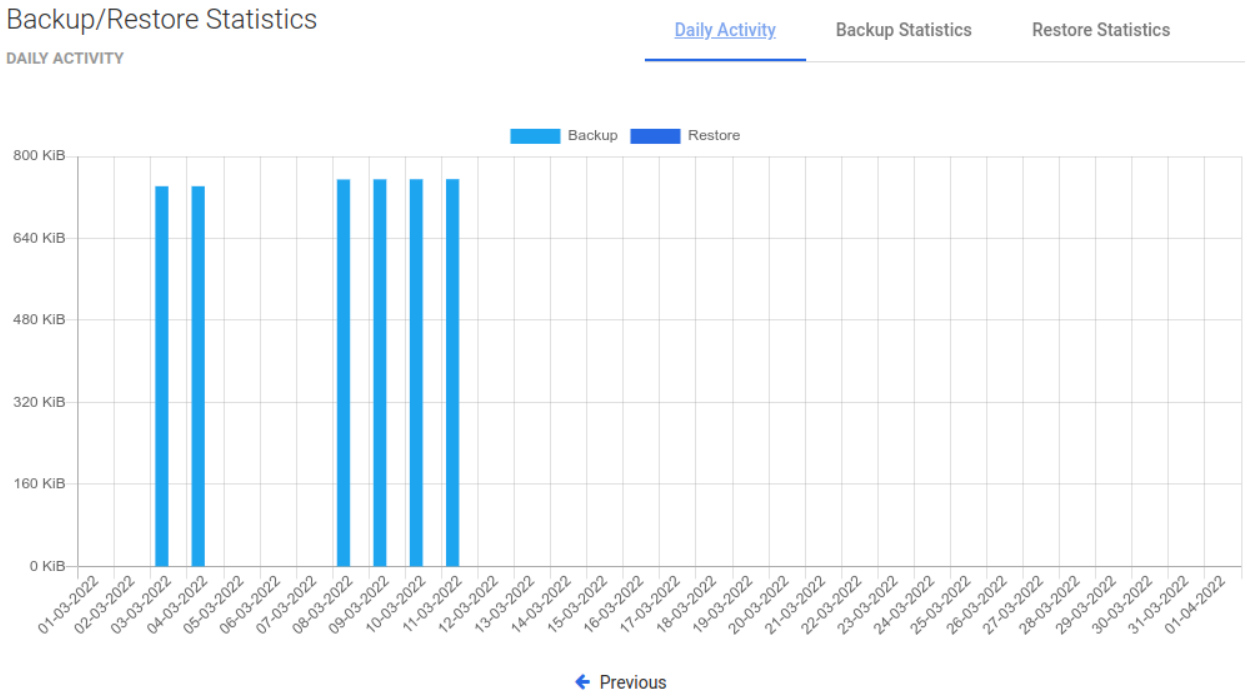

First, you'll see a daily summary of the backup and restore operations for the last month. This view is called "Daily Summary" and is the default view. You can switch the report between four views.

Backup Size

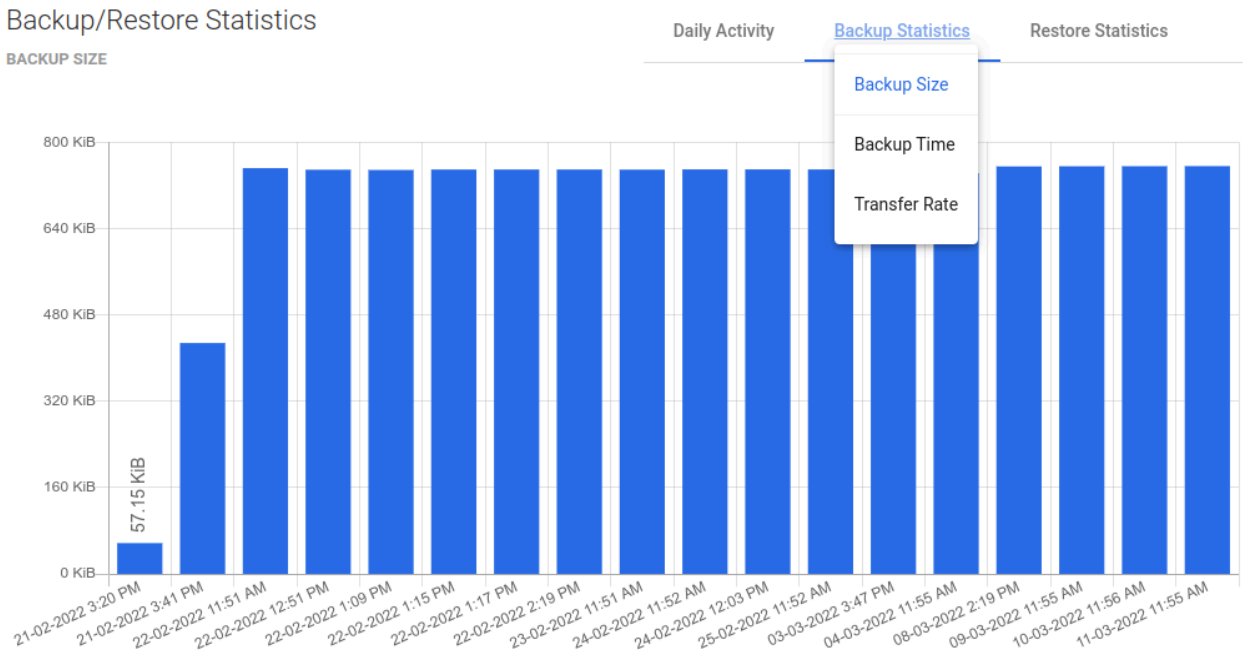

This view shows separate columns for each backup made to the virtual machine. Thanks to this, you can easily determine what data increase.

Backup Time

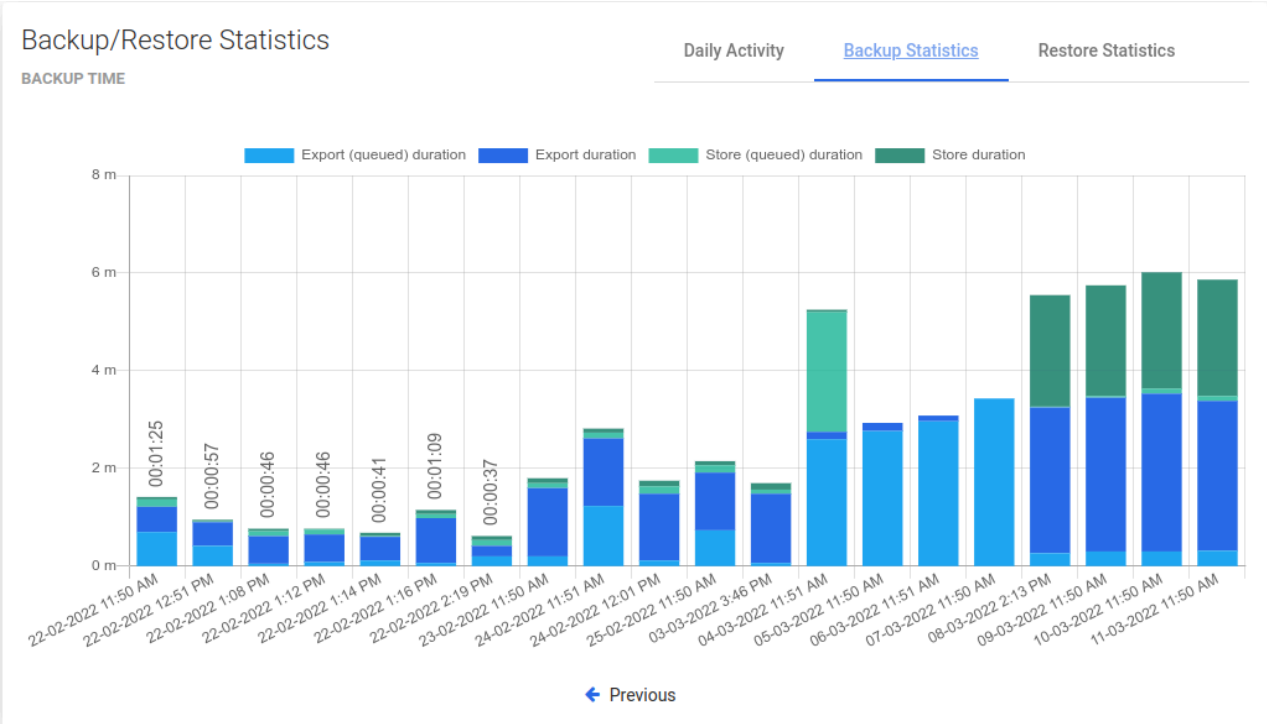

A very useful report. It allows you to determine the required window length for backups or, based on the time of individual phases, it is easy to deduce the cause of slow backups.

Transfer Rate

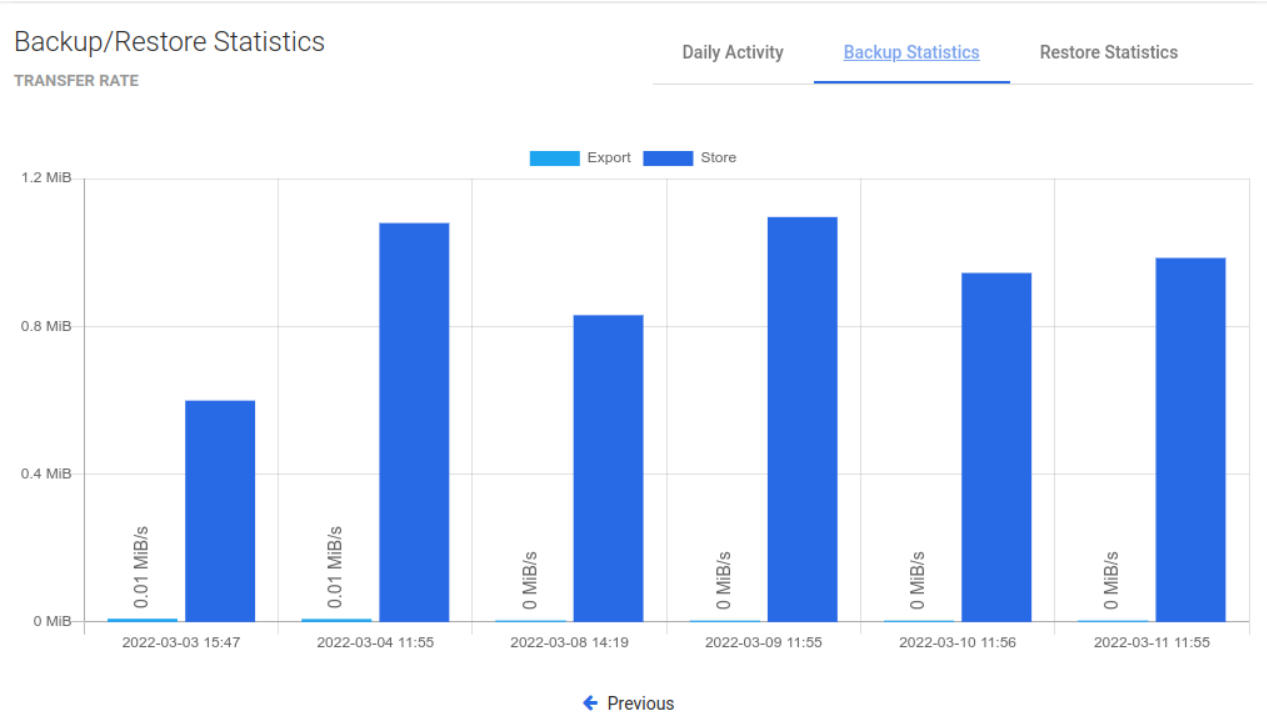

One of the latest reports, now you can easily see how fast data transfer is.

Restore Time

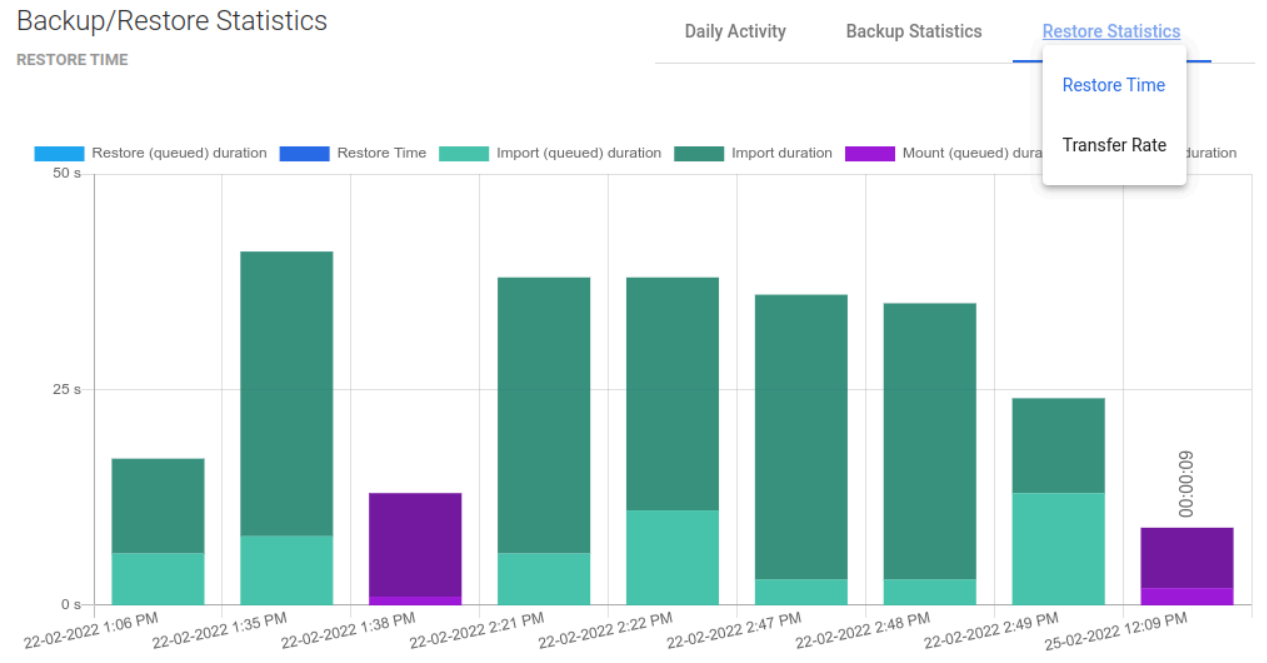

A view with the same properties as "Backup Time". It allows us to estimate how long it will take to restore the machine in the event of a failure.

Transfer Rate

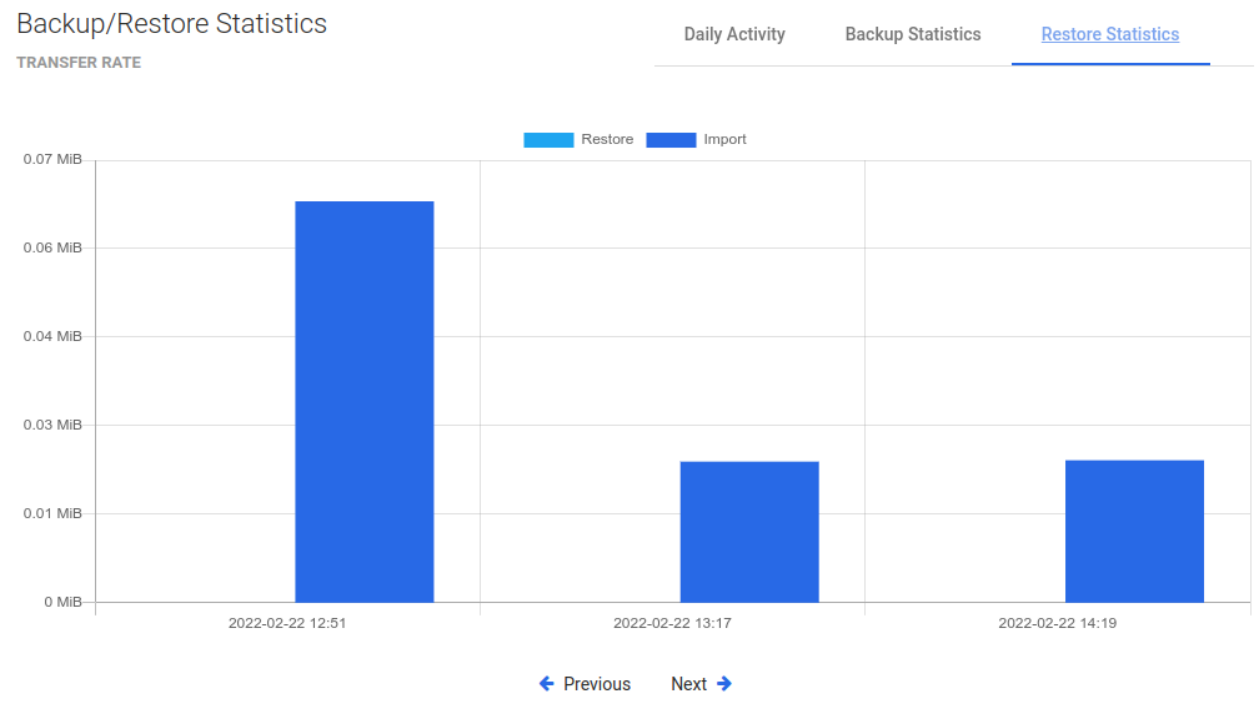

As in the previous case, we also have the transfer speed for the restore job.

Events Calendar

The calendar extends the possibilities of adjacent statistics. It allows you to neatly define the range of days you want to see, additionally makes a quick summary of the number of backups and restores (top right corner).

Blue - the sum of all backups, **Green** - the sum of successes, **Red** - the sum of failures.

Bottom menu

In the bottom menu, you can find a large number of tabs, each of which will present different information or will allow you to restore one or more objects from the selected user, site, team.

Backup History

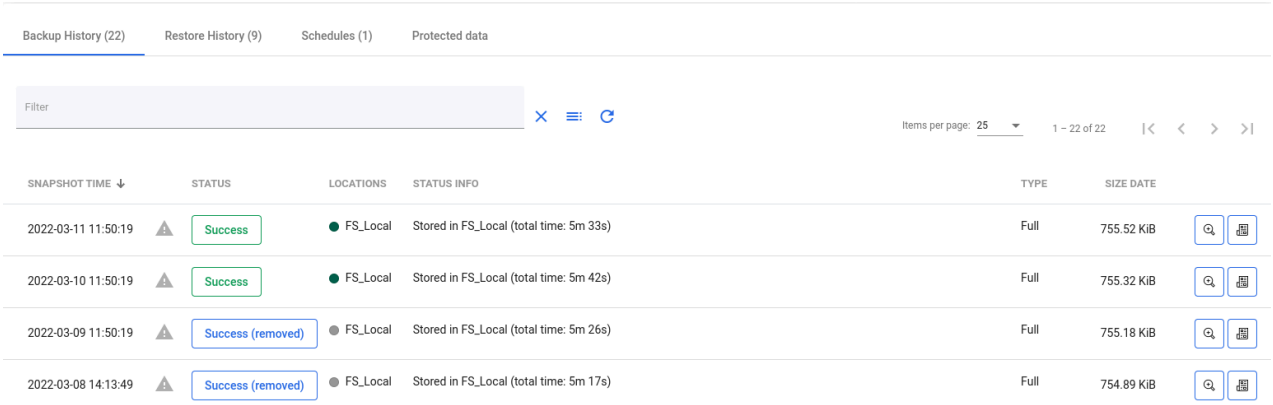

This tab shows information about all backups made, as well as information about failed, removed (because of retention), or currently executing backups.

Restore History

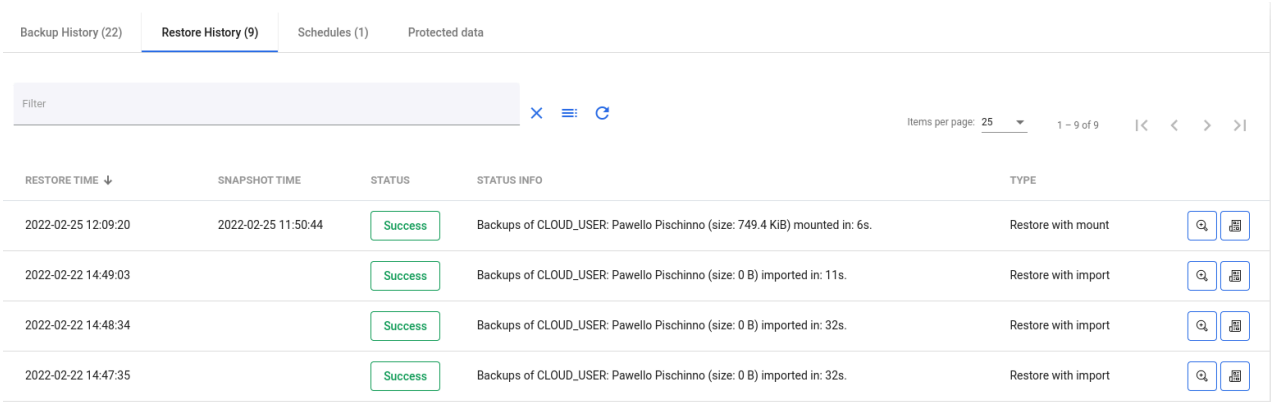

This tab is similar to "Backup History". This is a list with basic information about restores performed. When you open the details of the selected restore, you will see

more detailed information.

Schedules

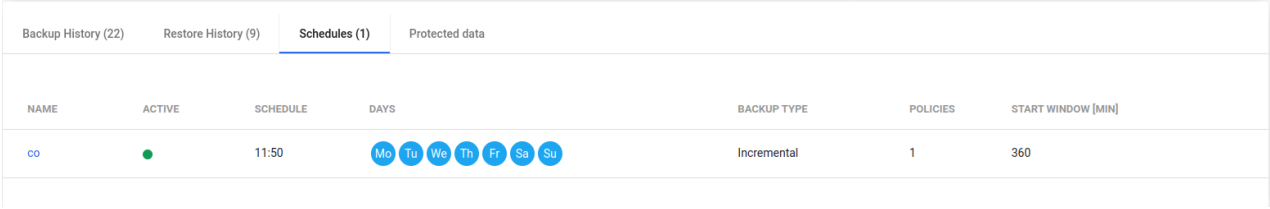

In this tab, you can see all the schedules assigned to this entity.

Protected data

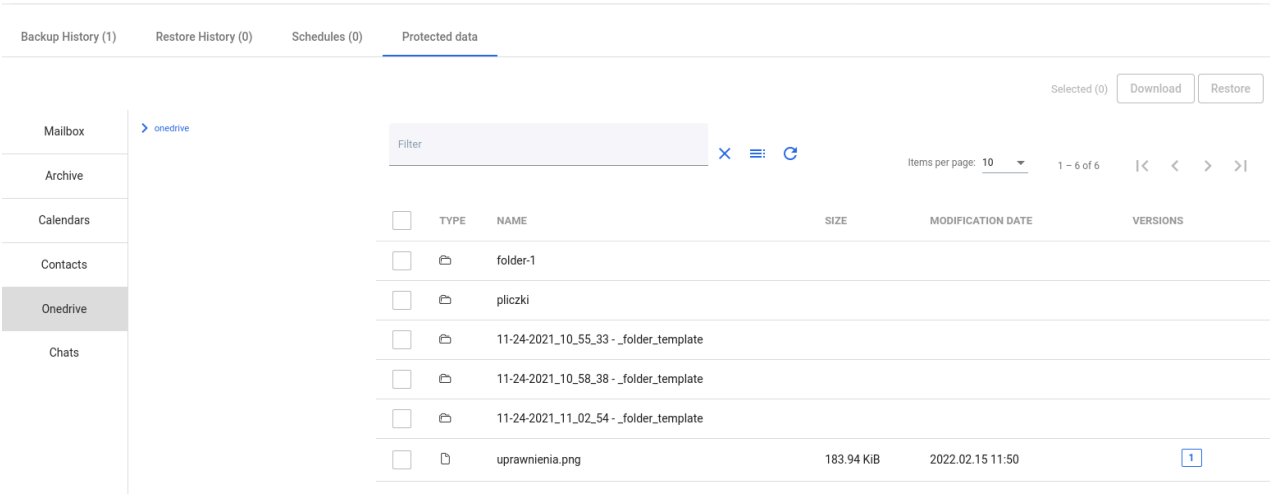

In this tab, you can see all backup data which you can select for restore or download.

You can always use the filters to move back in time to select the exact moment from the calendar and restore or download data.

> **Note:**
>
> If you are restoring to another path, provide the name of the library first.

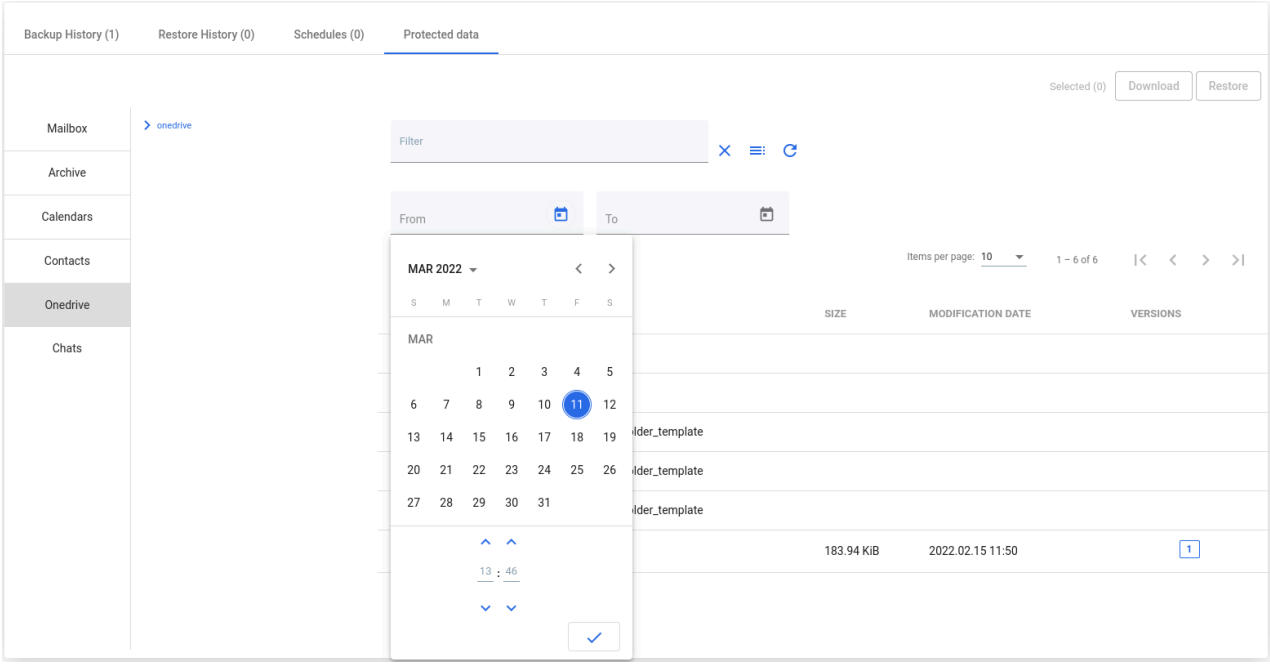

Export to PST

This feature allow you to download data and save it as PST file. Currently available only for Mailboxes.

Requirements

- Machine with Windows OS
- Microsoft Outlook 64-bit installed on mentioned machine

You need to enable WinRM and add firewall rules on the Windows machine. You can use ready scripts to do it. You can find these scripts in `/opt/vprotect/` directory:

- scripts/winrm_firewall_windows.ps1 - copy from node to target Windows host and execute
- scripts/winrm.ps1 - copy from node to target Windows host and execute

Next, download from Micro Focus website and install CLOUD2PST Converter application.

Exporting to PST

1. Go to **Nodes** → **Node Configurations** and edit existing node configuration. Click on **Cloud Agent** tab and type access credentials to your Windows machine which will be used to generate PST.

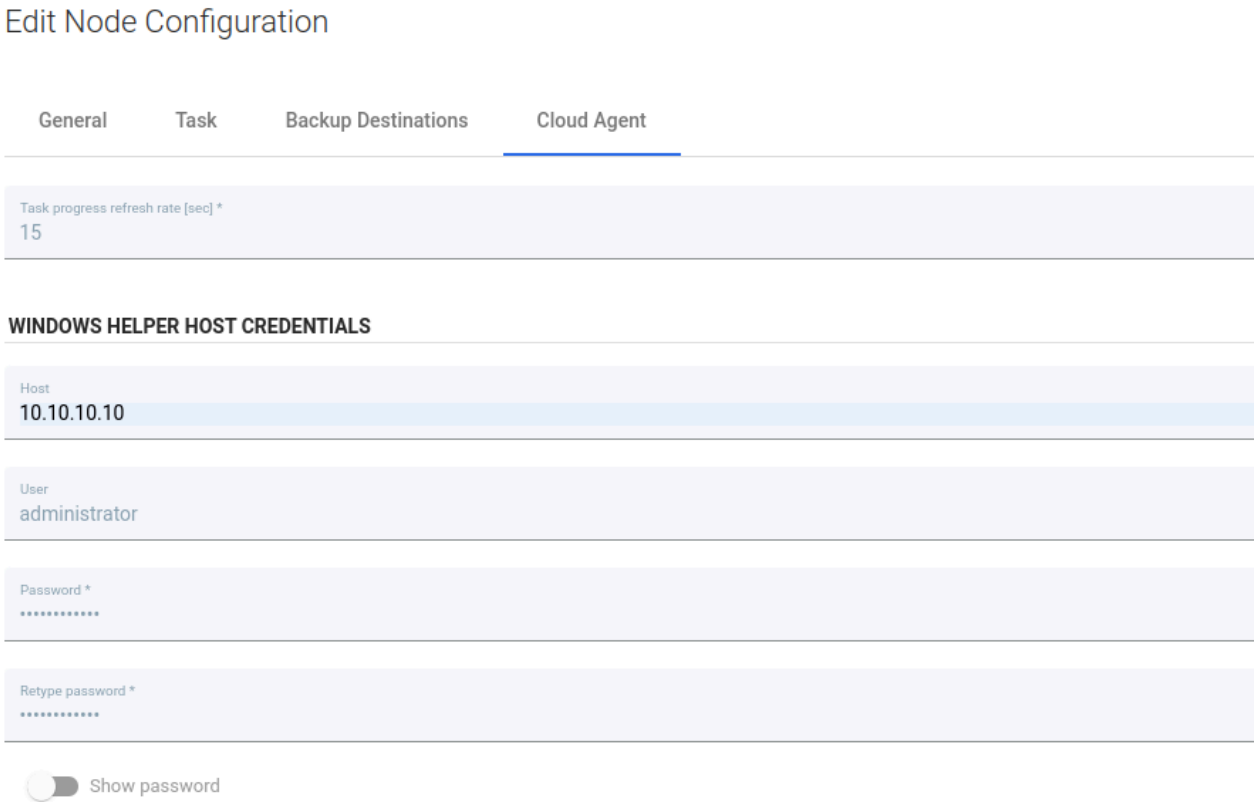

2. In the **Protected data** tab of the selected M365 user, click on item you want to export and then click **Export to PST** button. The download task will be created.

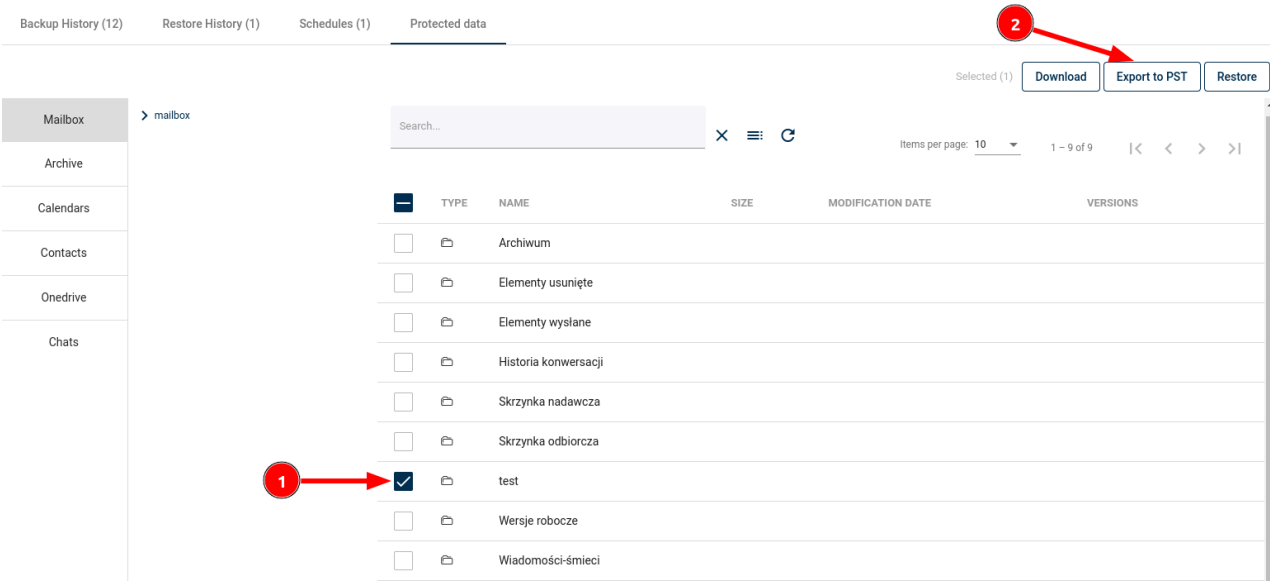

3. After download task is complete, go to **Cloud** → **Download**. Here, click on the **Download** button to download your exported PST file.

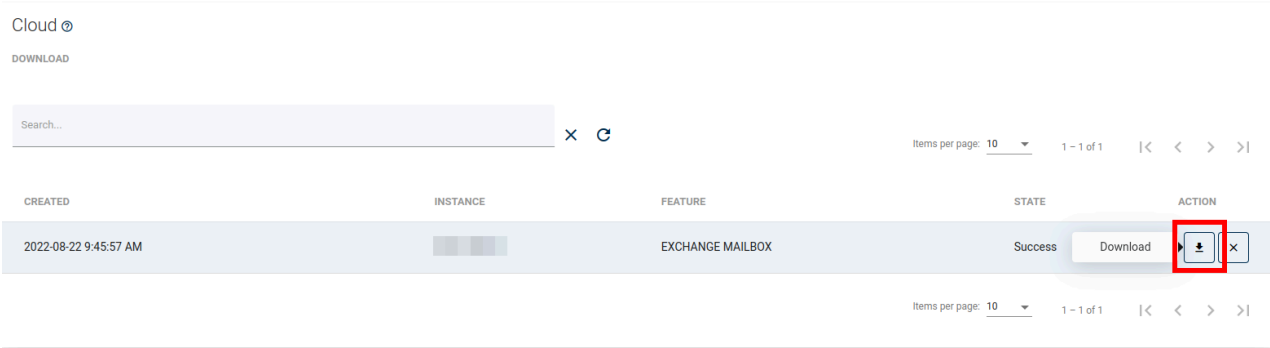

Service Providers

Service Providers

On this page, you can add, edit, or delete Microsoft 365 organizations.

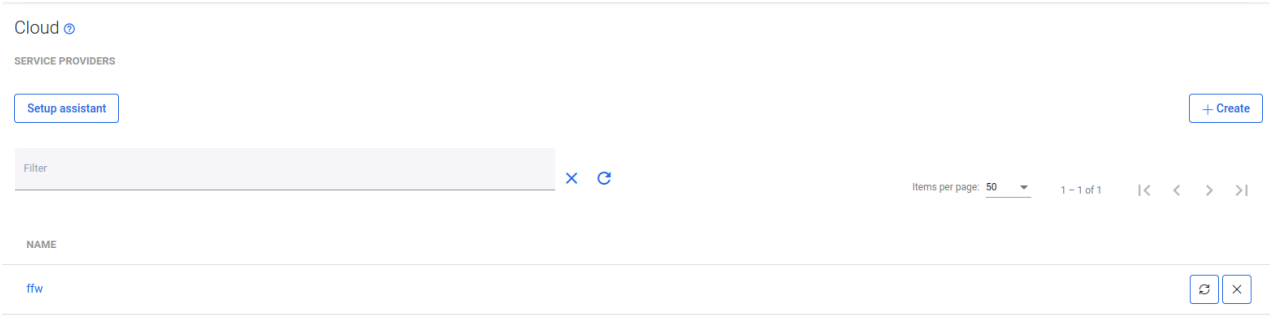

You can add more than one Microsoft 365 cloud service provider.

You can also on this tab do manual synchronization.

> **Note:** How to add new service providers we describe in the chapters manual or by using the setup assistant

Backup SLAs

Backup SLAs

Backup SLAs policy allows you to configure the following settings:

- Set backup retention in day(s), month(s), year(s)
- Number of protected item(s) or file(s) versions
- Select the desired backup destination
- Protected Microsoft 365 services (SLA):
  - Exchange Online Mailboxes
  - Exchange Online Archives
  - Exchange Online Calendars
  - Exchange Online Contacts
  - OneDrive for Business
  - User Chats
  - Sharepoint Online
  - Microsoft Teams

If you need to add organization user(s) or site(s)to the policy go to Users, Groups, Sites, or Teams Tabs and select required instances.

> **Note:** There must be at least one Backup SLAs policy on Data Protector for Cloud Workloads Server to do a backup.

Use `Create` button to create a new policy.

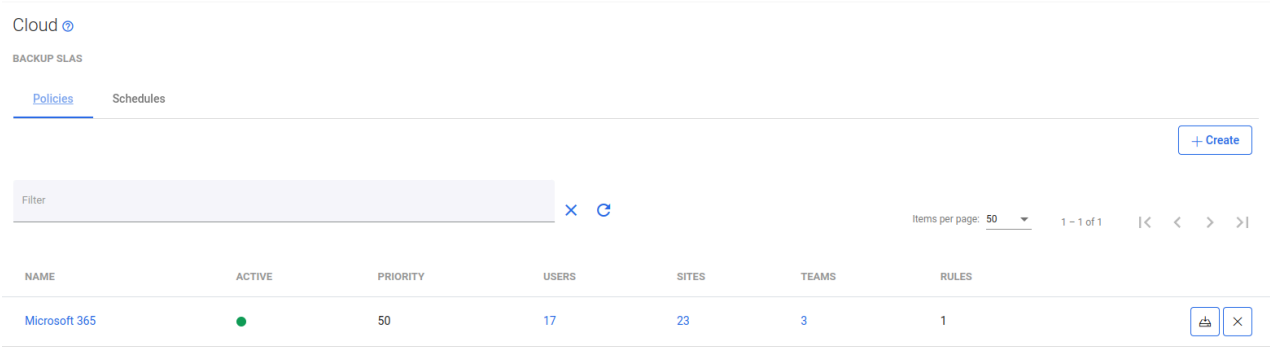
Cloud
BACKUP SLAS
Policies
Schedules
+ Create
Filter
Items per page: 50
1 – 1 of 1
NAME
ACTIVE
PRIORITY
USERS
SITES
TEAMS
RULES
Microsoft 365
50
17
23
3
1

Policies

Policies

Policies allow you to group cloud instances in many ways. For example, based on the cloud features.

To create a new backup policy, open the Backup SLAs tab under the Cloud section and click on + Create the button on the right.

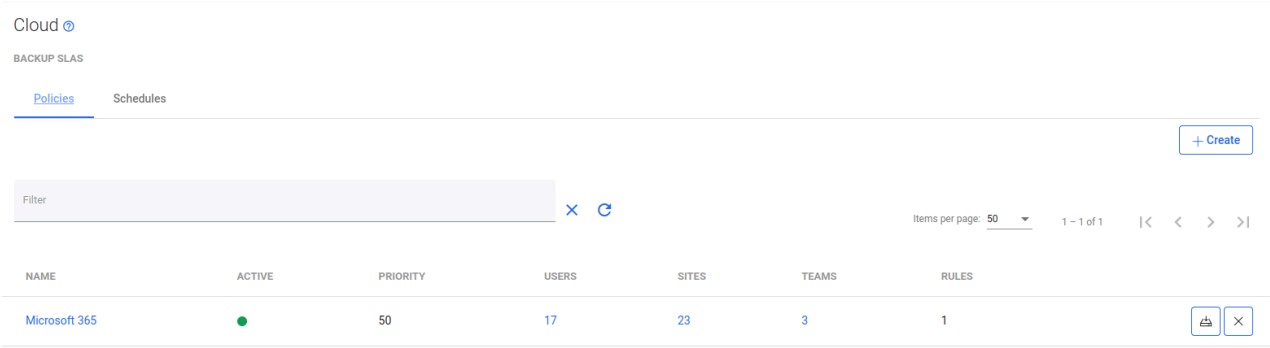

Now you should see the policy wizard with 8 main sections.

General

Under this section you can set up:

- The policy name
- Switch on/off scheduled backups enabled
- Set the priority for tasks

Auto-assignment

In this section, you can switch on/off automatic policy assignments to cloud service providers

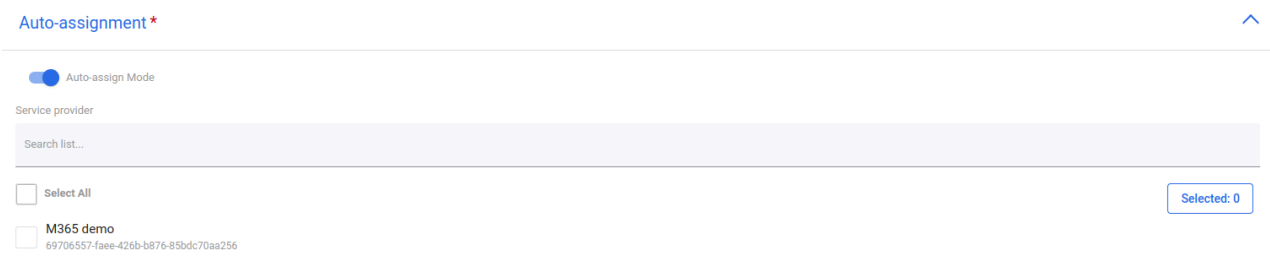

> **Note:** Go to the Account auto-synchronization to learn how to synchronize Microsoft 365 accounts.

Features

Here you can easily select features that will be backup according to policy.

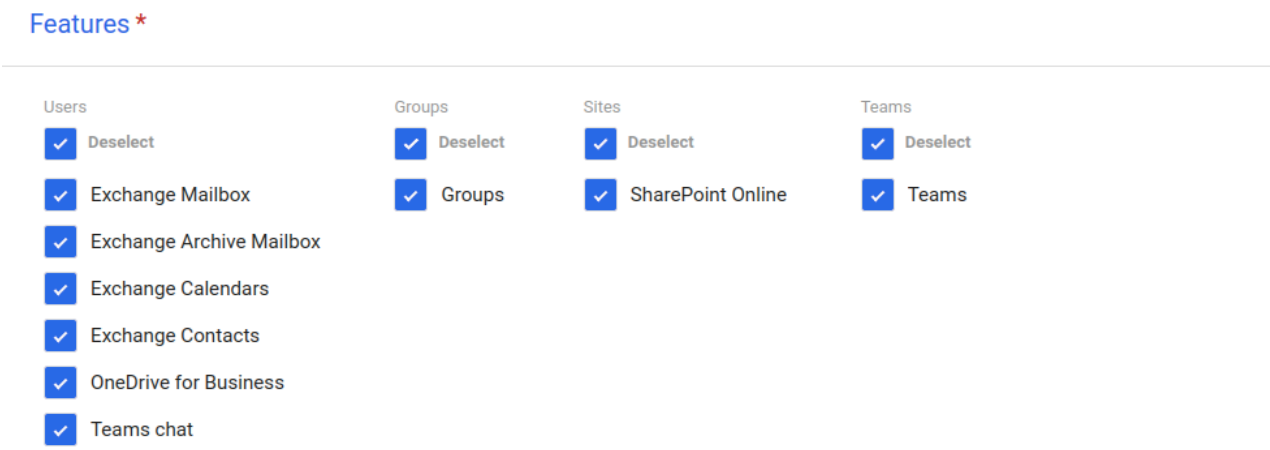

Users / Groups / Sites / Teams

These sections are used to select instances that will be protected in policy.

Rule

This section is used to select the backup destination.

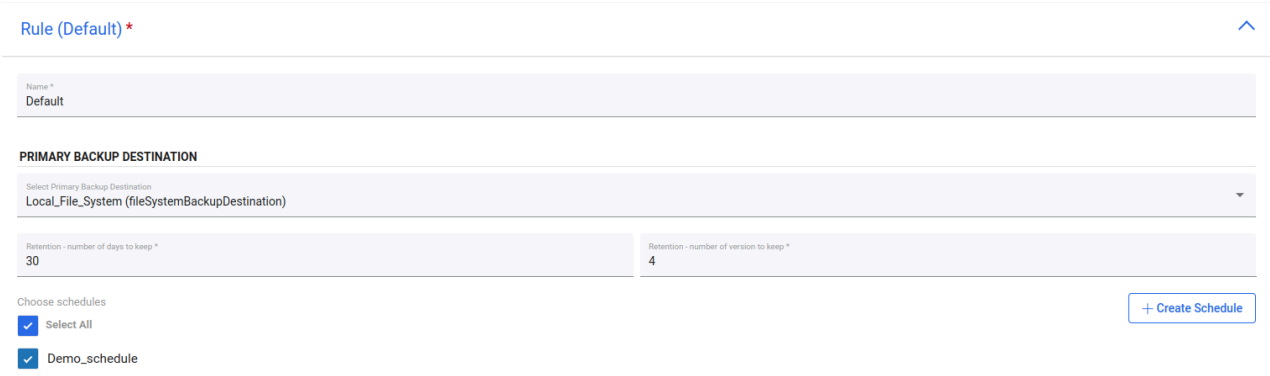

> **Note:** You can use only File System, Microsoft Azure Blob Storage, Micro Focus Data Protector as a backup destination.

You can also set here **Retention** settings for your backups. You can use a number of days and versions to keep.

If you have already created a schedule, you can also select it or Create New Schedule.

Retention

Data Protector for Cloud Workloads handles retention for all backup destinations. There are 2 properties that define how long backup should be kept in the backup destination:

- `Retention - no. of versions to keep` - number of versions to keep
- `Retention - no. of days to keep` - number of days to keep

Other

This is an optional section with two switches:

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

Two examples when using switches is useful It is very likely that if 30% of the backup tasks fail, the remaining tasks will also fail because the environment has failed. Or you are backing up a set of machines, and if even one is not secured, there is no point in backing up the rest.

In the end, save settings.

Schedules

Schedules

The schedules allow you to invoke specific policies periodically. This allows you to back up multiple Cloud instances automatically.

A schedule defines when and on which days cloud instances should be backup. To define a new schedule, open Backup SLAs under the Cloud section and go to the Schedules tab, then use the + Create button.

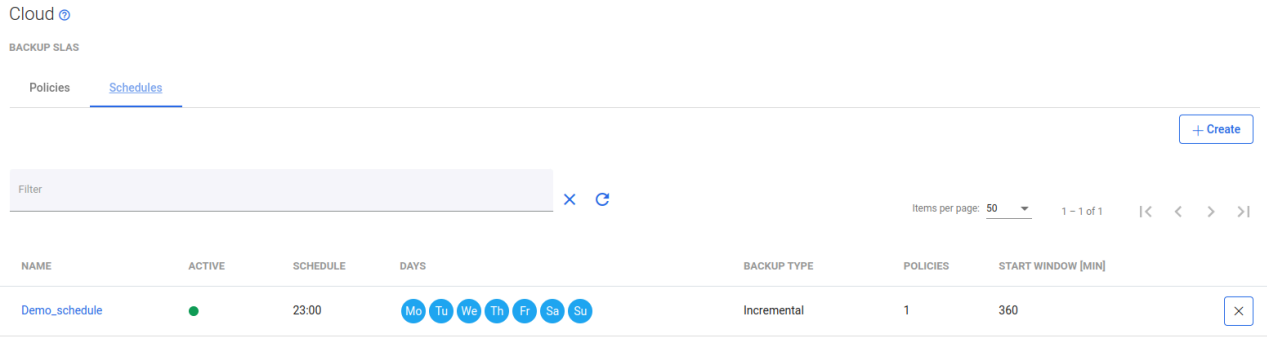

Now enter the properties:

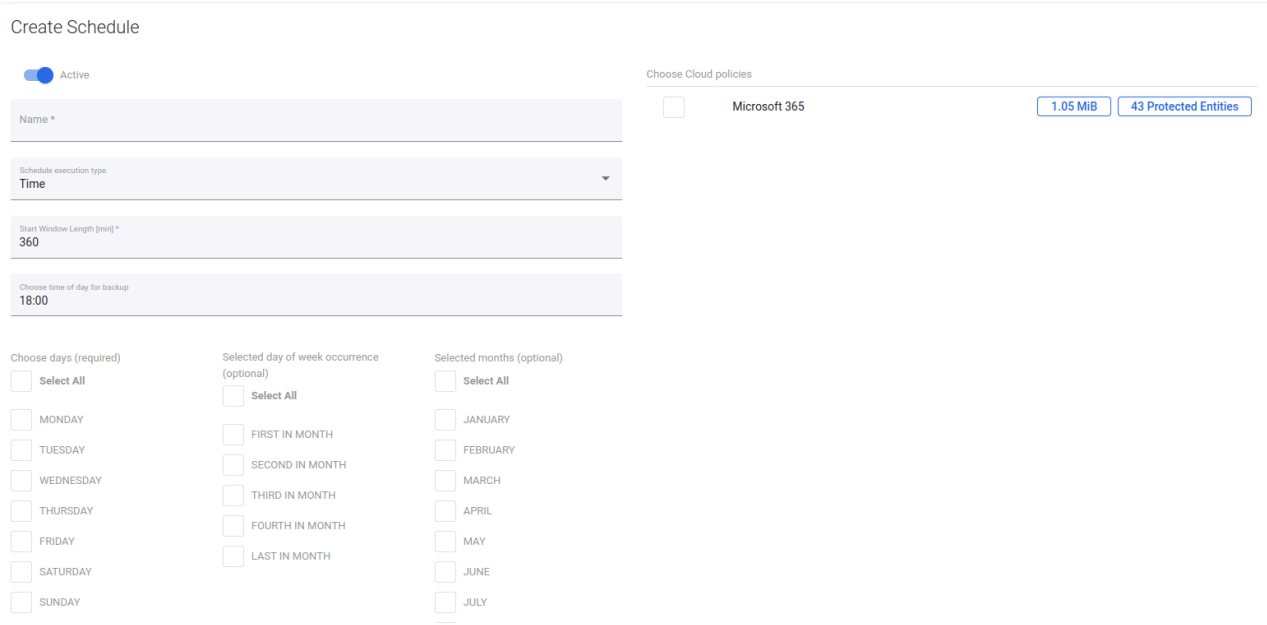

- `Schedule Active` - enable or disable executing schedule

- `Name` - schedule name
- `Execution Type` - choose the time or interval mode
- `Start Window` - defines for how long since the task start time scheduled tasks are allowed to be executed
- `Choose time of day` - for the time execution mode, this defines when the task should be added to the queue
- `Choose time of interval start` - for the interval execution mode, this defines when tasks should start
- `Choose time of interval end` - for the interval execution mode, this defines when tasks should end
- `Frequency` - defines how often the task will be executed during the interval
- `Choose days` - the last required parameter, select the days of the week on which the task will be performed

You can also use optional parameters to further personalize the backup time or select a cloud policy if it has been previously created.

> When you set the time with a user in a certain time zone, you specify a point in time at which you want the schedule to start. Changing the timezone doesn't change this point in time, it's converted to your timezone. The time displayed to the user is calculated based on the server time.

Download

Download

All downloads tasks that have been will be, or are being performed on the Data Protector for Cloud Workloads Cloud are collected in the **Downloads** view.

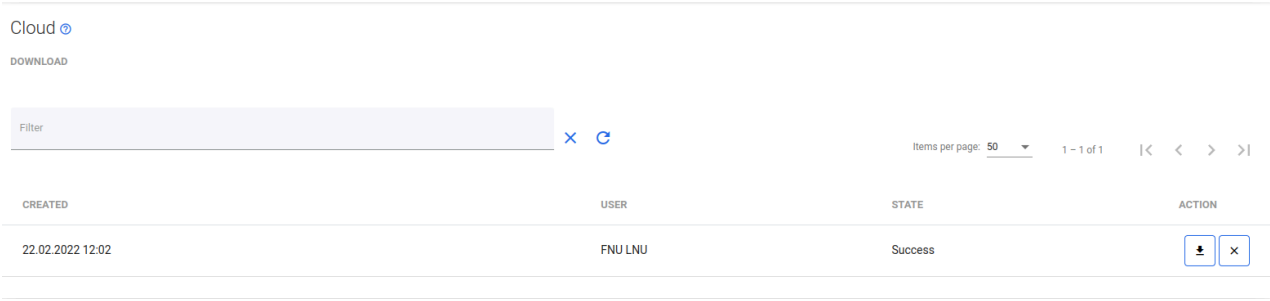

Here you can download the elements you selected previously and that has been prepared by the Data Protector for Cloud Workloads agent.

Applications

Applications

Data Protector for Cloud Workloads allows you to setup periodic backup with your own scripts or application-native backup commands. These can be executed either on the Node or remotely over SSH.

The Application can reside anywhere - it can be in a VM, Kubernetes deployment, or on a physical box.

> **Note:** The only requirement is to allow Data Protector for Cloud Workloads to execute a set of commands over SSH or on the node to access data remotely.

In this section are:

Instances - Create Application definition (at least its name and Command Execution Configuration and select node which is going to do the work)

Execution Configurations - Prepare script or commands (this is a description of how your script is going to be invoked)

Backups SLAs - allow you to setup a correlation between applications, backup destinations, and schedules. It also allows you to configure policy schedules (schedules tab).

Instances

Instances

General

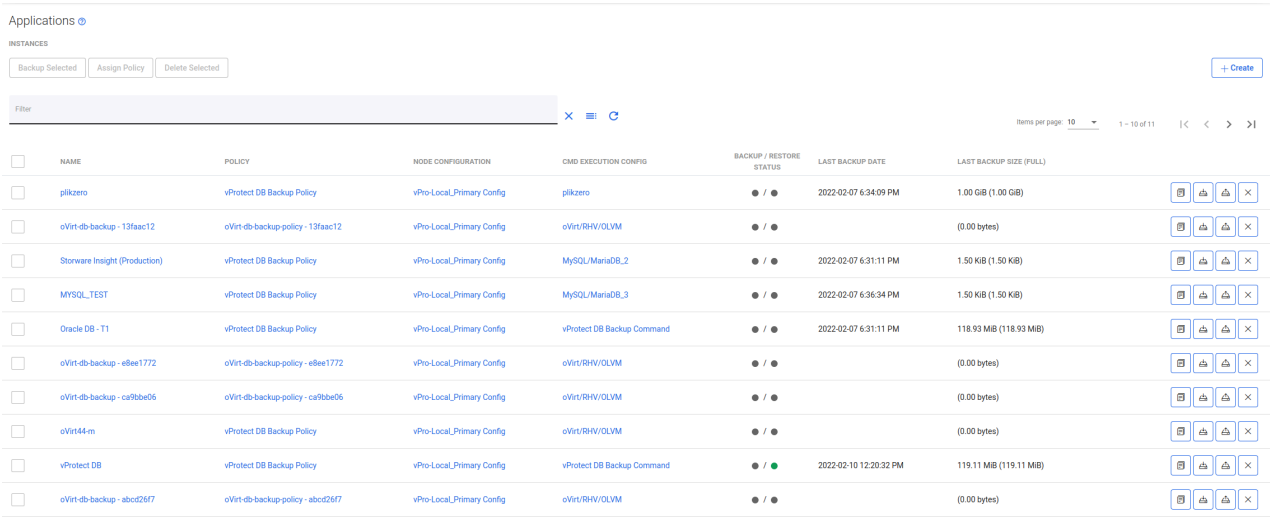

This tab allows you to create an "Application definition" and execute on-demand actions like:

- Clone
- Backup
- Restore
- Delete

Details page

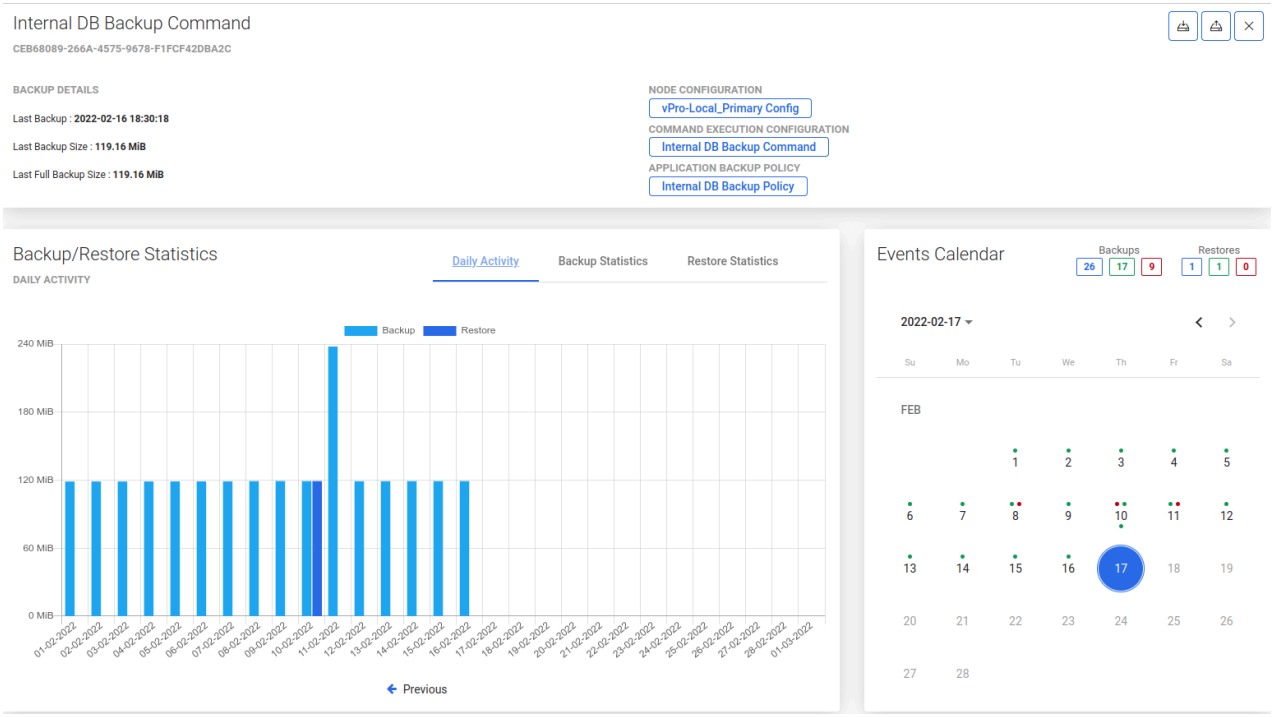

As you can see, the window is divided into several areas like

Application summary

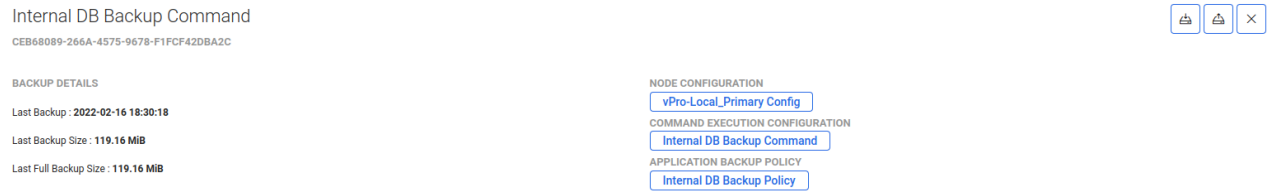

At the top, you can see summarized pieces of information about the Application, such as:

- Name of application object in Data Protector for Cloud Workloads
- To which command the execution is assigned
- Which node is backing up this application
- Short information about the last backup actions

You can also use several function buttons, such as:

- refresh
- back to instances list

- backup
- restore
- delete

Backup/Restore Statistics

Daily activity

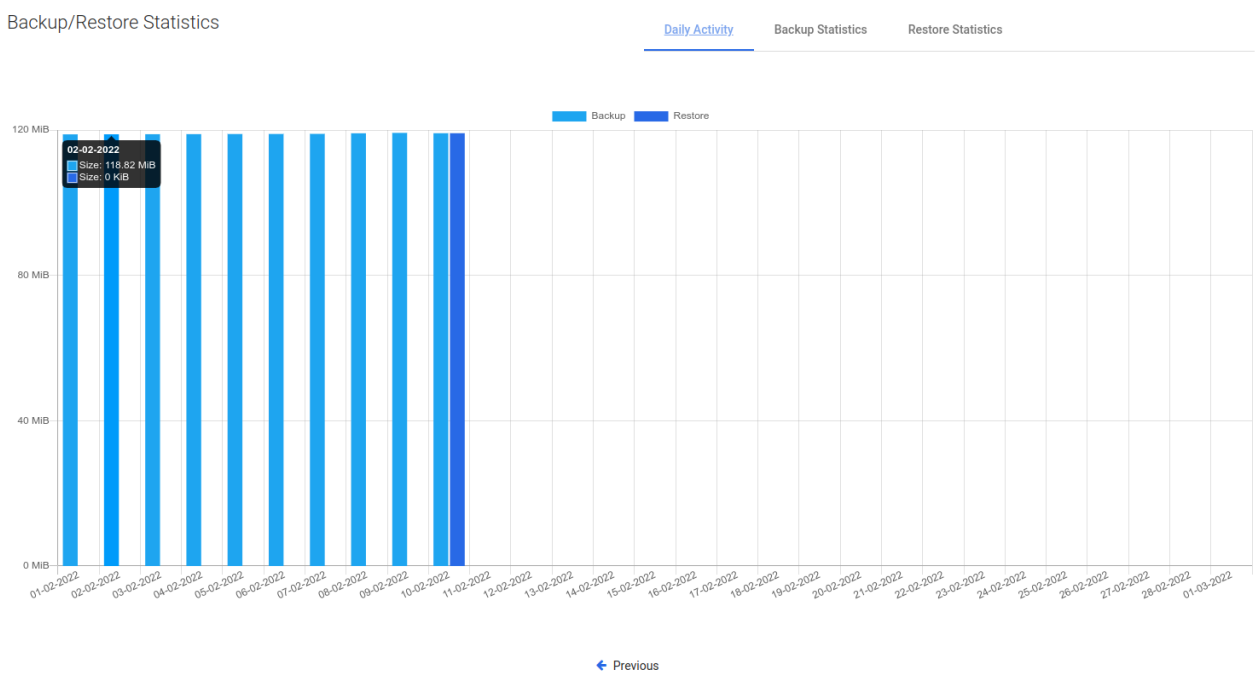

First, you'll see a daily summary of the backup and restore operations for the last month. This view is called "Daily Summary" and is the default view. You can switch the report between multiple views.

Backup Size

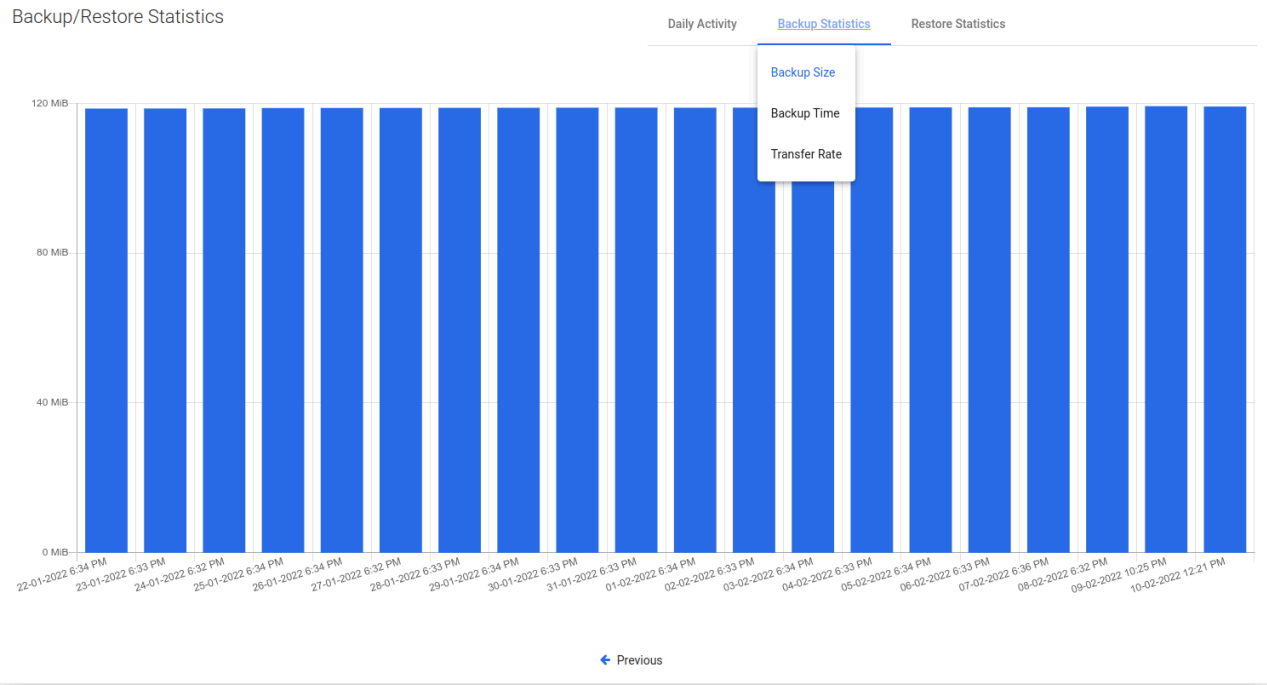

This view shows separate columns for each backup made to the application. Thanks to this, you can easily determine what data increase has occurred on a given app.

Backup Time

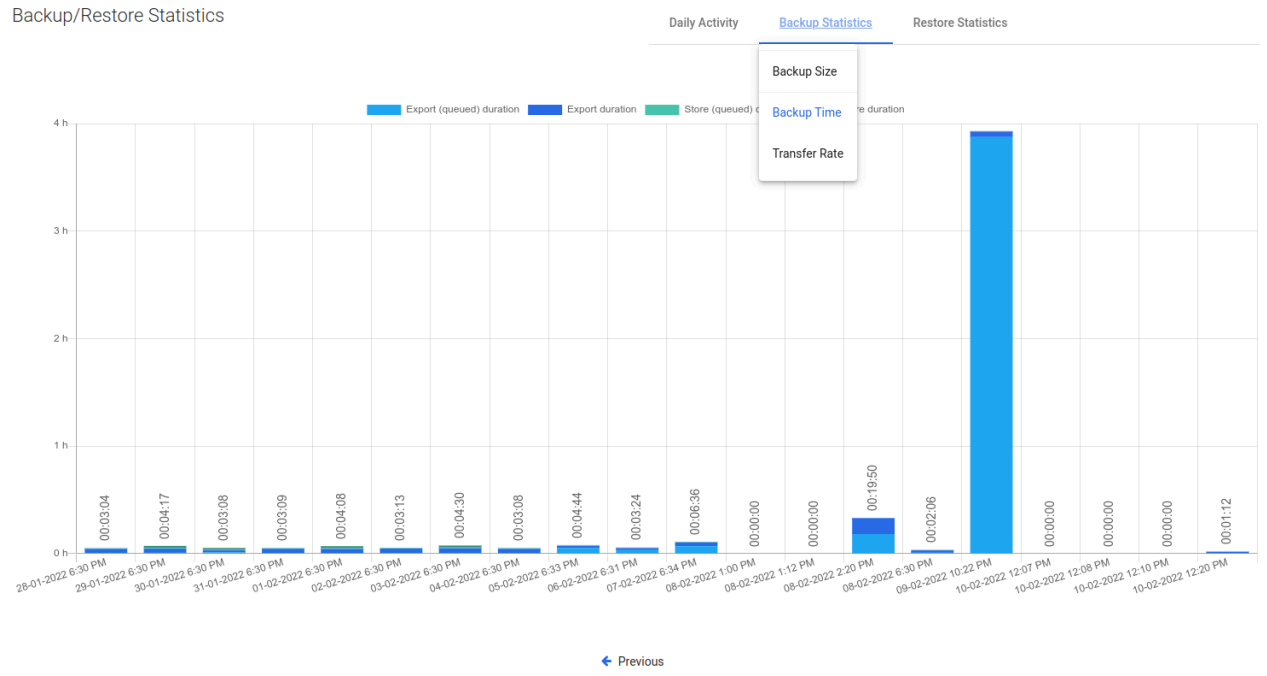

A very useful report. It allows you to determine the required window length for backups or, based on the time of individual phases, it is easy to deduce the cause

of slow backups.

Transfer Rate

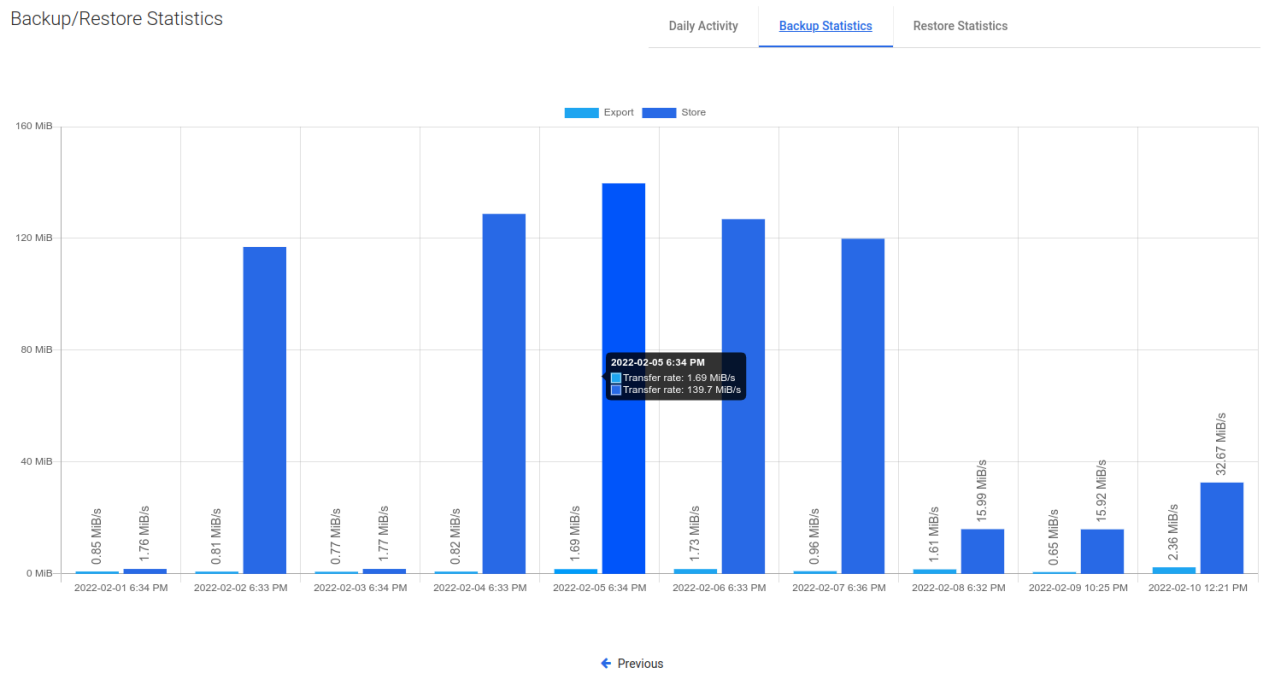

One of the newest reports - now you can easily see how fast the data transfer is.

Restore Duration

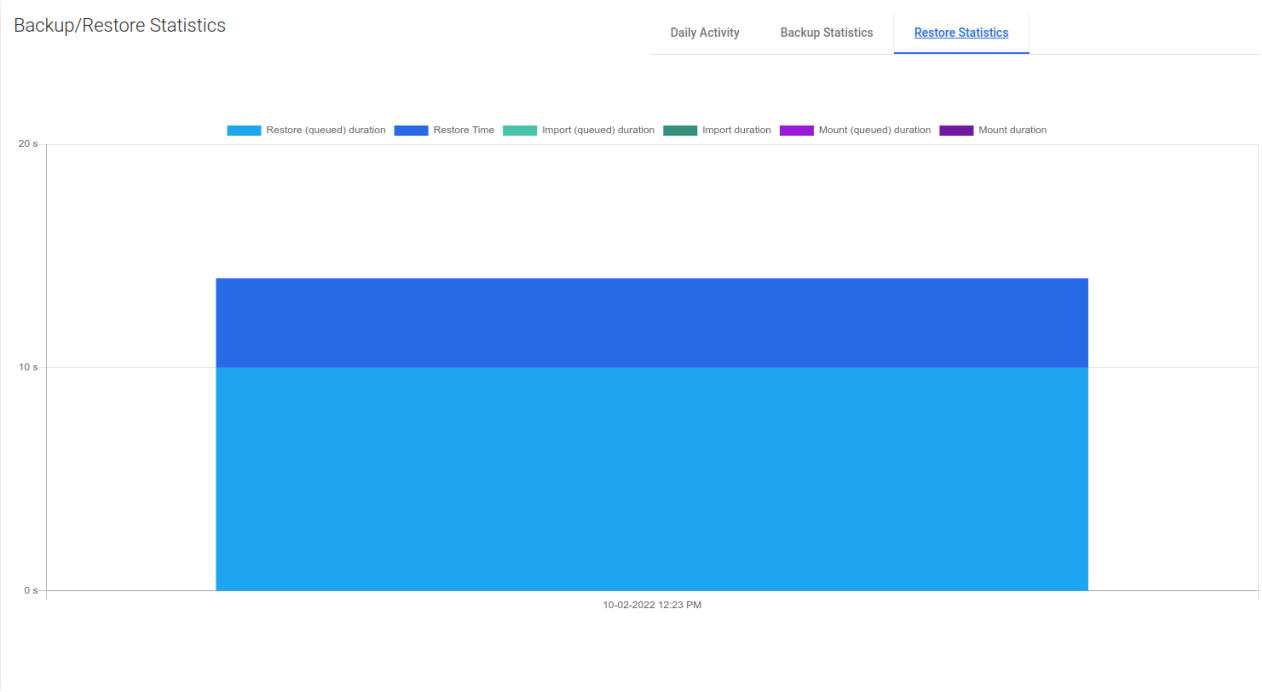

A view with the same properties as "Backup Time". It allows us to estimate how long it will take to restore the application backup files in the event of a failure.

Restore Rate

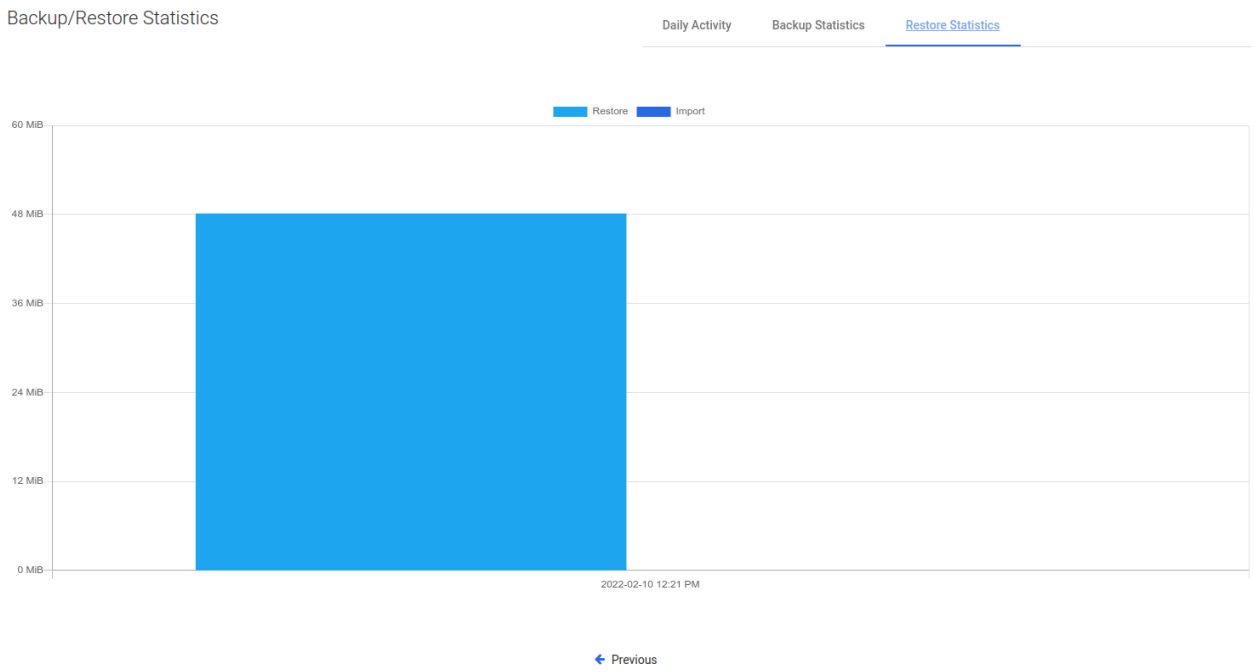

As in the previous case, we also have the transfer speed for the restore job.

Events Calendar

The calendar extends the possibilities of adjacent statistics. It allows you to neatly define the range of days you want to see, and additionally makes a quick summary of the number of backups and restores (top right corner).

Blue - the sum of all backups, **Green** - the sum of successes, **Red** - the sum of failures.

Bottom menu

In the bottom menu, you can find a large number of tabs, each of which will present different information or will allow you to change the configuration of the particular application.

Backup

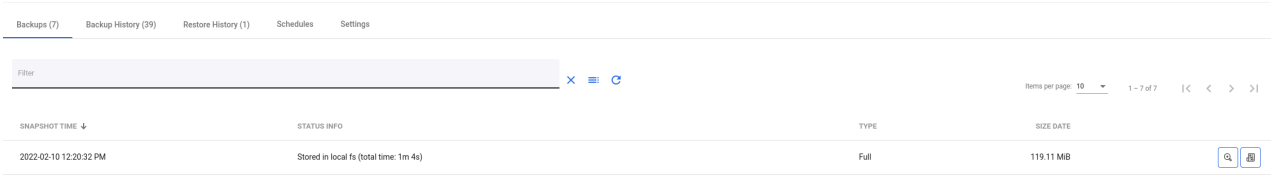

The first tab shows all application backups that are currently available and all basic information about them in a list. After pressing the magnifying glass button, you will see additional information. The button next to it allows you to download logs in the form of a .txt file.

Backup History

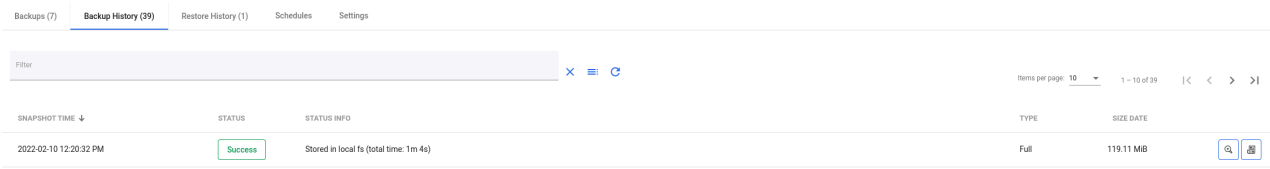

This tab shows information about all backups made for this application, including those failed, removed (because of retention), or currently executing.

Restore History

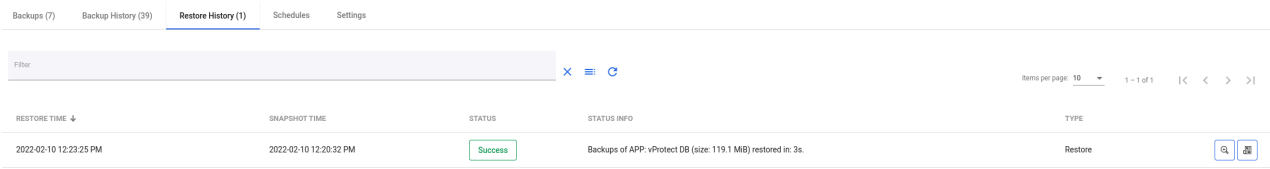

This tab is similar to "Backup History". This is a list with basic information about the application restores performed. When you open the details of the selected restore, you will see much more detailed information.

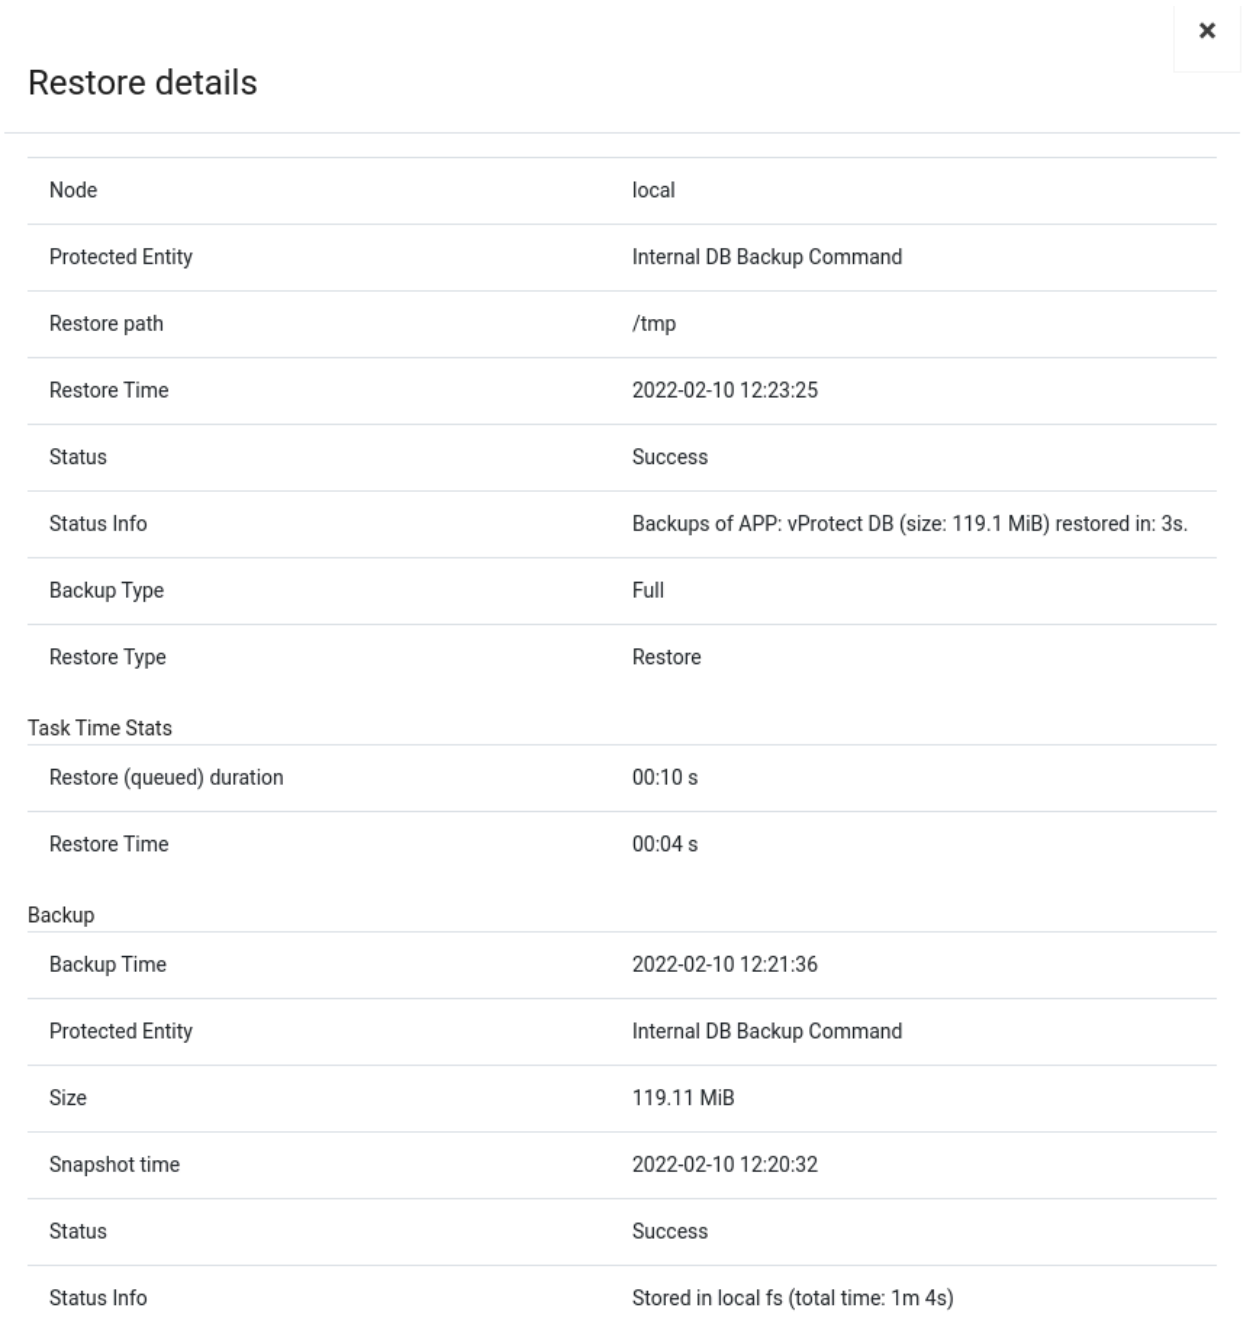

Schedules

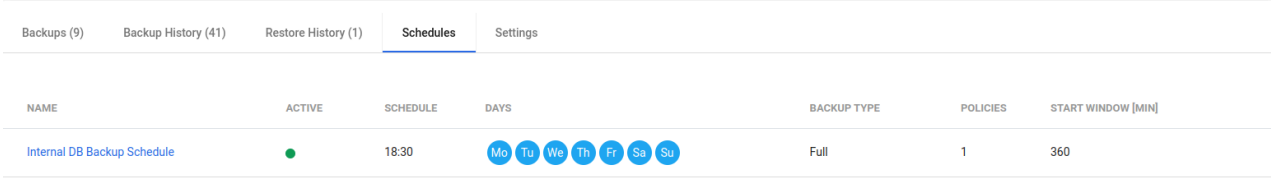

In this tab, you can see all the schedules assigned to the application.

Settings

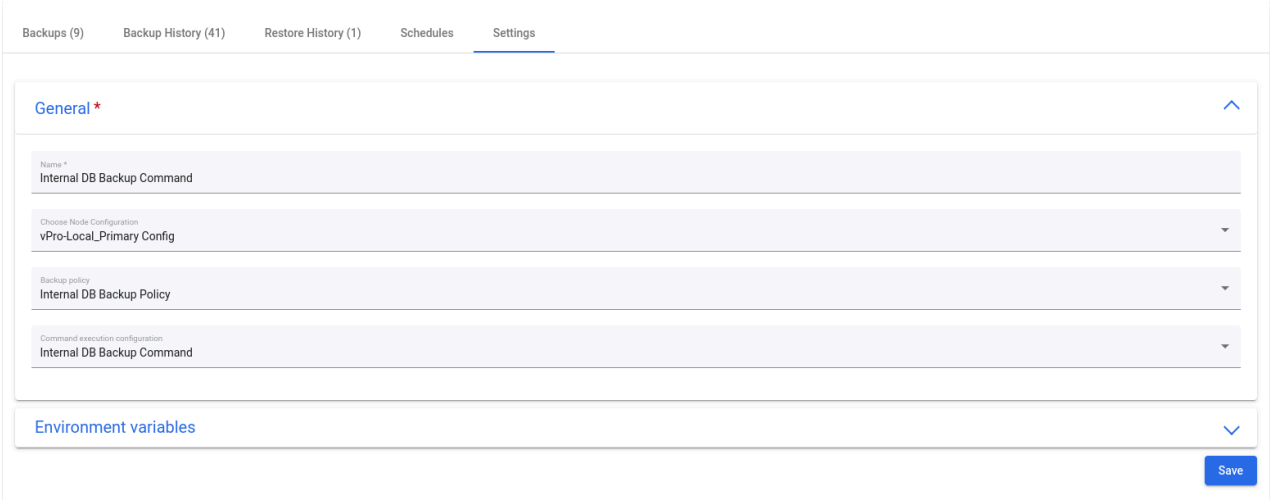

Finally, the last tab. The first two options allow you to change the node selected to perform backups and policies assigned to the application. The third allows you to choose the execution configuration.

Example - How to create an application definition

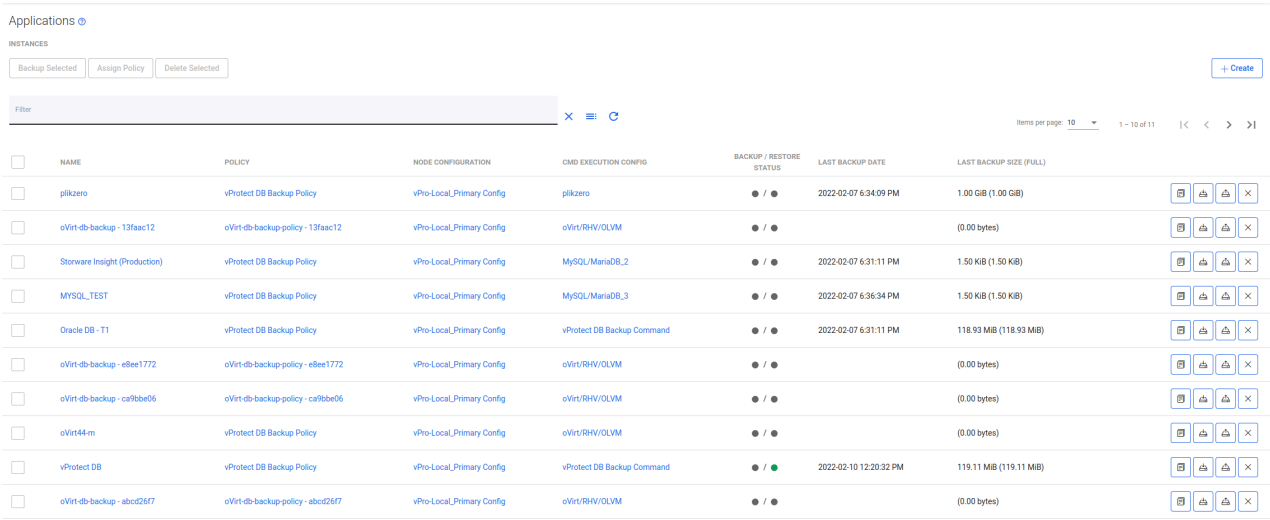

You need to provide its **name** and the **Command Execution Configuration** and select the **node** which is going to do the work - for "node" type command executions.
If the command execution type is "remote ssh", you also need to provide ssh access.
To schedule application backup, you need to select the backup policy.

> **Note:** You can find more about policies and schedules in this article: Backup SLAs.

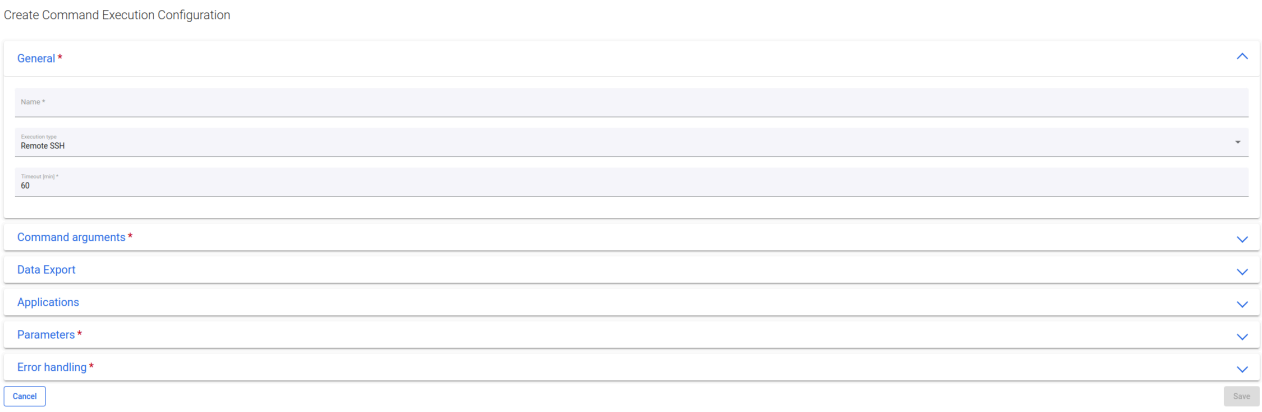

Example - How to restore an application

To perform an on-demand application restore, click the restore button on the right side of the application line.

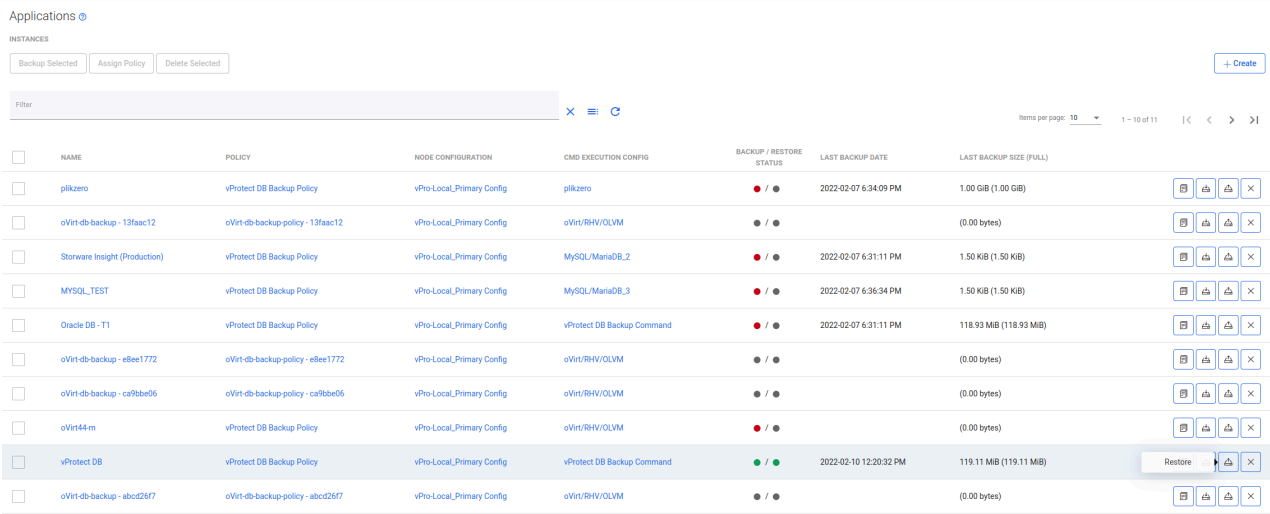

After clicking on it, you will see a pop-up window where you can customize the restore options.

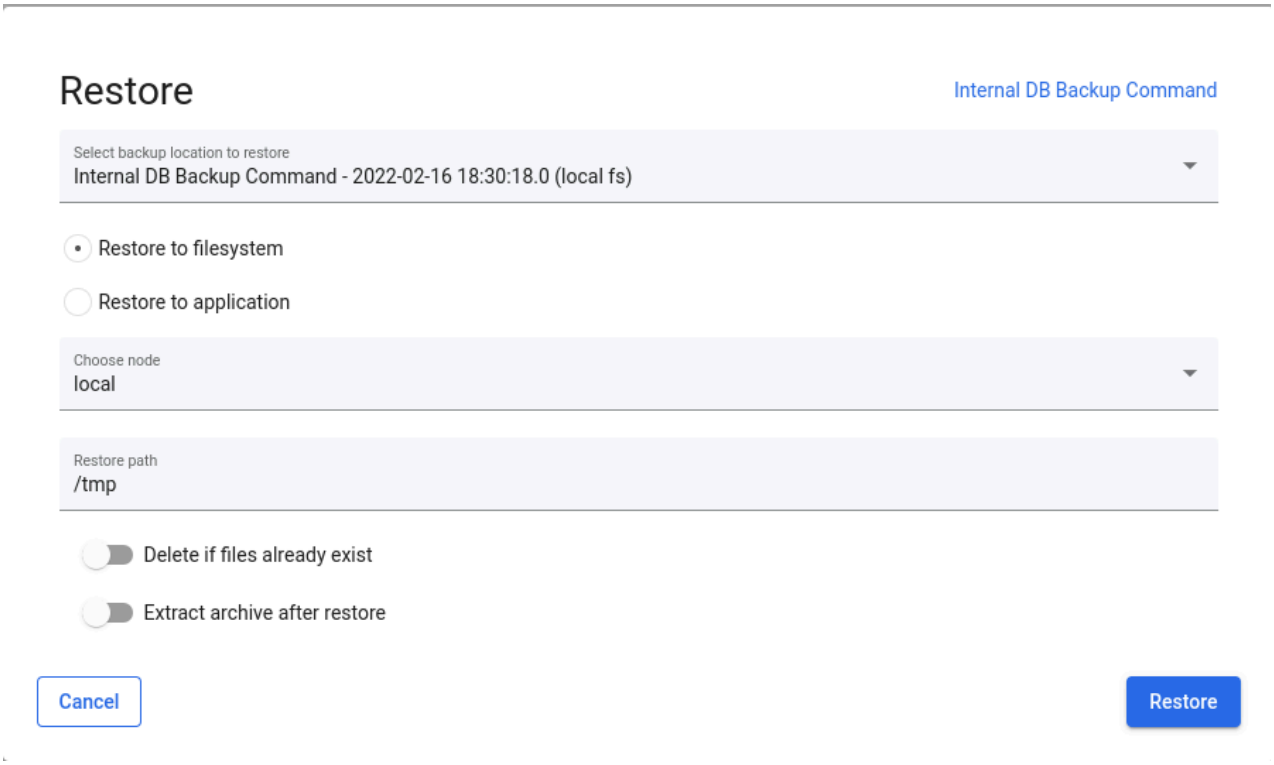

If you choose the "Restore to filesystem" option, you can select things such as which backup you want to restore, on which node you want to restore files, etc.

The "Restore to application" option differs in that it allows you to restore files to the server where the application is located.**

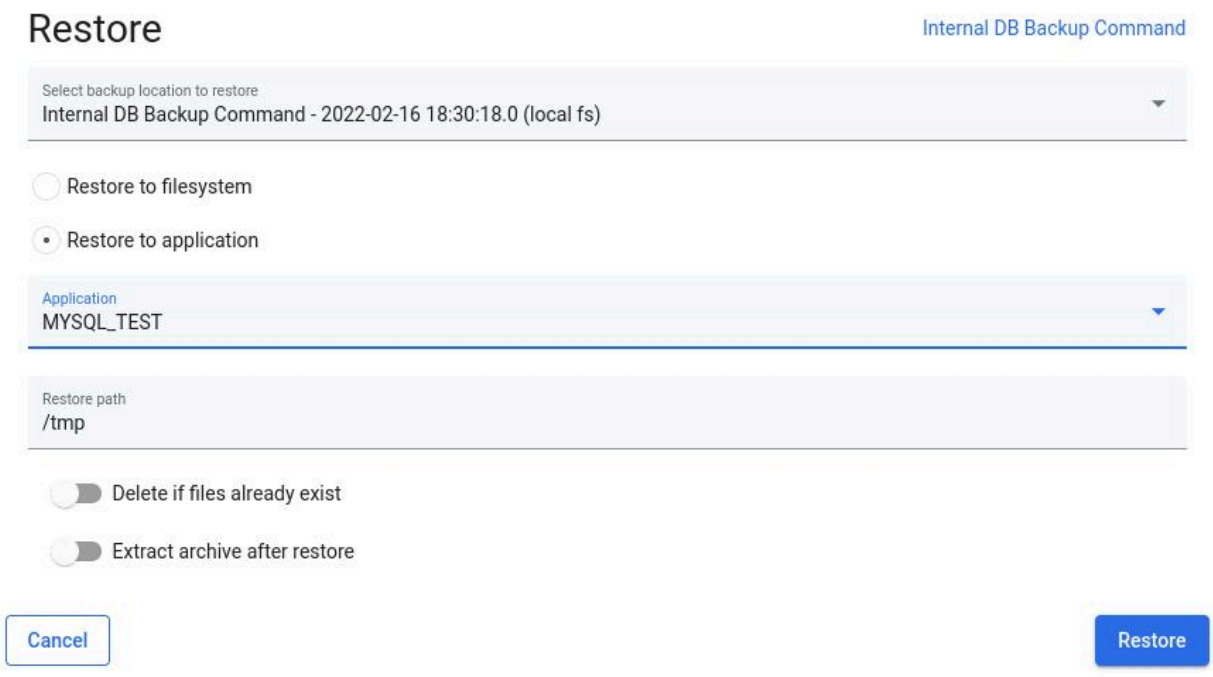

**This option is available for applications using "Remote SSH" in the command execution configuration

Execution Configurations

Execution Configurations

Using this tab, you can create a command execution configuration, for example commands executed for the proper backup of your application (or directed to the script file).

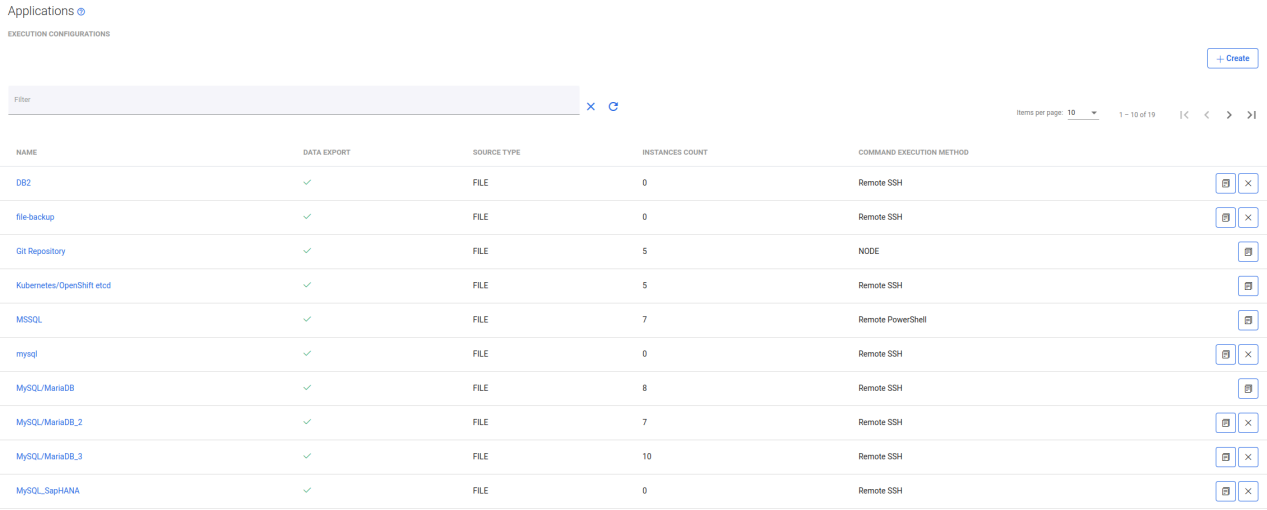

1. Prepare the script or commands, and (if remote SSH execution is required) put them on the remote machine with your application.
2. In the Applications section, create a new **Command Execution Configuration.**

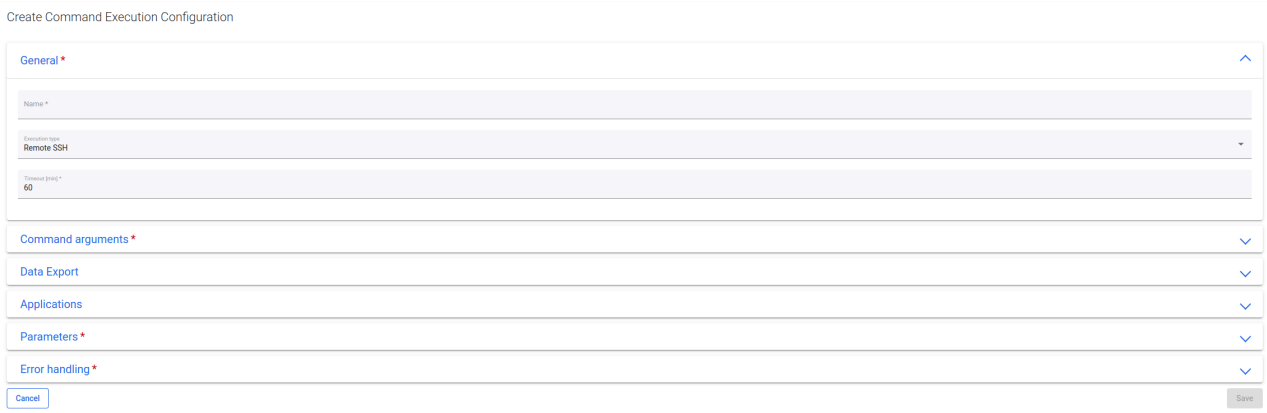

This is a description of how your script is going to be invoked - you need to specify:

- **Execution Type -** either on the Node or a command to be executed remotely.

- **Timeout** - for script/command execution - make sure to have this timeout set to a value high enough to allow your command to complete its execution.
- **Command arguments** (the first one is the actual executable) - all arguments should be provided separately - Data Protector for Cloud Workloads assumes that they can contain spaces.
- **Data export** - this switch allows you to export (or not) data that will be generated by the command (sometimes you may want Data Protector for Cloud Workloads just to execute commands without transferring anything).
- **Source type** (if you export data)
    - **FILE** - select if your command produces a single file/directory (or multiple files) - if you provide a directory or path with wildcards, Data Protector for Cloud Workloads will create a TAR archive - you have to specify the **Source Path** for this source type.
    - **STREAM** - select if your command generates backup to the standard output.
- Select existing **Applications** (or skip this for now - it can be assigned later).
- You can define the **Parameters** that you'll later be able to use in your commands/scripts. Note that the variable name should not contain white-space characters.**
- **Error handling** - you can decide what do you want to do with errors in your output stream

Backup SLAs

Backup SLAs

To backup your application periodically:

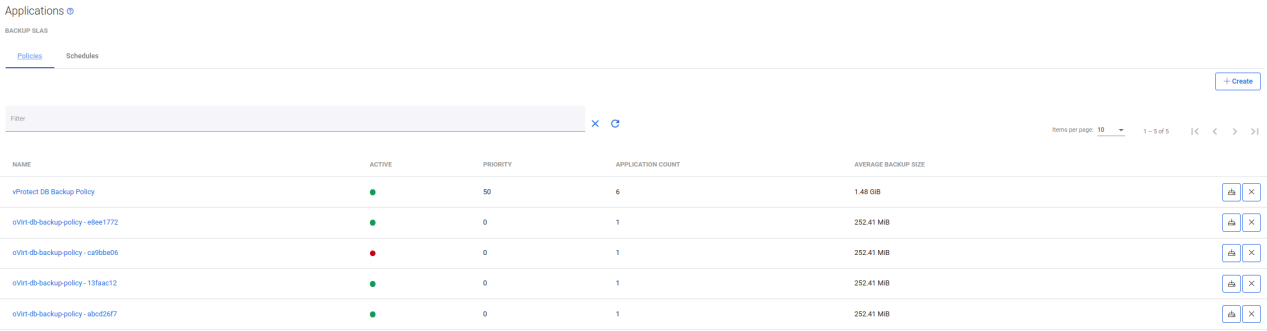

- You need to create a Policy:
  - Go to Applications from the left side menu and then to Backup SLAs
  - Provide a policy name
  - Select your application from the list
  - Specify backup rule details - especially backup destination and schedules

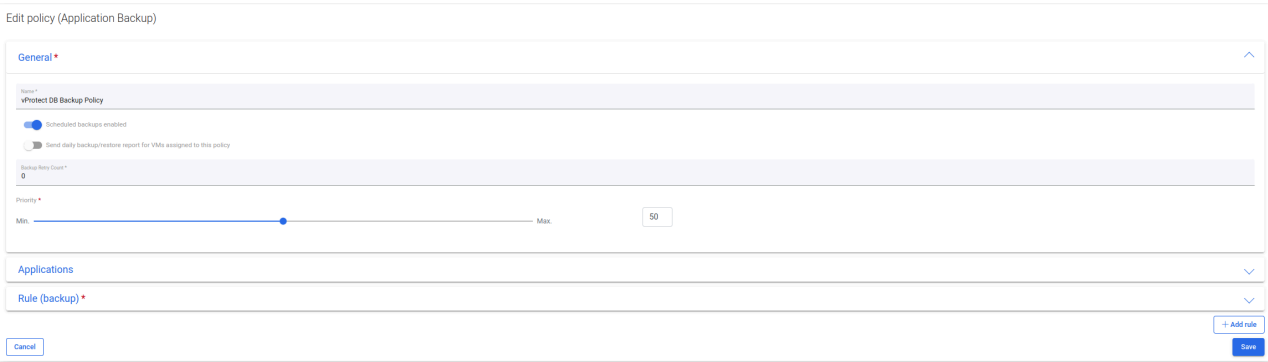

- The second part is a schedule for policy:
  - Change tab from Policies to Schedules
  - Create a new schedule for application policy - create it just like other schedules, enter a name, choose execution time (time or interval) and days of the week. *optionally you can select policy if already exists

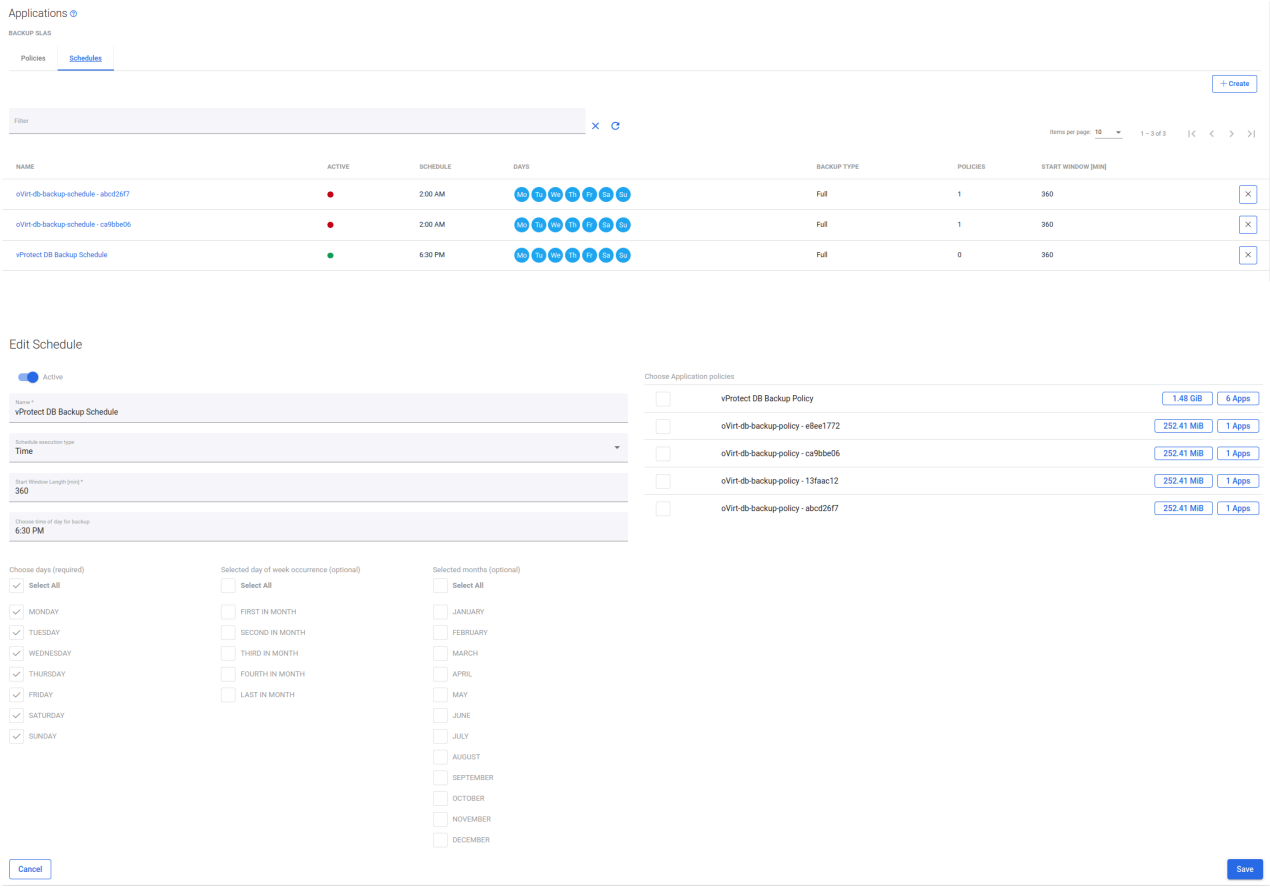

Now your application backups will be done periodically according to your policy.

Reporting

Reporting

Reporting features allow users to view statistics, especially for backup and restore tasks. They also provide the possibility to view what has happened lately in the Data Protector for Cloud Workloads environment.

Reporting is divided into five sections:

- Virtual Environment - for general statistics about data from virtual machines
- Storage - for general statistics about data from storage providers
- Cloud - for general statistics about data from Cloud (Microsoft 365)
- Applications - for general statistics about data from applications
- Audit logs - for important events in terms of security and system administration.

Detailed information on these individual sections can be found directly in the related articles.

Virtual Environments

Virtual Environments

General

Data Protector for Cloud Workloads provides a reporting feature for obtaining information about historical activities. You can reach this from the main dashboard under the Activity section via the **Reporting** button.

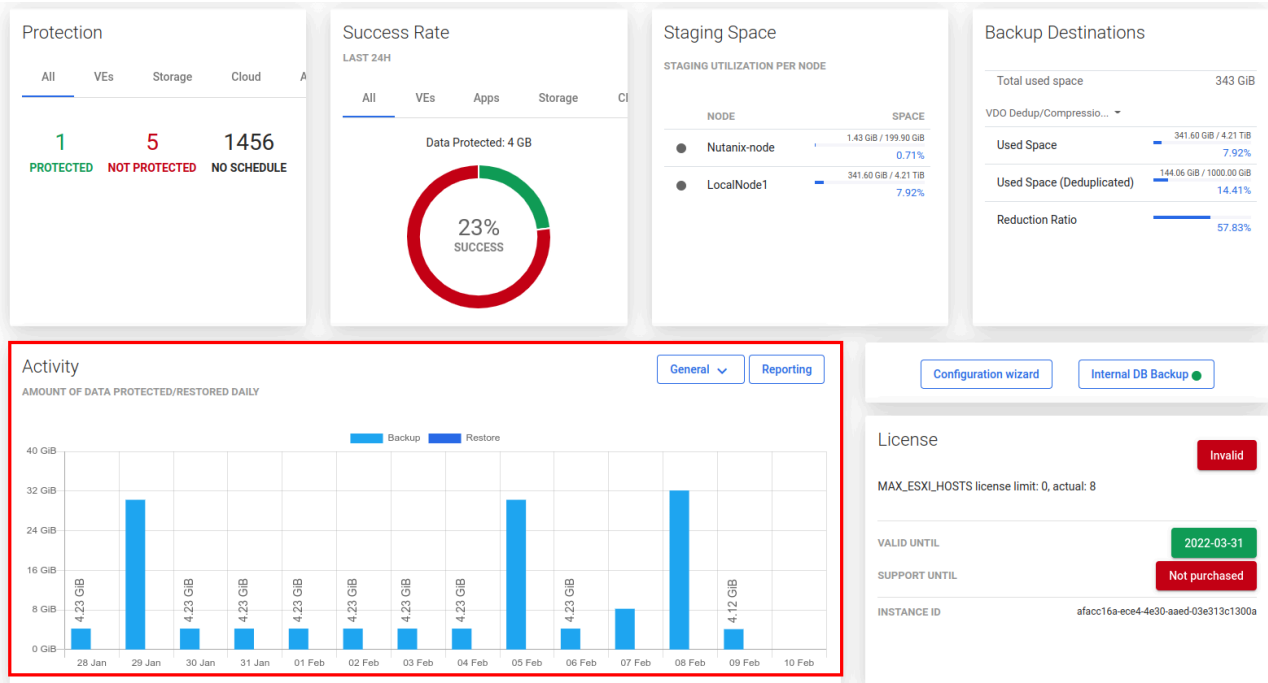

Or on the left side menu - **Reporting** → **Virtual Environments**.

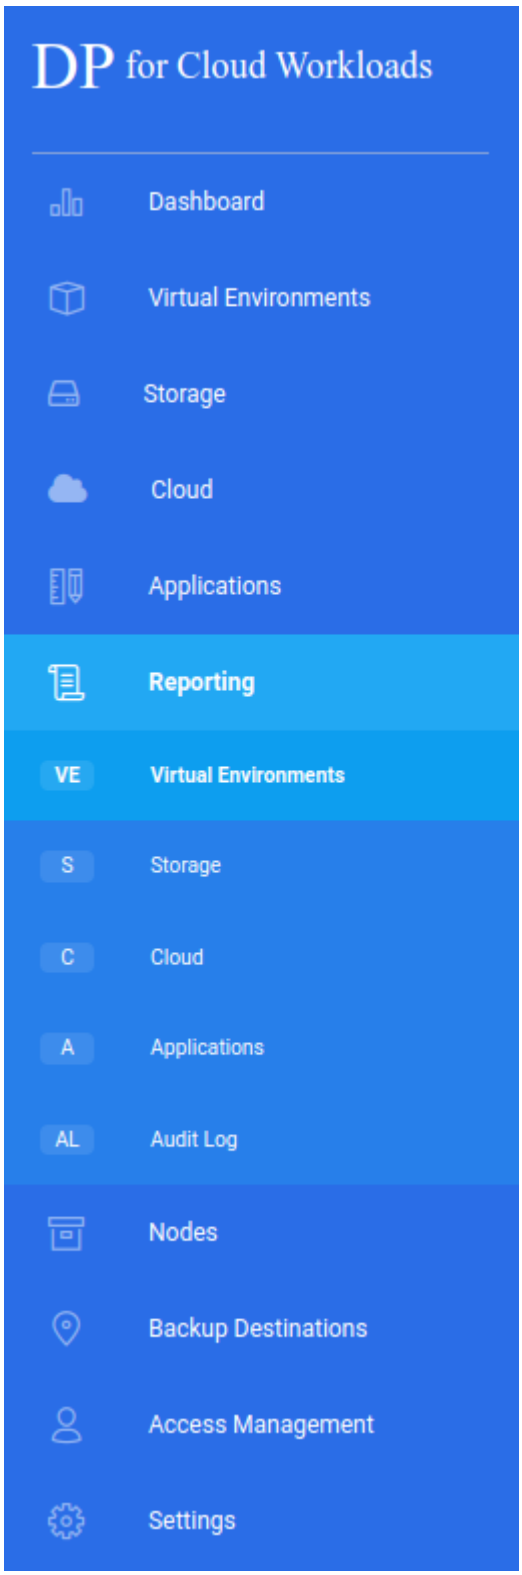

Summary

At the beginning, you should see the summary page.

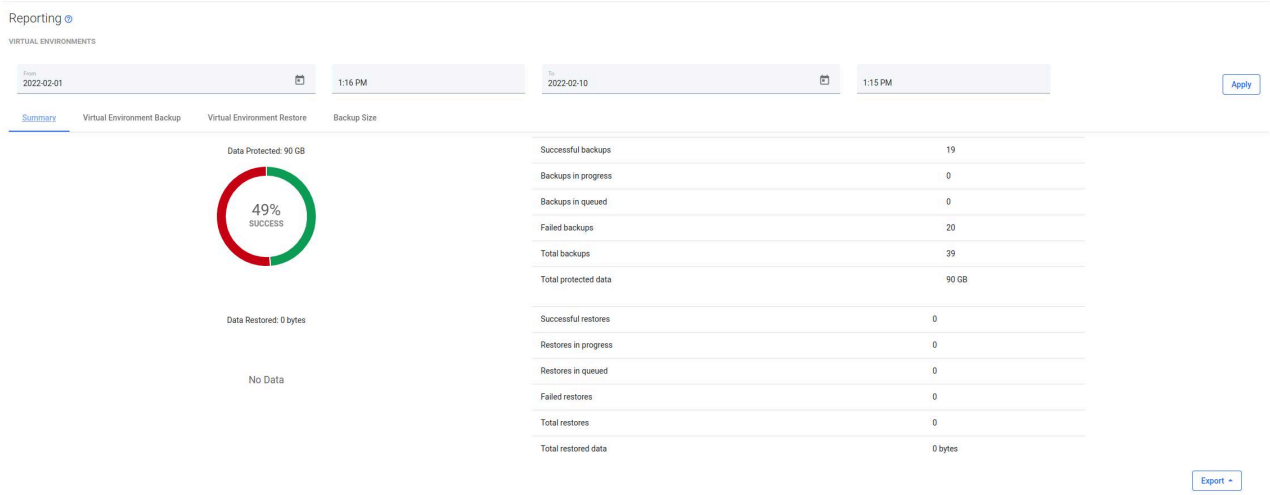

From here, you can set a date range or go to a detailed summary of backup, restore jobs, and backup size.

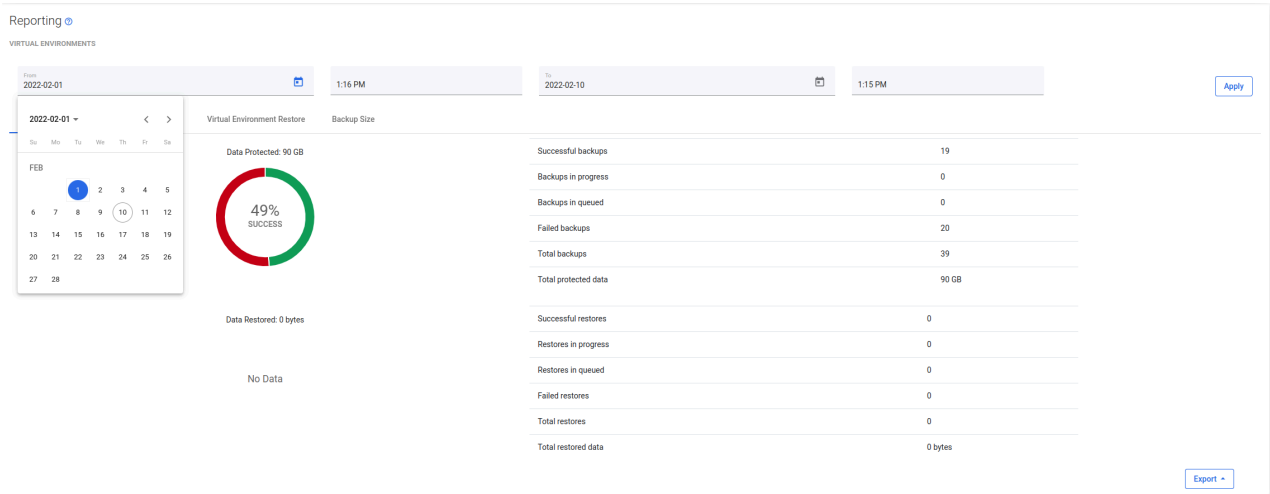

Now that you have set the date you are interested in, you can move to the next tabs to view details of each task performed by Data Protector for Cloud Workloads.

Backup Statistics

Now you can see all the tasks completed within the set date, or ...

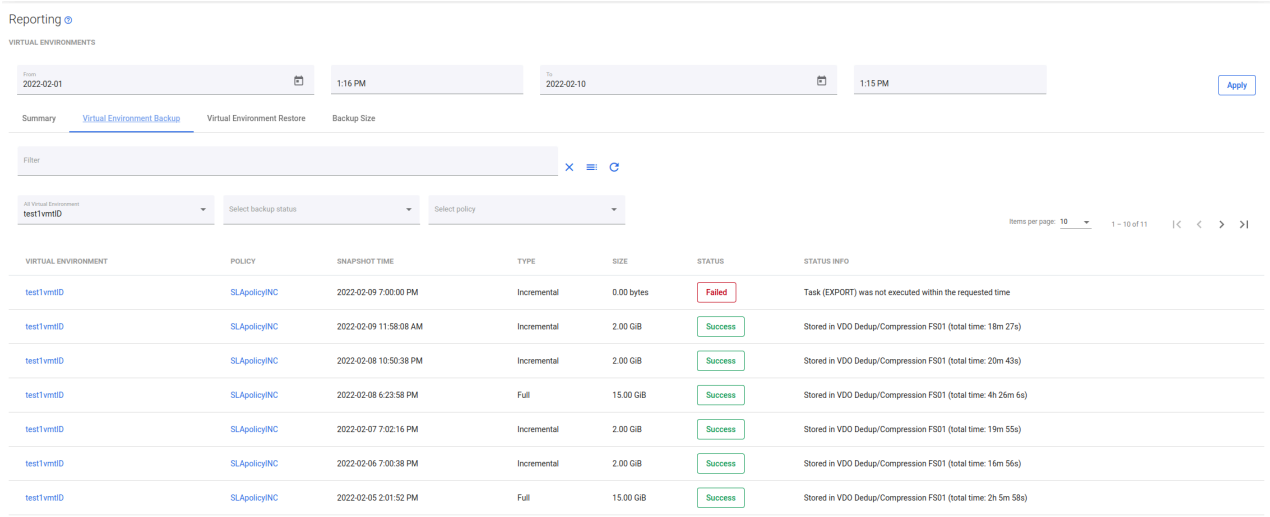

additionally filter them by status, rules or instance.

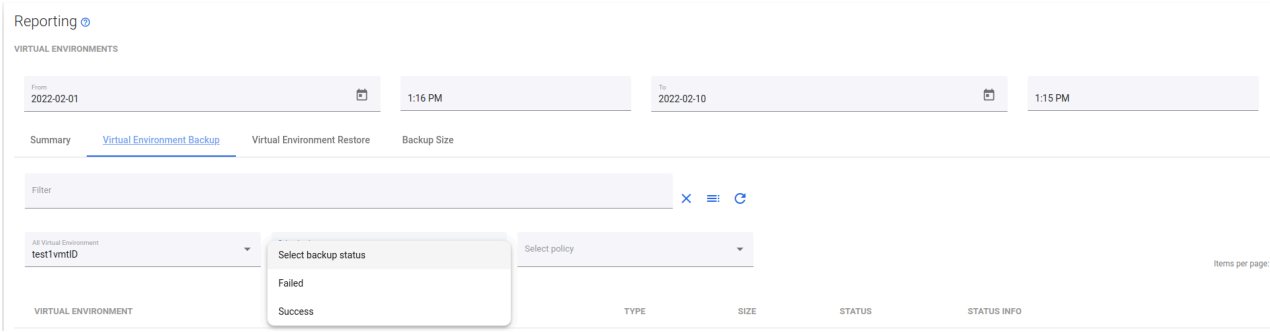

Restore Statistics

The functionality is the same as for the backup statistics, you can see all the tasks completed within the set deadline or filter them by status, policy, or instance.

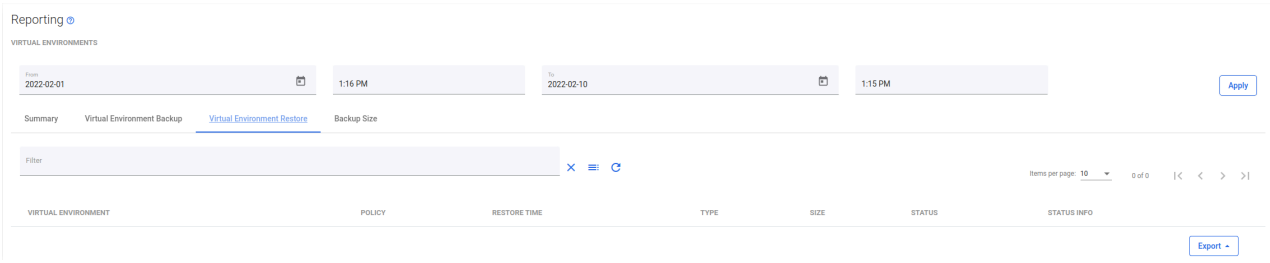

Report sections that you might need for chargeback reporting:

Transfer Size

Transfer size tab shows the amount of data transferred during the backup operations.

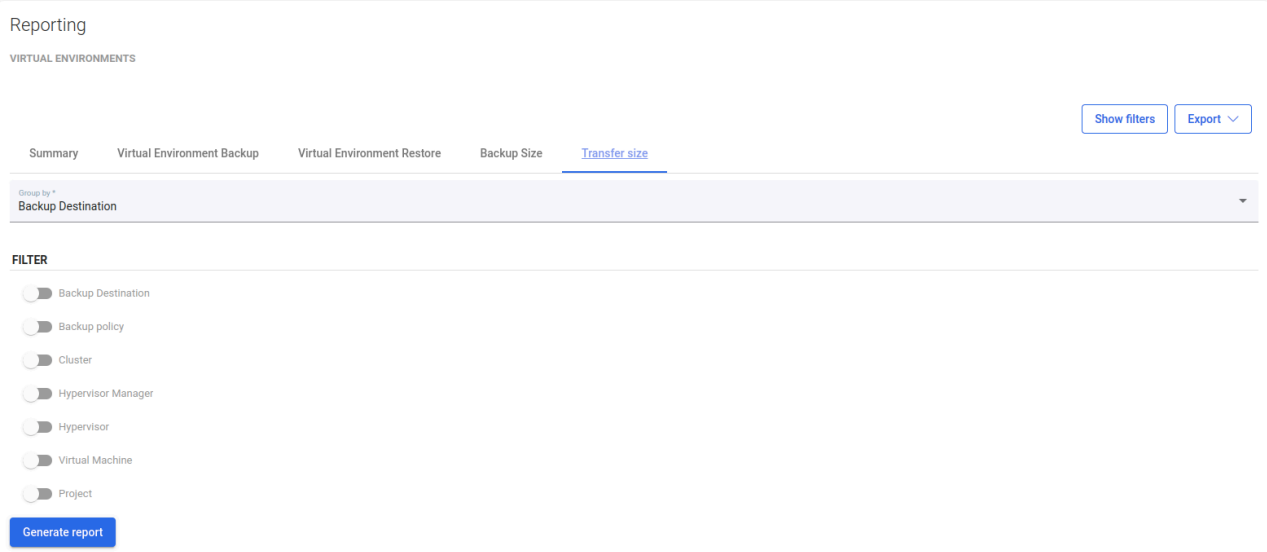

Backup Size

Backup Size section shows the current exported data.

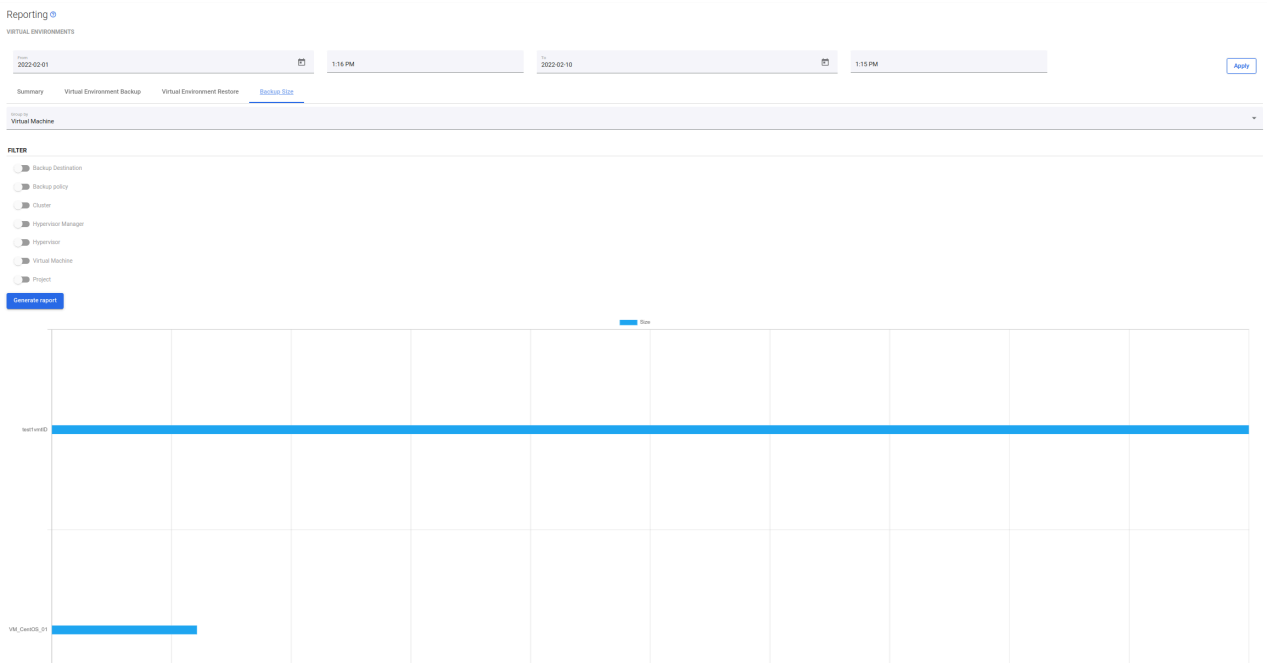

> **Note:** You can use the **Export** button to send reports with backup and restore statistics by email or export them as PDF or HTML.

Storage

Storage

General

Data Protector for Cloud Workloads provides reporting functionality for obtaining information on historical activity.
You can reach this from the main dashboard in the Activity section via the button `Reporting`.

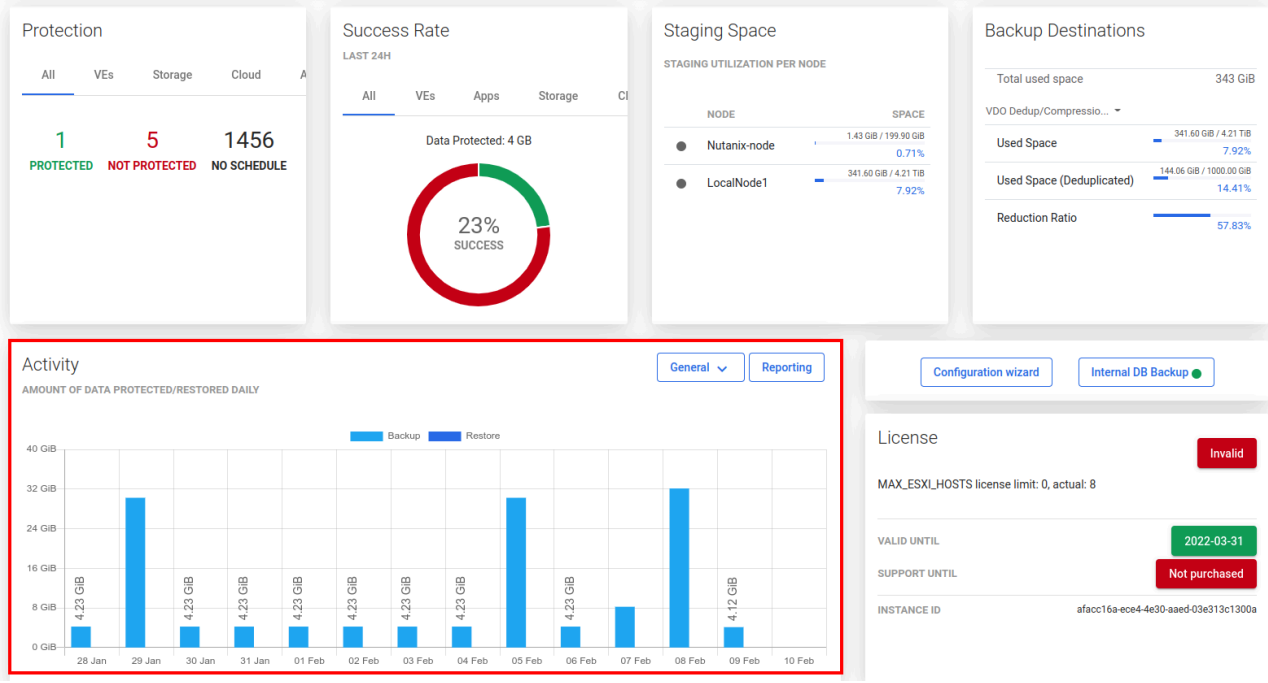

You can also open the storage reports in the reporting section using the left menu.

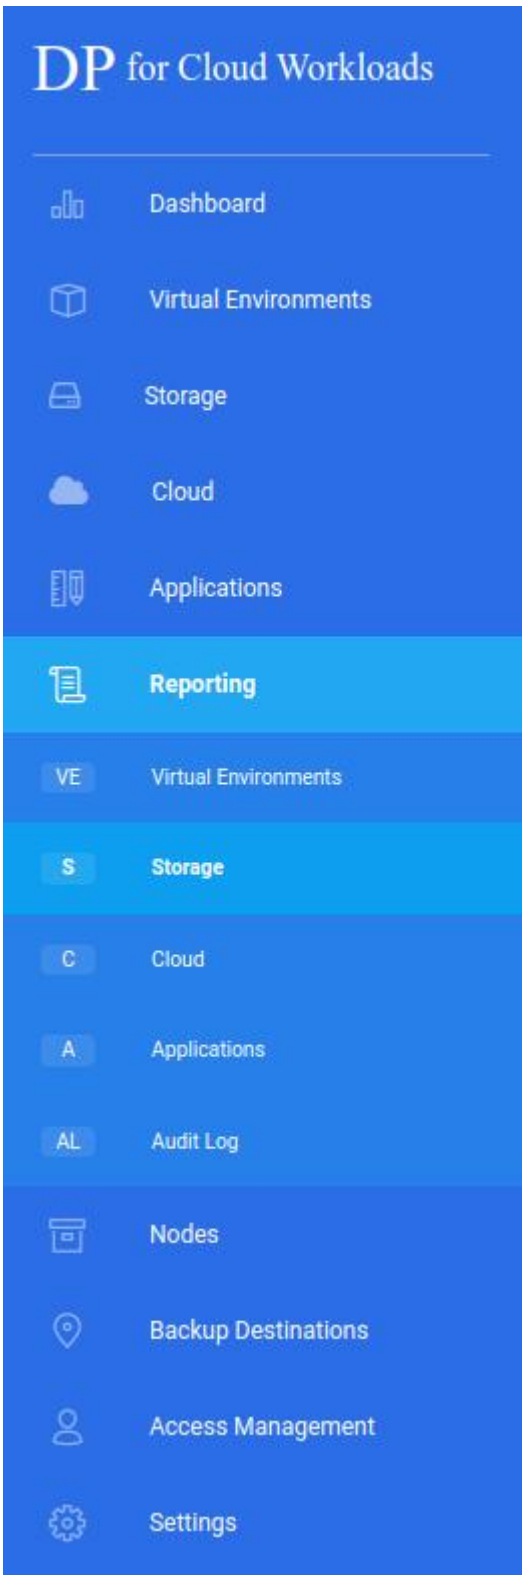

Summary

In the beginning, you should see the summary page.

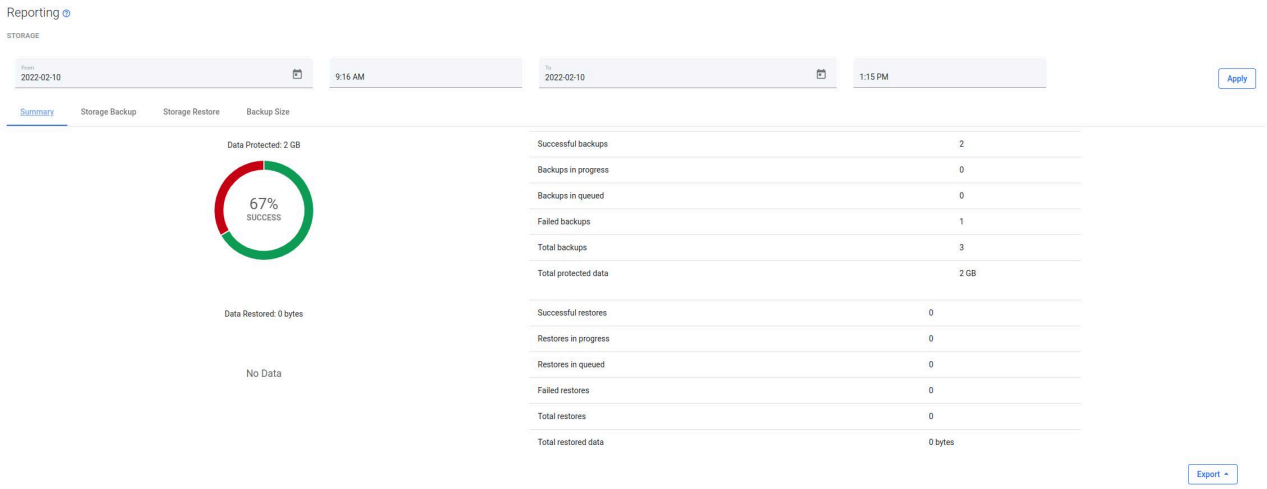

From here, you can set a date range or go to a detailed summary of backup, restore jobs and backup size.

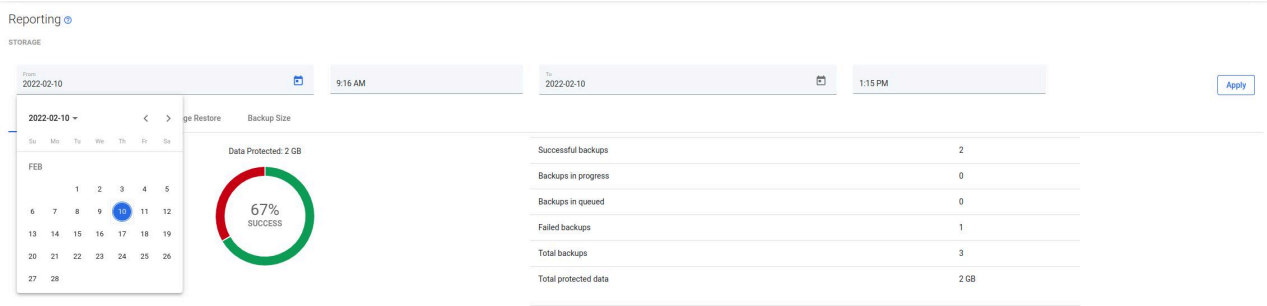

Now that you have set the date you are interested in, you can move to the next tabs to view the details of each task performed by Data Protector for Cloud Workloads.

Backup Statistics

Now you can see all the tasks completed within the set date, or ...

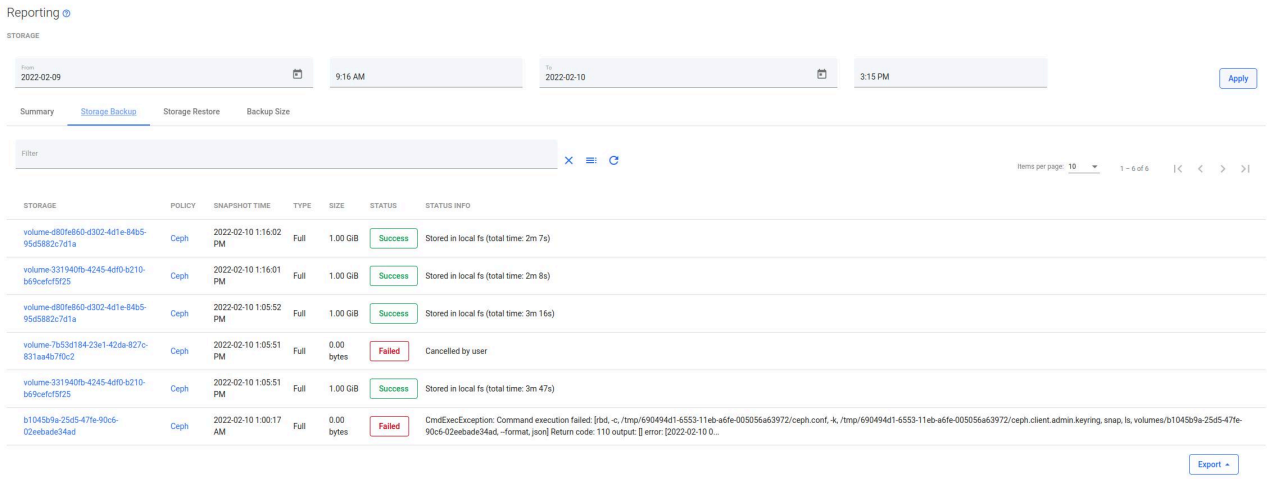

additionally filter them by status, rules or instance.

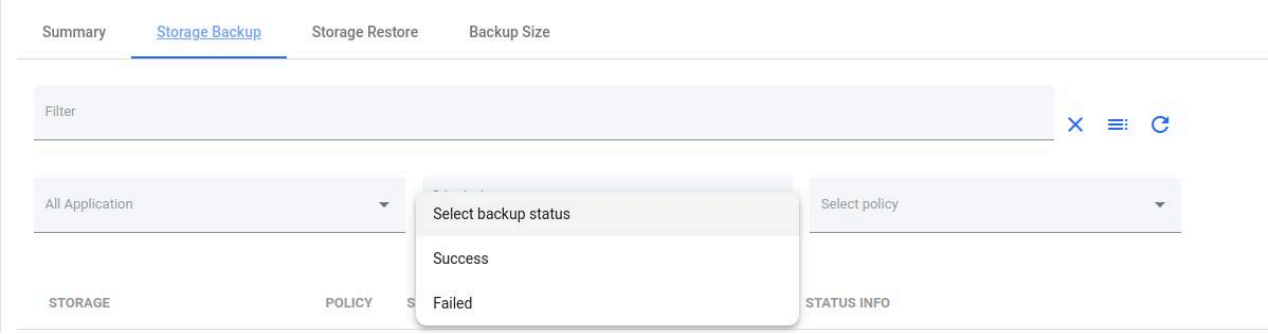

Restore Statistics

The functionality is the same as for the backup statistics, you can see all the tasks completed within the set deadline, or filter them by status, policy or instance.

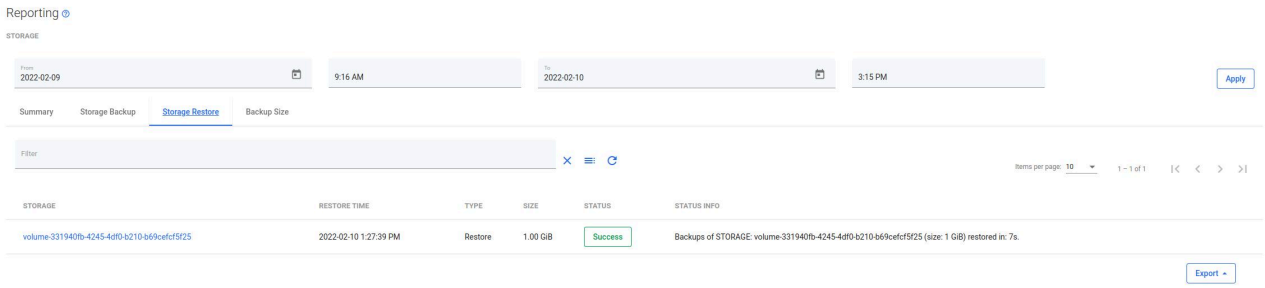

Report sections that you might need for chargeback reporting:

Transfer Size

Transfer size tab shows the amount of data transferred during the backup operations.

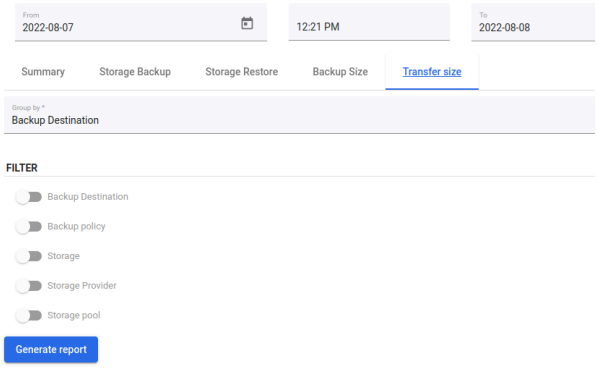

Backup Size

Backup Size section shows the current exported data.

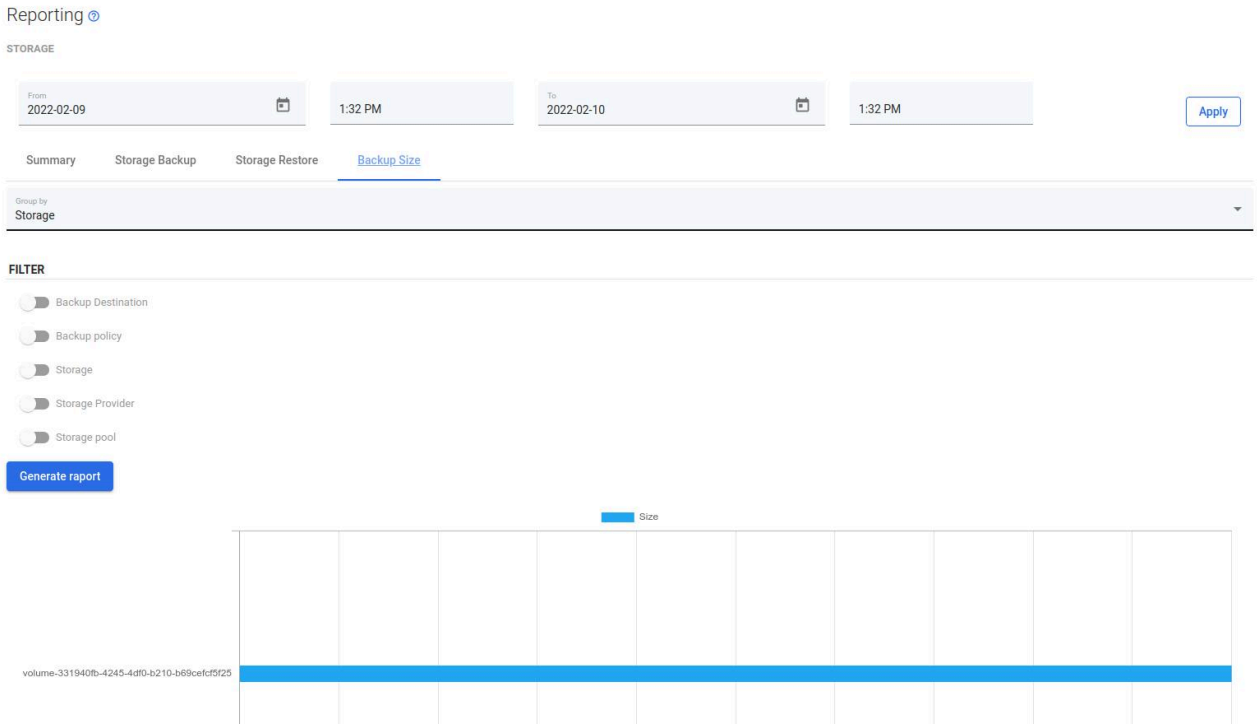

> **Note:** You can use the **Export** button to send reports with backup and restore statistics by email or export them as PDF or HTML.

Cloud

Cloud

General

Data Protector for Cloud Workloads provides a reporting feature for obtaining information about historical activities. You can reach this from the main dashboard under the Activity section via the **Reporting** button.

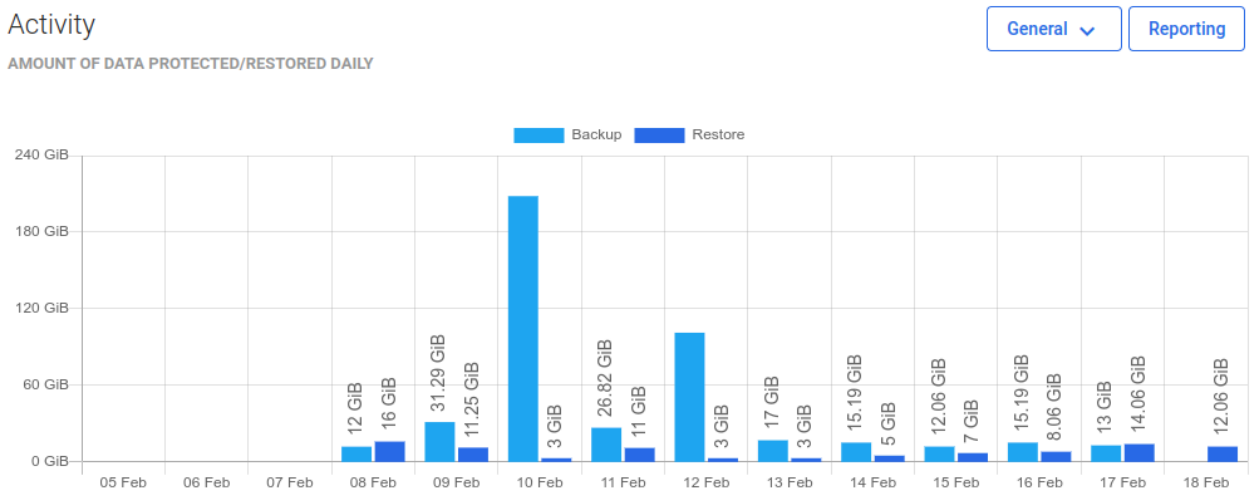

Or on the left side menu - **Reporting** → **Cloud**.

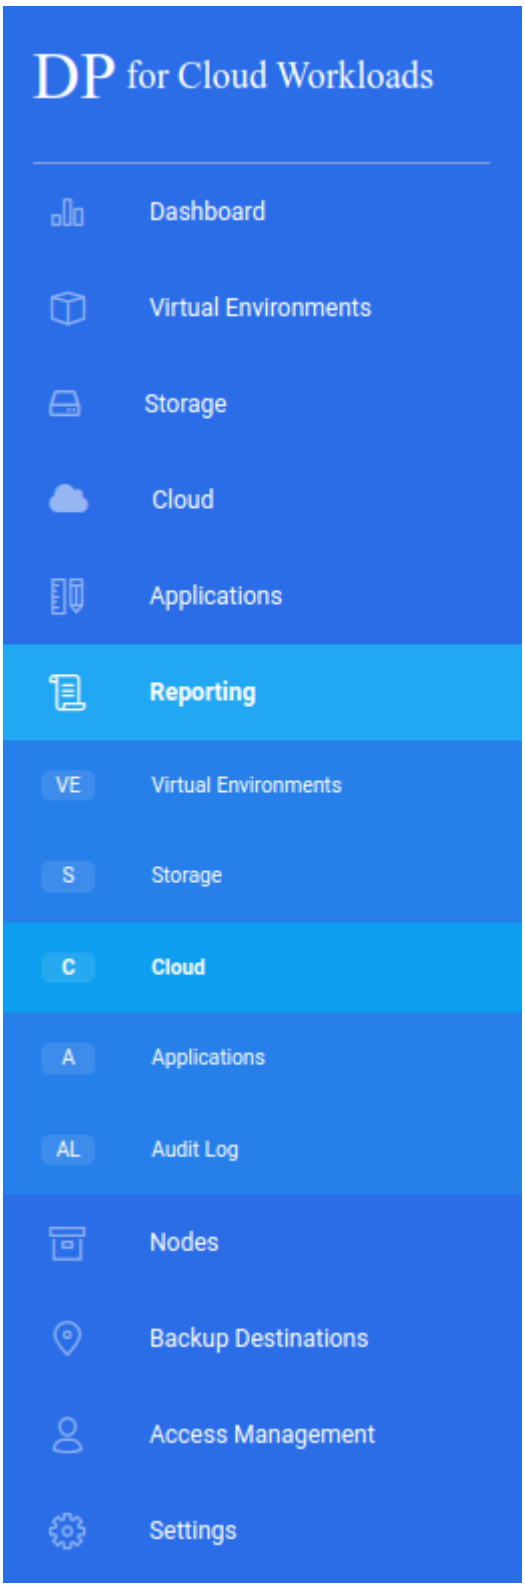

Summary

In the beginning, you should see the summary page.

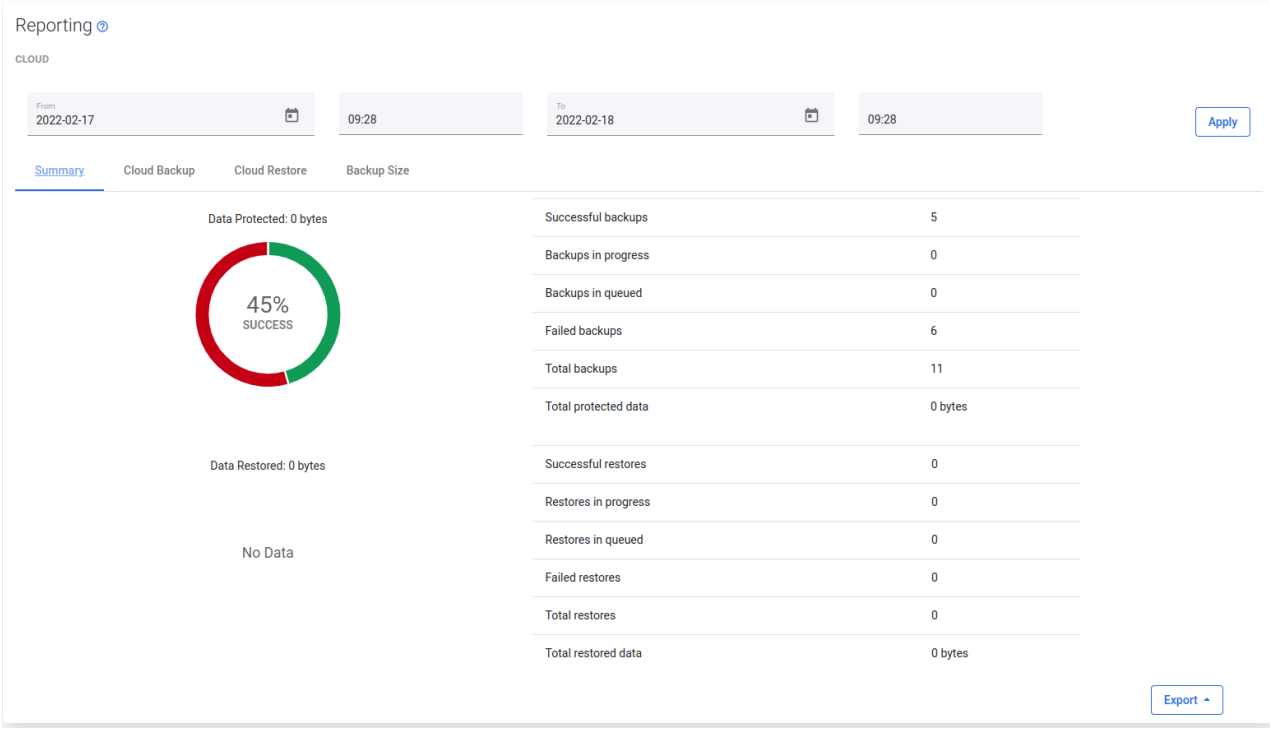

From here, you can set a date range or go to a detailed summary of backup, restore jobs, and backup size.

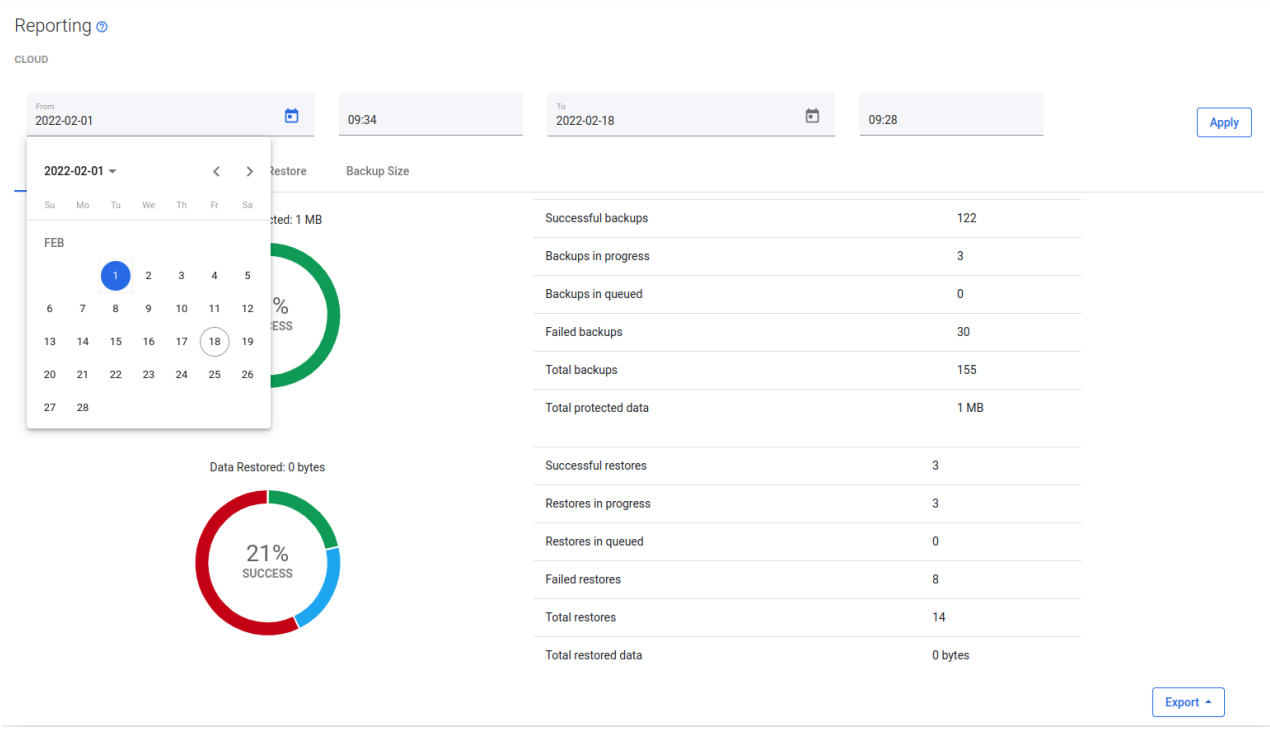

Now that you have set the date you are interested in, you can move to the next tabs to view details of each task performed by Data Protector for Cloud Workloads.

Backup Statistics

Now you can see all the tasks completed within the set date, or ...

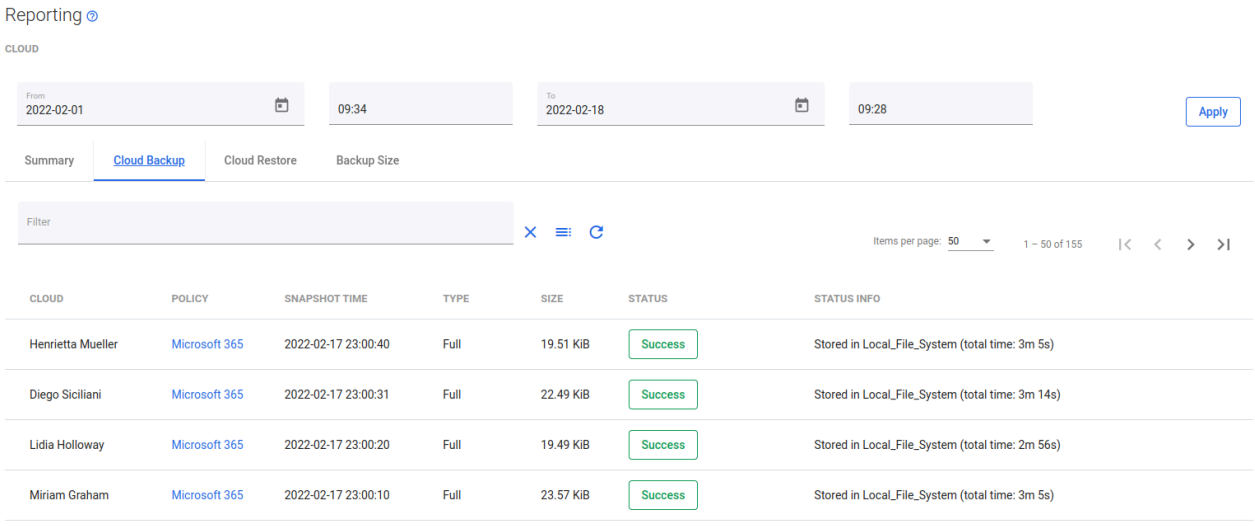

additionally filter them by status, rules, or instance.

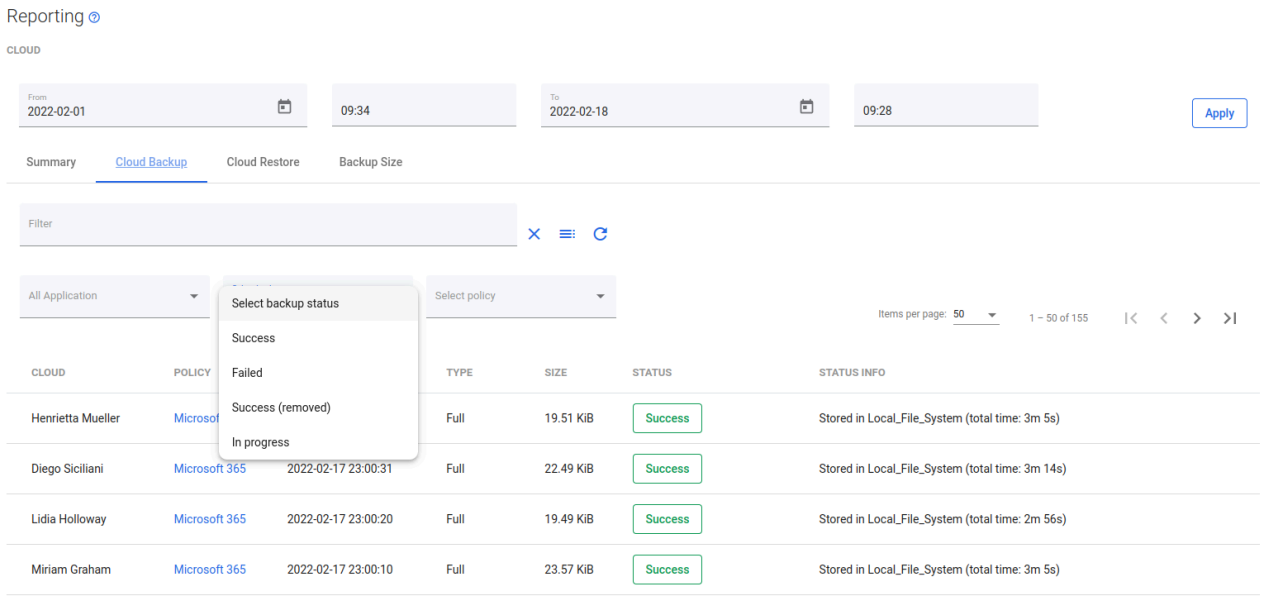

Restore Statistics

The functionality is the same as for the backup statistics, you can see all the tasks completed within the set deadline or filter them by status, policy, or instance.

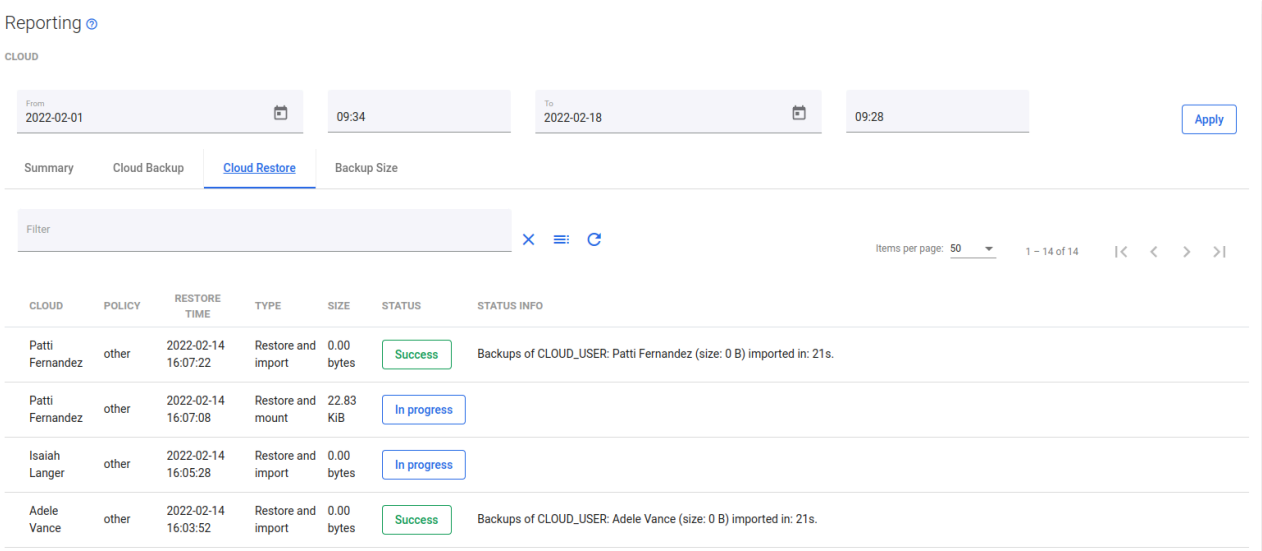

Report sections that you might need for chargeback reporting:

Transfer Size

Transfer size tab shows the amount of data transferred during the backup operations.

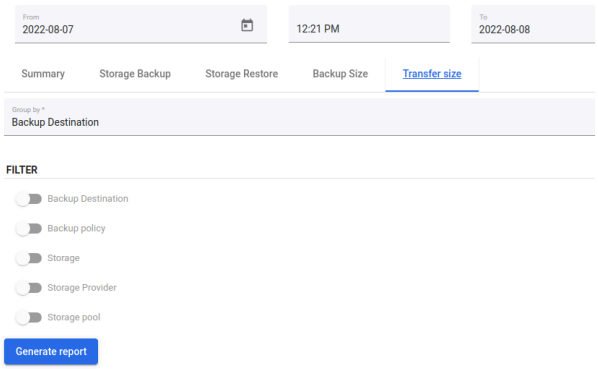

Backup Size

Backup Size section shows the current exported data.

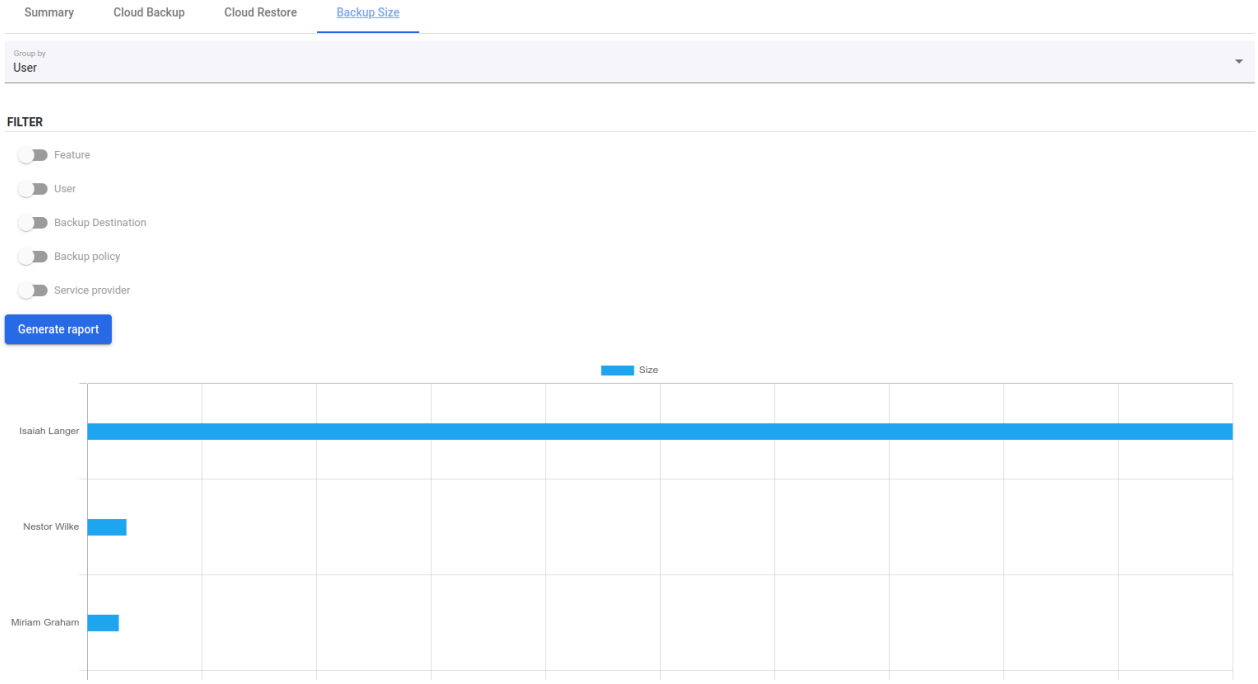

> **Note:** You can use the **Export** button to send reports with backup and restore statistics by email or export them as PDF or HTML.

Applications

Applications

General

Data Protector for Cloud Workloads provides reporting functionality to obtain information on historical activity.
You can reach them from the main dashboard in the "Activity" section with the button **Reporting**.

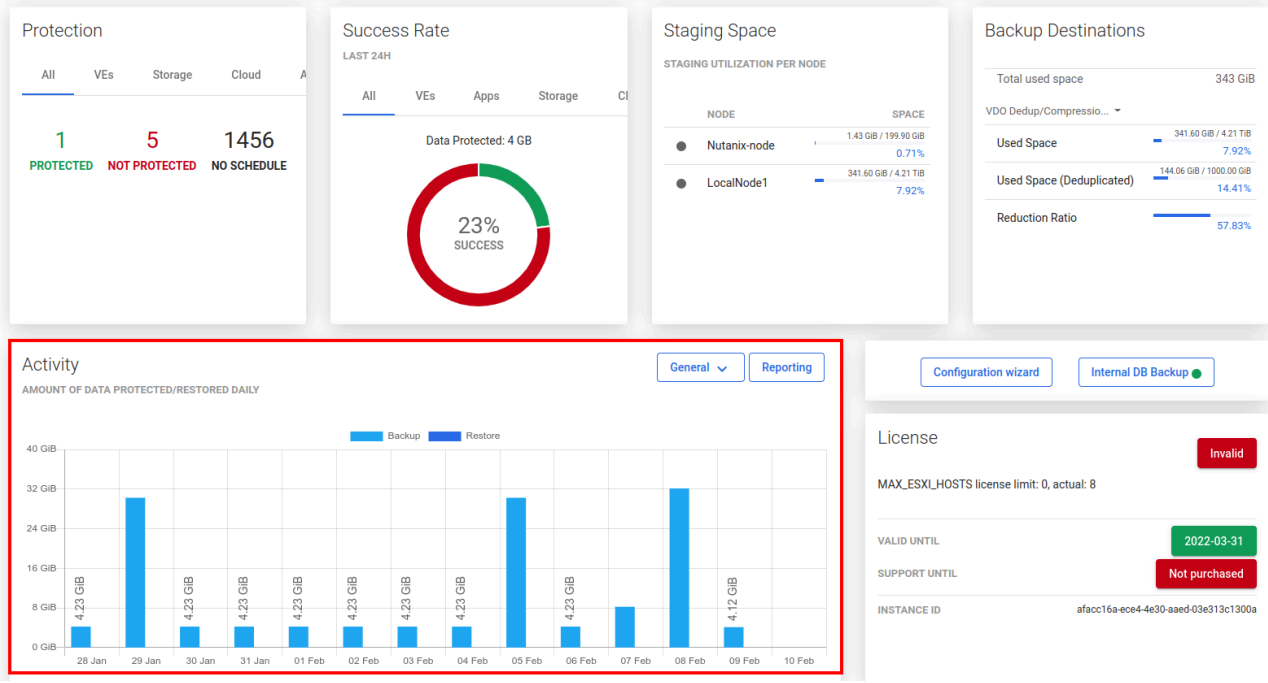

You can also open the Application reports in the reporting section from the left menu.

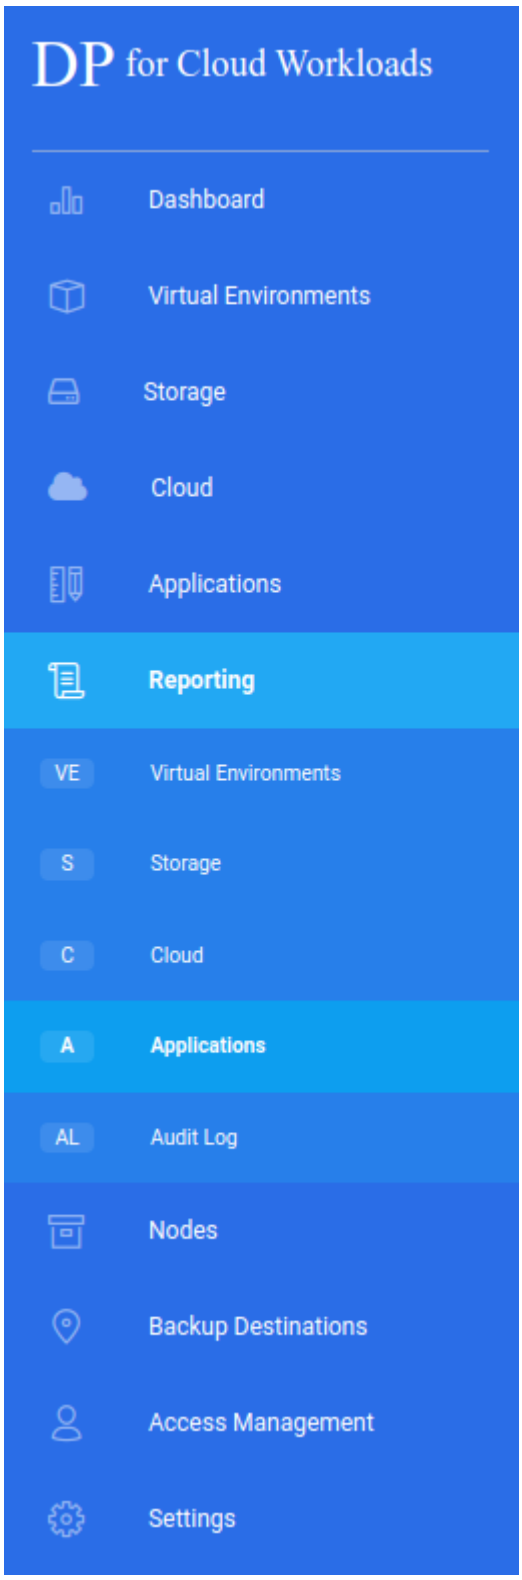

Summary

In the beginning, you should see the summary page

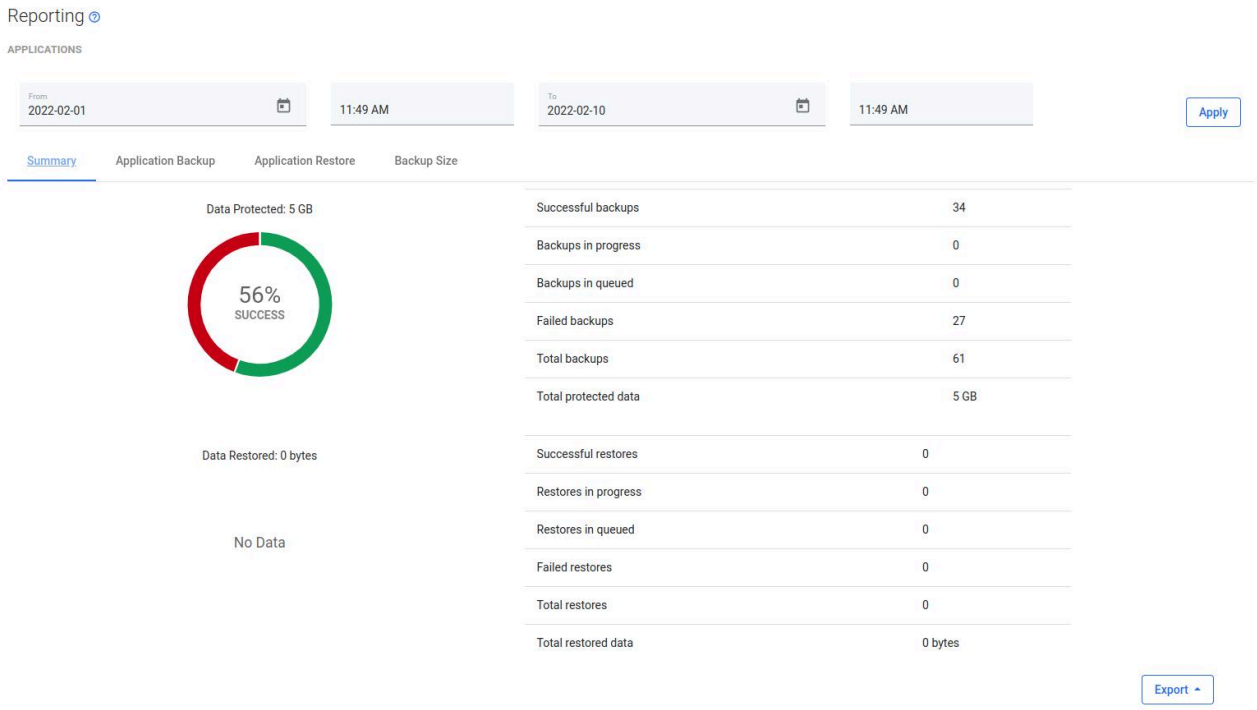

From here, you can set a date range or go to a detailed summary of backup, restore jobs, and backup size.

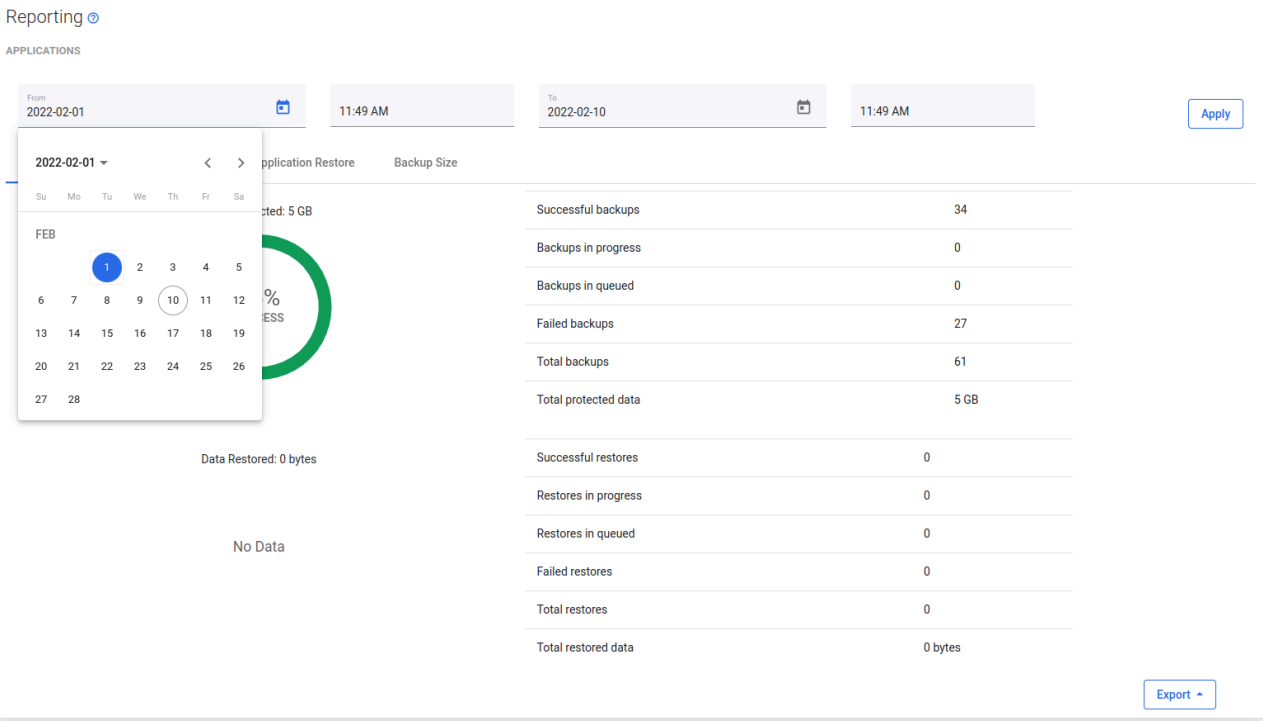

Now that you have set the date you are interested in, you can move to the next tabs to view details of each task performed byData Protector for Cloud Workloads.

Backup Statistics

Now you can see all tasks completed within the set date or ...

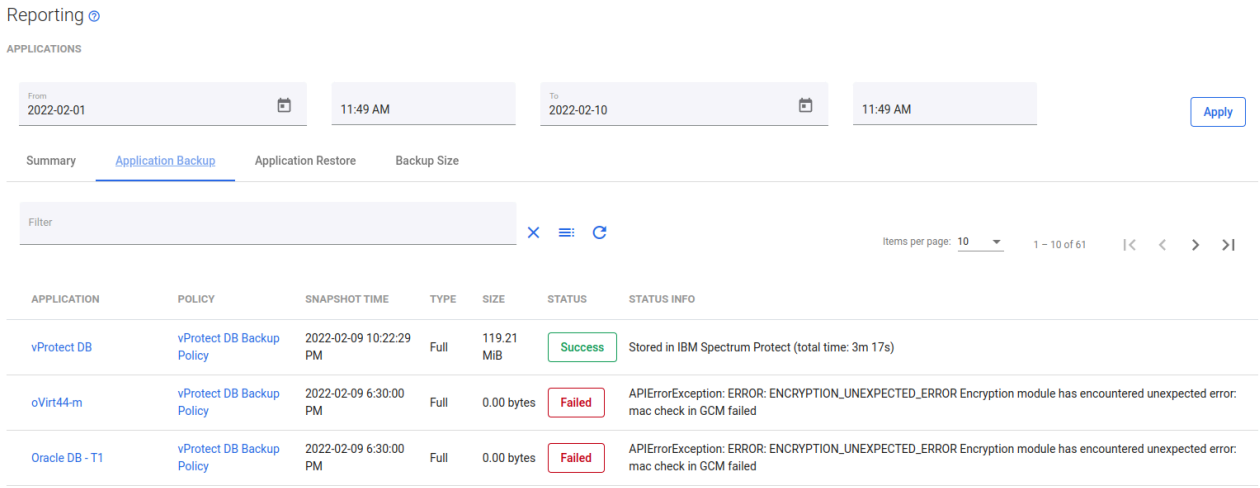

additionally filter them by status, rules, or instance.

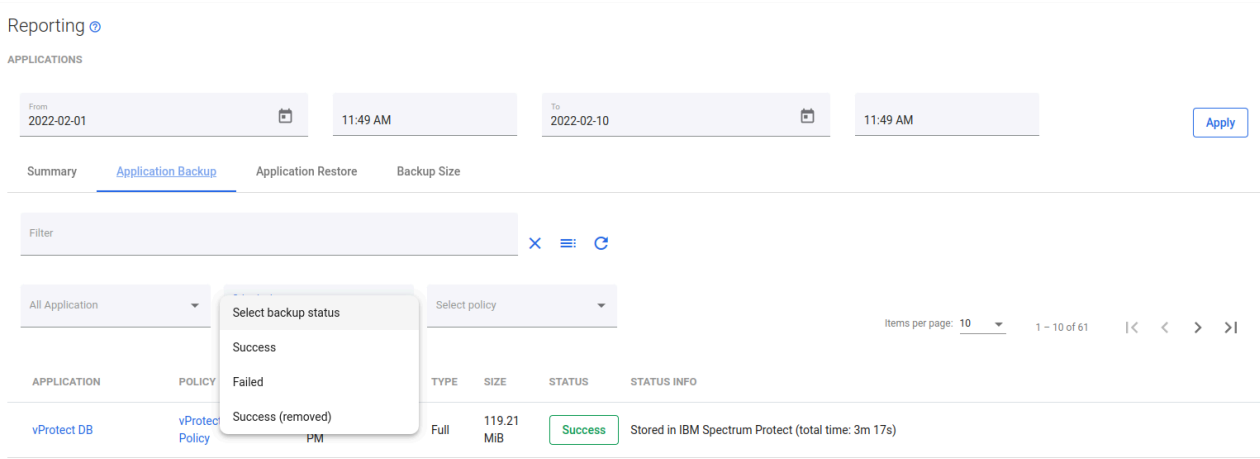

Restore Statistics

The functionality is the same as for the backup statistics, you can see all tasks completed within the set deadline or filter them by status, policy, or instance.

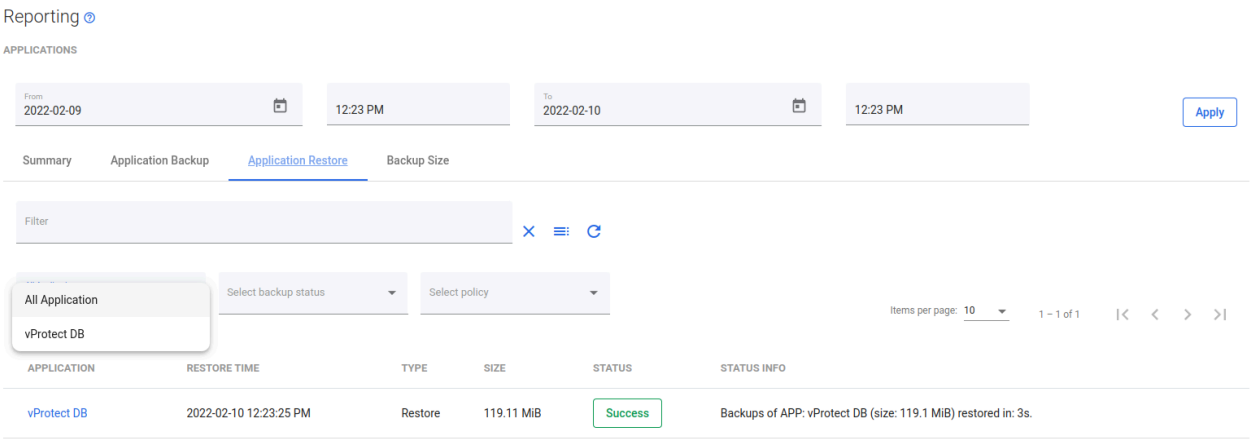

Backup Size

This is a backup size report that you might need for chargeback reporting.

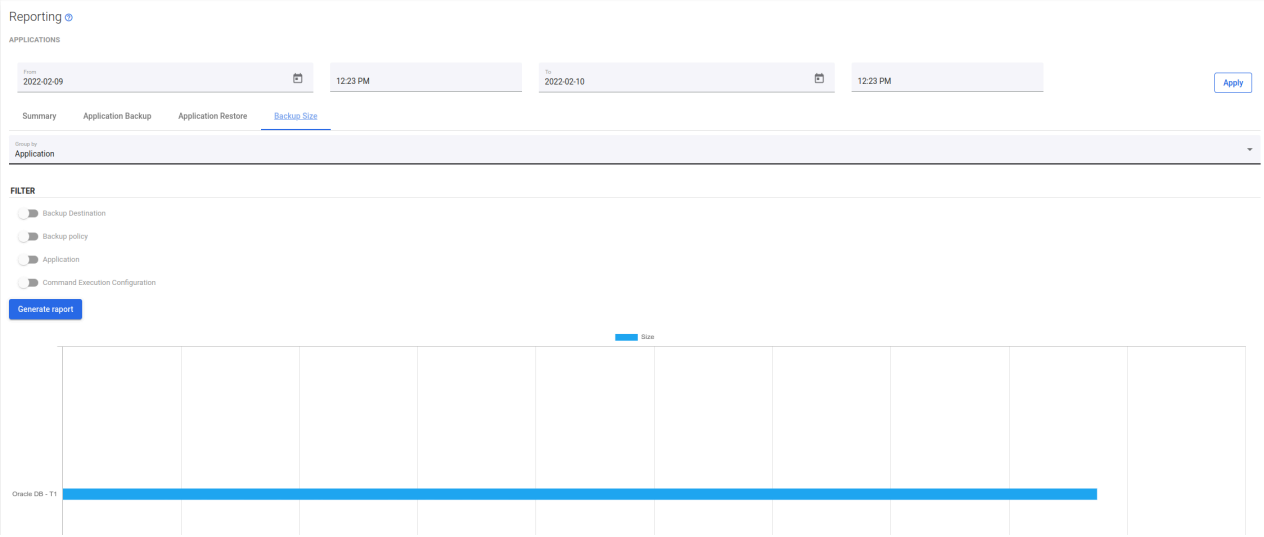

> **Note:** You can use the **Export** button to send reports with backup and restore statistics by email or export them as PDF or HTML.

Audit Log

Audit Log

You can open the Audit log under the reporting section using the left side menu to see what has happened lately in the Data Protector for Cloud Workloads environment.

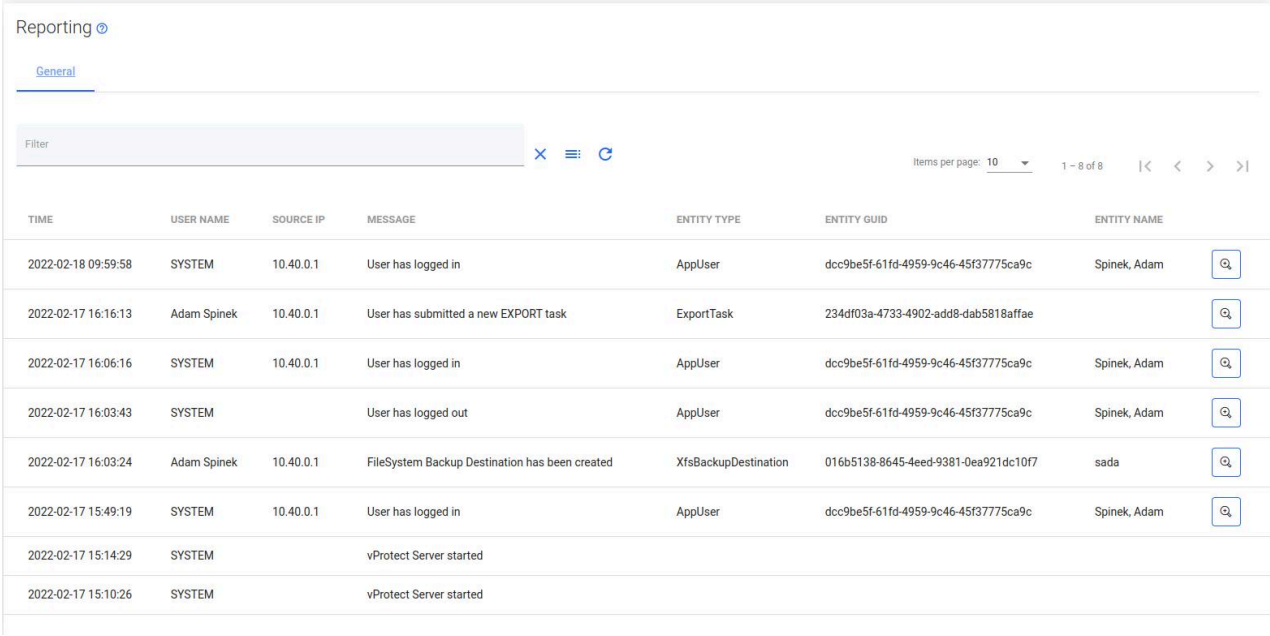

TIME	USER NAME	SOURCE IP	MESSAGE	ENTITY TYPE	ENTITY GUID	ENTITY NAME
2022-02-18 09:59:58	SYSTEM	10.40.0.1	User has logged in	AppUser	dcc9be5f-61fd-4959-9c46-45f37775ca9c	Spinek, Adam
2022-02-17 16:16:13	Adam Spinek	10.40.0.1	User has submitted a new EXPORT task	ExportTask	234df03a-4733-4902-add8-dab5818affae	
2022-02-17 16:06:16	SYSTEM	10.40.0.1	User has logged in	AppUser	dcc9be5f-61fd-4959-9c46-45f37775ca9c	Spinek, Adam
2022-02-17 16:03:43	SYSTEM		User has logged out	AppUser	dcc9be5f-61fd-4959-9c46-45f37775ca9c	Spinek, Adam
2022-02-17 16:03:24	Adam Spinek	10.40.0.1	FileSystem Backup Destination has been created	XfsBackupDestination	016b5138-8645-4eed-9381-0ea921dc10f7	sada
2022-02-17 15:49:19	SYSTEM	10.40.0.1	User has logged in	AppUser	dcc9be5f-61fd-4959-9c46-45f37775ca9c	Spinek, Adam
2022-02-17 15:14:29	SYSTEM		vProtect Server started			
2022-02-17 15:10:26	SYSTEM		vProtect Server started			

As you can see, you can easily check who logged in / out or what task was done.

Thanks to the filter section at the top of the page, you can easily choose a date range or enter a keyword.

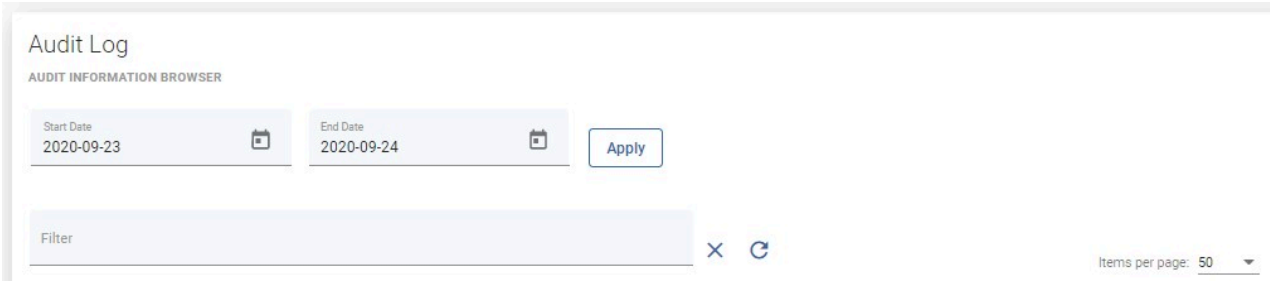

You can also open the task details page using the magnifier icon on the right.

Details


```
guid: "9e96b226-ba7d-44af-9f85-4d3914b766e4"
parentTask: null
type: "EXPORT"
restoreType: null
state: "QUEUED"
statusInfo: null
windowStart: 1644491280000
windowEnd: 1644509280000
progressChange: null
creationTime: 1644491300000
startTime: null
finishTime: null
progress: 0
priority: 50
retryCount: 0
processedObjectCount: null
totalObjectCount: null
▸ protectedEntity: Object {"guid":"ceb68089-266a-4575-9678-f1fcf42dba2c","name":"vProtect DB","typ
protectedEntityDisplayName: "vProtect DB"
dstProtectedEntity: null
protectedEntitySnapshot: null
hypervisor: null
hypervisorManager: null
storageProvider: null
cloudServiceProvider: null
▸ backup: Object {"guid":"e50b8d80-611a-40f3-88f9-5235150748f2","name":"vProtect DB - 2022-02-10 1
```

Nodes

Nodes

In this chapter, You will know how to manage Data Protector for Cloud Workloads Node and how to create Data Protector for Cloud Workloads Node configurations

Instances

Instances

Node Instances

Node instances are a list of the Data Protector for Cloud Workloads Nodes that currently exist in your environment. You can easily check basic information about the nodes or change the assigned node configuration.

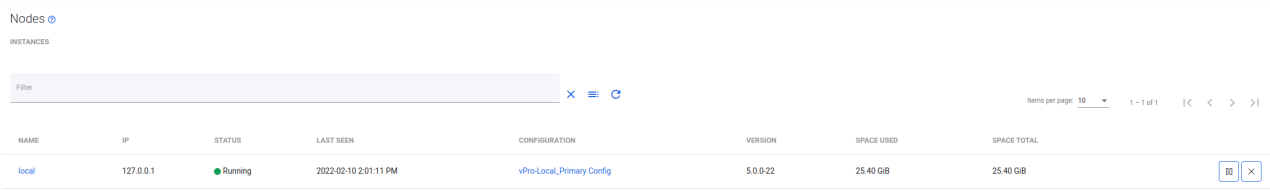

Node Configurations

Node Configurations

Node configurations are groups of settings assignable to a set of nodes. This is so that you don't have to change them on every node separately.

> **Note:**
>
> - The task-to-node assignment will be executed via the task's hypervisor/hypervisor manager/storage provider/application/cloud related to the node configuration where the system will decide on the specific node assigned to the aforementioned configuration responsible for executing the task.

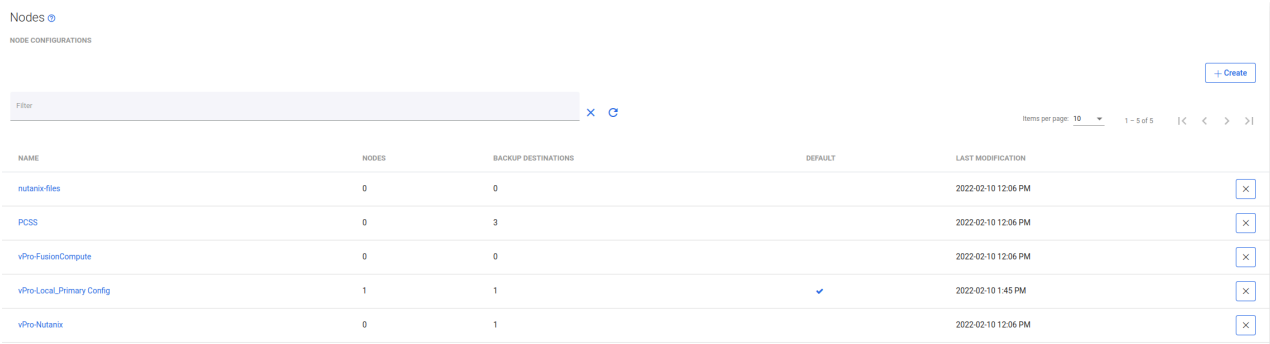

Available settings

General

- `Name` - unique name identifying the configuration
- `Set as default` - set default configuration to be assigned to new nodes

Storage

- `Export path (staging space)` - staging path (must be owned by the `vprotect` user)

- `Restore path for import tasks` - in rare cases, you may want to restore backups to a custom location within the node before the import process begins
- `Restore path for mount tasks` - in rare cases, you may want to restore backups to a custom location within the node before the mounting process begins
- `Min. free space for export [%]` - the amount of storage space left in order to force a node to wait before starting another export task

Backup Process

- `Max. allowed backup timestamp difference [s]`- maximum time difference between the timestamp of a backup in Data Protector for Cloud Workloads and the backup destination in order to match the local version of the backup with the remote one
- `Minimal free space for snapshot [GB]` - the amount of storage space left in order to force a node to wait before starting another task
- `Minimal free space for snapshot [%]` - the amount of storage space left in order to force a node to wait before starting another task
- `Dynamically attached disks slot offset` - this setting forces a shift of a disk slot number that the node reads/writes from when the disk-attachment method is used - currently used in the Nutanix disk-attachment method when you have block devices not reported by the hypervisor API, such as iSCSI mounted block devices. When set to 0, Data Protector for Cloud Workloads will mount drives just after the last occupied (and reported by the hypervisor API) slot (which means that block device number 3 in API will be /dev/sdc in OS). In general, N means that Data Protector for Cloud Workloads will shift N slots, so 1 will make the 3rd device be treated as the 4th in OS /dev/sdd)
- `Netcat min. port` - min. Netcat port range
- `Netcat max. port` - max. Netcat port range
- `Number of netcat transfer attempts` - maximum number of attempts
- `BZIP2 Compression` - compress backup files with bzip2

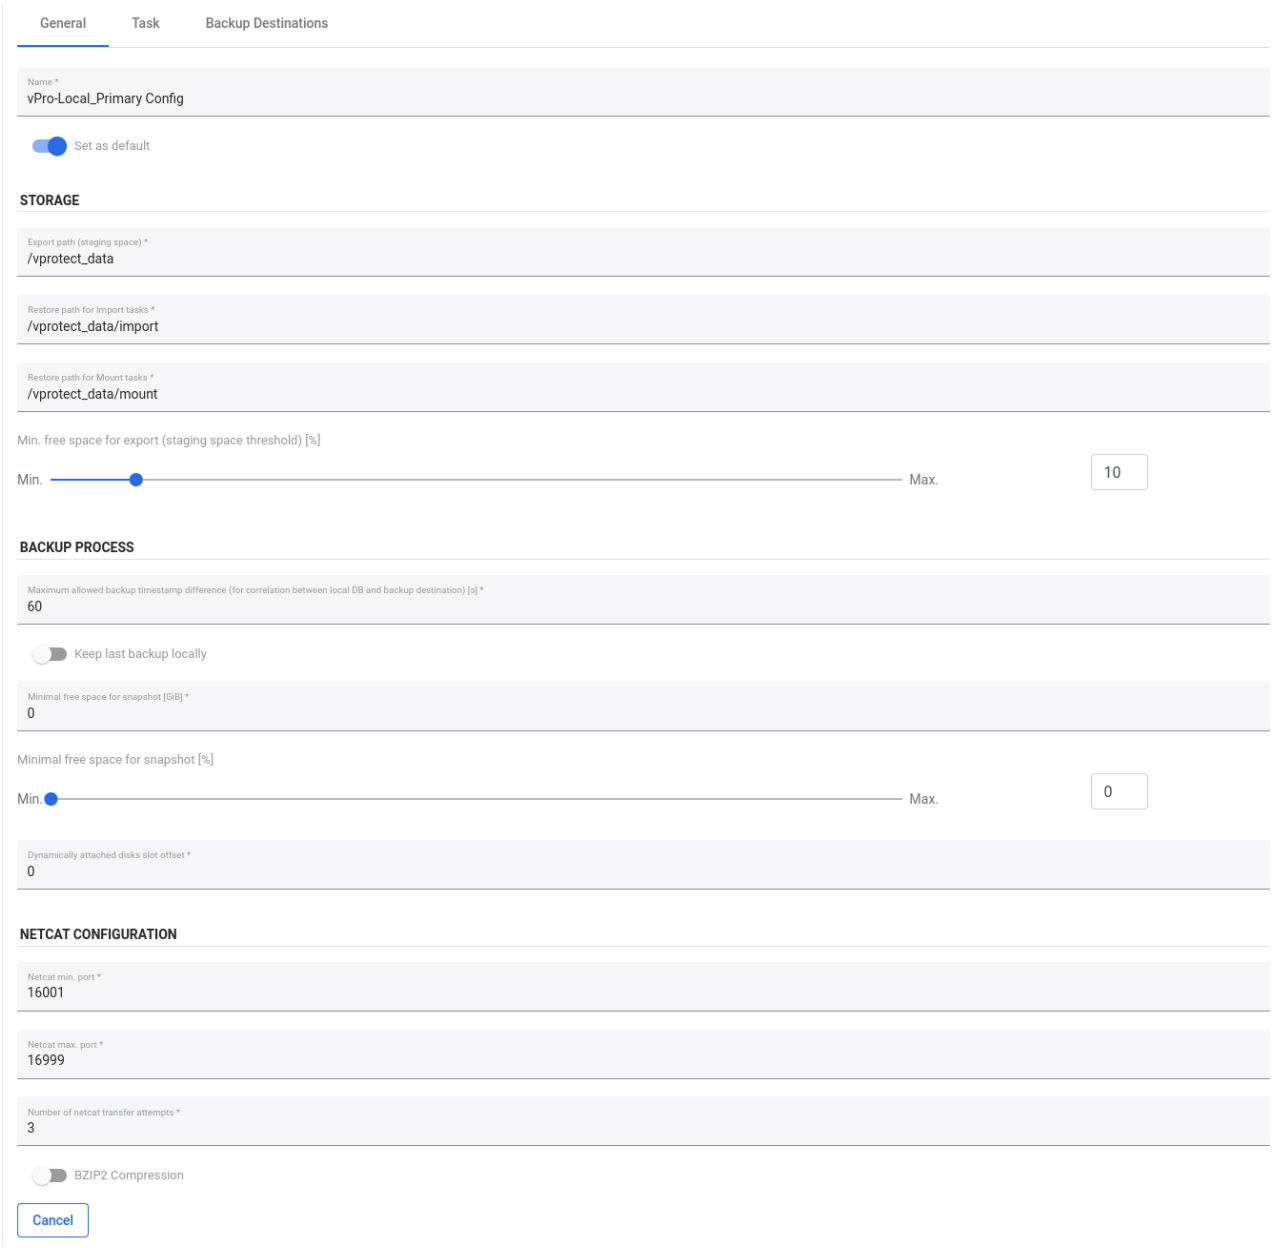
General
Task
Backup Destinations
Name *
vPro-Local_Primary Config
Set as default
STORAGE
Export path (staging space) *
/vprotect_data
Restore path for Import tasks *
/vprotect_data/import
Restore path for Mount tasks *
/vprotect_data/mount
Min. free space for export (staging space threshold) [%]
Min.
Max.
10
BACKUP PROCESS
Maximum allowed backup timestamp difference (for correlation between local DB and backup destination) [s] *
60
Keep last backup locally
Minimal free space for snapshot [GiB] *
0
Minimal free space for snapshot [%]
Min.
Max.
0
Dynamically attached disks slot offset *
0
NETCAT CONFIGURATION
Netcat min. port *
16001
Netcat max. port *
16999
Number of netcat transfer attempts *
3
BZIP2 Compression
Cancel

Task

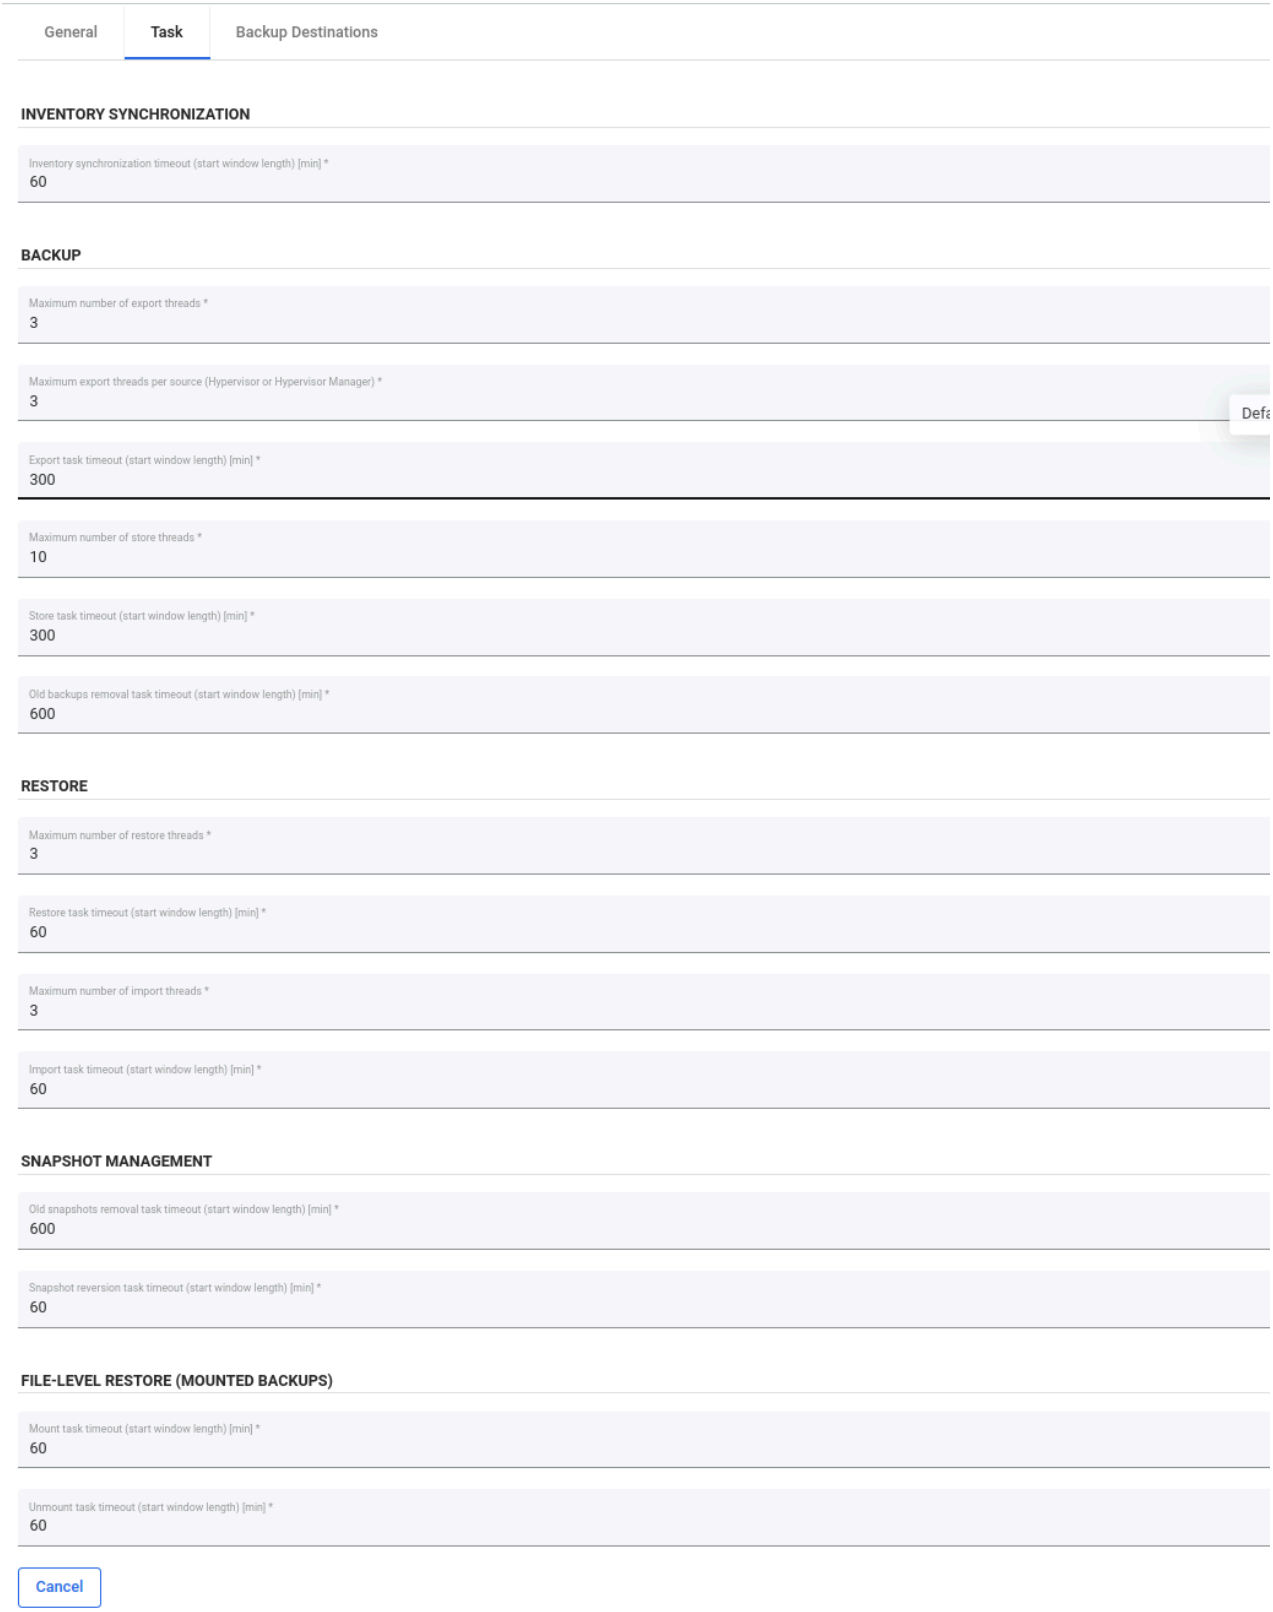

INVENTORY SYNCHRONIZATION

- `Inventory synchronization timeout (start window length) [min]` - default length of the start window for index tasks

BACKUP

- `Maximum number of export threads` - max. number of export tasks per node (total)
- `Maximum export threads per source (HV or HVM)` - max. number of export tasks per node and per HV/HVM
- `Export task timeout (start window length) [min]` - default length of the start window for export tasks
- `Maximum number of store threads` - max. number of store tasks per node
- `Store task timeout (start window length) [min]` - default length of the start window for store tasks
- `Old backups removal task timeout (start window length) [min]` - default length of start window for old backup removal tasks

RESTORE

- `Maximum number of restore threads` - max. number of restore tasks per node (total)
- `Restore task timeout (start window length) [min]` - default length of start window for restore tasks
- `Maximum number of import threads` - max. number of import tasks per node (total)
- `Import task timeout (start window length) [min]` - default length of the start window for import tasks

SNAPSHOT MANAGEMENT

- `Old snapshot removal task timeout (start window length) [min]` - default length of start window for old snapshot removal tasks
- `Snapshot reversion task timeout (start window length) [min]` - default length of start window for snapshot reversion tasks

FILE-LEVEL RESTORE (MOUNTED BACKUPS)

- `Mount task timeout (start window length) [min]` - default length of start window for mount tasks

- `Unmount task timeout (start window length) [min]` - default length of start window for unmount tasks

Backup destinations

This section is used to add/remove backup destinations to the nodes using this configuration. Only backup destinations enabled here can be used by the nodes.

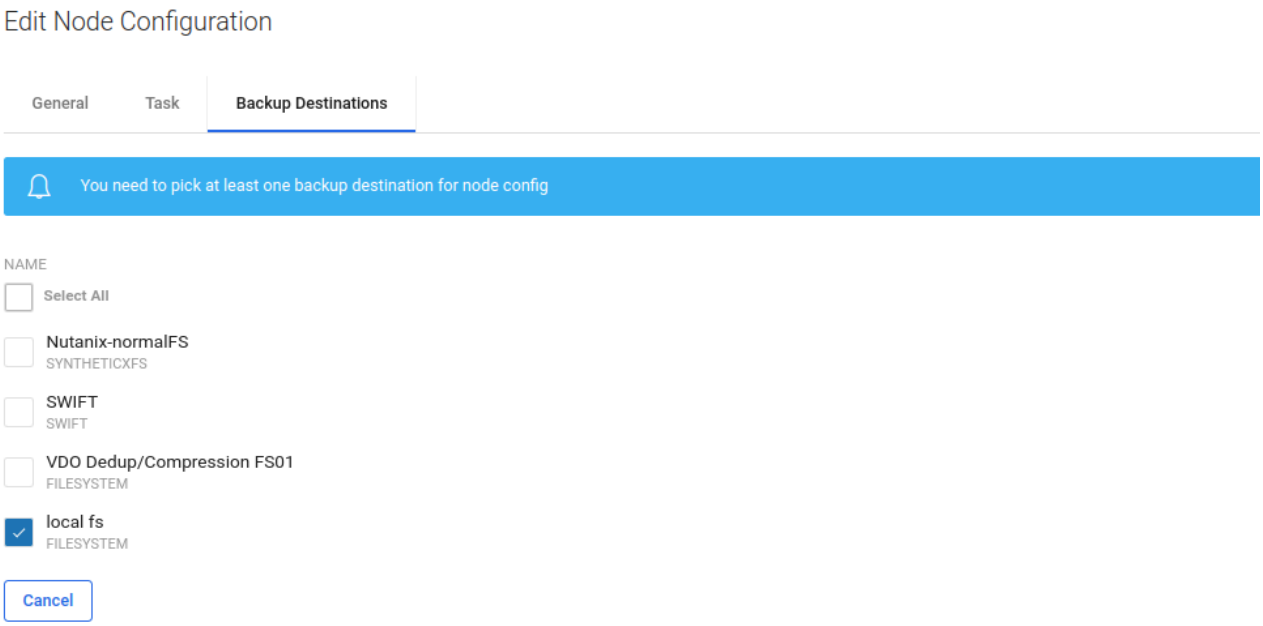

Cloud Agent

This section is used to provide connection details to the Windows Helper Host which is used in export to PST. You can find more information about requirements for Export to PST here.

Fill in the indicated fields:

- Host - IP address of the Windows host where cloud2pst converter is installed
- User - Windows host user
- Password - Windows host user password

Edit Node Configuration

General | Task | Backup Destinations | Cloud Agent

Task progress refresh rate [sec] *
15

WINDOWS HELPER HOST CREDENTIALS

Host
10.10.10.10

User
administrator

Password *
••••••••••••

Retype password *
••••••••••••

Show password

Access Management

Access Management

This section allows you to manage Users, Groups, Roles and OS Credentials.

The role is a set of permissions in Dell vProtect.

Users can belong to multiple groups.

Each group can have multiple roles assigned.

OS Credentials are used in connections with external systems, for example in application backup.

Users

Users

There are two types of users:

- local users
- LDAP users (refer to how setup LDAP authentication)

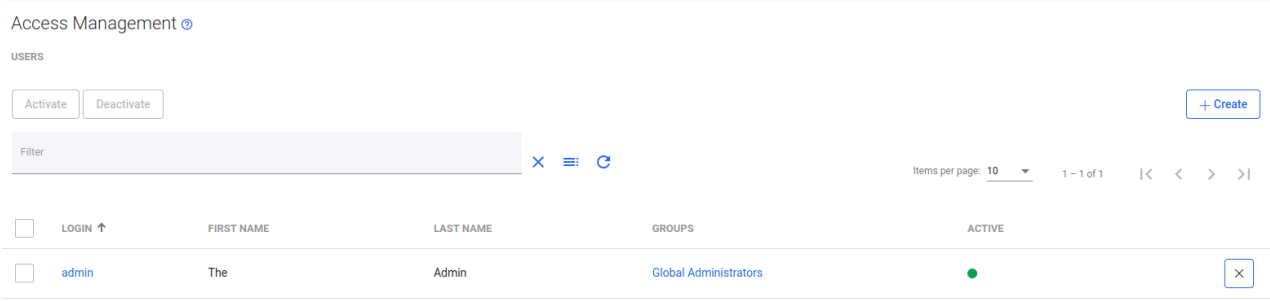

To create a new user, click the button on the right 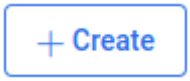

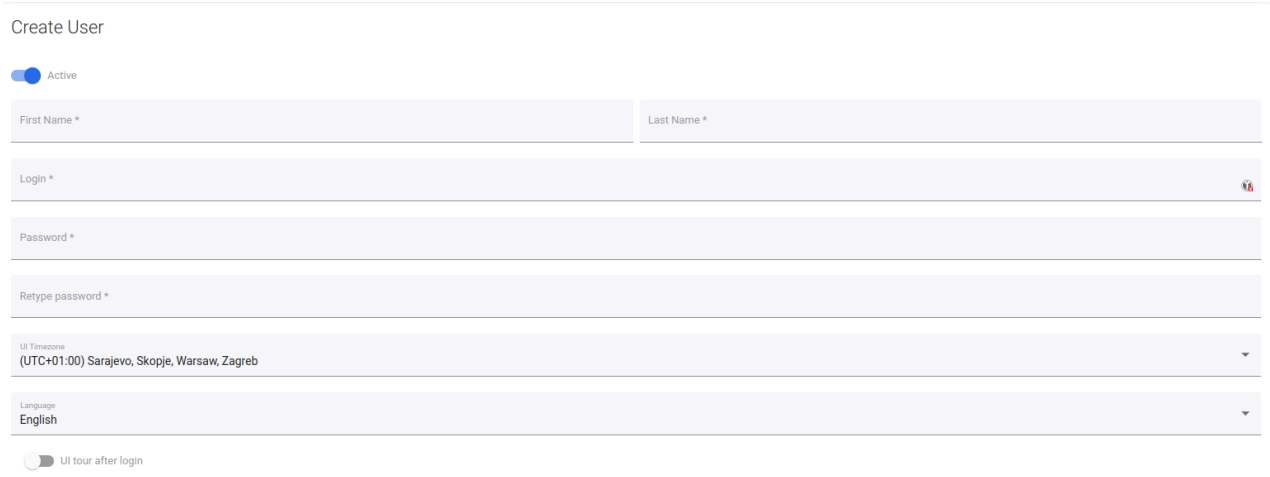

Complete the following parameters:

- First Name
- Last Name
- Password

- Timezone
- Language

Save, and now you can add this account to the specific group. In the Groups section, click on the name of the group and add a user.

Groups

Groups

Group is a central place to bind together a set of roles with users. There are already a set of predefined groups in Data Protector for Cloud Workloads:

Global Administrator - user have all possible privileges.

Virtual Environment Administrator - the user has full permissions to the **Virtual Environment** menu. Additionally, he can read information about nodes, nodes config and backup destinations.

Application Administrator - the user has full permissions to the **Applications** menu. Additionally, he can read information about nodes, nodes config and backup destinations.

Storage Administrator - the user has full permissions to the **Storage** menu. Additionally, he can read information about nodes, nodes config and backup destinations.

Backup Destination Administrator - the user has full permissions to the **Backup Destination** menu. Additionally, he can read information about nodes, nodes config and virtual environment infrastructure.

User Administrator - the user has full permissions to the **Access Management** menu. He can create new users, roles, groups or OS credential. User in this group can modify other users except their passwords.

Reporting Administrator - the user can browse and export all reports.

Read Only - the user has full read permissions to all menu sections.

Operator - the user has full read permissions to all menu sections. Additionally, he can execute backup, restore, mount and snapshot tasks for all kinds of instances.

> **Note:**
>
> - Automatically created accounts based on LDAP logins will have the `Operator` role by default.
> - Current group management is available in Data Protector for Cloud Workloads WebUI only (they are not mapped to LDAP groups)

Roles

Roles

The Role is a set of permissions to the different sections and actions in Data Protector for Cloud Workloads. When multiple roles are assigned to the Group, the resulting permissions will be the sum of all permissions from all assigned roles.

> **Note:**
>
> - It is mandatory to have always at least one account with permission to manage Users
> - in other words: at least one account must be assigned to the group, which has a role with this permission.

Permissions

Section (Name)	Permission	Allows
Nodes	Instances	Read - allows you to view the list of instances and basic information Register - allows you to register new nodes in the environment Write - allows you to change selected node config Delete - allows you to remove a node from the environment
Nodes	Node Configurations	Read - allows you to browse the list of configs and open them Write - allows you to modify existing configs and create

		new ones
Virtual Environments	Instances	Read - allows you to browse the list of instances and open them Write - allows you to modify virtual machine settings Backup - allows you to perform a manual backup Restore - allows you to perform a manual restore Snapshot - allows you to perform a manual snapshot Mount - allows you to mount a backup for file-level restore Clean old Snapshot - allows you to perform the "Clean old snapshot" task
Virtual Environments	Infrastructure	Read - allows you to browse the list of infrastructure objects for all tabs Write - allows you to modify settings for existing objects and to add new ones Test Connectivity - This permission allows you to enable or disable access to the connectivity test
Virtual Environments	Backup SLAs	Read - allows you to browse the list of policies and schedules Write - allows you to modify the existing and add new ones
Virtual Environments	Snapshot SLAs	Read - allows you to browse the list of policies and schedules

		Write - allows you to modify the existing and add new ones
Virtual Environments	Recovery Plans	Read - allows you to browse the list of policies and schedules Write - allows you to modify the existing and add new ones
Virtual Environments	Mounted Backups	Read - allows you to browse the list of mounted backups and also to download files Unmount - allows you to unmount a backup
Applications	Instances	Read - allows you to browse the list of instances and open them Write - allows you to modify virtual machine settings Backup - allows you to perform a manual backup Restore - allows you to perform a manual restore
Applications	Execution Configurations	Read - allows you to browse the list of execution configs Write - allows you to modify the existing and add new ones
Applications	Backup SLAs	Read - allows you to browse the list of policies and schedules Write - allows you to modify the existing and add new ones
		Read - allows you to browse the list of instances and open them

Storage	Instances	Write - allows you to modify storage instance settings Backup - allows you to perform a manual backup Restore - allows you to perform a manual restore Snapshot - allows you to perform a manual snapshot Mount - allows you to mount a backup for file-level restore Clean old Snapshot - allows you to perform the "Clean old snapshot" task
Storage	Infrastructure	Read - allows you to browse the list of infrastructure objects for all tabs Write - allows you to modify settings for existing objects and to add new ones Test Connectivity - This permission allows you to enable or disable access to the connectivity test
Storage	Backup SLAs	Read - allows you to browse the list of policies and schedules Write - allows you to modify the existing and add new ones
Storage	Snapshot SLAs	Read - allows you to browse the list of policies and schedules Write - allows you to modify the existing and add new ones
Storage	Mounted Backups	Read - allows you to browse the list of mounted backups and also to download files

		Unmount - allows you to unmount a backup
Reporting	Reporting	Read - allows you to view report data for virtual machines and storage providers Send report - This permission allows you to enable or disable the ability to send the report by mail
Reporting	Audit Log	Read - This permission allows you to enable or disable access to the audit log report
Backup Destinations	Backup Destination	Read - allows you to browse the list of backup destinations Write - allows you to modify settings for existing objects and to add new ones Test Connectivity - This permission allows you to enable or disable access to the connectivity test Clean old backups - This permission allows you to enable or disable the ability to perform this task
Access Management	Access Management	Read - allows you to browse the list of objects for all access management tabs Write - allows you to modify existing objects and to add new ones Change passwords - allows you to change user passwords
Settings	Settings	Read - allows you to view selected settings for all tabs in the "settings" section

		Write - allows you to modify settings for all tabs in the "settings" section
Tasks Console	Tasks Console	Read - this permission allows you to hide or show the "Task Console" floating panel Write - allows you to remove/cancel tasks from the list
Restore Job	Restore Job	Read - This permission allows you to enable or disable access to the information about restore tasks - under tasks console, details of virtual machine or storage instance and from recovery plans

Security contexts

The security context defines the set of system objects that can be accessed with defined permissions.

The object hierarchy is constructed in a way that any defined privileges will apply to the specified object and to all the downstream objects, therefore the `System Level` security context defines the access to all objects across the platform and disables the choice of lower-level objects.

Security contexts have been aggregated into the specified hierarchy. Different system object types can have different parental objects (or none) and so will appear in the different sections of the tree.

Selected objects will be visible in the table on the right-hand side.

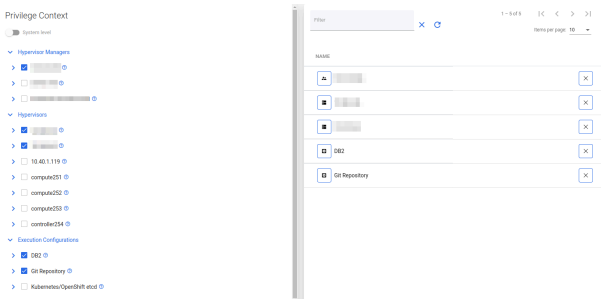

The full hierarchy:

- Hypervisor Managers
  - Projects
    - Virtual Environments
  - Data Centers
    - Hypervisors
      - Virtual Environments
    - Clusters
      - Hypervisors
      - Virtual Environments
    - Virtual Environments
  - Clusters
    - Hypervisors
    - Virtual Environments
  - Hypervisors
    - Virtual Environments
  - Virtual Environments
- Hypervisors
  - Virtual Environments
- Application Configurations
  - Applications
- Storage Providers
  - Storages

OS Credentials

OS Credentials

OS Credentials are center managed settings, that let you using them for :

- executing pre or post snapshot commands
- uploading files from mounted backup

Main place when you can mange OS Credentials is **Access Management** → **OS Credentials**. To add new credentials click **Create** button.

In the form provide:

- Name
- User
- Password
- SSH key path (optional)

Settings

Settings

In this section, you can find various general settings for Data Protector for Cloud Workloads as:

Global Settings

Internal DB Backup

Notification Rules

Mailing Lists

Global Settings

Global Settings

Global

Global value settings for some retentions and schedules:

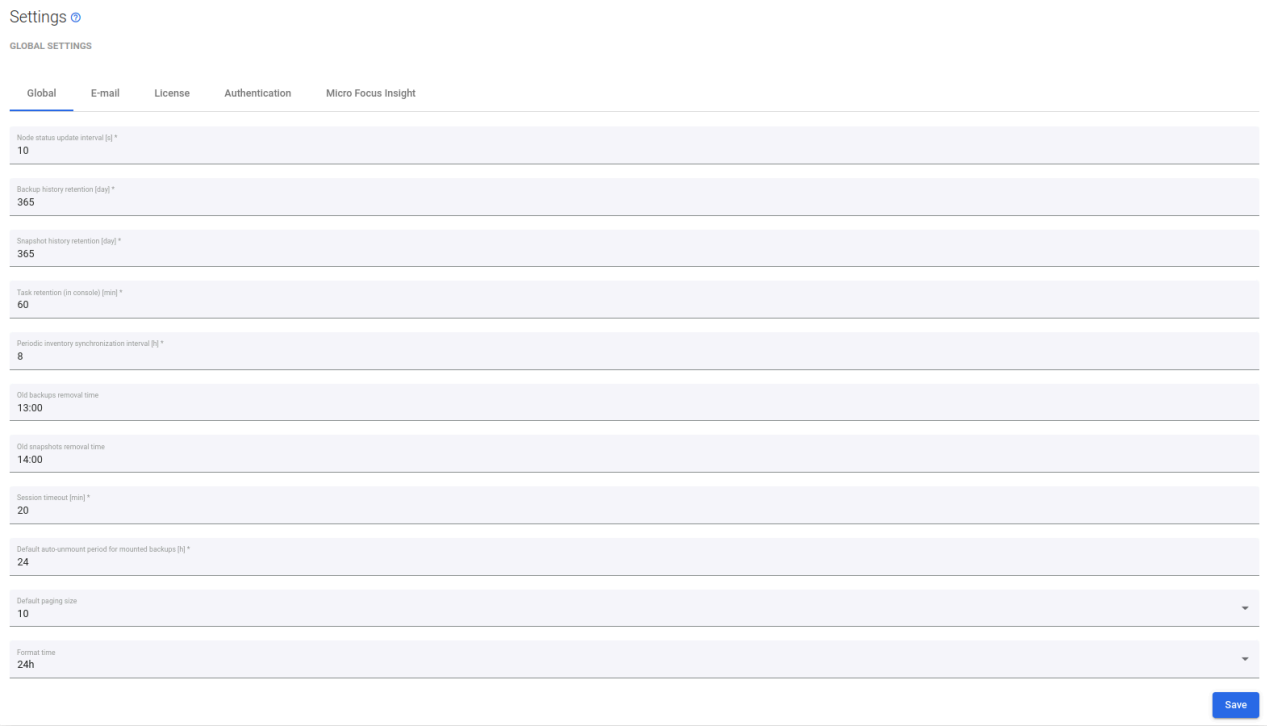

- `Node status update interval` - how often nodes should update their status
- `Backup history retention` - how long should the history of backups be kept (even removed from backup provider)
- `Task retention (in console)` - how long finished/failed tasks should be kept in the console in UI/CLI
- `Periodic inventory synchronizarion interval` - how often Data Protector for Cloud Workloads should scan for changes in VM inventory on HV/HVMs and Microsoft 365 users accounts, sites, and teams
- `Old backups removal time` - time, when daily backup destination cleanup should be invoked (for all backup destinations)

- `Old snapshots removal time` - time, when daily snapshots cleanup should be invoked (for all VMs with any policy assigned)
- `Session timeout [min]` - Session timeout [min] - the time after which you will be logged out of the WebUI
- `Default paging size` - Default value of items shown on lists
- `Format time` - you can choose between 12h and 24h time format

E-mail

Email configuration for reports purposes

- `Sender e-mail` - address from which should e-mail be sent
- `SMTP server` - SMTP server address
- `SMTP port` - SMTP server port
- `SMTP SSL port` - SMTP SSL port (if enabled)
- `SMTP user` - SMTP account used to send e-mails
- `E-mail recipients (comma-separated)` - list of recipients of daily backup report
- `Daily backup report (sending time)` - time when daily backup report should be sent
- `Daily backup report` - sending time for summary report

- `Enable virtual environment/application grouping` - you can group environment by selected parameter

License

This section enables you to view current license status and upload a new license if necessary.

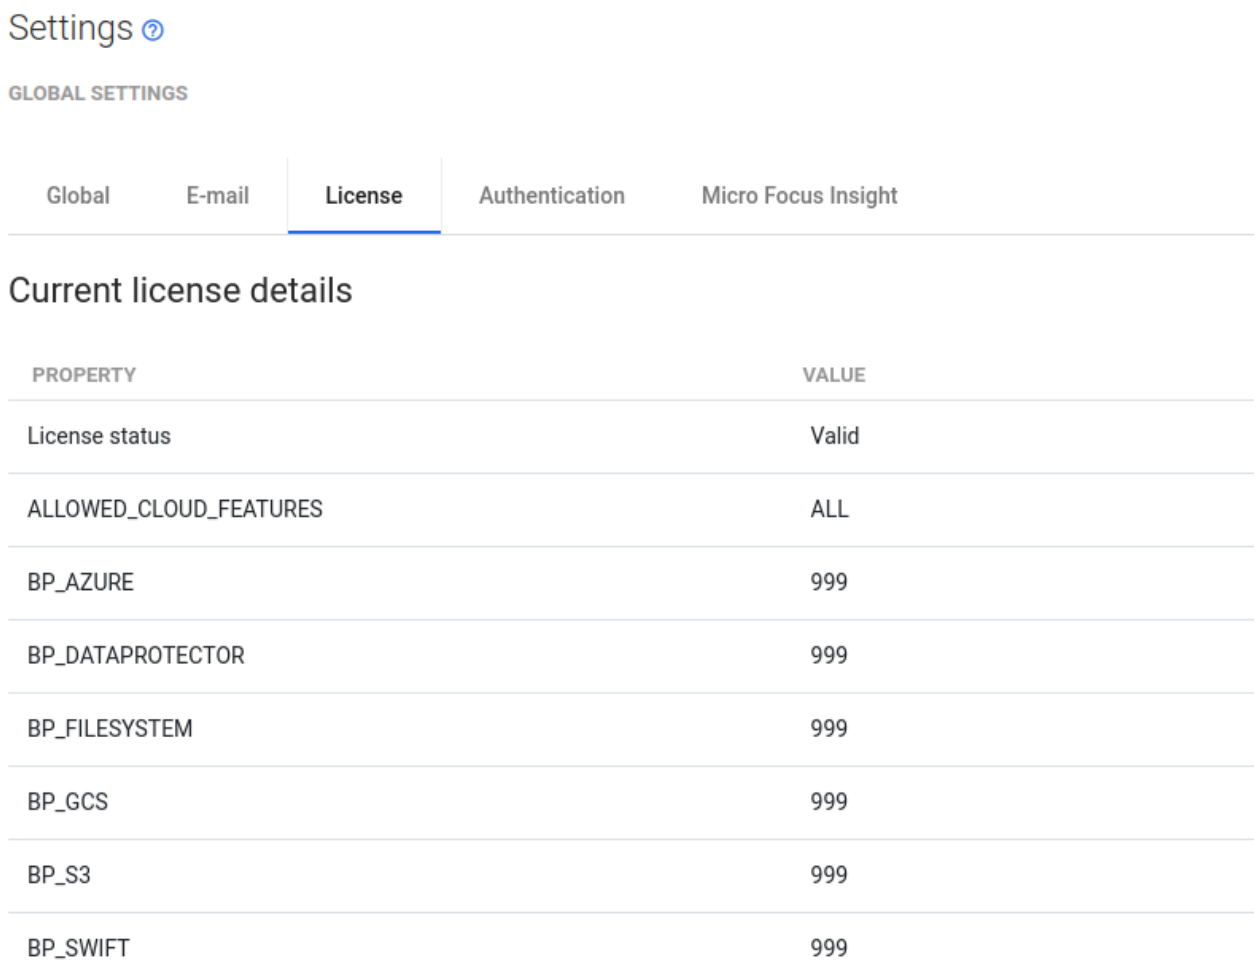

License details:

- `MAX_xxx_HOSTS` - maximum number of hosts for given platform
- `BP_xxx` - maximum number of backup destinations per backup provider type
- `EXPIRE_DATE` - trial period expiration date

Authentication

This section enables you to set up single sign-on between Keycloak or LDAP, and product. This section assume you have installed and are using Keycloak or LDAP.

In each of the configurations, you can select a group to which the user will be automatically assigned.

Keycloak

> **Note:** Supported Keycloak versions: 15.1 and newer

- `Server URL` - Keycloak server URL (if you want to use fqdn of server url you need to use DNS on Server machine or add line in `/etc/hosts`, in example: '1.2.3.4 dc.fqdn.address')
- `Realm` - Name of the realm configured in Keycloak
- `Resource` - Name of the client configured in keycloak
- `Secret` (optional) - Credential secret, if configured in Keycloak

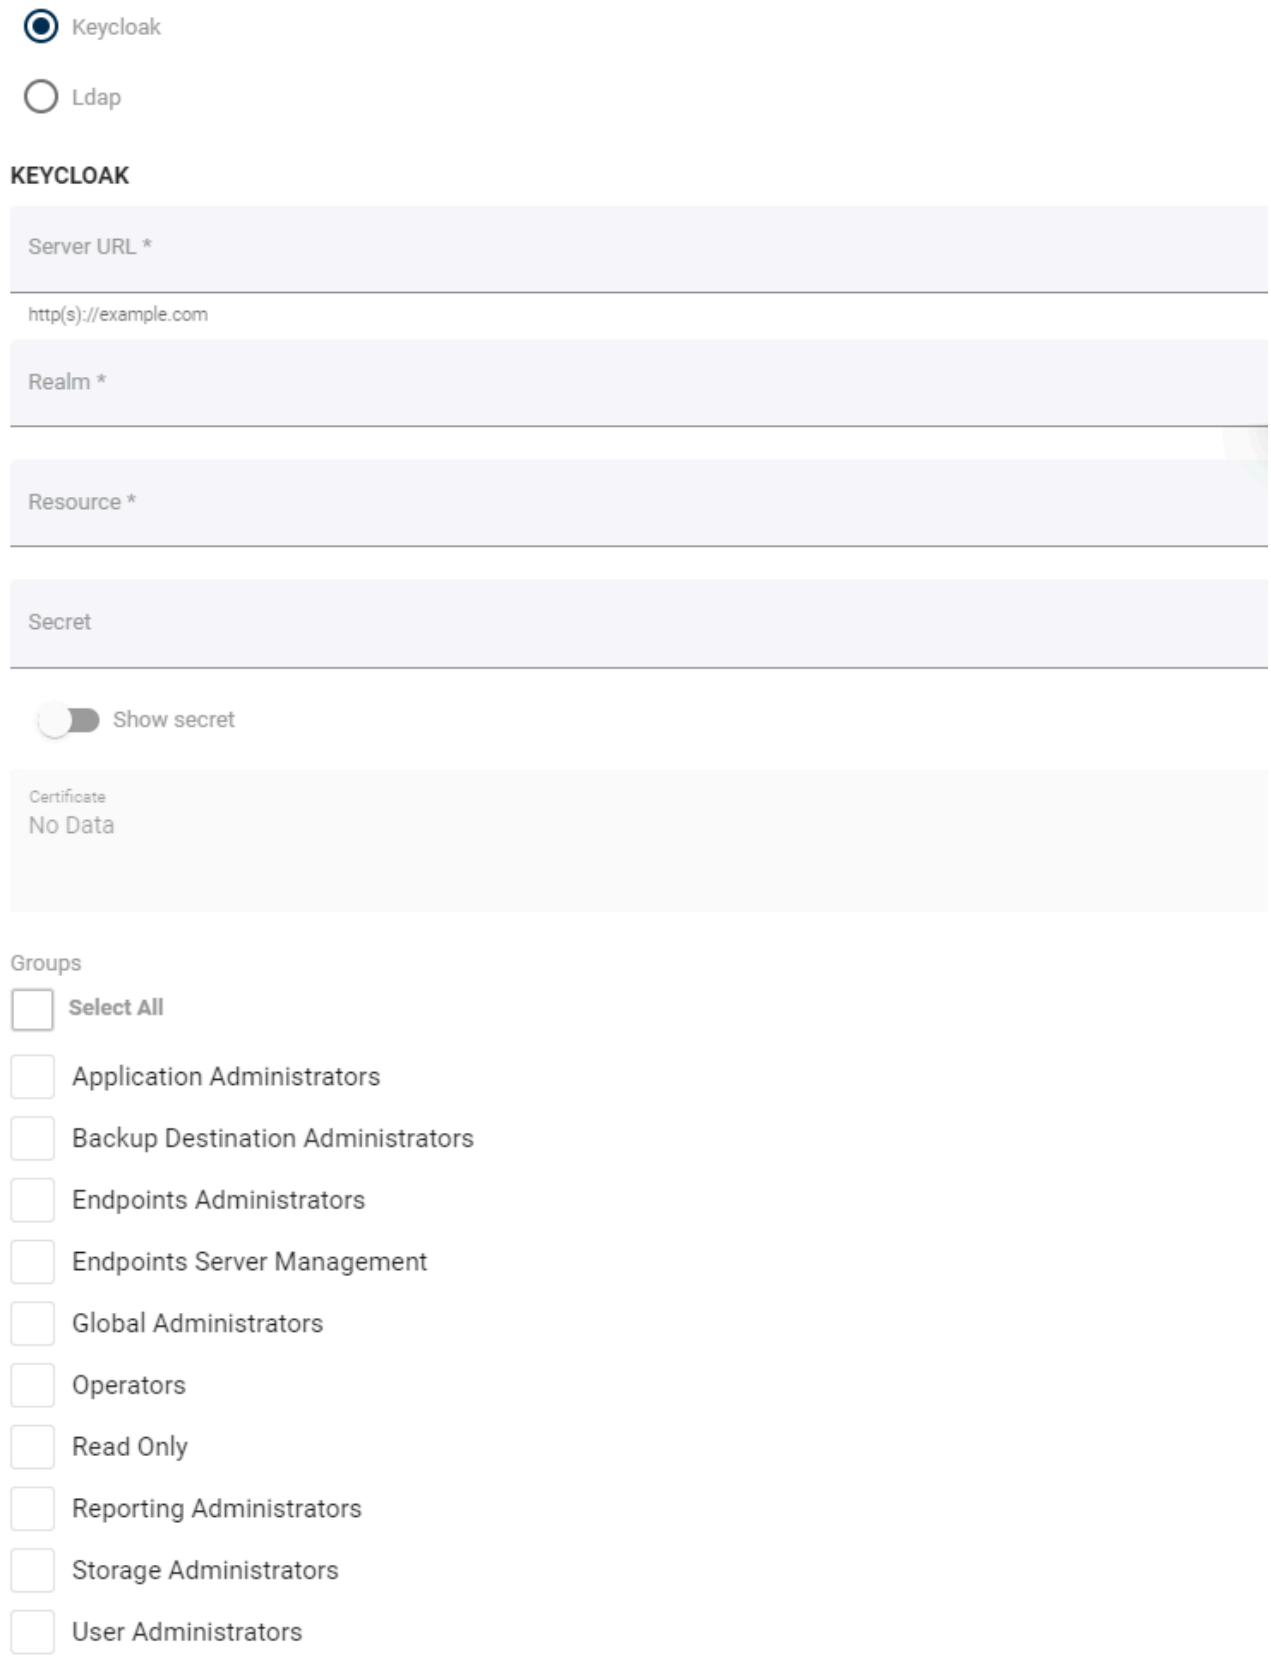
Keycloak
Ldap
KEYCLOAK
Server URL *
http(s)://example.com
Realm *
Resource *
Secret
Show secret
Certificate
No Data
Groups
Select All
Application Administrators
Backup Destination Administrators
Endpoints Administrators
Endpoints Server Management
Global Administrators
Operators
Read Only
Reporting Administrators
Storage Administrators
User Administrators

LDAP

> **Note:**
>
> - accounts will be added to Data Protector for Cloud Workloads automatically with the first successful login
> - all LDAP variables are case sensitive

- `Server URL` - LDAP server URL (if you want to use fqdn of server url you need to use DNS on Server machine or add line in `/etc/hosts`, in example: '1.2.3.4 dc.fqdn.address')
- `Base DN` - Base DN (Distinguished Name) that needs to be searched (it need full chain to OU with Users which you want to log into)
- `User filter` - filter to be used to authenticate only users in a specific group:

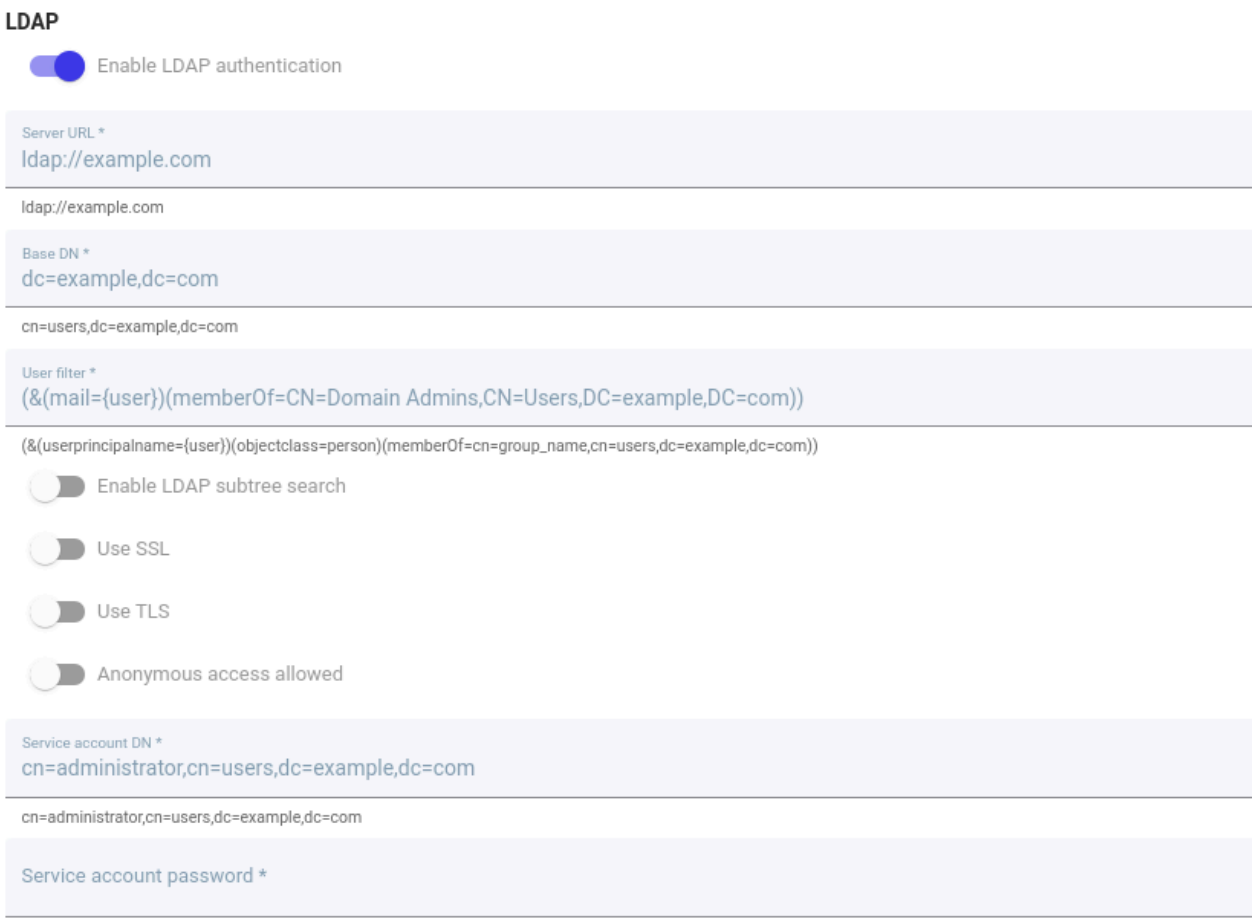

In example:

```
(&(mail={user})(memberOf=CN=Domain Admins,CN=Users,DC=example,DC=com))
```

Filter is combined from two sections:

- First section - `(mail={user})` - is a variable from LDAP account, which will be use as login
- Second section - `(memberOf=CN=Domain Admins,CN=Users,DC=example,DC=com)` - is a actual filter, which define who can log into WebUI, you need to define here variable name and DN of specific variable

Explanation of other options:

- `Enable LDAP subtree search` - when disabled, only 1 level below base DN is being searched
- `Use SSL` - enables SSL for LDAP connection ('ldaps://')
- `Use TLS` - enables TLS for LDAP connection
- `Anonymous access allowed` - if users are not allowed to anonymously browse LDAP directory you need to provide an account that has that privilege:
  - `Service account DN` - DN of the user
  - `Service account password` - password of that user

Internal DB Backup

Internal DB Backup

This allows you to control the process of creating an internal database backup from one place. Clicking on the magnifying glass icon will take you to the appropriate menu.

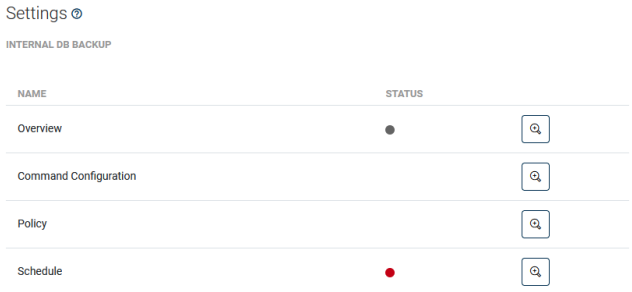

Notification Rules

Notification Rules

General

Data Protector for Cloud Workloads allows a Notification Rules feature to create a semi-custom rule that will send notification triggered by an event specified in rule. Go to Settings on a left menu and then click on Notification Rules button.

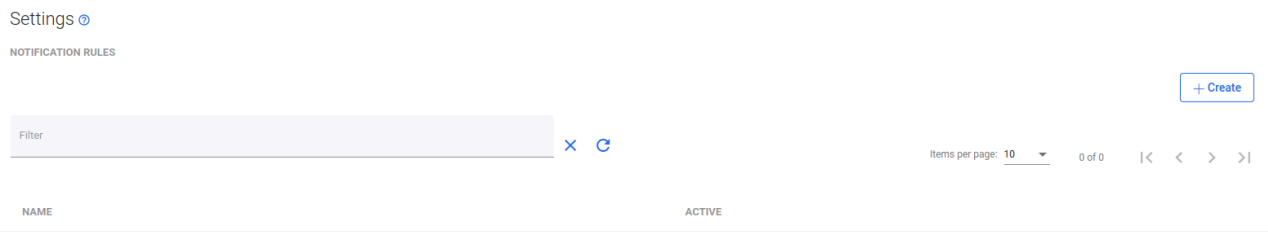

Creating new Notification Rule

To create new rule, simply click on **Create** button. A new window will appear.

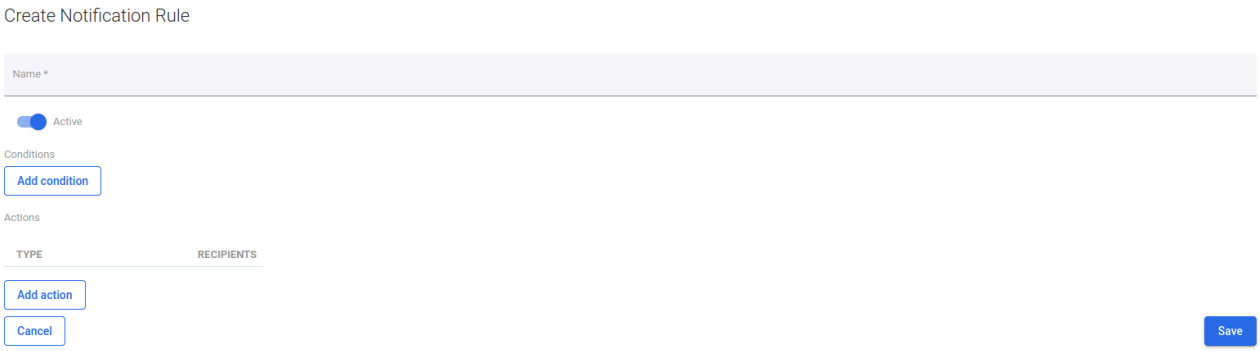

Fill in a name for your rule, set your notification to active and add one or more conditions from a list.

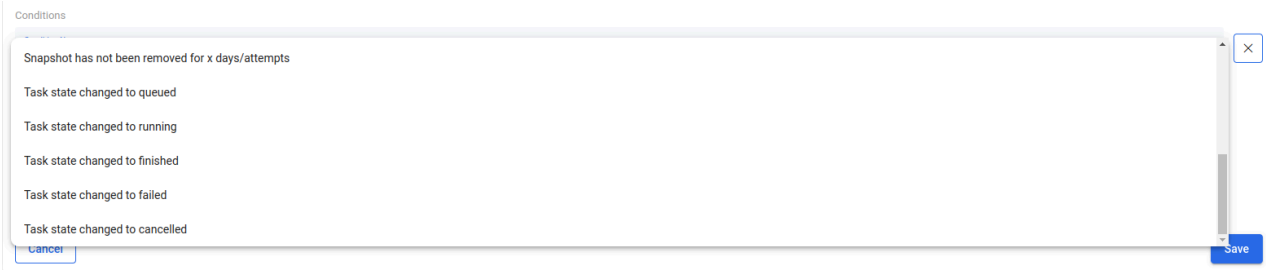

For some conditions, you will need to provide more information, for example, choose TaskType.

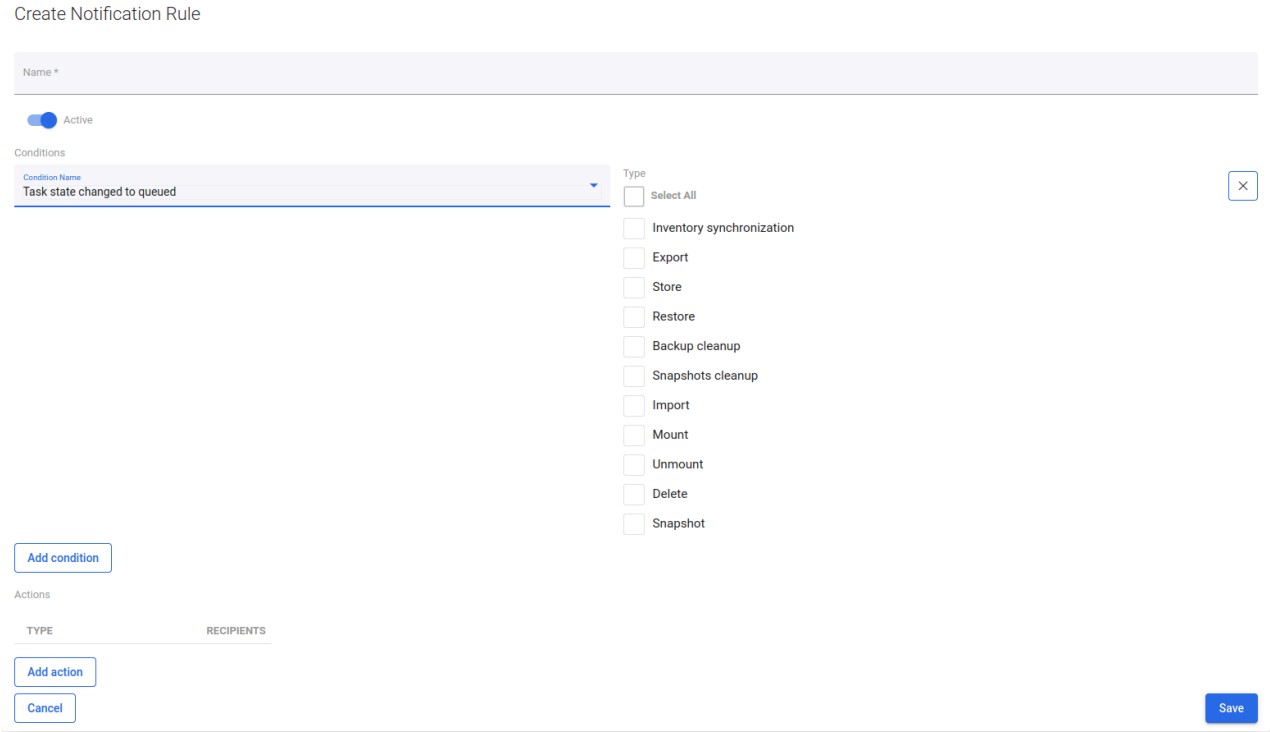

The last step is to add an action with related recipients.
Click **Add Action** button, choose **Action Type** from a list, add one or more recipients passing their email addresses.

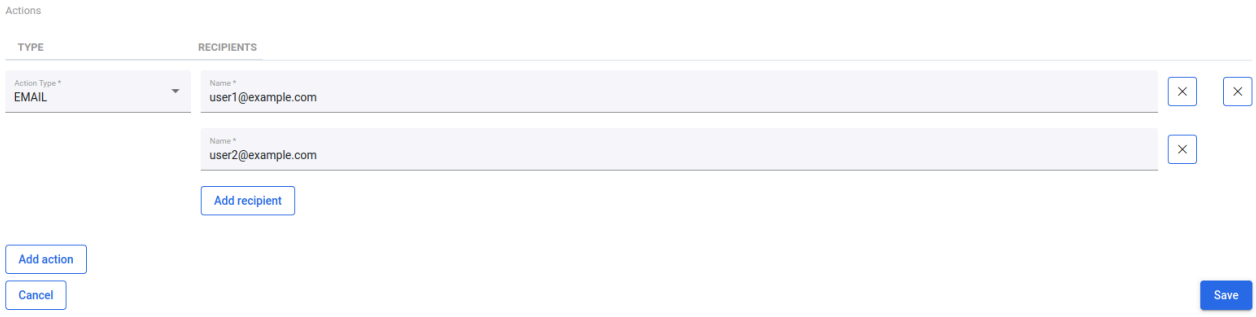

After an event that will trigger a notification, all recipients who have been added to the notification rule will receive an email notification.

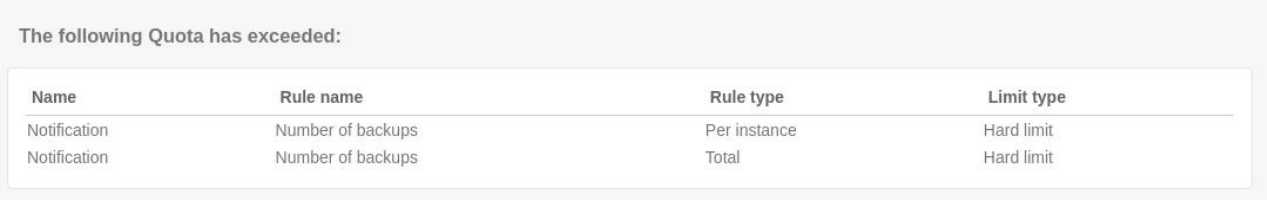

Slack integration

From Data Protector for Cloud Workloads version 5.0, there is a possibility to send notifications to the Slack channel. Follow these steps to authorize platform in your Slack workspace:

1. Go to this page ↗ and create a new Slack app in the workspace where you want to post messages. Select From scratch option and Enter the **App Name** and **Workspace** in which you want to use this app.

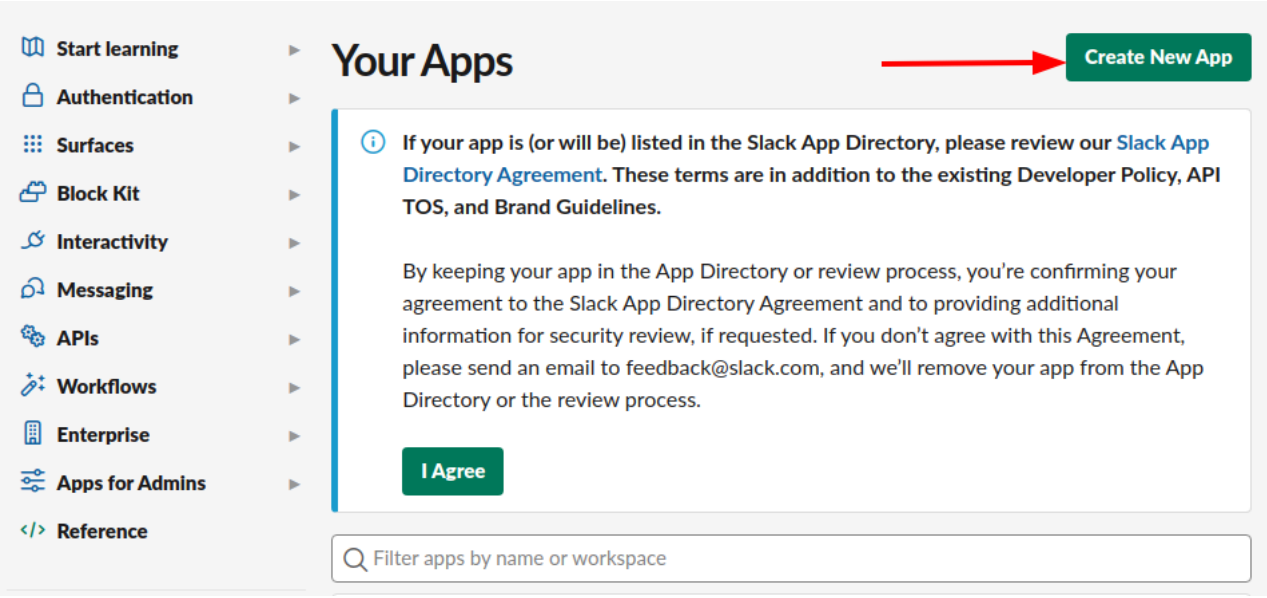

2. From the Features page, toggle Activate Incoming Webhooks on. Click Add New Webhook to Workspace.

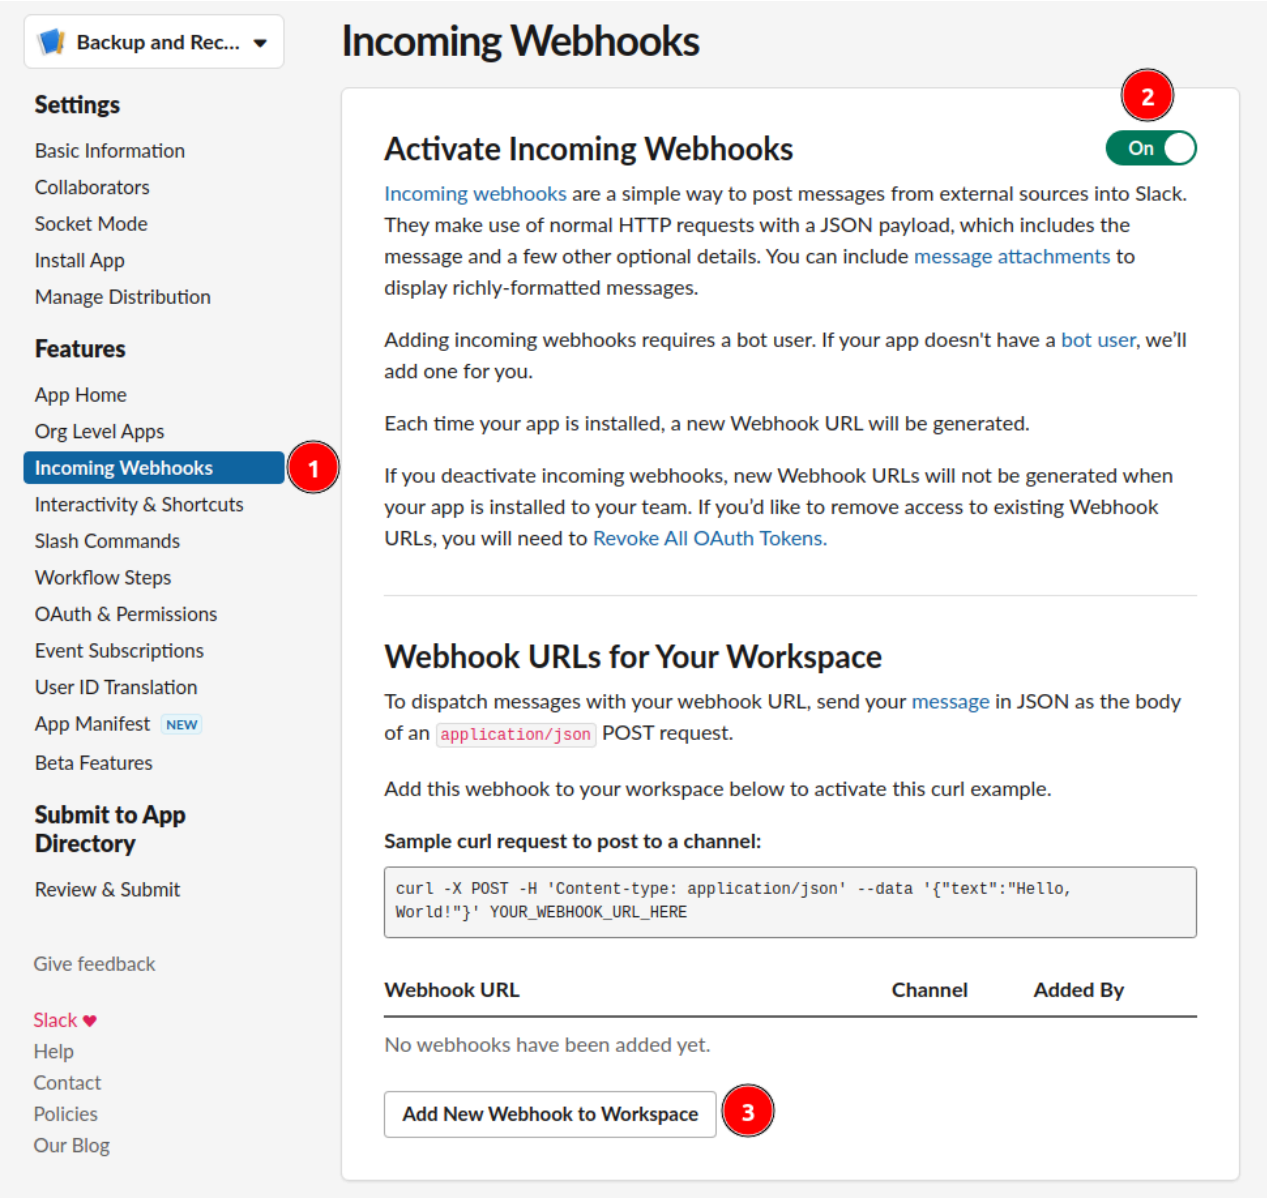

3. Pick a channel that the app will post to, then click Allow.

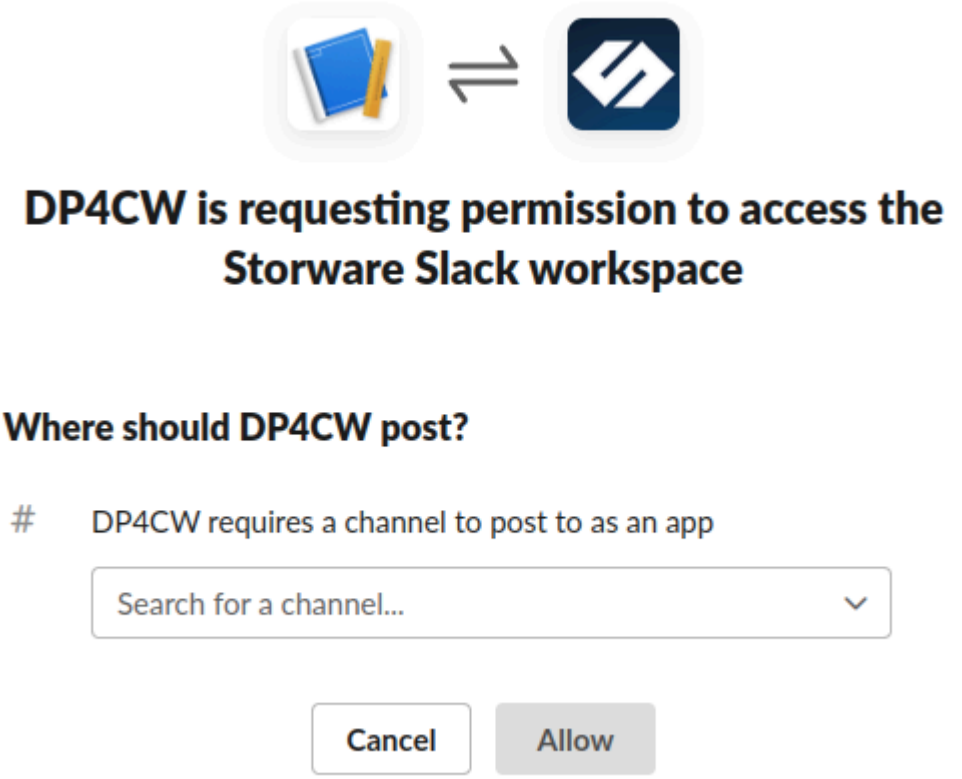

4. Copy the Webhook URL and use its Notification rules. Select Action Type as SLACK and paste the URL in the **Name** field.

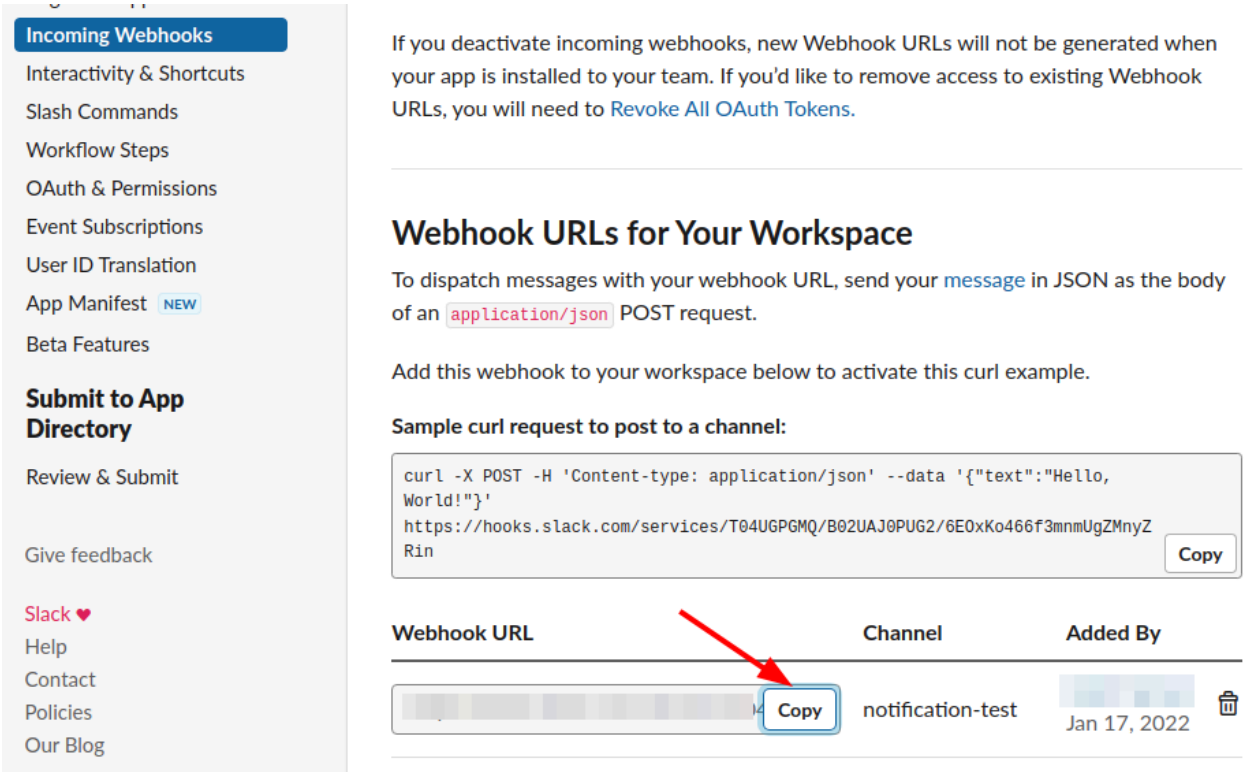

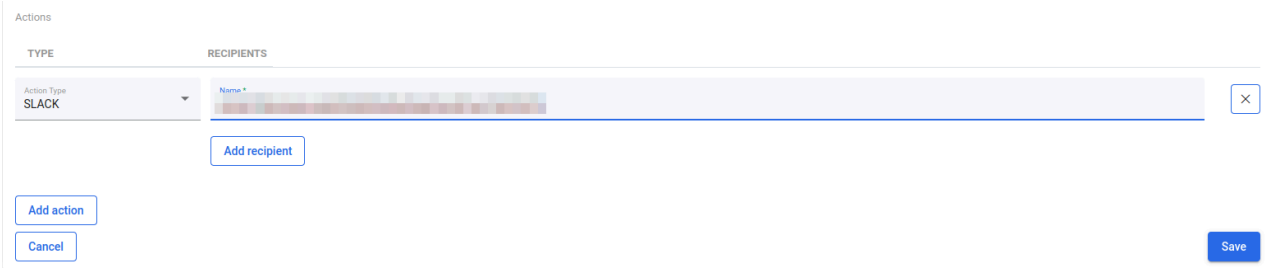

Configuration of Push Notification

To configure a push notification, provide URL of endpoint address, where you want to receive notifications and add if needed custom headers that will be added to POST request with notification.

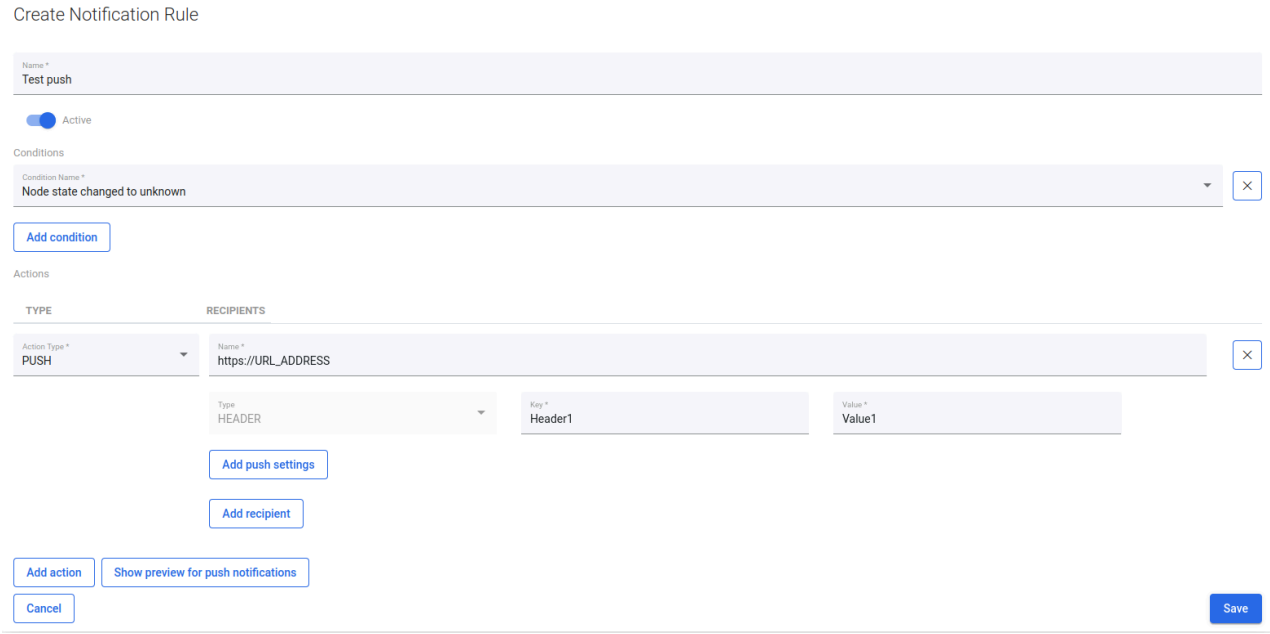

Body of notification will contain two fields:

- **subject**
- **messages** - list of notification messages that have been generated base on notification policy and triggered during run of program

To check how your notification will look, click on button Show preview for push notifications

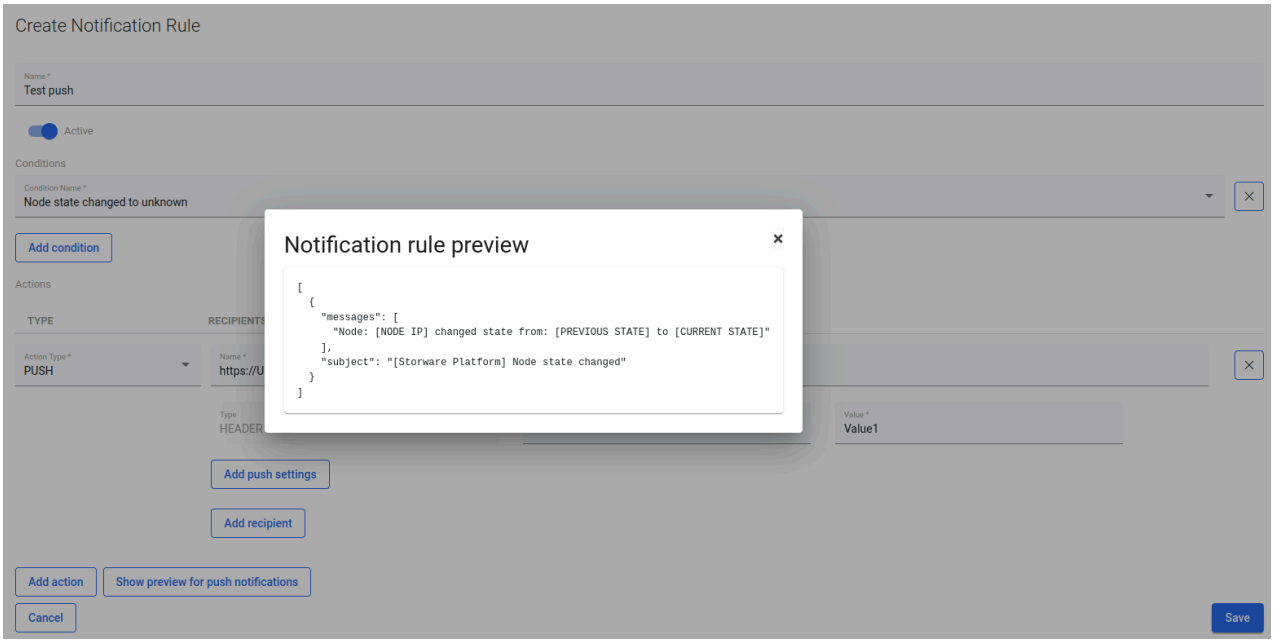
Create Notification Rule
Name *
Test push
Active
Conditions
Condition Name *
Node state changed to unknown
Add condition
Actions
TYPE
RECIPIENTS
Action Type *
PUSH
Name *
https://U
Type
HEADER
Value *
Value1
Add push settings
Add recipient
Add action
Show preview for push notifications
Cancel
Save
Notification rule preview
[
{
"messages": [
"Node: [NODE IP] changed state from: [PREVIOUS STATE] to [CURRENT STATE]"
],
"subject": "[Storware Platform] Node state changed"
}
]

Mailing Lists

Mailing Lists

This allows you to create a mailing list which can be used for sending group report e-mails.

Creating a new mailing list

To create a new list, click on **Create** button. A new window will appear.

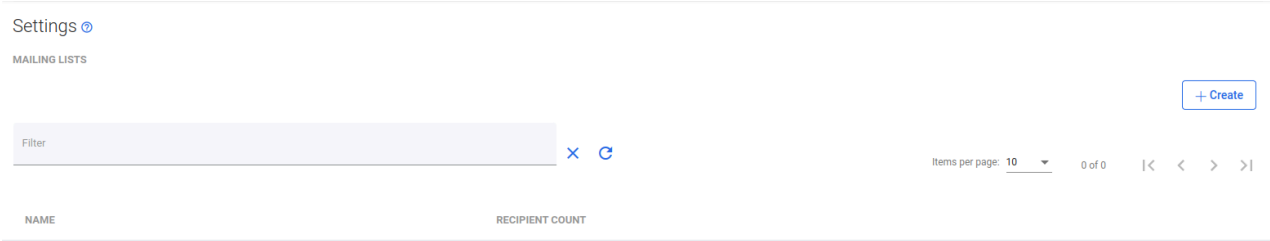

Fill the name for your list and add recipients e-mails by clicking **Add recipient** button.

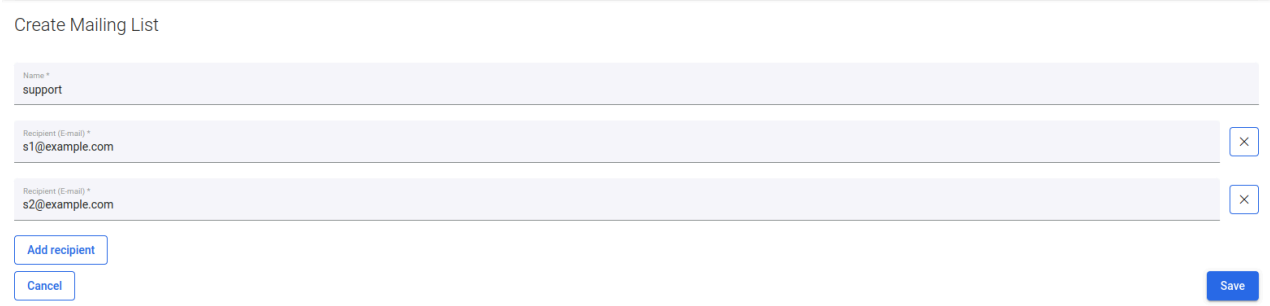

After you have finished adding recipients, click on the **Save** button.

Mailing list usage

Find the backup policy in which you want to use the mailing list and click on it to enter the **Edit policy** page. Enable the **Send daily backup/restore report for VMs assigned to this policy** option and select the mailing list.

Edit policy (Virtual Environment Backup)

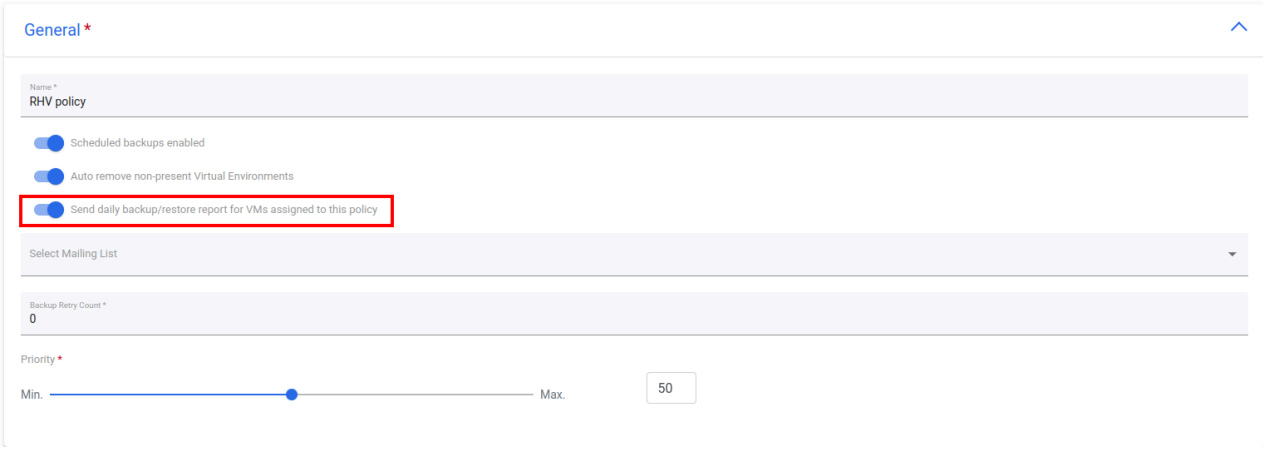

Upgrade

Upgrade

Before every update, check the version of installed packages. The database version is particularly important.

```
yum info vprotect-server vprotect-node mariadb-server
# Or
rpm -qa | egrep -e "vprotect|Maria"
```

If the host computer has an internet connection, use the yum command - you'll also see the new package versions provided by the repositories.

Server Upgrade

- Make sure you have the Data Protector for Cloud Workloads database backup.
    - You can use this command manually to back it up on-demand on the Data Protector for Cloud Workloads Server: `/opt/vprotect/scripts/backup_db.sh /path/to/backup/file.sql.gz`
- If Data Protector for Cloud Workloads was installed on a virtual machine (not a physical one), it would be a good move to take a snapshot.
- After backing up the database, you should carefully stop the Data Protector for Cloud Workloads service to make sure that you don't have any tasks running (a running task may cause problems updating the database).
    - View all tasks, if you see even one on the list, clear it (wait for the ongoing tasks to finish)
    - You can do this from the WebUI (it's faster)

```
[root@vprotect ~]# vprotect task -L

                GUID                    Type    State    [%]    Window
start       Window end     Pri.     Node       VM/APP
------------------------------------  ------  --------  ---  ----------
------  ---------------  ----  ----------  -----------
e3bb2496-3928-417c-a604-8c61b64df90e  Export  Running   0    2020-06-19
12:27  2020-06-19 17:27  50    vPro-Local  VM_01_Apine
05c1d6cc-fe3b-40fb-9811-94b976571d8e  Store   Finished  100  2020-06-19
12:10  2020-06-19 17:10  50    vPro-Local  VM_01_Apine
cb47190d-cf10-4cf9-8d1d-418eed5accf9  Export  Finished  100  2020-06-19
12:09  2020-06-19 17:09  50    vPro-Local  VM_01_Apine

#To delete a task from the list
[root@vprotect ~]# vprotect task -d cb47190d-cf10-4cf9-8d1d-
418eed5accf9
```

- Now, if you don't have any tasks on the list, you can stop the service.

```
[root@vprotect ~]# systemctl stop vprotect-server --now
```

- To make sure that no scheduler has started a task before stopping the service, let's query the database.
  - If the table is not empty, start the Data Protector for Cloud Workloads-Server service and clear the tasks again.

```
mysql -u root -p -e "Select * FROM vprotect.task;"
```

- Make sure you have MariaDB up-to-date - currently Data Protector for Cloud Workloads by default uses version **10.4**, while 10.2.31 is the minimum version supported.
  - If you need to migrate between versions (for example 10.3 to 10.4) - we recommend updating it as described here ↗, but when you uninstall MariaDB packages you **SHOULD NOT** remove the Data Protector for Cloud Workloads Server package (as a dependency) try the `--noautoremove` option: *As centos/rhel 7 do not have the --noautoremove option natively, please use the rpm method.*
  - Otherwise, minor MariaDB versions should be updated with `yum update`

- `rpm -e --nodeps "MariaDB-server-YOUR_VERSION_OF_PACKAGE"`
- Update the MariaDB repository to the correct version `vi /etc/yum.repos.d/MariaDB.repo`
- Install the new MariaDB-Server `yum install -y mariadb-server`
- Update all other components of MariaDB `yum update -y mariadb`
- Start the MariaDB engine `systemctl enable mariadb --now`
- Run mysql_upgrade to update the Data Protector for Cloud Workloads Database `mysql_upgrade --user=root --password`

- If the database update is successful, now you can start with the Data Protector for Cloud Workloads Update.
- Extract this package on the hosts with Data Protector for Cloud Workloads Server or Node:

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

- 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
```

- If the server service was not running before update, you may also need to execute:

```
systemctl start vprotect-server --now
```

Node Upgrade

1. Copy the Node RPM to all hosts with Data Protector for Cloud Workloads Node installed.
2. Run the script to configure the OS for Node:

```
vprotect-node-configure
```

3. If the node service was not running before the update, you may also need to execute:

```
systemctl enable vprotect-node --now
```

4. Log in to the web UI and check if the nodes are running.

- **Note:** You may need to refresh your browser cache after update - for Chrome use `CTRL+SHIFT+R` (Windows/Linux) / `CMD+SHIFT+R` (MacOS)

Cloud Server Upgrade

1. Copy the Cloud Server RPM to all hosts with Data Protector for Cloud Workloads Node installed.
2. Update each Cloud Server:

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

Cloud Agent Upgrade

1. Copy the Cloud Agent RPM to all hosts with Data Protector for Cloud Workloads Node installed.
2. Update each Cloud Agent:

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

Integration

Integration

In this section, we're going to highlight key aspects necessary to integrate 3rd party solutions with Data Protector for Cloud Workloads. We assume a typical scenario, where you want to invoke Data Protector for Cloud Workloads operations on behalf of a user of a self-service portal. We assume that the end-user uses the above-mentioned portal for a subset of administrative actions, and we want to give a user ability to perform basic backup-related operations.

These include:

- listing VMs in inventory (including already non-existing)
- getting VM details
- performing backup of a VM on the list
- browse backup history
- restore specified VM
- monitor progress of each operation
- manage policies and schedules

Architecture

The architecture below shows key components, communication, and data flows. All communication between 3rd party systems goes via RESTful API exposed by Data Protector for Cloud Workloads Server. Tasks are being performed by the Node behind the scenes. End-user is going to use only a 3rd party system to invoke and monitor the status of the tasks.

Multi-tenancy and permission handling are on the 3rd-party system side. There is however `tenantID` field in several cases which can be used to assign objects to the tenant on the 3rd-party system side, later to be used in filter criteria.

A 3rd-party system must use a dedicated Data Protector for Cloud Workloads account to perform operations.

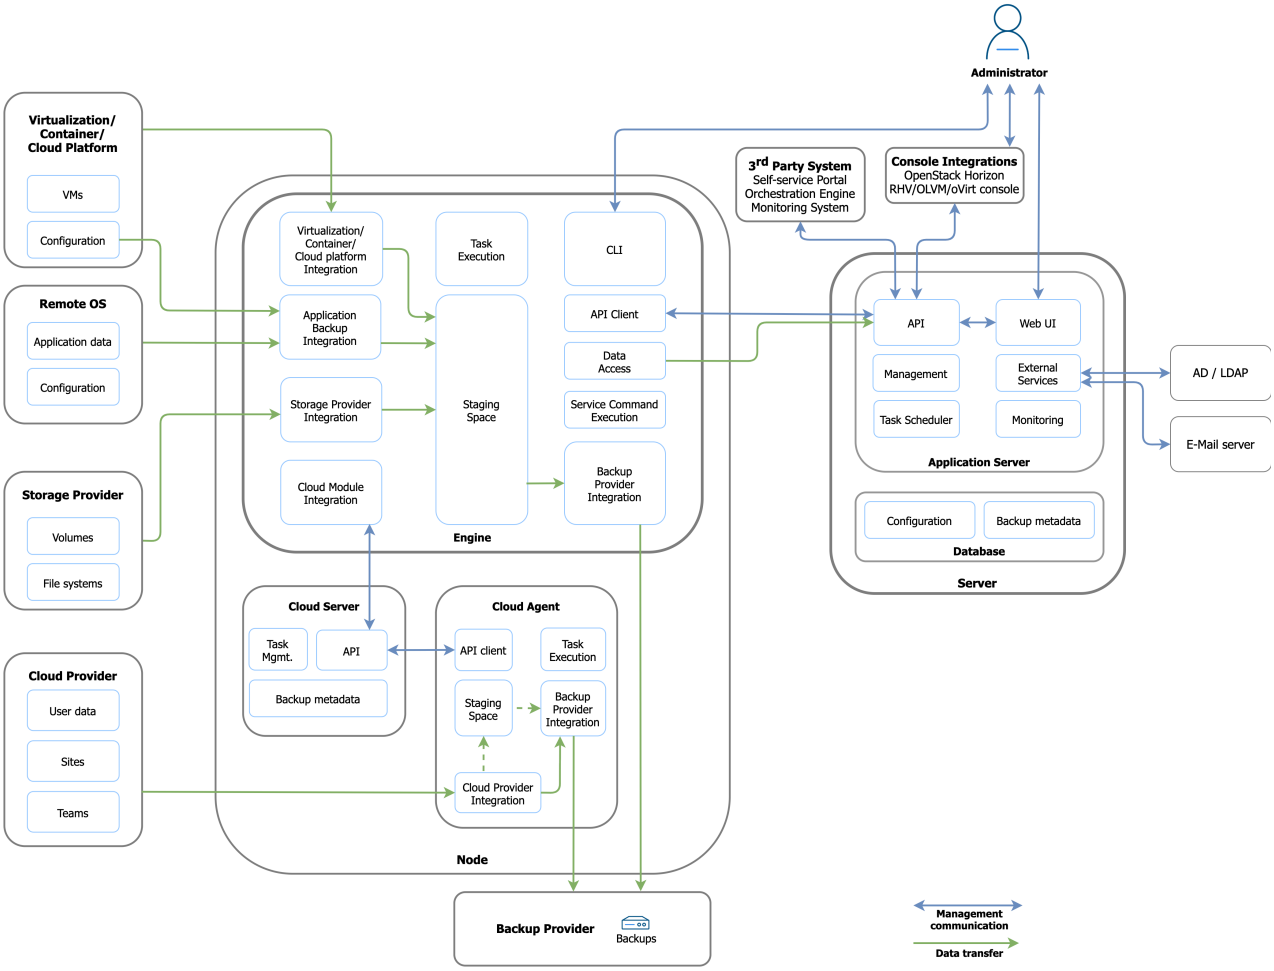

Integration steps

Setup

Your system is going to communicate over HTTPS which by default runs on port 8181, but during the setup can optionally be exposed on 443 as well. You need to generate an SSL certificate as described here ↗.

Data Protector for Cloud Workloads can invoke operations only on VMs that exist in its inventory. It is being synced periodically so that it reflects changes in the virtualization platform.

REST API can either be invoked directly or using a generated client for Java. Contact us to receive the current client version matching your language.

The base URL for API calls is: `https://<DP4CW_server_IP>:PORT/api`. We'll assume that all endpoints are prefixed with the base URL in the rest of this guide. In this guide we'll focus on the integration process, and skip actual requests and responses here and - check API docs for details.

Quite often in requests or responses, Data Protector for Cloud Workloads requires `NameAndGuid` to refer to other objects. `GUID` is the ID that you can use later to get additional information about hypervisor or backup. For convenience, we provide `name` to present it on the list views.

When you need to provide `NameAndGuid` in requests, you actually need to pass the object that has just `GUID` provided.

A similar concept applies to enum types - Data Protector for Cloud Workloads uses `EnumNameAndDescription` with `name` as enum name and `description` to show it to the end-user in a more user-friendly way. Anywhere when you are required to use enums such as type, state, days of the week, etc. you need to use `EnumNameAndDescription` object. In requests, you need to provide just `name`.

Login

When you want to invoke APIs, you need to be authenticated first. Data Protector for Cloud Workloads API exposes an endpoint that you need to send login and password first to the `POST /session/login` endpoint. Save cookies so that you'll be able to invoke the next calls.

If you ever call any end-point without having a valid session, you'll receive 401 Unauthorized responses. You need to re-login then and repeat your request.

In the rest of this guide, we assume that you have a valid session before you call any end-point mentioned.

Listing VMs in inventory (including already non-existing)

If you have a dedicated view of the VMs that are available for restore you need to call: `/virtual-machines`. This will retrieve all VMs visible by Data Protector for Cloud Workloads. `GUID` is the ID that you should refer to when invoking any operation on the VM. `UUID` is the ID that your infrastructure uses to identify objects.

For a multi-tenant environment make sure to filter out VMs on your side according to your ACLs or ownership of the VMs. For OpenStack Data Protector for Cloud Workloads records project ID and allow to use it to filter VMs like this: `/virtual-machines?tenantid={PROJECTID}`

Getting VM details

You may want to show VM details to the end-user. Call `GET /virtual-machines/{guid}`. Some useful information includes assigned policies, protection status or last backup sizes, and timestamps.

Performing backup of a VM on the list

The Backup consists of 2 phases in Data Protector for Cloud Workloads - export and store. End-user will initiate "backup", but your system should create an "export" task. Store task will be created automatically once export succeeds. To create an export task you need to provide the following information by the `POST /tasks/export` endpoint:

- `windowStart` and `windowEnd` - start and end of a time window for an export task - a task will fail if it is not started within this time range; both values provided as UNIX time in milliseconds
- `priority` - 0-100 - higher priority tasks are executed first; 50 by default
- `backupType` - `FULL` or `INCREMENTAL`; note that incremental backups are supported only for some platforms and they require at least one schedule of type incremental assigned to the policy that VM uses; if a snapshot for incremental backup is not found, a full backup will be done instead
- `backupDestination` - provided as `NameAndGuid` - target where the backup is going to be stored

- `protectedEntities` - collection of `NameAndGuid` referring to VMs that you want to backup - Data Protector for Cloud Workloads will create one export task for each referred VM

In the response, you'll receive task details, and you may want to record `GUID` to monitor later its progress or status. A new `backup` entry is going to be created automatically - you may also want to record this number if you want to present its details.

Progress monitoring of each operation

You can show the status and progress of each operation by calling `GET /tasks` (retrieves all tasks) and filtering the results or monitoring a particular task by calling `GET /tasks/{guid}`. There are also several useful query parameters that you can use to retrieve a filtered list:

- `protectedEntity` - GUID of VM that you refer to
- `backup` - GUID of a backup
- `schedule` - GUID of a schedule that invoked this task
- `state` - task state as `EnumNameAndDescription`: `QUEUED`, `RUNNING`, `FINISHED`, `FAILED`, `CANCELLED`
- `type` - task type as `EnumNameAndDescription`: `INDEX`, `EXPORT`, `STORE`, `RESTORE`, `OLD_BACKUPS_REMOVAL`, `OLD_SNAPSHOTS_REMOVAL`, `IMPORT`, `MOUNT`, `UNMOUNT`, `DELETE`, `SNAPSHOT`, `SNAPSHOT_REVERSION`
- `tenantId` - to filter out tasks only belonging to VMs owned by a particular tenantID (currently OpenStack only)

Browsing backup history

You can retrieve backup history including statuses, sizes, and time stats for each backup by calling `GET /backups/?protected-entity={guid}`, where you provide `guid` of your VM.

Restoring specified VM

Similar to backup, restore typically consists of several tasks. If you restore VM to the file system on the node, then it is just one task: `restore`. However, you usually want the user to restore and import VM automatically to the virtualization platform or mount it for file-level restore.

Let's focus on the restore with import case. You need to submit a task to `POST /tasks/restore-and-import` endpoint and provide:

- `backup` - GUID of a backup to be restored
- `hypervisor` or `hypervisorManager` - specify either one or the other - your target HV or HV manager depending on the virtualization platform
- `restoredPeName` - optional name of a restored VM
- `restoreStorageId` - some virtualization platforms require this to select the storage to which VM has to be restored;
- `restoreClusterId` - some virtualization platforms require this to select the cluster to which VM has to be restored;
- `restoreProject` - some virtualization platforms require this to select a project to which VM has to be restored
- `dataCenter` - some virtualization platforms require this to select datacenter to which VM has to be restored; this is a `DataCenterDTO` (currently having only `name` property)

Once the restore task completes, an import task will be created. Monitor tasks to show user current progress.

VM backup policy and schedule management

In order to allow users to have an automatic, scheduled backup you need to create a schedule and policy and make sure that both VMs and schedules are assigned to the policy. A Policy can have multiple schedules. VM can have exactly one backup policy assigned. Schedules are not assigned directly to the schedules - they are actually a part of a Policy Rule. And can be assigned to multiple rules, however in an external system this maybe not be convenient and we recommend using dedicated schedules for each policy. Currently, Data Protector for Cloud Workloads supports one rule per policy, but from API perspective - there is a collection of rules in the policy.

Schedules can be active or not, they have to specify which type of backup needs to be done, and when. Policies specify options to automatically assign VMs based on certain criteria. It is not recommended to expose these criteria to the end-user, as currently multi-tenancy support doesn't cover this case.

When you're exposing schedules and policies you may want to filter only these owned by a particular tenant. When you create/update them you need to provide additional `tenantId` to be able later to use it when listing objects.

Listing schedules

To list schedules and retrieve basic info use `GET /schedules` with optional query param `tenantid`.

Creating schedules

To create a **** schedule use `POST /schedules` endpoint and provide the following information:

- `name` - has to be globally unique, however, you can handle uniqueness on your site or generate names if you don't need to present them to end-user
- `backupType` - `EnumNameAndDescription` - a type of backup to perform: `FULL` or `INCREMENTAL`
- `type` - `EnumNameAndDescription` - for VM backup it is `VM_BACKUP` (other options are `APP_BACKUP` and `SNAPSHOT`); this type must match policy type
- `executionType` - `EnumNameAndDescription` - schedules can be executed at given `TIME` (based on `hour` field) or on `INTERVAL` basis
- `hour` - a time when schedule should be invoked - it is a time offset from UTC midnight in milliseconds, for example 3600 means 1:00 am UTC
- `active` - a boolean flag to activate or deactivate the schedule
- `startWindowLength` - used to assign window end to export tasks in milliseconds (which will be set to `hour` + `startWindowLength`)
- `daysOfWeek` - collection of `EnumNameAndDescription` - days of the week `MONDAY`, ..., `SUNDAY` when schedule needs to be run
- `months` - collection of `EnumNameAndDescription` - days of the week `JANUARY`, ..., `DECEMBER` when schedule needs to be run; if empty - only during specified

months schedule is executed

- `dayOfWeekOccurrences` - collection of `EnumNameAndDescription` - days of the week occurrences `FIRST_IN_MONTH`, `SECOND_IN_MONTH`, `THIRD_IN_MONTH`, `FOURTH_IN_MONTH`, `LAST_IN_MONTH`, when schedule needs to be run; if empty - only during first, ..., the last occurrence of specified days of week schedule is executed
- `rules` - collection of `NameAndGuid` - policy rules to which assign schedule to
- `interval` - object containing `startHour`, `endHour` (time offset from UTC midnight in milliseconds) and `frequency` (also in milliseconds)
- `tenantId` - string identifying tenant to which assign the schedule - this is used only by 3rd party system to filter out the listing

Getting schedule details

To get schedule details use `GET /schedules/{guid}`.

Updating schedule

To update the schedule use `PUT /schedules/{guid}` endpoint and provide the same information as in the creation request.

Deleting schedule

To delete the schedule use `DELETE /schedules/{guid}`.

Listing VM backup policies

To list policies and retrieve basic info use `GET /policies/vm-backup` with optional query param `tenantid`.

Creating VM backup policy

Policy creation is a 2 step process. It requires creating policy itself and then setting rules set on the policy. This ruleset currently must be a 1 element set.

To create a schedule use `POST /policies/vm-backup` endpoint and provide the following information:

- `name` - has to be globally unique, however, you can handle uniqueness on your site or generate names if you don't need to present them to end-user* `priority` - priority assigned to backup tasks (0-100)
- `autoRemoveNonPresent` - a boolean flag to automatically remove from policy non-existing VMs
- `autoAssignSettings` - object specifying details for VM auto-assignment mechanism - not supported in the multi-tenant environment; you need to set `mode` variable of this object to `DISABLED`
- `vms` - collection of `NameAndGuid` containing GUIDs of VMs to be assigned to the policy
- `tenantId` - string identifying tenant to which assign the policy - this is used only by 3rd party system to filter out the listing

You'll receive details including GUID of a newly created policy. The second step is to create a rule and assign it to the policy. Use `POST /rules/vm-backup` with the following information:

- `name` - we recommend generating it if not used by the user
- `schedules` - collection of `NameAndGuid` you want to be assigned to this policy
- `policy` - `NameAndGuid` of a policy that you want this rule to be assigned to.

Getting VM backup policy details

To get policy details use `GET /policies/vm-backup/{guid}`.

Updating VM backup policy

To update policy use `PUT /policies/vm-backup/{guid}` endpoint and provide the same information as in the creation request. Keep in mind, that you may need to update rules as well. You can either use `DELETE /rules/vm-backup/{guid}` and later `POST /rules/vm-backup` and re-create the policy rule or use `PUT /rules/vm-backup/{guid}` to update specific settings of a specific rule. GUID of a rule can be found in policy details (`rules` field).

Deleting VM backup policy

To delete policy use `DELETE /policies/vm-backup/{guid}`. Policy rules will be removed automatically.

Integration Plugins

Integration Plugins

Data Protector for Cloud Workloads can also be managed by using external plugins. Currently, there are several available, most of which support a subset of the available options of Data Protector for Cloud Workloads, with more and more functionalities being added with each release. Data Protector for Cloud Workloads currently provides integration for the UI

- Red Hat Virtualization
- oVirt
- Oracle Linux Virtualization Manager
- OpenStack Horizon

Red Hat Virtualization UI Plugin

Red Hat Virtualization UI Plugin

General

Integration with the Red Hat Virtualization interface allows administrators to perform most of the basic operations without logging into the Data Protector for Cloud Workloads dashboard.

After installation (which is described at the end of this article) you will see a new tab "VM Backups" in the RHV menu.

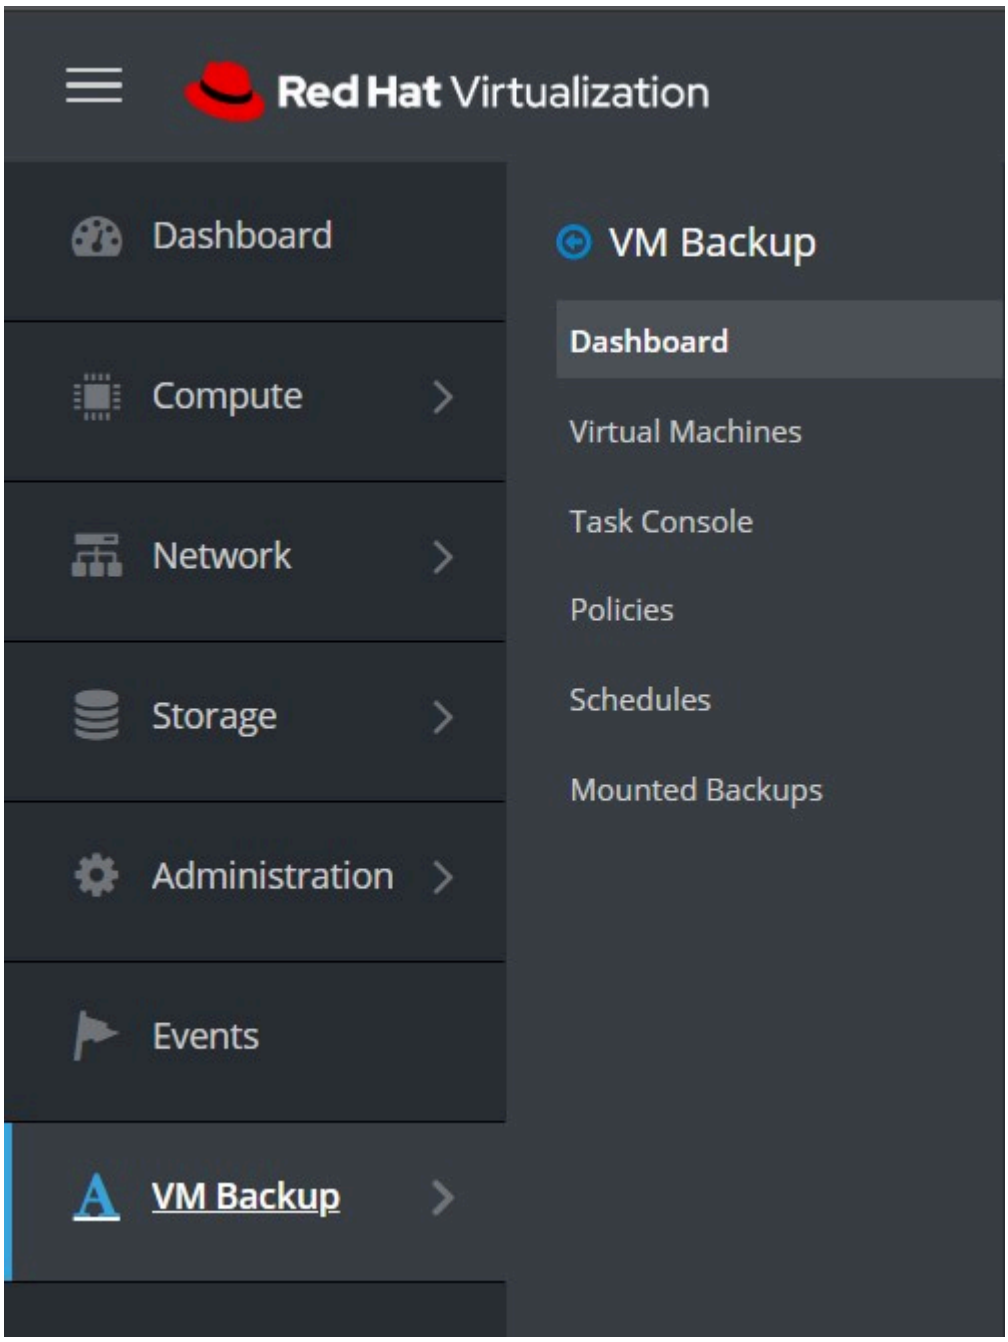

Dashboard

As usual, it contains a short summary of the environment along with a handful of statistics.

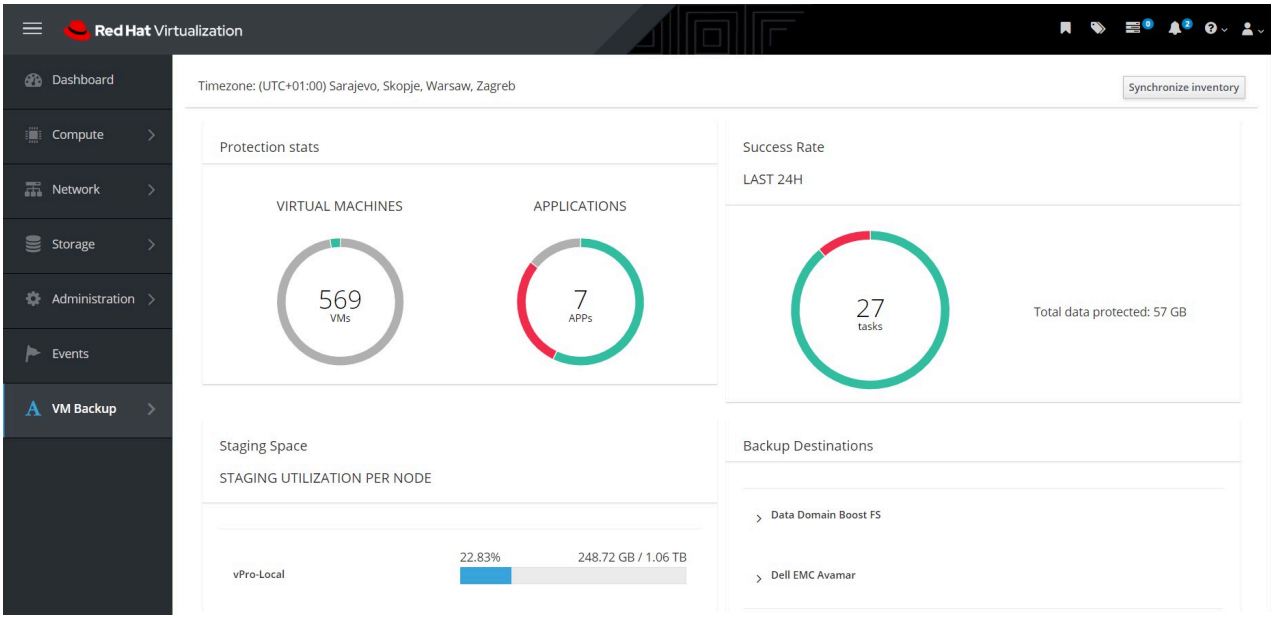

You can also see the data summary of the backup destination.

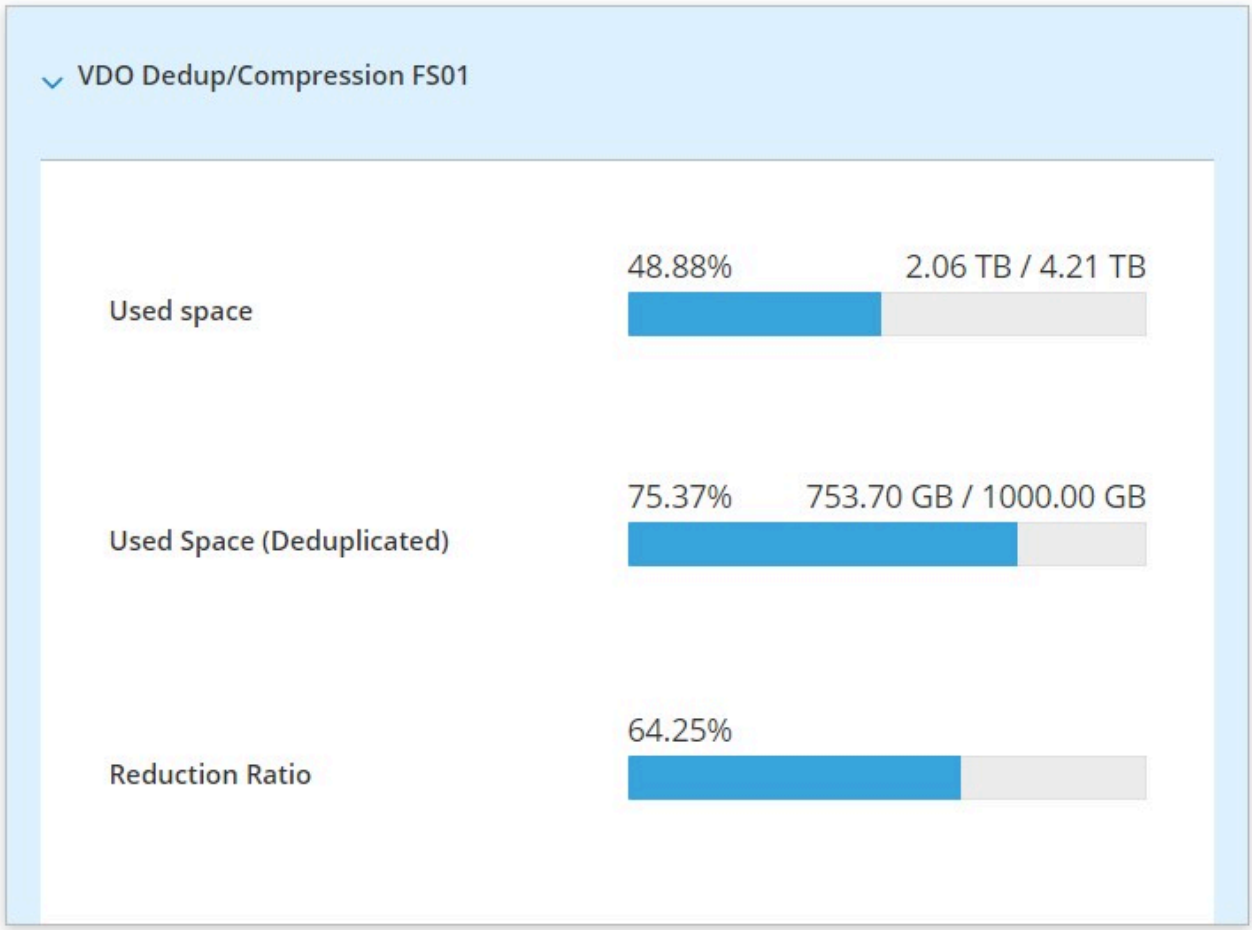

Virtual Machines

This tab shows all inventoried virtual machines in your RHV environment. In addition, you can also perform a basic backup or restore operations.

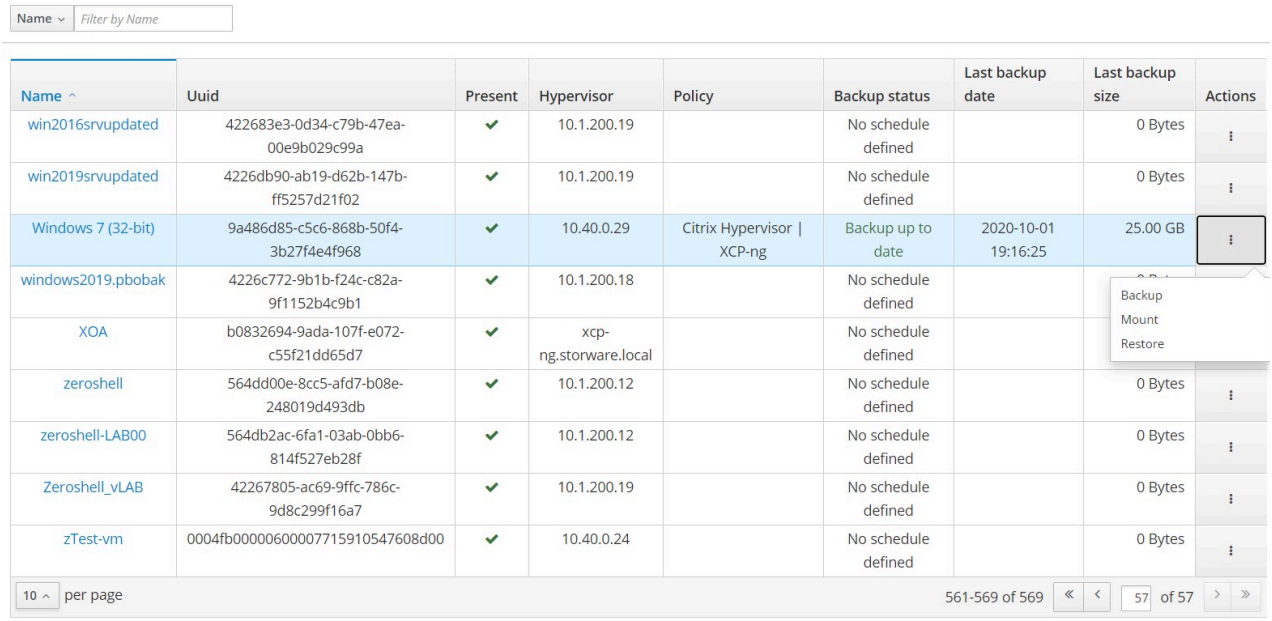

Name	Uuid	Present	Hypervisor	Policy	Backup status	Last backup date	Last backup size	Actions
win2016srvupdated	422683e3-0d34-c79b-47ea-00e9b029c99a	✔	10.1.200.19		No schedule defined		0 Bytes	
win2019srvupdated	4226db90-ab19-d62b-147b-ff5257d21f02	✔	10.1.200.19		No schedule defined		0 Bytes	
Windows 7 (32-bit)	9a486d85-c5c6-868b-50f4-3b27f4e4f968	✔	10.40.0.29	Citrix Hypervisor | XCP-ng	Backup up to date	2020-10-01 19:16:25	25.00 GB	
windows2019.pbobak	4226c772-9b1b-f24c-c82a-9f1152b4c9b1	✔	10.1.200.18		No schedule defined			
XOA	b0832694-9ada-107f-e072-c55f21dd65d7	✔	xcp-ng.storware.local		No schedule defined			
zeroshell	564dd00e-8cc5-afd7-b08e-248019d493db	✔	10.1.200.12		No schedule defined		0 Bytes	
zeroshell-LAB00	564db2ac-6fa1-03ab-0bb6-814f527eb28f	✔	10.1.200.12		No schedule defined		0 Bytes	
Zeroshell_vLAB	42267805-ac69-9ffc-786c-9d8c299f16a7	✔	10.1.200.19		No schedule defined		0 Bytes	
zTest-vm	0004fb0000060000771591054760 8d00	✔	10.40.0.24		No schedule defined		0 Bytes	

But that's not all, you can also go into the details of the virtual machine by clicking on its name:

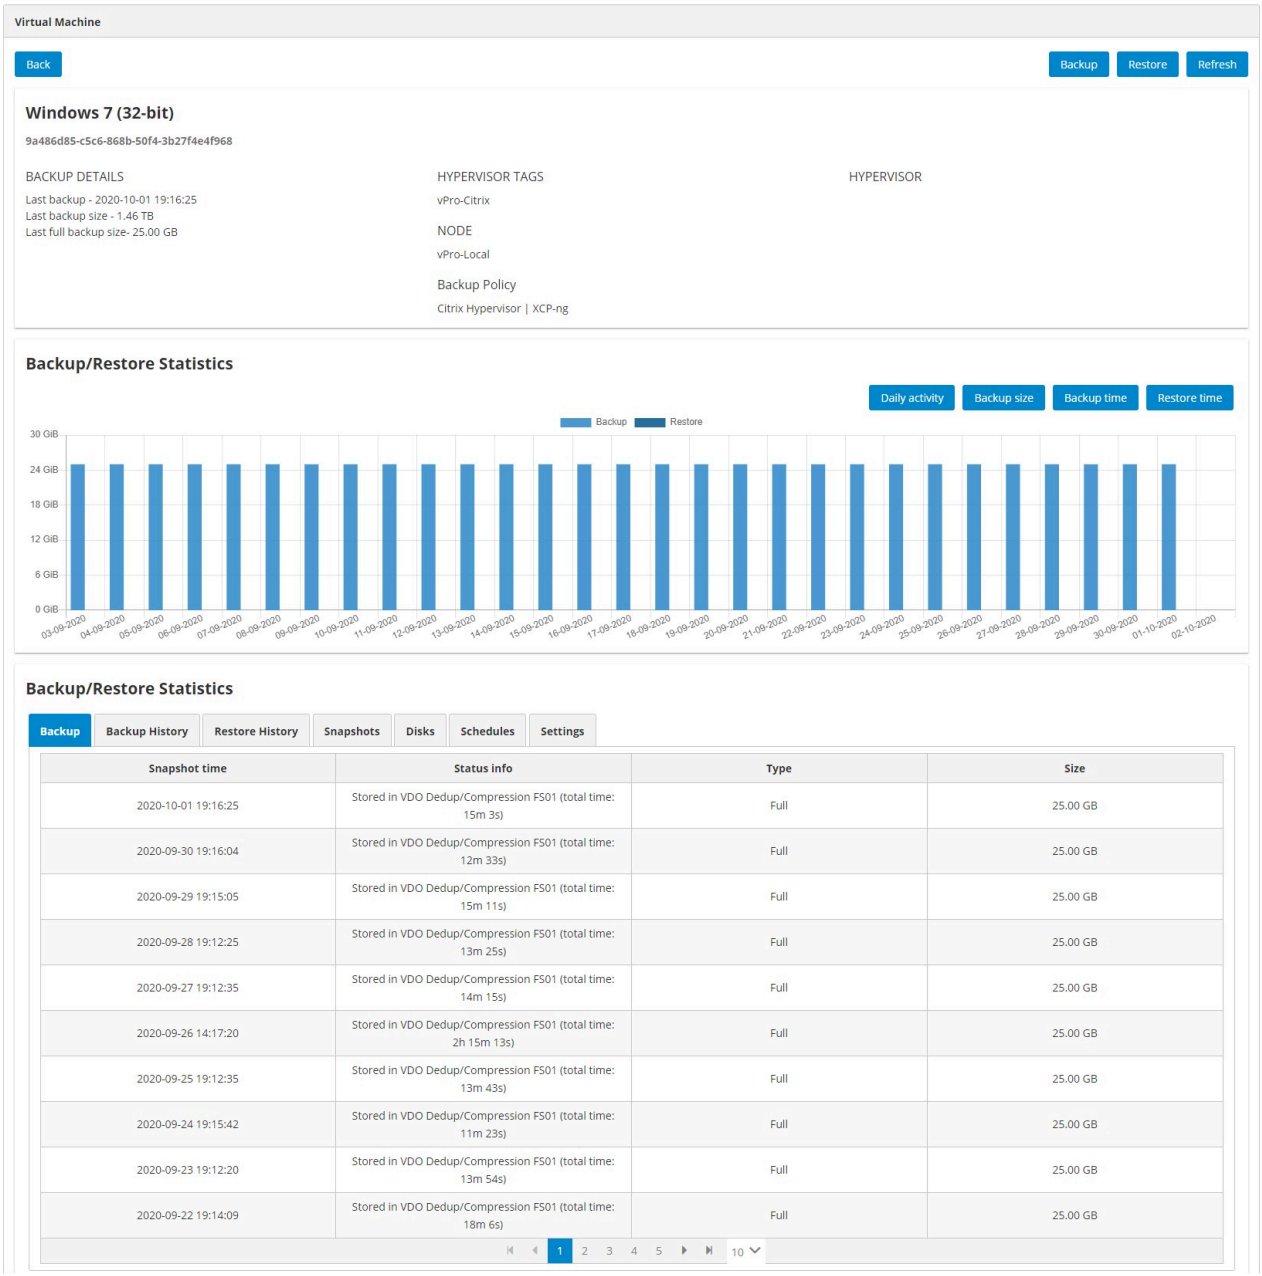

Backup window view:

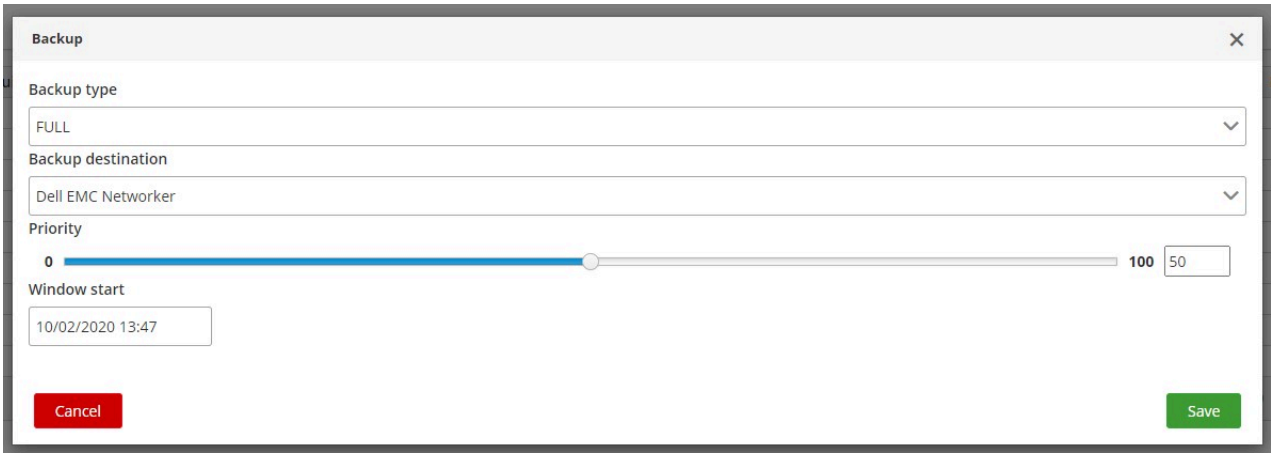

Task console

Basic information about current tasks performed by Data Protector for Cloud Workloads.

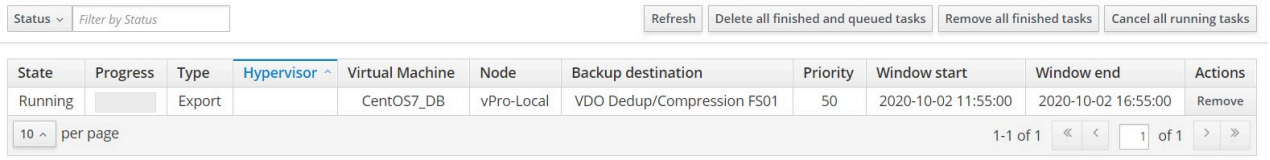

Policies

Allows you to create policies for VM's backups and snapshots. These basically have the same functionalities as Data Protector for Cloud Workloads WebUI.

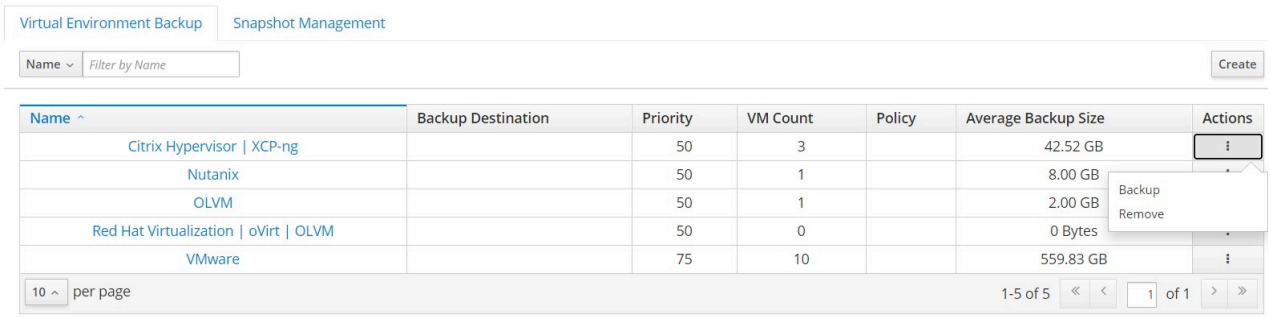

Schedules

As with the policies tab, it allows you to create schedules for the created rules.

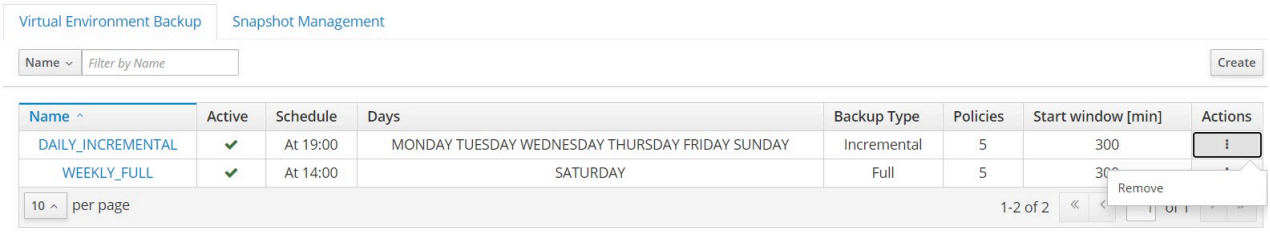

Mounted Backups

Finally, you can also browse your mounted backups from the RHV dashboard. You only need to enter the backup details using the menu on the right.

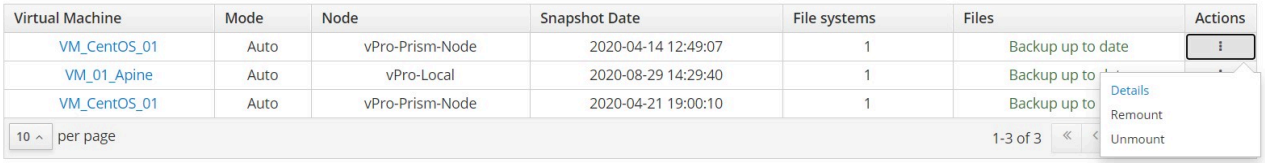

From here we can see the basic information about the backup and start browsing the files using the "Browse" button on the right.

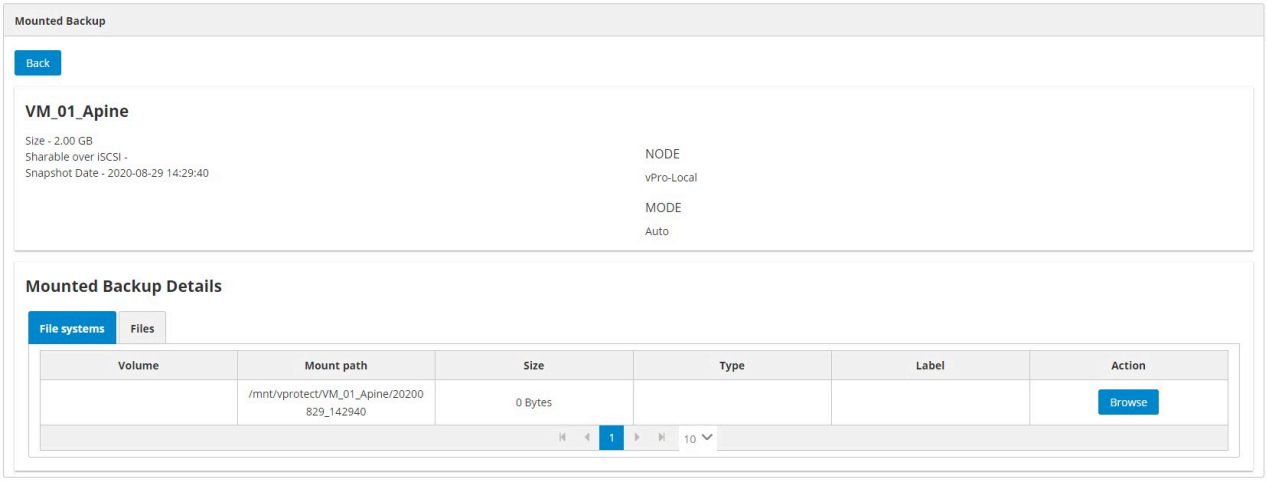

Just select a folder or file and then click on the "download selected" button to have the files on your computer.

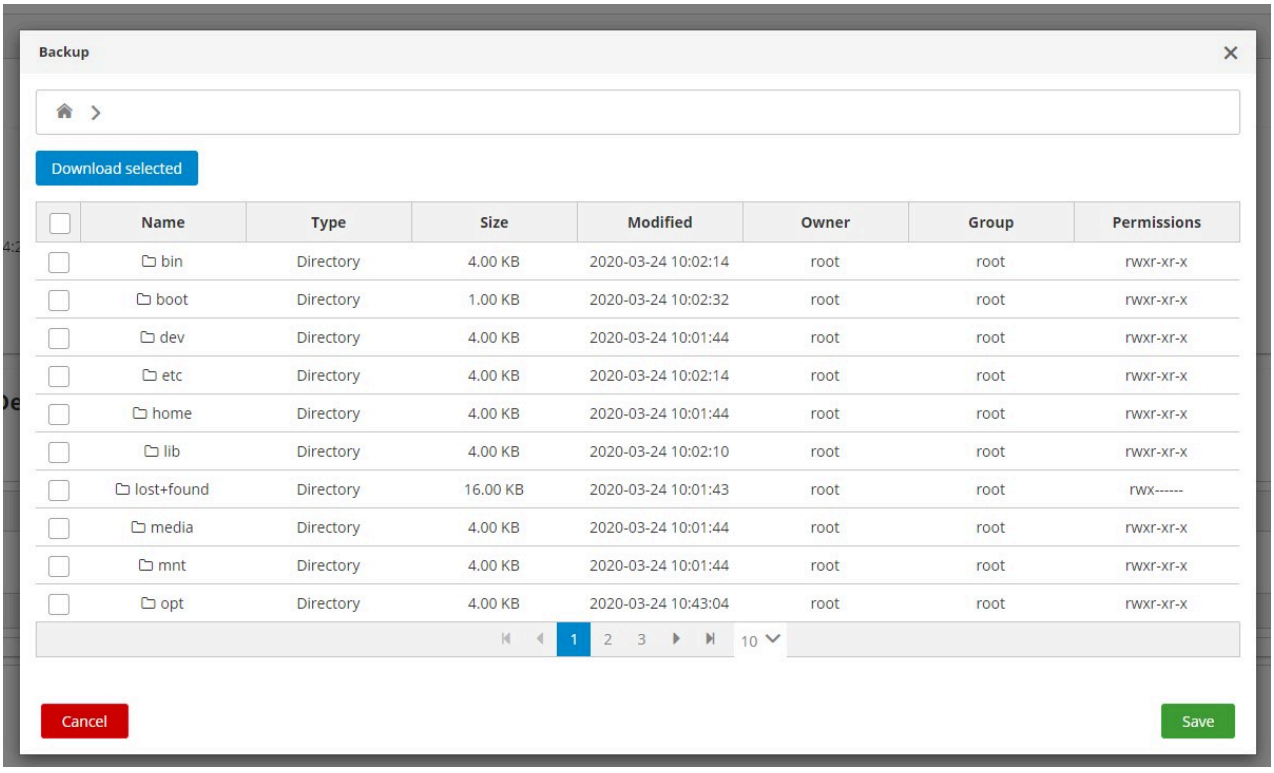

Installation

> You can find minimum requirements for Red Hat Virtualization UI Plugin in Support Matrix

1. You can find the add-on in the GitHub repository ↗. Follow the instructions in the README to build or the download plugin. Then extract the provided archive to your RHV manager.
2. In the file `vprotect.json` edit the lines in the config part:
   - `vProtectURL` - the URL to Data Protector for Cloud Workloads API
   - `username` - the name of the administrator in Data Protector for Cloud Workloads
   - `password` - the administrator password in Data Protector for Cloud Workloads

Example:

```
{
   "name": "vprotect",
   "url": "plugin/vprotect/plugin.html",
   "resourcePath": "vprotect-resources",
   "lazyLoad": false,

   "config": {
       "vProtectURL": "http://10.40.0.55:8080/api",
       "username": "admin",
       "password": "vPr0tect"
    }
}
```

1. Put the `vprotect.json` file and `vprotect-resources` directory in the `/usr/share/ovirt-engine/ui-plugins` directory in the RHV engine.

oVirt UI Plugin

oVirt UI Plugin

General

Integration with the oVirt Virtualization interface allows administrators to perform most of the basic operations without logging into the Data Protector for Cloud Workloads dashboard.

After installation (which is described at the end of this article) you will see a new tab "VM Backups" in the oVirt menu.

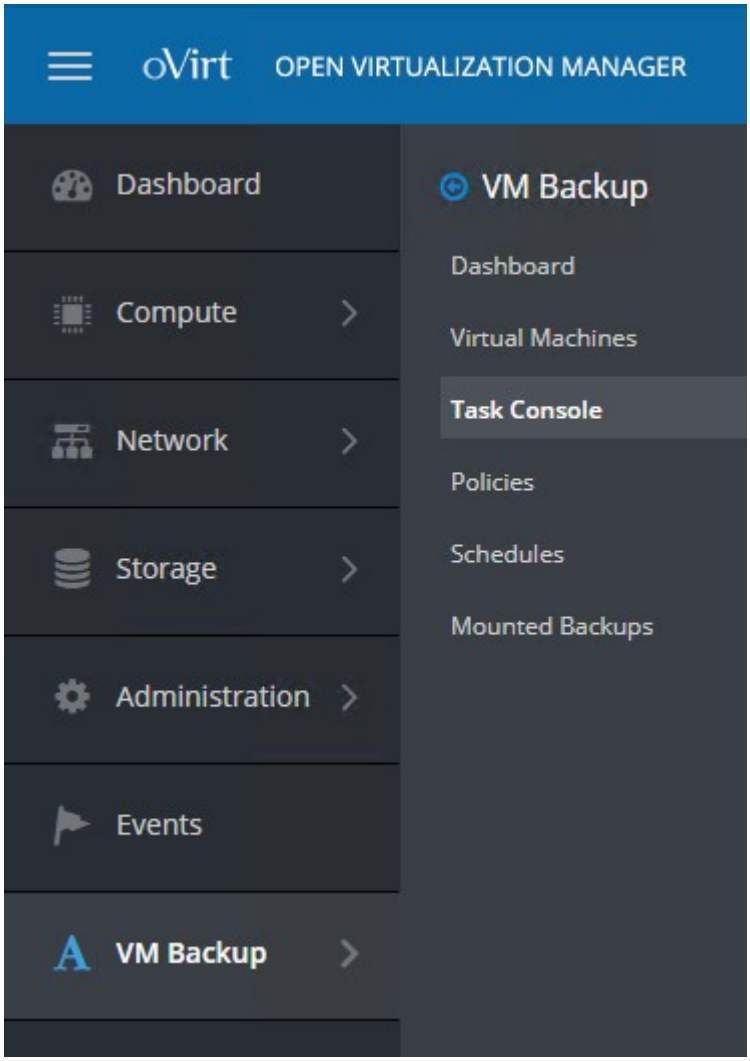

Dashboard

As usual, it contains a short summary of the environment along with a handful of statistics.

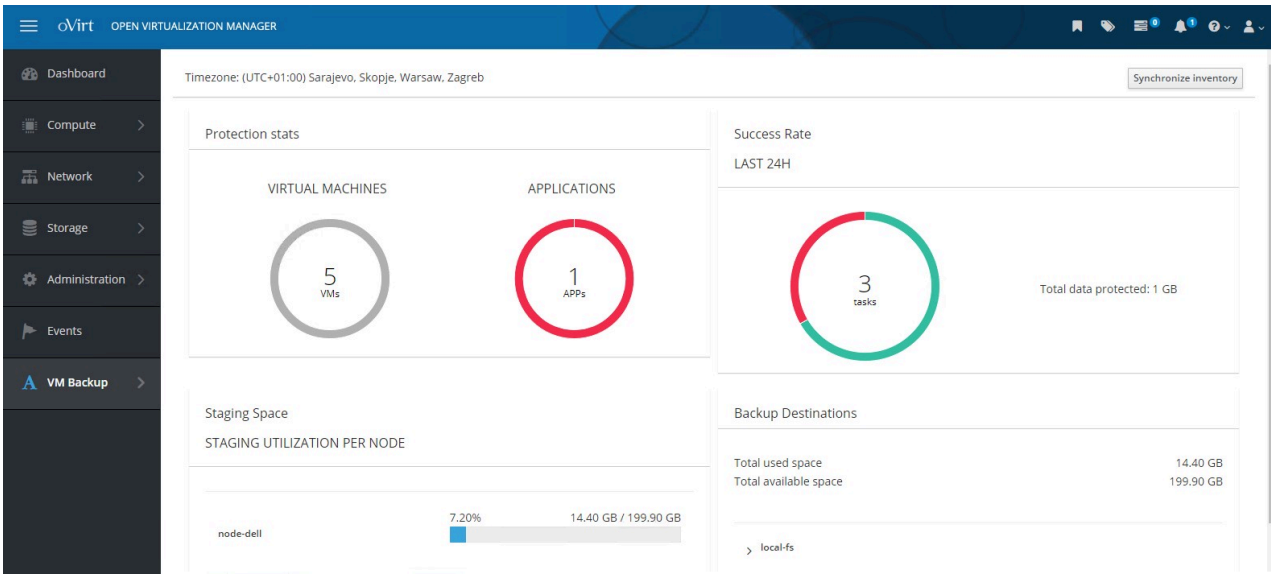

You can also see the data summary of the backup destination.

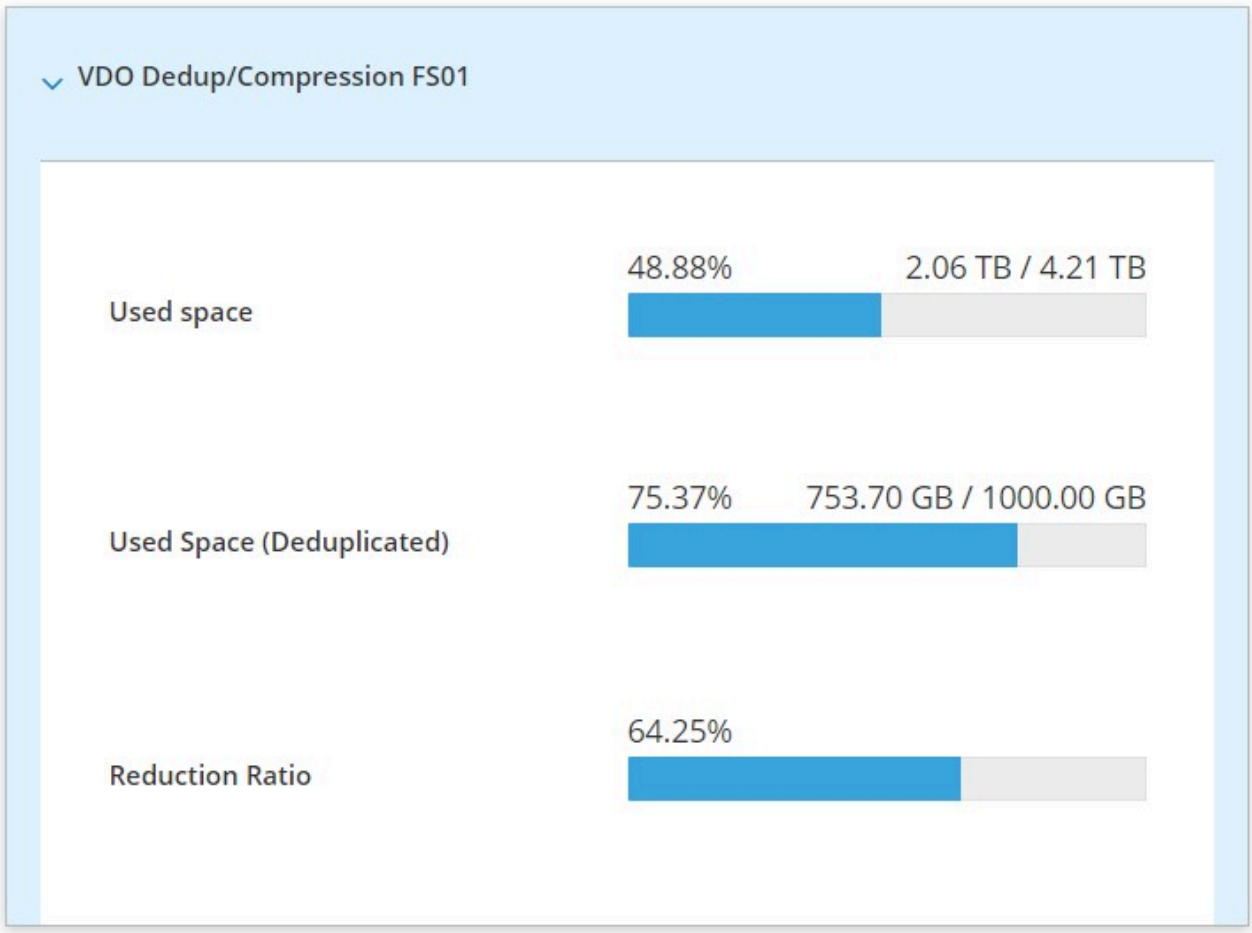

Virtual Machines

This tab shows all inventoried virtual machines in your oVirt environment. Besides, you can also perform a basic backup or restore operations.

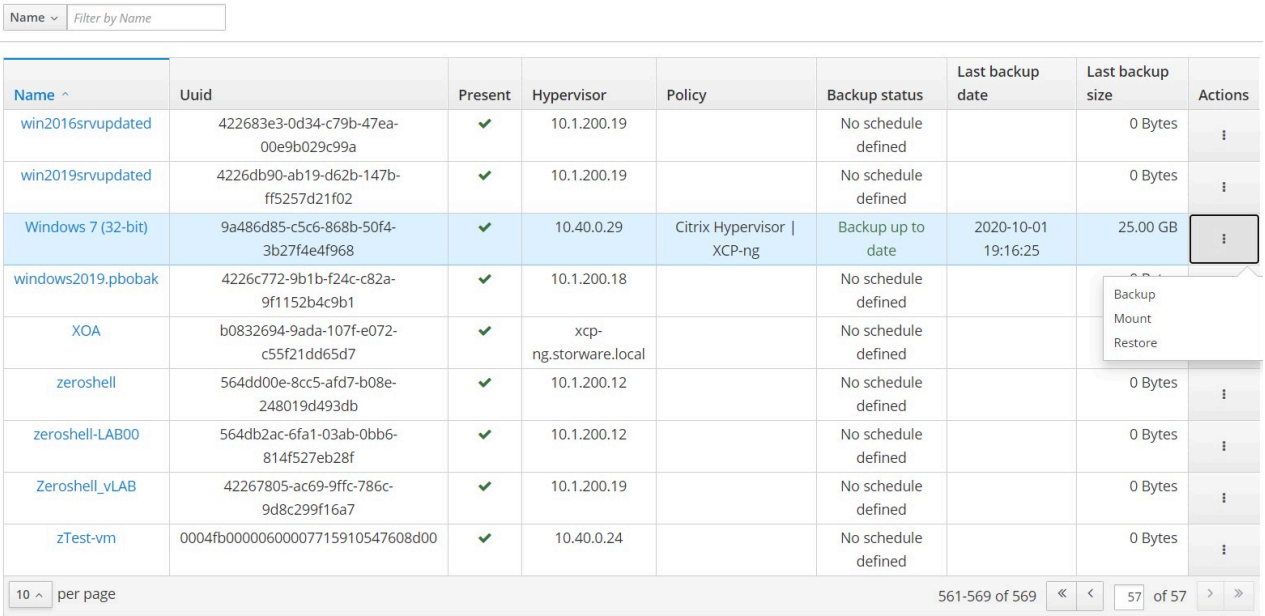

Name	Uuid	Present	Hypervisor	Policy	Backup status	Last backup date	Last backup size	Actions
win2016srvupdated	422683e3-0d34-c79b-47ea-00e9b029c99a	✔	10.1.200.19		No schedule defined		0 Bytes	
win2019srvupdated	4226db90-ab19-d62b-147b-ff5257d21f02	✔	10.1.200.19		No schedule defined		0 Bytes	
Windows 7 (32-bit)	9a486d85-c5c6-868b-50f4-3b27f4e4f968	✔	10.40.0.29	Citrix Hypervisor | XCP-ng	Backup up to date	2020-10-01 19:16:25	25.00 GB	
windows2019.pbobak	4226c772-9b1b-f24c-c82a-9f1152b4c9b1	✔	10.1.200.18		No schedule defined			
XOA	b0832694-9ada-107f-e072-c55f21dd65d7	✔	xcp-ng.storware.local		No schedule defined			
zeroshell	564dd00e-8cc5-afd7-b08e-248019d493db	✔	10.1.200.12		No schedule defined		0 Bytes	
zeroshell-LAB00	564db2ac-6fa1-03ab-0bb6-814f527eb28f	✔	10.1.200.12		No schedule defined		0 Bytes	
Zeroshell_vLAB	42267805-ac69-9ffc-786c-9d8c299f16a7	✔	10.1.200.19		No schedule defined		0 Bytes	
zTest-vm	0004fb0000060000771591054760 8d00	✔	10.40.0.24		No schedule defined		0 Bytes	

But that's not all, you can also go into the details of the virtual machine by clicking on its name:

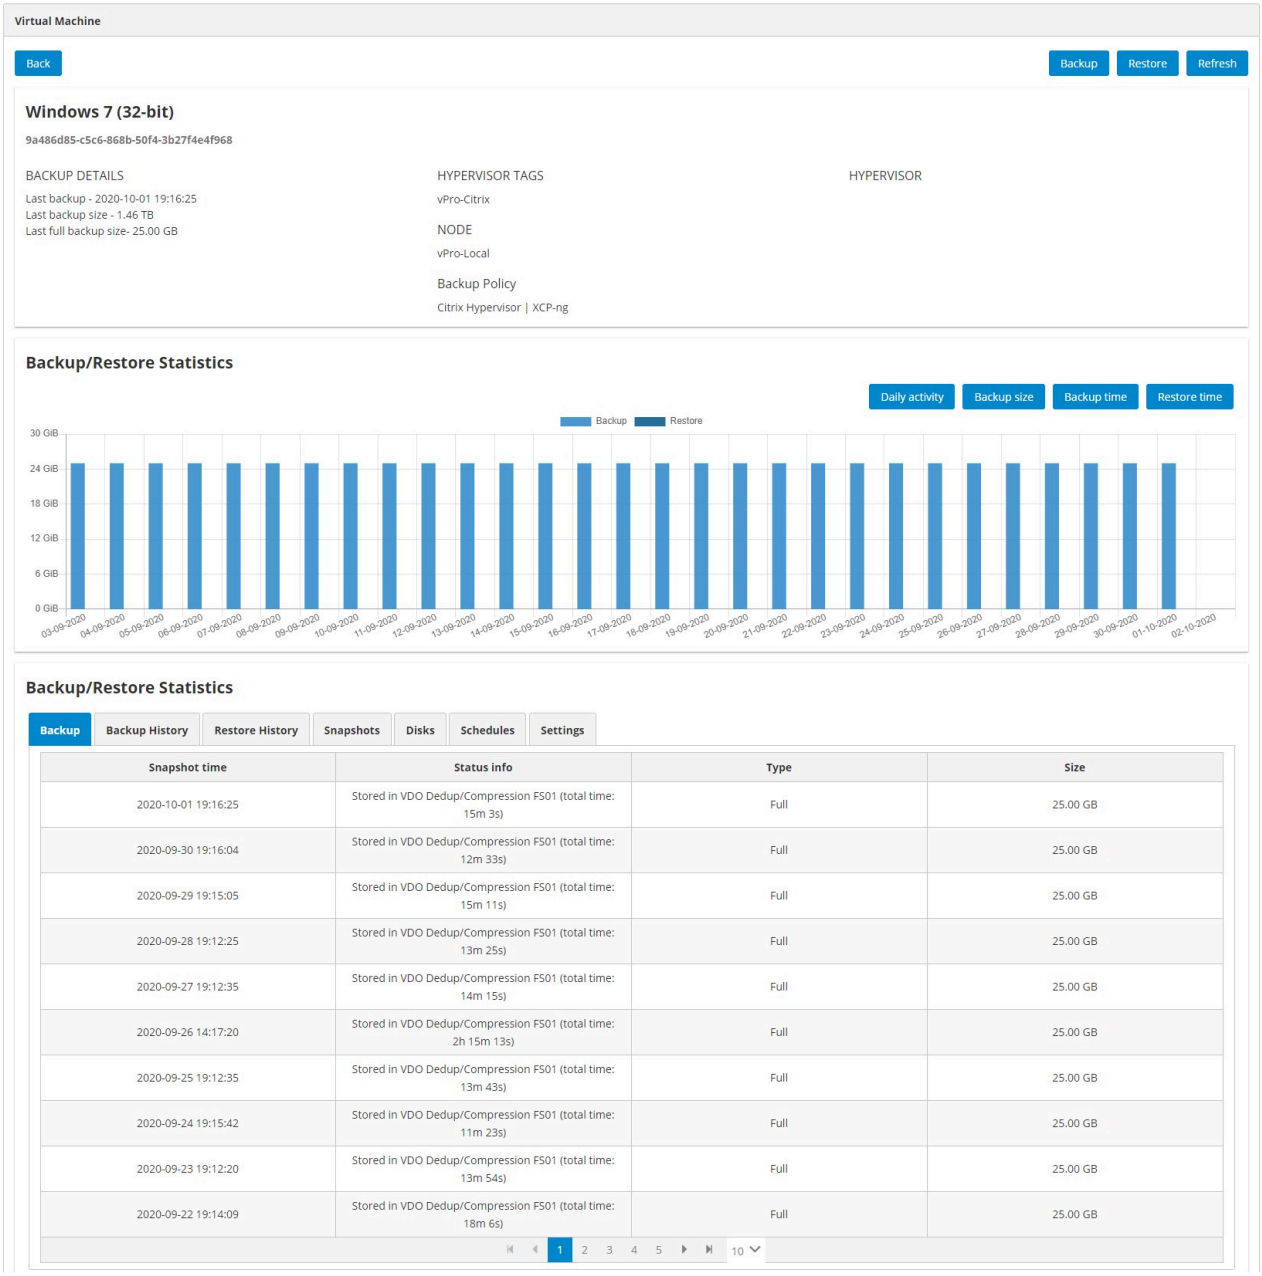

Backup window view:

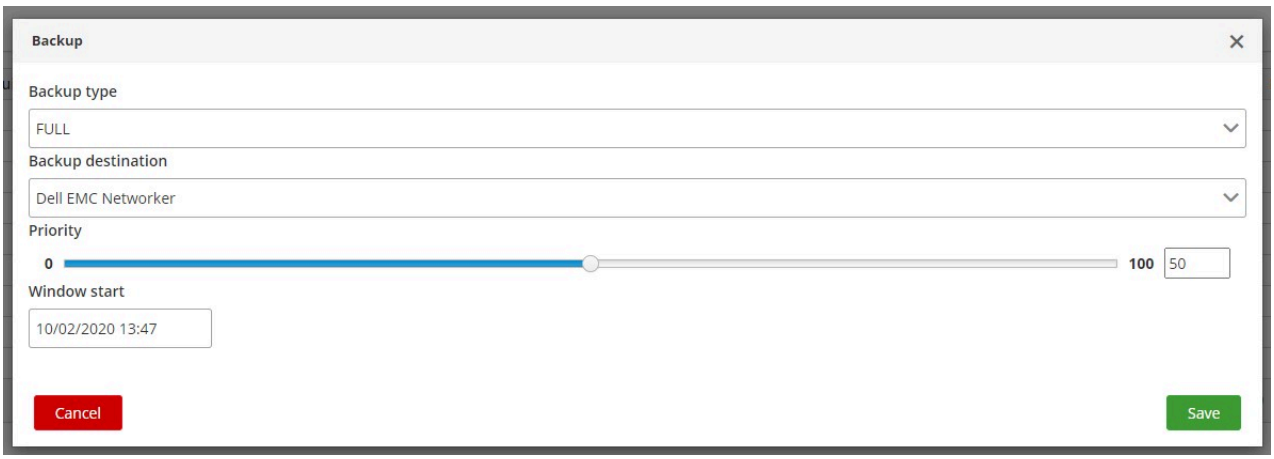

Task console

Basic information about current tasks performed by Data Protector for Cloud Workloads.

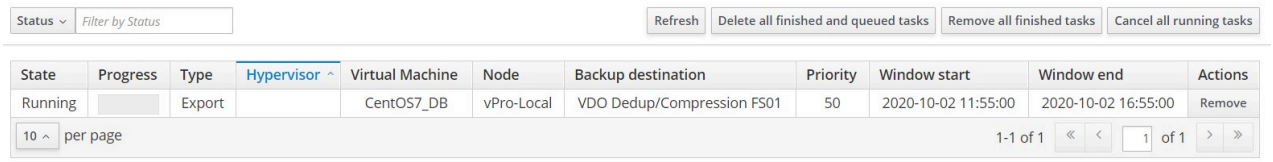

Policies

Allows you to create policies for VM's backups and snapshots. It basically has the same functionalities as Data Protector for Cloud Workloads WebUI.

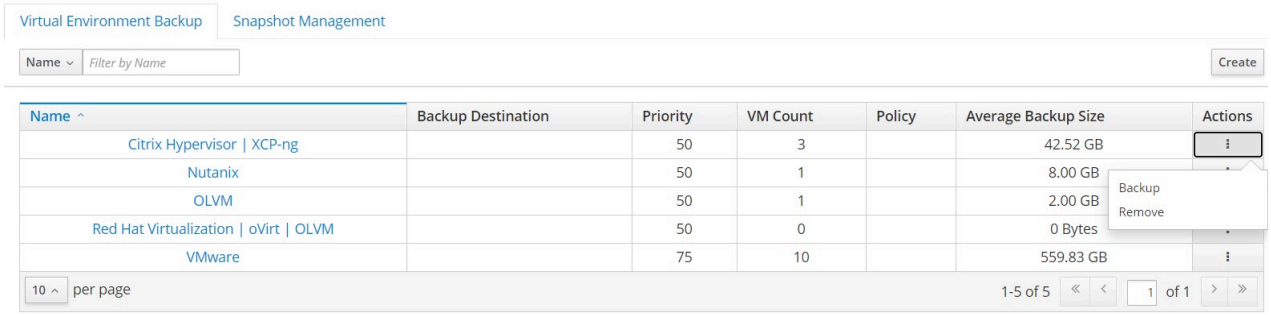

Schedules

As with the policies tab, it allows you to create schedules for the created rules.

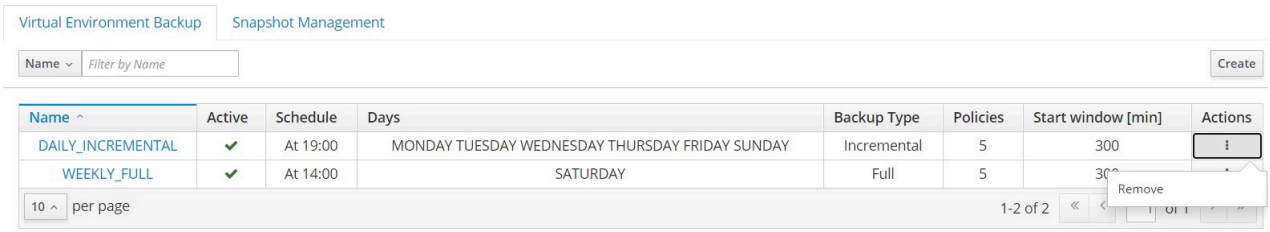

Mounted Backups

Finally, you can also browse your mounted backups from the oVirt dashboard. You only need to enter the backup details using the menu on the right.

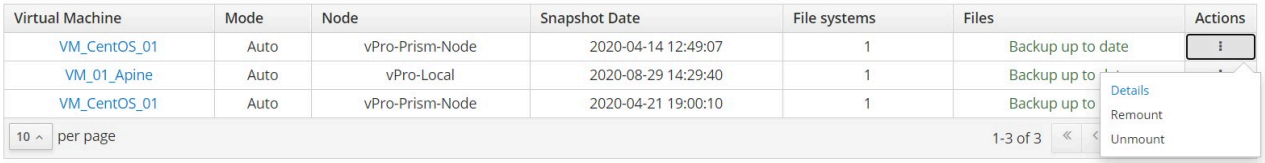

From here we can see the basic information about the backup and start browsing the files using the "Browse" button on the right.

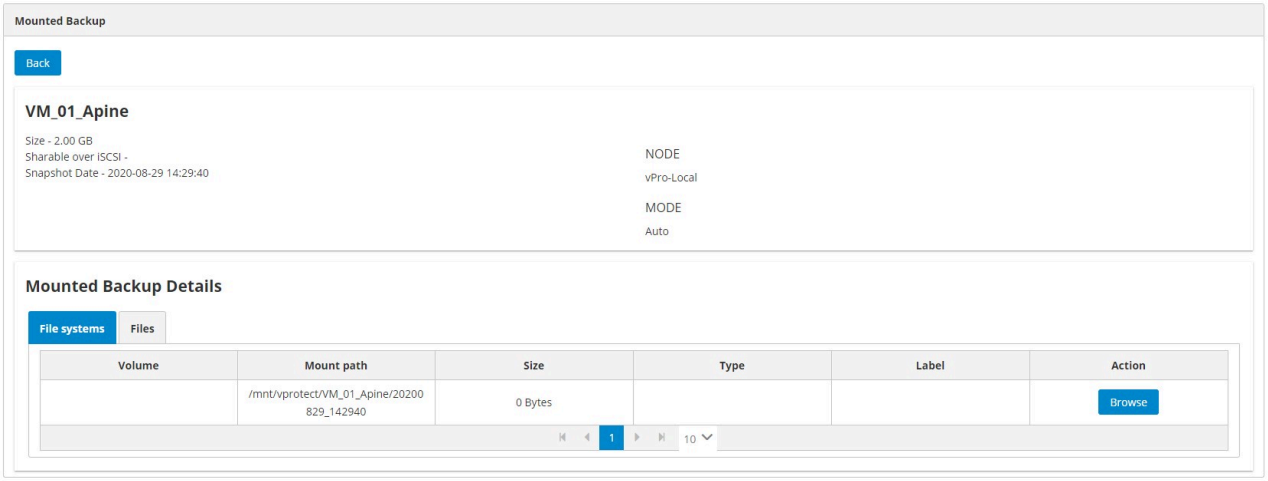

Just select a folder or file and then press the "download selected" button to have the files on your computer.

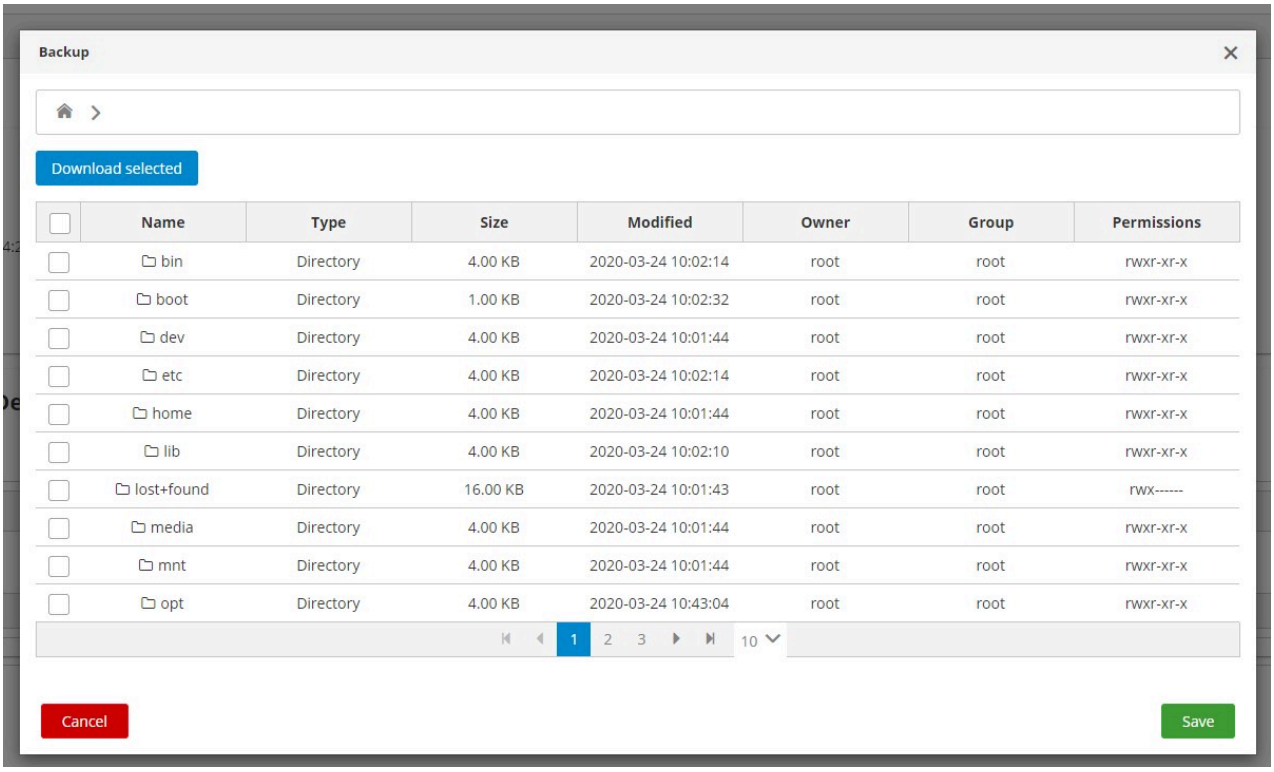

Installation

> You can find minimum requirements for oVirt UI Plugin in Support Matrix

1. You can find the add-on in the GitHub repository ↗. Follow the instructions in the README to build or download the plugin. Then extract the provided archive onto your oVirt manager.
2. In the file `vprotect.json` edit these lines in the config part:
   - `vProtectURL` - the URL to Data Protector for Cloud Workloads API
   - `username` - the name of admin in Data Protector for Cloud Workloads
   - `password` - the admin password in Data Protector for Cloud Workloads

Example:

```
{
   "name": "vprotect",
   "url": "plugin/vprotect/plugin.html",
   "resourcePath": "vprotect-resources",
   "lazyLoad": false,

   "config": {
       "vProtectURL": "http://10.40.0.55:8080/api",
       "username": "admin",
       "password": "vPr0tect"
    }
}
```

1. Put the `vprotect.json` file and `vprotect-resources` directory in the `/usr/share/ovirt-engine/ui-plugins` directory in the oVirt Engine.

Oracle Linux Virtualization Manager UI Plugin

Oracle Linux Virtualization Manager UI Plugin

General

Integration with the OLVM Virtualization interface allows administrators to perform most of the basic operations without logging into the Data Protector for Cloud Workloads dashboard.

After installation (which is described at the end of this article) you will see a new tab "VM Backups" in the OLVM menu.

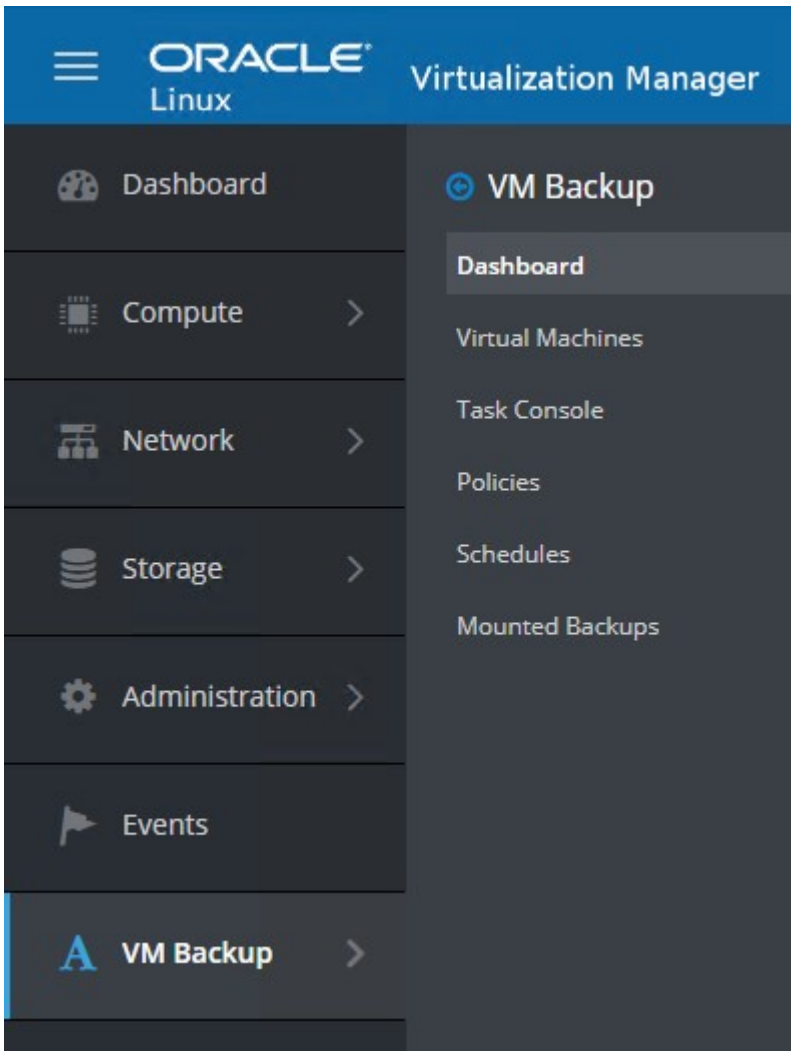

Dashboard

As usual, it contains a short summary of the environment along with a handful of statistics.

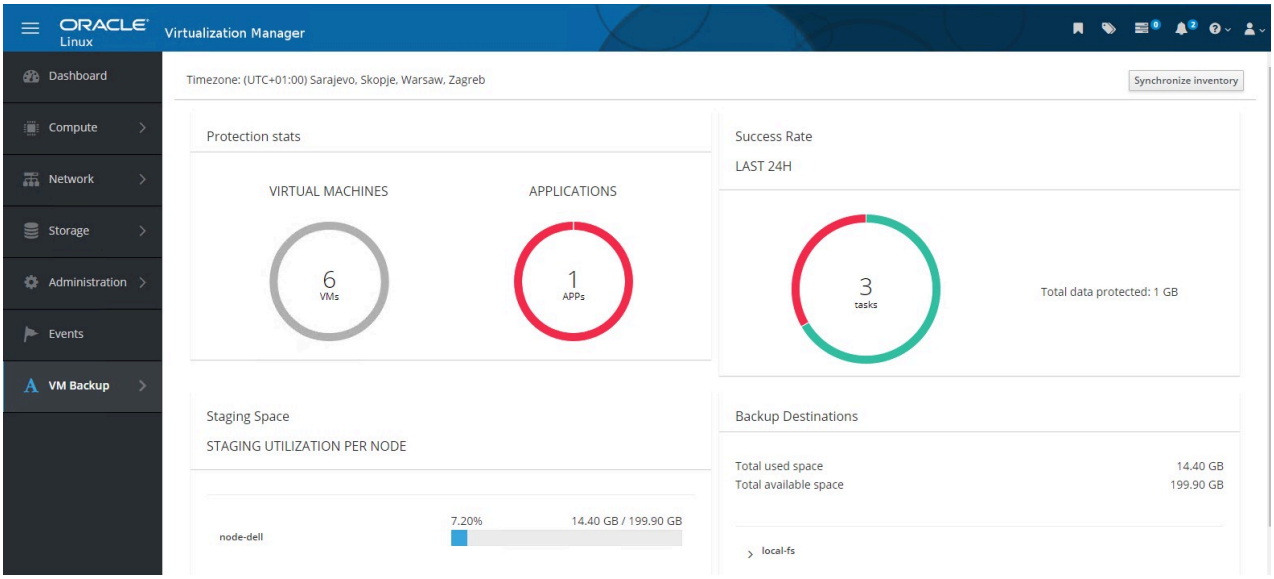

You can also see the data summary of the backup destination.

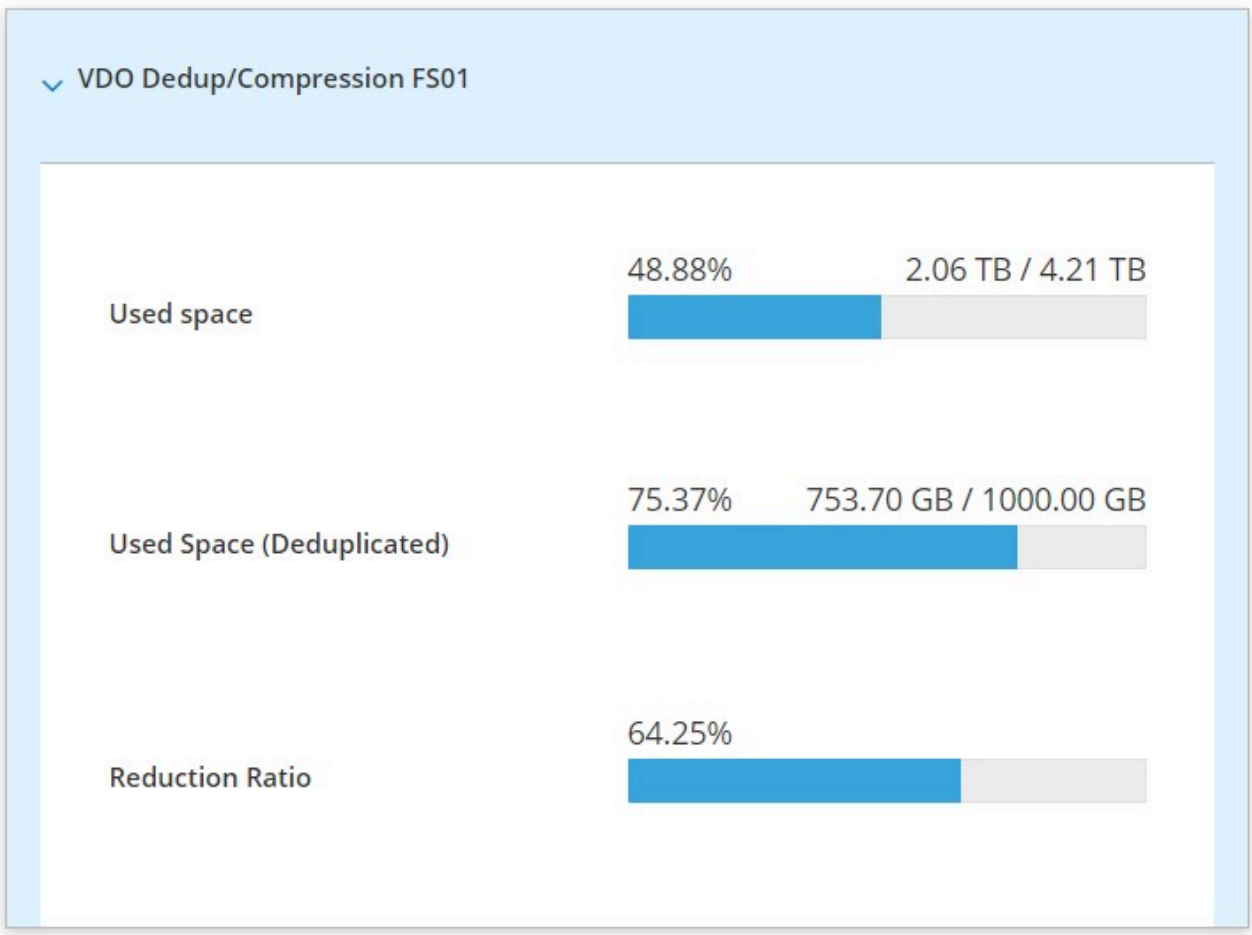

Virtual Machines

This tab shows all inventoried virtual machines in your OLVM environment. Besides, you can also perform a basic backup or restore operations.

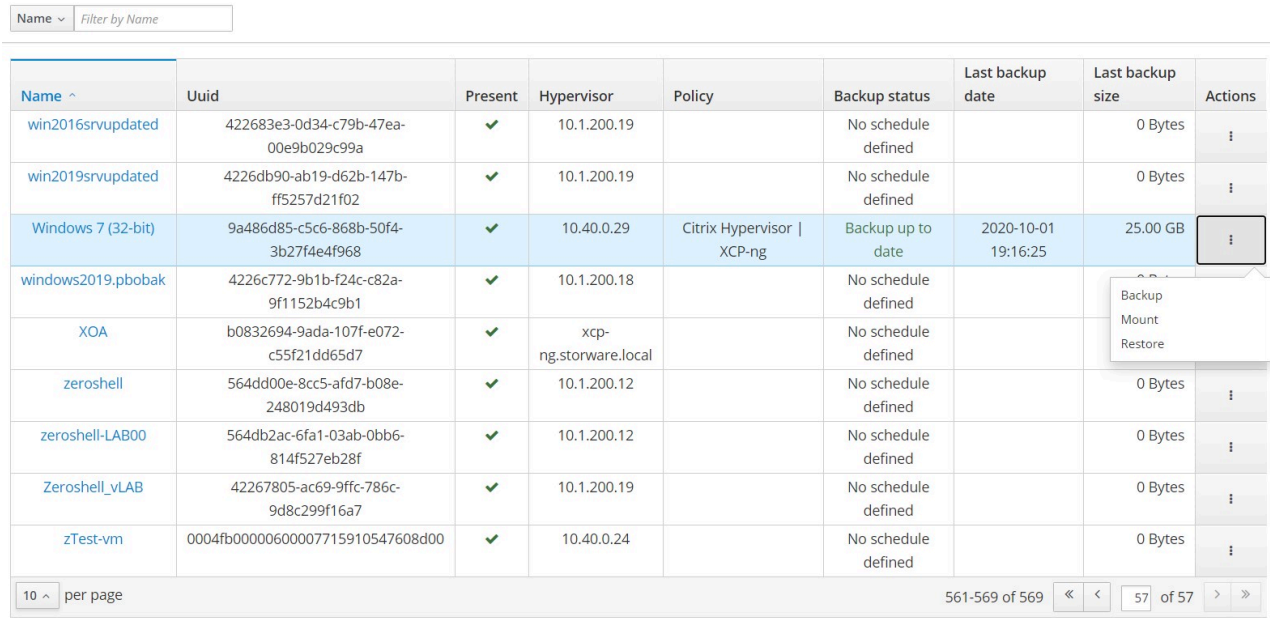

Name	Uuid	Present	Hypervisor	Policy	Backup status	Last backup date	Last backup size	Actions
win2016srvupdated	422683e3-0d34-c79b-47ea-00e9b029c99a	✔	10.1.200.19		No schedule defined		0 Bytes	
win2019srvupdated	4226db90-ab19-d62b-147b-ff5257d21f02	✔	10.1.200.19		No schedule defined		0 Bytes	
Windows 7 (32-bit)	9a486d85-c5c6-868b-50f4-3b27f4e4f968	✔	10.40.0.29	Citrix Hypervisor | XCP-ng	Backup up to date	2020-10-01 19:16:25	25.00 GB	
windows2019.pbobak	4226c772-9b1b-f24c-c82a-9f1152b4c9b1	✔	10.1.200.18		No schedule defined			
XOA	b0832694-9ada-107f-e072-c55f21dd65d7	✔	xcp-ng.storware.local		No schedule defined			
zeroshell	564dd00e-8cc5-afd7-b08e-248019d493db	✔	10.1.200.12		No schedule defined		0 Bytes	
zeroshell-LAB00	564db2ac-6fa1-03ab-0bb6-814f527eb28f	✔	10.1.200.12		No schedule defined		0 Bytes	
Zeroshell_vLAB	42267805-ac69-9ffc-786c-9d8c299f16a7	✔	10.1.200.19		No schedule defined		0 Bytes	
zTest-vm	0004fb0000060000771591054760 8d00	✔	10.40.0.24		No schedule defined		0 Bytes	

But that's not all, you can also go into the details of the virtual machine by clicking on its name:

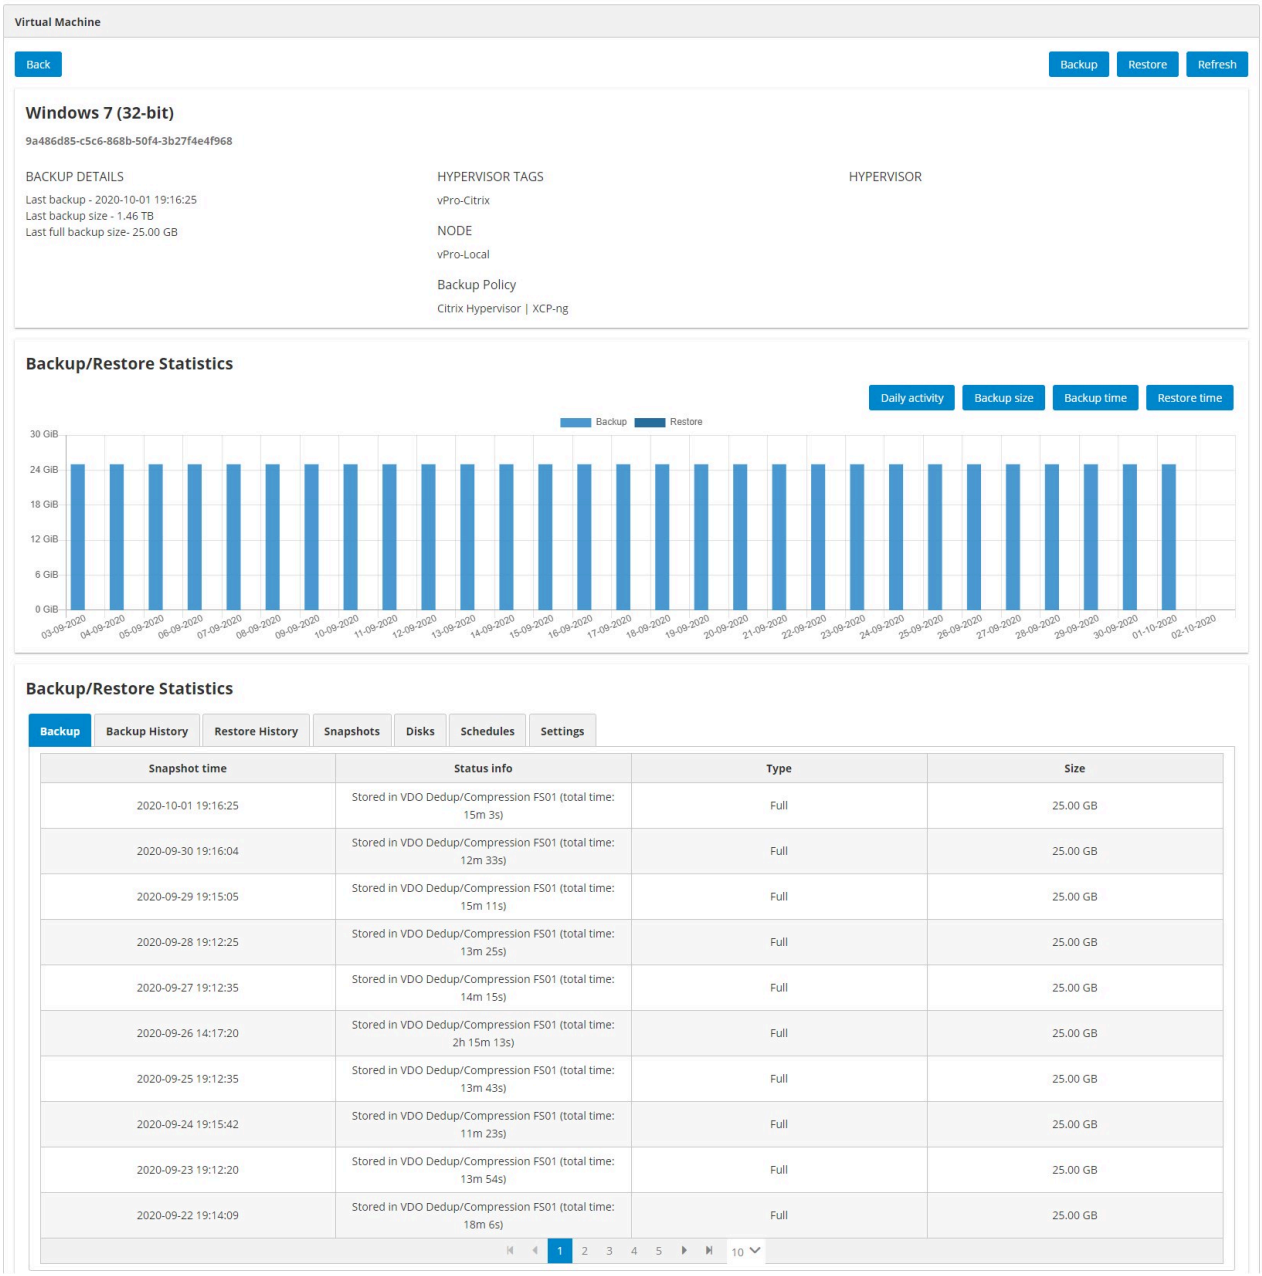

Backup window view:

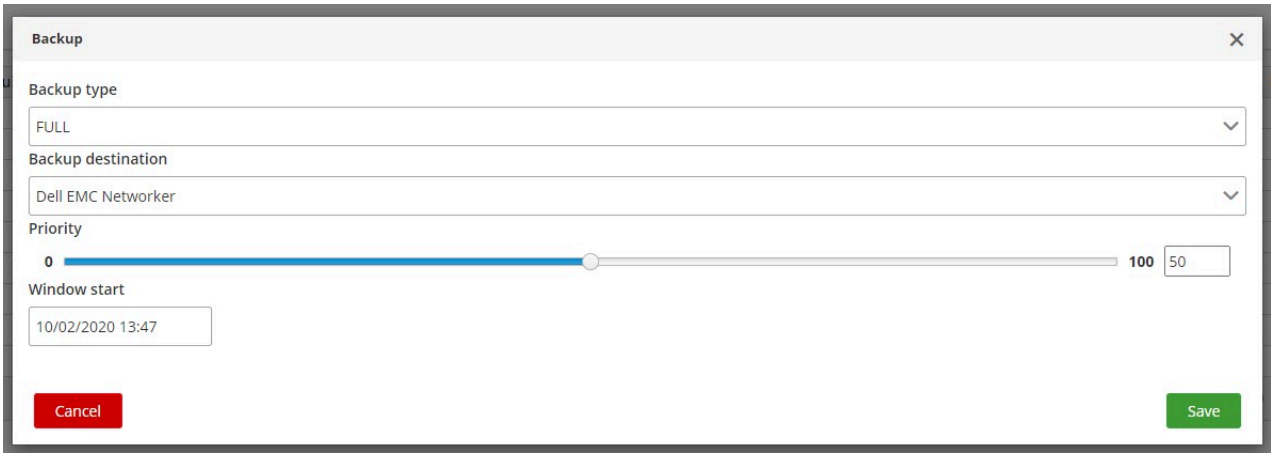

Task console

Basic information about current tasks performed by Data Protector for Cloud Workloads.

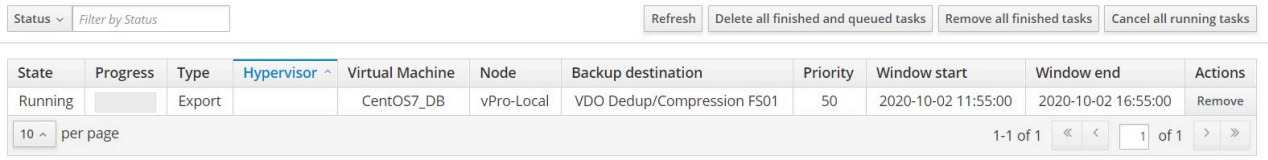

Policies

Allows you to create policies for VM's backups and snapshots. It basically has the same functionalities as Data Protector for Cloud Workloads WebUI.

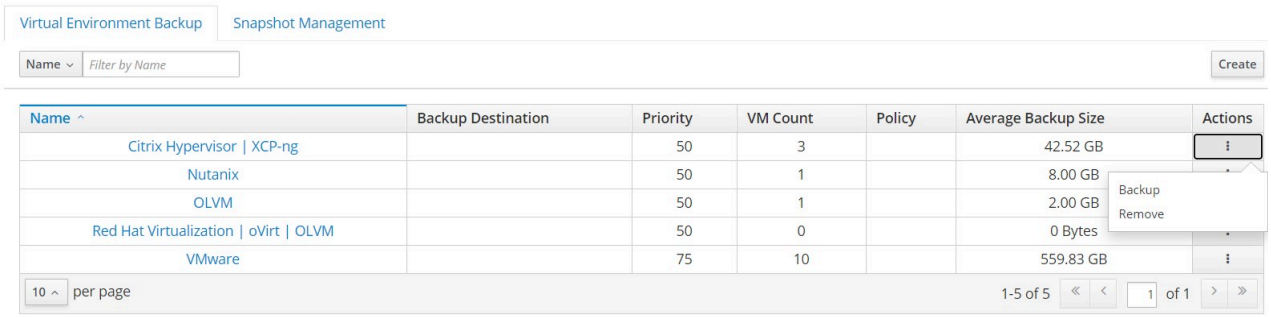

Schedules

As with the policies tab, it allows you to create schedules for the created rules.

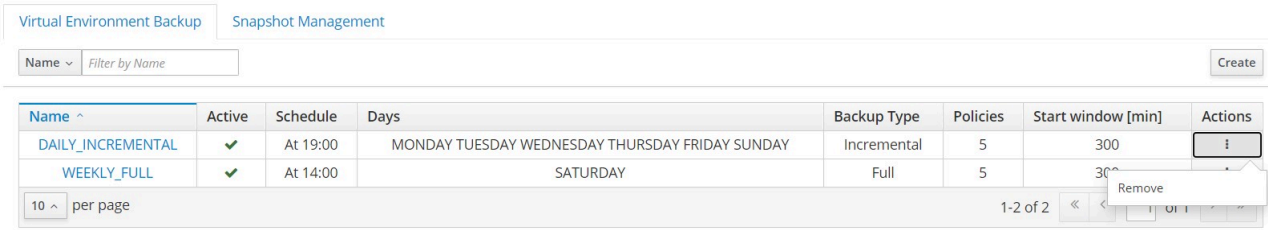

Mounted Backups

Finally, you can also browse your mounted backups from the OLVM dashboard. You only need to enter the backup details using the menu on the right.

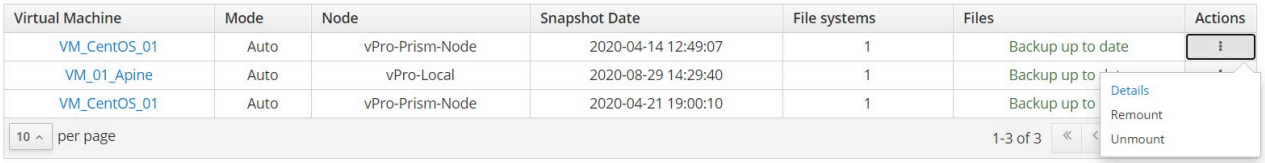

From here we can see the basic information about the backup and start browsing the files using the "Browse" button on the right.

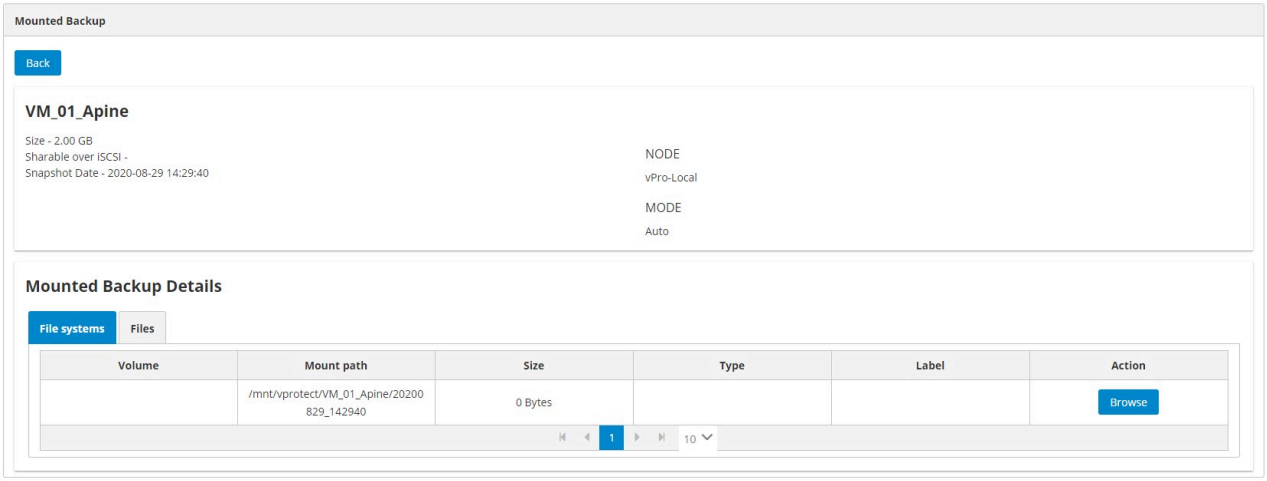

Just select a folder or file and then press the "download selected" button to have the files on your computer.

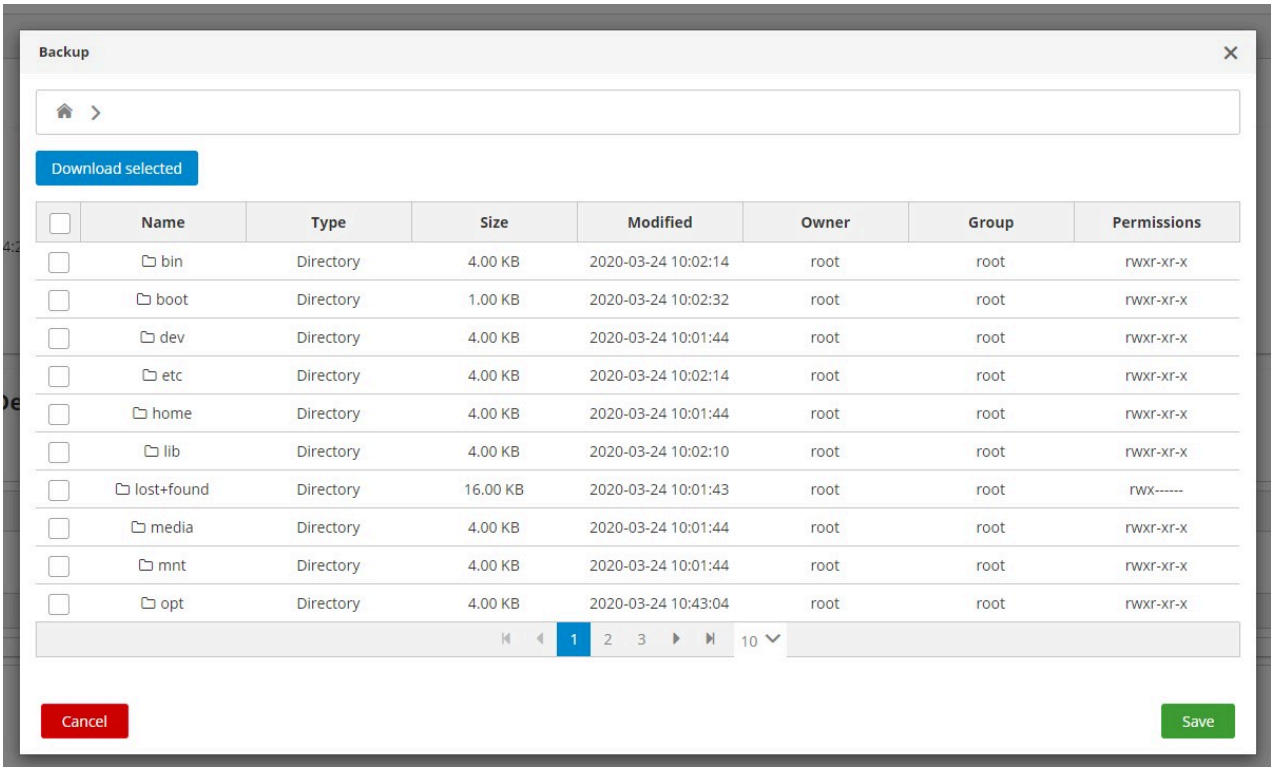

Installation

> You can find minimum requirements for Oracle Linux Virtualization Manager UI Plugin in Support Matrix

1. You can find the add-on in the GitHub repository ↗. Follow the instructions in the README to build or download the plugin. Then extract the provided archive onto your Oracle Linux virtualization manager.
2. In the file `vprotect.json` edit the lines in the config part:
   - `vProtectURL` - the URL to Data Protector for Cloud Workloads API
   - `username` - the name of admin in Data Protector for Cloud Workloads
   - `password` - the admin password in Data Protector for Cloud Workloads

Example:

```
{
   "name": "vprotect",
   "url": "plugin/vprotect/plugin.html",
   "resourcePath": "vprotect-resources",
   "lazyLoad": false,

   "config": {
       "vProtectURL": "http://10.40.0.55:8080/api",
       "username": "admin",
       "password": "vPr0tect"
    }
}
```

1. Put the `vprotect.json` file and `vprotect-resources` directory in the `/usr/share/ovirt-engine/ui-plugins` directory in the OLVM engine.

OpenStack UI Plugin

Overview

Integration with the Openstack interface is our second plugin alongside the oVirt/RHV virtualization family. Thanks to it, you can perform most of the basic operations without logging into the Data Protector for Cloud Workloads dashboard.

After installation (which is described at the end of this article) you will see a new tab "vProtect" in the OpenStack menu. This consists of several sub-tabs that allow you to perform basic actions such as backup, restore or create a new schedule.

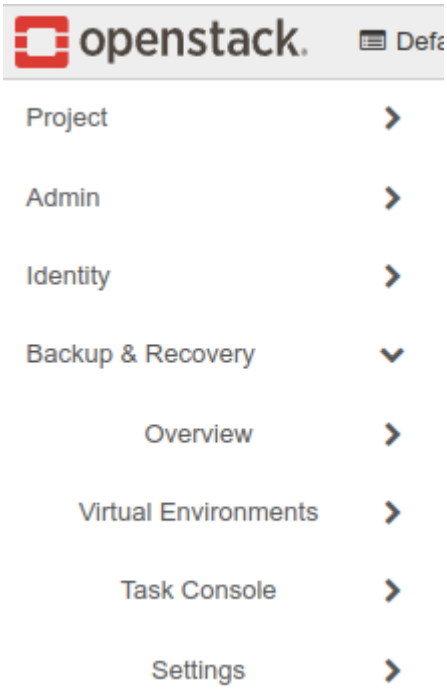

Dashboard

Divided into a few sections, it makes it possible to view and set the most vital options related to management, monitoring, and reporting.

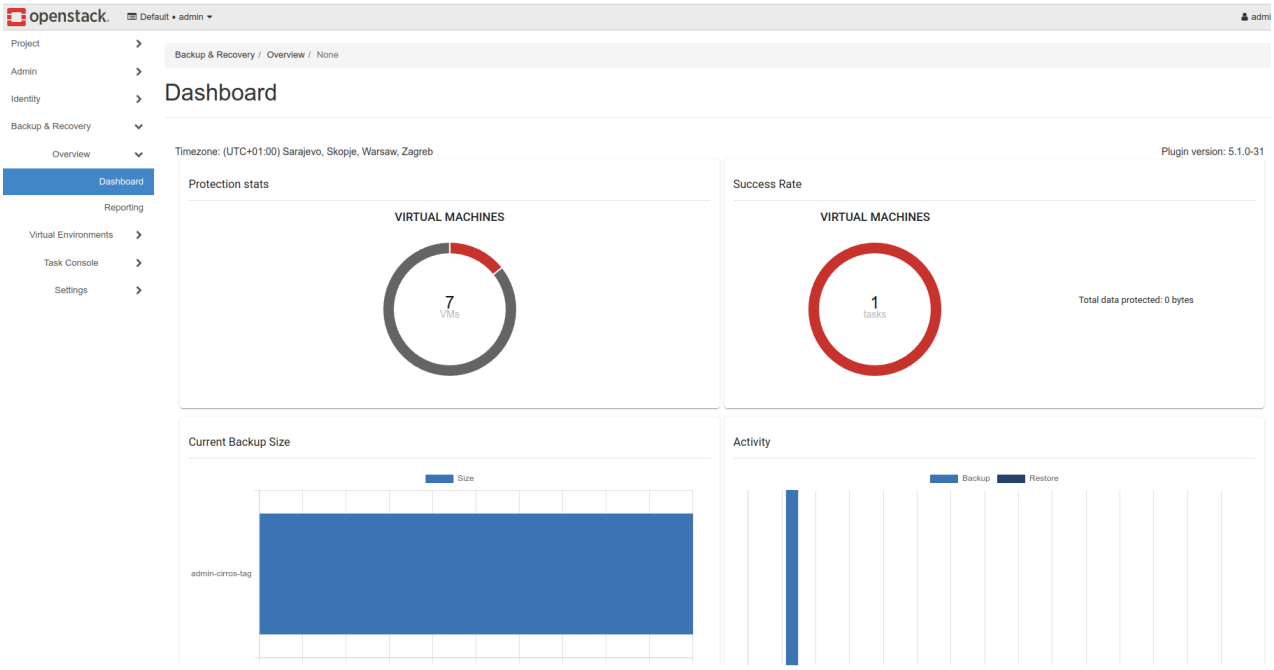

Reporting

Reporting allow users to view statistics, especially for backup and restore tasks. They also provide the possibility to view what has happened lately in the Data Protector for Cloud Workloads environment.

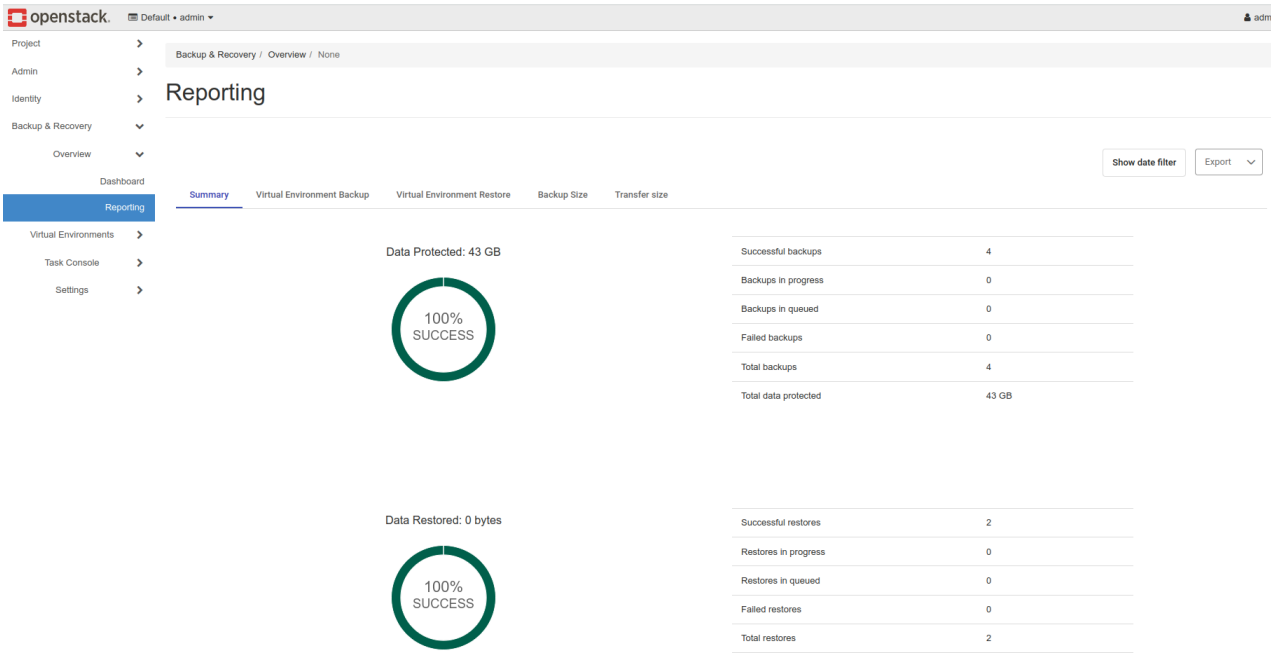

Virtual Environments

Instances

This tab shows all inventoried instances in your OpenStack environment.

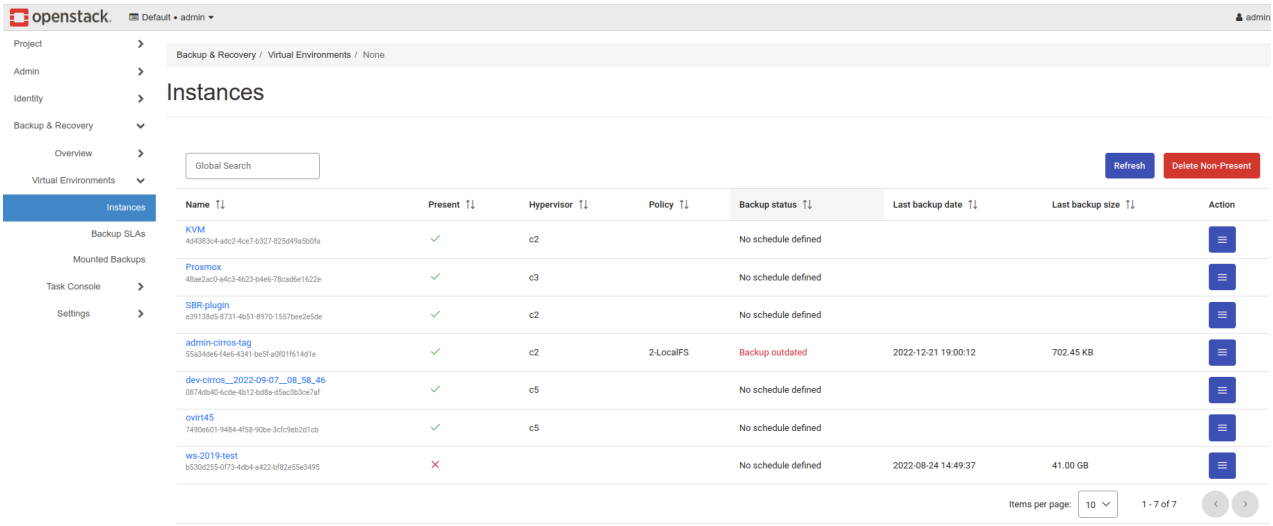

Besides, you can also perform basic backup operations.

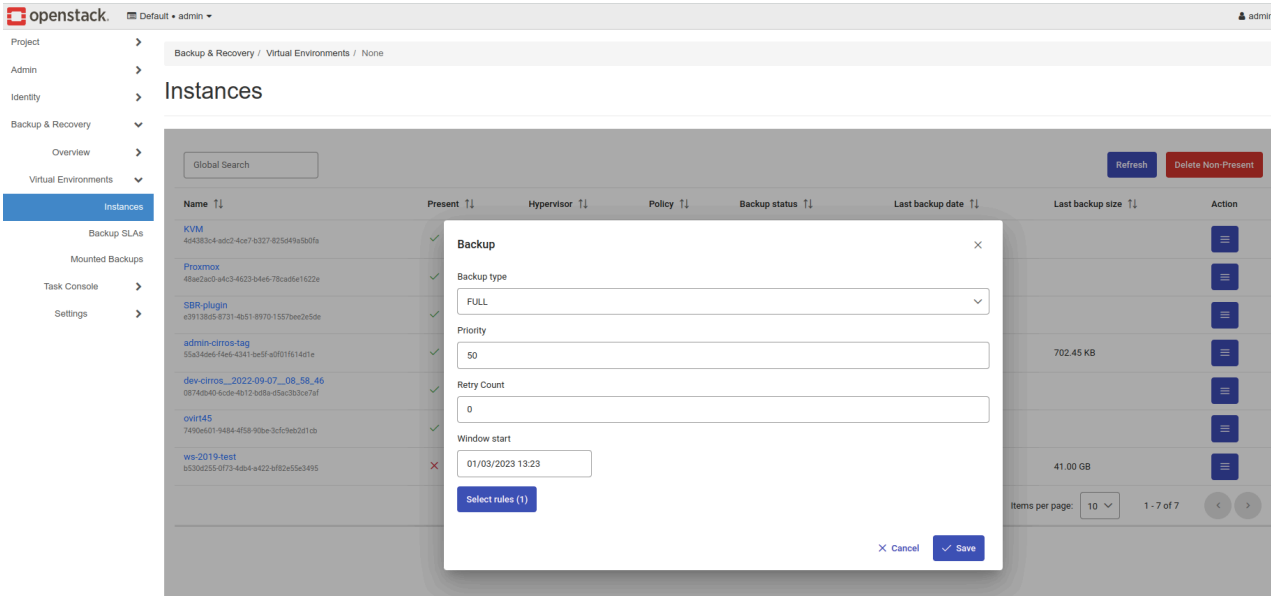

Backup SLAs

Our plugin also allows you to create or manage backup policies

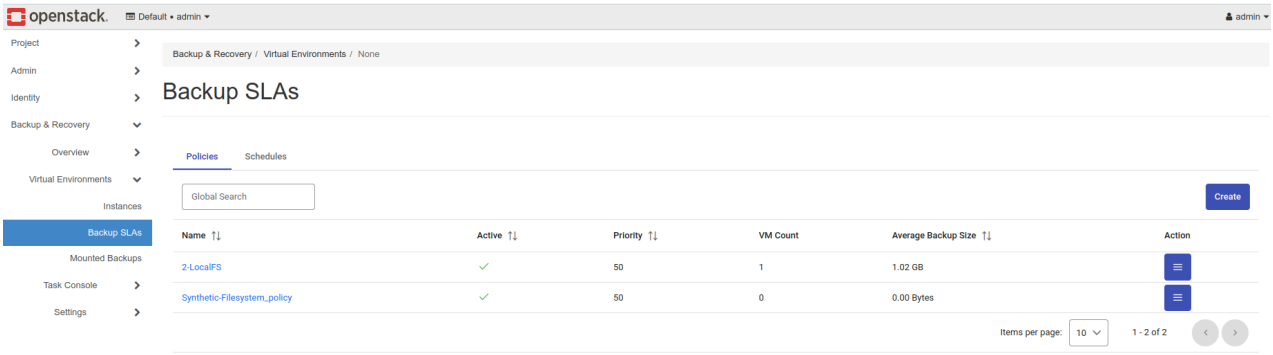

and schedules.

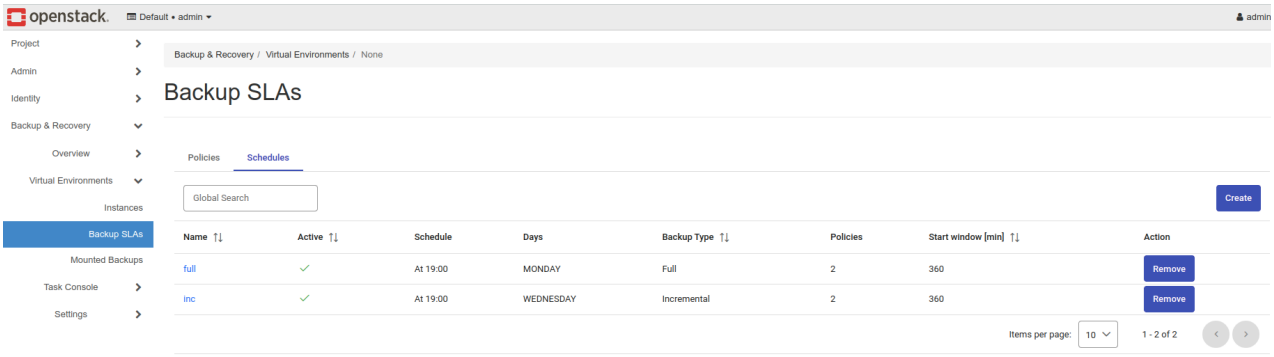

Mounted Backups

Finally, you can also browse your mounted backups from the OpenStack dashboard. You only need to enter the backup details using the menu on the right.

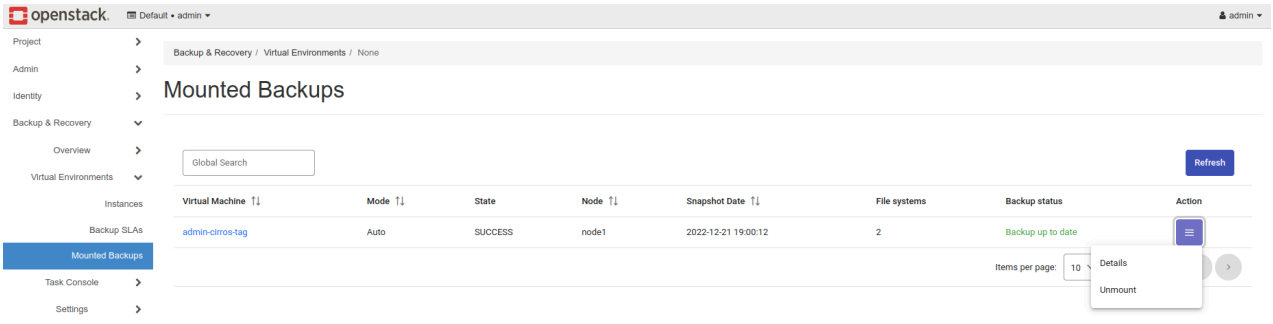

From here we can see the basic information about the backup and start browsing the files using the "Browse" button on the right.

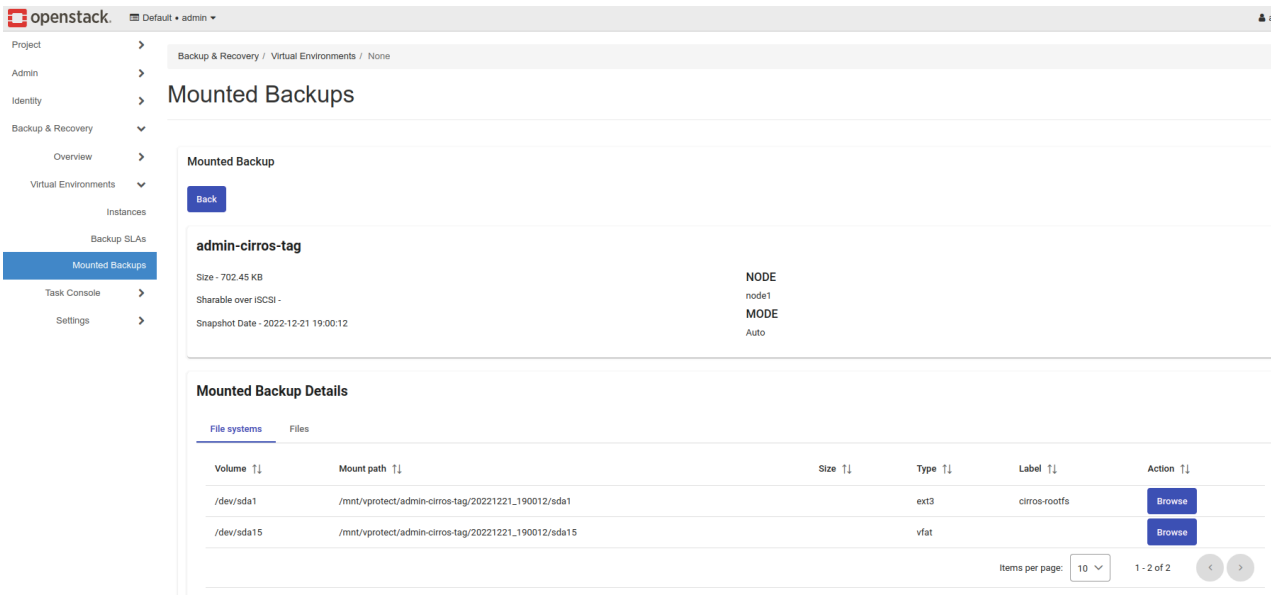

Just select a folder or file and then click on the "download selected" button to have the files on your computer.

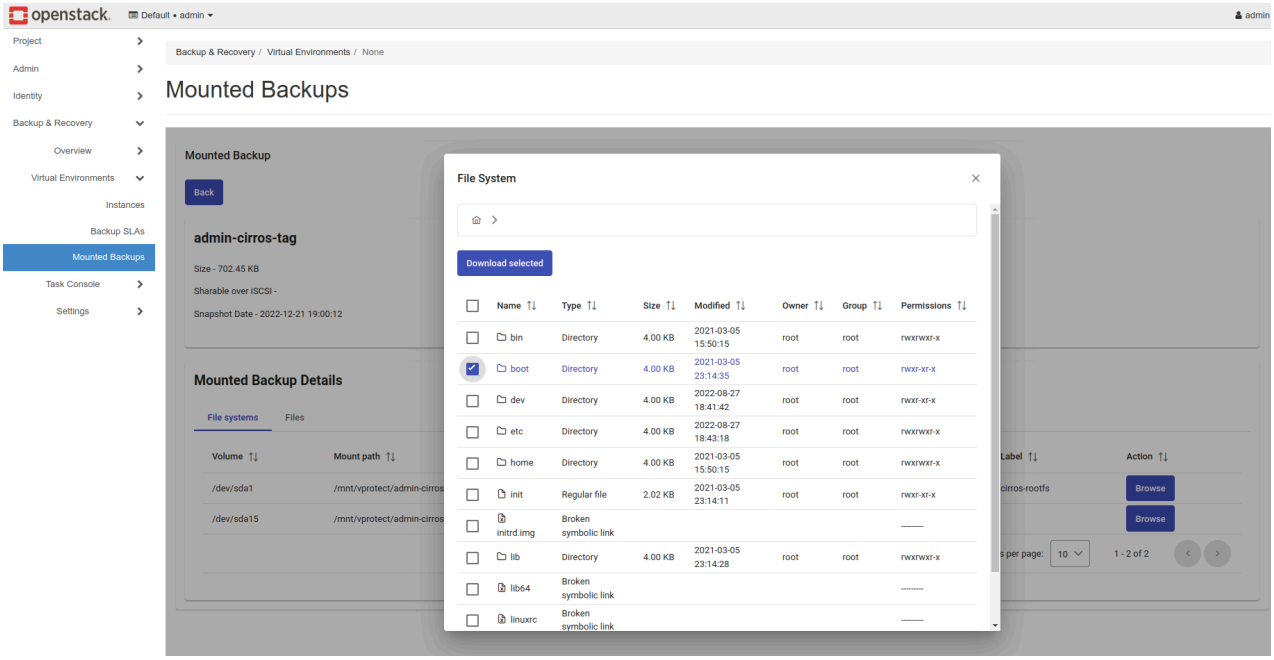

Task Console

Basic information about current tasks performed by Data Protector for Cloud Workloads.

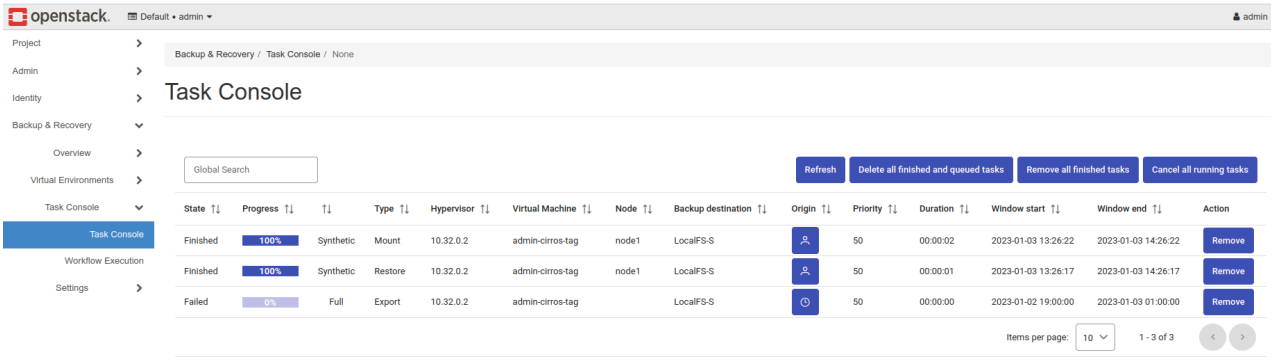

Settings

Mailing

It allows you to create a mailing list that can be used for sending group report e-mails.

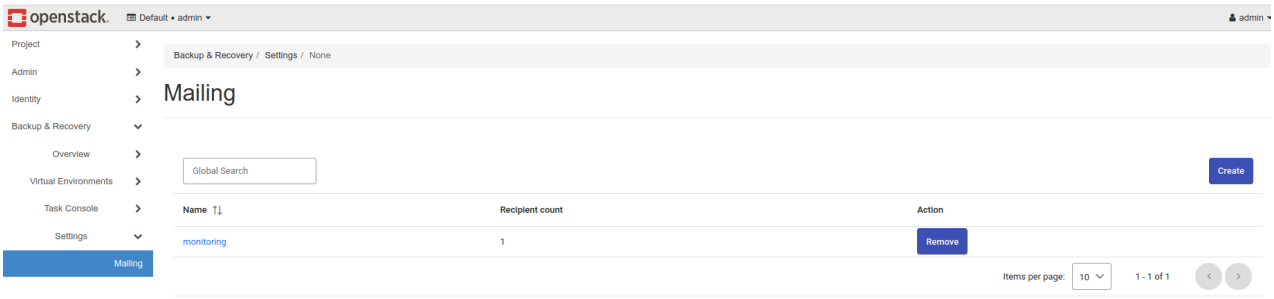

OpenStack general integration setup

> You can find minimum requirements for OpenStack UI Plugin in Support Matrix

You can find the add-on in the GitHub repository ↗. Extract the provided archive onto your Horizon host and execute `python install.py DP4CW_API_URL USER PASSWORD`

Example: `python install.py http://localhost:8080/api admin vprotect`.

Note: you need to restart your Horizon HTTP server after this

The above-mentioned script will copy the plug-in files to the following folders:

- `/usr/share/openstack-dashboard/openstack_dashboard/dashboards/vprotect` - plugin files
- `/usr/share/openstack-dashboard/openstack_dashboard/enabled` - file to enable the plugin

In order to **uninstall** it, remove the `vprotect` subfolder and `enabled/_50_vprotect.py` file and restart your Horizon HTTP server.

Integrate Data Protector for Cloud Workloads dashboard plugin to OpenStack (LXC)

Requirements:

- git, python3-yaml packages
- internet connection

1. Check the name of the horizon container:

```
lxc-ls -f | grep horizon

example:

[root@aio1 ~]# lxc-ls -f | grep horizon  aio1_horizon_container-
b2daccaa RUNNING 1 onboot, openstack 10.255.255.213, 172.29.239.229
- false
```

2. Enter horizon container:

```
[root@aio1 ~]# lxc-attach aio1_horizon_container-b2daccaa
```

3. Install requirements packages:

```
root@aio1-horizon-container-b2daccaa:~# apt install python3-yaml git
-y
```

4. Clone from github installations files:

```
root@aio1-horizon-container-b2daccaa:~# git clone
https://github.com/Storware/openstack-horizon-ui-vprotect-extensions
```

5. Change owner of the plugin directory to horizon:horizon

```
root@aio1-horizon-container-b2daccaa:~# chown -R horizon:horizon
openstack-horizon-ui-vprotect-extensions
```

6. Enter plugin directory:

```
root@aio1-horizon-container-b2daccaa:~# cd openstack-horizon-ui-
vprotect-extensions
```

7. Optionally you can ping Data Protector for Cloud Workloads server by ping

```
root@aio1-horizon-container-b2daccaa:~# ping dp4cw-server-IP-ADDRESS
```

8. Next, install the plugin

```
root@aio1-horizon-container-b2daccaa:~# python3 install.py
http://dp4cw-ip:8080/api admin_user admin_password
```

When the installation process is completed, plugin files should be placed in /usr/share/openstack-dashboard/openstack_dashboard directory. If your path to the dashboard directory is different, create symbolic links from plugin install directories to non-standard directories.

Example:

```
root@aio1-horizon-container-b2daccaa:~# ln -s /usr/share/openstack-
dashboard/openstack_dashboard/dashboards/vprotect
/openstack/venvs/horizon-23.1.0.dev65/lib/python3.8/dist-
packages/openstack_dashboard/dashboards/

root@aio1-horizon-container-b2daccaa:~#  ln -s /usr/share/openstack-
dashboard/static/vprotect /openstack/venvs/horizon-
23.1.0.dev65/lib/python3.8/dist-packages/static/

root@aio1-horizon-container-b2daccaa:~#  ln -s /usr/share/openstack-
dashboard/openstack_dashboard/enabled/_50_vprotect.py
/openstack/venvs/horizon-23.1.0.dev42/lib/python3.8/dist-
packages/openstack_dashboard/enabled/
```

9. Edit /etc/apache2/sites-available/openstack-dashboard.conf file:

- Add alias for static files

```
Alias /dashboard/static /openstack/venvs/horizon-
23.1.0.dev65/lib/python3.8/dist-packages/static/
```

- Directory tag informs you, where dashboards directories should be placed.
- Second Directory tag informs where static directory from plugin should be placed.

Example: configuration file should look like this:

```
# Ansible managed

# If horizon is being served via SSL from this web server,
# then we must redirect HTTP requests to HTTPS.

# If horizon is being served via SSL via a load balancer, we
# need to listen via HTTP on this web server. If SSL is not
# enabled, then the same applies.
<VirtualHost 172.29.239.229:80>
ServerName aio1-horizon-container-b2daccaa.openstack.local
LogLevel info
ErrorLog syslog:daemon
CustomLog "|/usr/bin/env logger -p [daemon.info]
(http://daemon.info/) -t apache2" "%h %l %u \"%r\" %>s %b \"%
{Referer}i\" \"%{User-agent}i\""
Options +FollowSymLinks
RequestHeader set X-Forwarded-Proto "https"

WSGIScriptAlias / /openstack/venvs/horizon-
23.1.0.dev65/lib/python3.8/dist-packages/openstack_dashboard/wsgi.py
WSGIDaemonProcess horizon user=horizon group=horizon processes=1
threads=1 python-path=/openstack/venvs/horizon-
23.1.0.dev65/lib/python3.8/site-packages

WSGIProcessGroup horizon
WSGIApplicationGroup %{GLOBAL}

<Directory /openstack/venvs/horizon-23.1.0.dev65/lib/python3.8/dist-
packages/openstack_dashboard>
<Files wsgi.py >
<IfVersion < 2.4>
Order allow,deny
Allow from all
</IfVersion>
<IfVersion >= 2.4>
Require all granted
</IfVersion>
</Files>
</Directory>

Alias /static /openstack/venvs/horizon-
23.1.0.dev65/lib/python3.8/dist-packages/static/
Alias /dashboard/static /openstack/venvs/horizon-
23.1.0.dev65/lib/python3.8/dist-packages/static/

<Directory /openstack/venvs/horizon-23.1.0.dev65/lib/python3.8/dist-
packages/static/>
Options -FollowSymlinks
```

```
<IfVersion < 2.4>
AllowOverride None
Order allow,deny
Allow from all
</IfVersion>
<IfVersion >= 2.4>
Require all granted
</IfVersion>
</Directory>
</VirtualHost>
Edit /openstack/venvs/horizon-23.1.0.dev65/lib/python3.8/dist-
packages/openstack_dashboard/urls.py and add in urlPatterns
following line
url(r'^dashboard/', horizon.base._wrapped_include(horizon.urls))
Your urls.py should looks like:
"""
URL patterns for the OpenStack Dashboard.
"""

from django.conf import settings
from django.conf.urls import include
from django.conf.urls.static import static
from django.conf.urls import url
from django.contrib.staticfiles.urls import staticfiles_urlpatterns
from django.views import defaults

import horizon
import horizon.base
from horizon.browsers import views as browsers_views
from horizon.decorators import require_auth

from openstack_dashboard.api import rest
from openstack_dashboard import views

urlpatterns = [
url(r'^$', views.splash, name='splash'),
url(r'^api/', include(rest.urls)),
url(r'^header/', views.ExtensibleHeaderView.as_view()),
url(r'', horizon.base._wrapped_include(horizon.urls)),
**url(r'^dashboard/',
horizon.base._wrapped_include(horizon.urls)),**
]

# add URL for ngdetails
```

10. Restart httpd service

```
/etc/init.d/apache2 restart
```

11. After refreshing the dashboard site, you should see Data Protector for Cloud Workloads tab in the Openstack menu.

Installing Horizon Plugin in a Juju Deployment

This guide provides detailed instructions for installing the Horizon plugin in a Juju-deployed OpenStack environment.

Prerequisites

Before you begin, make sure you have the following:

- Access to the Juju controller and the model where the OpenStack services are deployed.
- The IP address of the Data Protector for Cloud Workloads server

Installation

1. SSH into the leader unit of the `openstack-dashboard` application:

```
juju ssh openstack-dashboard/leader
```

2. Clone the Horizon plugin repository from GitHub:

```
git clone https://github.com/Storware/openstack-horizon-ui-vprotect-extensions
```

3. Change into the cloned directory:

```
cd openstack-horizon-ui-vprotect-extensions
```

4. Install the required Python packages using `pip`:

```
pip3 install -r requirements.txt
```

5. Run the plugin installer script with the following command:

```
python3 install.py http://DP4CW_SERVER_IP:8080/api admin vPr0tect
```

- Replace `http://DP4CW_SERVER_IP:8080` with the actual address of your vProtect server.
- `admin` is the username.
- `vPr0tect` is the password.

6. Navigate to the static files directory of the OpenStack Dashboard:

```
cd /usr/share/openstack-dashboard/openstack_dashboard/static/
```

7. Create a symbolic link to the vProtect static files:

```
ln -s /var/lib/openstack-dashboard/static/vprotect/ vprotect
```

8. Check if the following aliases are present in the Apache configuration file:

```
nano /etc/apache2/conf-enabled/openstack-dashboard.conf
```

Ensure the file contains the following aliases:

```
Alias /dashboard /usr/share/openstack-
dashboard/openstack_dashboard/dashboards
Alias /dashboard/static /var/lib/openstack-dashboard/static/
Alias /static /var/lib/openstack-dashboard/static/
Alias /horizon/static /var/lib/openstack-dashboard/static/
```

9. Edit the default Apache site configuration to add a rewrite rule:

```
vi /etc/apache2/sites-available/000-default.conf
```

Add the following lines under `DocumentRoot /var/www`:

```
RewriteEngine on
RewriteRule ^/dashboard/vprotect/(.*)$ http://%
{HTTP_HOST}/horizon/vprotect/$1 [R=301,L]
```

> **Note:** This rewrite rule may need adjustment based on your HAProxy/ReverseProxy setup.

10. Edit the `urls.py` file to include a specific path:

```
vi /usr/lib/python3/dist-packages/openstack_dashboard/urls.py
```

11. Add the following line to the `urlpatterns` section:

```
re_path(r'^horizon/', horizon.base._wrapped_include(horizon.urls)),
```

12. Restart the Apache service to apply the changes:

```
sudo service apache2 restart
```

The Horizon plugin should now be installed and configured in your Juju-deployed OpenStack environment. Verify that the plugin appears in the OpenStack Dashboard and is functioning correctly.

Troubleshooting

Troubleshooting

Data Protector for Cloud Workloads log directory path is `/opt/vprotect/logs` - this is the first place to check for the root cause of errors. All log files are also accessible from `Logs` sections in web UI.

The CLI interface records messages in `vprotect_client.log` files under the subdirectory with the same name as the node. In the same directory you can also find the `vprotect_daemon.log` files, which contain all engine related messages.

The web UI has several directories where it stores log files:

- `appserver` which has all messages coming from the application server - which hosts web UI and API
- `api` which has all messages related to - the core Data Protector for Cloud Workloads Server application
- `cloudagent` which has all messages from Cloud Agent
- `cloudserver` which has all messages from Cloud core part

To verify if services are running, you can use:

- `systemctl status vprotect-server` for the server
- `systemctl status vprotect-node` for the node, or `vprotect status`

If you don't find the root cause of the problem here, you can switch Data Protector for Cloud Workloads to DEBUG mode, and recreate the task to generate logs in DEBUG mode.

How to enable Data Protector for Cloud Workloads DEBUG mode

Available log levels

LEVEL	DESCRIPTION
INFO	Provides general operational messages to confirm that the software is functioning as expected
DEBUG	Offers detailed diagnostic information to help developers troubleshoot issues
TRACE	Provides the most granular level of detail, often used to trace the execution of specific code paths

Changing log level in the components

Server

1. Connect to the server host using SSH
2. Open the following file for editing
   `/opt/vprotect/server/quarkus.properties`

   Example:

   ```
   vi /opt/vprotect/server/quarkus.properties
   ```

3. Locate the line in the file that controls the logging mode
   `quarkus.log.level=INFO`

   The default value for the logging level is INFO

4. Set the value to INFO, DEBUG or TRACE
5. Save the file
6. Restart the service:

```
systemctl restart vprotect-server
```

Node

1. Connect to the node host using SSH
2. Open the following file for editing
`/opt/vprotect/node/log4j2-node.xml`

   Example:

   ```
   vi /opt/vprotect/node/log4j2-node.xml
   ```

3. Locate the line in the file that controls the logging mode
`<Property name="logLevel">INFO</Property>`

   The default value for the logging level is INFO
4. Set the value to INFO, DEBUG or TRACE
5. Save the file
6. Restart the service:

   ```
   systemctl restart vprotect-node
   ```

Cloud agent

1. Connect to the node host using SSH
2. Open the following file for editing
`/opt/vprotect/cloudagent/AgentConfig.json`

   Example:

```
vi /opt/vprotect/cloudagent/AgentConfig.json
```

3. In the JSON file, find the `logs` section nested under the `agent` key
4. Add the following key to the `logs` section
   `"level": "INFO"`

   The default value for the logging level is INFO.

   Set the value to INFO, DEBUG or TRACE.

   After the modification, the section should look like this:

   ```
   agent": {
       "logs": {
           "path": "../logs/cloudagent",
           "level": "DEBUG"
       }
   }
   ```

5. Save the file
6. Restart the node service:

   ```
   systemctl restart vprotect-node
   ```

Cloud server

1. Connect to the node host using SSH
2. Open the following file for editing
   `/opt/vprotect/tapemanager/config/application.properties`

   Example:

   ```
   vi /opt/vprotect/cloudserver/config/application.properties
   ```

3. Locate the line in the file that controls the logging mode
   `quarkus.log.level=INFO`

   The default value for the logging level is INFO
4. Set the value to INFO, DEBUG or TRACE
5. Save the file
6. Restart the node service:

   ```
   systemctl restart vprotect-node
   ```

Collecting logs

Collecting logs

Using the Web UI

General logs

1. Go to the `Logs` section (top bar)
2. You'll see directories for `api` , `cloudserver` , `cloudagent` and one directory for each node (same as their names)
3. To download all logs - click the button `Download all logs` which will generate an archive containing all logs
4. To download a specific file - browse to the directory and click the name of the log

> **Note:** By downloading logs from WebUI you can also download logs for Data Protector for Cloud Workloads Cloud server and node.

Backup-related logs

1. Go to the VM details and `backup history` tab.
2. Click on the second icon on the right of the backup you want to collect logs for.
3. The downloaded file is `vprotect_daemon.log` filtered, so it should contain only entries related to the tasks of backup/restore operations related to this backup.

Directly from the operating system

Data Protector for Cloud Workloads Node

Log files are stored in this folder: `/opt/vprotect/logs/<node_name>`:

- `vprotect_client.log` - stores CLI-related messages
- `vprotect_daemon.log` - stores Data Protector for Cloud Workloads Node engine related message

Data Protector for Cloud Workloads Server

Log files are stored in:

- `/opt/vprotect/logs/api` - Data Protector for Cloud Workloads Server application logs

Disaster Recovery

Disaster Recovery

Internal Database Protection

Data Protector for Cloud Workloads stores all of the metadata in the local database. It is **highly recommended** to set up periodic DB backup on Data Protector for Cloud Workloads Server. Check Initial configuration for more information.

In case you need to restore Data Protector for Cloud Workloads DB:

- if you have a working Data Protector for Cloud Workloads - you can use it to restore this file to the specified location and then restore it as described in this section
- if you don't have working Data Protector for Cloud Workloads - you can try to use the last copy of the database (it is left by default in `/tmp/vprotect_db.sql.gz` on the server host.
- if it is not there and you don't have Data Protector for Cloud Workloads working - you may need to use external tools such as S3 browser or just browse through your file system backup destination directories to find and download this file and later restore it as described in this section.

Once you have your backup you can restore DB with the following command:

```
gunzip < PATH_TO_GZIPPED_BACKUP | mysql -u root -pDBPASSWORD vprotect
```

In case of a complete loss of the Data Protector for Cloud Workloads Server:

1. Reinstall Data Protector for Cloud Workloads Server
   - if you lost your license file, contact support
2. Before starting Data Protector for Cloud Workloads Server - restore the database

    - you can also restore it later (for example if you want to reinstall it with Ansible or all-in-one option), but remember to shutdown server first, then restore DB and start the server again
3. Replace all backup provider-specific files (install any binaries specific for required backup destinations)
4. Start Data Protector for Cloud Workloads Server service
5. Install Data Protector for Cloud Workloads Nodes
6. Make sure the staging path on each node is correct and available
7. Re-register and start nodes

At this point, Data Protector for Cloud Workloads should be ready to continue operation.

Known software issues and limitations

Known software issues and limitations

Issue ID	Product feature	Description	Workaround
0001	Task cancellation	Task cancellation process will not be reflected in immediate task failure - task state will be changed to cancelled and only when the engine checks its state again will it initiate the cancellation operation - some platforms may even require data transfer to be completed first	Allow the task to cancel and fail gracefully - this will allow Data Protector for Cloud Workloads to clean up temporary artifacts. If you click again, the task will be forced to be removed from the queue and artifacts such as snapshots will be removed as part of the daily snapshot cleanup job. In general, avoid forced removal of tasks.
0002	Storage usage statistics	Storage statistics are updated after each backup or the clean old backups job - this data may not be up-to-date all the time	To have current storage usage updated you can invoke the Clean Old backups job from the Backup Destinations tab
		Complete command cannot be provided as a single string. Commands need to have their arguments provided	To use shell-specific operators/commands etc., execute commands with 3

0003	Pre/post access storage command execution	as separate entries (by clicking on the **Add command arg** button). Commands are directly executed using OS-level calls so shell operators are not supported directly.	command arguments `/bin/bash`, `-c`, `your command-with-all-arguments-and-shell-operators`
0004	Tasks stuck in the queue in Queued state	Tasks will usually be executed according to the limits set on the node and only if node is running and has available space on the staging	Verify that the node has available space in the staging space path - there should be a warning message in `vprotect_daemon.log`
0005	OpenStack backup using disk attachment	OpenStack with disk-attachment backup strategy (cinder) - 3.9.2 only supports Ceph RBD as a storage backend	N/A
0006	KVM stand-alone - disk formats	VMs being backed up must have virtual disks as QCOW2/RAW files or LVM volumes	N/A
0007	KVM stand-alone - snapshots	Snapshots on KVM hypervisors are made using libvirt (QCOW2/RAW files) or LVM snapshots and are created per volume basis; this operation may not be atomic if multiple drives are used	Make sure the data is in the VM (especially that file systems reside on as few disks as possible to lower the risks of data inconsistency) or try to use pre-post remote command execution to quiesce application

			before snapshot is done
0008	KVM standalone incremental backup on QCOW2	Incremental backups will be performed only on running VMs. libvirt doesn't allow blockcommit on a power-down VM so snapshot wouldn't be removed.	Full backup will be performed instead
0009	Backup providers path	The backup provider's paths must be mounted and available ahead - you should provide the path just to the mount point, without any protocol specification	Mount remote file systems first and make sure these are available all the time - in Data Protector for Cloud Workloads configuration provide just the locally available mount point
0010	Backups marked as Success (removed)	When backup completes or when the clean old backups job is performed, Data Protector for Cloud Workloads marks non-present backups as removed (if any of the files that were part of the backup are not present). This may also happen if your storage was temporarily not available.	Make sure storage and all of the files are available and run the Clean Old Backups job - this job also attempts again to sync files present in the backup provider with the database - if all files for a particular backup are found again (and all previous backups that this particular backup depends on are also present) it will again have Success status

0011	Hypervisor storage usage statistics	Restore may fail due to insufficient storage space in the Hypervisor Storage used as a target because of usage information that is not up to date. Usage statistics are updated only with inventory synchronization job	Run the Inventory Synchronization job again on your hypervisor (or manager) to update storage statistics and try to restore again.
0012	RHV - SSH transfer rate drops after some time	The SSH transfer rate may drop in some environments when used intensively over a longer time.	If possible, and when the network used for transfers is trusted, please use the netcat option to transfer files outside of the SSH channel
0013	Amazon EC2 - AMIs left in the account	For Amazon EC2 some instances require the original base image to be restored - this is especially true for Windows-based clients where license relates to the original disk image. If an image is not left, Data Protector for Cloud Workloads can only restore such guests by creating a new one from the new image (as a root device) that is available and attach data volumes. AMIs are kept as long as	For such guests, we recommend to enable Windows (or Linux) image required option in your Hypervisor Manager details

		the particular backup is going to use them and will be removed together.	
0014	AWS additional costs	Notice that Data Protector for Cloud Workloads needs sometimes to transfer EBS volumes between AZ if it resides in a different AZ then the node - AWS charges for intra-AZ transfers.	Recommended deployment is in the same AZ as the VMs that node is going to protect to limit the number of transfers.
0015	Node tasks limits	The number of concurrent tasks are configured in Node Configuration → Tasks section. These limits apply to all nodes that use this particular configuration. Currently there is no global setting to limit the number of tasks for all of the nodes in the environment.	To limit the number of tasks globally, reduce the numbers in individual node configurations.
		All of the configuration parameters in Hypervisors tab in Node Configuration are applied to all nodes with this configuration - regardless of which hypervisor it is attached to. This implies that	To use these settings with different values for some hypervisors you need to assign

0016	Hypervisor-specific settings in Node Configuration	Proxmox settings such as compression will have to be the same on all hypervisors handled by nodes with the same configuration assigned and will have to be the same on all of these hypervisors.	separate nodes and define separate node configurations. Ultimately, assign separate nodes for these hypervisors.
0017	Inventory synchronization - duplicated UUID	In some cases, it may happen that the same storage was previously detected with a different setup and remained in the database.	Remove the unused hypervisor storage and try to invoke inventory synchronization again.
0018	Estimated backup size of policy	The estimated backup size of a policy is computed based only on known backup sizes and extrapolated to the rest of the VMs in the group. This implies that estimation will use average backup size and multiply it by the number of all VMs in the group. Even though disk sizes are known it is not always the same as the size of the backups (especially considering compression or the	Wait for a longer period of time, and once more backups are completed this estimation will be closer to the real value.

		fact that some strategies require chains of backup deltas to be exported)	
0019	Citrix Hypervisor/ xcp-ng - transfers	Transfer NIC is not used in incremental backups when the CBT strategy is invoked - Citrix/xcp-ng may require NBD to be exposed by the master - so Data Protector for Cloud Workloads has to read from the address provided by the CBT mechanism in order to connect to the NBD device. Also, in some cases, data can only be transferred from the master host (especially when it is powered down)	Allow network traffic between all hypervisors and corresponding nodes in the same pool, as sometimes actual transfer may occur from the master host instead of the one which hosts the VM.
0020	RHV - SSH Transfer permissions on the hypervisor	SSH Transfer for RHV usually requires root permissions on the hypervisor in order to activate/deactivate LVM volumes for the backup	You may try with a different backup strategy such as Disk-attachment or Disk Image Transfer
		Data Protector for Cloud Workloads using SSH Transfer for RHV environments needs to be able to	

0021	RHV - SSH Transfer - hypervisor access	access all hypervisors in the cluster - as it may happen that the created disk is available only on a subset of them and needs to be transferred or recovered only by using this specific hypervisor	Allow network traffic and provide valid credentials to access all hypervisors in the cluster over SSH.
0022	Nutanix VG support	Data Protector for Cloud Workloads 3.9.2 only supports volumes residing on the storage containers - VGs are not supported yet	N/A
0023	Nutanix Prism Element/Central connectivity	Data Protector for Cloud Workloads is able to perform backups by using APIs provided by Prism Element only - Prism Central doesn't offer a backup API	Connect to your Prism elements by specifying separate Hypervisor Managers in Data Protector for Cloud Workloads (not by pointing to the Prism Central)
0024	Nutanix backup consistency for intensively used VMs	Intensive workload on the VM may affect backup consistency when using crash consistent backup	If you need higher consistency, install Nutanix Guest Tools inside your VM and enable application consistent snapshots in the VM details in Data Protector for Cloud Workloads
		RAW backups allow Data Protector for Cloud Workloads to	Use automatic mount instead or

0025	iSCSI shares for RAW backups	share them over iSCSI. If backups are in other formats (such as QCOW2), these cannot currently be shared over iSCSI	restore a backup and mount them using external tools such as qemu-nbd for QCOW2 files
0026	Backup and snapshot policies assignment	Only 1 backup and 1 snapshot management policy can be assigned to a given VM	If you need a dedicated setting for a single VM, you need to create a separate policy for that VM
0027	1 schedule per rule	Currently, each schedule can only be assigned to a single rule within the same policy.	If you need to execute two rules at the same time, you need to create separate schedules and assign them to these rules
0028	Staging space	The staging space is an integral part of a node - this allows to mix backup strategies, especially based on the export storage domain/repository approach with other methods and file system scanning for future file-level restores. It needs to be available at all times and the `vprotect` user needs to be able to write to all subdirectories.	To save space and boost backup time (direct writes to the backup destination) you can, however, mount staging and your PowerProtect DD in the same directory - `/vprotect_data`. Remember to point the backup destination path to a subdirectory of this mount point, such as `/vprotect_data/backups` - still on the same FS, but the paths must be different.

0029	Node OS-level permissions	At the OS level, Data Protector for Cloud Workloads requires significant permissions to be able to manipulate disks, scan for file systems, mount them, expose resources over iSCSI, operate on block devices (NBD/iSCSI/RBD) and more. These unfortunately require multiple sudo entries and that SELinux is disabled at the same moment.	If some features are not required at the same moment - including NBD/iSCSI/NFS related - you may reduce the number of entries in `/etc/sudoers.d/01-vprotect_node` You also can try to enable SELinux, but later you need to track SELinux errors and add appropriate permissions when some of the functionality is blocked
0030	Proxmox VE	CBT backup strategy	Qcow2 virtual machines are required to use the new Proxmox VE backup strategy - Change block tracking (CBT)
0031	Proxmox VE	CBT backup strategy	At the moment, we do not support the "Dirty Bitmaps" function, therefore we require the last snapshot to be left for incremental backups.
		Data Protector for Cloud Workloads supports only one node assigned to the Storage Provider, which means that backup of significantly big	Install multiple Data Protector for Cloud Workloads Server+Node environments and

0032	Storage Providers and node assignment	volumes from bigger storage providers, for example Ceph RBD etc. will require a high performing node and cannot be scaled out by adding nodes	protect the non-overlapping set of volumes with each Data Protector for Cloud Workloads instance.
0033	RHV - Restore with SPARSE disk allocation format	Restore to RHV using the SPARSE disk allocation format is not supported if backup files are in RAW format and the destination storage domain type is in either Fibre Channel or iSCSI. If such configuration is detected, then the disk allocation format is automatically switched to PREALLOCATED	You can use other backup strategies that use QCOW2 files instead of RAW (like disk image transfer). Alternatively, you can select a different storage domain of a type that supports SPARSE disks with RAW files
0034	Microsoft 365 - Restore of site which has been deleted from Bin	If the site has been deleted from the Bin, only site logic can be restored. Links added in the deleted site are not restored	Restore site or subsites to recover it's logic. Next, download data (site content) manually and upload it to SharePoint Online. Begin download from second level of SharePoint protected data (list/pages/document libary)
			You can still download shared files by copying the

0035	Microsoft 365 - Restore 1:1 teams chat	Links shared in chat are not working after restore.	link address and pasting it in the different web browser tab.
0036	Microsoft 365 - Restore site from the template	Sometimes after restoring the site from the template, despite the lack of errors in the logs, the page template is not set.	This condition can be repaired by: - re-restore - manual set template
0037	Microsoft 365 - Restore site	Sometimes after restoring the site, you may see `You need permission to access this site` message	This is due to the lack of a site owner. This condition can be repaired by: - re-restore - set the site owner in the admin panel
0038	Microsoft 365 - Restore site	After restoring the site, the correct images are not visible everywhere	The reason is the change of links to images and they do not always reload correctly after restore. This can be fixed by selecting the desired image from the library again.
0042	Microsoft 365 - Protected data download	Downloading more than one object from Microsoft 365 protected data fails	N/A

Glossary

Glossary

- **Backup Destination** - backup provider or storage space holding backups on Data Protector for Cloud Workloads Node where backup files are copied to from **Staging Space** - default and the only currently supported is PowerProtect DD via Boost FS
- **Backup policy** - are responsible for the automation of backups of Virtual Environments or Storage instances. Backup SLA consists of the **policy** and **schedule**.
- **Cluster** - corresponds to server pools/clusters/availability zones that have been detected during inventory synchronization of **Hypervisors**
- **Hypervisors** - a list of hypervisors automatically discovered during inventory synchronization of **Hypervisor Manager** or manually added to Data Protector for Cloud Workloads
- **Hypervisor Managers** - a list of hypervisors managers added to Data Protector for Cloud Workloads.
- **Instances** - a list of currently known virtual machines/storage.
- **Inventory synchronization** - a task that index the contents of **Hypervisor Manager**, **Storage Provider,** or **Hypervisor** (if it's not managed by Hypervisor Manager)
- **Mounted backups** - a list of backups which has been mounted by Data Protector for Cloud Workloads and can be browsed.
- **Node** - machine or VM with installed Data Protector for Cloud Workloads Node, its main job is to execute a backup, restore, and mounting tasks. It should have access to the **backup destination** and **staging space**.
- **Node Configuration** - contains settings for nodes to describe their behavior during tasks execution such as maximum numbers of simultaneous backup tasks, timeouts, or **backup destinations** which can be used by nodes. One node configuration can be attached to many **nodes**.
- **Policy** - allow you to group virtual machines or storage instances. Each policy can have multiple **schedules** assigned.

- **Recovery Plans** - are used to automate the DR process, so that Data Protector for Cloud Workloads executes multiple restore operations to the target environment with preconfigured settings.
- **Schedule** - allow you to invoke specific policies periodically. This allows you to back up multiple VMs or storage instances automatically.
- **Snapshot policy**- are responsible for the automation of creating snapshots of Virtual Environments or storage instances. Backup SLA consists of the **policy** and **schedule**. The Instance has to be assigned to the **snapshot policy** in order to execute snapshot on demand.
- **Staging Space** - temporary space for backup files mounted on a Data Protector for Cloud Workloads Node
- **Storage** - corresponds to datastores/storage repositories/storage domains that have been detected during inventory synchronization of **Hypervisors**.
- **Storage Provider** - software storage platform that provides storage instances.